Politecnico di Torino

ScuDo
Scuola di Dottorato ~ Doctoral School
WHAT YOU ARE, TAKES YOU FAR

Doctoral Dissertation
Doctoral Program in Physics (28^{th}cycle)

Cavity-Based Approaches to Stochastic Dynamics on Sparse Graphs

From Ecological Systems to Epidemiological Inference

By

Mattia Tarabolo

******

Supervisor(s):
Prof. Luca Dall'Asta, Supervisor
Prof. Alfredo Braunstein, Co-Supervisor

Doctoral Examination Committee:
Prof. Maria Chiara Angelini, *Referee*, Sapienza Università di Roma, Italy
Prof. Tobias Galla, *Referee*, Institute for Cross-Disciplinary Physics and Complex Systems, Spain

Politecnico di Torino
2026

Declaration

I hereby declare that, the contents and organization of this dissertation constitute my own original work and does not compromise in any way the rights of third parties, including those relating to the security of personal data.

Mattia Tarabolo
2026

* This dissertation is presented in partial fulfillment of the requirements for **Ph.D. degree** in the Graduate School of Politecnico di Torino (ScuDo).

Ad Anna

Acknowledgements

I wish to express my deepest gratitude to all the people who contributed to the completion of this thesis.

To my supervisor, LDA, for encouraging me to pursue this research, for teaching me how to carry it out, and for always being present and supportive. To AB, without whom I would not be here, and who has always been ready to help whenever I needed it. To all the members of our group – the scientific discussions we had over these three years have deeply shaped me as a researcher.

To NDE, DA, FC, LP, FF, and all the people who passed through the Saletta: thank you for making our lunches a daily source of laughter and good humor.

To RM, JP, ABT, and all the Cubans who welcomed me so warmly on that incredible island.

To all my friends, who have always stood by my side. A special thanks to LB, AF, and the BBB group, for always being there – even when we are far apart.

To A, for being my constant source of strength and serenity. Thank you for your patience, for your understanding in the difficult moments, and for reminding me, every day, that there is life beyond research. None of this would have been possible without your love and support.

And finally, to my family, for always being there for me with unconditional love. Thank you for your encouragement, for your pride, and for never stopping believing in me – even if, to this day, my attempts to explain what I actually study and why it matters have not entirely succeeded. Everything I have achieved begins with you.

Abstract

This thesis develops a coherent and flexible framework to study stochastic complex systems on sparse graphs through controlled small-coupling approximations of the dynamic cavity method. The aim is to extend cavity-based approaches, traditionally confined to equilibrium settings or discrete processes, to a wider class of stochastic dynamics with continuous or discrete degrees of freedom, where heterogeneity, noise, and network structure play a crucial role.

Part I introduces the mathematical foundations underpinning the whole work. It provides a detailed presentation of the cavity method for both equilibrium and dynamical systems, highlighting the difficulties that arise in its dynamic formulation and motivating the need for systematic approximations. The formal connection between the dynamic cavity equations and the path-integral representation of stochastic dynamics is established, and key tools such as the Martin-Siggia-Rose-Janssen-De Dominicis formalism, Dynamical Mean-Field Theory (DMFT), and the Plefka expansion are reviewed.

Building on this groundwork, Part II introduces the Gaussian-Expansion Cavity Method (GECaM), a systematic approximation scheme obtained from a second-order expansion of the dynamic cavity equations in the coupling strength. The method assumes a Gaussian form for the cavity marginals, parametrized by local means, correlations, and response functions, and yields a closed set of integro-differential equations valid for linearly coupled stochastic differential equations on locally tree-like graphs. For additive noise and linear drift, these equations provide an exact description of the dynamics and recover classical results from random-matrix theory. A perturbative closure scheme extends the framework to nonlinear forces and multiplicative noise. The validity of the approach is tested on models exhibiting paramagnetic-ferromagnetic phase transitions driven by cubic drift terms and on models with noise-induced transitions, for which GECaM performs excellently. Applications to the spherical two-spin model reveal how finite connectivity alters relaxation and ageing, while the extension to the random generalized Lotka-Volterra dynamics produces cavity fixed-point equations for species abundances. Solved by a population-dynamics algorithm, these equations capture how sparsity and heterogeneity reshape stability, diversity, and coexistence in large disordered ecosystems.

Part III presents the Small-Coupling Dynamic Cavity (SCDC) method, a cavity-based framework for Bayesian inference in compartmental epidemic models. Derived as a first-order expansion of the observation-reweighted dynamic cavity equations, SCDC introduces conjugate fields that encode the backward propagation of information generated by partial observations. Its efficient transfer-matrix formulation yields a computational cost linear in the epidemic duration, enabling accurate reconstruction of infection probabilities on both synthetic and real-world contact networks. The method provides a principled and interpretable approach, maintaining high accuracy even in recurrent epidemic processes where belief propagation is not applicable due to its prohibitive computational cost.

Contents

List of Figures

Notation

We summarize here the notation used throughout this thesis. We consider graph instances $G = (V, E)$, where $V = 1, \ldots, N$ is the set of N nodes and $E \subseteq V \times V$ is the set of edges. The discrete indices corresponding to the nodes of the graph are indicated by letters i, j, k, and l.

Complex variables are denoted by z, and their real and imaginary parts by $x = \mathrm{Re}\, z$ and $y = \mathrm{Im}\, z$, respectively, such that $z = x + \mathrm{i}y$. The complex conjugate is written as $z^* = x - \mathrm{i}y$, where i is the imaginary unit. The Dirac delta distribution over the complex plane is denoted by $\delta(z) = \delta(x)\delta(y)$, and we use the complex derivatives $\partial_z = (\partial_x - \mathrm{i}\partial_y)/2$ and $\partial_{z^*} = (\partial_x + \mathrm{i}\partial_y)/2$. The integration measure over the complex plane is defined as $d^2 z = dx\, dy$.

Boldface letters $\boldsymbol{x} = (x^1, \ldots, x^n)$ denote temporal row vectors, whose indices n and m correspond to discretized times. Letters with an underline, $\underline{x} = (x_1, \ldots, x_N)$, denote large row vectors of size $\mathscr{O}(N)$, whose indices i, j, and k run over nodes of the graph. When each element of $\underline{x}$ is itself a time vector, we write $\underline{\boldsymbol{x}}$ to emphasize both the temporal and spatial structure. The transpose operation is denoted by the superscript $^\top$.

Capital letters with a double underline, $\underline{\underline{A}}$, denote large matrices of size $O(N)$. Roman capital letters, such as A, indicate 2×2 square matrices whose elements can themselves be matrices. Boldface capital letters, such as $\boldsymbol{A}$, denote time matrices whose indices correspond to discretized times.

Unless otherwise specified, summations over repeated indices are explicit, while time integrals are written in their continuous form only when taking the limit of vanishing discretization step. The notation $\langle \cdot \rangle$ denotes averages over dynamical or disorder realizations, depending on the context.

List of Acronyms

BM Bouchaud-Mézard.

BP Belief Propagation.

DMFT Dynamical Mean-Field Theory.

DMP Dynamic Message Passing.

ERG Erdős-Rényi Graph.

FC Fully-Connected.

FDT Fluctuation-Dissipation Theorem.

GECaM Gaussian Expansion Cavity Method.

GLV Generalized Lotka-Volterra.

IBMF Individual-Based Mean-Field.

MC Monte Carlo.

MPBP Matrix-Product Belief Propagation.

MSRJD Martin-Siggia-Rose-Janssen-De Dominicis.

OU Ornstein-Uhlenbeck.

PopDyn Population Dynamics.

RRG Random Regular Graph.

SAD Species Abundance Distribution.

SCDC Small-Coupling Dynamic Cavity.

SDE Stochastic Differential Equation.

SEIR Susceptible-Exposed-Infected-Recovered.

SEIRS Susceptible-Exposed-Infected-Recovered-Susceptible.

SFN Scale-free Network.

SI Susceptible-Infected.

SIR Susceptible-Infected-Recovered.

SIRS Susceptible-Infected-Recovered-Susceptible.

SIS Susceptible-Infected-Susceptible.

SMF Simple Mean-Field.

TTI Time-Translationally Invariant.

WS Watts-Strogatz.

Stochastic Dynamics in Complex Systems

The study of stochastic dynamical systems on networks is central to understanding a wide range of complex systems, from the spread of infectious diseases to the stability of ecological communities and the activity of neural circuits. When the underlying interaction network is sparse–as is the case for most real-world systems–local correlations and heterogeneities play a crucial role, and traditional fully connected mean-field approaches often fail to capture them. In statistical physics, two powerful frameworks have been developed to study such systems: the cavity method, which leverages the local tree-like structure of sparse graphs to produce accurate message-passing equations, and the path-integral formalism, which, in the fully connected limit, underlies DMFT and related Gaussian approximations.

This thesis builds on these two traditions to develop generalized cavity-based methods for stochastic dynamics on sparse graphs. We begin by revisiting the path-integral and DMFT approaches for continuous degrees of freedom, and extend them to sparse networks via the GECaM. GECaM yields efficient message-passing equations for the moments of local stochastic processes, enabling the study of linearly-coupled stochastic differential equations on large sparse graphs. Its application to Generalized Lotka-Volterra (GLV) models reveals the interplay between network structure and ecological stability, and opens the way to modeling noisy neural dynamics.

We then turn to discrete-state Markov processes, where dynamic cavity methods are well established but computationally expensive for recurrent dynamics. To address this, we introduce the Small-Coupling Dynamic Cavity (SCDC) method, a perturbative expansion that reduces the complexity of solving dynamic cavity equations for recurrent epidemic models. SCDC enables efficient inference of epidemic trajectories from partial observations on large networks. Together, GECaM and SCDC provide a unified and computationally tractable framework for studying a broad class of stochastic dynamical systems on sparse graphs, bridging the gap between continuous-variable mean-field theory and discrete-state network inference.

Part I

Introduction and Background

Chapter 1

A short introduction to graphs

Graphs provide a natural framework to represent interactions in complex systems. In statistical physics, they typically encode the local interaction structure among individual components. Graphs are widely used to model diverse systems, from social and biological networks to contact networks in epidemiology. This chapter introduces the basic concepts of graph theory, discusses the properties of sparse random graphs, and presents the main ensembles of sparse graphs used throughout this thesis. What follows is intended as a non-exhaustive introduction to the topic; for a more comprehensive treatment, the reader is referred to standard references such as Bollobás [1], Barrat, Barthélemy and Vespignani [2], and Newman [3].

1.1 Mathematical background on graphs

An *undirected graph* G is defined as an ordered pair $G = (V, E)$, where V is the set of vertices (or nodes) and E is the set of edges (or links). The set of vertices is denoted by $V = 1, 2, \ldots, N$, where $N = |V|$ is the total number of vertices. Edges connect unordered pairs of distinct vertices, so that each edge can be written as (i, j) with $i, j \in V$ and $i \neq j$. Accordingly, the edge set E is a subset of the Cartesian product of vertices,

$$E \subseteq V \times V = \{(i, j) \mid i, j \in V, i \leq j\}.$$

A convenient way to represent a graph is through its *adjacency matrix* $\underline{\underline{A}} = \{A_{ij}\}_{i,j=1}^{N}$, defined by

$$A_{ij} = \begin{cases} 1 & \text{if } (i,j) \in E, \\ 0 & \text{otherwise.} \end{cases} \tag{1.1}$$

For undirected graphs, the adjacency matrix is symmetric, $A_{ij} = A_{ji}$, and has zero diagonal entries when self-loops are excluded. For the purpose of this thesis, we will focus on undirected graphs without self-loops.

The set of vertices directly connected to a vertex i is called its *neighborhood* and is denoted by $\partial i = j \in V \mid (i,j) \in E$. The number of neighbors of i, i.e. the cardinality of ∂i, is called the *degree* of i and is denoted by $k_i = |\partial i|$, or equivalently as the sum of the entries in the i-th row of the adjacency matrix, $k_i = \sum_{j=1}^{N} A_{ij}$. The degree provides a simple local measure of connectivity, which will play an important role when describing sparse networks. The average degree of the graph is then

$$\langle k \rangle = \frac{1}{N} \sum_{i=1}^{N} k_i = \frac{2|E|}{N}. \tag{1.2}$$

Beyond local neighborhoods, one can also describe the global structure of a graph in terms of walks and connectivity. A *walk* is a sequence of vertices such that each consecutive pair is connected by an edge. A *path* is a walk in which all vertices are distinct. A *cycle* (or loop) is a path whose initial and final vertices coincide.

The *distance* between two vertices $i, j \in V$, denoted $d(i,j)$, is defined as the number of edges in the shortest path connecting them. Given a vertex i and a non-negative integer r, the *ball of radius r* around i, is the set of all vertices at distance at most r from i, together with the induced subgraph.

Two vertices $i, j \in V$ are said to be *connected* if there exists a path from i to j. A graph is connected if every pair of vertices is connected; otherwise it is called disconnected and decomposes into *connected components*, defined as maximal connected subgraphs. Another useful structure is a *bipartite graph*, in which the vertex set can be partitioned into two disjoint subsets U and W such that all edges connect a vertex in U with a vertex in W, and never within the same subset.

A graph in which all possible pairs of vertices are connected by an edge, i.e. with

$$|E| = \binom{N}{2} = \frac{N(N-1)}{2},$$

is called a *complete graph* (or *fully-connected graph*). In general, the number of edges of a connected graph can range from the minimal value $|E| = N-1$, corresponding to a tree, up to the maximal value $\binom{N}{2}$ of the complete graph. A graph is said to be *dense* if the number of edges is close to the maximal number of edges. The opposite, a graph with a number of edges close to the minimal value, is said to be *sparse*. In this thesis we will define a graph as sparse if its number of edges grows at most as $|E| \sim O(N \log N)$, which implies that the *connectance* (or *density*)

$$D = \frac{|E|}{\binom{N}{2}} = \frac{2|E|}{N(N-1)} \tag{1.3}$$

tends to zero as $N \to \infty$. Equivalently, it follows that the average degree must be much smaller than the number of vertices, $\langle k \rangle \ll N$. Conversely, when $|E| \sim N^2$ the graph is called *dense*. In this case the average degree grows linearly with N, $\langle k \rangle \sim N$. Sparse graphs are of particular interest in this thesis, as they capture the finite connectivity of real interaction networks and exhibit structural properties absent in dense limits.

An important example of sparse graphs is given by *trees*, i.e. connected graphs with no cycles. Trees satisfy the relation $N = |E| + 1$, and the removal of any edge breaks the graph into two disconnected components. The loop-free structure of trees provides a useful idealization when analyzing dynamical processes on sparse networks. Finally, sparse graphs can be further classified according to the heterogeneity of their degree distribution. In *regular graphs*, also called *homogeneous* networks, all vertices have the same degree $k_i = K$. In contrast, *heterogeneous* graphs are characterized by broad degree fluctuations across vertices, a feature often encountered in real-world networks.

1.2 Sparse random graphs and their properties

The theory of random graphs, initiated by Erdős and Rényi in the late 1950s [4], brought a probabilistic viewpoint into graph theory, shifting attention from the enumeration of particular structures to the study of typical properties of graphs drawn from well-defined probability spaces. A *random graph* is defined by endowing the set of all graphs with a given number of vertices with a probability measure, so that the presence or absence of edges becomes a stochastic variable. This can be formulated in different but related ways: one may fix a probability for edges to appear independently between each pair of vertices; alternatively, one may fix the total number of edges and assign equal weight to all graphs compatible with this constraint; or, finally, one may prescribe a *degree distribution* for the vertices, denoted $p_{\deg}(k)$, and construct graphs consistent with it. Each of these formulations highlights a different aspect

of randomness in networks, but all share the central idea of characterizing the large-scale properties of a typical graph by probabilistic means.

A particularly relevant class for applications in statistical physics is that of sparse random graphs, in which the average degree remains finite as the number of vertices grows. In this regime, unlike in dense graphs where every vertex is connected to a finite fraction of the system, local neighborhoods remain finite even in the thermodynamic limit.

In this thesis we will focus mostly on the standard ensembles of sparse random graphs most commonly used in statistical physics and network science, such as Erdős-Rényi random graphs with fixed edge probability, random regular graphs with fixed degree, and configuration models with degree distributions of finite second moment. A key property of these ensembles is that they are *locally tree-like*. More precisely, let us consider a graph with N vertices and select a vertex uniformly at random. The ball of radius r around this vertex contains with probability tending to one (as $N \to \infty$) no loops whenever r is fixed. In other words, in the thermodynamic limit any finite neighborhood looks like a tree, and cycles become asymptotically rare: the number of short loops of any given finite length remains of order one, typically following a Poisson distribution [5]. This property is crucial for the analytical tractability of many models, as it underpins message-passing methods such as the cavity approach.

At finite sizes, however, deviations from the idealized thermodynamic behavior are significant. Real networks often exhibit a high density of short loops, which can strongly affect the behavior of dynamical processes. Moreover, finite-size effects can lead to substantial fluctuations in local properties, such as degree correlations and clustering coefficients. Understanding how these finite-size corrections influence the performance of algorithms and the dynamics on networks is an important aspect of network science, with implications for both theoretical studies and practical applications.

1.3 Main ensembles of sparse random graphs

As discussed above, sparse random graphs are characterized by the locally tree-like property, which is central to the analytical tractability of many models. In the following, we briefly recall the main ensembles that will be used throughout this thesis, together with their most relevant structural features.

- **Random Regular Graphs (RRGs).** Graphs chosen uniformly at random among all K-regular graphs with N vertices. All nodes have degree K, the ensemble is homogeneous, and local neighborhoods are tree-like as $N \to \infty$. The degree distribution is a Kronecker

delta function concentrated in $k = K$, i.e. $p_{\text{deg}}(k) = \delta_{k,K}$ [1]. RRGs are often used as a baseline model in spin glasses and dynamical processes on networks.

- **Erdős-Rényi Graphs (ERGs).** In the $G(N,p)$ ensemble, each possible edge is present independently with probability p (or exactly M edges in the $G(N,M)$ variant, where $M = pN(N-1)/2$). Degrees are binomially distributed, i.e.

$$p_{\text{deg}}(k) = \binom{N-1}{k} p^k (1-p)^{N-1-k}.$$

  ERGs are locally tree-like in the thermodynamic limit [1, 3, 4].

- **Configuration Model (CMs).** Graphs constructed by assigning to each vertex a number of stubs drawn from a prescribed degree distribution $p_{\text{deg}}(k)$, then matching them uniformly at random. If the second moment $\langle k^2 \rangle$ is finite, local neighborhoods are tree-like [2, 3]. The CM allows one to model heterogeneity while retaining analytic tractability.

- **Scale-free Networks (SFNs).** A special case of the CM with power law degree distribution $p_{\text{deg}}(k) \sim k^{-\gamma}$. For $2 < \gamma \leq 3$, the second moment diverges, producing hubs that strongly affect percolation and dynamics. These heavy-tailed networks are commonly observed in real-world systems. In this case the locally tree-like structure breaks down due to the presence of hubs even in the thermodynamic limit [2, 3].

- **Watts-Strogatz (WS) model.** Starting from a regular ring lattice, each edge is rewired with probability β. The WS model interpolates between regular lattices ($\beta = 0$) and random graphs ($\beta = 1$), producing small-world networks with short average path lengths for intermediate β. The WS model is not locally tree-like due to the high density of short loops [2, 3].

These ensembles have been chosen because they cover a broad range of structures relevant to the study of stochastic dynamics on networks: homogeneous and sparse graphs that are locally tree-like (RRGs, ERGs), as well as heterogeneous or partially clustered networks where short loops appear (scale free networks, Watts-Strogatz), allowing us to explore the effects of network structure on dynamical processes.

Chapter 2

The Cavity Method: From Statics to Dynamics

The cavity method is a cornerstone of statistical physics, providing a powerful framework to analyze high-dimensional probabilistic systems defined on networks. Its origins trace back to the Bethe approximation in the study of spin systems on lattices, later formalized through message-passing techniques and the development of belief propagation in computer science. Over time, the method has been extended from static problems, such as computing equilibrium properties of spin glasses, to dynamical processes on networks, where it enables the study of stochastic dynamics in complex systems.

This presentation is deliberately concise and non-exhaustive; for a more thorough exposition, the reader is referred to Mézard and Montanari [6], Mézard, Parisi and Virasoro [7], and Charbonneau et al. [8].

2.1 Graphical models

Many complex systems in statistical physics are characterized by local interactions. In such cases, their probability distributions can often be represented through *graphical models*, which provide a concise framework to encode the dependencies among random variables. This representation is particularly useful for the application of the cavity method, as it makes the structure of interactions explicit.

Let us consider a system composed of N variables $\underline{x} = \{x_i\}_{i=1}^N$, each taking values in a set $\mathscr{X}$, which may be discrete or continuous (for example, $\mathscr{X} = \{-1, +1\}$ in the case of Ising

spins). The joint probability distribution of these variables can typically be written in the factorized form

$$P(\underline{x}) = \frac{1}{Z} \prod_{a \in F} \psi_a(\underline{x}_{\partial a}), \tag{2.1}$$

where Z is the normalization constant (partition function), ψ_a denotes the factor functions encoding the interactions, and $\underline{x}_{\partial a}$ represents the subset of variables associated with factor $a \in F$, with F denoting the set of all factors.

The structure of these interactions can be represented by a bipartite graph $G = (V \cup F, E)$, known as a *factor graph*. Here, $V = \{1, 2, \ldots, N\}$ indexes the variable nodes, F the factor nodes, and the edges E connect variable i to factor a whenever the variable i is involved in the factor ψ_a, i.e.

$$E = \{(i,a) \mid i \in V,\ a \in F,\ x_i \in \underline{x}_{\partial a}\}.$$

As an illustration, consider a higher-order Ising model composed of three spins $\underline{x} = \{x_1, x_2, x_3\}$ with $x_i \in \{-1, +1\}$. The system is characterized by an external field h_1 acting on spin 1, a pairwise coupling J_{23} between spins 2 and 3, and a three-body interaction J_{123} involving all three spins. The Hamiltonian is given by

$$H(\underline{x}) = -h_1 x_1 - J_{23} x_2 x_3 - J_{123} \left(2\delta(x_1, x_2, x_3) - 1\right), \tag{2.2}$$

where $\delta(x_1, \ldots, x_n)$ is the Kronecker delta for an arbitrary number of binary arguments. The joint probability distribution naturally factorizes according to these interaction terms:

$$P(\underline{x}) = \frac{1}{Z} \psi_1(x_1)\, \psi_{23}(x_2, x_3)\, \psi_{123}(x_1, x_2, x_3), \tag{2.3}$$

where the factors are defined as

$$\psi_1(x_1) = e^{\beta h_1 x_1}, \tag{2.4}$$

$$\psi_{23}(x_2, x_3) = e^{\beta J_{23} x_2 x_3}, \tag{2.5}$$

$$\psi_{123}(x_1, x_2, x_3) = e^{\beta J_{123}(2\delta(x_1, x_2, x_3) - 1)}. \tag{2.6}$$

The corresponding factor graph representation is shown in Fig. 2.1.

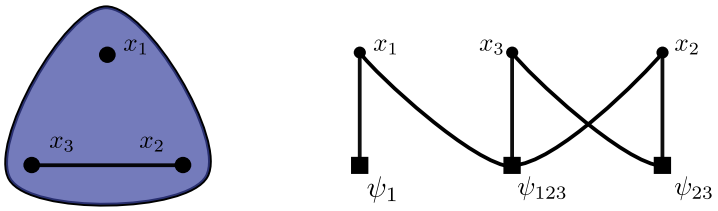

Fig. 2.1 **Graphical representation of a higher-order Ising model. Left panel**: Pictorial representation of the higher-order interactions of the model with Hamiltonian Eq. (2.2). **Right panel**: Factor graph representation of the joint distribution of the higher-order Ising model Eq. (2.3). Variable nodes (circles) represent the spins x_i, while factor nodes (squares) encode the interactions: the single-node factor ψ_1 (field h_1), the pairwise factor ψ_{23} (coupling J_{23}), and the hyperedge factor ψ_{123} (three-body coupling J_{123}).

2.2 Static cavity method

A common challenge in probabilistic models defined on graphs is the evaluation of marginal distributions of individual variables, or equivalently the computation of the partition function. For general models, this task is computationally intractable, as the number of configurations grows exponentially with the system size. Yet, if the underlying interaction network has a tree structure, one can take advantage of the locality of the interactions to perform the calculation efficiently.

This observation is at the foundation of the cavity method, first formulated in statistical physics by Bethe in the study of spin systems [9], known as *Bethe-Peierls approximation*. Variants of the same idea were later rediscovered in other disciplines, where it became known as Belief Propagation (BP): in information theory it underpins modern coding schemes such as low-density parity check codes [10], while in computer science it was proposed in the context of Bayesian inference by Pearl [11]. Despite the differences in language and applications, all these perspectives rely on the same principle of iterative message passing across the graph.

In this section, we introduce the static version of the cavity method, where no explicit notion of time is involved and variables are assumed to be static. In what follows, we will use the terms *cavity method* and *Belief Propagation* interchangeably.

We consider a system of N variables $\underline{x} = \{x_i\}_{i=1}^N$ with factorized joint distribution as described in Eq. (2.1). For each edge $(i,a) \in E$ of the factor graph, we introduce two types of messages: $c_{i\setminus a}(x_i)$ from variable node i to factor node a, and $c_{a\setminus i}(x_i)$ from factor node a to variable node i. Messages are probability distributions over the single-variable space $\mathscr{X}$, i.e. they are non-negative and normalized. Intuitively, $c_{i\setminus a}(x_i)$ represents the marginal distribution

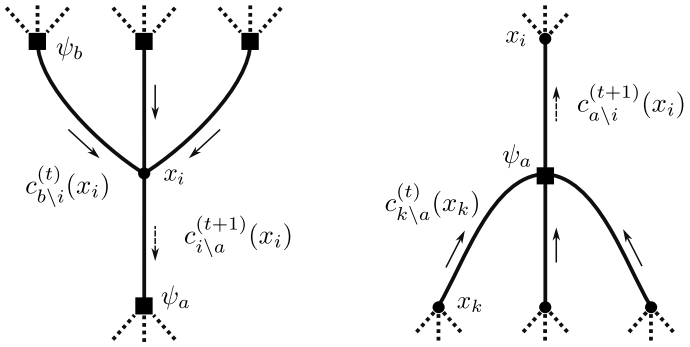

Fig. 2.2 **Pictorial representation of the Belief Propagation update rules. Left panel**: Portion of the factor graph involved in the update rule Eq. (2.7) for the variable-to-factor message. **Right panel**: Portion of the factor graph involved in the update rule Eq. (2.8) for the factor-to-variable message. Variable-to-factor messages are represented by dashed arrows, while factor-to-variable messages are represented by solid arrows.

of variable x_i in the cavity graph where factor a has been removed, while $c_{a\setminus i}(x_i)$ represents the marginal of x_i in the cavity graph where all factors $b \in \partial i$ except a have been removed.

The BP algorithm provides a way to update these messages iteratively using only local information, thereby enabling efficient computation of marginals in large graphical models. Assuming the messages are initialized at iteration $t = 0$ (e.g., uniformly), the *BP update rules* (or *cavity update rules*) at iteration $t+1$ read:

$$c^{(t+1)}_{i\setminus a}(x_i) \propto \prod_{b\in\partial i\setminus a} c^{(t)}_{b\setminus i}(x_i), \tag{2.7}$$

$$c^{(t+1)}_{a\setminus i}(x_i) \propto \sum_{\underline{x}_{\partial a\setminus i}} \psi_a(\underline{x}_{\partial a}) \prod_{k\in\partial a\setminus i} c^{(t)}_{k\setminus a}(x_k), \tag{2.8}$$

where the sum $\sum_{\underline{x}_{\partial a\setminus i}}$ runs over all configurations of the neighboring variables of factor a except i, i.e. $\underline{x}_{\partial a\setminus i} = \{x_k : k \in \partial a \setminus i\}$. If the single-variable space $\mathscr{X}$ is discrete, this is a sum; if $\mathscr{X}$ is continuous, the sum should be replaced by an integral. A pictorial representation of these update rules is shown in Fig. 2.2.

2.2.1 Cavity method on tree graphs

If the factor graph is a tree, i.e., it contains no cycles, the BP algorithm is guaranteed to converge to the exact marginal distributions in a finite number of iterations, regardless of the initialization [6]. The fixed point messages, denoted as $c_{i\setminus a}(x_i)$ and $c_{a\setminus i}(x_i)$, can be obtained by iterating the BP update rules until convergence. These messages satisfy the *BP equations* (or *cavity equations*), which are obtained by removing the iteration index t from Eqs. (2.7) and (2.8).

Once convergence is achieved, the messages can be used to compute the marginal distributions of the variables in the original graphical model. This is done by combining the incoming messages from neighboring factors according to the structure of the factor graph:

$$p_i(x_i) = \sum_{\underline{x}_{\setminus i}} P(\underline{x}) = \frac{1}{Z_i} \prod_{a \in \partial i} c_{a \setminus i}(x_i), \tag{2.9}$$

where $\underline{x}_{\setminus i}$ denotes all variables except x_i, and Z_i is a normalization constant ensuring that $p_i(x_i)$ sums to one over all possible values of x_i.

More generally, for tree-structured graphical models, the marginal distribution of any subset of variables can be written in terms of the fixed-point messages. Let $U \subseteq V$ be a subset of variable nodes. Define $\partial U = \{a \in F \mid \partial a \cap U \neq \emptyset\}$ as the set of factor nodes connected to at least one variable in U, and $F_U = \{a \in \partial U \mid \partial a \subseteq U\}$ as the set of factor nodes whose neighbors are all in U. For each boundary factor $a \in \partial U \setminus F_U$, there is a unique variable $i(a) \in U$ such that $i(a) \in \partial a \cap U$. The marginal distribution of the variables in U is then given by:

$$p_U(\underline{x}_U) = \sum_{\underline{x}_{\setminus U}} P(\underline{x}) = \frac{1}{Z_U} \prod_{a \in F_U} \psi_a(\underline{x}_{\partial a}) \prod_{b \in \partial U \setminus F_U} c_{b \setminus i(b)}(x_{i(b)}), \tag{2.10}$$

where $\underline{x}_{\setminus U}$ denotes all variables except those in U, and Z_U is a normalization constant ensuring that $p_U(\underline{x}_U)$ sums to one over all possible values of $\underline{x}_U$.

Furthermore, the free energy of the system can be expressed as a sum of local contributions from variables and factors, minus corrections for double counting. Specifically, the free energy F can be written as a function of the $2|E|$ fixed point messages $\underline{c} = \{c_{a \setminus i}, c_{i \setminus a}\}_{(i,a) \in E}$, known as the *Bethe free energy*:

$$\mathbb{F}(\underline{c}) = \sum_{a \in F} \mathbb{F}_a(\underline{c}) + \sum_{i \in V} \mathbb{F}_i(\underline{c}) - \sum_{(i,a) \in E} \mathbb{F}_{ia}(\underline{c}), \tag{2.11}$$

where

$$-\beta \mathbb{F}_a(\underline{c}) = \log Z_a = \log \sum_{\underline{x}_{\partial a}} \psi_a(\underline{x}_{\partial a}) \prod_{i \in \partial a} c_{i \setminus a}(x_i), \tag{2.12}$$

$$-\beta \mathbb{F}_i(\underline{c}) = \log Z_i = \log \sum_{x_i} \prod_{b \in \partial i} c_{b \setminus i}(x_i), \tag{2.13}$$

$$-\beta \mathbb{F}_{ia}(\underline{c}) = \log \sum_{x_i} c_{i \setminus a}(x_i) c_{a \setminus i}(x_i). \tag{2.14}$$

Here, β denotes the inverse temperature, which allows the free energy to be expressed in physical units. However, the cavity method is not restricted to physical systems and can be applied to a wide range of probabilistic models. In cases where no notion of temperature is present, one simply sets $\beta = 1$.

2.2.2 Loopy cavity method

When the factor graph contains cycles, the BP algorithm can still be applied as a heuristic method, often referred to as *loopy belief propagation* (or *loopy cavity method*) [6]. In this case, the messages are updated iteratively using the same rules as in the tree case Eqs. (2.7) and (2.8), but convergence is not guaranteed. Moreover, even if the algorithm converges, the resulting messages may not correspond to the exact marginals of the original graphical model. Nevertheless, loopy BP has been found to perform well in practice for many applications, providing good approximations of marginal distributions and being computationally efficient.

In the case of large sparse random graphs with a locally tree-like structure, loopy BP typically provides accurate approximations, and in the thermodynamic limit $N \to \infty$ it can even become asymptotically exact. However, the locally tree-like property, while necessary, is not sufficient to guarantee convergence of the algorithm or correctness of the fixed point. A general theoretical framework for convergence and correctness of loopy BP is still lacking, although significant progress has been made in particular cases [12–17]. Intuitively, convergence and correctness are more likely when the removal of any variable node i (i.e. creating a cavity) results in its neighboring factors $a \in \partial i$ being, with high probability, far apart from one another.

Finally, loopy BP can also be understood from a variational perspective: the stationary points of the Bethe free energy Eq. (2.11) $\mathbb{F}(\underline{c})$ correspond to fixed points of the cavity equations, and vice versa. This variational interpretation provides additional insight into the nature of the solutions obtained by the algorithm.

2.2.3 Pairwise interactions

A particularly relevant case in statistical physics is that of systems with pairwise interactions. In this setting, the factor graph simplifies considerably, since each factor node connects exactly two variable nodes. The joint distribution can be written as

$$P(\underline{x}) = \frac{1}{Z} \prod_{(i,j)\in E} \psi_{ij}(x_i, x_j) \prod_{i \in V} \psi_i(x_i), \tag{2.15}$$

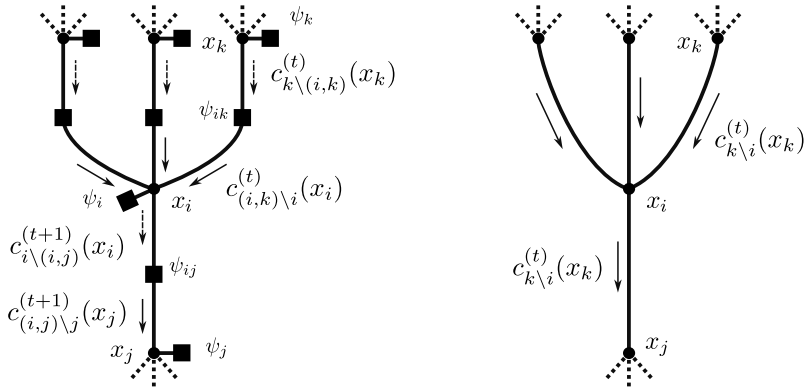

Fig. 2.3 **Pictorial representation of the cavity method for pairwise interactions. Left panel**: Original factor graph representation, where factor nodes (squares) connect pairs of variable nodes (circles). Variable-to-factor messages are represented by dashed arrows, while factor-to-variable messages are represented by solid arrows. **Right panel**: Simplified representation, where only variable nodes are shown, and edges represent pairwise interactions. The new variable-to-variable messages are represented by solid arrows, as they are a shorthand notation for the original factor-to-variable messages.

where Z is the partition function ensuring normalization, $\psi_{ij}(x_i,x_j)$ encodes the interaction between variables x_i and x_j for each edge $(i,j) \in E$, and $\psi_i(x_i)$ represents any local bias or external field acting on variable x_i. In this pairwise case, each factor node in the factor graph corresponds either to an edge (for ψ_{ij}) or to a node (for ψ_i) of the original interaction network.

Because each factor involves only two variables, the factor graph can be represented more compactly by retaining only variable nodes, with edges denoting pairwise interactions. In this representation one of the two messages per edge can be eliminated, leading to a simplified update rule. In what follows, we adopt a notation in terms of variable-to-variable messages, writing $c_{i\setminus j}(x_i)$ as shorthand for $c_{i\setminus(i,j)}(x_i)$. The BP update equations then take the form

$$c^{(t+1)}_{i\setminus j}(x_i) \propto \psi_i(x_i) \prod_{k\in\partial i\setminus j} \sum_{x_k\in\mathscr{X}} \psi_{ik}(x_i,x_k) c^{(t)}_{k\setminus i}(x_k). \tag{2.16}$$

A schematic representation of the BP updates in the pairwise case is shown in Fig. 2.3. The left panel depicts the original factor graph, with factor nodes (squares) connecting pairs of variable nodes (circles). The right panel illustrates the simplified representation, where edges encode pairwise interactions and messages are represented directly between variable nodes.

2.3 Dynamic cavity method for discrete-time processes

The idea of extending the cavity method to dynamics dates back to the pioneering works of Neri and Bollé on parallel Glauber dynamics [18], Kanoria and Montanari on majority dynamics [19], and Aurell and Mahmoudi on diluted kinetic Ising models [20]. These studies showed that the method yields exact recursion relations on trees, but also highlighted two major obstacles: the exponential growth of trajectory space with time and the lack of closure due to the dependence on neighboring trajectories. These difficulties motivate the use of auxiliary representations and suitable approximations, which we now illustrate through the mathematical formulation of the method.

2.3.1 Mathematical formulation

The dynamic cavity method treats entire *trajectories* of the variables as the relevant degrees of freedom. We consider here Markovian dynamics in discrete time, where the state of each variable at step $n+1$ depends only on its own state and those of its neighbors at step n. Let $\underline{x}^n = \{x_i^n\}_{i=1}^N$ denote the configuration of the system at time n. Given the joint probability distribution $P(\underline{x}^n)$, the distribution at the next step is obtained by applying the transition probability $W(\underline{x}^{n+1}|\underline{x}^n)$:

$$P(\underline{x}^{n+1}) = \sum_{\underline{x}^n} W(\underline{x}^{n+1}|\underline{x}^n) P(\underline{x}^n). \tag{2.17}$$

Introducing the trajectory vector $\boldsymbol{x}_i = (x_i^0, x_i^1, \ldots, x_i^T)$ of node i, and denoting by $\underline{\boldsymbol{x}} = \{\boldsymbol{x}_i\}_{i=1}^N$ the collection of all trajectories, the joint probability distribution of trajectories up to horizon T reads

$$P(\underline{\boldsymbol{x}}) = \prod_{n=1}^{T} W(\underline{x}^n|\underline{x}^{n-1}) P(\underline{x}^0). \tag{2.18}$$

For simplicity, let us assume independent initial conditions:

$$P(\underline{x}^0) = \prod_{i=1}^{N} p_i^0(x_i^0),$$

where $p_i^0(x_i^0)$ is the initial distribution of variable i.

As in the static case, we focus on pairwise interactions on a graph $G = (V, E)$, for which the transition probability factorizes over nodes:

$$W(\underline{x}^{n+1}|\underline{x}^n) = \prod_{i \in V} w_i(x_i^{n+1}|x_i^n, \underline{x}_{\partial i}^n), \tag{2.19}$$

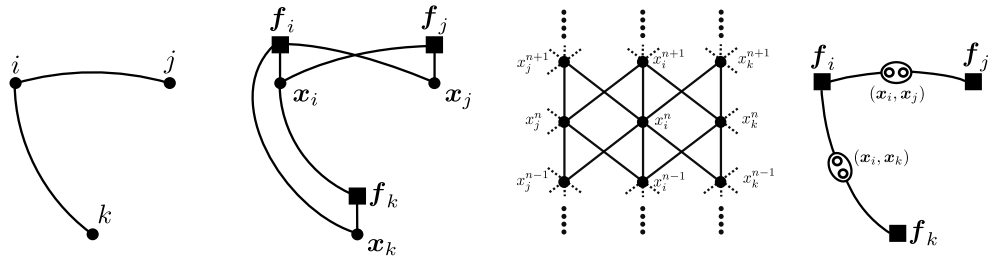

Fig. 2.4 **Dynamic factor graph representation of the system. Panel a:** Original contact graph G. **Panel b:** Naive factor graph associated to Eq. (2.20), which contains many short loops due to temporal dependencies. **Panel c:** Origin of these loops: each local transition couples neighbors at successive times. **Panel d:** Dual factor graph representation, where loops are removed at the price of enlarging variable nodes to contain the trajectories of pairs of neighbors.

where $w_i(x_i^{n+1}|x_i^n, \underline{x}_{\partial i}^n)$ is the local transition probability of i, depending on the state of node i and its neighbors $\underline{x}_{\partial i}$ at time n. The joint trajectory distribution therefore factorizes both over nodes and over timesteps:

$$P(\underline{\boldsymbol{x}}) = \prod_{i \in V} p_i^0(x_i^0) \prod_{n=1}^{T} w_i(x_i^{n+1}|x_i^n, \underline{x}_{\partial i}^n). \tag{2.20}$$

Even when G is locally tree-like, the factor graph of Eq. (2.20) contains many short loops caused by the temporal structure. This is illustrated in panels (b) and (c) of Fig. 2.4. As shown in [21, 22], such loops can be removed by introducing a *dual factor graph*, sketched in panel (d) of Fig. 2.4. One defines a factor f_i for each node $i \in G$, equal to the product over all timesteps of the local transition probabilities, and new variable nodes for each edge $(i, j) \in E$, containing the trajectories of the two neighboring variables. In this representation, the topology of the dual graph mirrors the original interaction graph.

The cavity method can then be applied to the dual graph, yielding recursive equations for messages defined as distributions over pairs of trajectories:

$$c_{i\backslash j}^{(t+1)}(\boldsymbol{x}_i, \boldsymbol{x}_j) \propto \sum_{\underline{\boldsymbol{x}}_{\partial i\backslash j}} f_i(\boldsymbol{x}_i, \boldsymbol{x}_j, \underline{\boldsymbol{x}}_{\partial i\backslash j}) \prod_{k \in \partial i\backslash j} c_{k\backslash i}^{(t)}(\boldsymbol{x}_i, \boldsymbol{x}_k). \tag{2.21}$$

At this stage the main limitation becomes apparent: the number of joint trajectories $(\boldsymbol{x}_i, \boldsymbol{x}_j)$ grows exponentially with T, making exact computations infeasible beyond very short transients. Moreover, the messages Eq. (2.21) depend parametrically on the trajectory $\boldsymbol{x}_i$. This originates from the dual construction, where each variable must be duplicated across all its incident edges,

preventing the equations from closing in terms of single-node trajectory distributions. These two issues – exponential growth of trajectory space and parametric dependence on neighbors – are the core difficulties of the dynamic cavity method.

Several approximations have been proposed to make the equations tractable. The *time factorization ansatz* [20, 23] assumes that messages factorize over time, yielding efficient update rules for single-time marginals. More sophisticated schemes include finite order *Markovian closures* [22] and *matrix product parametrizations* of trajectory distributions [24, 25], which provide good approximations to temporal correlations at the cost of considerably more involved mathematics. For *irreversible spreading processes*, such as cascades or epidemic models, the problem simplifies dramatically since trajectories can be parametrized by activation times, leading to exact dynamic message-passing equations on trees [26, 27].

Most of these works are formulated in discrete time, and often for discrete variables. A continuous-time extension was developed by Aurell and collaborators [28], who applied the cavity method directly to the master equation, obtaining a *cavity master equation* for discrete spins with asynchronous updates.

In this thesis we will build on these developments, but with a different perspective: our goal is to devise a framework that remains simple enough for practical use, while being general enough to encompass both discrete-time/discrete-variable systems and continuous-time/continuous-variable dynamics. This will be achieved through a *small-coupling expansion* of the dynamic cavity equations, which provides a controlled approximation while avoiding the technical complexity of the more elaborate approaches described above.

2.4 Cavity method with quenched disorder

In many relevant situations one does not know the precise form of the interactions, but only their statistical properties. This motivates the study of ensembles in which the interaction graph and the couplings are drawn at random once and remain *quenched* throughout the evolution. In such cases the central object of interest is not the set of messages for a single instance, but rather their typical distribution across the ensemble, which characterizes the large-scale properties of the system in the thermodynamic limit.

To illustrate the method, we focus on the simplest case of pairwise interactions in equilibrium systems. For a single instance on a tree, the cavity recursion for the message $c_{i\setminus j}(x_i)$ reads

$$c_{i\setminus j}(x_i) \propto \psi_i(x_i;h_i) \prod_{k\in\partial i\setminus j} \sum_{x_k} \psi_{ik}(x_i,x_k;J_{ik})\, c_{k\setminus i}(x_k) \tag{2.22a}$$

$$= \mathscr{F}\left(\{c_{k\setminus i}\}_{k\in\partial i\setminus j};h_i,\{J_{ik}\}_{k\in\partial i\setminus j}\right), \tag{2.22b}$$

where ψ_i and ψ_{ik} denote respectively the local and pairwise interaction factors, where we explicited the dependence on the external fields h_i and couplings J_{ik}. We denoted the right-hand side of this relation as an update map $\mathscr{F}$.

When the system is defined on a random-graph ensemble with quenched disorder in the couplings, the *replica symmetric (RS) cavity method* consists in studying the distribution $P(c)$ of cavity messages rather than individual messages. This distribution is the fixed point of a distributional recursion,

$$P(c) = \sum_k p_k\, \mathbb{E}_{h,J}\left[\int \prod_{l=1}^{k-1} \mathscr{D}P(c_l)\, \delta\left(c - \mathscr{F}\left(\{c_l\}_{l=1}^{k-1};h_l,\{J_l\}_{l=1}^{k-1}\right)\right)\right], \tag{2.23}$$

where p_k is the degree distribution of the random graph, and the expectation $\mathbb{E}_{h,J}$ averages over the quenched disorder in the system, i.e., the random external fields h and couplings J that enter the update map $\mathscr{F}$. Here, $\mathscr{D}P(c_l)$ denotes integration (or summation) over the distribution $P(c)$ of messages, and δ is the generalized delta function enforcing the update constraint Eq. (2.22a). The recursion Eq. (2.23) is in general analytically intractable, but it admits a numerical solution via population dynamics: one represents $P(c)$ by a large sample $\{c^\alpha\}_{\alpha=1}^M$ of messages; at each iteration a random degree k is drawn from p_k, $k-1$ incoming messages are sampled from the current population and the disorder parameters are drawn from their distribution, a new message is computed through $\mathscr{F}$ and renormalized, and then replaces a randomly-chosen element of the population.

The replica symmetric assumption is valid in some regimes but can fail in others. In the *replica symmetric* (RS) phase, the distribution of cavity messages converges to a unique fixed point, and belief propagation provides asymptotically exact predictions. At the *dynamical one-step replica symmetry breaking* (d1RSB) transition, the space of solutions becomes clustered: the RS fixed point still exists, but the system organizes into an exponential number of distant clusters of solutions. In this regime the RS cavity method gives correct averages, but local iterative algorithms such as BP typically fail to converge to the relevant fixed point when initialized randomly. A further transition may lead to a *static one-step replica symmetry breaking* (s1RSB) phase, where the RS description itself becomes inconsistent and must be replaced by more refined methods such as the 1RSB cavity method or survey propagation

[6, 29]. The terminology distinguishes the two situations: the d1RSB transition is revealed dynamically, as it affects the behavior of iterative algorithms, while the s1RSB transition is of static nature, as it signals the failure of the RS description at the level of ensemble-averaged quantities.

Chapter 3

Path-integral approaches

The cavity method provides a powerful framework for studying stochastic processes on graphs, both in equilibrium and out of equilibrium, and is particularly well-suited for sparse structures. However, it is not the only tool available. An alternative and complementary viewpoint is provided by path-integral approaches, which describe stochastic dynamics in terms of generating functionals and field-theoretic techniques. These methods, rooted in statistical physics, have been developed to treat interacting stochastic systems with a large number of degrees of freedom and are especially effective in the presence of quenched disorder and in the fully-connected limit.

The basic idea is to reformulate the stochastic dynamics, typically given in the form of Langevin or master equations, as a functional integral over trajectories. This representation makes it possible to compute dynamical averages and correlations systematically, and to perform approximations inspired by field theory, such as saddle point evaluations or systematic expansions. Path-integral approaches thus provide the natural setting for the derivation of dynamical mean-field equations and of perturbative schemes such as the Plefka expansion. At the same time, they also make transparent the limitations that arise when dealing with sparse interaction networks.

In this chapter we present the main ingredients of the path-integral approach to stochastic dynamics. We begin in section 3.1 by recalling the description of stochastic dynamics through Langevin equations. In section 3.2 we introduce the *Martin-Siggia-Rose-Janssen-De Dominicis (MSRJD) formalism*, which provides the functional integral representation of stochastic processes. We then discuss how quenched disorder can be incorporated in this framework, showing in section 3.3 how the *Dynamical Mean-Field Theory (DMFT)* emerges in the fully-connected limit. In section 3.4 we introduce the Plefka expansion and its dynamical generalization, and

finally, in section 3.5, we comment on the limitations of these approaches when applied to sparse graphs.

3.1 Stochastic dynamics and Langevin equations

A standard framework to describe the time evolution of stochastic systems with continuous variables is provided by *Langevin equations*[1], which represent the dynamics of degrees of freedom evolving in continuous time under the combined action of deterministic forces and stochastic fluctuations [30, 31]. This formalism is widely used in statistical physics to model systems where randomness originates from an effective interaction with an environment, and it provides a natural starting point for field-theoretic representations of stochastic dynamics.

Consider a system of N interacting degrees of freedom $x_i(t)$, whose dynamics is governed by

$$\frac{dx_i(t)}{dt} = f_i(\underline{x}(t)) + \xi_i(t), \tag{3.1}$$

where $f_i(\underline{x}(t))$ is the deterministic drift term, depending in general on the full configuration $\underline{x}(t) = \{x_j(t)\}_{j=1}^N$, and $\xi_i(t)$ is a stochastic noise term. The most common choice is Gaussian white noise with zero mean and covariance

$$\langle \xi_i(t) \rangle = 0, \qquad \langle \xi_i(t)\, \xi_j(t') \rangle = D_{ij}\, \delta(t-t'), \tag{3.2}$$

where D_{ij} is the noise covariance matrix. In the simplest case one assumes $D_{ij} = 2T\, \delta_{ij}$, with T playing the role of an effective temperature. More general choices, with non-diagonal D_{ij}, allow for correlated noise sources across different degrees of freedom. It is worth noting that the stochastic process defined by Eqs. (3.1) and (3.2) is Markovian, as the future evolution depends only on the current state $\underline{x}(t)$ and not on the past history. If instead the noise has temporal correlations, i.e. $\langle \xi_i(t)\, \xi_j(t') \rangle = D_{ij}(t-t')$ with a non-zero correlation time, the process becomes non-Markovian. Such noises are usually referred to as *colored noises*. Another important distinction is between *additive noise*, as in Eq. (3.1), where the noise term is independent of the state $\underline{x}(t)$, and *multiplicative noise*, where the noise amplitude depends on the state.

[1] In this thesis we shall use the term "Langevin equation" to refer specifically to the overdamped case. Thinking of the classical example of a particle diffusing in a fluid subject to friction, the viscous term dominates the inertial one and the acceleration $\propto d^2x/dt^2$ can be neglected. Over sufficiently long-time-scale inertial effects are irrelevant in the presence of dissipation, and the overdamped description is appropriate.

To make the discussion concrete, consider the case of a single variable $x(t)$ evolving according to

$$\frac{dx(t)}{dt} = -\lambda\, x(t) + \xi(t), \tag{3.3}$$

with $\langle \xi(t) \rangle = 0$ and $\langle \xi(t)\, \xi(t') \rangle = 2T\, \delta(t-t')$. This is the *Ornstein-Uhlenbeck (OU)*, the simplest non-trivial Langevin dynamics. The drift term $f(x) = -\lambda x$ drives the system toward the stable fixed point at $x = 0$, while the noise term induces fluctuations around it. The competition between the two terms results in a stationary Gaussian distribution with variance T/λ, which provides a simple illustration of how Langevin dynamics generates non-trivial stationary states.

Equations of the form Eqs. (3.1) and (3.2) define the class of continuous-time Markov processes that will serve as the starting point for the path-integral representations introduced in the next section. While extensions to multiplicative or non-Gaussian noise are possible, here we restrict ourselves to the additive Gaussian case, which already encompasses a wide range of models and admits a transparent functional formulation.

3.2 Path-integral formulation

Langevin equations provide a trajectory-based description of stochastic dynamics, but computing averages and correlations directly from them can be cumbersome. A powerful alternative is offered by the path-integral representation of stochastic processes, originally introduced by Martin and Siggia [32] and subsequently reformulated by Rose, Janssen [33], De Dominicis and Peliti [34, 35], and now commonly known as the MSRJD formalism. The essence of the method is to enforce the stochastic dynamics through functional delta constraints, thereby turning averages over trajectories into functional integrals from which correlation and response functions can be systematically obtained.

To keep the presentation concrete, let us focus on the OU process already introduced in section 3.1,

$$\frac{dx(t)}{dt} = -\lambda x(t) + \xi(t), \qquad \langle \xi(t) \rangle = 0, \qquad \langle \xi(t)\, \xi(t') \rangle = 2T\, \delta(t-t'). \tag{3.4}$$

For later convenience, time is discretized as $t_n = n\Delta$ with $n = 0, \ldots, T$ and $T\Delta = t_f$. The trajectories $x(t)$ and $\xi(t)$ are then represented as discrete-time vectors $\boldsymbol{x} = (x^0, \ldots, x^T)$ and

$\boldsymbol{\xi} = (\xi^0, \ldots, \xi^T)$, respectively. In the Itô convention the discretized dynamics reads

$$x^{n+1} = x^n - \Delta\lambda x^n + \Delta\xi^n, \tag{3.5}$$

where the noise increments $\Delta\xi^n = \int_{n\Delta}^{(n+1)\Delta} dt\, \xi(t)$ are independent Gaussian random variables with zero mean and variance $2T\Delta$. The corresponding path probability of a noise history is therefore

$$P[\boldsymbol{\Delta\xi}] \propto \prod_{n=0}^{T} \exp\left[-\frac{(\Delta\xi^n)^2}{4T\Delta}\right]. \tag{3.6}$$

The average of any trajectory observable $\mathscr{O}(\boldsymbol{x})$ is obtained by integrating over all paths $\boldsymbol{x}$ and noise realizations $\boldsymbol{\Delta\xi}$, weighted by $P[\boldsymbol{\Delta\xi}]$ and constrained by Eq. (3.5). The dynamical constraint can be expressed as a product of Dirac deltas, each of which is represented in Fourier form by introducing conjugate variables $\hat{x}^n$. After integrating out the Gaussian noise variables, one arrives at the discrete-time expression

$$\langle \mathscr{O}(\boldsymbol{x}) \rangle \propto \int d\boldsymbol{x}\, d\hat{\boldsymbol{x}}\, \mathscr{O}(\boldsymbol{x})\, p_0(x^0) \exp\big[-S(\boldsymbol{x},\hat{\boldsymbol{x}})\big], \tag{3.7}$$

with action

$$S(\boldsymbol{x},\hat{\boldsymbol{x}}) = \sum_{n=0}^{T-1} \Big[\mathrm{i}\hat{x}^n\big(x^{n+1} - x^n + \Delta\lambda x^n\big) + \Delta T\,(\hat{x}^n)^2\Big]. \tag{3.8}$$

Passing to the continuum limit, the sums become time integrals and the finite differences converge to time derivatives, so that

$$S[x,\hat{x}] = \int_0^{t_f} dt\, \Big[\mathrm{i}\hat{x}(t)\big(\dot{x}(t) + \lambda x(t)\big) + T\,\hat{x}(t)^2\Big], \tag{3.9}$$

and averages of path observables $\mathscr{O}[x,\hat{x}]$ take the path integral form

$$\langle \mathscr{O}[x,\hat{x}] \rangle = \int \mathscr{D}[x,\hat{x}]\; \mathscr{O}[x,\hat{x}]\, p_0\big(x(0)\big)\, e^{-S[x,\hat{x}]}. \tag{3.10}$$

Here $\mathscr{D}[x,\hat{x}]$ denotes the functional integration over all trajectories $x(t)$ and auxiliary fields $\hat{x}(t)$, and $p_0(x)$ is the distribution of initial conditions.

Unless stated otherwise, we henceforth work directly in the continuum limit and omit the discretization step; the correspondence with the discrete- and continuous-time formulations is reviewed in Appendix A.1.1.

The *auxiliary field* $\hat{x}(t)$, commonly referred to as the response field, enforces the stochastic dynamics and naturally generates linear response functions. The MSRJD formalism thus

recasts the problem of stochastic dynamics into a field-theoretic language, where all statistical properties of trajectories are encoded in the *effective action* $S[x,\hat{x}]$. This representation will serve as the starting point for the subsequent developments of this chapter.

The construction above generalizes straightforwardly to systems of many interacting degrees of freedom with Langevin dynamics as in Eq. (3.1), leading to a path-integral representation where all statistical properties of the dynamics are encoded in an effective action. In practice, it is often convenient to summarize these properties in terms of the generating functional

$$Z[\psi] = \left\langle \exp\left(\int_0^{t_f} dt\, \psi(t) x(t) \right) \right\rangle, \tag{3.11}$$

from which all correlation functions can be obtained by functional differentiation with respect to the source $\psi(t)$. For the OU process, assuming uniform initial conditions, the quadratic nature of the action Eq. (3.9) allows one to evaluate the functional integral exactly, leading to

$$Z[\psi] = \exp\left[\frac{T}{2\lambda} \int_0^{t_f} dt \int_0^{t_f} dt'\, e^{-\lambda|t-t'|}\, \psi(t)\, \psi(t') \right], \tag{3.12}$$

A detailed derivation of this result can be found in Appendix A.1.3. While the OU dynamics is exactly solvable due to its linear drift and additive noise, the computation of $Z[\psi]$ is no longer straightforward in the presence of nonlinear drift terms, state-dependent noise, or many interacting degrees of freedom, where approximations or alternative methods become necessary.

Correlation functions The simplest observables accessible within the MSRJD framework are *correlation functions*. For the Ornstein-Uhlenbeck process, the two-time (disconnected) correlation function $C^{\mathrm{dc}}(t,t') = \langle x(t)\, x(t') \rangle$ is obtained by differentiating the generating functional Eq. (3.12) with respect to the source $\psi(t)$,

$$C^{\mathrm{dc}}(t,t') = \frac{\delta^2}{\delta\psi(t)\,\delta\psi(t')} Z[\psi]\bigg|_{\psi=0} = \frac{T}{\lambda} e^{-\lambda|t-t'|}, \tag{3.13}$$

which is the well-known result for the OU process. The path-integral formalism thus reproduces the stationary Gaussian statistics of $x(t)$, and more generally allows for systematic computation of higher-order correlations in interacting systems where direct solutions are not available.

Response functions A distinctive advantage of the MSRJD formalism is that it provides a natural way to compute *response functions*. Consider perturbing the Langevin equation

Eq. (3.4) by a weak external driving field $h(t)$ coupled linearly to $x(t)$,

$$\frac{dx(t)}{dt} = -\lambda x(t) + \xi(t) + h(t). \tag{3.14}$$

The linear response of an observable $\mathscr{O}[x]$ to the perturbation is defined as

$$R_{\mathscr{O}}(t,t') = \frac{\delta}{\delta h(t')} \langle \mathscr{O}[x] \rangle \Big|_{h=0}, \tag{3.15}$$

where the average has to be taken with the perturbed dynamics. Within the path-integral formulation, the external driving contributes an additional term $-\int dt\, \mathrm{i}\hat{x}(t)h(t)$ to the action Eq. (3.9). Differentiating with respect to $h(t')$ then amounts to inserting a factor $\mathrm{i}\hat{x}(t')$ in the functional average. Hence,

$$R_{\mathscr{O}}(t,t') = \langle \mathscr{O}[x]\, \mathrm{i}\hat{x}(t') \rangle. \tag{3.16}$$

This identity shows that the auxiliary field $\hat{x}(t)$, often called the response field, directly generates linear response functions. In other words, the average of $x(t)$ multiplied by $\hat{x}(t')$ measures how $x(t)$ responds to an infinitesimal perturbation applied at time t'. Intuitively, one can think of $\hat{x}(t)$ as generating an infinitesimal external "kick" of the dynamics. Multiplying an observable $\mathscr{O}[x]$ by $\hat{x}(t')$ corresponds to computing its change under a perturbation $h(t) = h\delta(t-t')$ in the Langevin equation–where h is the infinitesimal amplitude of the perturbation–since

$$\langle \mathscr{O}[x] \rangle = \int \mathscr{D}[x,\hat{x}]\, \mathscr{O}[x]\, e^{h\,\mathrm{i}\hat{x}(t') - S[x,\hat{x}]}, \tag{3.17}$$

and thus

$$\langle \mathscr{O}[x]\, \mathrm{i}\hat{x}(t') \rangle = \frac{\partial}{\partial h} \langle \mathscr{O}[x] \rangle \Big|_{h=0}. \tag{3.18}$$

This provides a precise meaning to the terminology "response field." For the OU process, the linear response function of $x(t)$ with respect to an impulse applied at t' is obtained from Eq. (3.16) with $\mathscr{O}[x] = x(t)$, yielding

$$R(t,t') = \langle x(t)\, \mathrm{i}\hat{x}(t') \rangle = \Theta(t-t')\, e^{-\lambda(t-t')}, \tag{3.19}$$

where Θ is the Heaviside step function. The exponential relaxation reflects the stability of the deterministic drift term and agrees with the solution of the directly perturbed Langevin equation. The derivations above have been presented in a compact form. More detailed calculations are collected in Appendix A.1.

Vanishing of averages of response fields A useful property of the MSRJD formalism is that averages involving only response fields vanish identically. This follows from the role of $\hat{x}$ as an auxiliary field enforcing the stochastic constraint. A simple way to see this is to consider the generating functional with source $\hat{\psi}(t)$ coupled to $\mathrm{i}\hat{x}(t)$

$$Z[\hat{\psi}] = \left\langle \exp\left(\int dt\, \hat{\psi}(t) \mathrm{i}\hat{x}(t) \right) \right\rangle . \tag{3.20}$$

The functional integral over $\hat{x}$ can be performed explicitly, yielding a delta constraint that enforces the Langevin equation with an additional external field $\psi(t)$. The value of the generating functional is therefore independent of $\hat{\psi}(t)$, being always equal to one due to the normalization of the path probability. Hence, all functional derivatives with respect to $\hat{\psi}(t)$ vanish,

$$\langle \mathrm{i}\hat{x}(t_1) \cdots \mathrm{i}\hat{x}(t_n) \rangle = 0. \tag{3.21}$$

A detailed derivation is provided in Appendix A.1.6.

3.3 Dynamical Mean-Field Theory for quenched disorder

The MSRJD formalism provides a natural framework to study stochastic systems with quenched disorder, where the interactions between degrees of freedom are random variables drawn from a specified distribution. A paradigmatic example is provided by fully-connected networks with random couplings, where each degree of freedom interacts with all others through couplings J_{ij} drawn from a specific distribution. In the thermodynamic limit $N \to \infty$, such systems can be analyzed using DMFT, which reduces the many-body problem to an effective single-site stochastic process subject to self-consistent noise and memory kernels.

To illustrate the DMFT construction, consider N linearly-coupled OU processes, each evolving according to the Langevin equation

$$\frac{dx_i(t)}{dt} = -\lambda x_i(t) + \sum_{j \neq i} J_{ij} x_j(t) + \xi_i(t), \tag{3.22}$$

where $\lambda > 0$ is a damping coefficient, $\xi_i(t)$ are independent Gaussian white noises with zero mean and correlations given by

$$\langle \xi_i(t)\, \xi_j(t') \rangle = 2T\, \delta_{ij}\, \delta(t - t'). \tag{3.23}$$

For each ordered pair (i,j), the couplings (J_{ij},J_{ji}) are jointly Gaussian with

$$\overline{J_{ij}} = \overline{J_{ji}} = \frac{m}{N}, \qquad \mathrm{Var}(J_{ij}) = \mathrm{Var}(J_{ji}) = \frac{\sigma^2}{N}, \qquad \mathrm{Corr}(J_{ij},J_{ji}) = \gamma. \tag{3.24}$$

The $1/N$ scaling ensures a well-defined thermodynamic limit. The parameter $\gamma \in [-1,1]$ controls the degree of symmetry: $\gamma = 1$ yields symmetric interactions ($J_{ij} = J_{ji}$), $\gamma = -1$ antisymmetric ($J_{ij} = -J_{ji}$), and $\gamma = 0$ fully asymmetric couplings.

3.3.1 Generating functional at fixed disorder

For a single realization of the couplings $\{J_{ij}\}_{i,j=1}^N$, we introduce the generating functional with sources $\psi_i(t)$ and $\hat{\psi}_i(t)$ coupled to $x_i(t)$ and $\mathrm{i}\hat{x}_i(t)$, respectively,

$$Z_J\left[\underline{\psi},\underline{\hat{\psi}}\right] = \left\langle \exp\left(\sum_{i=1}^N \int dt \left[\psi_i(t)\,x_i(t) + \hat{\psi}_i(t)\mathrm{i}\hat{x}_i(t)\right]\right)\right\rangle, \tag{3.25}$$

where the average is taken over the noise realizations $\underline{\xi}$. The generating functional allows one to compute any correlation or response function by functional differentiation with respect to the sources. Using the MSRJD formalism, Z_J can be expressed as a path integral over all trajectories $\underline{x}(t)$, $\underline{\hat{x}}(t)$, with the action enforcing the dynamics Eq. (3.22). The explicit expression for the coupled OU processes reads

$$Z_J\left[\underline{\psi},\underline{\hat{\psi}}\right] \propto \int \mathscr{D}^0[\underline{x},\underline{\hat{x}}]\,\exp\Big\{-S^{\mathrm{loc}}[\underline{x},\underline{\hat{x}},\underline{\psi},\underline{\hat{\psi}}] - S_J^{\mathrm{int}}[\underline{x},\underline{\hat{x}}]\Big\}, \tag{3.26}$$

where the local part of the action is

$$\begin{aligned} S^{\mathrm{loc}}[\underline{x},\underline{\hat{x}},\underline{\psi},\underline{\hat{\psi}}] = \sum_{i=1}^N \int dt \Big[&\mathrm{i}\hat{x}_i(t)\left(\dot{x}_i(t) + \lambda x_i(t)\right) + T\hat{x}_i(t)^2 \\ &- \psi_i(t)\,x_i(t) - \hat{\psi}_i(t)\mathrm{i}\hat{x}_i(t)\Big], \end{aligned} \tag{3.27}$$

and the interaction part is

$$S_J^{\mathrm{int}}[\underline{x},\underline{\hat{x}}] = -\sum_{i\neq j} \int dt\, \mathrm{i}\hat{x}_i(t)\,J_{ij}\,x_j(t). \tag{3.28}$$

We assumed for simplicity that the initial conditions are independent and identically distributed, and we denoted $\mathscr{D}^0[\underline{x},\underline{\hat{x}}] = \prod_{i=1}^N \mathscr{D}^0[x_i,\hat{x}_i] = \prod_{i=1}^N \mathscr{D}[x_i,\hat{x}_i]p_0(x_i(0))$.

3.3.2 Average over disorder

The generating functional Eq. (3.26) depends on the specific realization of the couplings $\{J_{ij}\}_{i,j=1}^{N}$. Typical properties follow from an average over the disorder, which is exact here due to Gaussianity:

$$Z[\underline{\psi},\underline{\hat{\psi}}] = \overline{Z_J[\underline{\psi},\underline{\hat{\psi}}]} \propto \int \mathscr{D}^0[\underline{x},\underline{\hat{x}}] \exp\Big\{ - S^{\mathrm{loc}}[\underline{x},\underline{\hat{x}},\underline{\psi},\underline{\hat{\psi}}] \Big\} \overline{\exp\Big\{ - S_J^{\mathrm{int}}[\underline{x},\underline{\hat{x}}] \Big\}}. \tag{3.29}$$

The average over the disorder affects only the interaction part of the action, which is linear in the couplings J_{ij}. Using the properties of Gaussian integrals, one finds

$$\begin{aligned} Z[\underline{\psi},\underline{\hat{\psi}}] \propto \int \mathscr{D}^0[\underline{x},\underline{\hat{x}}] \exp\Big\{ & - S^{\mathrm{loc}}[\underline{x},\underline{\hat{x}},\underline{\psi},\underline{\hat{\psi}}] + \frac{m}{N} \sum_{i\neq j} \int dt\, \mathrm{i}\hat{x}_i(t) x_j(t) \\ & + \frac{\sigma^2}{2N} \int dt\, dt' \sum_{i\neq j} \Big[\mathrm{i}\hat{x}_i(t) \mathrm{i}\hat{x}_i(t') x_j(t) x_j(t') + \gamma \mathrm{i}\hat{x}_i(t) x_i(t') x_j(t) \mathrm{i}\hat{x}_j(t') \Big] \Big\}. \end{aligned} \tag{3.30}$$

Details of the average are deferred to Appendix A.2.1.

3.3.3 Order parameters

The disorder-averaged generating functional still couples all degrees of freedom through the interaction terms. To proceed, we introduce *macroscopic order parameters* that summarize the collective behavior of the system. In fully-connected models, these order parameters typically involve averages over all sites, reflecting the mean-field nature of the interactions. Specifically, we introduce the dynamical macroscopic observables

$$\mu(t) = \frac{1}{N} \sum_i x_i(t), \tag{3.31}$$

$$P(t) = \frac{1}{N} \sum_i \mathrm{i}\hat{x}_i(t), \tag{3.32}$$

$$C(t,t') = \frac{1}{N} \sum_i x_i(t)\, x_i(t'), \tag{3.33}$$

$$R(t,t') = \frac{1}{N} \sum_i x_i(t)\, \mathrm{i}\hat{x}_i(t'), \tag{3.34}$$

$$B(t,t') = \frac{1}{N} \sum_i \mathrm{i}\hat{x}_i(t)\, \mathrm{i}\hat{x}_i(t'), \tag{3.35}$$

representing the empirical average magnetization, response-field magnetization, two-time correlation function, linear response function, and response field two-time correlation function, respectively. These quantities capture the essential features of the system's dynamics and will serve as the basis for the DMFT equations.

To enforce the definitions Eqs. (3.31) to (3.35) in the path integral, we introduce functional delta constraints for each order parameter, represented in Fourier form by conjugate fields $\hat{\mu}(t)$, $\hat{P}(t)$, $\hat{C}(t,t')$, $\hat{R}(t,t')$, and $\hat{B}(t,t')$. For compactness we write $Q = \{\mu, P, C, R, B, \hat{\mu}, \hat{P}, \hat{C}, \hat{R}, \hat{B}\}$.

$$\begin{aligned} Z[\underline{\psi}, \underline{\hat{\psi}}] \propto \int \mathscr{D}[Q] \exp\left\{ -N\left(S^{\text{o.p.}}[Q] + S^{\text{avg}}[Q] + O(N^{-1})\right)\right\} \\ \times \exp\left\{ \sum_{i=1}^{N} \log \int \mathscr{D}^0[x_i, \hat{x}_i]\, e^{-S_1[x_i,\hat{x}_i,\psi_i,\hat{\psi}_i|Q]} \right\}, \end{aligned} \tag{3.36}$$

where $\mathscr{D}[Q]$ denotes integration over all order parameters and their conjugate fields. The action is decomposed into three contributions as shown.

$$\begin{aligned} S^{\text{o.p.}}[Q] = \mathrm{i} \int dt\, dt' \left[\hat{C}(t,t')C(t,t') + \hat{R}(t,t')R(t,t') + \hat{B}(t,t')B(t,t')\right] \\ + \mathrm{i} \int dt \left[\hat{\mu}(t)\,\mu(t) + \hat{P}(t)P(t)\right], \end{aligned} \tag{3.37}$$

arises from the Fourier representation of the delta constraints;

$$S^{\text{avg}}[Q] = -m \int dt\, P(t)\mu(t) - \frac{\sigma^2}{2} \int dt\, dt' \left[B(t,t')C(t,t') + \gamma R(t',t)R(t,t')\right], \tag{3.38}$$

collects the terms resulting from the disorder average; and

$$\begin{aligned} S_1[x, \hat{x}, \psi, \hat{\psi} \mid Q] = \int dt \Big[\mathrm{i}\hat{x}(t)\left(\dot{x}(t) + \lambda x(t)\right) + T\hat{x}(t)^2 - \psi(t)x(t) \\ - \hat{\psi}(t)\mathrm{i}\hat{x}(t) - \mathrm{i}\hat{\mu}(t)x(t) - \mathrm{i}\hat{P}(t)\mathrm{i}\hat{x}(t)\Big] \\ - \mathrm{i} \int dt dt' \left[\hat{C}(t,t')x(t)x(t') + \hat{R}(t,t')x(t)\mathrm{i}\hat{x}(t')\right] \\ - \mathrm{i} \int dt dt' \hat{B}(t,t')\mathrm{i}\hat{x}(t)\mathrm{i}\hat{x}(t'), \end{aligned} \tag{3.39}$$

is the effective local action depending on the order parameters. We refer to Appendix A.2.2 for details of the derivation.

The effective local action Eq. (3.39) describes a single-site problem where the dynamics of $x(t)$ and $\hat{x}(t)$ are influenced by the macroscopic order parameters and their conjugate fields.

We therefore introduce the single-site effective generating functional

$$Z_1[\psi,\hat{\psi}\mid Q]\propto\int\mathscr{D}^0[x,\hat{x}]\;e^{-S_1[x,\hat{x},\psi,\hat{\psi}|Q]}, \tag{3.40}$$

which encodes the statistics of the effective single-site dynamics given the order parameters Q. The constant of proportionality is fixed by the normalization condition $Z_1[0,0\mid Q]=1$. The average of any single-site observable $\mathscr{O}[x,\hat{x}]$ is then computed as

$$\langle\mathscr{O}[x,\hat{x}]\rangle_1\propto\int\mathscr{D}^0[x,\hat{x}]\;\mathscr{O}[x,\hat{x}]\,e^{-S_1[x,\hat{x},0,0|Q]}. \tag{3.41}$$

3.3.4 Thermodynamic limit and saddle point

The site index i appears only in the local action Eq. (3.39); aside from the site-specific sources ψ_i and $\hat{\psi}_i$ this action is identical for every site. This factorization, a direct consequence of the fully-connected (mean-field) interaction structure, makes all sites statistically equivalent. Hence the generating functional Eq. (3.36) can be written purely as an integral over the macroscopic order parameters and their conjugate fields. To leading order in N,

$$Z[\underline{\psi},\underline{\hat{\psi}}]\sim\int\mathscr{D}[Q]\,e^{-N\left(S^{\text{o.p.}}[Q]+S^{\text{avg}}[Q]\right)}\;\left(Z^{\text{loc}}_{\text{eff}}[\underline{\psi},\underline{\hat{\psi}}\mid Q]\right)^N, \tag{3.42}$$

where $Z^{\text{loc}}_{\text{eff}}$ is the single-site effective generating functional

$$\log Z^{\text{loc}}_{\text{eff}}[\underline{\psi},\underline{\hat{\psi}}\mid Q]=\frac{1}{N}\sum_{i=1}^{N}\log Z_1[\psi_i,\hat{\psi}_i\mid Q]. \tag{3.43}$$

In the thermodynamic limit $N\to\infty$, the integral Eq. (3.36) is dominated by a saddle point $Q^\star$, yielding

$$Z[\underline{\psi},\underline{\hat{\psi}}]\overset{N\to\infty}{\sim}\exp\Big\{-N\big(S^{\text{o.p.}}[Q^\star]+S^{\text{avg}}[Q^\star]\big)\Big\}\;\left(Z^{\text{loc}}_{\text{eff}}[\underline{\psi},\underline{\hat{\psi}}\mid Q^\star]\right)^N. \tag{3.44}$$

At the saddle point, the order parameters satisfy the self-consistent equations

$$\begin{aligned}
\mu^\star(t) &= \langle x(t)\rangle_1^\star, & \mathrm{i}\hat{\mu}^\star(t) &= mP^\star(t) = 0,\\
P^\star(t) &= \langle \mathrm{i}\hat{x}(t)\rangle_1^\star = 0, & \mathrm{i}\hat{P}^\star(t) &= m\mu^\star(t),\\
C^\star(t,t') &= \langle x(t)\,x(t')\rangle_1^\star, & \mathrm{i}\hat{C}^\star(t,t') &= \frac{\sigma^2}{2}B^\star(t,t') = 0,\\
R^\star(t,t') &= \langle x(t)\,\mathrm{i}\hat{x}(t')\rangle_1^\star, & \mathrm{i}\hat{R}^\star(t,t') &= \gamma\sigma^2 R^\star(t',t),\\
B^\star(t,t') &= \langle \mathrm{i}\hat{x}(t)\,\mathrm{i}\hat{x}(t')\rangle_1^\star = 0, & \mathrm{i}\hat{B}^\star(t,t') &= \frac{\sigma^2}{2}C^\star(t,t'),
\end{aligned}$$

where $\langle\cdot\rangle_1^\star$ denotes the average with respect to the effective single-site dynamics at the saddle point $Q^\star$. The vanishing of $P^\star(t)$ and $B^\star(t,t')$ follows from causality and normalization conditions, as detailed in Appendix A.2.3. By substituting the conjugate order parameters back into the effective local action Eq. (3.39), one obtains

$$\begin{aligned}
S_1^\star[x,\hat{x}] = &\int dt\,\mathrm{i}\hat{x}(t)\left(\dot{x}(t) + \lambda x(t) - m\mu(t) - \gamma\sigma^2\int dt'\,R(t,t')x(t')\right)\\
&+ \int dt\,T\hat{x}(t)^2 - \frac{\sigma^2}{2}\int dt dt'\,\mathrm{i}\hat{x}(t)C(t,t')\mathrm{i}\hat{x}(t').
\end{aligned} \tag{3.45}$$

The order parameters $\mu(t)$, $C(t,t')$, and $R(t,t')$ must be determined self-consistently from the single-site dynamics defined by Eq. (3.45). They can be interpreted as the disorder-averaged magnetization, correlation function, and response function of a typical site in the network. It is straightforward to verify that, for example for the magnetization,

$$\mu(t) = \frac{1}{N}\sum_{i=1}^{N}\overline{\langle x_i(t)\rangle} = \langle x(t)\rangle_1^\star, \tag{3.46}$$

where the overline denotes the average over the disorder distribution. The same holds for the correlation and response functions.

3.3.5 Effective single-site process

The effective local action Eq. (3.45) describes a single-site stochastic process where the dynamics of $x(t)$ is influenced by self-consistent memory and noise terms. To make this explicit, we can follow backwards the steps of the MSRJD construction to identify the corresponding Langevin equation. The quadratic term in $\hat{x}(t)$ can be decoupled by introducing an auxiliary

Gaussian noise $\eta(t)$ with zero mean and correlations

$$\langle \eta(t)\,\eta(t')\rangle = 2T\,\delta(t-t') + \sigma^2 C(t,t'). \tag{3.47}$$

The effective single-site dynamics is then governed by the non-Markovian Langevin equation

$$\frac{dx(t)}{dt} = -\lambda x(t) + m\mu(t) + \gamma\sigma^2 \int_0^t dt'\, R(t,t')x(t') + \eta(t), \tag{3.48}$$

where the integral term represents a memory effect due to the interactions. The order parameters $\mu(t)$, $C(t,t')$, and $R(t,t')$ must be determined self-consistently from the statistics of $x(t)$ under this effective dynamics. The effective noise $\eta(t)$ captures both the thermal fluctuations and the quenched-disorder effects, leading to a rich dynamical behavior. The DMFT equations thus reduce the original N-body problem to a single-site stochastic process with self-consistent memory and noise, providing a powerful framework to analyze the dynamics of disordered systems in the thermodynamic limit.

3.4 The Plefka expansion and its dynamical generalization

The *Plefka expansion* provides a perturbative route to mean-field theory by expanding the Gibbs free energy in the interaction strength. In its original formulation by Plefka [36] for equilibrium spin glasses, the small-coupling expansion recovers mean field at first order and the Thouless-Anderson-Palmer (TAP) equations at second order. A key advantage is that it does not rely on averaging over random couplings and thus applies to specific networks. A dynamical analog for Ising spins was developed by Roudi and Hertz [37]. They obtained a dynamical equivalent of the Gibbs free energy by Legendre transforming the logarithm of the dynamical generating functional, constraining only single-time magnetizations, and showed that the expansion in the couplings recovers (naive) dynamical mean field and dynamical TAP at first and second order, respectively. Bravi, Sollich and Opper [38] extended the construction to Langevin dynamics of continuous variables by constraining not only the means but also local two-time responses and correlations; expanding the dynamical Gibbs free energy to second order yields an effective single-site process with a memory kernel and coloured noise. For linear interactions and fully-connected interaction graphs the predictions coincide exactly with the thermodynamic limit solution of the DMFT, independently of the symmetry of the couplings.

To illustrate the method, consider a system of N linearly-coupled OU processes

$$\frac{dx_i(t)}{dt} = -\lambda_i x_i(t) + \alpha \sum_{j\neq i} J_{ij} x_j(t) + \xi_i(t), \quad \langle \xi_i(t)\xi_j(t')\rangle = 2T\delta_{ij}\delta(t-t'), \tag{3.49}$$

where $\lambda_i > 0$ are local decay rates, J_{ij} are fixed couplings, and α is a parameter controlling the interaction strength. The Plefka expansion treats α as small and expands the dynamical Gibbs free energy around $\alpha = 0$. The non-interacting case $\alpha = 0$ corresponds to independent OU processes, which are exactly solvable. The expansion is performed by constraining the means $\mu_i(t) = \langle x_i(t)\rangle$, the local two-time (connected) correlations $C_i(t,t') = \langle x_i(t)x_i(t')\rangle - \mu_i(t)\mu_i(t')$, and the local two-time linear responses $R_i(t,t') = \langle x_i(t)\mathrm{i}\hat{x}_i(t')\rangle$.

At first order one recovers the (naive) dynamical mean-field equations for the local magnetizations

$$\frac{d\mu_i(t)}{dt} = -\lambda_i \mu_i(t) + \alpha \sum_{j\neq i} J_{ij}\mu_j(t). \tag{3.50}$$

At second order one obtains a decoupled effective description for each degree of freedom, with the effective single-site dynamics

$$\frac{dx_i(t)}{dt} = -\lambda_i x_i(t) + \sum_{j\neq i} J_{ij}\mu_j(t) + \sum_{j\neq i} J_{ij}J_{ji}\int_0^t dt' R_j(t,t')x_i(t') + \eta_i(t), \tag{3.51}$$

with memory kernel and coloured Gaussian noise with zero mean and correlations

$$\langle \eta_i(t)\eta_i(t')\rangle = \sum_{j\neq i} J_{ij}^2 C_j(t,t'). \tag{3.52}$$

The self-consistent order parameters are the local magnetization, response and connected correlation,

$$\mu_i(t) = \langle x_i(t)\rangle, \quad R_i(t,t') = \left.\frac{\delta\langle x_i(t)\rangle}{\delta h_i(t')}\right|_{h=0}, \quad C_i(t,t') = \langle x_i(t)x_i(t')\rangle - \mu_i(t)\mu_i(t'), \tag{3.53}$$

whose evolution is governed by the closed equations given in Eqs. (3.52)-(3.53) of Ref. [38].

The Plefka expansion thus yields an effective single-site description with memory and coloured noise, similar in structure to the DMFT result Eq. (3.48). However, unlike DMFT, the Plefka expansion does not require averaging over random couplings and applies to specific networks. For fully-connected networks with random couplings, the Plefka expansion at second order recovers the exact DMFT results in the thermodynamic limit, as shown in [38]. This equivalence holds for any symmetry of the couplings J_{ij}.

3.5 Limitations for sparse graphs

The effectiveness of both DMFT and the Plefka small-coupling expansion relies on self-averaging and central limit mechanisms that are naturally realized in dense or fully-connected graphs. On sparse networks with $k_i = O(1)$ neighbors, local fields are sums over a few heterogeneous inputs and remain non-Gaussian and strongly site dependent; degree fluctuations and quenched neighborhood variability are not suppressed in the thermodynamic limit, and memory kernels acquire an explicit dependence on the local environment beyond degree.

For DMFT, a pragmatic high-connectivity regime leads to *degree-based* closures–often termed Heterogeneous DMFT (HDMFT)–where dynamics are coarse-grained by degree classes. This approximation is classical in epidemic modeling [39, 40] and has been exported to statistical physics and random-matrix settings [41–45]. While informative at large k, its accuracy deteriorates for moderate degrees, degree correlations, clustering, or strongly nonlinear dynamics–precisely because annealed by degree closures erase quenched neighborhood fluctuations.

For Plefka, second-order exactness is established in the weak, long-range scaling (dense graphs with appropriately scaled couplings), e.g. in the linear case independently of coupling symmetry [38]. On sparse graphs the second-order truncation is no longer controlled: effective fields are aggregates of a few heterogeneous inputs and remain strongly node specific, so the induced memory and coloured noise terms depend sensitively on the local neighborhood and are not faithfully captured by constraining only single-node order parameters. In this regime the extended Plefka approximation is best regarded as an informative small-coupling baseline rather than a quantitatively accurate predictor. This is precisely why, in what follows, we pair a cavity/message-passing formulation–suited to finite connectivity–with a Plefka-style organization of weak-interaction corrections; the Gaussian closures we employ there are a deliberate modeling choice to obtain closed equations, not a prerequisite for the dense limit guarantees of [38].

These limitations motivate the framework developed in this thesis: we combine the cavity method–well-suited to sparse, locally tree-like networks–with a small-coupling expansion. The resulting scheme retains the structural resolution of dynamic cavity/message-passing at finite connectivity while remaining consistent with dense graph limits where DMFT (and second-order Plefka) become exact. As for the Plefka expansion, our approach does not rely on averaging over random couplings and can be employed for specific networks.

Part II

Continuous Degrees of Freedom

Continuous Degrees of Freedom

Many systems of interest in physics, biology, and economics are naturally described in terms of variables evolving in continuous time under the combined effect of deterministic interactions and stochastic fluctuations. In this context, the appropriate mathematical framework is provided by overdamped Langevin equations, which capture the interplay between local deterministic drifts and noise.

In Part I we have discussed how path-integral approaches and mean-field techniques, such as Dynamical Mean-Field Theory and the Plefka expansion, provide valuable insights in the dense limit but lose accuracy on sparse networks, where finite connectivity and heterogeneity remain essential. The aim of this part is to extend cavity-based ideas to the setting of continuous degrees of freedom, introducing a Gaussian expansion that yields tractable equations for local moments of the dynamics.

This framework will allow us to address both analytically and numerically the behavior of linearly-coupled stochastic processes on sparse graphs, and to explore their perturbative extensions to nonlinear models. We will apply these methods to paradigmatic examples, including the spherical 2-spin model and ecological communities, thereby illustrating how network structure shapes the stability and dynamical phases of continuous-variable systems.

The chapters included in this part are based on published and ongoing work. In particular, chapters 4 to 6 are adapted from the article [46], where the Gaussian Expansion Cavity Method was first introduced and applied to stochastic dynamics with continuous degrees of freedom. All numerical data and codes used for the analyses presented in these chapters are openly available in [47]. The final chapter, 7, presents original results currently being finalized for publication and therefore constitutes an article in preparation. The code and data used for this chapter can be found in [48].

Chapter 4

The Gaussian-Expansion Cavity Method

In this chapter we introduce the *Gaussian Expansion Cavity Method (GECaM)*, a framework designed to extend the cavity approach to continuous-time stochastic dynamics. The method is based on a Gaussian parametrization of the cavity messages, which becomes exact for linearly-coupled Langevin equations with additive noise and provides a natural starting point for perturbative extensions to nonlinear cases. We begin by formulating the dynamic cavity equations for linearly-coupled processes (section 4.1), and show how the Gaussian closure yields exact equations for local moments (section 4.2) and how it relates to random-matrix theory (section 4.3). We then analyze specific graph ensembles, first considering random regular graphs (section 4.4) and subsequently heterogeneous networks (section 4.5), in order to highlight the role of topology. The chapter sets the basis for later applications to nonlinear extensions, the spherical 2-spin model, and ecological dynamics.

4.1 Dynamic cavity for linearly-coupled dynamics

To make the discussion concrete, we focus on the simplest non-trivial case: a system of interacting Ornstein-Uhlenbeck processes. The same model was already introduced in section 3.3 in the context of Dynamical Mean-Field Theory, where we considered fully-connected interaction networks with couplings drawn as quenched Gaussian random variables and averaged over disorder. Here we instead study a single instance of the couplings defined on a sparse interaction graph, and derive the corresponding cavity equations.

Let $G=(V,E)$ denote the underlying interaction network with $|V|=N$ nodes, and adjacency matrix $\boldsymbol{A}$. The state of node $i \in V$ evolves according to the linear Langevin equation

$$\frac{dx_i(t)}{dt} = -\lambda_i x_i(t) + \alpha \sum_{j\in\partial i} J_{ij} x_j(t) + \xi_i(t), \tag{4.1}$$

where $\lambda_i > 0$ is a local relaxation rate, J_{ij} are coupling coefficients associated with edges $(i,j) \in E$, and $\xi_i(t)$ is a Gaussian white noise with zero mean and covariance

$$\langle \xi_i(t)\xi_{i'}(t')\rangle = 2T\,\delta_{ii'}\,\delta(t-t'). \tag{4.2}$$

We have introduced a parameter α to control the overall interaction strength. The initial conditions $x_i(0)$ are drawn independently from a distribution $p_0(x_i(0))$.

As discussed in section 3.2, the statistical properties of the process can be encoded in a dynamical generating function, obtained via the MSRJD formalism. Introducing response fields $\{\hat{x}_i(t)\}_{i=1}^N$ and collecting trajectories over the time interval $[0,t_f]$, one arrives at

$$Z[\underline{\psi},\underline{\hat{\psi}}] = \int \mathscr{D}^0[\underline{x},\underline{\hat{x}}] \exp\left[-S[\underline{x},\underline{\hat{x}}] + \sum_{i=1}^N \int_0^{t_f} dt \left(\psi_i(t)x_i(t) + \hat{\psi}_i(t)\mathrm{i}\hat{x}_i(t)\right)\right], \tag{4.3}$$

where $\psi_i(t)$ and $\hat{\psi}_i(t)$ are source terms and the action reads

$$S[\underline{x},\underline{\hat{x}}] = \sum_{i=1}^N S_i^{\mathrm{loc}}[x_i,\hat{x}_i] + \alpha \sum_{(i,j)\in E} S_{ij}^{\mathrm{int}}[x_i,\hat{x}_i,x_j,\hat{x}_j], \tag{4.4}$$

where we defined the local and interaction contributions

$$S_i^{\mathrm{loc}}[x_i,\hat{x}_i] = \int_0^{t_f} dt \left[\mathrm{i}\hat{x}_i(t)\left(\frac{dx_i(t)}{dt} + \lambda_i x_i(t)\right) + T\hat{x}_i(t)^2\right], \tag{4.5}$$

$$S_{ij}^{\mathrm{int}}[x_i,\hat{x}_i,x_j,\hat{x}_j] = -\int_0^{t_f} dt \left(J_{ij}\,\mathrm{i}\hat{x}_i(t)\,x_j(t) + J_{ji}\,\mathrm{i}\hat{x}_j(t)\,x_i(t)\right). \tag{4.6}$$

The derivation of this path-integral representation follows directly from the standard MSRJD construction and is reported in Appendix B.1.1.

The important feature of Eq. (4.1) is that the action factorizes over the nodes and edges of the interaction graph. This factorization allows for a simplified treatment of the system's dynamics, as each node can be analyzed in the context of its local environment. The overall structure is reminiscent of a graphical model, with nodes representing pairs of trajectories $(x_i,\hat{x}_i)$ and edges representing interactions. The corresponding factor graph has the same topology of

the original interaction graph G. This analogy motivates the use of message-passing techniques, such as the cavity method, to study the dynamics. Assuming the interaction graph G to be locally tree-like, we can put forward a *dynamic cavity ansatz* to approximate the local trajectory distribution. Denoting by $c_{i\setminus j}[x_i,\hat{x}_i]$ the cavity message from i to j, i.e. the probability weight of the trajectory of node i in the absence of node j, the cavity equations read

$$c_{i\setminus j}[x_i,\hat{x}_i] = \frac{p_0(x_i(0))}{Z_{i\setminus j}} e^{-S_i^{\mathrm{loc}}[x_i,\hat{x}_i]} \prod_{k\in\partial i\setminus j} \int \mathscr{D}[x_k,\hat{x}_k]\, e^{-\alpha S_{ik}^{\mathrm{int}}[x_i,\hat{x}_i,x_k,\hat{x}_k]}\, c_{k\setminus i}[x_k,\hat{x}_k]. \tag{4.7}$$

As already mentioned in chapter 2, the normalization of cavity messages through the partition function $Z_{i\setminus j}$ is not critical and can be adjusted after convergence. For later convenience, we choose to properly normalize the cavity messages as quasi probability distributions over trajectories[1], i.e. such that

$$\int \mathscr{D}[x_i,\hat{x}_i]\, c_{i\setminus j}[x_i,\hat{x}_i] = 1.$$

The normalization constant $Z_{i\setminus j}$ is then fixed by this condition

$$\begin{aligned} Z_{i\setminus j} = \int \mathscr{D}[x_i,\hat{x}_i]\, p_0(x_i(0))\, e^{-S_i^{\mathrm{loc}}[x_i,\hat{x}_i]} \\ \times \prod_{k\in\partial i\setminus j} \int \mathscr{D}[x_k,\hat{x}_k]\, e^{-\alpha S_{ik}^{\mathrm{int}}[x_i,\hat{x}_i,x_k,\hat{x}_k]}\, c_{k\setminus i}[x_k,\hat{x}_k]. \end{aligned} \tag{4.8}$$

An analogous expression holds for the single-node marginals

$$c_i[x_i,\hat{x}_i] = \frac{p_0(x_i(0))}{Z_i} e^{-S_i^{\mathrm{loc}}[x_i,\hat{x}_i]} \prod_{k\in\partial i} \int \mathscr{D}[x_k,\hat{x}_k]\, e^{-\alpha S_{ik}^{\mathrm{int}}[x_i,\hat{x}_i,x_k,\hat{x}_k]}\, c_{k\setminus i}[x_k,\hat{x}_k], \tag{4.9}$$

with Z_i fixed by normalization

$$\begin{aligned} Z_i = \int \mathscr{D}[x_i,\hat{x}_i]\, p_0(x_i(0))\, e^{-S_i^{\mathrm{loc}}[x_i,\hat{x}_i]} \\ \times \prod_{k\in\partial i} \int \mathscr{D}[x_k,\hat{x}_k]\, e^{-\alpha S_{ik}^{\mathrm{int}}[x_i,\hat{x}_i,x_k,\hat{x}_k]}\, c_{k\setminus i}[x_k,\hat{x}_k]. \end{aligned} \tag{4.10}$$

Such a normalization ensures that the marginals $c_i[x_i,\hat{x}_i]$ are also properly normalized quasi-probability distributions over trajectories. We can therefore introduce local averages of any

[1] A quasi probability distribution is a generalization of a probability distribution that may take on negative values, but still integrates to one.

observable $\mathscr{O}_i[x_i, \hat{x}_i]$ as

$$\langle \mathscr{O}_i[x_i, \hat{x}_i] \rangle_{i \setminus j} = \int \mathscr{D}[x_i, \hat{x}_i]\, c_{i \setminus j}[x_i, \hat{x}_i]\, \mathscr{O}_i[x_i, \hat{x}_i], \tag{4.11}$$

$$\langle \mathscr{O}_i[x_i, \hat{x}_i] \rangle_{i} = \int \mathscr{D}[x_i, \hat{x}_i]\, c_{i}[x_i, \hat{x}_i]\, \mathscr{O}_i[x_i, \hat{x}_i]. \tag{4.12}$$

The dynamic cavity Eqs. (4.7) and (4.9) are exact on trees, and become asymptotically exact on locally tree-like graphs in the thermodynamic limit. However–similarly to the discrete case discussed in section 2.3–they are of very limited use, being their direct numerical implementation infeasible due to the exponential growth of the trajectory space with the horizon time t_f. To make progress, we need to introduce further approximations. A natural choice is to assume a specific functional form for the cavity messages, parametrized by a finite set of parameters. Given the linearity of the dynamics in Eq. (4.1), a natural choice is to assume that the cavity messages are Gaussian distributions over trajectories. This leads to the GECaM, which we introduce in the next section.

4.2 Exact Gaussian Expansion for Linear Dynamics

As discussed at the end of the previous section, the cavity equations are formally exact but intractable, since the trajectory space grows exponentially with the time horizon. A natural way forward is to approximate cavity messages by distributions with a restricted functional form, parametrized by a finite set of moments.

To motivate the Gaussian Ansatz, it is useful to follow the same strategy adopted for the Plefka expansion described in section 3.4. We introduce the parameter α as in Eq. (4.1) and treat the interaction terms in the cavity recursion Eq. (4.7) perturbatively around the non-interacting limit. In the weak-coupling regime $\alpha \ll 1$, the exponential of the interaction action can be expanded as a power series in α. Keeping terms up to second order one finds that the effective action contains only linear and quadratic contributions of the fields x and $\hat{x}$.

At first order in α, the correction simply shifts the local drift of node i by contributions proportional to the averages of its neighbors. At second order, averages of products of neighboring trajectories appear, leading to quadratic terms that renormalize the noise covariances through

the neighbors' cavity correlations and responses. Introducing the cavity averages

$$\mu_{i\setminus j}(t) = \langle x_i(t)\rangle_{i\setminus j}, \tag{4.13}$$
$$C_{i\setminus j}(t,t') = \langle x_i(t)x_i(t')\rangle_{i\setminus j} - \mu_{i\setminus j}(t)\mu_{i\setminus j}(t'), \tag{4.14}$$
$$R_{i\setminus j}(t,t') = \langle x_i(t)\,\mathrm{i}\hat{x}_i(t')\rangle_{i\setminus j}, \tag{4.15}$$

the cavity measure up to $O(\alpha^2)$ is approximated by

$$c_{i\setminus j}(x_i,\hat{x}_i) = \frac{p_0(x_i(0))}{Z^{\mathrm{eff}}_{i\setminus j}} \exp\left(-S^{\mathrm{eff}}_{i\setminus j}[x_i,\hat{x}_i]\right), \tag{4.16}$$

where $S^{\mathrm{eff}}_{i\setminus j}$ is an effective action quadratic in $(x_i,\hat{x}_i)$,

$$S^{\mathrm{eff}}_{i\setminus j}[x_i,\hat{x}_i] = \int_0^{t_f} dt \left[\mathrm{i}\hat{x}_i(t)\left(\dot{x}_i(t) + \lambda_i x_i(t) + \alpha \sum_{k\in\partial i\setminus j} J_{ik}\mu_{k\setminus i}(t)\right) + T\hat{x}_i^2(t)\right]$$
$$+ \frac{\alpha^2}{2}\int_0^{t_f} dt dt' \sum_{k\in\partial i\setminus j}\left[\mathrm{i}\hat{x}_i(t)J_{ik}^2 C_{k\setminus i}(t,t')\mathrm{i}\hat{x}_i(t') + 2\,\mathrm{i}\hat{x}_i(t)J_{ik}J_{ki}R_{k\setminus i}(t,t')x_i(t')\right]. \tag{4.17}$$

The normalization constant $Z^{\mathrm{eff}}_{i\setminus j}$ is fixed by the condition that $c_{i\setminus j}$ integrates to one. Further details on the derivation of Eqs. (4.16) and (4.17) are provided in Appendix B.1.2. The effective action Eq. (4.17) is quadratic in the fields, and therefore the cavity measure Eq. (4.16) is a Gaussian distribution over trajectories. Since α has served its purpose in the derivation, from this point onward we set $\alpha = 1$ for simplicity.

4.2.1 Gaussian Ansatz in discrete time

To make the Gaussian Ansatz explicit, we start from the discrete-time representation of the dynamics introduced in Eq. (4.1). Discretizing the interval $[0,t_f]$ into $T+1$ steps of size $\Delta = t_f/T$, we assume that the cavity messages are multivariate Gaussian distributions over the $(T+1)$-dimensional vectors $\boldsymbol{x}_i$ and $\hat{\boldsymbol{x}}_i$. For compactness, we collect the two fields into a single $2(T+1)$-dimensional vector $\boldsymbol{X}_i = (\boldsymbol{x}_i,\hat{\boldsymbol{x}}_i)$. The Gaussian Ansatz for the cavity message from i to j reads

$$c_{i\setminus j}(\boldsymbol{X}_i) = \frac{1}{Z_{i\setminus j}} \exp\left[-\tfrac{1}{2}(\boldsymbol{X}_i - \boldsymbol{M}_{i\setminus j})\boldsymbol{G}^{-1}_{i\setminus j}(\boldsymbol{X}_i - \boldsymbol{M}_{i\setminus j})^\top\right], \tag{4.18}$$

where $\boldsymbol{M}_{i\backslash j} = (\boldsymbol{\mu}_{i\backslash j}, \boldsymbol{P}_{i\backslash j})$ is the $2(T+1)$-dimensional mean vector and $\boldsymbol{G}_{i\backslash j}$ the corresponding $2(T+1)\times 2(T+1)$ covariance matrix,

$$\boldsymbol{G}_{i\backslash j} = \begin{pmatrix} \boldsymbol{C}_{i\backslash j} & \boldsymbol{R}_{i\backslash j} \\ \boldsymbol{R}_{i\backslash j}^\top & \boldsymbol{B}_{i\backslash j} \end{pmatrix}. \tag{4.19}$$

The vectors $\boldsymbol{\mu}_{i\backslash j}$ and $\boldsymbol{P}_{i\backslash j}$ contain the discrete-time cavity means,

$$\mu_{i\backslash j}^n = \langle x_i^n \rangle_{i\backslash j}, \qquad P_{i\backslash j}^n = \langle \mathrm{i}\hat{x}_i^n \rangle_{i\backslash j},$$

while the blocks $\boldsymbol{C}_{i\backslash j}$, $\boldsymbol{R}_{i\backslash j}$, and $\boldsymbol{B}_{i\backslash j}$ collect the correlation and response functions,

$$\begin{aligned} C_{i\backslash j}^{nm} &= \langle x_i^n x_i^m \rangle_{i\backslash j} - \mu_{i\backslash j}^n \mu_{i\backslash j}^m, \\ R_{i\backslash j}^{nm} &= \langle x_i^n \mathrm{i}\hat{x}_i^m \rangle_{i\backslash j} - \mu_{i\backslash j}^n \hat{\mu}_{i\backslash j}^m, \\ B_{i\backslash j}^{nm} &= \langle \mathrm{i}\hat{x}_i^n \mathrm{i}\hat{x}_i^m \rangle_{i\backslash j} - \hat{\mu}_{i\backslash j}^n \hat{\mu}_{i\backslash j}^m. \end{aligned}$$

As discussed in Appendix A.1.6, averages involving only response fields vanish identically in the MSRJD formalism, implying $\boldsymbol{P}_{i\backslash j} = \boldsymbol{0}$ and $\boldsymbol{B}_{i\backslash j} = \boldsymbol{0}$[2]. The normalization constant ensuring unit integral is

$$Z_{i\backslash j} = (2\pi)^{T+1} \sqrt{\det \boldsymbol{G}_{i\backslash j}}. \tag{4.20}$$

Inserting the Gaussian Ansatz Eq. (4.18) into the cavity equation Eq. (4.7), we can derive self-consistent equations for the local moments. In order to make the Gaussian integrals explicit, we rewrite the local and interaction actions Eqs. (4.5) and (4.6) in discrete time as

$$S_i^{\mathrm{loc}}(\boldsymbol{X}_i) = \frac{1}{2} \boldsymbol{X}_i \left[\boldsymbol{G}_i^{\mathrm{loc}}\right]^{-1} \boldsymbol{X}_i^\top, \tag{4.21}$$

$$S_{ij}^{\mathrm{int}}(\boldsymbol{X}_i, \boldsymbol{X}_j) = -\boldsymbol{X}_i \boldsymbol{J}_{ij} \boldsymbol{X}_j^\top, \tag{4.22}$$

where we introduced the $2(T+1)\times 2(T+1)$ matrices

$$\left[\boldsymbol{G}_i^{\mathrm{loc}}\right]^{-1} = \begin{pmatrix} \boldsymbol{0} & \left((\boldsymbol{R}_i^{\mathrm{loc}})^{-1}\right)^\top \\ (\boldsymbol{R}_i^{\mathrm{loc}})^{-1} & 2\Delta T \boldsymbol{I} \end{pmatrix}, \tag{4.23}$$

$$\boldsymbol{J}_{ij} = \begin{pmatrix} \boldsymbol{0} & J_{ij}\Delta \boldsymbol{I} \\ J_{ji}\Delta \boldsymbol{I} & \boldsymbol{0} \end{pmatrix}. \tag{4.24}$$

[2] In the MSRJD formalism this property follows from the normalization of the path-integral measure. Here the normalization depends on the external field, yet we can safely assume the same result and verify it *a posteriori*.

Here $\boldsymbol{I}$ is the $(T+1)\times(T+1)$ identity matrix and $\boldsymbol{R}_i^{\mathrm{loc}}$ is the bare response function of node i in the absence of interactions, i.e. $(\boldsymbol{R}_i^{\mathrm{loc}})^{-1} = \boldsymbol{\partial}_t + \lambda_i \Delta \boldsymbol{I}$, where $\boldsymbol{\partial}_t$ is the discrete-time derivative operator

$$\boldsymbol{\partial}_t = \begin{pmatrix} \ddots & 1 & 0 & 0 \\ 0 & -1 & 1 & 0 \\ 0 & 0 & -1 & 1 \\ 0 & 0 & 0 & \ddots \end{pmatrix}. \tag{4.25}$$

By substituting these expressions into Eq. (4.7) and performing the Gaussian integrals, we obtain self-consistent equations for the cavity moments. The details of the derivation are provided in Appendix B.1.3. The final result is a set of coupled equations for the cavity means $\boldsymbol{\mu}_{i\setminus j}$, correlations $\boldsymbol{C}_{i\setminus j}$, and responses $\boldsymbol{R}_{i\setminus j}$, which can be solved iteratively. We refer to them as the *discrete-time GECaM equations*, which are expressed in matrix form as

$$\left[\boldsymbol{G}_i^{\mathrm{loc}}\right]^{-1} \boldsymbol{G}_{i\setminus j} = \boldsymbol{I} + \sum_{k\in\partial i\setminus j} \boldsymbol{J}_{ik} \boldsymbol{G}_{k\setminus i} \boldsymbol{J}_{ik}^{\top} \boldsymbol{G}_{i\setminus j}, \tag{4.26}$$

$$\left[\boldsymbol{G}_i^{\mathrm{loc}}\right]^{-1} \boldsymbol{M}_{i\setminus j}^{\top} = \sum_{k\in\partial i\setminus j} \boldsymbol{J}_{ik} \boldsymbol{M}_{k\setminus i}^{\top} + \sum_{k\in\partial i\setminus j} \boldsymbol{J}_{ik} \boldsymbol{G}_{k\setminus i} \boldsymbol{J}_{ik}^{\top} \boldsymbol{M}_{i\setminus j}^{\top}, \tag{4.27}$$

where $\boldsymbol{I}$ is the $2(T+1)\times 2(T+1)$ identity matrix.

By expanding the matrix products in Eqs. (4.26) and (4.27), one obtains explicit equations for the cavity moments. In particular, the responses and the correlations satisfy the following equations

$$\left(\boldsymbol{R}_i^{\mathrm{loc}}\right)^{-1} \boldsymbol{R}_{i\setminus j} = \boldsymbol{I} + \Delta^2 \sum_{k\in\partial i\setminus j} J_{ik} J_{ki} \boldsymbol{R}_{k\setminus i} \boldsymbol{R}_{i\setminus j}, \tag{4.28}$$

$$\left(\boldsymbol{R}_i^{\mathrm{loc}}\right)^{-1} \boldsymbol{C}_{i\setminus j} = 2\Delta T \boldsymbol{R}_{i\setminus j}^{\top} + \Delta^2 \sum_{k\in\partial i\setminus j} J_{ik}^2 \boldsymbol{C}_{k\setminus i} \boldsymbol{R}_{i\setminus j}^{\top} + \Delta^2 \sum_{k\in\partial i\setminus j} J_{ik} J_{ki} \boldsymbol{R}_{k\setminus i} \boldsymbol{C}_{i\setminus j}, \tag{4.29}$$

while the cavity means satisfy

$$\left(\boldsymbol{R}_i^{\mathrm{loc}}\right)^{-1} \boldsymbol{\mu}_{i\setminus j}^{\top} = \Delta \sum_{k\in\partial i\setminus j} J_{ik} \boldsymbol{\mu}_{k\setminus i}^{\top} + \Delta^2 \sum_{k\in\partial i\setminus j} J_{ik} J_{ki} \boldsymbol{R}_{k\setminus i} \boldsymbol{\mu}_{i\setminus j}^{\top}. \tag{4.30}$$

The single-node marginals can be obtained similarly, leading to equations analogous to Eqs. (4.28) to (4.30) but with sums over all neighbors ∂i instead of $\partial i \setminus j$. These equations can be solved by direct integration assuming initial conditions for the means and correlations, and enforcing causality for the responses. The assumption that averages involving only response fields vanish, as used above, is justified *a posteriori*: a general derivation that allows for

non-vanishing response averages is presented in Appendix B.1.4, where we show that the self-consistent solution indeed yields vanishing averages. However, for non-causal dynamics–such as those arising in locally constrained dynamics relevant to inference or control theory (see chapter 10)–this property no longer holds. In such cases, the response fields play a crucial role in propagating information backward in time, and the more general formalism becomes necessary.

4.2.2 Continuous-time limit

The derivation of the discrete-time GECaM equations in section 4.2.1 can be straightforwardly extended to the continuous-time limit $\Delta \to 0$. In this limit, the discrete-time vectors and matrices become functions of continuous-time variables, and the sums over time steps are replaced by integrals. The discrete-time derivative operator $\boldsymbol{\partial}_t/\Delta$ becomes the continuous-time derivative ∂_t. As for the discrete-time case, we define the macro trajectory $\mathrm{X}_i(t) = (x_i(t), \hat{x}_i(t))$ and assume that the cavity messages are Gaussian distributions over the space of continuous trajectories,

$$c_{i\backslash j}[\mathrm{X}_i] = \frac{1}{Z_{i\backslash j}} \exp\left[-\frac{1}{2}\int_0^{t_f} dt dt' \left(\mathrm{X}_i(t) - \mathrm{M}_{i\backslash j}(t)\right) \mathrm{G}^{-1}_{i\backslash j}(t,t') \left(\mathrm{X}_i(t') - \mathrm{M}_{i\backslash j}(t')\right)^\top\right], \quad (4.31)$$

where $\mathrm{M}_{i\backslash j}(t) = (\mu_{i\backslash j}(t), 0)$ is the mean function and $\mathrm{G}_{i\backslash j}(t,t')$ the covariance function,

$$\mathrm{G}_{i\backslash j}(t,t') = \begin{pmatrix} C_{i\backslash j}(t,t') & R_{i\backslash j}(t,t') \\ R_{i\backslash j}(t',t) & 0 \end{pmatrix}. \quad (4.32)$$

The normalization constant ensures unit integral of the cavity message. The local and interaction actions in continuous time can be expressed as

$$S_i^{\mathrm{loc}}[\mathrm{X}_i] = \frac{1}{2}\int_0^{t_f} dt dt' \, \mathrm{X}_i(t) \left[\mathrm{G}_i^{\mathrm{loc}}\right]^{-1}(t,t') \, \mathrm{X}_i(t')^\top, \quad (4.33)$$

$$S_{ij}^{\mathrm{int}}[\mathrm{X}_i, \mathrm{X}_j] = -\int_0^{t_f} dt \, \mathrm{X}_i(t) \, \mathrm{J}_{ij} \, \mathrm{X}_j(t)^\top, \quad (4.34)$$

where the continuous-time operators are defined as

$$[\mathrm{G}_i^{\mathrm{loc}}]^{-1}(t,t') = \begin{pmatrix} 0 & -\partial_t + \lambda_i \\ \partial_t + \lambda_i & 2T \end{pmatrix} \delta(t-t'), \tag{4.35}$$

$$\mathrm{J}_{ij} = \begin{pmatrix} 0 & J_{ij} \\ J_{ji} & 0 \end{pmatrix} \delta(t-t'). \tag{4.36}$$

Inserting the Gaussian Ansatz Eq. (4.31) into the cavity equation Eq. (4.7) and performing the functional integrals, we obtain the generalization to continuous time of the GECaM equations. The *continuous-time GECaM equations* for the cavity responses, correlations, and means read

$$\frac{\partial R_{i\setminus j}(t,t')}{\partial t} = -\lambda_i R_{i\setminus j}(t,t') + \sum_{k\in\partial i\setminus j} J_{ik}J_{ki}\int_{t'}^{t} ds\, R_{k\setminus i}(t,s)R_{i\setminus j}(s,t') + \delta(t-t'), \tag{4.37}$$

$$\begin{aligned}\frac{\partial C_{i\setminus j}(t,t')}{\partial t} &= -\lambda_i C_{i\setminus j}(t,t') + \sum_{k\in\partial i\setminus j} J_{ik}J_{ki}\int_{0}^{t} ds\, R_{k\setminus i}(t,s)C_{i\setminus j}(s,t') \\ &\quad + \sum_{k\in\partial i\setminus j} J_{ik}^2\int_{0}^{t'} ds\, C_{k\setminus i}(t,s)R_{i\setminus j}(t',s) + 2TR_{i\setminus j}(t',t),\end{aligned} \tag{4.38}$$

$$\frac{d\mu_{i\setminus j}(t)}{dt} = -\lambda_i\mu_{i\setminus j}(t) + \sum_{k\in\partial i\setminus j} J_{ik}\mu_{k\setminus i}(t) + \sum_{k\in\partial i\setminus j} J_{ik}J_{ki}\int_{0}^{t} ds\, R_{k\setminus i}(t,s)\mu_{i\setminus j}(s), \tag{4.39}$$

where causality is enforced by the boundary condition $R_{i\setminus j}(t,t') = 0$ for $t \leq t'$. The single-node marginals can be obtained similarly, leading to equations analogous to Eqs. (4.37) to (4.39) but with sums over all neighbors ∂i instead of $\partial i \setminus j$.

$$\frac{\partial R_i(t,t')}{\partial t} = -\lambda_i R_i(t,t') + \sum_{k\in\partial i} J_{ik}J_{ki}\int_{t'}^{t} ds\, R_{k\setminus i}(t,s)R_i(s,t') + \delta(t-t'), \tag{4.40}$$

$$\begin{aligned}\frac{\partial C_i(t,t')}{\partial t} &= -\lambda_i C_i(t,t') + \sum_{k\in\partial i} J_{ik}J_{ki}\int_{0}^{t} ds\, R_{k\setminus i}(t,s)C_i(s,t') \\ &\quad + \sum_{k\in\partial i} J_{ik}^2\int_{0}^{t'} ds\, C_{k\setminus i}(t,s)R_i(t',s) + 2TR_i(t',t),\end{aligned} \tag{4.41}$$

$$\frac{d\mu_i(t)}{dt} = -\lambda_i\mu_i(t) + \sum_{k\in\partial i} J_{ik}\mu_{k\setminus i}(t) + \sum_{k\in\partial i} J_{ik}J_{ki}\int_{0}^{t} ds\, R_{k\setminus i}(t,s)\mu_i(s). \tag{4.42}$$

Relation with Plefka expansion We note here that the GECaM equations are an intuitive cavity generalization of dynamical TAP equations derived with the *extended Plefka expansion*

(Eqs. (3.52) and (3.53) of [38]). The main difference is that the Plefka equations involve only single-node quantities, while the GECaM equations involve cavity quantities that depend on the local neighborhood of each node. This makes the GECaM approach more accurate, as it captures local fluctuations and correlations that are neglected in the Plefka approximation. However, this increased accuracy comes at the cost of higher computational complexity, as the number of cavity messages grows with the number of edges in the graph. The original Plefka approach can be recovered from the GECaM equations by assuming that cavity quantities are independent of the removed neighbor, i.e. $\mu_{i\setminus j}(t) \approx \mu_i(t)$, $C_{i\setminus j}(t,t') \approx C_i(t,t')$, and $R_{i\setminus j}(t,t') \approx R_i(t,t')$. This assumption is valid in fully-connected graphs or in graphs with weak correlations between nodes, but it breaks down in sparse or strongly correlated systems.

4.3 Equilibrium and connection with random-matrix theory

In equilibrium conditions, the GECaM equations derived in section 4.2.2 simplify considerably and reveal a direct connection with the spectral theory of sparse random matrices. We focus on the stationary regime reached after an initial transient, where all moments become Time-Translationally Invariant (TTI). In this regime, we assume that the system has lost memory of its initial conditions, so that the initial time can be taken as $-\infty$ and all time integrals extend over the infinite past. The TTI assumption implies that the correlations and responses depend only on time differences,

$$\lim_{t,t'\to\infty} C(t,t') = C(\tau = t-t') = C(\tau),$$
$$\lim_{t,t'\to\infty} R(t,t') = R(\tau = t-t') = R(\tau).$$

The TTI GECaM equations for the cavity responses and correlations become

$$\frac{dR_{i\setminus j}(\tau)}{d\tau} = -\lambda_i R_{i\setminus j}(\tau) + \sum_{k\in\partial i\setminus j} J_{ik}J_{ki}\int_0^\tau ds\, R_{k\setminus i}(\tau-s)R_{i\setminus j}(s) + \delta(\tau), \tag{4.43}$$

$$\begin{aligned}\frac{dC_{i\setminus j}(\tau)}{d\tau} = &-\lambda_i C_{i\setminus j}(\tau) + \sum_{k\in\partial i\setminus j} J_{ik}J_{ki}\int_{-\infty}^\tau ds\, R_{k\setminus i}(\tau-s)C_{i\setminus j}(s) \\ &+ \sum_{k\in\partial i\setminus j} J_{ik}^2\int_\tau^\infty ds\, C_{k\setminus i}(s)R_{i\setminus j}(s-\tau) + 2TR_{i\setminus j}(-\tau).\end{aligned} \tag{4.44}$$

It is often convenient to rewrite Eq. (4.43) in the Laplace domain, by introducing the (one-sided) Laplace transformed response $\tilde{R}(z)$ and the (two-sided) Laplace transformed correlation $\tilde{C}(z)$

as

$$\tilde{R}(z) = \int_0^\infty \frac{d\tau}{2\pi} e^{-z\tau} R(\tau), \tag{4.45}$$

$$\tilde{C}(z) = \int_{-\infty}^\infty \frac{d\tau}{2\pi} e^{-z\tau} C(\tau). \tag{4.46}$$

Taking the Laplace transform of Eqs. (4.43) and (4.44) we obtain second-order algebraic equations for the transformed cavity response and correlation functions, which can be solved explicitly. The result is

$$\tilde{R}_{i\setminus j}(z) = \left(z + \lambda_i - \sum_{k\in\partial i\setminus j} J_{ik} J_{ki} \tilde{R}_{k\setminus i}(z) \right)^{-1}, \tag{4.47}$$

$$\tilde{C}_{i\setminus j}(z) = \tilde{R}_{i\setminus j}(z)\tilde{R}_{i\setminus j}(-z) \left(2T + \sum_{k\in\partial i\setminus j} J_{ik}^2 \tilde{C}_{k\setminus i}(z) \right). \tag{4.48}$$

The single-node marginals can be obtained similarly, leading to equations analogous to Eqs. (4.47) and (4.48) but with sums over all neighbors ∂i instead of $\partial i \setminus j$

$$\tilde{R}_i(z) = \left(z + \lambda_i - \sum_{k\in\partial i} J_{ik} J_{ki} \tilde{R}_{k\setminus i}(z) \right)^{-1}, \tag{4.49}$$

$$\tilde{C}_i(z) = \tilde{R}_i(z)\tilde{R}_i(-z) \left(2T + \sum_{k\in\partial i} J_{ik}^2 \tilde{C}_{k\setminus i}(z) \right). \tag{4.50}$$

In principle, these equations can be solved iteratively by initializing the cavity responses and correlations for each value of z in the complex plane. However, in practice, this approach is of limited utility, as recovering the time domain functions requires a numerical inversion of the Laplace transform, which is often computationally demanding and numerically delicate. When the interaction matrix $\underline{\underline{J}}$ is symmetric, i.e. $J_{ij} = J_{ji}$, the system satisfies detailed balance and reaches thermal equilibrium. In this case, the Fluctuation-Dissipation Theorem (FDT) relates the correlation and response functions as

$$TR^{\mathrm{eq}}(\tau) = -\frac{dC^{\mathrm{eq}}(\tau)}{d\tau}\Theta(\tau), \tag{4.51}$$

where $\Theta(\tau)$ is the Heaviside step function. Inserting Eq. (4.51) into Eq. (4.44), one finds that correlations satisfy a set of closed equations, which can be solved directly in the time domain without the need for Laplace transforms. The resulting equations for the cavity and single-node

correlations are

$$\operatorname{sgn}(\tau)\frac{dC^{\mathrm{eq}}_{i\backslash j}(\tau)}{d\tau} = -\lambda_i C^{\mathrm{eq}}_{i\backslash j}(\tau) + \sum_{k\in\partial i\backslash j}\frac{J^2_{ik}}{T}C^{\mathrm{eq}}_{k\backslash i}(\tau)C^{\mathrm{eq}}_{i\backslash j}(0) + \sum_{k\in\partial i\backslash j}\frac{J^2_{ik}}{T}\int_0^\tau ds\,\frac{dC^{\mathrm{eq}}_{k\backslash i}(s)}{ds}C^{\mathrm{eq}}_{i\backslash j}(\tau-s), \tag{4.52}$$

$$\operatorname{sgn}(\tau)\frac{dC^{\mathrm{eq}}_{i}(\tau)}{d\tau} = -\lambda_i C^{\mathrm{eq}}_{i}(\tau) + \sum_{k\in\partial i}\frac{J^2_{ik}}{T}C^{\mathrm{eq}}_{k\backslash i}(\tau)C^{\mathrm{eq}}_{i}(0) + \sum_{k\in\partial i}\frac{J^2_{ik}}{T}\int_0^\tau ds\,\frac{dC^{\mathrm{eq}}_{k\backslash i}(s)}{ds}C^{\mathrm{eq}}_{i}(\tau-s). \tag{4.53}$$

Further details on the derivation of the equilibrium GECaM equations, along with their numerical implementation, are provided in Appendix B.1.5.

4.3.1 Relation with random-matrix theory

In the preceding derivation, the GECaM equations were formulated for a fixed instance of the coupling matrix $\underline{\underline{J}}$, with elements $A_{ij}J_{ij}$ defined on a given sparse graph. However, in most disordered systems [8], either the graph structure or the couplings themselves are random, and $\underline{\underline{J}}$ must be regarded as a realization from an ensemble of sparse random matrices [49, 50]. We denote by $p_{\mathrm{dis}}(\underline{\underline{J}})$ the probability distribution defining this ensemble.

To highlight the connection between the equilibrium dynamics and random-matrix theory, it is convenient to rewrite the linear stochastic dynamics in vector form,

$$\frac{d\underline{x}(t)}{dt} = -\underline{\underline{\lambda}}\,\underline{x}(t) + \underline{\underline{J}}\,\underline{x}(t) + \underline{\xi}(t) + \underline{h}(t), \tag{4.54}$$

where $\underline{\underline{\lambda}}$ is the diagonal matrix of relaxation rates, with elements $\{\underline{\underline{\lambda}}\}_{ij} = \lambda_i\delta_{ij}$, and $\underline{h}(t)$ is a weak external field introduced to compute response functions. For an initial condition $\underline{x}(0) = \underline{x}_0$, the formal solution of Eq. (4.54) reads

$$\underline{x}(t) = e^{-\underline{\underline{\lambda}}t}\underline{x}_0 + \int_0^t dt'\,e^{(\underline{\underline{J}}-\underline{\underline{\lambda}})(t-t')}\left[\underline{\xi}(t') + \underline{h}(t')\right]. \tag{4.55}$$

In the long-time limit, provided that all $\lambda_i > 0$, the contribution from the initial condition vanishes. The dynamics is stable as long as all eigenvalues of the matrix $\underline{\underline{\tilde{J}}} = \underline{\underline{J}} - \underline{\underline{\lambda}}$ have

negative real parts. Under these conditions, the linear response matrix is obtained as

$$\underline{\underline{R}}(t,t') = \frac{\delta}{\delta \underline{h}(t')} \langle \underline{x}(t) \rangle \Big|_{\underline{h}=0} = \Theta(t-t')\, e^{(\underline{\underline{J}}-\underline{\underline{\lambda}})(t-t')}, \tag{4.56}$$

which is explicitly time-translationally invariant. Its one-sided Laplace transform is

$$\tilde{\underline{\underline{R}}}(z) = \int_0^\infty d\tau\, e^{-z\tau} e^{(\underline{\underline{J}}-\underline{\underline{\lambda}})\tau} = \left[z - (\underline{\underline{J}} - \underline{\underline{\lambda}})\right]^{-1}. \tag{4.57}$$

Hence, $\tilde{\underline{\underline{R}}}(z)$ is formally the resolvent of the random matrix $\tilde{\underline{\underline{J}}} = \underline{\underline{J}} - \underline{\underline{\lambda}}$. The diagonal entries $\tilde{R}_i(z)$ correspond to the Laplace transformed local response functions derived in Eq. (4.47), and satisfy the same recursive equations as the diagonal elements of the resolvent of sparse symmetric random matrices [51, 52]. For symmetric coupling matrices, $\tilde{\underline{\underline{J}}}$ has real eigenvalues ν_α and orthogonal eigenvectors $\underline{v}_\alpha$ satisfying $\tilde{\underline{\underline{J}}}\,\underline{v}_\alpha = \nu_\alpha \underline{v}_\alpha$. The spectral density of $\tilde{\underline{\underline{J}}}$ is obtained from the imaginary part of the trace of its resolvent via the inverse Stieltjes transform [49, 53, 54],

$$\rho_{\tilde{\underline{\underline{J}}}}(x) = \frac{1}{N}\sum_{\alpha=1}^{N} \delta\big(x - \nu_\alpha(\tilde{\underline{\underline{J}}})\big) = \frac{1}{\pi N} \lim_{\varepsilon\to 0^+} \mathrm{Im} \sum_{i=1}^{N} \tilde{R}_i(x - \mathrm{i}\varepsilon). \tag{4.58}$$

Since the eigenvalues of $\tilde{\underline{\underline{J}}}$ are related to those of $\underline{\underline{J}}$ by $\nu_\alpha(\tilde{\underline{\underline{J}}}) = \lambda_\alpha(\underline{\underline{J}}) - \lambda_\alpha$, the spectral density of $\underline{\underline{J}}$ follows from a simple shift,

$$\rho_{\underline{\underline{J}}}(x) = \frac{1}{N}\sum_{\alpha=1}^{N} \delta\big(x - \lambda_\alpha - \nu_\alpha(\tilde{\underline{\underline{J}}})\big) = \frac{1}{\pi N} \lim_{\varepsilon\to 0^+} \mathrm{Im} \sum_{i=1}^{N} \tilde{R}_i(x - \lambda_i - \mathrm{i}\varepsilon). \tag{4.59}$$

In disordered systems, one is typically interested not in the spectrum of a single realization of $\underline{\underline{J}}$, but in the spectral density averaged over the random-matrix ensemble defined by $p_{\mathrm{dis}}(\underline{\underline{J}})$. This ensemble-averaged spectral distribution is given by

$$\rho(x) = \overline{\rho_{\underline{\underline{J}}}(x)} = \frac{1}{\pi N} \lim_{\varepsilon\to 0^+} \mathrm{Im} \sum_{i=1}^{N} \overline{\tilde{R}_i(x - \lambda_i - \mathrm{i}\varepsilon)}, \tag{4.60}$$

where the overline denotes the average over disorder realizations. The stability condition derived above translates into a spectral constraint: the dynamics is stable if and only if the spectral density $\rho(x)$ is supported entirely on the left half of the complex plane, i.e. $\mathrm{Re}\, x < 0$.

For non-Hermitian coupling matrices, the same correspondence holds but the resolvent must be defined through a hermitization procedure [55, 56]. The adaptation of this technique to the present setting is discussed in Appendix B.2.

The connection between GECaM and random-matrix theory is therefore twofold. On the one hand, the cavity response functions derived from the GECaM equations encode the spectral properties of the effective interaction matrix $\underline{\underline{\tilde{J}}}$. On the other hand, once the spectral density $\rho(x)$ of the ensemble is known, one can reconstruct the disorder-averaged response and correlation functions of the system. To make this explicit, let us consider again the vectorial formulation of the linear dynamics in Eq. (4.54), whose solution is given by Eq. (4.55). From this formal solution, we can express the (disconnected) correlation matrix as

$$\underline{\underline{C}}^{\mathrm{dc}}(t,t') = \langle \underline{x}(t)^{\top} \underline{x}(t') \rangle = \int_0^t dt_1 \int_0^{t'} dt_2 \, e^{(\underline{\underline{J}}-\underline{\underline{\lambda}})(t-t_1)+(\underline{\underline{J}}^{\top}-\underline{\underline{\lambda}})(t'-t_2)} \langle \underline{\eta}(t_1)\underline{\eta}(t_2) \rangle \tag{4.61a}$$

$$= 2T \int_0^{\min(t,t')} ds \, e^{(\underline{\underline{J}}-\underline{\underline{\lambda}})(t-s)+(\underline{\underline{J}}^{\top}-\underline{\underline{\lambda}})(t'-s)}. \tag{4.61b}$$

Its diagonal entries correspond to the local (disconnected) correlations $C_i^{\mathrm{dc}}(t,t') = \langle x_i(t) x_i(t') \rangle$. For symmetric coupling matrices $\underline{\underline{J}}$, the correlator $\underline{\underline{C}}^{\mathrm{dc}}(t,t')$ is symmetric as well and becomes explicitly time-translationally invariant in the stationary regime,

$$\underline{\underline{C}}^{\mathrm{dc}}(\tau) = T \int_{|\tau|}^{\infty} dw \, e^{(\underline{\underline{J}}-\underline{\underline{\lambda}})w}, \tag{4.62}$$

where $\tau = t - t'$ and $w = t + t' - 2s$. In this case, the FDT holds as a consequence of detailed balance, yielding

$$T \, \underline{\underline{R}}(\tau) = -\frac{d\underline{\underline{C}}^{\mathrm{dc}}(\tau)}{d\tau} \Theta(\tau), \tag{4.63}$$

which can be verified directly by comparing Eq. (4.56) and Eq. (4.62).

The ensemble-averaged response and correlation functions can be obtained directly from the spectral density of the random-matrix ensemble. Let $\rho(x)$ denote the average spectral distribution of $\underline{\underline{J}}$, to which the empirical spectral density $\rho_{\underline{\underline{J}}}(x)$ converges in the thermodynamic limit. Averaging over the disorder, one finds

$$\overline{R(\tau)} = \frac{1}{N} \sum_i \overline{R_i(\tau)} = \mathrm{Tr} \, \overline{\underline{\underline{R}}(\tau)} = \int dx \, \rho(x) \, \Theta(\tau) \, e^{(x-\lambda)\tau}, \tag{4.64}$$

$$\overline{C^{\mathrm{dc}}(\tau)} = \frac{1}{N} \sum_i \overline{C_i^{\mathrm{dc}}(\tau)} = \mathrm{Tr} \, \overline{\underline{\underline{C}}^{\mathrm{dc}}(\tau)} = T \int dx \, \rho(x) \int_{|\tau|}^{\infty} dw \, e^{(x-\lambda)w}, \tag{4.65}$$

where we assumed homogeneous relaxation rates $\lambda_i = \lambda$ for simplicity. These expressions show explicitly that the disorder-averaged linear dynamics is fully determined by the spectral density of the interaction matrix. In this sense, the equilibrium GECaM formalism and random-matrix theory provide two complementary perspectives on the same structure: the former describes

the stochastic dynamics in the time domain, while the latter characterizes the same process in the spectral domain.

4.4 Linear dynamics with thermal noise on random regular graphs

The GECaM equations for response and correlation functions can in principle be defined for any specific realization of the interaction network. Here we focus on the case of RRGs with homogeneous interactions and parameters, where each node has the same degree K, all couplings are equal, $J_{ij} = J_{ji} = J$, and relaxation rates are uniform, $\lambda_i = \lambda$. In this setting, it is impossible to distinguish between nodes or edges, as all local environments are statistically equivalent. To obtain analytical results, we consider the thermodynamic limit, where the graph becomes locally tree-like. In this limit, we can drop the edge and node indices and denote the disorder-averaged cavity quantities by R_c and C_c, and the single-node ones by R and C.

Under these assumptions, the Laplace transformed GECaM equations for a linear dynamics with additive thermal noise simplify to

$$\tilde{R}_c(z) = \left(z + \lambda - (K-1)J^2\tilde{R}_c(z)\right)^{-1}, \tag{4.66}$$

$$\tilde{C}_c(z) = \tilde{R}_c(z)\tilde{R}_c(-z)\left(2T + (K-1)J^2\tilde{C}_c(z)\right). \tag{4.67}$$

Solving Eq. (4.66) yields an explicit expression for the cavity response in Laplace space,

$$\tilde{R}_c^{\pm}(z) = \frac{z+\lambda \pm \sqrt{(z+\lambda)^2 - 4(K-1)J^2}}{2(K-1)J^2}. \tag{4.68}$$

The physical branch is fixed by the short-time limit of the response function, which requires $R(\tau) \to 1$ as $\tau \to 0$, corresponding to $\tilde{R}(z) \sim 1/z$ for large z. This condition selects the negative branch $\tilde{R}_c^{-}(z)$, which from now on we denote simply by $\tilde{R}_c(z)$.

The single-node response $\tilde{R}(z)$ follows from the corresponding self-consistency equation,

$$\begin{aligned}\tilde{R}(z) &= \frac{1}{z+\lambda - KJ^2\tilde{R}_c(z)} \\ &= \frac{1}{2}\frac{(K-2)(z+\lambda) - K\sqrt{(z+\lambda)^2 - 4(K-1)J^2}}{K^2J^2 - (z+\lambda)^2},\end{aligned} \tag{4.69}$$

where again the negative branch ensures the correct asymptotic behaviour.

The cavity correlation $\tilde{C}_c(z)$ can be obtained from Eq. (4.67) as

$$\tilde{C}_c(z) = \frac{2T\,\tilde{R}_c(z)\tilde{R}_c(-z)}{1-(K-1)J^2\tilde{R}_c(z)\tilde{R}_c(-z)}, \tag{4.70}$$

which, upon inserting Eq. (4.68), gives

$$\tilde{C}_c(z) = \frac{T}{(K-1)J^2}\left(\frac{\sqrt{(z+\lambda)^2-4(K-1)J^2}}{2z} - \frac{\sqrt{(z-\lambda)^2-4(K-1)J^2}}{2z} - 1\right). \tag{4.71}$$

The full correlation function $\tilde{C}(z)$ is then obtained from the single-node equation,

$$\tilde{C}(z) = \tilde{R}(z)\tilde{R}(-z)\left(2T + KJ^2\tilde{C}_c(z)\right). \tag{4.72}$$

The corresponding spectral distribution of the adjacency matrix can also be recovered directly from the GECaM solution. For a random regular graph of degree K, the eigenvalue density of the adjacency matrix converges in the thermodynamic limit to the *Kesten-McKay distribution* [57, 58],

$$\rho(x) = \frac{K\sqrt{4(K-1)-x^2}}{2\pi(K^2-x^2)}, \qquad |x| \le 2\sqrt{K-1}. \tag{4.73}$$

Using Eq. (4.59) and the GECaM response from Eq. (4.69), we can reconstruct the same spectral law as

$$\begin{aligned}\rho(x) &= \frac{J}{\pi}\lim_{\varepsilon\to 0^+}\operatorname{Im}\tilde{R}(Jx-\lambda-\mathrm{i}\varepsilon)\\ &= \frac{J}{\pi}\frac{\left[(x^2-4(K-1))^2\right]^{1/4}}{2J(K^2-x^2)}\sin\left[\frac{1}{2}\arg\left(x^2-4(K-1)\right)\right],\end{aligned} \tag{4.74}$$

which reduces exactly to the Kesten-McKay form for $|x| \le 2\sqrt{K-1}$.

Conversely, starting from the spectral density Eq. (4.73), the response function can be recovered through the integral representation of the resolvent,

$$\tilde{R}(z) = \int dx\,\rho(x)\,[z-(Jx-\lambda)]^{-1}, \tag{4.75}$$

which yields after direct evaluation

$$\tilde{R}(z) = \frac{1}{2}\frac{(K-2)(z+\lambda) - K\sqrt{(z+\lambda)^2-4(K-1)J^2}}{K^2J^2-(z+\lambda)^2}, \tag{4.76}$$

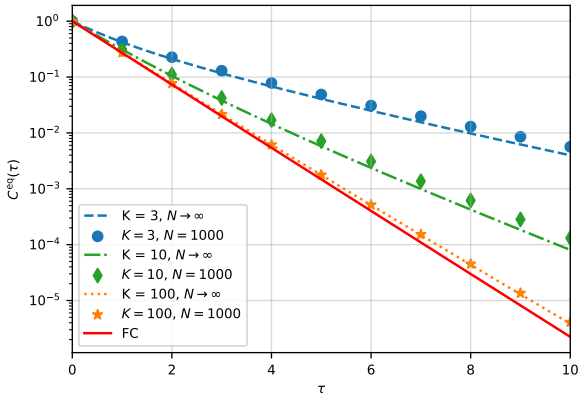

Fig. 4.1 **Equilibrium correlation on RRGs.** Comparison of the equilibrium correlation function C^{eq} in RRGs with different degrees K and in the Fully-Connected (FC) limit. Theoretical predictions from the Kesten-McKay spectral density and Eq. (4.65) (solid lines) are compared with numerical solutions of the GECaM equations on finite graphs with $N = 1000$ nodes (symbols). Parameters: $\lambda_i = \lambda = 1.3$, $J = 1/K$, $T = 1$.

identical to the expression derived via GECaM, thus confirming the consistency between the two approaches.

Figure 4.1 compares the average equilibrium correlation $C^{\text{eq}}(\tau)$ obtained from the analytic solution in the thermodynamic limit–i.e. by solving Eq. (4.65) with the Kesten-McKay spectral density–with the results from direct numerical integration of the equilibrium GECaM Eqs. (4.52) and (4.53) on finite RRGs. While the fully-connected case exhibits purely exponential relaxation, finite connectivity introduces a slower decay at long times, signaling weak non-exponential relaxation induced by network sparseness. This deviation is consistently reproduced by both the analytical and numerical solutions, confirming that the GECaM framework correctly captures the impact of finite connectivity on dynamical relaxation.

4.5 Linear dynamics with thermal noise on heterogeneous graphs

To assess the generality of the method beyond regular topologies, we applied GECaM to systems defined on heterogeneous sparse networks, considering both ferromagnetic and disordered couplings. In particular, we analysed graphs with bimodal interactions $\{J_{ij}\}_{(i,j)\in E}$

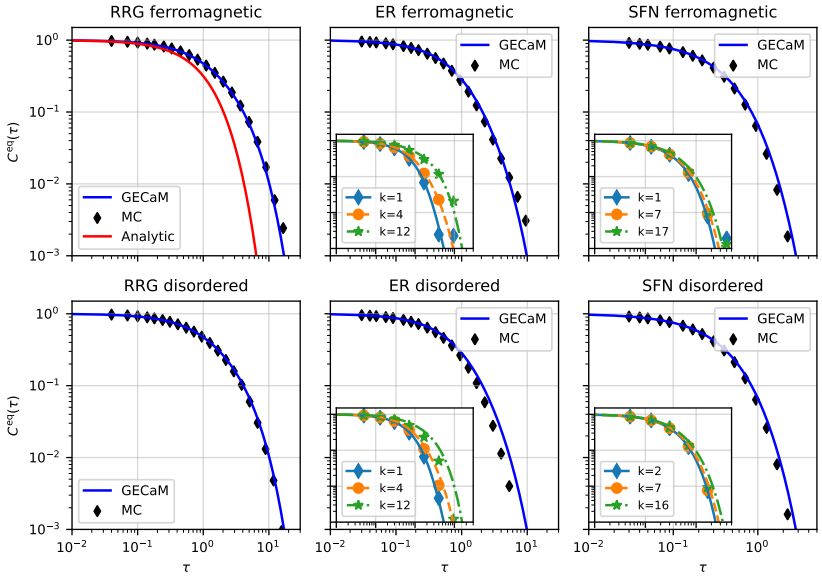

Fig. 4.2 **Equilibrium correlation on heterogeneous graph topologies.** Equilibrium correlation functions obtained with GECaM (solid lines) and Monte Carlo simulations (diamonds) for different topologies and interaction types. We compare RRGs ($N = 200$, $K = 3$, $\lambda = 1.2$), ERGs ($N = 100$, $K = 4$, $\lambda = 1.7$), and SFNs ($N = 200$, $K = 7$, $\alpha = 2.5$, $\lambda = 3.0$), all with $|J| = 1/K$ and $T = 1$. The analytical result for the ferromagnetic RRG in the thermodynamic limit (solid red line)–obtained from the Kesten-McKay spectral density and Eq. (4.65)–is also shown. Insets: equilibrium correlation for nodes with minimum, average, and maximum degree. Axes are on the same scale as the main plot.

independently drawn from the symmetric distribution

$$p(J) = \frac{1}{2}\delta\left(J - \frac{1}{K}\right) + \frac{1}{2}\delta\left(J + \frac{1}{K}\right),$$

as well as purely ferromagnetic systems with $J = 1/K$. We considered three classes of network topologies: RRGs with degree K, ERGs with average degree $\langle k \rangle = K$, and SFNs with degree distribution following a power law [59]. In ERGsthe degree distribution is Poissonian, whereas for SFNs we generated networks with exponent α and total number of edges $KN/2$, so that the average degree equals $\langle k \rangle = K$.

The equilibrium correlation functions were computed by solving the equilibrium GECaM Eqs. (4.52) and (4.53), and compared with Monte Carlo (MC) simulations of the corresponding Langevin dynamics. In the ferromagnetic RRG case, we also included the analytical result obtained in the thermodynamic limit from the Kesten-McKay spectral density and Eq. (4.65). The GECaM predictions reproduce the MC results with high accuracy, including finite-size corrections that are absent from the ensemble-averaged theoretical solution.

Results for heterogeneous graphs are shown in Fig. 4.2. While the overall agreement between GECaM and MC simulations is excellent, small deviations appear in the presence of disorder, particularly for the Erdős-Rényi case. To pinpoint their origin, we examined the equilibrium correlation functions of representative nodes with different degrees (minimum, average, and maximum). As illustrated in the inset panels, discrepancies are mainly associated with high-degree nodes, i.e. those belonging to the tail of the degree distribution. This behavior is plausibly due to the locally tree-like assumption underlying the cavity approximation: while it holds asymptotically in the thermodynamic limit, finite-size graphs with hubs or large-degree nodes generate additional correlations not captured by the approximation.

Chapter 5

Perturbative Extensions for Nonlinear Dynamics

The GECaM approach provides an exact description of linear stochastic dynamics with thermal noise, but it becomes approximate when nonlinearities in the drift or multiplicative noise are present. In these cases, the effective cavity measure is no longer Gaussian, and a systematic expansion beyond the Gaussian level is required to capture the leading non-Gaussian corrections. In this chapter we develop a perturbative closure scheme that extends the Gaussian approximation to weakly nonlinear stochastic processes. Section 5.1 introduces the theoretical framework for this perturbative expansion, while the following sections apply it to two representative cases: a cubic nonlinearity in the local drift (section 5.2), and a multiplicative noise model inspired by the Bouchaud-Mézard dynamics of wealth exchange (section 5.3). The presentation of perturbation theory in this chapter is intentionally non-exhaustive, focusing on the essential concepts and methods relevant for the subsequent applications. Readers interested in a more detailed and rigorous treatment of statistical field theory and perturbative techniques are referred to the standard texts [60, 61].

5.1 Perturbative Expansion for Non-Gaussian Dynamics

The dynamics, and consequently the cavity measure, ceases to be Gaussian when nonlinear drift terms or multiplicative noise are present in the underlying stochastic equations. To make

this explicit, we consider the generic coupled Langevin dynamics

$$\frac{dx_i(t)}{dt} = -\lambda_i x_i(t) + \varepsilon f_i(x_i(t)) + \alpha \sum_{j\neq i} J_{ij} x_j(t) + \xi_i(t), \tag{5.1}$$

with Gaussian noise $\xi_i(t)$ characterized by covariance

$$\langle \xi_i(t)\xi_{i'}(t')\rangle = 2\big[T + \varepsilon g_i(x_i(t))\big]\delta_{ii'}\delta(t-t'), \tag{5.2}$$

where the functions $f_i(x_i)$ and $g_i(x_i)$ describe, respectively, the nonlinear drift and multiplicative noise contributions. The small parameter ε controls the strength of these non-Gaussian terms and provides a natural perturbative parameter for expanding the cavity measure around its Gaussian limit[1].

5.1.1 Overview of the expansion strategy

Before proceeding with the detailed derivation of the perturbative expansion, we outline here the main steps of the procedure. The starting point is the decomposition of the local effective action into two components: a Gaussian part S^0, arising from the linear drift terms and thermal noise in the Langevin equation, and a perturbative part S^{p}, which encodes the contributions from nonlinearities and non-thermal (multiplicative) noise. A small parameter ε is introduced to track these non-Gaussian terms and systematically control the expansion.

Crucially, the framework relies on the approximation that the cavity messages remain Gaussian throughout the recursion. This implies they are fully parametrized by their first two moments (means, response, and correlation functions), which are iteratively updated to capture the effects of S^{p}. The resulting self-consistent procedure performed for every edge (i,j) is summarized as follows:

1. **Input (Perturbed Moments):** The incoming cavity messages are assumed to be Gaussian distributions parametrized by *renormalized* propagators $\{\mathrm{G}_{k\setminus i}\}_{k\in\partial i\setminus j}$ and means $\{\mathrm{M}_{k\setminus i}\}_{k\in\partial i\setminus j}$. These propagators already contain the corrections due to nonlinearities computed in the previous iteration on the neighboring branches.

2. **Linear Combination (Unperturbed State):** These renormalized inputs are combined using the standard linear GECaM Eqs. (4.26) and (4.27). This step yields the *unperturbed*

[1] In principle, two distinct expansion parameters, ε_f and ε_g, should be introduced to separately account for the perturbations in f and g. For simplicity of notation, we use a single parameter ε, without any loss of generality.

cavity propagator $\mathrm{G}^0_{i\backslash j}$ and mean $\mathrm{M}^0_{i\backslash j}$, representing the effective Gaussian process experienced by the node before the action of its local nonlinearity.

3. **Local Renormalization (Perturbed State):** The unperturbed propagator and mean are locally corrected to account for the perturbative action S_i^{p}. Using a diagrammatic expansion and the Dyson equation (closed at the Hartree-Fock level), we calculate the self-energy contributions that transform $\mathrm{G}^0_{i\backslash j}$ into the final *renormalized* propagator $\mathrm{G}_{i\backslash j}$ and $\mathrm{M}^0_{i\backslash j}$ into the *renormalized* mean $\mathrm{M}_{i\backslash j}$. This updated Gaussian parametrization, which now effectively includes the local non-Gaussian effects, becomes the message passed to the next node.

5.1.2 Derivation of the perturbative cavity equations

The dynamic cavity equations for the dynamics Eq. (5.1) can be derived following the same steps as in section 4.1, leading to the recursive equations

$$c_{i\backslash j}[\mathrm{X}_i] = \frac{p_0(x_i(0))}{Z_{i\backslash j}} e^{-S_i^{\mathrm{loc}}[\mathrm{X}_i] - \varepsilon S_i^{\mathrm{p}}[\mathrm{X}_i]} \prod_{k\in\partial i\backslash j} \int \mathscr{D}[\mathrm{X}_k]\, c_{k\backslash i}[\mathrm{X}_k]\, e^{-\alpha S_{ij}^{\mathrm{int}}[\mathrm{X}_i,\mathrm{X}_k]}, \tag{5.3}$$

where we have used the shorthand notations $\mathrm{X} = (x(t), \hat{x}(t))$ and $\mathscr{D}[\mathrm{X}] = \mathscr{D}[x,\hat{x}]$ to denote, respectively, the trajectory $x(t)$ and its conjugate variable $\hat{x}(t)$ and their functional measure. The local action $S_i^{\mathrm{loc}}[\mathrm{X}_i]$ and the interaction action $S_{ij}^{\mathrm{int}}[\mathrm{X}_i,\mathrm{X}_j]$ are given by Eqs. (4.5) and (4.6), while the perturbative action $S_i^{\mathrm{p}}[\mathrm{X}_i]$ reads

$$S_i^{\mathrm{p}}[\mathrm{X}_i] = -\int dt \left(\mathrm{i}\hat{x}_i(t) f_i(x_i(t)) + g_i(x_i(t))(\mathrm{i}\hat{x}_i(t))^2 \right). \tag{5.4}$$

By comparing Eq. (5.3) with the Gaussian cavity equation Eq. (4.7), we see that the perturbative action $S_i^{\mathrm{p}}[\mathrm{X}_i]$ introduces non-Gaussian contributions to the cavity measure, making it analytically intractable. Expanding the interacting exponential in Eq. (5.3) to second order in α we obtain as usual a non-interacting[2] cavity measure

$$c_{i\backslash j}[\mathrm{X}_i] \approx \frac{p_0(x_i(0))}{Z^{\mathrm{eff}}_{i\backslash j}} \exp\left(-S^0_{i\backslash j}[\mathrm{X}_i] - \varepsilon S_i^{\mathrm{p}}[\mathrm{X}_i] \right), \tag{5.5}$$

[2]The expression non-interacting does not refer to the global unperturbed limit as in conventional perturbation theory. In this context, the expansion is applied locally to each cavity measure, which after the weak-coupling expansion corresponds to an effective non-interacting process, subsequently corrected by local perturbations due to nonlinearities.

where the effective Gaussian action $S^0_{i\setminus j}[\mathrm{X}_i]$ is given by Eq. (4.17), or in a compact form as

$$S^0_{i\setminus j}[\mathrm{X}_i] = \frac{1}{2}\int dtdt' \left(\mathrm{X}_i(t) - \mathrm{M}_{i\setminus j}(t)\right)\left[\mathrm{G}^0_{i\setminus j}\right]^{-1}(t,t')\left(\mathrm{X}_i(t') - \mathrm{M}_{i\setminus j}(t')\right)^\top, \tag{5.6}$$

The effective normalization constant $Z^{\mathrm{eff}}_{i\setminus j}$ ensures that the cavity measure is normalized to one.

For notational simplicity, in what follows we set $\mathrm{M}_{i\setminus j} = 0$. When local averages are non-zero, the same construction applies by shifting to centered fields, $\delta x_i(t) = x_i(t) - \mu_{i\setminus j}(t)$ and $\delta\mathrm{X}_i(t) = (\delta x_i(t), \mathrm{i}\hat{x}_i(t))$, so that all propagators are understood as covariances of $\delta\mathrm{X}_i$ rather than raw second moments. No other modification is required.

After integrating over neighboring trajectories, the effective cavity measure in Eq. (5.5) retains the same non-interacting structure as in the linear case, but it is no longer quadratic in the fields. The presence of the perturbative action S^{p}_i renders the cavity measure non-Gaussian, so that averages of observables can no longer be computed in closed form. A systematic expansion in powers of ε can nevertheless be carried out, expressing the non-Gaussian corrections as cumulants of the Gaussian measure defined by the quadratic action $S^0_{i\setminus j}$. For a generic observable $\mathscr{O}[\mathrm{X}_i]$, its average with respect to the cavity measure can be expressed as

$$\langle\mathscr{O}[\mathrm{X}_i]\rangle_{i\setminus j} = \frac{\int\mathscr{D}[\mathrm{X}_i]\, e^{-S^0_{i\setminus j}[\mathrm{X}_i] - \varepsilon S^{\mathrm{p}}_i[\mathrm{X}_i]}\,\mathscr{O}[\mathrm{X}_i]}{\int\mathscr{D}[\mathrm{X}_i]\, e^{-S^0_{i\setminus j}[\mathrm{X}_i] - \varepsilon S^{\mathrm{p}}_i[\mathrm{X}_i]}} = \frac{\langle\mathscr{O}[\mathrm{X}_i]\, e^{-\varepsilon S^{\mathrm{p}}_i[\mathrm{X}_i]}\rangle^0_{i\setminus j}}{\langle e^{-\varepsilon S^{\mathrm{p}}_i[\mathrm{X}_i]}\rangle^0_{i\setminus j}}, \tag{5.7}$$

where averages $\langle\cdot\rangle^0_{i\setminus j}$ are taken with respect to the Gaussian measure. The average in Eq. (5.7) can be expanded in powers of ε using the cumulant expansion, yielding

$$\langle\mathscr{O}\rangle_{i\setminus j} = \langle\mathscr{O}\rangle^0_{i\setminus j} + \sum_{n=1}^{\infty}\frac{(-\varepsilon)^n}{n!}\kappa^0_{i\setminus j}\Big(\mathscr{O}, \underbrace{S^{\mathrm{p}}_i, \ldots, S^{\mathrm{p}}_i}_{n\text{ times}}\Big), \tag{5.8}$$

where $\kappa^0_{i\setminus j}\left(\mathscr{O}, S^{\mathrm{p}}_i, \ldots, S^{\mathrm{p}}_i\right)$ denotes the joint cumulant of order $n+1$ between the observable $\mathscr{O}$ and n copies of the perturbative action S^{p}_i, all computed with respect to the Gaussian measure. They correspond to connected correlation functions involving $\mathscr{O}$ and S^{p}_i. For example, stopping the expansion at first order in ε

$$\langle\mathscr{O}\rangle_{i\setminus j} \approx \langle\mathscr{O}\rangle^0_{i\setminus j} - \varepsilon\left(\langle\mathscr{O}S^{\mathrm{p}}_i\rangle^0_{i\setminus j} - \langle\mathscr{O}\rangle^0_{i\setminus j}\langle S^{\mathrm{p}}_i\rangle^0_{i\setminus j}\right). \tag{5.9}$$

The cumulants in Eq. (5.8) can be computed using Wick's theorem, which expresses higher order moments of Gaussian variables in terms of their two-point correlations. The perturbative expansion in Eq. (5.8) thus provides a systematic way to compute corrections to Gaussian aver-

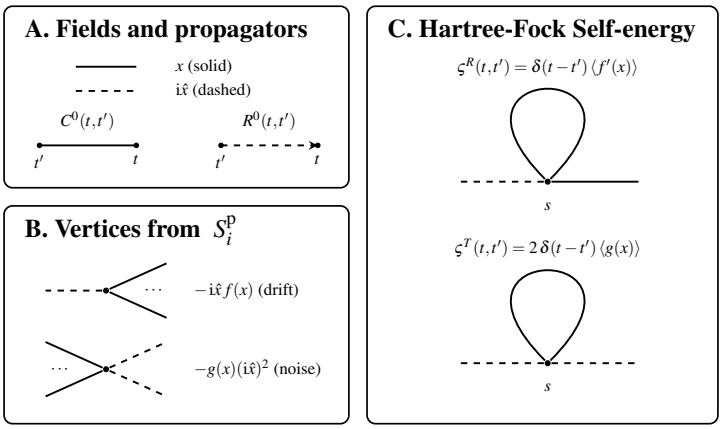

Fig. 5.1 **Feynman diagrams for the perturbative closure. Panel A**: Diagrammatic representation of the fields and bare propagators of the Gaussian theory. The response function is represented by a dashed arrow pointing from time t' to t, visually indicating causality. **Panel B**: Vertices generated by the perturbative action S_i^{p}, with the many x legs of $f(x)$ and $g(x)$ represented by two explicit x legs and dots. **Panel C**: Hartree-Fock (tadpole) self-energies: the drift tadpole renormalizes the $(x, \mathrm{i}\hat{x})$ block via ς^R, and the noise tadpole renormalizes the $(\hat{x}, \hat{x})$ block via ς^T. The tadpole diagrams shown correspond to the case of a cubic nonlinearity $f(x) \propto x^3$ and a quadratic multiplicative noise $g(x) \propto x^2$.

ages due to weak nonlinearities and multiplicative noise. These corrections can be represented diagrammatically using Feynman diagrams, as we now discuss.

Feynman diagrams In the MSRJD formalism, for each degree of freedom i there are two types of fields: the original variable $x_i(t)$ and its conjugate $\hat{x}_i(t)$. In the diagrammatic representation, these fields are depicted as lines with different styles: solid lines represent $x_i(t)$, while dashed lines represent $\mathrm{i}\hat{x}_i(t)$ (the factor i is included to ensure that the response function is real). The bare propagators of the Gaussian theory are represented as follows:

- The correlation function $C^0_{i\setminus j}(t,t') = \langle x_i(t) x_i(t') \rangle^0_{i\setminus j}$ is represented by a solid line connecting two points at times t and t'.
- The response function $R^0_{i\setminus j}(t,t') = \langle x_i(t) \mathrm{i}\hat{x}_i(t') \rangle^0_{i\setminus j}$ is represented by a dashed line with an arrow pointing from t' to t, indicating causality (the response at time t depends only on perturbations at earlier times $t' < t$).
- The conjugate-conjugate correlator $\langle \mathrm{i}\hat{x}_i(t) \mathrm{i}\hat{x}_i(t') \rangle^0_{i\setminus j}$ vanishes and is not represented.

The perturbative action $S_i^{\mathrm{p}}[\mathrm{X}_i]$ introduces vertices in the diagrams, corresponding to the nonlinear terms in the dynamics. Each vertex is associated with a specific time s and has legs corresponding to the fields involved in the interaction:

- The drift vertex $-\mathrm{i}\hat{x}_i(s) f_i(x_i(s))$ has one dashed leg (for $\mathrm{i}\hat{x}_i(s)$) and multiple solid legs (for the powers of $x_i(s)$ appearing in f_i when Taylor expanded).
- The noise vertex $-g_i(x_i(s))(\mathrm{i}\hat{x}_i(s))^2$ has two dashed legs (for $\mathrm{i}\hat{x}_i(s)^2$) and multiple solid legs (for the powers of $x_i(s)$ in g_i when Taylor expanded).

Figure 5.1 illustrates the basic elements of the diagrammatic representation, including the propagators and vertices.

self-energy and Dyson equation The perturbative expansion in Eq. (5.8) can be viewed as a systematic sum over all possible contractions of the vertices generated by the perturbative action S_i^{p} with the fields in the observable $\mathscr{O}$. Each term corresponds to a specific diagram, obtained by connecting vertices through the bare propagators of the Gaussian theory. When the observable is the two-point function $\mathscr{O} = \mathrm{X}_i(t)^\top \mathrm{X}_i(t')$–whose average over the perturbed distribution is the renormalized propagator–the diagrams contributing to its expansion can be reorganized in terms of the *self-energy* $\Sigma_{i\setminus j}$.

The self-energy is defined as the sum of all *one-particle irreducible* (1PI) two-point diagrams–those that remain connected when any single propagator line is cut. These diagrams represent the elementary irreducible corrections to the Gaussian propagator. All remaining (reducible) contributions, obtained by chaining 1PI blocks with Gaussian propagators, can be resummed algebraically, yielding

$$\mathrm{G}_{i\setminus j} = \mathrm{G}^0_{i\setminus j} + \mathrm{G}^0_{i\setminus j}\Sigma_{i\setminus j}\mathrm{G}^0_{i\setminus j} + \mathrm{G}^0_{i\setminus j}\Sigma_{i\setminus j}\mathrm{G}^0_{i\setminus j}\Sigma_{i\setminus j}\mathrm{G}^0_{i\setminus j} + \cdots, \tag{5.10}$$

where each product represents a time convolution over intermediate arguments, e.g.

$$[\mathrm{G}\Sigma\mathrm{G}](t,t') = \int ds\, ds'\, \mathrm{G}(t,s)\,\Sigma(s,s')\,\mathrm{G}(s',t').$$

Factoring out $\mathrm{G}^0_{i\setminus j}$ on the left gives the operator identity

$$\mathrm{G}_{i\setminus j} = \mathrm{G}^0_{i\setminus j} + \mathrm{G}^0_{i\setminus j}\Sigma_{i\setminus j}\mathrm{G}_{i\setminus j}, \tag{5.11}$$

which, upon inversion, leads to the Dyson equation

$$[\mathrm{G}_{i\setminus j}]^{-1} = [\mathrm{G}^0_{i\setminus j}]^{-1} - \Sigma_{i\setminus j}. \tag{5.12}$$

The Dyson equation compactly expresses the infinite resummation of all reducible two-point diagrams, with $\Sigma_{i\setminus j}$ acting as the kernel that encodes the irreducible corrections to the Gaussian dynamics.

The self-energy itself can be decomposed in the $(x, \mathrm{i}\hat{x})$ basis as

$$\Sigma_{i\setminus j}(t,t') = \begin{pmatrix} \varsigma^C_{i\setminus j}(t,t') & \varsigma^R_{i\setminus j}(t,t') \\ \varsigma^R_{i\setminus j}(t',t) & \varsigma^T_{i\setminus j}(t,t') \end{pmatrix}, \tag{5.13}$$

where the off diagonal term ς^R renormalizes the drift (response sector), while ς^T acts as an additional noise kernel driving correlations. Projecting Eq. (5.12) onto the response and correlation sectors, and using the causality of the response functions, it turns out that ς^C must vanish identically. The Dyson equations for the response and correlation functions thus read

$$\left[R_{i\setminus j}\right]^{-1}(t,t') = \left[R^0_{i\setminus j}\right]^{-1}(t,t') - \varsigma^R_{i\setminus j}(t,t'), \tag{5.14}$$

$$\begin{aligned} C_{i\setminus j}(t,t') = &\int ds_1 ds_2 ds_3 ds_4 \, R_{i\setminus j}(t,s_1)\left[R^0_{i\setminus j}\right]^{-1}(s_1,s_3) C_{i\setminus j}(s_3,s_4)\left[R^0_{i\setminus j}\right]^{-1}(s_2,s_4) \\ &+ \int ds_1 ds_2 \, R_{i\setminus j}(t,s_1)\varsigma^T_{i\setminus j}(s_1,s_2) R_{i\setminus j}(t',s_2). \end{aligned} \tag{5.15}$$

Causality is enforced by $R_{i\setminus j}(t,t') = 0$ for $t \leq t'$. In the Gaussian limit, i.e. $\varepsilon = 0$, $\varsigma^R = \varsigma^T = 0$ and one recovers $R = R^0$ and $C = C^0$. Further details on the derivation of the Dyson equation and its components are provided in Appendix C.1.

Hartree-Fock Approximation The self-energy expansion naturally inherits the perturbative structure in powers of ε, as each vertex in S^{p}_i contributes one factor of ε. Thus,

$$\Sigma_{i\setminus j} = \varepsilon \Sigma^{(1)}_{i\setminus j} + \varepsilon^2 \Sigma^{(2)}_{i\setminus j} + \cdots. \tag{5.16}$$

At first order the only 1PI contributions are *tadpoles*, which are local in time, proportional to $\delta(t-t')$, and renormalize the coefficients of the linear dynamics. At second order one encounters *sunset* diagrams, which are nonlocal in time and generate memory kernels and colored effective noise. Higher orders follow the same logic, each contributing increasingly complex nonlocal corrections to the effective propagator. Figure 5.1 shows the first-order tadpole diagrams.

Truncating the expansion to the tadpole level defines the *Hartree-Fock (or mean-field) approximation*, in which the self-energy is time-local and depends only on equal-time correlators. Its components ς^R and ς^T can be computed explicitly from the first-order (tadpole) diagrams,

yielding

$$\varsigma^R_{i\setminus j}(t,t') = \delta(t-t')\left\langle f_i'(x_i(t))\right\rangle_{i\setminus j}, \tag{5.17}$$

$$\varsigma^T_{i\setminus j}(t,t') = 2\,\delta(t-t')\left\langle g_i(x_i(t))\right\rangle_{i\setminus j}, \tag{5.18}$$

where $f_i'(x) = df_i(x)/dx$. The averages are evaluated self-consistently with the renormalized measure defined by Eq. (5.12). The presence of the derivative $f_i'(x)$ in ς^R can be understood directly from the diagrammatic expansion. In the drift contribution to the perturbative action, $\int dt\, \mathrm{i}\hat{x}_i(t)\, f_i(x_i(t))$, the function $f_i(x)$ must be interpreted as its Taylor series in x, so that each power of x represents one solid leg emerging from the vertex. To obtain a two-point (1PI) correction to the propagator, one must leave exactly two external legs uncontracted. The only surviving contractions–consistent with causality, i.e. $R(t,t') = 0$ for $t \leq t'$–are those that leave one $\mathrm{i}\hat{x}$ and one x leg from the Taylor expansion of f. Removing one x field from the expansion is therefore equivalent to taking the derivative of f, which explains the appearance of $f_i'(x)$ in Eq. (5.17). Equivalently, when expanding S_f to identify bilinear terms in $\mathrm{i}\hat{x}$ and x, the effective coupling is $-\mathrm{i}\hat{x}\, f_i'(x)$. The nonlinear drift function thus renormalizes the response function through the average of its derivative. A similar reasoning applies to the noise vertex $-g_i(x)(\mathrm{i}\hat{x}_i)^2$, which must also be expanded in powers of x. Here, the only non-vanishing two-point diagram compatible with causality and the constraint $\langle \mathrm{i}\hat{x}(t)\mathrm{i}\hat{x}(t')\rangle = 0$ is the one where the two $\mathrm{i}\hat{x}$ legs remain external and all the x legs are internally contracted. This yields a local contribution to the $(\hat{x},\hat{x})$ block, $\varsigma^T \propto \langle g_i(x)\rangle\, \delta(t-t')$, renormalizing the effective noise amplitude.

Figure 5.1 summarizes the main elements of the diagrammatic expansion and the Hartree-Fock approximation.

Renormalization of the mean In the derivation above we assumed vanishing averages, $\mu_{i\setminus j} = 0$. When the Gaussian measure has nonzero mean, however, the nonlinear drift term induces a perturbative correction to it. At first order in ε, the mean becomes

$$\mu_{i\setminus j}(t) = \mu^0_{i\setminus j}(t) + \varepsilon \int dt'\, R^0_{i\setminus j}(t,t')\langle f_i(x_i(t'))\rangle^0_{i\setminus j} + O(\varepsilon^2). \tag{5.19}$$

At the Hartree-Fock level, we can sum the infinite series of first-order corrections by replacing the bare response R^0 with the renormalized one R, and the average with respect to the unperturbed measure with the full average, yielding

$$\mu_{i\setminus j}(t) = \mu^0_{i\setminus j}(t) + \varepsilon \int dt'\, R_{i\setminus j}(t,t')\langle f_i(x_i(t'))\rangle_{i\setminus j}. \tag{5.20}$$

The response-field average $\langle \mathrm{i}\hat{x}_i \rangle_{i\setminus j}$ instead remains zero at all orders, as required by causality.

The perturbative equations for the mean can be written in a compact form by introducing the vector

$$\Omega_{i\setminus j}(t) = \left(0 \quad , \quad \int dt' \, R_{i\setminus j}(t,t') \langle f_i(x_i(t')) \rangle_{i\setminus j}\right). \tag{5.21}$$

The mean then reads

$$\mathrm{M}_{i\setminus j}^{\top} = \left[\mathrm{M}_{i\setminus j}^{0}\right]^{\top} + \varepsilon \mathrm{G}_{i\setminus j} \Omega_{i\setminus j}^{\top}, \tag{5.22}$$

5.1.3 Self-consistent perturbative closure within the GECaM framework

The perturbative expansion described above can be naturally incorporated into the Gaussian-Expansion Cavity Method to obtain a closed set of self-consistent equations for the local moments. In this formulation, the cavity messages retain their Gaussian structure, but with renormalized propagators $\mathrm{G}_{i\setminus j}$ that include the perturbative corrections through the self-energy $\Sigma_{i\setminus j}$ defined in Eq. (5.12). The resulting procedure extends the Gaussian closure of GECaM to weakly nonlinear or multiplicative stochastic dynamics.

At each message update, the incoming cavity distributions $\{c_{k\setminus i}\}_{k\in\partial i\setminus j}$ are assumed Gaussian and are fully characterized by their mean, response and correlation functions, collected into the mean vectors $\mathrm{M}_{k\setminus i}$ and propagators $\mathrm{G}_{k\setminus i}$. Given these inputs, the unperturbed cavity mean vector $\mathrm{M}_{i\setminus j}^{0}$ and propagator $\mathrm{G}_{i\setminus j}^{0}$ are obtained by the usual Gaussian recursion of GECaM,

$$\left[\mathrm{G}_{i\setminus j}^{0}\right]^{-1} = \left[\mathrm{G}_{i}^{\mathrm{loc}}\right]^{-1} - \sum_{k\in\partial i\setminus j} \mathrm{J}_{ik} \, \mathrm{G}_{k\setminus i} \, \mathrm{J}_{ik}^{\top}, \tag{5.23}$$

$$\left[\mathrm{M}_{i\setminus j}^{0}\right]^{\top} = \mathrm{G}_{i\setminus j} \sum_{k\in\partial i\setminus j} \mathrm{J}_{ik} \, \mathrm{M}_{k\setminus i}^{\top}, \tag{5.24}$$

where $\mathrm{G}_{i}^{\mathrm{loc}}$ is the bare single-node propagator and J_{ik} the interaction matrix defined in Eqs. (4.35) and (4.36). The nonlinear and multiplicative contributions are then accounted for by adding the Hartree-Fock self-energy, according to

$$\left[\mathrm{G}_{i\setminus j}\right]^{-1} = \left[\mathrm{G}_{i\setminus j}^{0}\right]^{-1} - \varepsilon \Sigma_{i\setminus j}^{(1)}, \tag{5.25}$$

where

$$\Sigma_{i\setminus j}^{(1)}(t,t') = \begin{pmatrix} 0 & \delta(t-t') \, \langle f_i'(x_i(t)) \rangle_{i\setminus j} \\ \delta(t-t') \, \langle f_i'(x_i(t)) \rangle_{i\setminus j} & 2\,\delta(t-t') \, \langle g_i(x_i(t)) \rangle_{i\setminus j} \end{pmatrix}. \tag{5.26}$$

The mean is updated according to Eq. (5.20), which at the Hartree-Fock level reads

$$\mathrm{M}_{i\backslash j}^{\top} = \left[\mathrm{M}_{i\backslash j}^{0}\right]^{\top} + \varepsilon \mathrm{G}_{i\backslash j}\Omega_{i\backslash j}^{\top}, \tag{5.27}$$

where

$$\Omega_{i\backslash j}(t) = \begin{pmatrix} 0 & , & \int dt' R_{i\backslash j}(t,t')\langle f_i(x_i(t'))\rangle_{i\backslash j} \end{pmatrix}. \tag{5.28}$$

Substituting Eqs. (5.25) and (5.27) into the Gaussian cavity recursion gives a set of closed, self-consistent equations for the local moments,

$$\left[\mathrm{G}_i^{\mathrm{loc}}\right]^{-1}\mathrm{G}_{i\backslash j} = \mathrm{I} + \sum_{k\in\partial i\backslash j} \mathrm{J}_{ik}\mathrm{G}_{k\backslash i}\mathrm{J}_{ik}^{\top}\mathrm{G}_{i\backslash j} + \varepsilon\Sigma_{i\backslash j}^{(1)}\mathrm{G}_{i\backslash j}, \tag{5.29}$$

$$\left[\mathrm{G}_i^{\mathrm{loc}}\right]^{-1}\mathrm{M}_{i\backslash j}^{\top} = \sum_{k\in\partial i\backslash j}\left(\mathrm{J}_{ik}\mathrm{M}_{k\backslash i}^{\top} + \mathrm{J}_{ik}\mathrm{G}_{k\backslash i}\mathrm{J}_{ik}^{\top}\mathrm{M}_{i\backslash j}^{\top}\right) + \varepsilon\left(\Omega_{i\backslash j}^{\top} + \Sigma_{i\backslash j}^{(1)}\mathrm{M}_{i\backslash j}^{\top}\right). \tag{5.30}$$

The compact matrix notation used in Eqs. (5.29) and (5.30) implicitly includes time convolutions, e.g. $[\Sigma\mathrm{G}](t,t') = \int ds\,\Sigma(t,s)\,\mathrm{G}(s,t')$. Expanding these products yields explicit equations for the response, correlation, and mean fields. At order $\mathscr{O}(\varepsilon)$ (Hartree-Fock level), the additional local self-energy contributions modify the GECaM equations Eqs. (4.37) to (4.39) as follows:

$$\frac{\partial R_{i\backslash j}(t,t')}{\partial t} = -\lambda_{i\backslash j}^{\mathrm{r}}(t)R_{i\backslash j}(t,t') + \sum_{k\in\partial i\backslash j} J_{ik}J_{ki}\int_{t'}^{t} ds\,R_{k\backslash i}(t,s)R_{i\backslash j}(s,t') + \delta(t-t'), \tag{5.31}$$

$$\begin{aligned}\frac{\partial C_{i\backslash j}(t,t')}{\partial t} &= -\lambda_{i\backslash j}^{\mathrm{r}}(t)C_{i\backslash j}(t,t') + \sum_{k\in\partial i\backslash j} J_{ik}J_{ki}\int_{0}^{t} ds\,R_{k\backslash i}(t,s)C_{i\backslash j}(s,t') \\ &\quad + \sum_{k\in\partial i\backslash j} J_{ik}^2\int_{0}^{t'} ds\,C_{k\backslash i}(t,s)R_{i\backslash j}(t',s) + 2T_{i\backslash j}^{\mathrm{r}}(t)R_{i\backslash j}(t',t)\end{aligned} \tag{5.32}$$

$$\begin{aligned}\frac{d\mu_{i\backslash j}(t)}{dt} &= -\lambda_{i\backslash j}^{\mathrm{r}}(t)\mu_{i\backslash j}(t) + \sum_{k\in\partial i\backslash j} J_{ik}\mu_{k\backslash i}(t) + \sum_{k\in\partial i\backslash j} J_{ik}J_{ki}\int_{0}^{t} ds\,R_{k\backslash i}(t,s)\mu_{i\backslash j}(s) \\ &\quad + \varepsilon\int_{0}^{t} ds\,R_{i\backslash j}(t,s)\left\langle f_i(x_i(s))\right\rangle_{i\backslash j},\end{aligned} \tag{5.33}$$

where we have defined the renormalized relaxation rate and temperature

$$\lambda_{i\backslash j}^{\mathrm{r}}(t) = \lambda_i - \varepsilon\langle f_i'(x_i(t))\rangle_{i\backslash j}, \tag{5.34}$$

$$T_{i\backslash j}^{\mathrm{r}}(t) = T + \varepsilon\langle g_i(x_i(t))\rangle_{i\backslash j}. \tag{5.35}$$

The perturbative drift $f_i(x)$ thus renormalizes the effective relaxation rate through the local average $\langle f_i'(x)\rangle$, while the multiplicative noise $g_i(x)$ modifies the effective temperature via $\langle g_i(x)\rangle$. The last term in Eq. (5.33) accounts for the direct first-order correction to the mean induced by the nonlinear drift, while multiplicative noise affects μ only indirectly through its impact on C.

In summary, the perturbative extension of GECaM provides a controlled framework to include weak nonlinearities and multiplicative noise within the cavity formalism. By introducing a local self-energy that renormalizes the drift and noise amplitudes, the Gaussian closure is promoted to a self-consistent, field-theoretically grounded approximation. The following sections apply this construction to specific cases, illustrating how the perturbative scheme captures nontrivial dynamical effects beyond the Gaussian regime.

An equivalent formulation can be obtained for the local single-site propagators G_i and means M_i, by including the contributions from all neighbors $k \in \partial i$ in Eqs. (5.29) and (5.30).

5.2 Gaussian perturbative closure for a cubic nonlinearity

We begin by applying the perturbative GECaM framework to the simplest nonlinear case: a cubic correction to the local drift. This example illustrates how weak nonlinearities renormalize the effective relaxation rate and induce a noise-driven phase transition in an otherwise linear system of coupled UO processes.

We consider the dynamics

$$\frac{dx_i(t)}{dt} = -\lambda x_i(t) - \varepsilon x_i^3(t) + \sum_{j\in\partial i} J_{ij}x_j(t) + \xi_i(t), \tag{5.36}$$

corresponding to a nonlinear drift $f_i(x_i) = -x_i^3$ and additive noise $g_i(x_i) = 0$. This model describes overdamped particles diffusing in symmetric double-well potentials and coupled through linear interactions on a network.

The perturbative expansion follows the general formulation of section 5.1, with ε controlling the strength of the cubic nonlinearity. The associated perturbative action reads

$$S_i^{\mathrm{p}}[\mathrm{X}_i] = \int_0^{t_f} ds\, \mathrm{i}\hat{x}_i(s)\, x_i^3(s). \tag{5.37}$$

The unperturbed dynamics admits a stable stationary state centered at zero for $\lambda \geq KJ$, so that $\mathrm{M}_{i\setminus j} = 0$ can be assumed without loss of generality. When the cubic term is introduced, a

transition emerges between a symmetric phase with $\langle x_i \rangle = 0$ and a symmetry broken phase where $\langle x_i \rangle \neq 0$. This is analogous to the *para-ferro transition* in Ising-like systems with quartic potentials, where spontaneous magnetization develops below a critical temperature. Within the perturbative GECaM framework, this transition can be analyzed by studying the stability of the symmetric solution $\mu_{i\setminus j} = 0$ in the stationary, TTI regime.

At weak coupling ($\varepsilon \ll 1$), the leading correction to the Gaussian theory is captured by the Hartree-Fock self-energy terms Eqs. (5.17) and (5.18). Since $g(x) = 0$, the noise self-energy vanishes, $\varsigma^T = 0$, and the effective temperature remains unchanged ($T^{\mathrm{r}} = T$). The cubic nonlinearity, however, renormalizes the relaxation rate through ς^R, which at the Hartree-Fock level reads

$$\varsigma^R_{i\setminus j}(t,t') = \delta(t-t') \left\langle f_i'(x_i(t)) \right\rangle_{i\setminus j} = -3\,\delta(t-t') \left\langle x_i^2(t) \right\rangle_{i\setminus j} \tag{5.38}$$

$$= -3\,C_{i\setminus j}(t,t')\,\delta(t-t'). \tag{5.39}$$

The corresponding renormalized relaxation rate is

$$\lambda^{\mathrm{r}}_{i\setminus j}(t) = \lambda + 3\varepsilon\, C_{i\setminus j}(t,t). \tag{5.40}$$

The effective relaxation rate increases with the local equal-time correlation, reflecting that the cubic term enhances the restoring force near the origin. At stationarity, the equal-time cavity correlation becomes time independent ($q^{\mathrm{st}}_{i\setminus j}$), and the perturbed system is equivalent to N linearly-coupled OU processes with renormalized rates $\lambda^{\mathrm{r}}_{i\setminus j} = \lambda + 3\varepsilon q^{\mathrm{st}}_{i\setminus j}$. The paramagnetic phase with $\mu_{i\setminus j} = 0$ is stable as long as $\lambda^{\mathrm{r}}_{i\setminus j} > KJ$ for every edge (i,j), or equivalently

$$\lambda > \lambda^{\mathrm{crit}}(\varepsilon) = KJ - 3\varepsilon \min_{(i,j)\in E} q^{\mathrm{st}}_{i\setminus j}. \tag{5.41}$$

The critical line $\lambda^{\mathrm{crit}}(\varepsilon)$ depends on the stationary variances $q^{\mathrm{st}}_{i\setminus j}$, which can be computed self-consistently within the GECaM framework.

In principle, $q^{\mathrm{st}}_{i\setminus j}$ can be obtained by inverting the Laplace transform of the correlation function Eq. (4.48) with λ replaced by λ^{r}. However, due to the branch cuts and poles of the Laplace kernel, it is more convenient to exploit the relation between GECaM and random matrix theory discussed in section 4.3. The stationary correlation can be expressed in terms of the cavity spectral density $\rho_{\mathrm{cav}}(x)$ of the interaction matrix $\underline{\underline{J}}$, which can be computed from the stationary response function.

For simplicity, we consider a homogeneous RRG with uniform couplings $J_{ij} = J_{ji} = J$ and degree K. All cavity quantities are identical, $q^{\mathrm{st}}_{i\setminus j} = q^{\mathrm{st}}_c$, and the Laplace transformed cavity response function reads

$$\tilde{R}_c(z) = \frac{z + \lambda^{\mathrm{r}}_c - \sqrt{(z+\lambda^{\mathrm{r}}_c)^2 - 4J^2(K-1)}}{2(K-1)J^2}, \tag{5.42}$$

where $\lambda^{\mathrm{r}}_c = \lambda + 3\varepsilon q^{\mathrm{st}}_c$. The corresponding spectral density, obtained from the Stieltjes transform relation Eq. (4.60), is

$$\rho_c(x) = \frac{J}{\pi} \lim_{\eta\to 0^+} \mathrm{Im}\, \tilde{R}_c(Jx - \lambda^{\mathrm{r}}_c - \mathrm{i}\eta) = \frac{\sqrt{4(K-1)-x^2}}{2\pi(K-1)}, \tag{5.43}$$

supported on $x \in [-2\sqrt{K-1}, 2\sqrt{K-1}]$. The stationary correlation follows from Eq. (4.65) as

$$q^{\mathrm{st}}_c = T \int_{-2\sqrt{K-1}}^{2\sqrt{K-1}} \frac{\sqrt{4(K-1)-x^2}}{2\pi(K-1)}\, dx \int_0^\infty e^{(Jx-\lambda^{\mathrm{r}}_c)w}\, dw \tag{5.44}$$

$$= \frac{T\,\lambda^{\mathrm{r}}_c}{2(K-1)J^2} \left(1 - \sqrt{1 - 4\frac{(K-1)J^2}{(\lambda^{\mathrm{r}}_c)^2}}\right). \tag{5.45}$$

Substituting this into Eq. (5.41) yields the critical line

$$\lambda^{\mathrm{crit}}(\varepsilon) = KJ - 3\varepsilon \frac{T}{(K-1)J}. \tag{5.46}$$

The transition line decreases linearly with both ε and T, indicating that stronger nonlinearities and higher noise levels promote the symmetry-broken phase. Increasing the connectivity K also lowers λ^{crit}, as stronger coupling enhances collective ordering.

Figure 5.2 compares the theoretical prediction with numerical simulations of Eq. (5.36) on a RRG with $N = 200$ nodes, for $K = 3$ and $K = 10$, at fixed $J = 1$ and $\varepsilon = 0.01$. Numerical results were obtained by integrating the stochastic dynamics using the Euler-Maruyama scheme with homogeneous initial conditions $x_i(0) = 1$ for all nodes, and measuring the stationary average

$$\langle x \rangle = \frac{1}{N}\sum_i x_i(t)$$

at $t = 1000$, i.e. well beyond the equilibration time. Results were averaged over 100 independent realizations of the noise and the graph. The theoretical transition line from Eq. (5.41) (solid white line) matches well the simulation data, confirming that the perturbative closure accurately captures the onset of symmetry breaking.

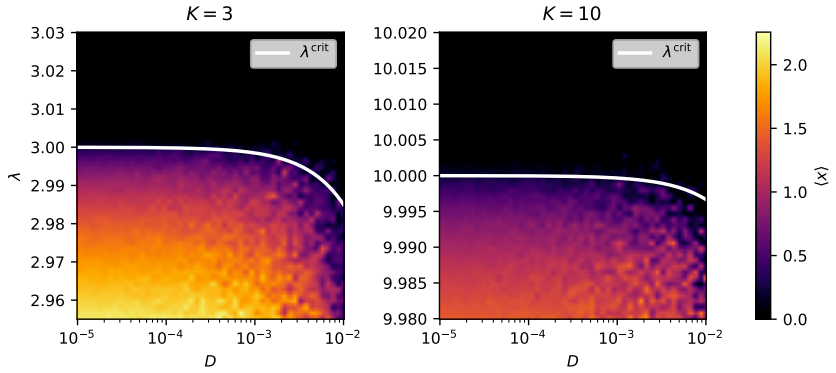

Fig. 5.2 **Stationary magnetization for a cubic nonlinearity.** Stationary magnetization $\langle x \rangle$ from numerical simulations of Eq. (5.36) on a RRG with $N = 200$ nodes, for $K = 3$ (left) and $K = 10$ (right), at fixed $J = 1$ and $\varepsilon = 0.01$. The stationary values were computed at $t = 1000$ from homogeneous initial conditions $x_i(0) = 1$. The theoretical transition line from Eq. (5.41) (solid white line) matches well the simulation data, confirming that the perturbative closure accurately captures the onset of symmetry breaking.

5.3 Noise-driven phase transition in the Bouchaud-Mézard model

Stochastic differential equations with state-dependent, or *multiplicative*, noise terms naturally arise in the modeling of physical, biological, and socio-economic systems, as they describe the effect of both intrinsic and environmental fluctuations acting proportionally to the system's state [62, 63]. In contrast to additive noise, multiplicative fluctuations can qualitatively modify the dynamical stability and give rise to noise-induced transitions, where the stationary properties of the system change as a consequence of the coupling between the noise amplitude and the local state variables.

A paradigmatic example of such behavior is provided by the Bouchaud-Mézard (BM) model of wealth exchange [64]. Originally introduced to describe the emergence of heavy-tailed wealth distributions in economic systems, it also serves as a prototype for a broad class of multiplicative stochastic processes encountered in diverse contexts, ranging from surface growth to population dynamics. Empirical studies show that real-world wealth distributions often display a Pareto tail [65], reflecting the coexistence of many poor agents with a small number of extremely wealthy ones. The BM model captures this feature through a simple networked dynamics that couples deterministic exchange of wealth among connected agents with stochastic, multiplicative fluctuations.

On a general graph, the evolution of the wealth $x_i(t)$ of agent i is governed by the set of Langevin equations

$$\frac{dx_i(t)}{dt} = J \sum_{j \in \partial i} \big(x_j(t) - x_i(t)\big) + \varepsilon\, \sigma\, x_i(t)\, \xi_i(t), \tag{5.47}$$

where J controls the exchange rate between neighboring nodes, σ is the amplitude of the noise, and $\xi_i(t)$ is a Gaussian white noise. The first term on the right-hand side represents a diffusion of wealth across the network due to interactions with neighboring agents, while the second term models random multiplicative fluctuations in wealth, arising from investment returns or market volatility. The parameter ε will serve as a small perturbative parameter controlling the strength of the multiplicative contribution in the following analysis.

The stochastic Eq. (5.47) is invariant under uniform rescaling of all wealths, $x_i \to \kappa x_i$ for any $\kappa > 0$. Hence, without loss of generality, one can fix the average wealth to unity in the theoretical analysis (i.e. $\langle x_i(t) \rangle_i = 1$) and study the fluctuations of the normalized quantity $x_i/\bar{x}$, where $\bar{x}$ denotes the population average. In the original fully connected formulation, Bouchaud and Mézard showed that the stationary distribution of $x_i/\bar{x}$ exhibits a power law tail with exponent $\alpha > 0$, and that a phase transition occurs at $\sigma > \sigma^{\text{crit}} = \sqrt{2J}$, where the variance of the normalized wealth diverges. This phenomenon, referred to as *wealth condensation*, corresponds to the emergence of a small number of agents accumulating a finite fraction of the total wealth, while the rest of the population remains poor. For $\sigma < \sigma^{\text{crit}}$, instead, the variance remains finite and wealth is homogeneously distributed around the mean. Alternatively, the same transition can be interpreted as a crossover from a weak-noise to a strong-noise regime obtained by varying σ^2 at fixed J [63].

The generalization of the BM model to sparse random graphs was later addressed in [66], where the author combined adiabatic and independence assumptions to construct a self-consistent mean-field closure. The adiabatic approximation allows slow variations of local averages to be treated as quasi static, while the independence assumption enables the use of the central limit theorem to approximate the sum over neighbors. This approach yields improved predictions for the critical point on sparse topologies, such as random regular graphs, where degree heterogeneity and finite connectivity affect the onset of condensation. In what follows, we develop an alternative derivation based on the perturbative closure introduced in Sec. 5.1.3, which provides a systematic expansion in powers of the noise amplitude ε and recovers, as limiting cases, both the fully connected mean-field result and the sparse graph correction obtained in [66].

The BM dynamics Eq. (5.47) can be viewed as a special case of coupled OU processes with vanishing temperature in the unperturbed limit, $T = 0$, and site-dependent linear relaxation $\lambda_i = J|\partial i|$. The perturbation is purely multiplicative noise: there is no nonlinear drift, $f_i(x_i) = 0$, while $g_i(x_i) = \sigma^2 x_i^2/2$. The associated perturbative action is

$$S_i^{\mathrm{p}}[\mathrm{X}_i] = -\frac{\sigma^2}{2}\int_0^{t_f} ds\, x_i^2(s)\left(\mathrm{i}\hat{x}_i(s)\right)^2, \tag{5.48}$$

and it enters the cavity measure multiplied by the small parameter ε as in Eqs. (5.3) and (5.5). In other words, the perturbative expansion proceeds in powers of ε while σ sets the bare noise scale of the multiplicative vertex.

As already discussed in section 5.1, when the dynamics has a non-zero mean $\mu_{i\setminus j}$, one should center the fields by defining $\delta x = x - \mu$. In the present case, we can exploit the scale invariance of Eq. (5.47) to set $\mu_{i\setminus j} = 1$ for all i, j. The perturbative action then reads

$$\begin{aligned} S_i^{\mathrm{p}}[\mathrm{X}_i] = -\frac{\sigma^2}{2}\int_0^{t_f} ds\,(1+\delta x_i(s))^2\left(\mathrm{i}\hat{x}_i(s)\right)^2 = -\frac{\sigma^2}{2}\int_0^{t_f} ds\left(\mathrm{i}\hat{x}_i(s)\right)^2 \\ -\sigma^2\int_0^{t_f} ds\,\delta x_i(s)\left(\mathrm{i}\hat{x}_i(s)\right)^2 - \frac{\sigma^2}{2}\int_0^{t_f} ds\,\delta x_i^2(s)\left(\mathrm{i}\hat{x}_i(s)\right)^2. \end{aligned} \tag{5.49}$$

At the Hartree-Fock level, the only non-vanishing self-energy is ς^T, since $f_i(x_i) = 0$ implies $\varsigma^R = 0$. The multiplicative noise therefore renormalizes the effective temperature, which becomes

$$T_{i\setminus j}^{\mathrm{r}}(t) = \varepsilon\,\frac{\sigma^2}{2}\left(1 + C_{i\setminus j}(t,t)\right). \tag{5.50}$$

We focus again on the stationary regime, where the system is time-translationally invariant. Since $\varsigma^R = 0$, the response function remains unperturbed and satisfies in Laplace space

$$\tilde{R}_{i\setminus j}(z) = \left(z + \lambda_i - J^2\sum_{k\in\partial i\setminus j}\tilde{R}_{k\setminus i}(z)\right)^{-1}. \tag{5.51}$$

The correlation function, instead, becomes non-zero due to the perturbation and can be written as

$$C_{i\setminus j}(\tau) = \left(\sigma_{i\setminus j}^{\mathrm{r}}\right)^2\int_\tau^\infty du\, R_{i\setminus j}(u)\, R_{i\setminus j}(u-\tau), \tag{5.52}$$

where $\sigma_{i\setminus j}^{\mathrm{r}}$ is a renormalized noise amplitude obtained by resumming the contributions from the first and last vertex in S_i^{p}. The cubic vertex does not contribute, since all its contractions involve equal time or advanced responses, which vanish by causality. As we show in Appendix C.2.1, this resummation can be expressed as a geometric series of one-loop corrections, yielding at

stationarity

$$\left(\sigma^{\mathrm{r}}_{i\setminus j}\right)^2 = \frac{\sigma^2}{1-\varepsilon\sigma^2 I_{i\setminus j}}, \qquad I_{i\setminus j} = \int_{-\infty}^{\infty} \frac{d\omega}{2\pi}\, \tilde{R}_{i\setminus j}(\omega)\tilde{R}_{i\setminus j}(-\omega), \tag{5.53}$$

with $\tilde{R}_{i\setminus j}(\omega) = \tilde{R}_{i\setminus j}(z = \mathrm{i}\omega)$.

For random regular graphs, where $\lambda_i = KJ$, all nodes and edges become statistically equivalent and indices can be dropped. In this case, it can be shown (see Appendix C.2.2) that

$$\begin{aligned} I &= \int_{-\infty}^{\infty} \frac{d\omega}{2\pi} \frac{\mathrm{i}\omega + KJ - \sqrt{(\mathrm{i}\omega + KJ)^2 - 4(K-1)J^2}}{2(K-1)J^2} \\ &\quad \times \frac{-\mathrm{i}\omega + KJ - \sqrt{(-\mathrm{i}\omega + KJ)^2 - 4(K-1)J^2}}{2(K-1)J^2} \qquad (5.54a) \\ &= \frac{2}{\pi\sqrt{K-1}J} \int_{-1}^{1} du \, \frac{\sqrt{1-u^2}}{\frac{K}{\sqrt{K-1}} - u + \sqrt{\left(\frac{K}{\sqrt{K-1}} - u\right)^2 - 1}}. \qquad (5.54b) \end{aligned}$$

The fully connected limit is easily recovered by setting $K = N-1$, rescaling $J \to J/K$, and taking $K \to \infty$, obtaining

$$I_{\mathrm{FC}} \approx \frac{2}{\pi J} \frac{K}{\sqrt{K-1}} \int_{-1}^{1} du \, \frac{\sqrt{1-u^2}}{2K/\sqrt{K-1}} \longrightarrow \frac{1}{2J}. \tag{5.55}$$

The renormalized noise amplitude diverges at the critical value $(\sigma^{\mathrm{crit}})^2 = I^{-1}$, signaling a transition to the condensed phase. In the fully connected limit, one recovers $(\sigma^{\mathrm{crit}}_{\mathrm{FC}})^2 = 2J$, in agreement with the mean-field result of [64].

To assess the accuracy of our method, we compared its predictions with numerical simulations. We rescaled the coupling as $J \to J/K$ to facilitate comparison across different connectivity values. Simulations of the BM model on RRGs with $N = 5000$ nodes were performed using the Milstein algorithm [67] with initial condition $x_i(0) = 1$ and $J = 1$. The stationary distribution was estimated at $t = 1500$, and the normalized wealth $x/\bar{x}$ was found to follow a Pareto tail

$$p\left(\frac{x}{\bar{x}}\right) \overset{x/\bar{x}\to\infty}{\sim} \frac{1}{(x/\bar{x})^{\alpha+1}}, \tag{5.56}$$

whose variance diverges for $\alpha \leq 2$. The tail exponent α was estimated using a rank-1/2 ordinary least squares log-log regression with bias correction [68]. Figure 5.3 shows the results for $K = 3$, 10, and 500, together with the theoretical thresholds $(\sigma^{\mathrm{crit}}_{\mathrm{I}})^2$ from [66] and $(\sigma^{\mathrm{crit}}_{\mathrm{G}})^2 = I^{-1}$ from our perturbative closure. Our method provides a more accurate estimate of the critical point

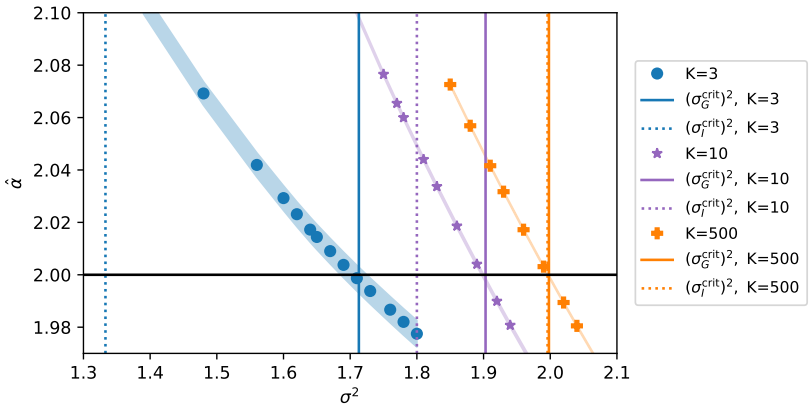

Fig. 5.3 **Tail exponent near the condensation transition of the BM model.** Estimates of the tail exponent α obtained by simulation for $K = 3$ (dots), $K = 10$ (stars), and $K = 500$ (crosses) at different values of σ^2 around the phase transition. Simulations were performed for RRGs with $N = 5000$ nodes, using the Milstein algorithm [67], averaged over 100 independent realizations. Vertical lines indicate the theoretical critical values predicted by our method ($(\sigma_G^{crit})^2$, solid) and by the self-consistent approach of [66] ($(\sigma_I^{crit})^2$, dotted). The horizontal black line marks $\alpha_c = 2$: a tail exponent below this value signals wealth condensation.

for sparse graphs, while the two predictions coincide in the dense limit. Interestingly, the self-consistent approach of [66] can be recovered from GECaM when further approximating the cavity quantities by their full counterparts evaluated at zero time difference (see Appendix C.2.3 for details).

Chapter 6

Dynamics of the Spherical 2-Spin Model

6.1 The Spherical 2-Spin model

The spherical p-spin model, originally introduced as a paradigmatic example of mean-field glassy systems, has played a central role in the theoretical understanding of slow relaxation and aging phenomena in disordered systems. For $p \geq 3$, the model exhibits a highly non-trivial free energy landscape, characterized by an extensive number of metastable states and the emergence of dynamical glass transitions [69]. The case $p = 2$, instead, is structurally much simpler: it does not possess a glassy landscape but undergoes a continuous transition into a phase that can be interpreted as a disguised ferromagnet [70]. Nevertheless, the $p = 2$ spherical model still displays rich dynamical properties. After a quench into the low-temperature phase, the system fails to reach equilibrium, and its relaxation becomes dependent on the waiting time, a hallmark of aging dynamics.

Because of its linear structure, the model is exactly solvable through diagonalization of the interaction matrix, provided the spectral properties of the underlying graph are known [71–73]. Here, we show that GECaM allows one to describe the same non-equilibrium relaxation behavior without relying on explicit diagonalization or spectral information. This model therefore provides an ideal benchmark to test the accuracy and consistency of our approximation scheme.

The model can be seen as relaxation dynamics of N continuous spins on the vertices of a graph G, whose equilibrium properties are described by the Hamiltonian

$$H[\underline{x}] = -\frac{1}{2}\sum_{i\neq j} A_{ij}J_{ij}x_i x_j, \tag{6.1}$$

where $x_i(t)$ are continuous spin variables satisfying the spherical constraint

$$\sum_{i=1}^{N} x_i^2(t) = N.$$

The elements A_{ij} are the entries of the adjacency matrix of the interaction graph G, while the couplings J_{ij} are independent, symmetric random variables drawn from the bimodal distribution

$$P(J_{ij}) = \frac{1}{2}\left[\delta(J_{ij} - J) + \delta(J_{ij} + J)\right]. \tag{6.2}$$

The natural scaling $J \propto 1/\sqrt{K}$ ensures a well-defined thermodynamic limit and allows the recovery of the fully connected model as $K \to \infty$.

The relaxation dynamics follows the Langevin equation

$$\frac{dx_i(t)}{dt} = -\lambda(t)x_i(t) - \frac{\delta H[\underline{x}(t)]}{\delta x_i(t)} + \xi_i(t) \tag{6.3}$$

$$= -\lambda(t)x_i(t) + \sum_{j\in\partial i} J_{ij}x_j(t) + \xi_i(t), \tag{6.4}$$

where $\lambda(t)$ is a Lagrange multiplier enforcing the spherical constraint at all times, and $\xi_i(t)$ is a Gaussian white noise with zero mean and correlations

$$\langle \xi_i(t)\xi_{i'}(t')\rangle = 2T\,\delta_{ii'}\delta(t-t'), \tag{6.5}$$

where the noise strength T plays the role of temperature.

6.2 Cavity formulation

The GECaM equations can be derived for any specific realization of the interaction graph. However, in order to extract general dynamical properties and facilitate comparison with known results, it is convenient to consider the homogeneous limit appropriate for random regular graphs. Because of the statistical symmetry of the ensemble and the absence of replica symmetry breaking in the $p = 2$ case, all disorder-averaged quantities become site- and edge-independent. We therefore assume

$$\overline{R_i} = R, \quad \overline{C_i} = C, \quad \overline{R_{i\setminus j}} = R_c, \quad \overline{C_{i\setminus j}} = C_c.$$

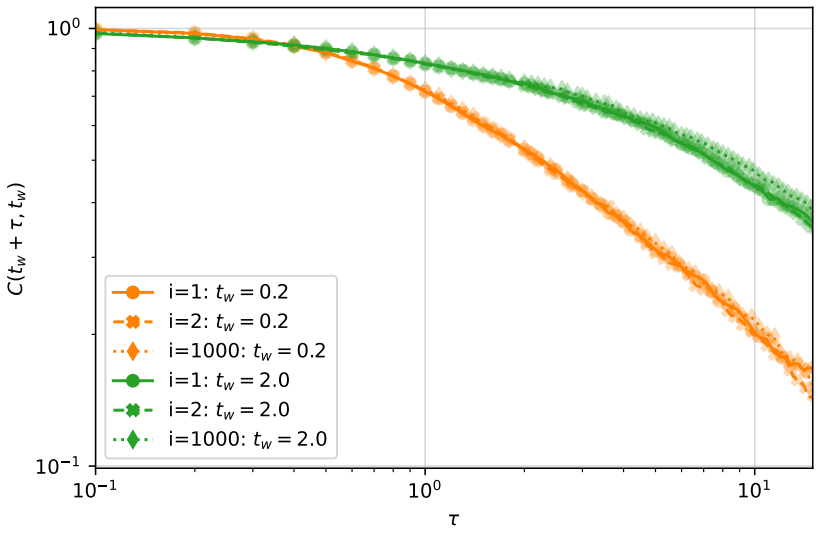

Fig. 6.1 **Disorder-averaged correlations on random regular graphs.** Disorder-averaged correlation functions for three representative nodes of a RRG with $N = 1000$ and degree $K = 5$. Each node is shown for two different waiting times, $t_w = 0.2$ (orange curves) and $t_w = 2.0$ (green curves). Markers correspond to results from numerical simulations, while lines represent running averages of the markers computed on a logarithmic scale of the abscissas. Parameters: $J = 1.0$, $T = 0.3$. The dynamics is integrated with Euler-Maruyama, and averages are taken over 10^4 realizations of noise, coupling disorder, and initial conditions. The near-perfect overlap confirms the working assumption of site/edge independence.

The spherical constraint implies $C(t,t) = 1$ at all times, thus reducing the number of independent functions.

Before solving the self-consistent equations, we verified numerically that disorder-averaged correlations are effectively homogeneous across sites on random regular graphs. In Fig. 6.1 we show that correlation functions measured on three representative nodes coincide within numerical accuracy on a RRG with $N = 1000$, degree $K = 5$, coupling $J = 1.0$, and temperature $T = 0.3$. Trajectories were integrated with Euler–Maruyama and averaged over 10^4 independent realizations of noise, coupling disorder, and uniform initial conditions. This supports treating both cavity and full quantities as site-independent in the thermodynamic limit.

Under these assumptions, and for a sudden quench from the high temperature phase, the corresponding GECaM equations for the cavity response and correlation functions read

$$\frac{\partial R_c(t,t')}{\partial t} = -\lambda(t)R_c(t,t') + (K-1)J^2\int_{t'}^{t} ds R_c(t,s)R_c(s,t') + \delta(t-t'), \tag{6.6}$$

$$\frac{\partial C_c(t,t')}{\partial t} = -\lambda(t)C_c(t,t') + (K-1)J^2\int_{0}^{t} ds R_c(t,s)C_c(s,t') + 2TR_c(t',t) + (K-1)J^2\int_{0}^{t'} ds R_c(t',s)C_c(t,s). \tag{6.7}$$

Similarly, for the full response and correlation functions one finds

$$\frac{\partial R(t,t')}{\partial t} = -\lambda(t)R(t,t') + KJ^2\int_{t'}^{t} ds R_c(t,s)R(t,s) + \delta(t-t'), \tag{6.8}$$

$$\frac{\partial C(t,t')}{\partial t} = -\lambda(t)C(t,t') + KJ^2\int_{0}^{t} ds R_c(t,s)C(s,t') + 2TR(t',t) + KJ^2\int_{0}^{t'} ds R(t',s)C_c(t,s). \tag{6.9}$$

The Lagrange multiplier $\lambda(t)$ is determined by enforcing the spherical constraint. Since the equal time correlation must be equal to unity, we can set[1]

$$\left(\frac{\partial C(t,t')}{\partial t} + \frac{\partial C(t,t')}{\partial t'}\right)\Big|_{t,t'=s} = 0, \tag{6.10}$$

from which we obtain

$$\lambda(t) = KJ^2\int_0^t ds\left(R_c(t,s)C(s,t) + R(t,s)C_c(s,t)\right) + T. \tag{6.11}$$

When setting $J^2 \propto 1/K$ and taking the limit $K \to \infty$, the cavity and full quantities coincide, and the above equations reduce to the well-known dynamical mean-field equations for the fully connected spherical model, as originally derived by Cugliandolo and Kurchan [71].

[1] In the case of RRGs, the global spherical constraint translates into a local condition for the disorder-averaged correlation function, owing to the homogeneity of the graph. In heterogeneous networks, however, this simplification no longer holds, as one would need to impose the vanishing of the time derivative of $C(t,t') = \sum_i C_i(t,t')$ explicitly.

To make further analytical progress, it is useful to introduce the parametric ansatz

$$R_c(t,t') = \frac{f(t')}{f(t)} R_c^{\mathrm{st}}(t-t'), \tag{6.12}$$

$$C_c(t,t') = \frac{1}{f(t)f(t')} C_c^{1}(t,t'), \tag{6.13}$$

with the scaling function

$$f(t) = \exp\left[\int_0^t ds\, \lambda(s)\right], \tag{6.14}$$

normalized as $f(0) = 1$. Substituting this ansatz into Eqs. (6.6) and (6.7) leads to a decoupling between stationary and aging components. The stationary part, $R_c^{\mathrm{st}}(t-t')$, satisfies the TTI equation

$$\frac{\partial R_c^{\mathrm{st}}(t-t')}{\partial t} = (K-1)J^2 \int_{t'}^{t} ds\, R_c^{\mathrm{st}}(t-s) R_c^{\mathrm{st}}(s-t') + \delta(t-t'), \tag{6.15}$$

corresponding to the response of a linear system with additive noise. The non-stationary function $C_c^{1}(t,t')$ instead obeys

$$\begin{aligned}\frac{\partial}{\partial t} C_c^{1}(t,t') = {} & 2T f(t)^2 R_c^{\mathrm{st}}(t'-t) + (K-1)J^2 \int_0^t ds\, R_c^{\mathrm{st}}(t-s) C_c^{1}(s,t') \\ & + (K-1)J^2 \int_0^{t'} ds\, R_c^{\mathrm{st}}(t'-s) C_c^{1}(t,s).\end{aligned} \tag{6.16}$$

Analogous equations hold for the full functions R^{st} and C^{1}, obtained by replacing $K-1 \to K$ in the kernels.

By direct substitution, one finds that the coupled system admits the integral solutions

$$C_c^{1}(t,t') = R_c^{\mathrm{st}}(t+t') + 2T \int_0^{\min(t,t')} ds\, f^2(s) R_c^{\mathrm{st}}(t+t'-2s), \tag{6.17}$$

$$C^{1}(t,t') = R^{\mathrm{st}}(t+t') + 2T \int_0^{\min(t,t')} ds\, f^2(s) R^{\mathrm{st}}(t+t'-2s), \tag{6.18}$$

while the scaling function $f(t)$ obeys the linear Volterra equation

$$f^2(t) = R^{\mathrm{st}}(2t) + 2T \int_0^t ds\, f^2(t_1) R^{\mathrm{st}}(2t-2s). \tag{6.19}$$

Equation (6.19) is formally identical to the exact result obtained in Ref. [71], confirming that the GECaM equations reproduce the correct dynamical behaviour of the spherical 2-spin model without invoking any diagonalization of the coupling matrix. In practice, an explicit solution

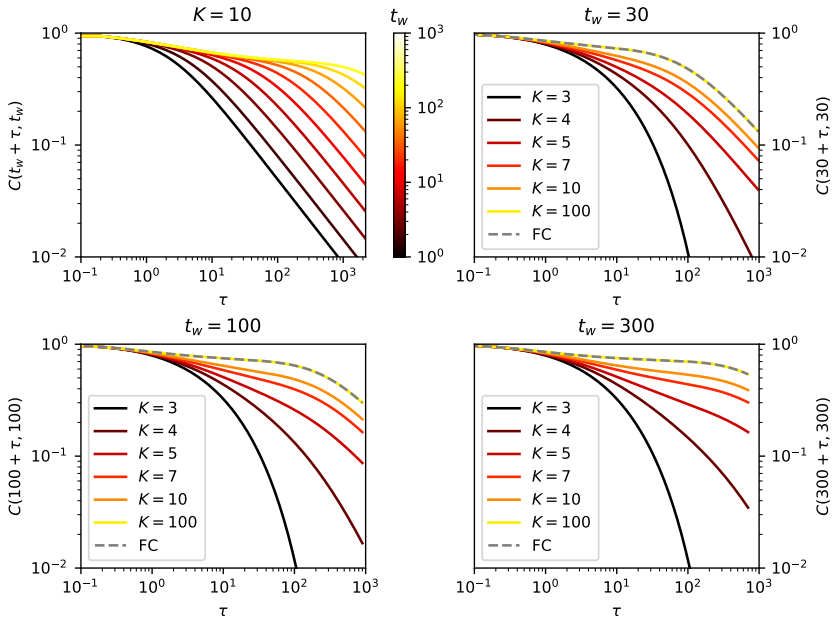

Fig. 6.2 **Ageing and correlation decay in the spherical 2-spin model.** Decay of the correlation $C(t_w + \tau, t_w)$ as a function of τ for the spherical 2-spin model on a random regular graph with coupling constant $J = 1.0$ and temperature $T = 0.3$. **Top-left panel**: Decay at fixed connectivity $K = 10$ for increasing waiting times t_w (from bottom to top). **Top-right and bottom panels**: Decay at fixed $t_w = 30, 100, 300$ for $K = 3, 4, 5, 7, 10, 100$ (from bottom to top), compared with the fully connected limit (dashed gray line).

for $f(t)$ can be obtained numerically by Laplace transform techniques or by direct forward integration of Eqs. (6.6) and (6.11).

6.3 Ageing and Non-Equilibrium Dynamics

The numerical integration of the GECaM equations reveals a characteristic two-time behaviour of the correlation function $C(t, t')$, which encodes the slow relaxation dynamics of the model. In Fig. 6.2 we show the decay of $C(t_w + \tau, t_w)$ for temperature $T = 0.3$ and coupling constant $J = 1.0$, as a function of τ for several waiting times t_w. The correlation exhibits a clear separation of time scales: an initial rapid decay for $\tau \ll t_w$, followed by a much slower decay for $\tau \gtrsim t_w$. This dependence of the relaxation rate on the system age demonstrates the presence of aging, in full agreement with the exact analytical solution and numerical simulations of the model.

At fixed temperature, reducing the connectivity K accelerates decorrelation and suppresses aging, as the effective number of interaction channels per site decreases. In the dense limit $K \to$

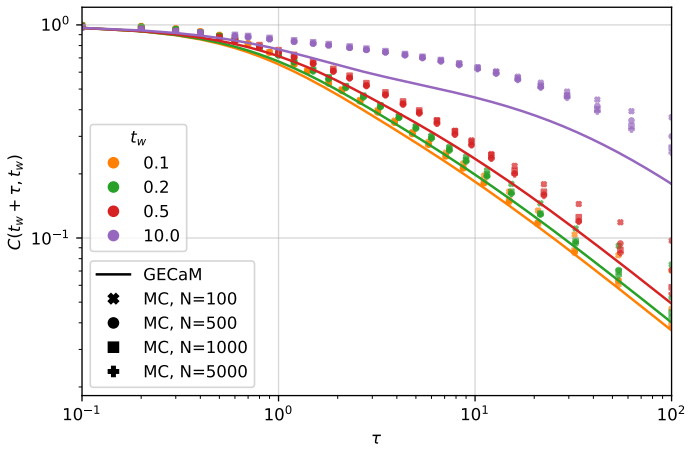

Fig. 6.3 **Comparison between GECaM and MC correlations.** Correlation function $C(t_w+\tau,t_w)$ for the spherical 2-spin model computed via GECaM (solid lines) compared with MC simulations (markers) on a RRG with $K=4$, $J=1.0$, and $T=0.3$. The comparison is shown for waiting times $t_w = 0.1, 0.2, 0.5, 10$. To highlight finite-size effects, MC results are reported for increasing system sizes $N = 100, 500, 1000, 5000$. As N increases, the microscopic simulations systematically approach the GECaM prediction, confirming that the discrepancy observed at finite N is a finite-size effect and that the cavity theory correctly captures the thermodynamic limit.

∞, the behavior converges smoothly to that of the fully connected model, as expected from the analytical correspondence discussed above. The numerical solutions of the GECaM equations reproduce the Monte Carlo results with good qualitative agreement across all timescales. Quantitative deviations appear at large waiting times, as shown in Fig. 6.3, where the correlation $C(t_w+\tau,t_w)$ obtained from GECaM (solid lines) is compared with Monte Carlo data (markers) for an RRG with $K=4$, $J=1.0$, and $T=0.3$. The figure reports microscopic simulations for increasing system sizes $N \in \{100, 500, 1000, 5000\}$. As N grows, the MC curves shift systematically towards the GECaM prediction, confirming that the observed deviations are indeed due to finite-size effects in the particle simulations. The discrepancy remains significant at the largest waiting time ($t_w = 10$) even for $N = 5000$; this persistence is a consequence of the slow decorrelation dynamics in the aging regime, which implies that extremely large system sizes would be required to effectively sample the thermodynamic limit captured by the cavity theory.

6.3.1 Ageing in the Spherical Ferromagnet

Interestingly, the same equations remain valid for the spherical ferromagnet, differing only in the scaling of the couplings. While the disordered (bimodal) model requires $J \sim 1/\sqrt{K}$ to

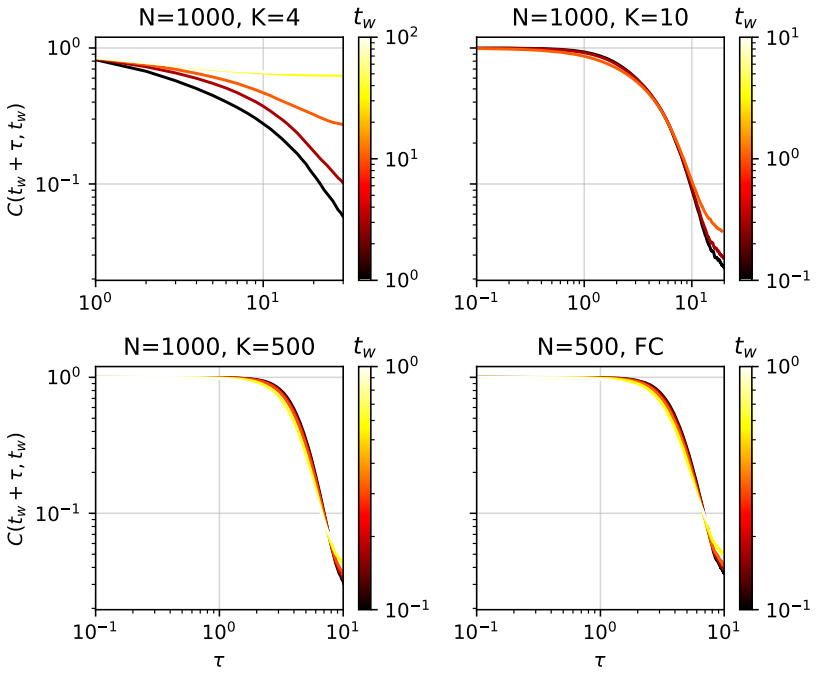

Fig. 6.4 **Finite-size and connectivity effects on ageing for 2-spin ferromagnet.** Monte Carlo results for $C(t_w+\tau,t_w)$ on RRGs with varying size and connectivity for homogeneous coupling $J=1.0$ and temperature $T=0.3$. Ageing is present at small K and vanishes as K increases, consistently with the scaling $J\sim 1/K$.

preserve extensivity, the ferromagnetic case requires $J\sim 1/K$. Consequently, the ferromagnet exhibits ageing in the sparse regime (small K) but relaxes exponentially in the dense limit, where J^2 contributions are suppressed. This behavior is confirmed numerically in Fig. 6.4, where ageing is visible at small K and progressively disappears as K increases.

The spherical 2-spin model thus provides a stringent test for the Gaussian-expansion cavity formalism. Without requiring explicit spectral information or global diagonalization, the method correctly reproduces both stationary and aging regimes, including their dependence on temperature and connectivity. This validates the GECaM as a powerful framework for studying non-equilibrium stochastic dynamics of linearly coupled degrees of freedom on sparse graphs.

Chapter 7

Complex Ecological Communities

Ecological communities are paradigmatic examples of complex systems, composed of a large number of interacting species whose collective dynamics emerge from the interplay between nonlinear interactions and environmental variability. A central question in theoretical ecology concerns how the structural properties of the interaction network affect the stability and diversity of the community, a problem that has been at the heart of theoretical investigations since the pioneering work of May [74, 75]. Within this framework, the Generalized Lotka-Volterra (GLV) equations [76–79] have emerged as a central model for describing the coupled population dynamics of many interacting species, where the interaction coefficients are treated as quenched random variables. This stochastic formulation enables the use of analytical methods from statistical physics to study macroscopic properties such as diversity, stability, and susceptibility of large ecosystems.

Most analytical progress so far has been achieved in the fully connected limit, where dynamical theory and path integral approaches provide a detailed description of the equilibria, extinction patterns, and stability boundaries of random ecosystems [78, 80, 81]. However, empirical studies indicate that real ecological networks are far from fully connected: empirical food webs and mutualistic networks display sparse and heterogeneous structures, with each species interacting with only a limited number of partners [82–86].

Sparsity is not merely a realistic structural detail, but rather a fundamental ingredient that qualitatively alters the collective dynamics. In dense random systems, central limit arguments lead to effective Gaussian noise and mean field-like behavior, allowing compact analytical treatments. When the number of interactions per species remains finite in the thermodynamic limit, this simplification breaks down: fluctuations no longer self-average, correlations between neighboring species persist, and local heterogeneity plays a crucial

dynamical role. Sparse interactions may thus generate non-Gaussian statistics of abundances even when the couplings themselves are Gaussian distributed [87], induce localized instabilities confined to small substructures of the network [88], or even give rise to collective transitions driven purely by the network topology [89]. In dense ecosystems, several works have shown that non-Gaussian abundance statistics can also emerge through different mechanisms: either by assuming non-Gaussian interaction distributions within the fully connected limit of DMFT [90], by introducing heterogeneous generalizations of DMFT in the dense regime [43, 45], or by considering Gaussian time-dependent interactions corresponding to annealed disorder [91]. While these studies demonstrate how deviations from Gaussianity can be rationalized within the tools of DMFT, they do not account for the intrinsic sparseness of real ecological interaction networks. Such effects cannot be captured within dense mean field theories and call for analytical frameworks explicitly designed for sparse interaction topologies.

In this chapter, we extend the cavity-based framework developed in the previous chapters to the study of the generalized Lotka-Volterra model on sparse random graphs. Building on the GECaM chapter 4, we derive effective stochastic equations describing single-species dynamics subject to self-consistent random fields and colored noise. This approach provides a tractable description of the fixed points of the sparse GLV dynamics and unveils how finite connectivity and interaction-disorder reshape the stability landscape, giving rise to rich dynamical regimes and non-Gaussian species abundance distributions.

7.1 The generalized Lotka-Volterra model

The GLV model provides a minimal yet versatile framework to describe the dynamics of ecological communities composed of many interacting species. Each species $i = 1, \dots, N$ is characterized by its abundance $x_i(t)$, whose temporal evolution is governed by a set of coupled nonlinear differential equations

$$\frac{dx_i(t)}{dt} = r_i x_i(t) \left[1 - \frac{x_i(t)}{K_i} + \sum_{j \in \partial i} J_{ij} x_j(t) \right]. \tag{7.1}$$

The first two terms describe the intrinsic growth of species i in isolation, with intrinsic growth rate r_i and carrying capacity K_i, while the last term accounts for pairwise interactions. The coefficients J_{ij} quantify the per-capita effect of species j on species i and are treated as quenched random variables. Negative couplings represent competitive interactions, while positive ones correspond to mutualistic or beneficial effects.

To isolate the effects of network sparsity, we consider undirected random graphs with a prescribed degree distribution $p_{\text{deg}}(k)$ with average value K. Throughout this chapter we focus primarily on RRGs, which eliminates degree heterogeneity while preserving the essential feature of finite connectivity, allowing analytical tractability and a clear interpretation of the emerging phenomena.

The interaction strengths $\{J_{ij}\}_{i,j=1}^{N}$ are independently drawn for each existing edge (i,j) from a Gaussian distribution with mean $\overline{J_{ij}} = m/K$, variance $\text{Var}(J_{ij}) = \sigma^2/K$, and correlation between couplings on the same edge $\text{Corr}(J_{ij}, J_{ji}) = \gamma$, so that J_{ij} and J_{ji} are correlated random variables controlling the degree of reciprocity between pairs. The parameter $\gamma \in [-1,1]$ interpolates between fully antisymmetric interactions ($\gamma = -1$), corresponding to purely predator-prey relationships, and fully symmetric ones ($\gamma = 1$), describing purely competitive or mutualistic interactions. Intermediate values of γ produce mixed ecosystems combining both types of interactions in varying proportions. We will denote throughout the chapter $\bar{\cdot}$ as the average over the quenched disorder in the couplings and graph ensemble.

Without loss of generality, we assume homogeneous intrinsic parameters, $r_i = 1$ and $K_i = 1$ for all species. In these units, the dynamics reduces to

$$\frac{dx_i(t)}{dt} = x_i(t)\left[1 - x_i(t) + \sum_{j \in \partial i} J_{ij} x_j(t)\right], \tag{7.2}$$

where ∂i denotes the set of neighbors of node i in the interaction graph. Equation (7.2) defines a nonlinear dynamical system with multiplicative interactions and quenched disorder, whose collective behavior depends crucially on the network topology and on the statistics of the couplings.

In the fully connected limit ($K = N-1$), the system admits a description, and the statistical properties of its stationary states can be analyzed through dynamical mean-field theory [77–79]. In contrast, when the degree K remains finite in the thermodynamic limit, correlations induced by the sparse topology become dominant, and mean-field approximations fail. Understanding how such sparse interactions affect stability thus requires a different analytical approach, which we develop in the following sections using a small-coupling expansion of the dynamic cavity equations.

7.2 Small-Coupling Expansion of the Dynamic Cavity Method

To study the stochastic dynamics of the generalized Lotka-Volterra model on sparse networks, we apply the dynamic cavity formalism introduced in section 2.3, following the steps detailed in section 4.1. We consider a fixed realization of disorder, specified by the interaction graph and coupling strengths.

The dynamical generating functional for the GLV dynamics Eq. (7.2) is written as

$$Z[\underline{\psi},\underline{\hat{\psi}}] = \int \mathscr{D}^0[\underline{x}] \prod_{i=1}^{N} \delta\left[\frac{\dot{x}_i}{x_i} - \left(1 - x_i + \sum_{j\in\partial i} J_{ij}x_j\right)\right] e^{\int dt\,[\psi_i(t)x_i(t) + \hat{\psi}_i(t)\hat{x}_i(t)]}, \tag{7.3}$$

where we have divided both sides of Eq. (7.2) by $x_i(t)$ to express the dynamics in terms of the logarithmic growth rate $\dot{x}_i/x_i$. This form requires abundances to remain positive, which is not strictly true as species may go extinct ($x_i = 0$). A careful derivation avoiding this assumption is reported in Appendix D.1.1, where we show that the final expressions remain unaffected.

Using the Fourier representation of the delta function, we obtain the MSRJD functional integral

$$Z[\underline{\psi},\underline{\hat{\psi}}] = \int \mathscr{D}^0[\underline{x},\underline{\hat{x}}] \exp\left[-S[\underline{x},\underline{\hat{x}}] + \sum_{i=1}^{N} \int_0^{t_f} dt\,\left(\psi_i(t)x_i(t) + \hat{\psi}_i(t)\,\mathrm{i}\hat{x}_i(t)\right)\right], \tag{7.4}$$

with action

$$S[\underline{x},\underline{\hat{x}}] = \sum_{i=1}^{N} S_i^{\mathrm{loc}}[x_i,\hat{x}_i] + \sum_{(i,j)\in E} S_{ij}^{\mathrm{int}}[x_i,\hat{x}_i,x_j,\hat{x}_j], \tag{7.5}$$

where the local and interaction contributions read

$$S_i^{\mathrm{loc}}[x_i,\hat{x}_i] = \int_0^{t_f} dt\,\mathrm{i}\hat{x}_i(t)\left[\frac{\dot{x}_i(t)}{x_i(t)} - 1 + x_i(t)\right], \tag{7.6}$$

$$S_{ij}^{\mathrm{int}}[x_i,\hat{x}_i,x_j,\hat{x}_j] = -\int_0^{t_f} dt\,\left(\mathrm{i}\hat{x}_i(t)J_{ij}x_j(t) + \mathrm{i}\hat{x}_j(t)J_{ji}x_i(t)\right). \tag{7.7}$$

The functional Eq. (7.4) defines a high-dimensional path integral whose action decomposes into local contributions and pairwise interactions along the nodes and edges of the network. This structure naturally defines a graphical model over continuous paths, to which we can apply the dynamic cavity formalism.

Following the derivation of section 4.1, we introduce the cavity marginals $c_{i\setminus j}[x_i,\hat{x}_i]$, representing the probability weight of the trajectory of node i when its link to j is removed. These

obey the recursive relations

$$c_{i\setminus j}[x_i,\hat{x}_i] = \frac{1}{Z_{i\setminus j}} e^{-S_i^{\mathrm{loc}}[x_i,\hat{x}_i]} \prod_{k\in\partial i\setminus j} \int \mathscr{D}[x_k,\hat{x}_k]\, c_{k\setminus i}[x_k,\hat{x}_k]\, e^{-S_{ik}^{\mathrm{int}}[x_i,\hat{x}_i,x_k,\hat{x}_k]}, \tag{7.8}$$

where $Z_{i\setminus j}$ ensures normalization.

We now perform a small-coupling expansion of Eq. (7.8) up to second order in the interaction strengths J_{ij}. Neglecting averages involving only response fields (which vanish identically in the MSRJD formalism), we obtain a Gaussian approximation for the cavity marginals, fully determined by their first and second moments. These can be written as

$$c_{i\setminus j}[x_i,\hat{x}_i] \approx \frac{1}{Z_{i\setminus j}} e^{-S_{i\setminus j}^{\mathrm{eff}}[x_i,\hat{x}_i]}, \tag{7.9}$$

where the effective action reads

$$\begin{aligned} S_{i\setminus j}^{\mathrm{eff}}[x_i,\hat{x}_i] = &\int_0^{t_f} dt\, \mathrm{i}\hat{x}_i(t) \left[\frac{\dot{x}_i(t)}{x_i(t)} - 1 + x_i(t) - \sum_{k\in\partial i\setminus j} J_{ik}\mu_{k\setminus i}(t)\right] \\ &+ \frac{1}{2}\int_0^{t_f} dt\, dt' \sum_{k\in\partial i\setminus j} \left[\mathrm{i}\hat{x}_i(t) J_{ik}^2 C_{i\setminus j}(t,t')\, \mathrm{i}\hat{x}_i(t') - 2\,\mathrm{i}\hat{x}_i(t) J_{ik}J_{ki} R_{k\setminus i}(t,t')\, x_i(t')\right]. \end{aligned} \tag{7.10}$$

The functions $\mu_{i\setminus j}$, $C_{i\setminus j}$, and $R_{i\setminus j}$ denote respectively the cavity mean abundance, correlation, and response,

$$\mu_{i\setminus j}(t) = \langle x_i(t)\rangle_{i\setminus j}, \tag{7.11}$$

$$C_{i\setminus j}(t,t') = \langle x_i(t)x_i(t')\rangle_{i\setminus j} - \mu_{i\setminus j}(t)\mu_{i\setminus j}(t'), \tag{7.12}$$

$$R_{i\setminus j}(t,t') = \langle x_i(t)\,\mathrm{i}\hat{x}_i(t')\rangle_{i\setminus j} = \left.\frac{\delta\langle x_i(t)\rangle_{i\setminus j}}{\delta h_i(t')}\right|_{h=0}. \tag{7.13}$$

The effective action Eq. (7.10) describes a non-interacting single-species process subject to self-consistent colored noise and memory kernels originating from its neighbors on the cavity graph.

Reversing the steps of the MSRJD construction (see section 3.2), we obtain the corresponding effective stochastic dynamics for the cavity abundance $x_{i\setminus j}(t)$,

$$\frac{dx_{i\setminus j}(t)}{dt} = x_{i\setminus j}(t)\Bigg[1 - x_{i\setminus j}(t) + \sum_{k\in\partial i\setminus j} J_{ik}\mu_{k\setminus i}(t) + h_{i\setminus j}(t) + \sum_{k\in\partial i\setminus j} J_{ik}J_{ki}\int_0^t dt'\, R_{k\setminus i}(t,t')\, x_{i\setminus j}(t') + \eta_{i\setminus j}(t)\Bigg], \tag{7.14}$$

where $h_{i\setminus j}(t)$ is an external field used to compute response functions, and $\eta_{i\setminus j}(t)$ is a Gaussian noise with zero mean and covariance

$$\langle \eta_{i\setminus j}(t)\eta_{i\setminus j}(t')\rangle = \sum_{k\in\partial i\setminus j} J_{ik}^2 C_{k\setminus i}(t,t'). \tag{7.15}$$

The effective process Eq. (7.14) closely resembles the dynamical mean-field description obtained in the fully connected limit [81], with the crucial difference that here no disorder-average is performed. The stochasticity and memory terms emerge instead from the fluctuations of the neighboring species in the cavity graph. Consequently, the process depends explicitly on the quenched couplings $\{J_{ik}\}_{k\in\partial i\setminus j}$ and on the cavity moments of the neighbors, which must satisfy the self-consistency conditions

$$\mu_{i\setminus j}(t) = \langle x_{i\setminus j}(t)\rangle_{i\setminus j}, \tag{7.16}$$

$$C_{i\setminus j}(t,t') = \langle x_{i\setminus j}(t)x_{i\setminus j}(t')\rangle_{i\setminus j} - \mu_{i\setminus j}(t)\mu_{i\setminus j}(t'), \tag{7.17}$$

$$R_{i\setminus j}(t,t') = \left.\frac{\delta\langle x_{i\setminus j}(t)\rangle_{i\setminus j}}{\delta h_{i\setminus j}(t')}\right|_{h=0}. \tag{7.18}$$

A similar effective process follows from the full marginal $c_i[x_i,\hat{x}_i]$, describing the dynamics of species i in the original graph. Repeating the same procedure, we obtain

$$\frac{dx_i(t)}{dt} = x_i(t)\Bigg[1 - x_i(t) + \sum_{k\in\partial i} J_{ik}\mu_{k\setminus i}(t) + \sum_{k\in\partial i} J_{ik}J_{ki}\int_0^t dt'\, R_{k\setminus i}(t,t')\, x_i(t') + \eta_i(t)\Bigg], \tag{7.19}$$

where the Gaussian noise satisfies

$$\langle \eta_i(t)\eta_i(t')\rangle = \sum_{k\in\partial i} J_{ik}^2 C_{k\setminus i}(t,t'). \tag{7.20}$$

The self-consistent nature of the fields and the emergence of colored noise are unavoidable consequences of the cavity decoupling procedure. While numerical schemes exist for solving such self-consistent non-Markovian equations in fully connected systems [92], their extension to large sparse networks would entail a prohibitive computational cost. To obtain analytical insight, in the next section we focus on the stationary fixed point regime of the effective processes, from which closed equations for the local observables and stability properties can be derived.

7.3 Fixed point equations

The effective stochastic equations derived in the previous section describe non-Markovian single-species dynamics with colored noise and memory kernels determined self-consistently by the neighboring species. Solving these equations in full generality is computationally demanding. To make analytical progress, we focus on the stationary regime where the system relaxes to a fixed point independent of the initial conditions. In this limit, all time-dependent quantities become static, allowing for a closed set of self-consistent equations for the cavity moments and responses.

We thus assume that the dynamics converges to a time-independent state, such that

$$\lim_{t\to\infty} x_i(t) = x_i^\star. \tag{7.21}$$

An analogous assumption is made for the cavity effective processes,

$$\lim_{t\to\infty} x_{i\setminus j}(t) = x_{i\setminus j}^\star, \qquad \lim_{t\to\infty} \eta_{i\setminus j}(t) = \eta_{i\setminus j}^\star. \tag{7.22}$$

Since $x_i(t)$ and $x_{i\setminus j}(t)$ are stochastic processes, assuming they become time-independent is equivalent to assuming they become static random variables with time-independent distributions. Consequently, the mean cavity abundance and two-point correlation become time independent,

$$\lim_{t\to\infty} \mu_{i\setminus j}(t) = \mu_{i\setminus j}^\star = \langle x_{i\setminus j}^\star \rangle_{i\setminus j}^\star, \tag{7.23}$$

$$\lim_{t,t'\to\infty} C_{i\setminus j}(t,t') = q_{i\setminus j}^\star = \langle (x_{i\setminus j}^\star)^2 \rangle_{i\setminus j}^\star - (\mu_{i\setminus j}^\star)^2, \tag{7.24}$$

where $\langle\cdot\rangle_{i\setminus j}^\star$ denotes the average over the stationary cavity process. Under this assumption, the cavity response function becomes time-translationally invariant and depends only on the time

difference $\tau = t - t'$. We define its time integral,

$$\chi^{\star}_{i\backslash j} = \int_0^{\infty} d\tau R^{\star}_{i\backslash j}(\tau), \tag{7.25}$$

which represents the integrated response (or susceptibility) of the cavity process.

At the fixed point, both $x^{\star}_{i\backslash j}$ and the effective field $\eta^{\star}_{i\backslash j}$ become static random variables. The variable $x^{\star}_{i\backslash j}$ depends functionally on $\eta^{\star}_{i\backslash j}$, which is itself Gaussian distributed with zero mean and variance

$$\langle (\eta^{\star}_{i\backslash j})^2 \rangle^{\star}_{i\backslash j} = \sum_{k\in\partial i\backslash j} J_{ik}^2 q^{\star}_{k\backslash i} \equiv \Sigma^2_{i\backslash j}. \tag{7.26}$$

The fixed point abundances satisfy the self-consistent relations

$$0 = x^{\star}_{i\backslash j}\left[1 - x^{\star}_{i\backslash j} + \sum_{k\in\partial i\backslash j} J_{ik}\mu^{\star}_{k\backslash i} + \sum_{k\in\partial i\backslash j} J_{ik}J_{ki}\chi^{\star}_{k\backslash i}x^{\star}_{i\backslash j} + \eta^{\star}_{i\backslash j} + h_{i\backslash j}\right], \tag{7.27}$$

where $h_{i\backslash j}$ is a constant external field introduced to compute the susceptibility. By performing the change of variable $\eta^{\star}_{i\backslash j} = -\Sigma_{i\backslash j}s$, where s is a standard Gaussian variable with zero mean and unit variance, Eq. (7.27) admits a non-vanishing solution

$$x^{\star}_{i\backslash j}(s) = \frac{\Sigma_{i\backslash j}(\Delta_{i\backslash j} - s) + h_{i\backslash j}}{\Gamma_{i\backslash j}}\,\Theta\left(\frac{\Sigma_{i\backslash j}(\Delta_{i\backslash j} - s) + h_{i\backslash j}}{\Gamma_{i\backslash j}}\right), \tag{7.28}$$

where $\Theta(x)$ denotes the Heaviside step function enforcing positivity of the fixed point abundance. We have defined

$$\Delta_{i\backslash j} = \frac{1 + \sum_{k\in\partial i\backslash j} J_{ik}\mu^{\star}_{k\backslash i}}{\Sigma_{i\backslash j}}, \tag{7.29}$$

$$\Gamma_{i\backslash j} = 1 - \sum_{k\in\partial i\backslash j} J_{ik}J_{ki}\chi^{\star}_{k\backslash i}. \tag{7.30}$$

Averaging Eq. (7.28) over the Gaussian variable s yields a closed set of self-consistent equations for the cavity moments, which read

$$\mu^{\star}_{i\backslash j} = \mathscr{F}_{\mu}\left(\Delta_{i\backslash j}, \Sigma_{i\backslash j}, \Gamma_{i\backslash j}\right), \tag{7.31}$$

$$q^{\star}_{i\backslash j} = \mathscr{F}_{q}\left(\Delta_{i\backslash j}, \Sigma_{i\backslash j}, \Gamma_{i\backslash j}\right) - (\mu^{\star}_{i\backslash j})^2, \tag{7.32}$$

$$\chi^{\star}_{i\backslash j} = \mathscr{F}_{\chi}\left(\Delta_{i\backslash j}, \Gamma_{i\backslash j}\right), \tag{7.33}$$

where the update functions are given explicitly by

$$\mathscr{F}_\mu(\Delta,\Sigma,\Gamma) = \frac{\Sigma}{\Gamma}\Big[\Theta(\Gamma)\Big(\Delta\Phi(\Delta)+\varphi(\Delta)\Big)+\Theta(-\Gamma)\Big(\Delta\Phi(-\Delta)-\varphi(-\Delta)\Big)\Big] \tag{7.34}$$

$$\mathscr{F}_q(\Delta,\Sigma,\Gamma) = \frac{\Sigma^2}{\Gamma^2}\Big[\Theta(\Gamma)\Big((1+\Delta^2)\Phi(\Delta)+\Delta\varphi(\Delta)\Big) + \Theta(-\Gamma)\Big((1+\Delta^2)\Phi(-\Delta)-\Delta\varphi(-\Delta)\Big)\Big], \tag{7.35}$$

$$\mathscr{F}_\chi(\Delta,\Gamma) = \frac{1}{\Gamma}\big[\Theta(\Gamma)\,\Phi(\Delta)+\Theta(-\Gamma)\,\Phi(-\Delta)\big]. \tag{7.36}$$

We have introduced the shorthand notations $\varphi(x) = e^{-x^2/2}/\sqrt{2\pi}$ and $\Phi(x) = [1+\mathrm{erf}(x/\sqrt{2})]/2$ which denote, respectively, the standard normal probability density and cumulative distribution functions. A detailed derivation of Eqs. (7.31) to (7.36) is provided in Appendix section D.1.2.

An equivalent set of equations holds for the full marginals, describing the species abundances in the original graph. Denoting the fixed point moments by $\mu_i^\star$, $q_i^\star$, and $\chi_i^\star$, we have

$$\mu_i^\star = \mathscr{F}_\mu(\Delta_i,\Sigma_i,\Gamma_i), \tag{7.37}$$

$$q_i^\star = \mathscr{F}_q(\Delta_i,\Sigma_i,\Gamma_i) - (\mu_i^\star)^2, \tag{7.38}$$

$$\chi_i^\star = \mathscr{F}_\chi(\Delta_i,\Gamma_i), \tag{7.39}$$

where the update functions are the same as in Eqs. (7.34) to (7.36), and we have defined

$$\Delta_i = \frac{1+\sum_{k\in\partial i} J_{ik}\mu_{k\setminus i}^\star}{\Sigma_i}, \tag{7.40}$$

$$\Sigma_i^2 = \sum_{k\in\partial i} J_{ik}^2\, q_{k\setminus i}^\star, \tag{7.41}$$

$$\Gamma_i = 1 - \sum_{k\in\partial i} J_{ik}J_{ki}\chi_{k\setminus i}^\star. \tag{7.42}$$

The coupled set of nonlinear equations Eqs. (7.31) to (7.33) defines the fixed point statistics of the cavity processes in terms of their local fields. These relations can be solved numerically via a message-passing procedure that iteratively updates the cavity quantities by substituting the moments of neighboring species.

7.3.1 Population dynamics algorithm for quenched disorder

To compute disorder-averaged observables in the thermodynamic limit, we solve the self-consistent equations Eqs. (7.31) to (7.33) through a population dynamics algorithm. This stochastic iterative method propagates the distribution of the cavity fields under the joint randomness of network topology and interaction disorder, providing an efficient way to sample the typical stationary state of large random ecosystems.

The algorithm proceeds as follows:

1. Initialize a population $\mathscr{P}$ of P cavity variables, each represented by a triplet $(\mu^{\mathrm{c}}, q^{\mathrm{c}}, \chi^{\mathrm{c}})$ corresponding to the mean abundance, variance, and susceptibility of a cavity process.

2. At each iteration:

   (a) Sample a degree k from the cavity degree distribution

   $$p^{\mathrm{c}}_{\mathrm{deg}}(k) = \frac{k+1}{K}\, p_{\mathrm{deg}}(k+1).$$

   (b) randomly-draw k triplets $\{(\mu^{\mathrm{c}}_j, q^{\mathrm{c}}_j, \chi^{\mathrm{c}}_j)\}_{j=1}^{k}$ from the current population to represent the neighbors in the cavity graph.

   (c) Draw k pairs of random couplings $\{(J_j, J'_j)\}_{j=1}^{k}$ from the coupling distribution.

   (d) Compute the effective parameters:

   $$\Delta^{\mathrm{c}} = \frac{1 + \sum_{j=1}^{k} J_j \mu^{\mathrm{c}}_j}{\Sigma^{\mathrm{c}}},$$
   $$\Sigma^{\mathrm{c}} = \sqrt{\sum_{j=1}^{k} J_j^2 q^{\mathrm{c}}_j},$$
   $$\Gamma^{\mathrm{c}} = 1 - \sum_{j=1}^{k} J_j J'_j \chi^{\mathrm{c}}_j.$$

   (e) Update the new fields according to

   $$\mu^{\mathrm{c}} = \mathscr{F}_{\mu}(\Delta^{\mathrm{c}}, \Sigma^{\mathrm{c}}, \Gamma^{\mathrm{c}}),$$
   $$q^{\mathrm{c}} = \mathscr{F}_{q}(\Delta^{\mathrm{c}}, \Sigma^{\mathrm{c}}, \Gamma^{\mathrm{c}}, \mu^{\mathrm{c}}),$$
   $$\chi^{\mathrm{c}} = \mathscr{F}_{\chi}(\Delta^{\mathrm{c}}, \Gamma^{\mathrm{c}}).$$

(f) Replace softly a randomly-chosen element of the population with the new triplet (μ^c, q^c, χ^c).

3. Repeat until convergence of the population statistics, as assessed from the empirical averages of μ^c, q^c, and χ^c.

After convergence, single-site observables for the full graph can be estimated by performing one additional update step using degrees drawn from p_{deg} instead of p^c_{deg}, thus completing the cavity construction. In RRGs, the degree distribution is a Dirac delta centered at K, while the cavity degree distribution is centered at $K-1$. Hence, a randomly-sampled node has degree K, whereas a cavity node has degree $K-1$.

This algorithm captures the self-consistent statistical behavior of species in ensembles of random regular graphs with quenched disordered interactions, allowing the estimation of typical macroscopic observables in the thermodynamic limit. Further implementation details of the algorithm are provided in Appendix D.2.

7.4 Species abundance distributions

The fixed point equations derived in the previous section provide a complete characterization of the stationary state of the generalized Lotka-Volterra dynamics on sparse random graphs. In particular, they allow us to compute the *Species Abundance Distribution (SAD)* across the ecosystem, revealing how sparsity and interaction-disorder shape the statistical properties of the community.

To assess the accuracy of the theoretical predictions, we compare the species abundance distributions obtained from the *Population Dynamics (PopDyn)* algorithm with those measured from direct numerical integration of the microscopic dynamics defined in Eq. (7.2) (hereafter referred to as MC). The numerical integration is performed by evolving the system until convergence to a stationary state for $N = 500$ species interacting on random regular graphs with fixed degree K. For each parameter configuration (m, σ, γ), averages are computed over 20 independent realizations of the interaction graph and coupling matrix, and for each of them over 20 distinct random initial conditions uniformly sampled in the interval $x_i(0) \in (0, 1]$. The PopDyn results are instead obtained from the full marginal distribution of the population dynamics algorithm at convergence, using populations of size ranging from $P = 10^5$ to 10^7 depending on the convergence properties.

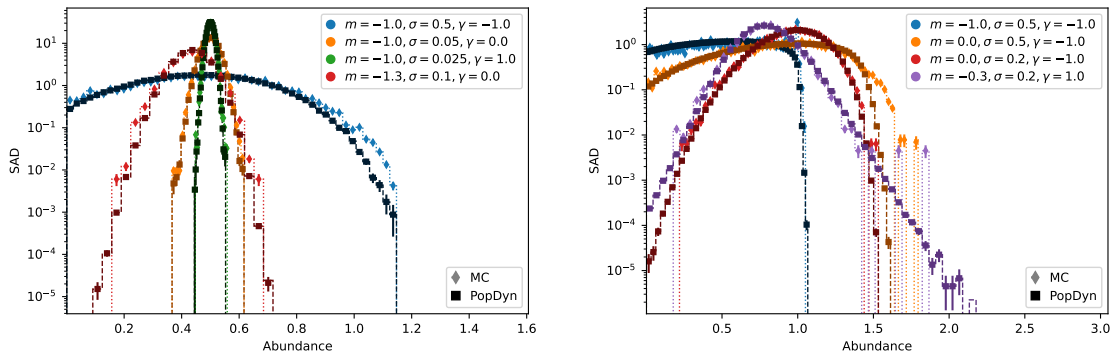

(a) $K = 10$ (b) $K = 3$

Fig. 7.1 **Species abundance distributions on sparse random graphs.** Comparison between stationary abundance distributions (SADs) obtained from direct numerical integration of the generalized Lotka-Volterra dynamics (MC, diamonds) and from the fixed point solution of the population dynamics equations (PopDyn, squares). Results correspond to RRGs with different connectivities and various combinations of coupling parameters (m, σ, γ). Each MC histogram is obtained by aggregating equilibrium abundances from 20 graph realizations and 20 initial conditions per realization. PopDyn results are obtained at convergence of populations of size $P \in [10^5, 10^7]$. Vertical bars indicate standard errors of the MC histograms.

As shown in Fig. 7.1, the PopDyn results reproduce with remarkable accuracy the stationary abundance distributions measured from microscopic simulations, confirming the validity of the fixed point equations derived from the cavity formalism. For $K = 10$, the SADs are well approximated by Gaussian curves, with only mild skewness visible for $(m, \sigma, \gamma) = (-1.3, 0.1, 0)$ and $(m, \sigma, \gamma) = (1.0, 0.5, 1.0)$. In these cases, the PopDyn predictions closely follow the MC results, accurately capturing the small asymmetries observed. The residual differences are minor and mostly confined to the tails, where finite-sample fluctuations are more relevant.

For smaller connectivities, such as $K = 3$, deviations from Gaussianity become more pronounced. The SADs exhibit broader, asymmetric shapes with heavy tails and a higher fraction of extinct species, reflecting the stronger heterogeneity induced by local fluctuations in sparse interaction networks. Despite these nontrivial features, the PopDyn predictions remain in close quantitative agreement with the MC distributions, faithfully reproducing both the central bulk and the asymmetry of the empirical curves. The slightly larger discrepancies observed in the tails can be attributed to limited statistics, which are particularly relevant when the number of surviving species is small.

At small disorder ($\sigma \ll 1$), the distributions are narrow and approach a Dirac delta function, corresponding to stable communities where all species maintain similar abundances. As the disorder increases, the distributions broaden and develop asymmetry, signaling the emergence

of heterogeneous communities in which a few species dominate while many persist at low abundance or go extinct.

Overall, the excellent agreement between PopDyn and MC results demonstrates that the Gaussian expansion of the dynamic cavity equations captures the essential statistical features of the stationary ecosystem, including the non-Gaussian distortions induced by finite connectivity and interaction disorder. These findings confirm that sparse interactions fundamentally reshape the macroscopic organization of complex ecological communities, even in the absence of external noise or environmental heterogeneity.

7.5 Stability of complex ecosystems with sparse interactions

The Gaussian expansion of the dynamic cavity equations provides a tractable framework to study not only stationary abundance distributions but also the global stability properties of large ecological networks. In this section, we focus on the phase diagram of the random generalized Lotka-Volterra (rGLV) dynamics on sparse random regular graphs, emphasizing how finite connectivity alters the onset of instability and the qualitative structure of the stationary phase space. To capture the distinct effects of sparsity, we perform our analysis on graphs with varying degrees of connectivity, specifically contrasting the behavior of moderately sparse networks ($K = 10$) with the regime of strong sparsity ($K = 3$), while varying the control parameters (m, σ, γ) that define the coupling statistics. A note on terminology is necessary here: throughout this section, we use the term *disorder* to refer strictly to the quenched randomness in the coupling coefficients $\{J_{ij}\}$. While the topology of the interaction network is also random, the use of RRGs ensures that the graph is locally homogeneous (all nodes have degree K); consequently, we treat the network structure as a source of specific topological constraints rather than a source of disorder.

A distinctive feature of sparse ecosystems is the emergence, even in the absence of disorder, of a collective transition between a phase with a single globally stable equilibrium and a phase with multiple alternative equilibria. This *topological phase transition* is not induced by heterogeneity in the couplings but by the combinatorial constraints imposed by finite connectivity. Its origin can be traced to the competitive exclusion between neighboring species: when the effective competition $|m|$ exceeds a critical threshold, loops in the sparse graph frustrate global coexistence, forcing the community to fragment into alternative stable states. This phenomenon, first observed numerically by Marcus and collaborators [89], is absent in fully connected systems, making it an ideal case study for our framework.

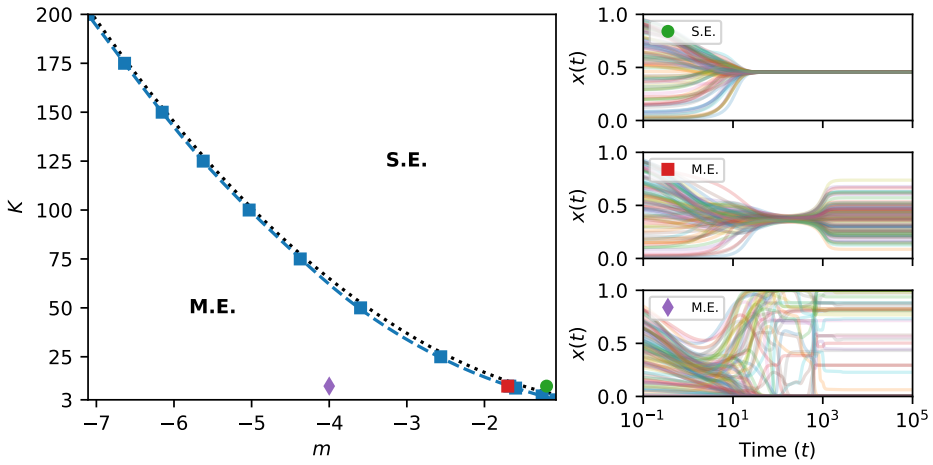

Fig. 7.2 **Topological phase transition at zero disorder.** Phase diagram in the (m, K) plane for RRGs with K neighbors and homogeneous couplings $J_{ij} = m/K$. *Markers* indicate the critical point $m_c(K)$ where PopDyn develops a non-vanishing variance q, signaling the transition from a globally stable Single Equilibrium (SE) to a fragmented Multiple Equilibria (ME) phase. The blue dashed line has been obtained by interpolation. The analytical prediction (black dotted line) corresponds to the loss of linear stability of the fully feasible equilibrium, determined by the Kesten-McKay spectral edge of the adjacency matrix. On the right, examples of microscopic trajectories (MC) are shown for three representative values of m: from top to bottom, $m = -1.2$ (single-equilibrium phase), $m = -1.7$ (multiple equilibria near the transition), and $m = -4$ (multiple equilibria with strong competition).

To disentangle this intrinsically topological mechanism from the effects of heterogeneity, we first analyze in section 7.5.1 the case of vanishing disorder ($\sigma = 0$), where all couplings take the deterministic value $J_{ij} = m/K$. In this regime, the system's behavior is governed solely by the mean interaction strength and the network connectivity, and the onset of instability can be predicted analytically from the spectrum of the interaction matrix. We then examine in section 7.5.2 the full phase diagram with disordered interactions ($\sigma > 0$). By varying the interaction symmetry γ from competitive to predator-prey types, we reveal how structural frustration and disorder synergize to reshape the stability landscape, particularly in the limit of low connectivity where mean-field intuitions break down.

7.5.1 Topological phase transition at zero disorder

The structural effects of sparseness can be isolated by setting the strength of the disorder $\sigma = 0$, so that the system is governed solely by the mean interaction m and connectivity K. A distinctive feature of sparse ecosystems is the emergence of a topological phase transition from a single globally stable equilibrium (SE) to a Multiple Equilibria (ME) phase. This transition is

identified within our cavity framework as the point where the distribution of the local variance q is no longer concentrated at zero.

Our derivation explicitly assumes that the system relaxes to a unique fixed point independent of initial conditions; this implies that the stationary state of the effective process is deterministic, and thus its self-consistent variance must vanish ($q = 0$). For m above a critical threshold $m_c(K)$, the population dynamics confirms this hypothesis, relaxing to a unique solution where the entire population satisfies $q = 0$. As m drops below $m_c(K)$, the assumption of a unique attractor breaks down. In this regime, the effective process develops a non-trivial variance ($q > 0$), reflecting the dependence on initial conditions characteristic of a landscape with multiple locally stable attractors. Consequently, the population dynamics iterations typically fail to converge to a stationary distribution, exhibiting persistent oscillations or divergence. Although the fixed-point theory cannot be used to quantitatively describe the interior of this phase, the breakdown of the assumption $q = 0$ (often coinciding with the loss of algorithmic convergence) accurately marks the stability boundary. We refer to section D.3.1 for the algorithmic details on how these critical points are identified from the PopDyn data.

Figure 7.2 shows the phase diagram obtained for $\sigma = 0$ on random regular graphs with K neighbors. The numerical transition points coincide perfectly with the analytical stability bound derived from the spectral properties of the interaction matrix. In the absence of disorder, the stability of the fully feasible equilibrium is determined by the condition that the matrix $\underline{\underline{I}} + (m/K)\underline{\underline{A}}$ is positive definite. The spectral distribution of the adjacency matrix $\underline{\underline{A}}$ follows the Kesten-McKay law, supported in the interval $[-2\sqrt{K-1}, 2\sqrt{K-1}]$. Hence, the critical value of m is reached when the spectral edge crosses unity, i.e.,

$$m_c(K) = -\frac{K}{2\sqrt{K-1}}. \tag{7.43}$$

This result establishes that the instability is purely topological, driven by combinatorial exclusion constraints rather than random heterogeneity: when effective competition exceeds a threshold, loops in the sparse graph frustrate global coexistence, forcing the community to fragment.

7.5.2 Phase diagram with disordered interactions

We now turn to the case of disordered interactions, $\sigma > 0$, and study how the stability landscape of the rGLV model evolves when randomness is introduced in the couplings. The analysis is performed by systematically varying the parameters (m, σ) for different values of the interaction

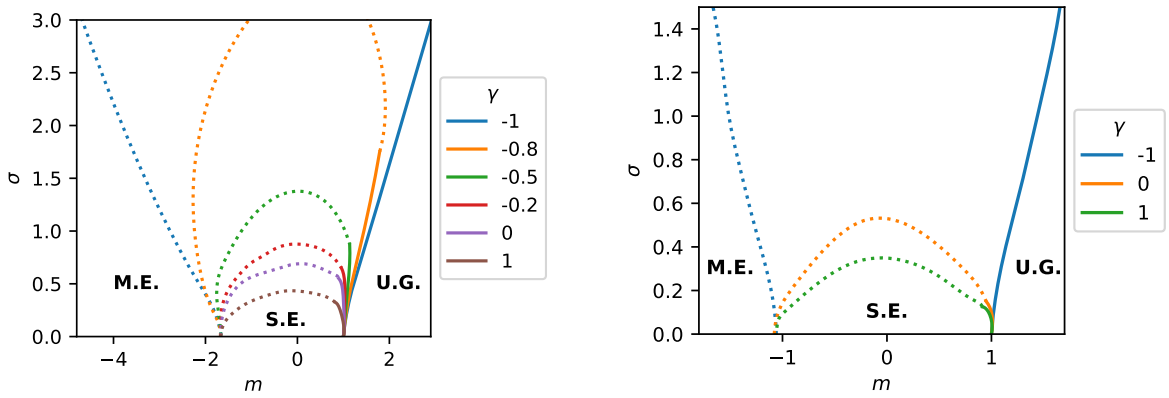

(a) $K = 10$ (b) $K = 3$

Fig. 7.3 **Stability phase diagrams in the (m, σ) plane.** Results for RRGs with $K = 10$ (left) and $K = 3$ (right) obtained via PopDyn. Dashed lines mark the transition from the Single Equilibrium (SE) to the Multiple Equilibria (ME) phase, identified by the onset of a non-vanishing variance distribution ($q > 0$). Solid lines mark the transition to the Unbounded Growth (UG) phase, identified by the divergence of the mean abundance μ. Different colors correspond to interaction symmetry values $\gamma \in [-1, 1]$ (see legend). Note the shrinking of the stable SE region as $\gamma \to 1$ and the severe compression of stability at low connectivity ($K = 3$).

symmetry $\gamma \in [-1, 1]$, and for two distinct connectivity regimes: a moderately sparse case ($K = 10$) and a high-sparsity case ($K = 3$).

The complete stability phase diagrams are shown in Fig. 7.3. Three distinct phases can be identified based on the behavior of the PopDyn algorithm:

- **Single Equilibrium (SE):** The algorithm converges to a unique fixed point where the distribution of the cavity variance q is concentrated at zero. This corresponds to a unique globally stable fixed point.
- **Multiple Equilibria (ME):** The distribution of q shifts towards strictly positive values and broadens; simultaneously, PopDyn iterations typically become oscillatory or diverge, signaling the fragmentation of the abundance landscape.
- **Unbounded Growth (UG):** The distribution of the mean abundance μ diverges during the iterations, indicating the absence of finite fixed points and corresponding to the existence of unbounded trajectories in the microscopic dynamics.

The phase boundaries shown in Fig. 7.3 are obtained from the interpolation of the critical points determined by PopDyn. The detailed phase diagrams displaying the calculated points, along with a full explanation of the interpolation procedure, are reported in Appendix D.3.2.

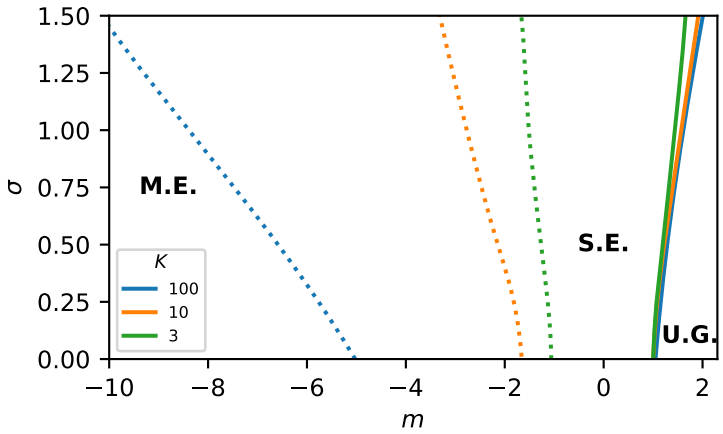

Fig. 7.4 **Phase diagrams for predator-prey interactions ($\gamma = -1$) at varying connectivity.** Results for RRGs with $K = 3$ (green), $K = 10$ (orange) and $K = 100$ (blueo) btained via PopDyn. Dashed lines mark the transition from the Single Equilibrium (SE) to the Multiple Equilibria (ME) phase, identified by the onset of a non-vanishing variance distribution ($q > 0$). Solid lines mark the transition to the Unbounded Growth (UG) phase, identified by the divergence of the mean abundance μ. As connectivity increases, the SE-ME critical line shifts systematically towards larger negative values of m, confirming that the ME phase disappears in the thermodynamic limit of fully connected models ($K \to \infty$).

The phase diagrams reveal a non-trivial synergy between topology and disorder. As seen in Fig. 7.3a ($K = 10$), the size of the stable SE phase is strongly dependent on the coupling correlation γ. As γ increases from -1 (predator-prey) to 1 (symmetric competition/mutualism), the SE region shrinks systematically. Crucially, the boundary between the SE and ME phases is not horizontal as in the fully connected limit [81]; instead, it bends markedly towards the $\sigma = 0$ axis, eventually intersecting it at the topological critical point $m_c(K)$. This bending indicates that structural frustration amplifies the effect of interaction heterogeneity, driving the system into a fragmented state at coupling variance σ^2 much weaker than those predicted by mean-field theory for dense systems.

This structural fragility is further amplified at low connectivity ($K = 3$, Fig. 7.3b). In this regime, the SE region is markedly narrower than for $K = 10$, underscoring the increased sensitivity of sparse networks to local competitive imbalances. The synergy between topology and disorder manifests as a stronger curvature of the SE-ME boundary, further constraining the stability domain. This contraction reflects the breakdown of self-averaging within small neighborhoods: with only three neighbors, local fluctuations and asymmetric motifs can disrupt the unique stable state more readily than in denser networks.

Finally, for fully antisymmetric interactions ($\gamma = -1$), the ME phase persists at finite connectivity, extending the topological phase observed at $\sigma = 0$ into the disordered regime.

This behavior is in sharp contrast to the infinite-connectivity limit, where the $\gamma = -1$ system remains in a single stable equilibrium up to the UG transition. To explicitly visualize this finite-size effect, Fig. 7.4 compares the phase diagrams obtained for varying degrees of sparsity ($K = 3, 10, 100$). While the Unbounded Growth boundary remains largely insensitive to connectivity, the critical line marking the onset of multistability shifts systematically towards larger negative values of m as K increases. This trend confirms that the ME region retreats in the dense limit ($K \to \infty$), eventually disappearing as predicted by mean-field theory. The existence of the ME phase for predator-prey communities is thus a hallmark of sparsity, arising from localized feedback loops that destabilize the unique stationary state once the mean interaction m becomes sufficiently negative.

Part III

Inference in Compartmental Epidemic Models

Discrete Compartmental Epidemic models

Epidemic processes on networks constitute a paradigmatic example of stochastic dynamics with discrete states. In individual-based compartmental models, each node represents an individual whose health state evolves according to stochastic transition rules, while edges encode potential transmission paths. Despite their Markovian nature–where the future state depends only on the present–the network structure induces complex spatiotemporal correlations that challenge analytical treatment.

Compartmental epidemic models can be broadly classified into two families. Non-recurrent models, such as Susceptible-Infected (SI), or Susceptible-Infected-Recovered (SIR), describe irreversible dynamics through compartments, and are suited to diseases conferring permanent recovery or immunity. Recurrent models, including Susceptible-Infected-Susceptible (SIS), or Susceptible-Infected-Recovered-Susceptible (SIRS), allow individuals to return to previously occupied compartments, thereby generating stationary or endemic states. While non-recurrent models can be solved exactly on tree-like networks using dynamic-message-passing or belief propagation, recurrent models pose significant computational challenges due to the exponential growth of possible trajectories over time.

A central practical problem is that of *epidemic reconstruction*: given partial and possibly noisy observations of a subset of individuals at scattered times, can one reconstruct the most probable epidemic trajectory, or estimate the infection probability of unobserved individuals? Within a Bayesian framework, this problem amounts to combining the stochastic epidemic dynamics, which defines the prior distribution over trajectories, with an observation model that specifies the likelihood of the observed data. The resulting posterior distribution encodes all available information on the underlying epidemic process, but computing its marginals is in general an intractable task due to the exponential growth of trajectories with time and system size.

Exact inference can, in principle, be carried out using *dynamic cavity equations* or *Belief Propagation (BP)*, which provide locally consistent approximations to the full posterior on

tree-like networks. However, the computational cost of these approaches scales exponentially with the duration of the dynamics when recurrent processes are considered, limiting their applicability to short time horizons or irreversible dynamics.

This part of the thesis introduces the *Small-Coupling Dynamic Cavity (SCDC)* method, a polynomial time approximation scheme derived from a perturbative expansion of the dynamic cavity equations in the weak-coupling regime. In close analogy with the Gaussian-Expansion Cavity Method developed for continuous stochastic dynamics, SCDC performs a systematic small-coupling expansion of the observation-reweighted dynamic cavity equations. At first order, the expansion yields a closed set of self-consistent equations for two-time-dependent cavity quantities: the probability of individual infection and the conjugate field encoding the conditioning effect of observations. The latter reintroduces non-causal, backward-in-time information flow, a hallmark of constrained dynamical systems. Unlike the continuous cases presented in Part II, averages of response fields do not vanish in this setting, since the presence of observations explicitly breaks causality.

The SCDC method thus provides a computationally efficient and physically interpretable approximation to the full Bayesian inference problem. It retains the local structure and message-passing nature of belief propagation while drastically reducing its complexity.

In what follows, we first introduce individual-based epidemic models on networks and formalize the Bayesian inference problem. We then derive the SCDC equations for the SI model and present their efficient algorithmic formulation. Finally, we assess the method's accuracy and generality by comparing it with belief propagation and Monte Carlo simulations on both irreversible and recurrent epidemic models.

All chapters in this part are adapted from the article [93], in which the Small-Coupling Dynamic Cavity method was developed and applied to epidemic inference and risk assessment on contact networks. The data and code used for the analyses presented here are openly available in [94].

Chapter 8

Markovian Compartmental Epidemic Models

Mathematical models of epidemic spreading offer a principled way to describe how infections propagate in a population and to reason about the effect of contact structure on the evolution of an outbreak. Classical compartmental descriptions provide a macroscopic view based on the fractions of individuals in each compartment under the assumption of homogeneous mixing. Here we adopt an individual-based perspective that keeps the network of contacts explicit and models the state of each agent at discrete times. This framework will serve as the prior dynamical model in the next chapters.

8.1 Discrete-time individual-based dynamics

We consider a population of N individuals, indexed by $i = 1, \dots, N$, observed at discrete times $n = 0, 1, \dots, T$ (e.g. daily). The epidemiological state of individual i at time n is a random variable

$$x_i^n \in \mathscr{X},$$

where $\mathscr{X}$ denotes the set of compartments relevant for the disease under study (for instance $\mathscr{X} = \{S, I\}$ in the SI and SIS models, $\mathscr{X} = \{S, I, R\}$ in SIR and SIRS. The configuration of the population at time n is $\underline{x}^n = \{x_i^n\}_{i=1}^N$, and the full trajectory of individual i over the temporal window $n = 0, \dots, T$ is $\boldsymbol{x}_i = (x_i^0, \dots, x_i^T)$. The trajectory of the whole population is $\underline{\boldsymbol{x}} = \{\boldsymbol{x}_i\}_{i=1}^N$.

Individuals interact on a contact network G that may vary with time. At each time step n we define λ_{ij}^n as the probability that an infectious individual j transmits the infection to a

susceptible individual i during step n. We set $\lambda_{ij}^n = 0$ if no contact occurs from j to i during step n. The temporal neighborhood of node i at time n is $\partial i^n = \{j : \lambda_{ij}^n > 0\}$, and its degree is $k_i^n = |\partial i^n|$.

We focus on Markovian dynamics, where the state of each individual at time $n+1$ depends only on its own state and on the states of its neighbors at time n. The transition probability for individual i is therefore

$$w_i^n\left(x_i^{n+1}|x_i^n, \underline{x}_{\partial i}^n\right),$$

where $\underline{x}_{\partial i}^n$ is the set of states of the neighbors of i at time n. The transition probability may depend on time and on the node, allowing for full heterogeneity in the population and in the contact structure. The transition probability of the whole population factorizes over nodes,

$$w^n(\underline{x}^{n+1}|\underline{x}^n) = \prod_{i=1}^N w_i^n\left(x_i^{n+1}|x_i^n, \underline{x}_{\partial i}^n\right). \tag{8.1}$$

Due to the Markov property, the probability of the configuration at time $n+1$ follows from that at time n through the master equation

$$P(\underline{x}^{n+1}) = \sum_{\underline{x}^n \in \mathscr{X}^N} W^n(\underline{x}^{n+1}|\underline{x}^n)\, P(\underline{x}^n), \tag{8.2}$$

where the initial distribution of states is assumed to factorize over nodes

$$P(\underline{x}^0) = \prod_{i=1}^N p_i^0(x_i^0). \tag{8.3}$$

It follows that the probability of a full trajectory $\underline{\boldsymbol{x}}$ is

$$P(\underline{\boldsymbol{x}}) = P(\underline{x}^0) \prod_{n=0}^{T-1} \prod_{i=1}^N w_i^n\left(x_i^{n+1}|x_i^n, \underline{x}_{\partial i}^n\right). \tag{8.4}$$

In what follows we focus on models with a single infectious state I, acting upon susceptible individuals S through the contact network. An individual i in state S (i.e. $x_i^n = S$) can become infected only if some of its neighbors are in the infectious state at time n. The infection pressure felt by i therefore increases both with the number of infected neighbors and with the magnitude of the infection couplings λ_{ij}^n along the corresponding links. It is therefore convenient to define

the total infection pressure acting on i as

$$h_i^n = \sum_{j \in \partial i^n} \log\left(1 - \lambda_{ij}^n \delta_{x_j^n, I}\right) = \sum_{j \in \partial i^n} \log(1 - \lambda_{ij}^n)\, \delta_{x_j^n, I}, \tag{8.5}$$

so that the probability that none of the infected neighbors of i transmits the infection during step n (i.e. i remains susceptible) is given by

$$e^{h_i^n} = \prod_{j \in \partial i^n} \left(1 - \lambda_{ij}^n \delta_{x_j^n, I}\right). \tag{8.6}$$

In the simplest specifications, a susceptible individual remains susceptible with probability $e^{h_i^n}$, while leaving the susceptible state with complementary probability $1 - e^{h_i^n}$. Intrinsic transitions among non-susceptible compartments are described by model-specific parameters that may depend on the node and on time.

8.2 Examples and standard parameterizations

The general form above accommodates both irreversible and recurrent processes. Irreversible models such as SI and SIR admit a partial order of compartments along each trajectory, while recurrent models such as SIS or SIRS allow returns to previously visited compartments. Below we list the single-node kernels in a compact form; in all cases $e^{h_i^n}$ is defined by Eq. (8.5). For completeness, we also include a small spontaneous infection probability ε_i^n in the transition from S to I (or E), so that the probability of remaining susceptible becomes $e^{h_i^n}(1 - \varepsilon_i^n)$. In general, self-infections can model scenarios where the contact network is only partially known, in such a way to avoid incompatibility between the inference model and observations arising from unknown transmission channels; however, a small self-infection can also be useful to heal pathological cases that can arise in the presence of noiseless observations (more details in Appendix section E.2). We introduce for shortness of notation its complement $\alpha_i^n = 1 - \varepsilon_i^n$.

SI model $\mathscr{X} = \{S, I\}$.

$$w_i^n(x_i^{n+1} = S | x_i^n, \underline{x}_{\partial i}^n) = \delta_{x_i^n, S} \alpha_i^n e^{h_i^n}, \tag{8.7}$$

$$w_i^n(x_i^{n+1} = I | x_i^n, \underline{x}_{\partial i}^n) = \delta_{x_i^n, I} + \delta_{x_i^n, S}\left(1 - \alpha_i^n e^{h_i^n}\right) = 1 - \delta_{x_i^n, S} \alpha_i^n e^{h_i^n}. \tag{8.8}$$

SIR model $\mathscr{X} = \{S, I, R\}$. Individuals recover from infection with probability r_i^n. Once recovered, they are immune and cannot be infected again.

$$w_i^n(x_i^{n+1} = S | x_i^n, \underline{x}_{\partial i}^n) = \delta_{x_i^n, S} \alpha_i^n e^{h_i^n}, \tag{8.9}$$

$$w_i^n(x_i^{n+1} = I | x_i^n, \underline{x}_{\partial i}^n) = \delta_{x_i^n, I}(1 - r_i^n) + \delta_{x_i^n, S}\left(1 - \alpha_i^n e^{h_i^n}\right), \tag{8.10}$$

$$w_i^n(x_i^{n+1} = R | x_i^n, \underline{x}_{\partial i}^n) = \delta_{x_i^n, R} + \delta_{x_i^n, I} r_i^n. \tag{8.11}$$

SIS model $\mathscr{X} = \{S, I\}$. Individuals recover from infection with probability r_i^n. There is no immunity, so that recovered individuals return to the susceptible state.

$$w_i^n(x_i^{n+1} = S | x_i^n, \underline{x}_{\partial i}^n) = \delta_{x_i^n, I} r_i^n + \delta_{x_i^n, S} \alpha_i^n e^{h_i^n}, \tag{8.12}$$

$$\begin{aligned} w_i^n(x_i^{n+1} = I | x_i^n, \underline{x}_{\partial i}^n) &= \delta_{x_i^n, I}(1 - r_i^n) + \delta_{x_i^n, S}\left(1 - \alpha_i^n e^{h_i^n}\right) \\ &= 1 - \delta_{x_i^n, I} r_i^n - \delta_{x_i^n, S} \alpha_i^n e^{h_i^n}. \end{aligned} \tag{8.13}$$

SIRS model $\mathscr{X} = \{S, I, R\}$. Individuals recover from infection with probability r_i^n. Once recovered, they acquire temporary immunity, which is lost with probability σ_i^n.

$$w_i^n(x_i^{n+1} = S | x_i^n, \underline{x}_{\partial i}^n) = \delta_{x_i^n, R} \sigma_i^n + \delta_{x_i^n, S} \alpha_i^n e^{h_i^n}, \tag{8.14}$$

$$w_i^n(x_i^{n+1} = I | x_i^n, \underline{x}_{\partial i}^n) = \delta_{x_i^n, I}(1 - r_i^n) + \delta_{x_i^n, S}\left(1 - \alpha_i^n e^{h_i^n}\right), \tag{8.15}$$

$$w_i^n(x_i^{n+1} = R | x_i^n, \underline{x}_{\partial i}^n) = \delta_{x_i^n, R}(1 - \sigma_i^n) + \delta_{x_i^n, I} r_i^n. \tag{8.16}$$

In all the models above the transition from S to I depends on the states of the neighbors through the infection pressure h_i^n, while all other transitions are intrinsic and depend only on the state of the individual. This structure is common to many compartmental models, and it will be exploited in the derivation of the inference algorithms in the next chapters.

Chapter 9

Bayesian reconstruction of epidemics

In this chapter we introduce the general Bayesian framework for the reconstruction of epidemics from partial observations. We discuss the computational challenges arising in this context, and we review existing approaches based on Monte Carlo sampling, mean-field approximations and Belief Propagation. The limitations of these methods motivate the development of the novel SCDC algorithm presented in chapter 10.

9.1 Bayesian formulation of the inference problem

We consider a population of N individuals interacting on a (possibly time-varying) contact network, and we assume that the spreading of an infection is described by a discrete-time Markovian compartmental model as in chapter 8. The probability of a full trajectory $\underline{\boldsymbol{x}}$ is given by Eq. (8.4),

$$P(\underline{\boldsymbol{x}}) = p_0(\underline{x}^0) \prod_{n=0}^{T-1} \prod_{i=1}^{N} w_i^n \left(x_i^{n+1} | x_i^n, \underline{x}_{\partial i}^n\right), \tag{9.1}$$

where p_0 is the initial distribution of states, assumed to factorize over nodes as in Eq. (8.3), and w_i^n are the single-node transition probabilities.

We assume that partial observations on the state of some individuals are available at certain times. This is the typical situation in real-world scenarios, where the state of only a fraction of the population is known at a few discrete times, due to limited testing capacity. We denote by $\mathscr{O} = \{\boldsymbol{O}_i\}_{i=1}^N$ the set of site independent observations, where $\boldsymbol{O}_i = (O_i^0, \dots, O_i^T)$ is the vector of observations for individual i. Each O_i^n can either be a compartment label in $\mathscr{X}$, if the state of individual i has been observed at time n, or a special symbol $\emptyset$ if the state is unknown. We

assume that the tests can give false negatives and false positives, respectively with rates η^- and η^+. The likelihood contribution of an observation O_i^n given the true state x_i^n is defined as:

$$P(O_i^n|x_i^n) = \begin{cases} (1-\eta^+)(1-\delta_{x_i^n,I}) + \eta^- \delta_{x_i^n,I} & \text{if } O_i^n \in \mathscr{X} \setminus \{I\}, \\ \eta^+(1-\delta_{x_i^n,I}) + (1-\eta^-)\delta_{x_i^n,I} & \text{if } O_i^n = I, \\ 1 & \text{if } O_i^n = \emptyset. \end{cases} \tag{9.2}$$

Note that for $O_i^n = \emptyset$, the factor is unity, reflecting that the absence of a test provides no information about the state x_i^n. The likelihood of the full set of observations $\mathscr{O}$ given the trajectory $\underline{\boldsymbol{x}}$ factorizes over nodes and times,

$$P(\mathscr{O}|\underline{\boldsymbol{x}}) = \prod_{i=1}^{N}\prod_{n=0}^{T} P(O_i^n|x_i^n). \tag{9.3}$$

In the following, we adopt a Bayesian formulation of the inference problem. Within this framework, the epidemic dynamics plays the role of a generative model defining a prior distribution over the space of all possible trajectories $\underline{\boldsymbol{x}}$, while the partial observations $\mathscr{O}$ provide evidence that reweights the statistical weight of each trajectory according to its compatibility with the observed data. The goal of the inference task is to estimate marginal probabilities of infection or recovery from the posterior distribution conditioned on the available observations. Using Bayes' theorem, the posterior distribution over epidemic trajectories reads

$$P(\underline{\boldsymbol{x}}|\mathscr{O}) = \frac{P(\mathscr{O}|\underline{\boldsymbol{x}})P(\underline{\boldsymbol{x}})}{P(\mathscr{O})}, \tag{9.4a}$$

$$\propto \prod_{i=1}^{N} p_i^0(x_i^0) P(O_i^T|x_i^T) \prod_{n=0}^{T-1} w_i^n\left(x_i^{n+1}|x_i^n, \underline{x}_{\partial i}^n\right) P(O_i^n|x_i^n), \tag{9.4b}$$

where $p(\mathscr{O})$ is a normalization constant independent of $\underline{\boldsymbol{x}}$. The posterior distribution Eq. (9.4b) defines a high-dimensional graphical model with loops, whose structure is determined by the contact network and by the temporal sequence of observations. The main computational challenge in this context is to efficiently estimate marginal probabilities of the form

$$P(x_i^n = x|\mathscr{O}) = \sum_{\underline{\boldsymbol{x}}\setminus x_i^n} P(\underline{\boldsymbol{x}}\setminus x_i^n, x_i^n = x|\mathscr{O}), \tag{9.5}$$

where the sum runs over all trajectories except for x_i^n. Exact computation of Eq. (9.5) is intractable in general, due to the exponential growth of the configuration space with N and

T. In the next section we review existing approaches to approximate the posterior marginals Eq. (9.5), highlighting their strengths and limitations.

9.2 Simple Mean-Field approximation

A first class of methods to approximate the posterior marginals Eq. (9.5) relies on mean-field approximations, which replace the true posterior distribution with a simpler one factorized over nodes and times. From Eq. (8.2), the marginal probability that individual i is in state x at time $n+1$ can be written as

$$P(x_i^{n+1} = x) = \sum_{\underline{x}^n} w_i^n(x|x_i^n, \underline{x}_{\partial i}^n)\, P(\underline{x}^n). \tag{9.6}$$

The simplest mean-field approximation consists in replacing the full joint distribution $P(\underline{x}^n)$ with the product of its one-site marginals, $P(\underline{x}^n) \approx \prod_{i=1}^N p_{\mathrm{MF}}(x_i^n)$. This standard assumption in statistical physics neglects all correlations between neighboring nodes, leading to a closed set of recursion equations

$$p_{\mathrm{MF}}(x_i^{n+1} = x) = \sum_{x_i^n} p_{\mathrm{MF}}(x_i^n) \sum_{\underline{x}_{\partial i}^n} w_i^n(x|x_i^n, \underline{x}_{\partial i}^n) \prod_{k \in \partial i^n} p_{\mathrm{MF}}(x_k^n). \tag{9.7}$$

The Simple Mean-Field (SMF) equations Eq. (9.7) can be iterated forward in time from the initial condition $p_{\mathrm{MF}}(x_i^0)$ to obtain an approximate evolution of the marginal probabilities. Their intuitive content mirrors the transition equations of the stochastic process, yet the neglect of correlations among neighbors leads to systematic errors, typically overestimating infection levels. Improved approaches, such as the Dynamic Message Passing (DMP) equations [95, 96], partially restore these correlations and are exact on tree-like networks. We will discuss them in more detail in section 9.3.

As an illustration, consider the SI model, for which $\mathscr{X} = \{S, I\}$ and infection is transmitted from individual k to i with probability λ_{ik}^n at time n. The SMF equations read

$$p_{\mathrm{MF}}(x_i^{n+1} = S) = p_{\mathrm{MF}}(x_i^n = S)\, \alpha_i^n \prod_{k \in \partial i^n} \left(1 - \lambda_{ik}^n p_{\mathrm{MF}}(x_k^n = I)\right), \tag{9.8}$$

$$p_{\mathrm{MF}}(x_i^{n+1} = I) = 1 - p_{\mathrm{MF}}(x_i^n = S)\, \alpha_i^n \prod_{k \in \partial i^n} \left(1 - \lambda_{ik}^n p_{\mathrm{MF}}(x_k^n = I)\right). \tag{9.9}$$

For weak infection probabilities $\lambda_{ij}^n \ll 1$, one may expand the product in Eq. (9.8) to first order. Defining the infection probability $\mu_i^n = p_{\mathrm{MF}}(x_i^n = I)$, the small-coupling approximation

becomes

$$\mu_i^{n+1} \approx 1 - (1 - \mu_i^n)\, \alpha_i^n \exp\left(- \sum_{k \in \partial i^n} \lambda_{ik}^n \mu_k^n \right). \tag{9.10}$$

This expression allows each individual to update its infection probability locally, using only information from its own state and from the infection probabilities of its contacts weighted by the couplings λ_{ik}^n. Such an update can be implemented in a distributed manner, without requiring a central repository of information.

The SMF formulation in Eq. (9.7) does not account for the observations $\mathscr{O}$. To incorporate the information provided by test results, one can reweight the probabilities at the observation times according to the measured states. For simplicity, we consider here the case of *hard observations*, assuming ideal tests with vanishing false-negative and false-positive rates ($\eta^- = \eta^+ = 0$). Suppose that individual i is observed at time n_{obs} with outcome $O_i^{n_{\mathrm{obs}}}$. The SMF equations Eq. (9.7) are run forward in time from $n - n_{\mathrm{MF}}$, assuming all individuals are initially susceptible, and the observation is imposed as a constraint at n_{obs}:

$$\text{if } O_i^{n_{\mathrm{obs}}} = S: \quad p_{\mathrm{MF}}(x_i^{n'} = S) = 1, \qquad n' \in [n - n_{\mathrm{MF}}, n_{\mathrm{obs}}], \tag{9.11}$$

$$\text{if } O_i^{n_{\mathrm{obs}}} = I: \quad p_{\mathrm{MF}}(x_i^{n'} = I) = 1, \qquad n' \in [n_{\mathrm{obs}} - \Delta n_I, n_{\mathrm{obs}}]. \tag{9.12}$$

Here n_{MF} represents the number of time steps prior to the observation considered in the forward update, and Δn_I the typical delay between infection and testing. This heuristic procedure is straightforward and efficient, but it applies only to irreversible dynamics, such as SI and SIR, where infection times are ordered and uniquely determined. In models with recurrent dynamics, such as SIS and SIRS, backward propagation of observational information becomes intrinsically ambiguous, as multiple infection and recovery events can be compatible with a single observation.

9.3 Belief Propagation for epidemic trajectories

The posterior distribution Eq. (9.4b) defines a high-dimensional graphical model whose topology mirrors the contact network and the temporal sequence of observations. This structure closely resembles that of the trajectory distribution Eq. (2.20), with the difference that each factor is now reweighted by the observation likelihood Eq. (9.2). On a tree-like contact network, the posterior distribution Eq. (9.4b) can in principle be treated exactly using the dynamic cavity method–also known as dynamic Belief Propagation (BP) [27, 97]. However, as discussed in section 2.3, the explicit time dependence of the dynamics introduces short loops

between consecutive temporal layers, so the standard factor graph representation is no longer loop-free. To apply the dynamic cavity method consistently, one must instead adopt the dual graph representation, in which variable nodes encode pairs of trajectories on the edges of the contact network, thus eliminating these redundant temporal loops. When the underlying contact network is sparse and locally tree-like, the resulting dual factor graph is also locally tree-like, allowing BP to compute the posterior marginals Eq. (9.5) exactly.

Applying the BP update rule Eq. (2.21) to the posterior factors in Eq. (9.4b) yields the following recursion for the messages:

$$c_{i\setminus j}(\boldsymbol{x}_i,\boldsymbol{x}_j) \propto p_0(x_i^0)P(\boldsymbol{O}_i|\boldsymbol{x}_i) \sum_{\underline{\boldsymbol{x}}_{\partial i\setminus j}} \prod_{n=0}^{T-1} w_i^n(x_i^{n+1}|x_i^n,x_j^n,\underline{x}_{\partial i\setminus j}^n) \prod_{k\in\partial i\setminus j} c_{k\setminus i}(\boldsymbol{x}_k,\boldsymbol{x}_i), \tag{9.13}$$

where $P(\boldsymbol{O}_i|\boldsymbol{x}_i) = \prod_{n=0}^{T} P(O_i^n|x_i^n)$ is the likelihood of the observations for node i given its trajectory $\boldsymbol{x}_i$.

Despite their conceptual simplicity, the direct use of Eq. (9.13) is computationally prohibitive. The sum over neighboring trajectories grows exponentially with T, as each trajectory $\boldsymbol{x}_k$ spans $|\mathscr{X}|^T$ possible sequences. For irreversible processes such as SI or SIR, however, the trajectory of each node can be fully characterized by a small number of transition times, which reduces the complexity dramatically. In these cases, the BP recursion can be reformulated in terms of messages over transition times, leading to an algorithm that scales linearly in T. In contrast, for recurrent models (e.g., SIS or SIRS), multiple infection and recovery events can occur within the same trajectory, and the number of possible paths still grows exponentially with time, rendering exact BP intractable.

To illustrate this approach, let us focus again on the SI model, where each individual i can switch from the susceptible to the infected state only once. The trajectory $\boldsymbol{x}_i$ is therefore uniquely determined by the *infection time* n_i, defined as

$$n_i = \min\{n \mid x_i^n = I\}.$$

The infection time n_i can take $T+2$ values, corresponding to infection at times $0,1,\dots,T$, or remaining susceptible throughout the observation window ($n_i = T+1$). The state of individual i at time n follows directly from n_i:

$$x_i^n = \begin{cases} S, & n < n_i, \\ I, & n \geq n_i. \end{cases} \tag{9.14}$$

In the SI process, an infected node j transmits the infection to a neighbor i after a random *transmission delay* s_{ij}, drawn from a geometric distribution $\omega_{ij}(s)$ with parameter λ_{ij},

$$\omega_{ij}(s) = \lambda_{ij}(1-\lambda_{ij})^s, \qquad s = 0,1,2,\dots \tag{9.15}$$

Hence, the infection time of node i is determined by the minimum infection time of the neighbors plus their corresponding transmission delays:

$$n_i = 1 + \min_{k\in\partial i}\{n_k + s_{ik}\}. \tag{9.16}$$

The joint distribution of infection times $\underline{n} = \{n_i\}$ given the initial condition $\underline{x}^0$ can then be written as a sum over all the possible transmission delays $\underline{s} = \{s_{ij}\}_{i\neq j}$

$$P(\underline{n}|\underline{x}^0) = \sum_{\underline{s}} P(\underline{s}) \prod_{i=1}^{N} \delta\left(n_i, (1-\delta_{x_i^0,I})\left[1+\min_{k\in\partial i}\{n_k+s_{ik}\}\right]\right) \tag{9.17a}$$

$$= \sum_{\underline{s}} \prod_{i\neq j} \omega_{ij}(s_{ij}) \prod_{i=1}^{N} \delta\left(n_i, (1-\delta_{x_i^0,I})\left[1+\min_{k\in\partial i}\{n_k+s_{ik}\}\right]\right), \tag{9.17b}$$

where the prefactor $(1-\delta_{x_i^0,I})$ ensures that initially infected individuals have $n_i = 0$. For shortness of notation, we denote by ψ_i the local factor enforcing Eq. (9.16):

$$\psi_i(n_i, \{n_k, s_{ik}\}_{k\in\partial i}) = \delta\left(n_i, (1-\delta_{x_i^0,I})\left[1+\min_{k\in\partial i}\{n_k+s_{ik}\}\right]\right), \tag{9.18}$$

and by ψ_{ij} the pairwise factor associated with the transmission delay s_{ij}:

$$\psi_{ij}(s_{ij}, s_{ji}) = \omega_{ij}(s_{ij})\omega_{ji}(s_{ji}). \tag{9.19}$$

Conditioning on the available observations $\mathscr{O}$, the posterior distribution of infection times reads

$$P(\underline{n}|\mathscr{O}) \propto \sum_{\underline{s}} \prod_{(i,j)\in E} \psi_{ij}(s_{ij}, s_{ji}) \prod_{i=1}^{N} \psi_i(n_i, \{n_k, s_{ik}\}_{k\in\partial i}) \prod_{(i,n,O_i^n)\in\mathscr{O}} P(O_i^n|n_i), \tag{9.20}$$

where $P(O_i^n|n_i)$ is the likelihood of observing O_i^n at time n given the infection time n_i, obtained from Eq. (9.2):

$$P(O_i^n|n_i) = \begin{cases} (1-\eta^+)\,\Theta(n_i-n) + \eta^-\Theta(n-n_i), & O_i^n = S, \\ \eta^+\Theta(n_i-n) + (1-\eta^-)\,\Theta(n-n_i), & O_i^n = I. \end{cases} \tag{9.21}$$

As for the trajectory-based posterior, the infection time formulation Eq. (9.20) defines a graphical model with many short loops, whose structure follows the contact network and observation times. By adopting the dual factor graph representation of Fig. 2.4(d), one can again eliminate these loops and apply BP to compute the posterior marginals $P(n_i|\mathscr{O})$ exactly on tree-like networks, or approximately on sparse loopy ones. We get rid of the delays s_{ij} defining $n_{ij} = n_j + s_{ij}$ and $n_{ji} = n_i + s_{ji}$ as the transmission times from j to i and from i to j, respectively. The posterior distribution can thus be factorized as

$$P(\underline{n}|\mathscr{O}) \propto \sum_{\underline{n}_{ij}} \prod_{i=1}^{N} \psi_i(n_i, \{n_{ik}\}_{k\in\partial i}) \prod_{(i,j)\in E} \phi_{ij}(n_i, n_j, n_{ij}, n_{ji}), \tag{9.22}$$

where the sum runs over all transmission times $\underline{n}_{ij} = \{n_{ij}\}_{(i,j)\in E}$. This defines a dual factor graph that is locally tree-like if the contact network is sparse and tree-like. Applying BP to this representation yields an efficient algorithm for computing the posterior marginals $P(n_i|\mathscr{O})$ exactly on trees and approximately on sparse networks. Each BP update scales as $O(T|E|)$, where $|E|$ is the number of edges in the contact network [98].

Despite its accuracy, the BP approach presents important practical limitations. First, its computational complexity, although linear in T for irreversible dynamics, becomes prohibitive for recurrent models such as SIS or SIRS, where each node can undergo multiple infection and recovery events. In these cases, the number of possible trajectories grows exponentially with time, and BP can only be applied through uncontrolled approximations. These limitations motivate the development of alternative methods that retain the interpretability and Bayesian consistency of BP while ensuring tractable computational complexity for both recurrent and non-recurrent epidemic models. The SCDC method, introduced in chapter 10, represents a concrete realization of this idea.

Chapter 10

The Small-Coupling Dynamic Cavity Method

This chapter introduces a dynamic cavity approximation to overcome the high computational cost of Belief Propagation in recurrent dynamic models. The proposed method uses a small-coupling expansion to derive a closed system of equations for single-time marginals, achieving a uniform computational complexity that is linear in the time horizon. We derive the method using the SI model and benchmark its performance against Belief Propagation and Monte Carlo sampling. Finally, we explore its application to inference in recurrent models and discuss practical implications. All numerical results presented in this chapter were obtained using the open source `SmallCouplingDynamicCavity.jl` package [94], which implements the SCDC algorithm and the transfer matrix solvers described below.

10.1 Dynamic cavity for the SI model

In order to illustrate the derivation of the SCDC method, we focus on the SI model. As recalled in section 9.3, the observation reweighted dynamic cavity equations read

$$\tilde{c}_{i\backslash j}(\boldsymbol{x}_i,\boldsymbol{x}_j) = \frac{p_i^0(x_i^0)P(\boldsymbol{O}_i|\boldsymbol{x}_i)}{\tilde{Z}_{i\backslash j}} \sum_{\underline{\boldsymbol{x}}_{\partial i\backslash j}} \prod_{n=0}^{T-1} w_i^n\left(x_i^{n+1}|x_i^n,x_j^n,\underline{x}_{\partial i\backslash j}^n\right) \prod_{k\in\partial i\backslash j} \tilde{c}_{k\backslash i}(\boldsymbol{x}_k,\boldsymbol{x}_i), \tag{10.1}$$

where the tilde reminds that we will soon introduce an equivalent trajectory-field representation. The main challenge in solving these equations is the exponential growth of the number of

possible trajectories with the time horizon T. To obtain a tractable approximation, we expand in the edge couplings, assuming small infection probabilities.

Introducing the couplings $J_{ij}^n = \log(1-\lambda_{ij}^n)$, the product of single-step transitions can be written as

$$e^{\tilde{S}_{i\setminus j}} = \prod_{n=0}^{T-1} w_i^n\big(x_i^{n+1}|x_i^n, x_j^n, \underline{x}_{\partial i\setminus j}^n\big) \tag{10.2a}$$

$$= \prod_{n=0}^{T-1}\Big[\alpha_i^n \exp\Big(s_i^n + \sum_{k\in\partial i\setminus j} J_{ik}^n \delta_{x_k^n,I}\Big)\big(\delta_{x_i^{n+1},x_i^n} - \delta_{x_i^{n+1},I}\big) + \delta_{x_i^{n+1},I}\Big], \tag{10.2b}$$

where we introduced the auxiliary field $s_i^n = J_{ij}^n \delta_{x_j^n,I}$ that isolates the contribution of neighbor j along the edge $j\to i$. The dynamic cavity equation can thus be rewritten in terms of the pair $(\boldsymbol{x}_i, \boldsymbol{s}_i)$ as

$$\tilde{c}_{i\setminus j}(\boldsymbol{x}_i, \boldsymbol{s}_i) = \frac{p_i^0(x_i^0)P(\boldsymbol{O}_i|\boldsymbol{x}_i)}{\tilde{Z}_{i\setminus j}} \sum_{\underline{\boldsymbol{x}}_{\partial i\setminus j}} e^{\tilde{S}_{i\setminus j}} \prod_{k\in\partial i\setminus j} \tilde{c}_{k\setminus i}\big(\boldsymbol{x}_k, \boldsymbol{J}_{ki}\cdot\boldsymbol{x}_i\big), \tag{10.3}$$

with the shorthand $\boldsymbol{J}_{ki}\cdot\boldsymbol{x}_i = (J_{ki}^0\delta_{x_i^0,I},\dots,J_{ki}^T\delta_{x_i^T,I})$. The normalization constant $\tilde{Z}_{i\setminus j}$ can be neglected, as one can always re-normalize the messages at the end of the iterations.

To perform the small-coupling expansion it is convenient to pass to Fourier space with respect to $\boldsymbol{s}_i$. We define the Fourier transform of the cavity messages as

$$c(\boldsymbol{x}, \hat{\boldsymbol{h}}) = \int \left(\prod_{n=0}^{T} \frac{\mathrm{d}s^n}{2\pi} e^{\mathrm{i}\hat{h}^n s^n}\right) \tilde{c}(\boldsymbol{x}, \boldsymbol{s}), \tag{10.4}$$

where $\hat{\boldsymbol{h}} = (\hat{h}^0,\dots,\hat{h}^T)$ are the conjugate variables to the fields $\boldsymbol{s}$. Applying the inverse transform to Eq. (10.3) and recalling the Fourier representation of the delta function, we obtain

$$c_{i\setminus j}(\boldsymbol{x}_i, \hat{\boldsymbol{h}}_i) = \frac{p_i^0(x_i^0)P(\boldsymbol{O}_i|\boldsymbol{x}_i)}{Z_{i\setminus j}} e^{S_i^{\mathrm{loc}}} \prod_{k\in\partial i\setminus j} \sum_{\boldsymbol{x}_k} \int \mathrm{d}\hat{\boldsymbol{h}}_k\, c_{k\setminus i}(\boldsymbol{x}_k, \hat{\boldsymbol{h}}_k) e^{S_{ik}^{\mathrm{int}}}, \tag{10.5}$$

where the local and pair actions are

$$S_i^{\mathrm{loc}} = \sum_{n=0}^{T-1} \log\Big[\delta(\hat{h}_i^n + \mathrm{i})\alpha_i^n\big(\delta_{x_i^{n+1},x_i^n} - \delta_{x_i^{n+1},I}\big) + \delta(\hat{h}_i^n)\delta_{x_i^{n+1},I}\Big], \tag{10.6}$$

$$S_{ik}^{\mathrm{int}} = \mathrm{i}\sum_{n=0}^{T-1}\Big(\delta_{x_i^n,I} J_{ki}^n \hat{h}_k^n + \delta_{x_k^n,I} J_{ik}^n \hat{h}_i^n\Big). \tag{10.7}$$

The Fourier transformed equations mirror the structure of the continuous-time dynamic cavity, with the response fields replaced here by the conjugate variables $\hat{\boldsymbol{h}}$. This analogy allows us to apply the same small-coupling expansion strategy used in section 4.2 to obtain a closed system for single-time marginals. A complementary derivation based on a path integral representation is reported in Appendix E.1.

An equivalent equation is obtained for the single-node marginals

$$c_i(\boldsymbol{x}_i, \hat{\boldsymbol{h}}_i) = \frac{p_i^0(x_i^0)\, P(\boldsymbol{O}_i|\boldsymbol{x}_i)}{Z_i} e^{S_i^{\mathrm{loc}}} \prod_{k\in\partial i} \sum_{\boldsymbol{x}_k} \int \mathrm{d}\hat{\boldsymbol{h}}_k\, c_{k\setminus i}(\boldsymbol{x}_k, \hat{\boldsymbol{h}}_k) e^{S_{ik}^{\mathrm{int}}}. \tag{10.8}$$

10.2 Small-Coupling Dynamic Cavity Equations

Following the same steps as in section 4.2, we expand the interaction term to first order in the coupling strengths J_{ij}^n. Before proceeding, we fix the normalization constants $Z_{i\setminus j}$ and Z_i in Eqs. (10.5) and (10.8) by imposing

$$\sum_{\boldsymbol{x}_i} \int \mathrm{d}\hat{\boldsymbol{h}}_i\, c_{i\setminus j}(\boldsymbol{x}_i, \hat{\boldsymbol{h}}_i) = 1, \qquad \sum_{\boldsymbol{x}_i} \int \mathrm{d}\hat{\boldsymbol{h}}_i\, c_i(\boldsymbol{x}_i, \hat{\boldsymbol{h}}_i) = 1. \tag{10.9}$$

This choice ensures that messages and marginals are properly normalized quasi-probability distributions on the joint space of trajectories and conjugate variables. Consequently, for any function $f(\boldsymbol{x}_i, \hat{\boldsymbol{h}}_i)$ we set

$$\langle f(\boldsymbol{x}_i, \hat{\boldsymbol{h}}_i)\rangle_{i\setminus j} = \sum_{\boldsymbol{x}_i} \int \mathrm{d}\hat{\boldsymbol{h}}_i\, c_{i\setminus j}(\boldsymbol{x}_i, \hat{\boldsymbol{h}}_i)\, f(\boldsymbol{x}_i, \hat{\boldsymbol{h}}_i), \tag{10.10}$$

with an analogous definition for $\langle\cdot\rangle_i$. By inverse Fourier transforming the normalization, one finds that this is equivalent to fixing $\boldsymbol{s}_i = \boldsymbol{0}$ in Eq. (10.3), namely

$$\sum_{\boldsymbol{x}_i} \int \mathrm{d}\hat{\boldsymbol{h}}_i\, c_{i\setminus j}(\boldsymbol{x}_i, \hat{\boldsymbol{h}}_i) = \sum_{\boldsymbol{x}_i} \tilde{c}_{i\setminus j}(\boldsymbol{x}_i, \boldsymbol{0}) = 1, \tag{10.11}$$

which corresponds to treating the missing neighbor j as susceptible at all times in the cavity graph.

Within the small-coupling approximation, removing node j from the cavity graph implies that j exerts no infection pressure on i. In contrast, the original dynamic cavity method of section 9.3 conditions on the actual trajectory of j and updates edge messages in a self-consistent manner, which can become exponentially costly in the time horizon T for recurrent

dynamics. Choosing a normalization independent of the trajectory of j is convenient for developing an approximation that does not explicitly track the joint trajectories of i and j. A drawback is that certain observed trajectories of i can only be reconciled with the model by introducing a nonzero self-infection probability. Appendix E.2 discusses this issue on a simple example and shows how a small self-infection resolves the inconsistency.

Expanding to first order in J_{ij}^n the cavity equations Eq. (10.5) gives

$$c_{i\setminus j}(\boldsymbol{x}_i, \hat{\boldsymbol{h}}_i) \approx \frac{p_i^0(x_i^0)\, P(\boldsymbol{O}_i|\boldsymbol{x}_i)}{Z_{i\setminus j}} e^{S_i^{\mathrm{loc}}} \prod_{k\in\partial i\setminus j} \left[1 + \sum_{n=0}^{T-1} \delta_{x_i^n, I}\, J_{ki}^n \langle \mathrm{i}\hat{h}_k^n \rangle_{k\setminus i} + \mathrm{i} \sum_{n=0}^{T-1} J_{ik}^n \langle \delta_{x_k^n, I} \rangle_{k\setminus i} \hat{h}_i^n \right]. \tag{10.12}$$

Introduce the cavity infection probability and the average conjugate field

$$\mu_{i\setminus j}^n = \langle \delta_{x_i^n, I} \rangle_{i\setminus j} = \sum_{\boldsymbol{x}_i} \tilde{c}_{i\setminus j}(\boldsymbol{x}_i, \boldsymbol{0})\, \delta_{x_i^n, I}, \tag{10.13}$$

$$\hat{\mu}_{i\setminus j}^n = \langle \mathrm{i}\hat{h}_i^n \rangle_{i\setminus j} = \sum_{\boldsymbol{x}_i} \left. \frac{\partial \tilde{c}_{i\setminus j}(\boldsymbol{x}_i, \boldsymbol{s}_i)}{\partial s_i^n} \right|_{\boldsymbol{s}_i = \boldsymbol{0}}, \tag{10.14}$$

where we used the inverse transform to express the averages in terms of $\tilde{c}_{i\setminus j}$. The quantity $\mu_{i\setminus j}^n$ is the probability that i is infected at time n in the cavity where j is fixed susceptible. The quantity $\hat{\mu}_{i\setminus j}^n$ measures the sensitivity of the trajectory probability of i to an infinitesimal increase of the field s_i^n, which raises the infection pressure from j. It can be positive or negative since it is a response function. Direct calculations show that $\hat{\mu}$ vanishes in the absence of observations, while it becomes nonzero when observations are present. We will see in section 10.2.2 that $\hat{\mu}$ propagates information backward in time and is essential for inference.

Re exponentiating Eq. (10.12) yields a compact form for the cavity message in terms of a decoupled effective action

$$c_{i\setminus j}(\boldsymbol{x}_i, \hat{\boldsymbol{h}}_i) \approx \frac{p_i^0(x_i^0)\, P(\boldsymbol{O}_i|\boldsymbol{x}_i)}{Z_{i\setminus j}} \exp\left(S_{i\setminus j}^{\mathrm{eff}}(\boldsymbol{x}_i, \hat{\boldsymbol{h}}_i) \right), \tag{10.15}$$

$$S_{i\setminus j}^{\mathrm{eff}}(\boldsymbol{x}_i, \hat{\boldsymbol{h}}_i) = \sum_{n=0}^{T-1} \left\{ \log\left[\delta(\hat{h}_i^n + \mathrm{i})\alpha_i^n \left(\delta_{x_i^{n+1}, x_i^n} - \delta_{x_i^{n+1}, I} \right) + \delta(\hat{h}_i^n) \delta_{x_i^{n+1}, I} \right] + \sum_{k\in\partial i\setminus j} \left(\delta_{x_i^n, I}\, J_{ki}^n\, \hat{\mu}_{k\setminus i}^n + \mathrm{i}\, \hat{h}_i^n\, J_{ik}^n\, \mu_{k\setminus i}^n \right) \right\}. \tag{10.16}$$

This effective action describes a single node i under an effective time-dependent infection pressure from its neighbors in the cavity graph.

The SCDC closure assumes that cavity messages are parametrized by the one-time quantities $\mu^n_{i\backslash j}$ and $\hat{\mu}^n_{i\backslash j}$. Evaluating their expectations with the effective action $S^{\text{eff}}_{i\backslash j}$ gives

$$\mu^{n+1}_{i\backslash j} = \langle \delta_{x^{n+1}_i, I} \rangle_{i\backslash j} = \sum_{\boldsymbol{x}_i} \delta_{x^{n+1}_i, I} \frac{p^0_i(x^0_i)\, P(\boldsymbol{O}_i|\boldsymbol{x}_i)}{\mathscr{Z}_{i\backslash j}} \exp\left(S^{\mu}_{i\backslash j}(\boldsymbol{x}_i)\right), \tag{10.17}$$

$$\hat{\mu}^{n+1}_{i\backslash j} = \langle \mathrm{i}\, \hat{h}^{n+1}_i \rangle_{i\backslash j} = \sum_{\boldsymbol{x}_i} \frac{p^0_i(x^0_i)\, P(\boldsymbol{O}_i|\boldsymbol{x}_i)}{\mathscr{Z}_{i\backslash j}} \exp\left(S^{\hat{\mu}}_{i\backslash j}(\boldsymbol{x}_i)\right), \tag{10.18}$$

with

$$S^{\mu}_{i\backslash j}(\boldsymbol{x}_i) = \sum_{n=0}^{T-1} \Bigg\{ \log\left[\alpha^n_i e^{\sum_{k\in\partial i\backslash j} J^n_{ik}\mu^n_{k\backslash i}(\delta_{x^{n+1}_i, x^n_i} - \delta_{x^{n+1}_i, I}) + \delta_{x^{n+1}_i, I}}\right] + \sum_{k\in\partial i\backslash j} \delta_{x^n_i, I} J^n_{ki} \hat{\mu}^n_{k\backslash i} \Bigg\}, \tag{10.19}$$

$$S^{\hat{\mu}}_{i\backslash j}(\boldsymbol{x}_i) = \sum_{n=0}^{T-1} \Bigg\{ \log\left[\alpha^n_i (\delta_{x^{n+1}_i, x^n_i} - \delta_{x^{n+1}_i, I})\right] + \sum_{k\in\partial i\backslash j} \left(J^n_{ik}\mu^n_{k\backslash i} + \delta_{x^n_i, I} J^n_{ki} \hat{\mu}^n_{k\backslash i}\right) \Bigg\}. \tag{10.20}$$

The constant $\mathscr{Z}_{i\backslash j}$ enforces normalization. Single-node marginals follow analogously:

$$\mu^{n+1}_i = \langle \delta_{x^{n+1}_i, I} \rangle_i = \sum_{\boldsymbol{x}_i} \delta_{x^{n+1}_i, I} \frac{p^0_i(x^0_i)\, P(\boldsymbol{O}_i|\boldsymbol{x}_i)}{\mathscr{Z}_i} \exp\left(S^{\mu}_i(\boldsymbol{x}_i)\right), \tag{10.21}$$

$$\hat{\mu}^{n+1}_i = \langle -\mathrm{i}\, \hat{h}^{n+1}_i \rangle_i = \sum_{\boldsymbol{x}_i} \frac{p^0_i(x^0_i)\, P(\boldsymbol{O}_i|\boldsymbol{x}_i)}{\mathscr{Z}_i} \exp\left(S^{\hat{\mu}}_i(\boldsymbol{x}_i)\right), \tag{10.22}$$

with

$$S^{\mu}_i(\boldsymbol{x}_i) = \sum_{n=0}^{T-1} \Bigg\{ \log\left[\alpha^n_i e^{\sum_{k\in\partial i} J^n_{ik}\mu^n_{k\backslash i}(\delta_{x^{n+1}_i, x^n_i} - \delta_{x^{n+1}_i, I}) + \delta_{x^{n+1}_i, I}}\right] + \sum_{k\in\partial i} \delta_{x^n_i, I} J^n_{ki} \hat{\mu}^n_{k\backslash i} \Bigg\}, \tag{10.23}$$

$$S^{\hat{\mu}}_i(\boldsymbol{x}_i) = \sum_{n=0}^{T-1} \Bigg\{ \log\left[\alpha^n_i (\delta_{x^{n+1}_i, x^n_i} - \delta_{x^{n+1}_i, I})\right] + \sum_{k\in\partial i} \left(J^n_{ik}\mu^n_{k\backslash i} + \delta_{x^n_i, I} J^n_{ki} \hat{\mu}^n_{k\backslash i}\right) \Bigg\}. \tag{10.24}$$

The SCDC equations Eqs. (10.17), (10.18), (10.21) and (10.22) form a closed system of $4NT$ equations for the $4NT$ unknowns $\mu^n_{i\backslash j}$, $\hat{\mu}^n_{i\backslash j}$, μ^n_i and $\hat{\mu}^n_i$. They provide a non-causal

generalization of the mean-field equations of section 9.2, where the additional order parameter $\hat{\mu}$ captures the effect of observations propagating backward in time. A computational bottleneck is the explicit sum over trajectories of node i, which involves 2^T terms in the SI model. A naive full update would thus cost $O(|E|\,T\,2^T)$. In the next section we reduce this to $O(|E|\,T)$ by using a transfer matrix formulation.

10.2.1 Efficient implementation via transfer matrices

The sums in Eqs. (10.17), (10.18), (10.21) and (10.22) can be carried out efficiently with a transfer matrix approach. The normalization constant $\mathscr{Z}_{i\backslash j}$ can be written as

$$\mathscr{Z}_{i\backslash j} = \sum_{\boldsymbol{x}_i} p_i^0(x_i^0) \prod_{n=0}^{T-1} M^{i\backslash j,n}_{x_i^n,x_i^{n+1}}\, P(O_i^T|x_i^T) \tag{10.25a}$$

$$= \sum_{x_i^0} p_i^0(x_i^0) \prod_{n=0}^{T-1} \sum_{x_i^{n+1}} M^{i\backslash j,n}_{x_i^n,x_i^{n+1}} P(O_i^T|x_i^T), \tag{10.25b}$$

where

$$M^{i\backslash j,n} = \begin{pmatrix} M^{i\backslash j,n}_{SS} & M^{i\backslash j,n}_{SI} \\ M^{i\backslash j,n}_{IS} & M^{i\backslash j,n}_{II} \end{pmatrix} \tag{10.26a}$$

$$= \begin{pmatrix} \alpha_i^n e^{\sum_{k\in\partial i\backslash j} J_{ik}^n \mu_{k\backslash i}^n} P(O_i^n|S) & \left(1-\alpha_i^n e^{\sum_{k\in\partial i\backslash j} J_{ik}^n \mu_{k\backslash i}^n}\right) P(O_i^n|S) \\ 0 & e^{\sum_{k\in\partial i\backslash j} J_{ki}^n \hat{\mu}_{k\backslash i}^n} P(O_i^n|I) \end{pmatrix}. \tag{10.26b}$$

Here $M^{i\backslash j,n}_{xy}$ is the contribution to the path weight when $x_i^n = x$ and $x_i^{n+1} = y$. The normalization can be computed in $O(T)$ operations by iterating forward. The same matrices yield

$$\mu_{i\backslash j}^n = \frac{1}{\mathscr{Z}_{i\backslash j}} \rho^{i\backslash j}_{\to n}(I)\, \rho^{i\backslash j}_{n\leftarrow}(I), \tag{10.27}$$

$$\hat{\mu}_{i\backslash j}^n = \frac{1}{\mathscr{Z}_{i\backslash j}} \rho^{i\backslash j}_{\to n}(S)\, M^{i\backslash j,n}_{SS}\, [\rho^{i\backslash j}_{n+1\leftarrow}(S) - \rho^{i\backslash j}_{n+1\leftarrow}(I)], \tag{10.28}$$

with forward and backward messages

$$\rho^{i\backslash j}_{\to n}(x_i^n) = \sum_{x_i^{n-1}} \rho^{i\backslash j}_{\to n-1}(x_i^{n-1})\, M^{i\backslash j,n-1}_{x_i^{n-1},x_i^n}, \tag{10.29}$$

$$\rho^{i\backslash j}_{n\leftarrow}(x_i^n) = \sum_{x_i^{n+1}} M^{i\backslash j,n}_{x_i^n,x_i^{n+1}}\, \rho^{i\backslash j}_{n+1\leftarrow}(x_i^{n+1}), \tag{10.30}$$

initialized with $\rho^{i\backslash j}_{\to 0}(x_i^0) = p_i^0(x_i^0)$ and $\rho^{i\backslash j}_{T\leftarrow}(x_i^T) = P(O_i^T|x_i^T)$. It is worth noting that the normalization constant can be expressed as

$$\mathscr{Z}_{i\backslash j} = \sum_{x_i^n} \rho^{i\backslash j}_{\to n}(x_i^n)\, \rho^{i\backslash j}_{n\leftarrow}(x_i^n), \tag{10.31}$$

for any time n. An equivalent transfer matrix formulation holds for the single-node marginals, where the sum runs over all neighbors ∂i.

The messages $\rho^{i\backslash j}_{\to n}$ summarize the contribution of the past up to time n, while $\rho^{i\backslash j}_{n\leftarrow}$ summarize the contribution of the future from n to T. Hence a complete update of $\mu^n_{i\backslash j}$ and $\hat{\mu}^n_{i\backslash j}$ over all edges and times costs $O(|E|T)$. From Eq. (10.28) one reads that $\hat{\mu}^n_{i\backslash j}$ is zero whenever the backward in time cavity messages at time $n+1$ are uniform. As shown in sections 10.2.2 and 10.3.1, this occurs when there are no observations at later times, yielding a purely forward reduction of the SCDC equations.

Generalization to other models The SCDC equations can be straightforwardly adapted to other recurrent models. The main changes concern the single-step transition probabilities, which modify the transfer matrices $M^{i\backslash j,n}$. For instance, in the SIS model one has

$$M^{i\backslash j,n} = \begin{pmatrix} \alpha_i^n e^{\sum_{k\in\partial i\backslash j} J^n_{ik}\mu^n_{k\backslash i}} P(O_i^n|S) & \left(1-\alpha_i^n e^{\sum_{k\in\partial i\backslash j} J^n_{ik}\mu^n_{k\backslash i}}\right) P(O_i^n|S) \\ r_i^n e^{\sum_{k\in\partial i\backslash j} J^n_{ki}\hat{\mu}^n_{k\backslash i}} P(O_i^n|I) & (1-r_i^n)\, e^{\sum_{k\in\partial i\backslash j} J^n_{ki}\hat{\mu}^n_{k\backslash i}} P(O_i^n|I) \end{pmatrix}. \tag{10.32}$$

For the SIR and the SIRS models, the state space of each node is larger, and the transfer matrices become 3×3. The computational cost remains $O(|E|T)$, as it depends on the number of edges and time steps, but not on the size of the state space. The transfer matrices for the SIR and SIRS models are, respectively,

$$M^{i\backslash j,n} = \begin{pmatrix} M^{i\backslash j,n}_{SS} & M^{i\backslash j,n}_{SI} & M^{i\backslash j,n}_{SR} \\ M^{i\backslash j,n}_{IS} & M^{i\backslash j,n}_{II} & M^{i\backslash j,n}_{IR} \\ M^{i\backslash j,n}_{RS} & M^{i\backslash j,n}_{RI} & M^{i\backslash j,n}_{RR} \end{pmatrix} \tag{10.33a}$$

$$= \begin{pmatrix} g^n_{i\backslash j} P(O_i^n|S) & \left(1-g^n_{i\backslash j}\right) P(O_i^n|S) & 0 \\ 0 & (1-r_i^n)\hat{g}^n_{i\backslash j} P(O_i^n|I) & r_i^n \hat{g}^n_{i\backslash j} P(O_i^n|I) \\ 0 & 0 & P(O_i^n|R) \end{pmatrix}, \tag{10.33b}$$

and

$$M^{i\backslash j,n} = \begin{pmatrix} g^n_{i\backslash j} P(O^n_i|S) & \left(1 - g^n_{i\backslash j}\right) P(O^n_i|S) & 0 \\ 0 & (1-r^n_i)\hat{g}^n_{i\backslash j} P(O^n_i|I) & r^n_i \hat{g}^n_{i\backslash j} P(O^n_i|I) \\ \sigma^n_i P(O^n_i|R) & 0 & (1-\sigma^n_i) P(O^n_i|R) \end{pmatrix}, \tag{10.34}$$

where we introduced the shorthand notations

$$g^n_{i\backslash j} = \alpha^n_i e^{\sum_{k\in\partial i} J^n_{ik} \mu^n_{k\backslash i}}, \tag{10.35}$$

$$\hat{g}^n_{i\backslash j} = e^{\sum_{k\in\partial i} J^n_{ki} \hat{\mu}^n_{k\backslash i}}. \tag{10.36}$$

As long as further compartments are included with transitions being parametrized by individual-based rates, generalization of the above construction follows straightforwardly (e.g., Susceptible-Exposed-Infected-Recovered (SEIR) and Susceptible-Exposed-Infected-Recovered-Susceptible (SEIRS) models).

10.2.2 Effect of the response function

The additional order parameter $\hat{\mu}$ is the distinctive feature of SCDC. It captures the impact of observations that propagate backward in time and is therefore essential for inference. To illustrate its role, we consider a uniform SI model with $\lambda^n_{ij} = \lambda$ and $\alpha^n_i = \alpha$ on a contact graph built by randomly adding edges to a K-regular tree. When an observation is made at time n_{obs} on node i, the outgoing fields $\hat{\mu}^n_{i\backslash j}$ become nonzero for all $n < n_{\text{obs}}$ and start propagating through the network. As shown in Figure 10.1, these fields carry a signed signal: if a node is observed infected, the response tends to be small and negative as it flows backward in time from the observed node; if a node is observed susceptible, the response is positive and often close to one. The magnitude increases as one moves backward from the observation time, reflecting the accumulation of explanatory pathways, and it is strongly filtered by the network structure, so that some directions are less favored than others. This nontrivial anisotropy reflects the central role of topology in shaping the backpropagation of information.

A complementary, probabilistic view of the role of the response fields is obtained by monitoring the temporal behavior of the normalized time-forward and time-backward messages,

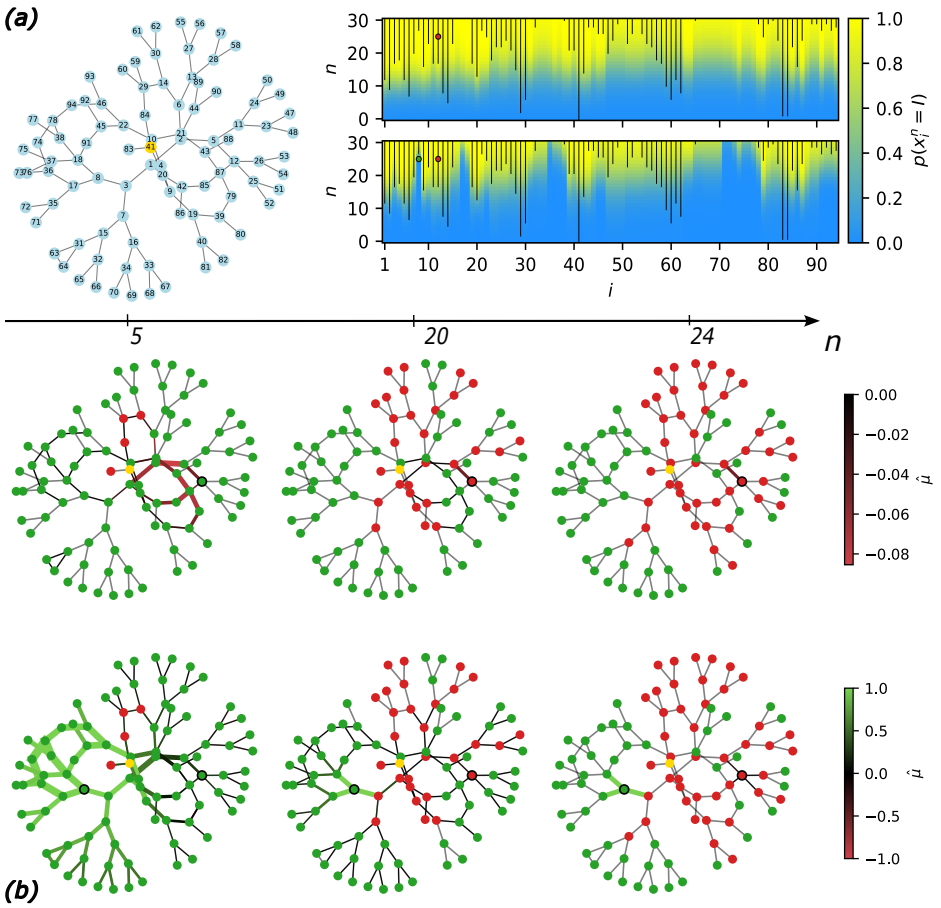

Fig. 10.1 **Backward propagation of response fields induced by observations. Panel a-left**: contact graph obtained by randomly adding links to a tree; the epidemic source is shown in yellow. **Panel a-right**: inference performance of SCDC with one observation on node 12 at time $n_{\text{obs}} = 25$ (top) and two observations on nodes 12 and 8 at the same time (bottom). The panels report the posterior probability of being infected as a function of time (vertical axis) for all nodes in a single epidemic outbreak generated by a uniform SI model with $\lambda = 0.19$. The same parameter is used for inference. Observed states are marked by colored dots (green $= S$, red $= I$). Black vertical lines indicate the ground truth infection intervals. **Panel b**: spatio-temporal propagation of the cavity response fields from a single observed node (top row) and two observed nodes (bottom row). For each edge (i, j) the line thickness is proportional to $\left|\hat{\mu}^n_{i\setminus j} + \hat{\mu}^n_{j\setminus i}\right|$, while the color encodes the value of the sum $\hat{\mu}^n_{i\setminus j} + \hat{\mu}^n_{j\setminus i}$ (see legend). Time flows from left to right. Node states are indicated by colored dots (green $= S$, red $= I$), while observed nodes are circled in black.

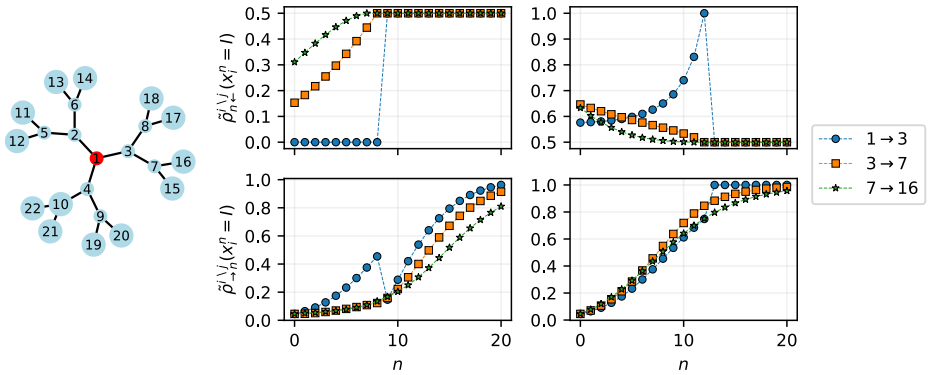

Fig. 10.2 **Time evolution of forward and backward cavity messages.** Normalized backward (top row) and forward (bottom row) cavity messages $\tilde{\rho}^{i\backslash j}_{n\leftarrow}(x_i^n = I)$ and $\tilde{\rho}^{i\backslash j}_{\rightarrow n}(x_i^n = I)$ on selected edges of a small tree. The outbreak is generated by a uniform SI model with $\lambda = 0.1$. The left panel shows the graph; the observed node is red, the others are blue. The observed node becomes infected at $n = 10$. Left column: observation $x = S$ at $n_{\text{obs}} = 8$. Right column: observation $x = I$ at $n_{\text{obs}} = 12$.

$\tilde{\rho}^{i\backslash j}_{\rightarrow n}(x_i^n) = \rho^{i\backslash j}_{\rightarrow n}(x_i^n)/Z^{i\backslash j}_{\rightarrow n}$ and $\tilde{\rho}^{i\backslash j}_{n\leftarrow}(x_i^n) = \rho^{i\backslash j}_{n\leftarrow}(x_i^n)/Z^{i\backslash j}_{n\leftarrow}$, where the normalization factors are

$$Z^{i\backslash j}_{\rightarrow n} = \sum_{x_i^{n-1}, x_i^n \in \{S,I\}} \rho^{i\backslash j}_{\rightarrow n-1}(x_i^{n-1})\, M^{i\backslash j}_{x_i^{n-1} x_i^n}, \tag{10.37}$$

$$Z^{i\backslash j}_{n\leftarrow} = \sum_{x_i^n, x_i^{n+1} \in \{S,I\}} \rho^{i\backslash j}_{n+1\leftarrow}(x_i^{n+1})\, M^{i\backslash j}_{x_i^n x_i^{n+1}}. \tag{10.38}$$

Consider a homogeneous SI realization on a small tree, shown in Figure 10.2 (left), where the root is infected at $n = 10$. An observation at time n_{obs} modifies both families of messages on all directed edges. If the root is observed susceptible, the time-backward message exiting the root is identically zero at all earlier times; moving away from the observed node the backward probability becomes nonzero (and typically increases with the graph distance) but decreases as one goes further back in time. If the root is observed infected, the backward message exhibits an instantaneous jump to one at the observation time and then decays when going backward. Moving away from the root, the backward messages tend to increase at earlier times, consistent with the fact that neighbors might have caused the infection of the observed node. In the absence of information, normalized time-backward messages relax to 0.5 at later times. The normalized time-forward messages display a complementary pattern: they increase with time to account for the observed cascade; they drop when a node is observed susceptible and recover afterward due to the forward spread, while for an infected observation they approach one after a short delay.

10.3 Benchmarking on the SI model

10.3.1 Time-forward dynamics

The apparent breaking of causality in the SCDC equations originates from the presence of observations at later times, which induce a backward propagation of information through the conjugate fields $\hat{\mu}$. In the absence of observations, all $\hat{\mu}^n_{i\backslash j}$ vanish identically, and the dynamics reduces to a causal, time-forward evolution of the standard mean-field type.

This reduction can be proven inductively by showing that the response fields $\hat{\mu}$ vanish whenever no observations are present. If $p(O^n_i|x^n_i) = 1$ for all i,n, the backward messages at the final time T are uniform, $\rho^{i\backslash j}_{T\leftarrow}(x^T_i) = 1$. From Eq. (10.28), this implies $\hat{\mu}^{T-1}_{i\backslash j} = 0$. The backward messages at the preceding time step are then

$$\rho^{i\backslash j}_{T-1\leftarrow}(x^{T-1}_i) = \sum_{x^T_i} M^{i\backslash j,T-1}_{x^{T-1}_i,x^T_i} = 1, \tag{10.39}$$

where we used that the columns of the transfer matrix sum to one in the absence of observations and with $\hat{\mu} = 0$. Assuming now that the backward messages are uniform at time $n+1$, one has $\hat{\mu}^n_{i\backslash j} = 0$ from Eq. (10.28), which in turn implies that the backward messages at time n are also uniform:

$$\rho^{i\backslash j}_{n\leftarrow}(x^n_i) = \sum_{x^{n+1}_i} M^{i\backslash j,n}_{x^n_i,x^{n+1}_i} = 1. \tag{10.40}$$

By induction, $\hat{\mu}^n_{i\backslash j} = 0$ for all n in the absence of observations, and backward messages remain uniform at all times.

It follows that, without observations, the SCDC equations for the cavity infection probabilities reduce to

$$\mu^n_{i\backslash j} = \frac{1}{\mathscr{Z}_{i\backslash j}} \rho^{i\backslash j}_{\to n}(I) = \frac{1}{\mathscr{Z}_{i\backslash j}} \sum_{x^{n-1}_i} \rho^{i\backslash j}_{\to n-1}(x^{n-1}_i) M^{i\backslash j,n-1}_{x^{n-1}_i,I}, \tag{10.41}$$

$$= \mu^{n-1}_{i\backslash j} M^{i\backslash j,n-1}_{II} + \frac{1}{\mathscr{Z}_{i\backslash j}} \sum_{x^{n-1}_i \neq I} \rho^{i\backslash j}_{\to n-1}(x^{n-1}_i) M^{i\backslash j,n-1}_{x^{n-1}_i,I}. \tag{10.42}$$

For the SI model, this gives

$$\mu^n_{i\setminus j} = \mu^{n-1}_{i\setminus j} + \frac{\rho^{i\setminus j}_{\to n-1}(S)}{\mathscr{Z}_{i\setminus j}}\left(1-\alpha^{n-1}_i e^{\sum_{k\in\partial i\setminus j} J^{n-1}_{ik}\mu^{n-1}_{k\setminus i}}\right), \tag{10.43}$$

$$= 1-(1-\mu^{n-1}_{i\setminus j})\alpha^{n-1}_i \exp\left(\sum_{k\in\partial i\setminus j} J^{n-1}_{ik}\mu^{n-1}_{k\setminus i}\right), \tag{10.44}$$

and similarly for the single-node marginals:

$$\mu^n_i = 1-(1-\mu^{n-1}_i)\alpha^{n-1}_i \exp\left(\sum_{k\in\partial i} J^{n-1}_{ik}\mu^{n-1}_{k\setminus i}\right). \tag{10.45}$$

Numerically, one can verify that in the absence of observations, the SCDC algorithm in Eqs. (10.25b) and (10.27) to (10.30) converges to the same results obtained by iterating Eqs. (10.44) and (10.45) forward in time from the initial conditions $\mu^0_{i\setminus j} = p^0_i(x^0_i = I)$.

For small infection probabilities λ^n_{ij}, Eqs. (10.44) and (10.45) can be linearized to yield

$$\mu^n_{i\setminus j} \approx 1-(1-\mu^{n-1}_{i\setminus j})\alpha^n_i \prod_{k\in\partial i\setminus j}(1-\lambda^n_{ik}\mu^{n-1}_{k\setminus i}), \tag{10.46}$$

$$\mu^n_i \approx 1-(1-\mu^{n-1}_i)\alpha^n_i \prod_{k\in\partial i}(1-\lambda^n_{ik}\mu^{n-1}_{k\setminus i}), \tag{10.47}$$

which represent a natural cavity generalization of the SMF equations in Eq. (9.9). The SMF equations are recovered by replacing the cavity marginals $\mu^{n-1}_{k\setminus i}$ with the single-node marginals μ^{n-1}_k in Eq. (10.47). The cavity formulation yields improved accuracy, particularly on sparse networks, as it partially accounts for local correlations by removing direct feedback between neighboring nodes.

Figure 10.3 illustrates the quality of the approximation obtained from Eqs. (10.44) and (10.45) for purely time-forward SI dynamics in the absence of observations. Simulations were performed on standard classes of static graphs, including regular trees, RRGs, and proximity random graphs (see section 10.3.2 for details). Results from SCDC are compared with those from the BP algorithm (section 9.3) and the SMF equations (section 9.2). For reference, averages over 10^4 Monte Carlo realizations of the exact time-forward SI dynamics are also shown. All mean-field-based approaches tend to overestimate the number of infected individuals, as correlations between neighboring nodes are neglected. BP, which is exact on trees, closely reproduces Monte Carlo results and remains highly accurate on sparse graphs. Both SCDC and SMF overestimate infection prevalence, though SCDC systematically provides a more

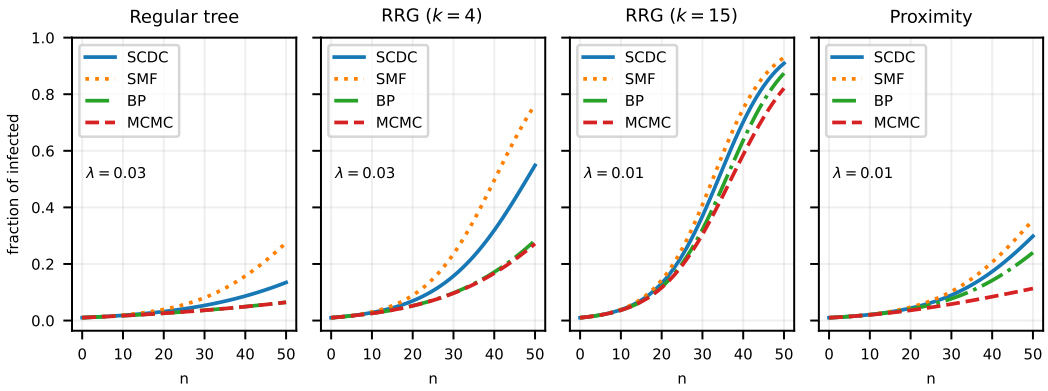

Fig. 10.3 **Time-forward epidemic dynamics on different contact networks.** Time-forward dynamics obtained from SCDC, SMF, BP, and Monte Carlo simulations (10^4 realizations). Each panel shows the fraction of infected individuals as a function of time for four different static contact graphs. From left to right: regular tree with degree $K = 4$ and $N = 485$; RRG with $N = 500$ and $K = 4$; RRG with $N = 500$ and $K = 15$; proximity graph with $N = 500$. The infection probability λ is reported in each panel. The initial infection probability is $\gamma^0 = 5/N$, and $\varepsilon_i^n = 0$.

accurate description. Performance is comparable on dense random graphs but deteriorates on spatially structured networks such as proximity graphs, where short loops can create additional correlations. In all tested cases, SCDC consistently outperforms SMF.

A detailed analysis of the role of λ is reported in Fig. E.1 of Appendix E.3. For small λ, SCDC yields significantly better estimates of the infected fraction than SMF across all graph ensembles. For larger λ, SMF performs slightly better on random regular and proximity graphs, though the difference remains marginal. Overall, the time-forward reduction of SCDC provides a robust and accurate alternative to SMF, with a computational cost that scales as $O(|E|T)$ instead of $O(NT)$, representing a modest increase for sparse graphs.

10.3.2 Epidemic inference and risk assessment

We examine a standard epidemic risk assessment scenario in which, given a contact network and a small number of observations collected at the final time step of a stochastic epidemic realization, the task is to infer the posterior infection probability of each unobserved individual. Epidemic outbreaks following the SI dynamics are simulated using the `EpiGen` Python package [99] on both synthetic and empirical contact networks. A random subset of individuals is observed at the last time T, and the inferred probabilities are compared with the ground truth states to evaluate inference accuracy. For each method, the classification performance is quantified through the Receiver Operating Characteristic (ROC) curve, which plots the true

positive rate versus the false positive rate. The area under the ROC curve (AUC) provides a quantitative measure of the probability of correctly identifying infected individuals.

Inference with the SCDC method is performed using its efficient transfer matrix formulation, and its accuracy is benchmarked against the SMF method described in section 9.2 and the BP algorithm presented in section 9.3. For both BP and SCDC, the individual infection probabilities at the final time are obtained from the corresponding single-node marginals once the iterative message-passing updates have converged, i.e., when the error on the cavity marginals falls below a prescribed tolerance threshold. In some instances, neither BP nor SCDC reaches convergence within a reasonable number of iterations (typically a few thousand in the analyzed epidemic instances). Such non-convergence is often related to the presence of strong loop correlations and long-range dependencies in the contact network. In these cases, the node marginals are estimated by averaging the messages over a sufficiently long sequence of updates (typically hundreds of iterations). A systematic analysis of the convergence properties of both algorithms and their impact on inference accuracy is reported later in this section. For reference, the SMF approach used here corresponds to the inference method proposed in [97], where the information provided by observations of susceptible or infected individuals is incorporated as external constraints on the forward equations.

Figure 10.4 summarizes the results obtained for various network ensembles. The top panels report results for synthetic contact structures. The first two panels correspond to static random graphs: Watts-Strogatz networks [100] and soft random geometric (or proximity) graphs [101]. In the Watts-Strogatz model, edges of an initially regular ring lattice are randomly rewired with probability p_{rw}, producing small-world networks characterized by short average path lengths and high clustering. In proximity graphs, individuals are uniformly distributed in the unit square and connected with probability decaying exponentially with their Euclidean distance l, up to a cutoff $l_{\max}$. Both ensembles display a strong local structure with short loops and clustering. The last two panels of the top row correspond to more realistic contact structures generated by agent-based models, namely OpenABM-Covid19 [102] and Covasim [103]. These models simulate daily interactions within households, workplaces, schools, and community environments, producing temporally varying weighted contact networks that mimic real social mixing patterns. In our analysis, we only use the static contact layers aggregated over a few weeks, on which the SI dynamics is simulated. For each link (i, j), a weight w_{ij}^n represents the cumulative contact duration between i and j on day n, and the corresponding infection probability is computed as $\lambda_{ij}^n = 1 - \exp(-\gamma w_{ij}^n)$, with γ denoting the per contact infection rate. Across all synthetic settings, when a limited number of observations is available at the final time, the SCDC method consistently outperforms SMF and attains accuracy comparable to BP.

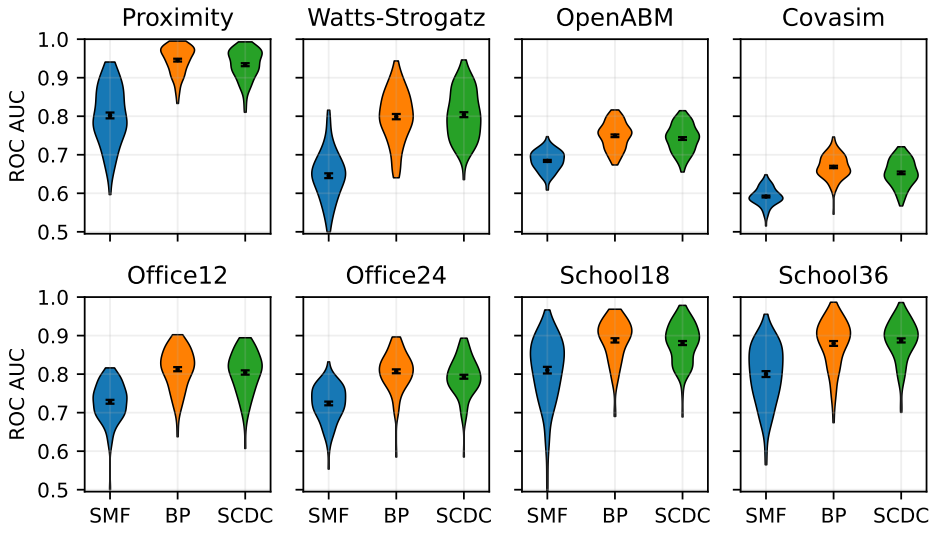

Fig. 10.4 **Inference performance on SI model across synthetic and real networks.** Inference performances for SMF, BP, and SCDC on synthetic (top row: proximity, Watts-Strogatz, OpenABM, Covasim) and real-world (bottom row: office and school) contact networks. Violin plots display the distribution of AUC values over 100 epidemic realizations; black markers indicate mean AUCs with standard errors. Simulation parameters: proximity ($N = 600$, $T = 28$, $\lambda = 0.08$), Watts-Strogatz ($N = 600$, $T = 25$, $\lambda = 0.16$), OpenABM ($N = 2000$, $T = 21$, $\gamma = 0.026$), Covasim ($N = 1000$, $T = 24$, $\gamma = 0.038$), office ($N = 219$, $T = 12$ or 24, $\gamma = 6\times10^{-4}$), school ($N = 328$, $T = 18$ or 36, $\gamma = 7\times10^{-5}$).

The same inference protocol is applied to two empirical contact networks collected via RFID sensors and made available in Ref. [104]: a high-school network (Thiers13) and an office network (InVS15). Both datasets record face-to-face interactions over multiple days with a 20-second temporal resolution. For computational tractability, contact events are aggregated into coarse-grained temporal windows of duration τ_w, ranging from 3 hours to 24 hours, yielding total time horizons T between 12 and 36 steps. Within each window n, the number of contacts c_{ij}^n between individuals i and j is used to define an effective infection probability $\lambda_{ij}^n = 1-(1-\gamma)^{c_{ij}^n}$, where γ quantifies the infectiousness of a single contact. The results obtained for these real-world datasets, shown in the bottom row of Figure 10.4, confirm that SCDC achieves performance nearly indistinguishable from BP and systematically superior to the SMF heuristic, across all considered network structures and temporal resolutions.

Effect of the infection probability and number of observations

To further characterize the inference performance, we studied how accuracy varies with the infection probability λ and the fraction of observed individuals. Figure 10.5a reports the AUC

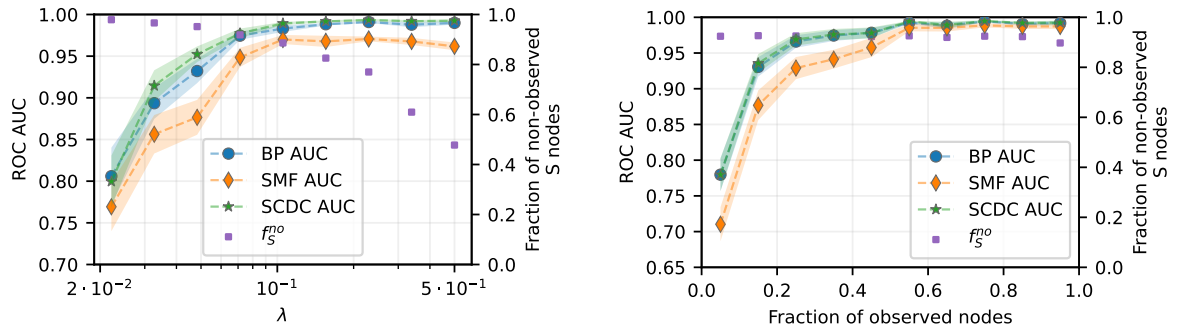

(a) Dependence on infection probability λ with 30% observed individuals.

(b) Dependence on observation fraction with $\lambda = 0.15$.

Fig. 10.5 **Effect of infection probability and observation rate on inference accuracy.** AUC comparison of SCDC, BP, and SMF on soft random geometric graphs. Results are averaged over 100 random networks ($N = 300$, $T = 15$, $l_{\max} = \sqrt{1.8/N}$). Shaded regions represent standard errors, while purple markers indicate the fraction of non-observed susceptible individuals at the final time.

as a function of λ for epidemics simulated on soft random geometric graphs with $N = 300$ nodes and cutoff distance $l_{\max} = \sqrt{1.8/N}$. Each curve is averaged over 100 epidemic realizations with two initially infected nodes and an observation fraction of 30% at the final time. The infection probability λ is varied between 0.02 and 0.5. Remarkably, the SCDC algorithm achieves accuracy comparable to BP even at relatively large λ values, despite its derivation being based on a small-coupling approximation. Throughout the entire range, both SCDC and BP outperform SMF.

Figure 10.5b shows the dependence of the AUC on the fraction of observed individuals for the same network ensemble and fixed $\lambda = 0.15$. For small fractions of observed nodes, SCDC exhibits accuracy nearly identical to BP and clearly superior to SMF. As the observation fraction increases, all methods tend to converge toward similar performances. The convergence frequencies of BP and SCDC under varying λ and observation fractions are summarized in Tables 10.1 and 10.2. While BP converges more often, the two methods maintain comparable accuracy even when SCDC convergence is partial. Notably, the convergence behavior of SCDC is non-monotonic with respect to the fraction of observed individuals. Convergence is high at low observation fractions, where the scarcity of constraints simplifies the search for a solution (despite lower inference quality), and at high fractions, where the system is rigidly constrained. A minimum occurs at intermediate regimes: here, the inference is most challenging due to the high degeneracy of possible dynamical paths compatible with the observations. Conversely, BP exhibits a constant convergence rate across all fractions, suggesting that the approximations inherent to SCDC result in a loss of stability precisely when the multiplicity of explanatory trajectories is maximal.

Table 10.1 Convergence rates of BP and SCDC as a function of the infection probability λ. Same setting as in Fig. 10.5a.

λ	0.02	0.03	0.05	0.07	0.1	0.15	0.23	0.34	0.5
BP	1.0	0.98	1.0	0.68	0.48	0.20	0.12	0.06	0.0
SCDC	0.7	0.5	0.48	0.20	0.16	0.16	0.16	0.08	0.06

Table 10.2 Convergence rates of BP and SCDC as a function of the fraction of observed nodes. Same setting as in Fig. 10.5b.

Frac. obs.	0.05	0.15	0.25	0.35	0.45	0.55	0.65	0.75	0.85	0.95
BP	0.8	0.8	0.8	0.8	0.8	0.8	0.8	0.8	0.8	0.8
SCDC	0.8	0.4	0.22	0.28	0.26	0.22	0.38	0.44	0.4	0.5

In some epidemic instances, BP and SCDC fail to converge within a few thousand iterations. This behavior is typically associated with pronounced loop structures or strong long-range correlations. In these cases, the marginals are obtained by averaging the iterative updates over hundreds of steps. Importantly, the convergence tests reported in Tables 10.1 and 10.2 correspond to the same epidemic realizations used in Fig. 10.5a and Fig. 10.5b. Although BP generally exhibits higher convergence rates, the overall inference accuracy of SCDC remains comparable, even in settings where convergence is only partial.

10.3.3 Epidemic reconstruction of recurrent models

The results presented so far have focused on irreversible epidemic dynamics, where individuals can only transition from susceptible to infected and never recover. We now turn to recurrent epidemic processes, where recovery and loss of immunity allow for reinfections, thus introducing new challenges for Bayesian inference due to the presence of cyclic trajectories and long-range temporal correlations. In this context, we test the capability of the SCDC framework to reconstruct dynamic states under recurrent dynamics, and compare its performance to that of the Matrix-Product Belief Propagation (MPBP) method [105]. MPBP is a powerful approximation of the full BP scheme, based on the ansatz that the cavity messages can be parametrized as a product of matrices, with one matrix associated with each time step. While this representation is theoretically exact in the limit of infinite matrix dimensions, algorithmic tractability is achieved via singular value decomposition. This technique is used to truncate the matrices to a fixed bond dimension M, effectively compressing the representation while retaining the most significant temporal correlations.

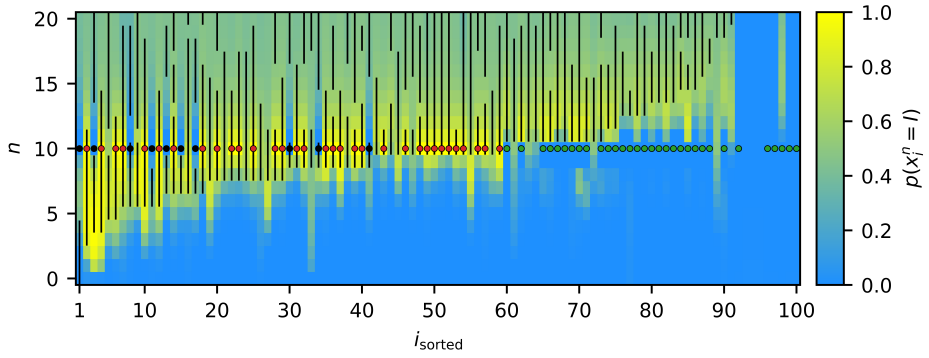

Fig. 10.6 **Posterior infection probabilities in SIRS epidemics.** Posterior infection probabilities for all nodes of an ERG ($N = 100$, $K = 3$) during a single SIRS outbreak with $\lambda = 0.4$ and $r = \sigma = 0.15$. At $n = 10$, 75% of nodes are observed and marked with colored dots (green $= S$, red $= I$, black $= R$). Black bars mark the true infection intervals, with nodes ordered by their first infection time.

Inference on the SIRS model. We start by analyzing a single realization of an SIRS process on an ER random graph with $N = 100$ and average degree $K = 3$. Each node can transition from the infected to the recovered state with probability r, and from recovered back to susceptible with probability σ. We set $\lambda = 0.4$ and $r = \sigma = 0.15$ uniformly across all nodes and times. Figure 10.6 displays the posterior infection probabilities μ_i^n inferred by SCDC as a function of time, superimposed on the true infection intervals. The correspondence between the inferred probabilities and the ground-truth trajectories is remarkably accurate, even for nodes that were not directly observed or for time intervals far from the observations. Several reinfection events are correctly identified, indicating that the SCDC fields $\hat{\mu}$ successfully propagate information backward in time, updating the probabilities of past events in a consistent way with the observed data.

To quantitatively assess performance, we compared SCDC and MPBP on the same class of SIRS processes. As a measure of inference accuracy, we considered the fraction of correct Maximum-a-Posteriori (MAP) predictions, defined as

$$f_{\text{correct}}^{\text{MAP}} = \frac{1}{NT} \sum_{i=1}^{N} \sum_{n=1}^{T} \delta\left(\text{argmax}_{x_i^n \in \{S,I,R\}}\, c_i(x_i^n), X_i^n\right), \tag{10.48}$$

where $c_i(x_i^n)$ is the inferred marginal probability for node i at time n, X_i^n is the ground truth, and δ denotes the Kronecker symbol. Figure 10.7 reports $f_{\text{correct}}^{\text{MAP}}$ as a function of λ, averaged over 50 network realizations and epidemic outbreaks. At small λ, all algorithms achieve nearly perfect reconstruction, as the slower spread makes the epidemic easier to infer. As λ increases, inference becomes harder, yet SCDC maintains a performance comparable to MPBP with bond

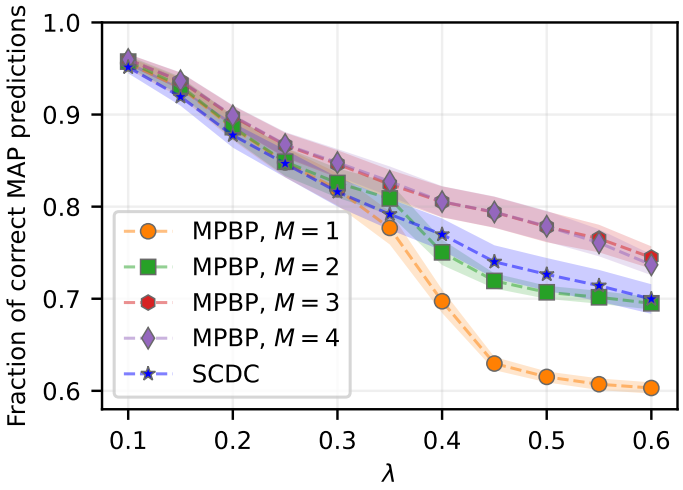

Fig. 10.7 **Comparison between SCDC and MPBP on SIRS epidemics.** Fraction of correct MAP predictions for SCDC and MPBP (bond dimensions $M = 1,2,3,4$) on ERGs ($N = 100$, $K = 2.5$). The SIRS dynamics has λ, $r = \sigma = 0.15$, $T = 20$, with two initially infected nodes and 75% of nodes observed at $n = 10$. Results are averaged over 50 instances; shaded ribbons indicate standard errors.

dimension $M = 2$. Furthermore, the improvement of MPBP saturates for $M > 3$, suggesting that higher bond dimensions do not yield additional benefits in this setting.

From a computational standpoint, SCDC is considerably more efficient: while the MPBP complexity scales as $O(S^6|E|TM^6)$ for a model with S states, SCDC only requires $O(S^2|E|T)$ operations. Although MPBP may achieve higher accuracy for very large M, the simplicity and interpretability of SCDC make it an appealing alternative. The fields μ and $\hat{\mu}$ retain a clear physical meaning as infection probabilities and response terms, offering valuable insight into how information propagates across space and time.

Inference on the SIS model. We next analyze the simpler recurrent case of the SIS model, where recovered individuals immediately return to the susceptible state and can be reinfected indefinitely. This model provides a convenient testbed for exploring how inference accuracy depends on the competing effects of infection and recovery. Figure 10.8 shows the fraction of correct MAP classifications for SCDC as a function of the infection probability λ and the recovery probability σ, on soft random geometric graphs with $N = 300$ and cutoff distance $l_{\max} = \sqrt{1.8/N}$. Epidemic realizations last for $T = 15$ steps, starting from two infected individuals, and a fraction of 60% of the nodes is observed at $n = 7$. Each data point is averaged over 30 independent networks and outbreak instances.

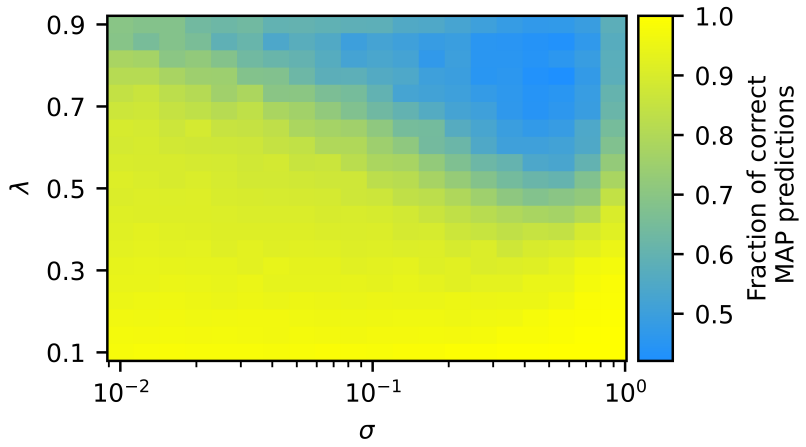

Fig. 10.8 **Inference accuracy for the SIS model in the (λ, σ) plane.** Fraction of correct MAP predictions for the SIS model on soft random geometric graphs ($N = 300$, $T = 15$, $l_{\max} = \sqrt{1.8/N}$). At $n = 7$, 60% of nodes are observed. Accuracy is shown as a function of infection λ and recovery σ, averaged over 30 instances. Lighter colors denote higher accuracy.

Across most of the parameter space, the SCDC reconstruction remains robust, with high fractions of correctly inferred states. Only in the upper-right region of the (λ, σ) plane—where both infection and recovery are fast and reinfections occur frequently—the inference accuracy deteriorates due to the highly entangled temporal structure of the dynamics. This regime is intrinsically difficult for any trajectory-based inference approach, as the epidemic trajectories fluctuate rapidly and the same individuals may switch state multiple times within the observation window.

Overall, these results demonstrate that the SCDC framework extends naturally to recurrent epidemic processes, maintaining a good balance between computational efficiency and predictive accuracy. While MPBP may reach slightly higher accuracy at increased computational cost, SCDC provides a principled, interpretable, and scalable approximation capable of handling both irreversible and recurrent epidemic dynamics on large, heterogeneous networks.

Chapter 11

Conclusions and Outlook

This thesis has developed a coherent set of cavity-based approximations for stochastic dynamics on sparse graphs, bridging continuous- and discrete-state variables within a unified framework. On the side of continuous degrees of freedom, the *Gaussian Expansion Cavity Method (GECaM)* was introduced as a message-passing formulation for linearly-coupled stochastic differential equations, obtained by combining the dynamic cavity construction with a second-order small-coupling expansion. At this order, the cavity messages are assumed to be Gaussian, leading to closed self-consistent equations for local means, variances, and two-time response and correlation functions. In the case of linear drift and additive noise, the method reproduces the exact dynamics and provides an efficient numerical scheme for large sparse systems. The GECaM equations for the response functions can be directly related to the spectral properties of the interaction matrix, establishing a bridge between temporal relaxation and random-matrix theory that underpins the analysis of stability. Beyond the Gaussian regime, a perturbative closure around the second-order expansion was developed to incorporate weak nonlinearities in the drift and in the noise covariance, which generate non-Gaussian fluctuations and noise-induced correlations. These extensions allow one to describe noise-driven phase transitions and slow relaxation phenomena such as ageing in the spherical two-spin model. Altogether, GECaM emerges as a versatile framework to analyze non-equilibrium stochastic processes on locally tree-like structures.

The same philosophy–retaining the cavity methods while organizing corrections in small couplings–has been extended to interacting ecological communities. For sparse Generalized Lotka-Volterra (GLV) systems, the cavity formulation produces effective single-species dynamics driven by self-consistent random fields and colored noise. This description captures how finite connectivity and disorder reshape species abundance distributions and stability,

producing rich regimes and non-Gaussian abundance statistics that were never observed in dense mean-field limits. The analysis clarifies when and how sparse topology alone can induce localized instabilities and collective transitions, and it offers a tractable route to quantify the interplay of reciprocity, heterogeneity, and degree in determining community resilience.

For discrete dynamics, the *Small-Coupling Dynamic Cavity (SCDC)* method was introduced as a cavity-based approximation tailored to Bayesian epidemic inference on networks. Starting from an observation reweighted dynamic cavity formulation, a first-order expansion delivers a closed system for one-time infection probabilities and their conjugate response fields. The latter encode the non-causal flow of information induced by observations and vanish in their absence, thereby smoothly reducing to a straightforward cavity generalization of standard mean-field approximations. In practice, SCDC achieves a computational complexity that scales linearly with the time horizon through a transfer matrix implementation. Its accuracy has been systematically benchmarked against Belief Propagation (BP) and mean-field approximations on the non-recurrent Susceptible-Infected (SI) model, where BP can be applied exactly. In this setting, SCDC attains inference performances comparable to BP and consistently superior to simple mean-field schemes. When extended to recurrent epidemic models such as Susceptible-Infected-Susceptible (SIS) and Susceptible-Infected-Recovered-Susceptible (SIRS), where the number of possible trajectories grows combinatorially and BP becomes computationally intractable, SCDC remains applicable and exhibits excellent reconstruction accuracy, providing a scalable approximation for inference in general recurrent dynamics.

Several limitations naturally delineate the frontier for future research. Although the present formulation of GECaM has proven effective for the classes of systems investigated in this thesis, the development of a non-perturbative closure appears indispensable to capture the local effects of nonlinearities without introducing further approximations. A promising extension would combine the analytical Gaussian expansion with a numerical sampling of on-site stochastic trajectories conditioned on incoming Gaussian cavity messages, in the spirit of recent approaches proposed for random ecosystems [92]. From these samples, updated estimates of cavity response and correlation functions could be obtained self-consistently. A major current limitation of the method lies in the assumption of linear coupling between neighboring degrees of freedom. In many relevant contexts, interactions cannot be reduced to a weighted sum of neighboring variables, owing to the presence of non-linear activation functions–as in neural or learning models–or higher-order interaction patterns, as in ecosystems and reaction networks. Extending the path-integral formulation to encompass these more complex forms of coupling would require a richer graphical representation and a more elaborate functional structure of the cavity messages, potentially involving the introduction of auxiliary variables to retain computational tractability.

The main limitation of the SCDC method lies in the fact that its efficient transfer matrix formulation is, at present, restricted to Markovian dynamics. Further investigation is required to devise an algorithmic scheme capable of treating non-Markovian recurrent epidemic models, where temporal correlations extend beyond a single time step. From a methodological perspective, an interesting direction for future work concerns the analysis of the second-order terms neglected in the small-coupling expansion, and the development of an improved version of the algorithm that consistently incorporates their contribution. Finally, the generality of the SCDC framework suggests the possibility of extending it beyond epidemic processes to other classes of dynamical models on networks, such as rumor spreading dynamics [19, 106, 107], information diffusion, or opinion formation.

Beyond these methodological avenues, the broader message is that sparse structure is not a mere complication to be averaged away, but a primary source of dynamical phenomena. By combining local, graph-aware closures with principled expansions, the methods developed here reconcile the descriptive power of message passing at finite connectivity with the limiting consistency of dense mean-field theories. The resulting toolkit–GECaM for continuous variables and SCDC for discrete-epidemic processes–offers a scalable and interpretable pathway to analyze, infer, and ultimately control stochastic dynamics on networks, from ecological stability to epidemiological forecasting.

In this sense, the thesis contributes both techniques and perspectives: techniques that turn path-integral and cavity ideas into practical algorithms on large graphs; perspectives that foreground the spectral and topological fingerprints of dynamics. Closing these gaps–non-perturbative treatments of nonlinearities, rigorous convergence theory on loopy graphs, and data assimilative extensions on time varying networks–promises a unified and predictive theory of stochastic dynamics on complex sparse systems.

Bibliography

[1] B. Bollobás. *Random Graphs*. 2nd ed. Cambridge Studies in Advanced Mathematics. Cambridge: Cambridge University Press, 2001. ISBN: 978-0-521-80920-7. DOI: 10.1017/CBO9780511814068.

[2] A. Barrat, M. Barthélemy, and A. Vespignani. *Dynamical Processes on Complex Networks*. Cambridge: Cambridge University Press, 2008. ISBN: 978-0-521-87950-7. DOI: 10.1017/CBO9780511791383.

[3] M. Newman. *Networks*. 2nd ed. Oxford University PressOxford, July 2018. ISBN: 978-0-19-880509-0. DOI: 10.1093/oso/9780198805090.001.0001.

[4] P. Erdős and A. Rényi. "On Random Graphs I". In: *Publ. math. debrecen* 6.3-4 (12) (1959), pp. 290–297. DOI: 10.5486/PMD.1959.6.3-4.12.

[5] D. Aldous and J. M. Steele. "The Objective Method: Probabilistic Combinatorial Optimization and Local Weak Convergence". In: *Probability on Discrete Structures*. Ed. by H. Kesten. Berlin, Heidelberg: Springer, 2004, pp. 1–72. ISBN: 978-3-662-09444-0. DOI: 10.1007/978-3-662-09444-0_1.

[6] M. Mézard and A. Montanari. *Information, Physics, and Computation*. Oxford University Press, Jan. 2009. ISBN: 978-0-19-857083-7. DOI: 10.1093/acprof:oso/9780198570837.001.0001.

[7] M. Mezard, G. Parisi, and M. Virasoro. *Spin Glass Theory and Beyond: An Introduction to the Replica Method and Its Applications*. Vol. 9. World Scientific Lecture Notes in Physics. WORLD SCIENTIFIC, Nov. 1986. ISBN: 978-9971-5-0116-7. DOI: 10.1142/0271.

[8] P. Charbonneau et al. *Spin Glass Theory and Far Beyond*. WORLD SCIENTIFIC, 2023. DOI: 10.1142/13341.

[9] H. A. Bethe. "Statistical Theory of Superlattices". In: *Proceedings of the Royal Society of London. Series A - Mathematical and Physical Sciences* 150.871 (Jan. 1997), pp. 552–575. DOI: 10.1098/rspa.1935.0122.

[10] R. Gallager. "Low-Density Parity-Check Codes". In: *IRE Transactions on Information Theory* 8.1 (Jan. 1962), pp. 21–28. ISSN: 2168-2712. DOI: 10.1109/TIT.1962.1057683.

[11] J. Pearl. *Reverend Bayes on Inference Engines: A Distributed Hierarchical Approach*.

[12] Y. Weiss. "Correctness of Local Probability Propagation in Graphical Models with Loops". In: *Neural Comput* 12.1 (Jan. 2000), pp. 1–41. ISSN: 0899-7667. DOI: 10.1162/089976600300015880.

[13] J. M. Mooij and H. J. Kappen. "Sufficient Conditions for Convergence of the Sum–Product Algorithm". In: *IEEE Transactions on Information Theory* 53.12 (Dec. 2007), pp. 4422–4437. ISSN: 1557-9654. DOI: 10.1109/TIT.2007.909166.

[14] M. Chertkov. "Exactness of Belief Propagation for Some Graphical Models with Loops". In: *J. Stat. Mech.* 2008.10 (Oct. 2008), P10016. ISSN: 1742-5468. DOI: 10.1088/1742-5468/2008/10/P10016.

[15] M. Bayati, D. Shah, and M. Sharma. "Max-Product for Maximum Weight Matching: Convergence, Correctness, and LP Duality". In: *IEEE Trans. Inform. Theory* 54.3 (Mar. 2008), pp. 1241–1251. ISSN: 0018-9448. DOI: 10.1109/TIT.2007.915695.

[16] M. Bayati, A. Braunstein, and R. Zecchina. "A Rigorous Analysis of the Cavity Equations for the Minimum Spanning Tree". In: *Journal of Mathematical Physics* 49.12 (Dec. 2008), p. 125206. ISSN: 0022-2488, 1089-7658. DOI: 10.1063/1.2982805.

[17] D. Gamarnik, D. Shah, and Y. Wei. "Belief Propagation for Min-Cost Network Flow: Convergence and Correctness". In: *Operations Research* (Apr. 2012). DOI: 10.1287/opre.1110.1025.

[18] I. Neri and D. Bollé. "The Cavity Approach to Parallel Dynamics of Ising Spins on a Graph". In: *J. Stat. Mech.* 2009.08 (Aug. 2009), P08009. ISSN: 1742-5468. DOI: 10.1088/1742-5468/2009/08/P08009.

[19] Y. Kanoria and A. Montanari. "Majority Dynamics on Trees and the Dynamic Cavity Method". In: *The Annals of Applied Probability* 21.5 (Oct. 2011), pp. 1694–1748. ISSN: 1050-5164, 2168-8737. DOI: 10.1214/10-AAP729.

[20] E. Aurell and H. Mahmoudi. "A Message-Passing Scheme for Non-Equilibrium Stationary States". In: *J. Stat. Mech.* 2011.04 (Apr. 2011), P04014. ISSN: 1742-5468. DOI: 10.1088/1742-5468/2011/04/P04014.

[21] F. Altarelli et al. "Optimizing Spread Dynamics on Graphs by Message Passing". In: *J. Stat. Mech.* 2013.09 (Sept. 2013), P09011. DOI: 10.1088/1742-5468/2013/09/P09011.

[22] G. Del Ferraro and E. Aurell. "Dynamic Message-Passing Approach for Kinetic Spin Models with Reversible Dynamics". In: *Phys. Rev. E* 92.1 (July 2015), p. 010102. DOI: 10.1103/PhysRevE.92.010102.

[23] E. Aurell and H. Mahmoudi. "Dynamic Mean-Field and Cavity Methods for Diluted Ising Systems". In: *Phys. Rev. E* 85.3 (Mar. 2012), p. 031119. DOI: 10.1103/PhysRevE.85.031119.

[24] T. Barthel, C. De Bacco, and S. Franz. "Matrix Product Algorithm for Stochastic Dynamics on Networks Applied to Nonequilibrium Glauber Dynamics". In: *Phys. Rev. E* 97.1 (Jan. 2018), p. 010104. DOI: 10.1103/PhysRevE.97.010104.

[25] T. Barthel. "The Matrix Product Approximation for the Dynamic Cavity Method". In: *J. Stat. Mech.* 2020.1 (Jan. 2020), p. 013217. ISSN: 1742-5468. DOI: 10.1088/1742-5468/ab5701.

[26] F. Altarelli et al. "Large Deviations of Cascade Processes on Graphs". In: *Phys. Rev. E* 87.6 (June 2013), p. 062115. DOI: 10.1103/PhysRevE.87.062115.

[27] F. Altarelli et al. "Bayesian Inference of Epidemics on Networks via Belief Propagation". In: *Phys. Rev. Lett.* 112.11 (Mar. 2014), p. 118701. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.112.118701.

[28] E. Aurell et al. “Cavity Master Equation for the Continuous Time Dynamics of Discrete-Spin Models”. In: *Phys. Rev. E* 95.5 (May 2017), p. 052119. DOI: 10.1103/PhysRevE.95.052119.

[29] A. Braunstein, M. Mézard, and R. Zecchina. “Survey Propagation: An Algorithm for Satisfiability”. In: *Random Struct. Algorithms* 27.2 (Sept. 2005), pp. 201–226. ISSN: 1042-9832. DOI: 10.1002/rsa.20057.

[30] P. Langevin. “Sur La Théorie Du Mouvement Brownien [On the Theory of Brownian Motion]”. In: *C. R. Acad. Sci.* 146 (1908), pp. 530–533.

[31] C. Gardiner. *Stochastic Methods: A Handbook for the Natural and Social Sciences*. 4th ed. Vol. 13. Springer Series in Synergetics. Heidelberg: Springer Berlin, 2009. ISBN: 978-3-540-70712-7.

[32] P. C. Martin, E. D. Siggia, and H. A. Rose. “Statistical Dynamics of Classical Systems”. In: *Phys. Rev. A* 8.1 (July 1973), pp. 423–437. DOI: 10.1103/PhysRevA.8.423.

[33] H.-K. Janssen. “On a Lagrangean for Classical Field Dynamics and Renormalization Group Calculations of Dynamical Critical Properties”. In: *Z Physik B* 23.4 (Dec. 1976), pp. 377–380. DOI: 10.1007/BF01316547.

[34] C. De Dominicis. “Dynamics as a Substitute for Replicas in Systems with Quenched Random Impurities”. In: *Phys. Rev. B* 18.9 (Nov. 1978), pp. 4913–4919. DOI: 10.1103/PhysRevB.18.4913.

[35] C. De Dominicis and L. Peliti. “Field-Theory Renormalization and Critical Dynamics above T c : Helium, Antiferromagnets, and Liquid-Gas Systems”. In: *Phys. Rev. B* 18.1 (July 1978), pp. 353–376. ISSN: 0163-1829. DOI: 10.1103/PhysRevB.18.353.

[36] T. Plefka. “Convergence Condition of the TAP Equation for the Infinite-Ranged Ising Spin Glass Model”. In: *J. Phys. A: Math. Gen.* 15.6 (June 1982), pp. 1971–1978. DOI: 10.1088/0305-4470/15/6/035.

[37] Y. Roudi and J. Hertz. “Dynamical TAP Equations for Non-Equilibrium Ising Spin Glasses”. In: *J. Stat. Mech.* 2011.03 (Mar. 2011), P03031. ISSN: 1742-5468. DOI: 10.1088/1742-5468/2011/03/P03031.

[38] B. Bravi, P. Sollich, and M. Opper. “Extended Plefka Expansion for Stochastic Dynamics”. In: *J. Phys. A: Math. Theor.* 49.19 (May 2016), p. 194003. DOI: 10.1088/1751-8113/49/19/194003.

[39] R. Pastor-Satorras and A. Vespignani. “Epidemic Dynamics and Endemic States in Complex Networks”. In: *Phys. Rev. E* 63.6 (May 2001), p. 066117. ISSN: 1063-651X, 1095-3787. DOI: 10.1103/PhysRevE.63.066117.

[40] R. Pastor-Satorras et al. “Epidemic Processes in Complex Networks”. In: *Rev. Mod. Phys.* 87.3 (Aug. 2015), pp. 925–979. ISSN: 0034-6861, 1539-0756. DOI: 10.1103/RevModPhys.87.925.

[41] F. L. Metz and J. D. Silva. “Spectral Density of Dense Random Networks and the Breakdown of the Wigner Semicircle Law”. In: *Phys. Rev. Research* 2.4 (Oct. 2020), p. 043116. ISSN: 2643-1564. DOI: 10.1103/PhysRevResearch.2.043116.

[42] J. W. Baron. “Eigenvalue Spectra and Stability of Directed Complex Networks”. In: *Phys. Rev. E* 106.6 (Dec. 2022), p. 064302. ISSN: 2470-0045, 2470-0053. DOI: 10.1103/PhysRevE.106.064302.

[43] F. Aguirre-López. "Heterogeneous Mean-Field Analysis of the Generalized Lotka–Volterra Model on a Network". In: *J. Phys. A: Math. Theor.* 57.34 (Aug. 2024), p. 345002. ISSN: 1751-8121. DOI: 10.1088/1751-8121/ad6ab2.

[44] J. I. Park et al. "Incorporating Heterogeneous Interactions for Ecological Biodiversity". In: *Phys. Rev. Lett.* 133.19 (Nov. 2024), p. 198402. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.133.198402.

[45] L. Poley, T. Galla, and J. W. Baron. "Interaction Networks in Persistent Lotka-Volterra Communities". In: *Phys. Rev. E* 111.1 (Jan. 2025), p. 014318. ISSN: 2470-0045, 2470-0053. DOI: 10.1103/PhysRevE.111.014318.

[46] M. Tarabolo and L. Dall'Asta. "Gaussian Approximation of Dynamic Cavity Equations for Linearly-Coupled Stochastic Dynamics". In: *SciPost Phys.* 19.1 (July 2025), p. 019. ISSN: 2542-4653. DOI: 10.21468/SciPostPhys.19.1.019.

[47] M. Tarabolo. *GaussianExpansionCavityMethod.jl*. Version v1.0.0. GitHub repository: GaussianExpansionCavityMethod.jl. 2025.

[48] M. Tarabolo. *RandomLotkaVolterraCavity.jl*. Version v1.0.0-DEV. GitHub repository: RandomLotkaVolterraCavity.jl. 2025.

[49] G. Livan, M. Novaes, and P. Vivo. *Introduction to Random Matrices*. 1st ed. Vol. 26. SpringerBriefs in Mathematical Physics. Cham: Springer International Publishing, 2018. ISBN: 978-3-319-70885-0. DOI: 10.1007/978-3-319-70885-0.

[50] T. Tao. *Topics in Random Matrix Theory*. Vol. 132. Graduate Studies in Mathematics. American Mathematical Society, 2012. DOI: 10.1090/gsm/132.

[51] T. Rogers et al. "Cavity Approach to the Spectral Density of Sparse Symmetric Random Matrices". In: *Phys. Rev. E* 78.3 (Sept. 2008), p. 031116. DOI: 10.1103/PhysRevE.78.031116.

[52] V. A. R. Susca, P. Vivo, and R. Kühn. "Cavity and Replica Methods for the Spectral Density of Sparse Symmetric Random Matrices". In: *SciPost Physics Lecture Notes* (Aug. 2021), p. 033. ISSN: 2590-1990. DOI: 10.21468/SciPostPhysLectNotes.33.

[53] R. Couillet and M. Debbah. *Random Matrix Methods for Wireless Communications*. Cambridge: Cambridge University Press, 2011. DOI: 10.1017/CBO9780511994746.

[54] G. Semerjian and L. F. Cugliandolo. "Sparse Random Matrices: The Eigenvalue Spectrum Revisited". In: *J. Phys. A: Math. Gen.* 35.23 (June 2002), pp. 4837–4851. ISSN: 03054470. DOI: 10.1088/0305-4470/35/23/303.

[55] F. Lucas Metz, I. Neri, and T. Rogers. "Spectral Theory of Sparse Non-Hermitian Random Matrices". In: *J. Phys. A: Math. Theor.* 52.43 (Oct. 2019), p. 434003. DOI: 10.1088/1751-8121/ab1ce0.

[56] W. Cui, J. W. Rocks, and P. Mehta. "An Elementary Mean-Field Approach to the Spectral Densities of Random Matrix Ensembles". In: *Physica A: Statistical Mechanics and its Applications* 637 (Mar. 2024), p. 129608. ISSN: 0378-4371. DOI: 10.1016/j.physa.2024.129608.

[57] H. Kesten. "Symmetric Random Walks on Groups". In: *Transactions of the American Mathematical Society* 92 (1959), pp. 336–354. DOI: 10.1090/S0002-9947-1959-0109367-6.

[58] B. D. McKay. “The Expected Eigenvalue Distribution of a Large Regular Graph”. In: *Linear Algebra and its Applications* 40 (Oct. 1981), pp. 203–216. DOI: 10.1016/0024-3795(81)90150-6.

[59] Y. S. Cho et al. “Percolation Transitions in Scale-Free Networks under the Achlioptas Process”. In: *Phys. Rev. Lett.* 103.13 (Sept. 2009), p. 135702. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.103.135702.

[60] M. Kardar. *Statistical Physics of Fields*. Cambridge: Cambridge University Press, 2007. ISBN: 978-0-521-87341-3. DOI: 10.1017/CBO9780511815881.

[61] U. C. Täuber. *Critical Dynamics: A Field Theory Approach to Equilibrium and Non-Equilibrium Scaling Behavior*. Cambridge: Cambridge University Press, 2014. ISBN: 978-0-521-84223-5. DOI: 10.1017/CBO9781139046213.

[62] W. Horsthemke. “Noise Induced Transitions”. In: *Non-Equilibrium Dynamics in Chemical Systems*. Ed. by C. Vidal and A. Pacault. Berlin, Heidelberg: Springer, 1984, pp. 150–160. ISBN: 978-3-642-70196-2. DOI: 10.1007/978-3-642-70196-2_23.

[63] M. A. Munoz. *Multiplicative Noise in Non-Equilibrium Phase Transitions: A Tutorial*. Comment: Review article to appear in “Advances in Condensed Matter and Statistical Mechanics", Ed. E. Korutcheva and R. Cuerno. To be published by Nova Science Publishers. 43 pages. 9 eps figures. References updated and new figure added. May 2003. DOI: 10.48550/arXiv.cond-mat/0303650.

[64] J.-P. Bouchaud and M. Mézard. “Wealth Condensation in a Simple Model of Economy”. In: *Physica A: Statistical Mechanics and its Applications* 282.3-4 (July 2000), pp. 536–545. DOI: 10.1016/S0378-4371(00)00205-3.

[65] V. Pareto. *Cours d’économie Politique. Œuvres Complètes Publiées Sous La Direction de Giovanni Busino. Tomes 1 et 2 En Un Volume*. Travaux de Sciences Sociales. Genève: Librairie Droz, 1964.

[66] T. Ichinomiya. “Bouchaud-Mézard Model on a Random Network”. In: *Phys. Rev. E* 86.3 (Sept. 2012), p. 036111. DOI: 10.1103/PhysRevE.86.036111.

[67] G. N. Mil’shtejn. “Approximate Integration of Stochastic Differential Equations”. In: *Theory Probab. Appl.* 19.3 (June 1975), pp. 557–562. ISSN: 0040-585X. DOI: 10.1137/1119062.

[68] X. Gabaix and R. Ibragimov. “Rank - 1 / 2: A Simple Way to Improve the OLS Estimation of Tail Exponents”. In: *Journal of Business & Economic Statistics* 29.1 (Jan. 2011), pp. 24–39. ISSN: 0735-0015. DOI: 10.1198/jbes.2009.06157.

[69] T. Castellani and A. Cavagna. “Spin-Glass Theory for Pedestrians”. In: *J. Stat. Mech.* 2005.05 (May 2005), P05012. ISSN: 1742-5468. DOI: 10.1088/1742-5468/2005/05/P05012.

[70] C. De Dominicis and I. Giardina. “The p = 2 Spherical Model”. In: *Random Fields and Spin Glasses: A Field Theory Approach*. Cambridge: Cambridge University Press, 2006, pp. 57–78. ISBN: 978-0-521-84783-4. DOI: 10.1017/CBO9780511534836.005.

[71] L. F. Cugliandolo and D. S. Dean. “Full Dynamical Solution for a Spherical Spin-Glass Model”. In: *J. Phys. A: Math. Gen.* 28.15 (Aug. 1995), p. 4213. ISSN: 0305-4470. DOI: 10.1088/0305-4470/28/15/003.

[72] W. Zippold, R. Kühn, and H. Horner. "Non-Equilibrium Dynamics of Simple Spherical Spin Models". In: *Eur. Phys. J. B* 13.3 (Feb. 2000), pp. 531–537. ISSN: 1434-6036. DOI: 10.1007/s100510050065.

[73] G. Semerjian and L. F. Cugliandolo. "Dynamics of Dilute Disordered Models: A Solvable Case". In: *EPL* 61.2 (Jan. 2003), p. 247. ISSN: 0295-5075. DOI: 10.1209/epl/i2003-00226-8.

[74] R. M. May. "Will a Large Complex System Be Stable?" In: *Nature* 238.5364 (Aug. 1972), pp. 413–414. ISSN: 1476-4687. DOI: 10.1038/238413a0.

[75] R. M. May. *Stability and Complexity in Model Ecosystems*. Princeton University Press, Dec. 2019. ISBN: 978-0-691-20691-2. DOI: 10.1515/9780691206912.

[76] R. May et al., eds. *Theoretical Ecology: Principles and Applications*. Third Edition, Third Edition. Oxford, New York: Oxford University Press, Feb. 2007. ISBN: 978-0-19-920999-6.

[77] G. Bunin. "Ecological Communities with Lotka-Volterra Dynamics". In: *Phys. Rev. E* 95.4 (Apr. 2017), p. 042414. DOI: 10.1103/PhysRevE.95.042414.

[78] G. Biroli, G. Bunin, and C. Cammarota. "Marginally Stable Equilibria in Critical Ecosystems". In: *New J. Phys.* 20.8 (Aug. 2018), p. 083051. ISSN: 1367-2630. DOI: 10.1088/1367-2630/aada58.

[79] A. Altieri et al. "Properties of Equilibria and Glassy Phases of the Random Lotka-Volterra Model with Demographic Noise". In: *Phys. Rev. Lett.* 126.25 (June 2021), p. 258301. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.126.258301.

[80] H. Rieger. "Solvable Model of a Complex Ecosystem with Randomly Interacting Species". In: *J. Phys. A: Math. Gen.* 22.17 (Sept. 1989), p. 3447. ISSN: 0305-4470. DOI: 10.1088/0305-4470/22/17/011.

[81] T. Galla. "Dynamically Evolved Community Size and Stability of Random Lotka-Volterra Ecosystems(a)". In: *EPL* 123.4 (Sept. 2018), p. 48004. ISSN: 0295-5075. DOI: 10.1209/0295-5075/123/48004.

[82] J. A. Dunne, R. J. Williams, and N. D. Martinez. "Network Structure and Biodiversity Loss in Food Webs: Robustness Increases with Connectance". In: *Ecology Letters* 5.4 (2002), pp. 558–567. ISSN: 1461-0248. DOI: 10.1046/j.1461-0248.2002.00354.x.

[83] J. A. Dunne, R. J. Williams, and N. D. Martinez. "Food-Web Structure and Network Theory: The Role of Connectance and Size". In: *Proceedings of the National Academy of Sciences* 99.20 (Oct. 2002), pp. 12917–12922. DOI: 10.1073/pnas.192407699.

[84] P. Jordano, J. Bascompte, and J. M. Olesen. "Invariant Properties in Coevolutionary Networks of Plant–Animal Interactions". In: *Ecology Letters* 6.1 (2003), pp. 69–81. ISSN: 1461-0248. DOI: 10.1046/j.1461-0248.2003.00403.x.

[85] E. Thébault and C. Fontaine. "Stability of Ecological Communities and the Architecture of Mutualistic and Trophic Networks". In: *Science* 329.5993 (Aug. 2010), pp. 853–856. DOI: 10.1126/science.1188321.

[86] B. Pichon et al. "Telling Mutualistic and Antagonistic Ecological Networks Apart by Learning Their Multiscale Structure". In: *Methods in Ecology and Evolution* 15.6 (2024), pp. 1113–1128. ISSN: 2041-210X. DOI: 10.1111/2041-210X.14328.

[87] T. Tonolo et al. *Generalized Lotka-Volterra Model with Sparse Interactions: Non-Gaussian Effects and Topological Multiple-Equilibria Phase*. Mar. 2025. DOI: 10.48550/arXiv.2503.20887.

[88] S. Marcus, A. M. Turner, and G. Bunin. "Local and Extensive Fluctuations in Sparsely Interacting Ecological Communities". In: *Phys. Rev. E* 109.6 (June 2024), p. 064410. ISSN: 2470-0045, 2470-0053. DOI: 10.1103/PhysRevE.109.064410.

[89] S. Marcus, A. M. Turner, and G. Bunin. "Local and Collective Transitions in Sparsely-Interacting Ecological Communities". In: *PLOS Computational Biology* 18.7 (July 2022), e1010274. ISSN: 1553-7358. DOI: 10.1371/journal.pcbi.1010274.

[90] S. Azaele and A. Maritan. "Generalized Dynamical Mean Field Theory for Non-Gaussian Interactions". In: *Phys. Rev. Lett.* 133.12 (Sept. 2024), p. 127401. DOI: 10.1103/PhysRevLett.133.127401.

[91] S. Suweis et al. "Generalized Lotka-Volterra Systems with Time Correlated Stochastic Interactions". In: *Phys. Rev. Lett.* 133.16 (Oct. 2024), p. 167101. DOI: 10.1103/PhysRevLett.133.167101.

[92] F. Roy et al. "Numerical Implementation of Dynamical Mean Field Theory for Disordered Systems: Application to the Lotka–Volterra Model of Ecosystems". In: *J. Phys. A: Math. Theor.* 52.48 (Nov. 2019). Publisher: IOP Publishing, p. 484001. ISSN: 1751-8121. DOI: 10.1088/1751-8121/ab1f32.

[93] A. Braunstein et al. "Small-Coupling Dynamic Cavity: A Bayesian Mean-Field Framework for Epidemic Inference". In: *Phys. Rev. Res.* 7.2 (Apr. 2025), p. 023089. DOI: 10.1103/PhysRevResearch.7.023089.

[94] M. Tarabolo. *SmallCouplingDynamicCavity.jl*. Version v4.0.2. Julia package: SmallCouplingDynamicCavity.jl. 2024.

[95] A. Y. Lokhov et al. "Inferring the Origin of an Epidemic with a Dynamic Message-Passing Algorithm". In: *Phys. Rev. E* 90.1 (July 2014), p. 012801. ISSN: 1539-3755, 1550-2376. DOI: 10.1103/PhysRevE.90.012801.

[96] A. Y. Lokhov, M. Mézard, and L. Zdeborová. "Dynamic Message-Passing Equations for Models with Unidirectional Dynamics". In: *Phys. Rev. E* 91.1 (Jan. 2015), p. 012811. DOI: 10.1103/PhysRevE.91.012811.

[97] A. Baker et al. "Epidemic Mitigation by Statistical Inference from Contact Tracing Data". In: *Proceedings of the National Academy of Sciences* 118.32 (2021), e2106548118. DOI: 10.1073/pnas.2106548118.

[98] F. Altarelli et al. "The Patient-Zero Problem with Noisy Observations". In: *J. Stat. Mech.* 2014.10 (Oct. 2014), P10016. ISSN: 1742-5468. DOI: 10.1088/1742-5468/2014/10/P10016.

[99] F. Mazza and Indaco Biazzo. *Sibyl-Team/Epigen: Release Epigen*. Zenodo. Apr. 2023. DOI: 10.5281/ZENODO.7852232.

[100] D. J. Watts and S. H. Strogatz. "Collective Dynamics of 'Small-World' Networks". In: *Nature* 393.6684 (1998), pp. 440–442. ISSN: 1476-4687. DOI: 10.1038/30918.

[101] M. D. Penrose. "Connectivity of Soft Random Geometric Graphs". In: *The Annals of Applied Probability* 26.2 (2016), pp. 986–1028. ISSN: 1050-5164, 2168-8737. DOI: 10.1214/15-AAP1110.

[102] R. Hinch et al. “OpenABM-Covid19—An Agent-Based Model for Non-Pharmaceutical Interventions against COVID-19 Including Contact Tracing”. In: *PLOS Computational Biology* 17.7 (2021), e1009146. ISSN: 1553-7358. DOI: 10.1371/journal.pcbi.1009146.

[103] C. C. Kerr et al. “Covasim: An Agent-Based Model of COVID-19 Dynamics and Interventions”. In: *PLOS Computational Biology* 17.7 (2021), e1009149. ISSN: 1553-7358. DOI: 10.1371/journal.pcbi.1009149.

[104] M. Génois and A. Barrat. “Can Co-Location Be Used as a Proxy for Face-to-Face Contacts?” In: *EPJ Data Science* 7.1 (May 2018), p. 11. ISSN: 2193-1127. DOI: 10.1140/epjds/s13688-018-0140-1.

[105] S. Crotti and A. Braunstein. “Matrix Product Belief Propagation for Reweighted Stochastic Dynamics over Graphs”. In: *Proceedings of the National Academy of Sciences* 120.47 (Nov. 2023), e2307935120. DOI: 10.1073/pnas.2307935120.

[106] D. J. Daley and D. G. Kendall. “Epidemics and Rumours”. In: *Nature* 204.4963 (Dec. 1964), pp. 1118–1118. ISSN: 1476-4687. DOI: 10.1038/2041118a0.

[107] S. Boccaletti et al. “Complex Networks: Structure and Dynamics”. In: *Physics Reports* 424.4 (2006), pp. 175–308. ISSN: 0370-1573. DOI: 10.1016/j.physrep.2005.10.009.

[108] T. Ichinomiya. “Wealth Distribution on Complex Networks”. In: *Phys. Rev. E* 86.6 (Dec. 2012), p. 066115. DOI: 10.1103/PhysRevE.86.066115.

[109] T. Ichinomiya. “Power-Law Exponent of the Bouchaud-Mézard Model on Regular Random Networks”. In: *Phys. Rev. E* 88.1 (July 2013), p. 012819. DOI: 10.1103/PhysRevE.88.012819.

Appendix A

Detailed derivations for the path-integral techniques

A.1 MSRJD path-integral construction

This section collects the detailed calculations underlying the compact derivations in section 3.2. We focus on the Ornstein-Uhlenbeck (OU) process as a guiding example, since its dynamics is simple enough to allow exact evaluation of path integrals while still illustrating the key ideas.

A.1.1 Path integrals as limits of discrete products

In the physicists' approach, a path integral is defined as the continuum limit of a product of ordinary integrals over discretized time steps. Concretely, if time is discretized as $t_n = n\Delta$ with $n = 0, \dots, T$ and $T\Delta = t_f$, a trajectory $x(t)$ is represented by $T+1$ integration variables. An integral over all paths is then

$$\int \mathscr{D}[x] \cdots = \lim_{\Delta \to 0} \prod_{n=0}^{T} \int dx^n \cdots , \tag{A.1}$$

where the dots denote the weight functional defined by the stochastic dynamics. Functional delta functions, such as $\delta[f(x(t))]$, are similarly defined as the limit of products of ordinary delta functions enforcing the discretized constraints. This construction gives path integrals their meaning without resorting to measure theory: they are limits of ordinary integrals over finely discretized trajectories.

A.1.2 Discrete-time MSRJD action

We start from the OU dynamics,

$$\frac{dx(t)}{dt} = -\lambda x(t) + \xi(t), \qquad \langle \xi(t) \rangle = 0, \qquad \langle \xi(t)\xi(t') \rangle = 2T\delta(t-t'). \tag{A.2}$$

In discrete time $t_n = n\Delta$, the Itô discretization gives

$$x^{n+1} = x^n - \Delta\lambda x^n + \Delta\xi^n, \tag{A.3}$$

where the integrated noise increments $\Delta\xi^n = \int_{n\Delta}^{(n+1)\Delta} \xi(t)\,dt$ are Gaussian with variance $\langle (\Delta\xi^n)^2 \rangle = 2T\Delta$. The noise probability weight is

$$P[\boldsymbol{\Delta\xi}] \propto \prod_{n=0}^{T} \exp\left[-\frac{(\Delta\xi^n)^2}{4T\Delta} \right]. \tag{A.4}$$

The dynamics imposes constraints

$$x^{n+1} - x^n + \Delta\lambda x^n - \Delta\xi^n = 0, \tag{A.5}$$

which are enforced via delta functions. The Martin-Siggia-Rose-Janssen-De Dominicis (MSRJD) path-integral representation of the average of an observable $\mathscr{O}(\boldsymbol{x})$ is then

$$\langle \mathscr{O}(\boldsymbol{x}) \rangle = \int d\boldsymbol{x}\, \mathscr{O}(\boldsymbol{x}) \prod_{n=0}^{T-1} \delta\left(x^{n+1} - x^n + \Delta\lambda x^n - \Delta\xi^n \right) \tag{A.6a}$$

$$\propto \int d\boldsymbol{x}\, d\hat{\boldsymbol{x}}\, d\boldsymbol{\Delta\xi}\, P[\boldsymbol{\Delta\xi}]\, \mathscr{O}(\boldsymbol{x}) \prod_{n=0}^{T-1} e^{-\mathrm{i}\hat{x}^n (x^{n+1} - x^n + \Delta\lambda x^n - \Delta\xi^n)}, \tag{A.6b}$$

where we made use of the Fourier representation $\delta(y) \propto \int d\hat{x}\, e^{-\mathrm{i}\hat{x}y}$ with conjugate fields $\hat{x}^n$. Integrating over the Gaussian noise variables gives

$$\langle \mathscr{O}(\boldsymbol{x}) \rangle \propto \int d\boldsymbol{x}\, d\hat{\boldsymbol{x}}\, \mathscr{O}(\boldsymbol{x}) \prod_{n=0}^{T-1} e^{-\mathrm{i}\hat{x}^n (x^{n+1} - x^n + \Delta\lambda x^n)} \int d\Delta\xi^n e^{-\frac{(\Delta\xi^n)^2}{4T\Delta} - \mathrm{i}\hat{x}^n \Delta\xi^n} \tag{A.7a}$$

$$\propto \int d\boldsymbol{x}\, d\hat{\boldsymbol{x}}\, \mathscr{O}(\boldsymbol{x}) \prod_{n=0}^{T-1} e^{-\mathrm{i}\hat{x}^n (x^{n+1} - x^n + \Delta\lambda x^n) - T\Delta(\hat{x}^n)^2} \tag{A.7b}$$

$$\propto \int d\boldsymbol{x}\, d\hat{\boldsymbol{x}}\, \mathscr{O}(\boldsymbol{x})\, e^{-\sum_{n=0}^{T-1} \left[\mathrm{i}\hat{x}^n (x^{n+1} - x^n + \Delta\lambda x^n) + T\Delta(\hat{x}^n)^2 \right]}, \tag{A.7c}$$

where we used the Gaussian integral $\int dy\, e^{-ay^2+by} \propto e^{b^2/(4a)}$. The discrete-time MSRJD action is therefore

$$S(\boldsymbol{x},\hat{\boldsymbol{x}}) = \sum_{n=0}^{T-1} \Big[i\hat{x}^n \big(x^{n+1} - x^n + \Delta\lambda x^n \big) + \Delta T (\hat{x}^n)^2 \Big]. \tag{A.8}$$

In the continuum limit $\Delta \to 0$ the sums $\Delta \sum_{n=0}^{T-1}$ become integrals $\int_0^{t_f} dt$, and the finite differences $(x^{n+1}-x^n)/\Delta$ converge to time derivatives $\dot{x}(t)$, yielding the continuous-time action

$$S[x,\hat{x}] = \int_0^{t_f} dt \, \Big[i\hat{x}(t) \big(\dot{x}(t) + \lambda x(t) \big) + T\hat{x}^2(t) \Big]. \tag{A.9}$$

A.1.3 Exact generating functional

To illustrate the power of the MSRJD representation, consider the generating functional

$$Z[\psi] = \left\langle e^{\int dt\, \psi(t) x(t)} \right\rangle, \tag{A.10}$$

where $\psi(t)$ is an external source. Inserting this into the path integral Eq. (3.10), assuming uniform initial conditions $p_0(x) \propto 1$, gives

$$\begin{aligned} Z[\psi] = \int \mathscr{D}[x,\hat{x}] \exp \Big\{ & - \int dt\, i\hat{x}(t) \big(\dot{x}(t) + \lambda x(t) \big) \\ & - T \int dt\, \hat{x}^2(t) + \int dt\, \psi(t) x(t) \Big\}. \end{aligned} \tag{A.11}$$

By integrating by parts the term $\int dt\, i\hat{x}(t)\dot{x}(t)$ and assuming boundary terms vanish, this can be rewritten as

$$Z[\psi] = \int \mathscr{D}[x,\hat{x}] \exp \left\{ \int dt\, \mathrm{i} x(t) \big(\dot{\hat{x}}(t) - \lambda \hat{x}(t) - \mathrm{i}\psi(t) \big) - T \int dt\, \hat{x}^2(t) \right\}. \tag{A.12}$$

Since the exponent is linear in $x(t)$, we can integrate over x exactly. This is equivalent to applying the Fourier representation of the delta functional in reverse:

$$Z[\psi] \propto \int \mathscr{D}[\hat{x}] \exp \left(-T \int dt\, \hat{x}^2(t) \right) \delta \Big[\dot{\hat{x}}(t) - \lambda \hat{x}(t) - \mathrm{i}\psi(t) \Big]. \tag{A.13}$$

The delta functional enforces the equation of motion for $\hat{x}(t)$

$$\frac{d\hat{x}(t)}{dt} = \lambda \hat{x}(t) + \mathrm{i}\psi(t), \tag{A.14}$$

Integrating this gives

$$\hat{x}(t) = -i \int_t^\infty ds\, e^{\lambda(t-s)} \psi(s), \tag{A.15}$$

from which we can solve the integral

$$-T \int dt\, \hat{x}^2(t) = T \int dt \int_t^\infty ds_1\, ds_2\, e^{\lambda(2t-s_1-s_2)} \psi(s_1)\psi(s_2) \tag{A.16}$$

$$= T \int_{-\infty}^\infty ds_1\, ds_2\, e^{-\lambda(s_1+s_2)} \psi(s_1)\psi(s_2) \int_{-\infty}^{\min(s_1,s_2)} dt\, e^{2\lambda t} \tag{A.17}$$

$$= \frac{T}{2\lambda} \int_{-\infty}^\infty ds_1\, ds_2\, e^{-\lambda|s_1-s_2|} \psi(s_1)\psi(s_2). \tag{A.18}$$

Substituting back, the generating functional becomes

$$Z[\psi] = \exp\left[\frac{T}{2\lambda} \int_{-\infty}^\infty ds_1 \int_{-\infty}^\infty ds_2\, e^{-\lambda|s_1-s_2|}\, \psi(s_1)\psi(s_2)\right], \tag{A.19}$$

where the overall normalization is fixed by requiring $Z[0] = 1$.

A.1.4 Correlation functions

Correlation functions follow from functional derivatives of $\ln Z[\psi]$. In particular,

$$C^{\mathrm{dc}}(t,t') = \langle x(t)x(t')\rangle = \frac{\delta^2}{\delta\psi(t)\delta\psi(t')} \ln Z[\psi]\bigg|_{\psi=0}. \tag{A.20}$$

Evaluating with the explicit form of $Z[\psi]$ given by Eq. (A.19) yields

$$C^{\mathrm{dc}}(t,t') = \frac{T}{2\lambda}\, e^{-\lambda|t-t'|}, \tag{A.21}$$

which is the familiar exponentially decaying correlation of the OU process.

A.1.5 Response functions

To derive response functions, perturb the Stochastic Differential Equation (SDE) with a source $h(t)$:

$$\frac{dx(t)}{dt} = -\psi x(t) + \xi(t) + h(t). \tag{A.22}$$

In the generating functional this contributes a term $\int dt\, \mathrm{i}\hat{x}(t)h(t)$ into the exponent

$$Z[\psi;h] \propto \int \mathscr{D}[\hat{x}] \exp\left(-T\int dt\, \hat{x}^2(t) + \mathrm{i}\int dt\, h(t)\hat{x}(t)\right) \times \delta\left[\dot{\hat{x}}(t) - \lambda\hat{x}(t) - \mathrm{i}\psi(t)\right] \tag{A.23a}$$

$$\propto \exp\left[\frac{T}{2\lambda}\int_{-\infty}^{\infty} ds_1 \int_{-\infty}^{\infty} ds_2\, e^{-\lambda|s_1-s_2|}\,\psi(s_1)\psi(s_2) + \int_{-\infty}^{\infty} ds_1\, e^{-\lambda s_1}\psi(s_1)\int_{-\infty}^{s_1} ds_2\, e^{\lambda s_2}h(s_2)\right] \tag{A.23b}$$

Differentiating with respect to $h(t')$ the average of $x(t)$ gives the linear response function

$$R(t,t') = \frac{\delta}{\delta h(t')}\langle x(t)\rangle\Big|_{h=0} = \frac{\delta^2}{\delta h(t')\delta\psi(t)} Z[\psi;h]\Big|_{\substack{h=0\\ \psi=0}} \tag{A.24a}$$

$$= \frac{\delta}{\delta h(t')} e^{-\lambda t}\int_{-\infty}^{t} ds\, e^{\lambda s}h(s)\Big|_{h=0} = \Theta(t-t')\, e^{-\lambda(t-t')}. \tag{A.24b}$$

This is the expected causal response function of the OU process Eq. (3.19), which vanishes for $t < t'$.

Equivalently, the response function can be computed by making use of the response field $\hat{x}(t)$. By defining the generating functional with sources $\psi(t)$ and $\hat{\psi}(t)$ coupled to $x(t)$ and $\hat{x}(t)$, respectively, we have Differentiating with respect to $h(t')$ the average of $x(t)$ gives the linear response function

$$Z[\psi,\hat{\psi}] = \left\langle e^{\int dt\,\left(\psi(t)x(t)+\hat{\psi}(t)\hat{x}(t)\right)}\right\rangle \tag{A.25a}$$

$$= \int \mathscr{D}[\hat{x}] \exp\left(-T\int dt\, \hat{x}^2(t) + \int dt\, \hat{\psi}(t)\hat{x}(t)\right) \times \delta\left[\dot{\hat{x}}(t) - \lambda\hat{x}(t) - \mathrm{i}\psi(t)\right] \tag{A.25b}$$

$$\propto \exp\left[\frac{T}{2\lambda}\int_{-\infty}^{\infty} ds_1 \int_{-\infty}^{\infty} ds_2\, e^{-\lambda|s_1-s_2|}\,\psi(s_1)\psi(s_2) - \mathrm{i}\int_{-\infty}^{\infty} ds_1\, e^{-\lambda s_1}\psi(s_1)\int_{-\infty}^{s_1} ds_2\, e^{\lambda s_2}\hat{\psi}(s_2)\right]. \tag{A.25c}$$

The response function is then Differentiating with respect to $h(t')$ the average of $x(t)$ gives the linear response function

$$R(t,t') = \frac{\delta}{\delta h(t')}\langle x(t)\rangle\Big|_{h=0} = \langle x(t)\mathrm{i}\hat{x}(t')\rangle = \mathrm{i}\frac{\delta^2}{\delta\hat{\psi}(t')\delta\psi(t)}Z[\psi,\hat{\psi}]\Big|_{\substack{\hat{\psi}=0\\ \psi=0}} \tag{A.26a}$$

$$= \mathrm{i}\,\frac{\delta}{\delta\hat{\psi}(t')}(-\mathrm{i})e^{-\lambda t}\int_{-\infty}^{t} ds\, e^{\lambda s}\hat{\psi}(s)\Big|_{\hat{\psi}=0} = \Theta(t-t')e^{-\lambda(t-t')}. \tag{A.26b}$$

A.1.6 Vanishing of response-field averages

We prove that all averages involving only response fields vanish identically in the MSRJD formalism. We consider a general SDE of the form

$$\dot{x}(t) = f(x(t)) + \xi(t) + h(t), \qquad \langle\xi(t)\xi(t')\rangle = 2g(x(t))\delta(t-t'), \tag{A.27}$$

where $h(t)$ is an external field. After averaging over the Gaussian noise, the action takes the form

$$S[x,\hat{x}] = \int_0^{t_f} dt\, \left[\mathrm{i}\hat{x}(t)\left(\dot{x}(t) - f(x(t)) - h(t)\right) + g(x(t))\hat{x}(t)^2\right]. \tag{A.28}$$

The path probability is normalized to one, independently of the external field $h(t)$

$$\int \mathscr{D}[x]\,P[x] = \int \mathscr{D}[x,\hat{x}]e^{-S[x,\hat{x}]} = 1. \tag{A.29}$$

It follows that any functional derivative with respect to h must vanish, i.e.

$$\langle\mathrm{i}\hat{x}(t_1)\cdots\mathrm{i}\hat{x}(t_n)\rangle = \frac{\delta^n}{\delta h(t_1)\cdots\delta h(t_n)}\int \mathscr{D}[x,\hat{x}]e^{-S[x,\hat{x}]}\Big|_{h=0} \tag{A.30}$$

$$= \frac{\delta^n}{\delta h(t_1)\cdots\delta h(t_n)}1 = 0. \tag{A.31}$$

A.2 Detailed derivations of DMFT

This appendix collects the detailed calculations underlying the compact derivations in section 3.3. We focus on the fully-connected coupled OU process as a guiding example, since its dynamics is simple enough to allow exact evaluation of path integrals while still illustrating the key ideas.

A.2.1 Disorder average

We write the couplings as

$$J_{ij} = \frac{m}{N} + \frac{\sigma}{\sqrt{N}} z_{ij} \tag{A.32}$$

where z_{ij} are zero-mean, unit-variance Gaussian variables with correlation $\overline{z_{ij} z_{ji}} = \gamma$. For shortness of notation we define

$$A_{ij} = \int dt \, \mathrm{i}\hat{x}_i(t) x_j(t).$$

The average over the disorder of the interaction part is therefore

$$\overline{e^{-S_J^{\mathrm{int}}}} = \overline{e^{\sum_{i\neq j} J_{ij} A_{ij}}} = e^{\frac{m}{N}\sum_{i\neq j} A_{ij}} \prod_{i<j} \overline{e^{\frac{\sigma}{\sqrt{N}}(z_{ij}A_{ij} + z_{ji}A_{ji})}} \tag{A.33}$$

$$= e^{\frac{m}{N}\sum_{i\neq j} A_{ij}} \prod_{i<j} e^{\frac{\sigma^2}{2N}(A_{ij}^2 + 2\gamma A_{ij}A_{ji} + A_{ji}^2)} \tag{A.34}$$

$$= e^{\frac{m}{N}\sum_{i\neq j} A_{ij} + \frac{\sigma^2}{2N}\sum_{i\neq j}(A_{ij}^2 + \gamma A_{ij}A_{ji})}. \tag{A.35}$$

Reinstating the time integrals gives the expression reported in the main text, Eq. (3.30).

A.2.2 Order parameters

To decouple the sites, we introduce the order parameters

$$\mu(t) = \frac{1}{N}\sum_i x_i(t), \qquad P(t) = \frac{1}{N}\sum_i \mathrm{i}\hat{x}_i(t),$$

$$C(t,t') = \frac{1}{N}\sum_i x_i(t)x_i(t'), \qquad R(t,t') = \frac{1}{N}\sum_i x_i(t)\mathrm{i}\hat{x}_i(t'),$$

$$B(t,t') = \frac{1}{N}\sum_i \mathrm{i}\hat{x}_i(t)\mathrm{i}\hat{x}_i(t').$$

and enforce their definitions via delta functionals. For example,

$$1 = \int \mathscr{D}[\mu]\, \delta\left[\mu(t) - \frac{1}{N}\sum_{i=1}^{N} x_i(t)\right]$$

$$\propto \int \mathscr{D}[\mu,\hat{\mu}] \exp\left\{-\mathrm{i}\int dt\, \hat{\mu}(t)\Big(N\mu(t) - \sum_i x_i(t)\Big)\right\},$$

where $\hat{\mu}(t)$ is a conjugate field. For shortness of notation, we define the set of order parameters as $W = \{\mu, P, C, R, B\}$ and the corresponding conjugate fields as $\hat{W} = \{\hat{\mu}, \hat{P}, \hat{C}, \hat{R}, \hat{B}\}$. Repeating for all order parameters and substituting into Eq. (3.26) gives

$$Z[\underline{\psi}, \underline{\hat{\psi}}] = \int \mathscr{D}^0[\underline{x}, \underline{\hat{x}}] \prod_{w \in W} \left(\mathscr{D}[w]\, \delta\left[\text{def. of } w\right]\right) e^{-S^{\text{loc}}[\underline{x}, \underline{\hat{x}}, \underline{\psi}, \underline{\hat{\psi}}] + \frac{m}{N} \sum_{i \neq j} \int dt\, \mathrm{i}\hat{x}_i(t) x_j(t)} \tag{A.36}$$

$$\times\, e^{\frac{\sigma^2}{2N} \int dt\, dt' \sum_{i \neq j} \left[\mathrm{i}\hat{x}_i(t) \mathrm{i}\hat{x}_i(t') x_j(t) x_j(t') + \gamma \mathrm{i}\hat{x}_i(t) x_i(t') x_j(t) \mathrm{i}\hat{x}_j(t')\right]}. \tag{A.37}$$

By using the properties of the delta functional

$$\int \mathscr{D}[f]\, \delta[f(t) - g(t)] f(t) = \int \mathscr{D}[f]\, \delta[f(t) - g(t)] g(t),$$

we can rewrite the sums over $i \neq j$ in terms of the order parameters. At leading order in N we notice that $\sum_{i \neq j} f_{ij} = \sum_{i,j} f_{ij} - \sum_i f_{ii} \sim \sum_{i,j} f_{ij}$, so that

$$Z[\underline{\psi}, \underline{\hat{\psi}}] \sim \int \mathscr{D}^0[\underline{x}, \underline{\hat{x}}] \prod_{w \in W} \left(\mathscr{D}[w]\, \delta\left[\text{def. of } w\right]\right) e^{-S^{\text{loc}}[\underline{x}, \underline{\hat{x}}, \underline{\psi}, \underline{\hat{\psi}}]} \tag{A.38}$$

$$\times\, e^{Nm \int dt\, P(t) \mu(t) + \frac{N\sigma^2}{2} \int dt\, dt' [B(t,t') C(t,t') + \gamma R(t,t') R(t',t)]}. \tag{A.39}$$

Introducing the conjugate fields and defininining $Q = W, \hat{W}$ gives

$$Z[\underline{\psi}, \underline{\hat{\psi}}] \sim \int \mathscr{D}[Q]\, e^{-N(S^{\text{o.p.}}[Q] + S^{\text{avg}}[Q])} \int \mathscr{D}^0[\underline{x}, \underline{\hat{x}}]\, e^{-S^{\text{loc}}[\underline{x}, \underline{\hat{x}}, \underline{\psi}, \underline{\hat{\psi}}]} \tag{A.40}$$

$$\times\, e^{\mathrm{i} \sum_{i=1}^N \int dt \left[\hat{\mu}(t) x_i(t) + \hat{P}(t) \mathrm{i}\hat{x}_i(t)\right]}$$

$$\times\, e^{\mathrm{i} \sum_{i=1}^N \int dt\, dt' \left[\hat{C}(t,t') x_i(t) x_i(t') + \hat{R}(t,t') x_i(t) \mathrm{i}\hat{x}_i(t') + \hat{B}(t,t') \mathrm{i}\hat{x}_i(t) \mathrm{i}\hat{x}_i(t')\right]} \tag{A.41}$$

$$= \int \mathscr{D}[Q]\, e^{-N(S^{\text{o.p.}}[Q] + S^{\text{avg}}[Q])}$$

$$\times\, e^{\sum_{i=1}^N \log \int \mathscr{D}^0[x_i, \hat{x}_i]\, \exp\{-S_1[x_i, \hat{x}_i, \psi_i, \hat{\psi}_i | Q]\}} \tag{A.42}$$

where $S^{\text{o.p.}}[Q]$ is the action enforcing the definitions of the order parameters and their conjugate fields, $S^{\text{avg}}[Q]$ is the action averaging over the disorder, and $S_1[\underline{x}, \underline{\hat{x}}, \underline{\psi}, \underline{\hat{\psi}}]$ is the action for the local effective theory, as defined in Eqs. (3.37) to (3.39).

A.2.3 Saddle-point equations

To leading order in N, the disorder-averaged generating functional Eq. (3.36) can be written as

$$Z[\underline{\psi},\underline{\hat{\psi}}] \sim \int \mathscr{D}[Q]\, e^{-N\left(S^{\text{o.p.}}[Q]+S^{\text{avg}}[Q]\right)} \left(Z^{\text{loc}}_{\text{eff}}[\underline{\psi},\underline{\hat{\psi}} \mid Q]\right)^N, \tag{A.43}$$

where $Z^{\text{loc}}_{\text{eff}}[\underline{\psi},\underline{\hat{\psi}} \mid Q]$ is the local effective partition function depending on the order parameters and their conjugate fields

$$\log Z^{\text{loc}}_{\text{eff}}[\underline{\psi},\underline{\hat{\psi}} \mid Q] = \frac{1}{N}\sum_{i=1}^{N} \log Z_1[\psi_i,\hat{\psi}_i \mid Q]. \tag{A.44}$$

In the thermodynamic limit $N \to \infty$, the integral over Q is dominated by the saddle point $Q^\star$ minimizing the exponent. The saddle-point equations are obtained by setting to zero the functional derivatives of the exponent with respect to each order parameter and conjugate field. The functional derivatives with respect to the order parameters give

$$\begin{aligned}
\frac{\delta}{\delta\mu(t)}(S^{\text{o.p.}}+S^{\text{avg}}) = 0, \quad &\Rightarrow \quad \mathrm{i}\hat{\mu}^\star(t) = -mP^\star(t),\\
\frac{\delta}{\delta P(t)}(S^{\text{o.p.}}+S^{\text{avg}}) = 0, \quad &\Rightarrow \quad \mathrm{i}\hat{P}^\star(t) = -m\mu^\star(t),\\
\frac{\delta}{\delta C(t,t')}(S^{\text{o.p.}}+S^{\text{avg}}) = 0, \quad &\Rightarrow \quad \mathrm{i}\hat{C}^\star(t,t') = -\frac{\sigma^2}{2}B^\star(t,t'),\\
\frac{\delta}{\delta R(t,t')}(S^{\text{o.p.}}+S^{\text{avg}}) = 0, \quad &\Rightarrow \quad \mathrm{i}\hat{R}^\star(t,t') = -\sigma^2\gamma R^\star(t',t),\\
\frac{\delta}{\delta B(t,t')}(S^{\text{o.p.}}+S^{\text{avg}}) = 0, \quad &\Rightarrow \quad \mathrm{i}\hat{B}^\star(t,t') = -\frac{\sigma^2}{2}C^\star(t,t').
\end{aligned}$$

The functional derivatives with respect to the conjugate fields give

$$\begin{aligned}
\frac{\delta}{\delta\hat{\mu}(t)}\left(S^{\text{o.p.}}-\log Z_{\text{eff}}^{\text{loc}}\right)=0, \qquad &\Rightarrow \qquad \mu^{\star}(t)=\langle x(t)\rangle_1^{\star},\\
\frac{\delta}{\delta\hat{P}(t)}\left(S^{\text{o.p.}}-\log Z_{\text{eff}}^{\text{loc}}\right)=0, \qquad &\Rightarrow \qquad P^{\star}(t)=\langle \mathrm{i}\hat{x}(t)\rangle_1^{\star},\\
\frac{\delta}{\delta\hat{C}(t,t')}\left(S^{\text{o.p.}}-\log Z_{\text{eff}}^{\text{loc}}\right)=0, \qquad &\Rightarrow \qquad C^{\star}(t,t')=\langle x(t)x(t')\rangle_1^{\star},\\
\frac{\delta}{\delta\hat{R}(t,t')}\left(S^{\text{o.p.}}-\log Z_{\text{eff}}^{\text{loc}}\right)=0, \qquad &\Rightarrow \qquad R^{\star}(t,t')=\langle x(t)\mathrm{i}\hat{x}(t')\rangle_1^{\star},\\
\frac{\delta}{\delta\hat{B}(t,t')}\left(S^{\text{o.p.}}-\log Z_{\text{eff}}^{\text{loc}}\right)=0, \qquad &\Rightarrow \qquad B^{\star}(t,t')=\langle \mathrm{i}\hat{x}(t)\mathrm{i}\hat{x}(t')\rangle_1^{\star},
\end{aligned}$$

where $\langle\cdots\rangle_1^{\star}$ denotes averages with respect to the local effective process with generating functional Z_1 evaluated at the saddle point $Q^{\star}$. The terms $P^{\star}(t)$ and $B^{\star}(t,t')$ are zero, since they involve averages of the response field $\hat{x}(t)$ alone. Substituting back the saddle-point values of the conjugate fields into the local effective action $S_1[x,\hat{x},0,0\,|\,Q]$ gives the effective action reported in Eq. (3.45).

Appendix B

Detailed derivations for the Gaussian-Expansion Cavity Method

B.1 Construction of the Gaussian-Expansion Cavity Method

This section provides detailed derivations for the Gaussian Expansion Cavity Method (GECaM) introduced in Chapter 4.

B.1.1 Derivation of the MSRJD action

The generating functional for the coupled Ornstein-Uhlenbeck (OU) processes defined in Eq. (4.1) is given by

$$Z[\underline{\psi},\underline{\hat{\psi}}] = \left\langle \int \mathscr{D}^0[\underline{x}] \prod_{i=1}^{N} e^{\int dt(\psi_i(t)x_i(t)+\hat{\psi}_i(t)\mathrm{i}\hat{x}_i(t))} \delta\left[\dot{x}_i - \lambda_i x_i + \alpha \sum_{j\in\partial i} J_{ij}x_j + \xi_i\right] \right\rangle \tag{B.1a}$$

$$\propto \int \mathscr{D}^0[\underline{x},\underline{\hat{x}}] \prod_{i=1}^{N} e^{\int dt[(\psi_i(t)x_i(t)+\hat{\psi}_i(t)\hat{x}_i(t))-\mathrm{i}\hat{x}_i(t)(\dot{x}_i(t)-\lambda_i x_i(t))]} \left\langle e^{-\int dt \mathrm{i}\hat{x}_i(t)\xi_i(t)} \right\rangle_i$$

$$\times \prod_{i=1}^{N}\prod_{j\in\partial i} e^{-\int dt \mathrm{i}\hat{x}_i(t)J_{ij}x_j(t)} \tag{B.1b}$$

where we used the integral representation of the delta functional and introduced auxiliary response fields $\hat{x}_i(t)$. The Gaussian average over the Gaussian white noise $\xi_i(t)$ has been

denoted by $\langle\cdots\rangle_i$. Using the properties of the Gaussian noise, we have

$$Z[\underline{\psi},\underline{\hat{\psi}}] \propto \int \mathscr{D}^0[\underline{x},\underline{\hat{x}}] \prod_{i=1}^{N} e^{\int dt\left[(\psi_i(t)x_i(t)+\hat{\psi}_i(t)\hat{x}_i(t))-\mathrm{i}\hat{x}_i(t)(\dot{x}_i(t)-\lambda_i x_i(t))+T\hat{x}_i(t)^2\right]} \times \prod_{(i,j)\in E} e^{-\int dt \mathrm{i}\left(\hat{x}_i(t)J_{ij}x_j(t)+\hat{x}_j(t)J_{ji}x_i(t)\right)}, \tag{B.2a}$$

$$= \int \mathscr{D}^0[\underline{x},\underline{\hat{x}}] e^{-\sum_i S_i^{\mathrm{loc}}[x_i,\hat{x}_i]-\sum_{(i,j)} S_{ij}^{\mathrm{int}}[x_i,\hat{x}_i,x_j,\hat{x}_j]+\int dt \sum_{i=1}^{N}(\psi_i(t)x_i(t)+\hat{\psi}_i(t)\hat{x}_i(t))}, \tag{B.2b}$$

Where we have defined the local action $S_i^{\mathrm{loc}}[x_i,\hat{x}_i]$ as

$$S_i^{\mathrm{loc}}[x_i,\hat{x}_i] = \int dt \left[\mathrm{i}\hat{x}_i(t)\left(\dot{x}_i(t)-\lambda_i x_i(t)\right)-T\hat{x}_i(t)^2\right], \tag{B.3}$$

and the interaction action $S_{ij}^{\mathrm{int}}[x_i,\hat{x}_i,x_j,\hat{x}_j]$ as

$$S_{ij}^{\mathrm{int}}[x_i,\hat{x}_i,x_j,\hat{x}_j] = \int dt \left(\mathrm{i}\hat{x}_i(t)J_{ij}x_j(t)+\mathrm{i}\hat{x}_j(t)J_{ji}x_i(t)\right). \tag{B.4}$$

The generating functional can therefore be written as

$$Z[\underline{\psi},\underline{\hat{\psi}}] = \int \mathscr{D}^0[\underline{x},\underline{\hat{x}}] e^{-S[\underline{x},\underline{\hat{x}}]+\int dt \sum_{i=1}^{N}(\psi_i(t)x_i(t)+\hat{\psi}_i(t)\hat{x}_i(t))}, \tag{B.5}$$

where the total action $S[\underline{x},\underline{\hat{x}}]$ is given by

$$S[\underline{x},\underline{\hat{x}}] = \sum_{i=1}^{N} S_i^{\mathrm{loc}}[x_i,\hat{x}_i] + \sum_{(i,j)\in E} S_{ij}^{\mathrm{int}}[x_i,\hat{x}_i,x_j,\hat{x}_j]. \tag{B.6}$$

B.1.2 Small-coupling expansion of the dynamic cavity equations

The dynamic cavity equations for the coupled OU processes defined in Eqs. (4.7) and (4.9) are of limited practical use, as they involve functional distributions over the entire trajectories of the cavity processes. To make them more tractable, we consider a perturbative expansion in the coupling strength α, assuming α to be small. We expand the interaction term in the cavity

measure $c_{i\setminus j}$ at second order in α:

$$c_{i\setminus j}[x_i,\hat{x}_i] = \frac{p_0(x_i(0))}{Z_{i\setminus j}} e^{-S_i^{\mathrm{loc}}[x_i,\hat{x}_i]} \prod_{k\in\partial i\setminus j} \int \mathscr{D}[x_k,\hat{x}_k]\, e^{-\alpha S_{ik}^{\mathrm{int}}[x_i,\hat{x}_i,x_k,\hat{x}_k]}\, c_{k\setminus i}[x_k,\hat{x}_k] \tag{B.7a}$$

$$\approx e^{-S_i^{\mathrm{loc}}[x_i,\hat{x}_i]} \prod_{k\in\partial i\setminus j} \left\langle 1 - \alpha S_{ik}^{\mathrm{int}} + \frac{\alpha^2}{2}\left(S_{ik}^{\mathrm{int}}\right)^2 + \cdots \right\rangle_{k\setminus i} \tag{B.7b}$$

where we defined the average over the cavity measure $c_{k\setminus i}$ as $\langle\cdots\rangle_{k\setminus i}$. The averages of the interaction terms can be computed explicitly:

$$\left\langle S_{ik}^{\mathrm{int}} \right\rangle_{k\setminus i} = -\int dt \left(\mathrm{i}\hat{x}_i(t) J_{ik} \langle x_k(t)\rangle_{k\setminus i} + \langle \mathrm{i}\hat{x}_k(t)\rangle_{k\setminus i} J_{ki} x_i(t)\right), \tag{B.8}$$

$$\begin{aligned}\left\langle \left(S_{ik}^{\mathrm{int}}\right)^2 \right\rangle_{k\setminus i} = \int dt dt' \Big(& J_{ik}^2 \mathrm{i}\hat{x}_i(t)\mathrm{i}\hat{x}_i(t')\langle x_k(t)x_k(t')\rangle_{k\setminus i} + J_{ki}^2 x_i(t)x_i(t')\langle \mathrm{i}\hat{x}_k(t)\mathrm{i}\hat{x}_k(t')\rangle_{k\setminus i} \\ & + 2J_{ik}J_{ki}\mathrm{i}\hat{x}_i(t)x_i(t')\langle x_k(t)\mathrm{i}\hat{x}_k(t')\rangle_{k\setminus i}\Big).\end{aligned} \tag{B.9}$$

We identify the cavity mean $\mu_{k\setminus i}(t)$, (disconnected) correlation $C^{\mathrm{dc}}_{k\setminus i}(t,t')$ and response $R_{k\setminus i}(t,t')$ functions as

$$\mu_{k\setminus i}(t) = \langle x_k(t)\rangle_{k\setminus i}, \tag{B.10}$$

$$C^{\mathrm{dc}}_{k\setminus i}(t,t') = \langle x_k(t)x_k(t')\rangle_{k\setminus i}, \tag{B.11}$$

$$R_{k\setminus i}(t,t') = \langle x_k(t)\mathrm{i}\hat{x}_k(t')\rangle_{k\setminus i}, \tag{B.12}$$

while the terms $\langle \mathrm{i}\hat{x}_k(t)\rangle_{k\setminus i}$ and $\langle \mathrm{i}\hat{x}_k(t)\mathrm{i}\hat{x}_k(t')\rangle_{k\setminus i}$ vanish due to causality, being averages of response fields only. The cavity measure can therefore be written as

$$\begin{aligned}c_{i\setminus j}[x_i,\hat{x}_i] \approx e^{-S_i^{\mathrm{loc}}[x_i,\hat{x}_i]} \prod_{k\in\partial i\setminus j} \Bigg[& 1 + \alpha \int dt\, \mathrm{i}\hat{x}_i(t) J_{ki} \mu_{k\setminus i}(t) \\ & + \frac{\alpha^2}{2}\int dt dt' \left(\mathrm{i}\hat{x}_i(t) J_{ik}^2 C^{\mathrm{dc}}_{k\setminus i}(t,t')\mathrm{i}\hat{x}_i(t') + 2\mathrm{i}\hat{x}_i(t) J_{ik}J_{ki} R_{k\setminus i}(t,t') x_i(t')\right)\Bigg].\end{aligned} \tag{B.13}$$

It is convenient to re-exponentiate the expression after having computed the averages. Noticing that

$$e^{-\alpha\langle x\rangle + \frac{\alpha^2}{2}\left(\langle x^2\rangle - \langle x\rangle^2\right)} \approx 1 - \alpha\langle x\rangle + \frac{\alpha^2}{2}\langle x^2\rangle + O(x^3)$$

in the exponentiation process an additional second-order term appears and can be taken into account by using the (connected) cavity correlation function

$$C_{i\setminus j}(t,t') = C^{\mathrm{dc}}_{i\setminus j}(t,t') - \mu_{i\setminus j}(t)\mu_{i\setminus j}(t').$$

The cavity measure can therefore be written as

$$\begin{aligned} c_{i\setminus j}[x_i,\hat{x}_i] \approx \frac{1}{Z^{\mathrm{eff}}_{i\setminus j}} e^{-S^{\mathrm{loc}}_i[x_i,\hat{x}_i] + \alpha \int dt\, \mathrm{i}\hat{x}_i(t) \sum_{k\in\partial i\setminus j} J_{ik}\mu_{k\setminus i}(t)} \\ \times e^{\frac{\alpha^2}{2}\int dt dt' \sum_{k\in\partial i\setminus j}\left(J^2_{ik}\mathrm{i}\hat{x}_i(t)\mathrm{i}C_{k\setminus i}(t,t')\mathrm{i}\hat{x}_i(t') + 2\mathrm{i}\hat{x}_i(t)J_{ik}J_{ki}R_{k\setminus i}(t,t')x_i(t')\right)}, \end{aligned} \tag{B.14}$$

where $Z_{i\setminus j}$ is a normalization constant, obtained by imposing unitary integral of the cavity measure. The cavity measure is therefore Gaussian, with an effective action

$$\begin{aligned} S^{\mathrm{eff}}_{i\setminus j}[x_i,\hat{x}_i] = \int_0^{t_f} dt \left[\mathrm{i}\hat{x}_i(t)\left(\dot{x}_i(t) + \lambda_i x_i(t) + \alpha \sum_{k\in\partial i\setminus j} J_{ik}\mu_{k\setminus i}(t)\right) + T\hat{x}_i^2(t)\right] \\ + \frac{\alpha^2}{2}\int_0^{t_f} dt dt' \sum_{k\in\partial i\setminus j}\left[\mathrm{i}\hat{x}_i(t)J^2_{ik}C_{k\setminus i}(t,t')\mathrm{i}\hat{x}_i(t') + 2\,\mathrm{i}\hat{x}_i(t)J_{ik}J_{ki}R_{k\setminus i}(t,t')x_i(t')\right]. \end{aligned} \tag{B.15}$$

B.1.3 Gaussian ansatz for the cavity measures

The Gaussian ansatz Eq. (4.18) for the cavity measures $c_{i\setminus j}$ is motivated by the small-coupling expansion derived above, which shows that the cavity measures are approximately Gaussian for weak couplings. If we plug the Gaussian ansatz into the dynamic cavity equations Eq. (4.7), we can use the properties of Gaussian integrals to derive self-consistent equations for the cavity means, correlations and response functions

$$c_{i\setminus j}[\boldsymbol{X}_i] \propto e^{-S^{\mathrm{loc}}_i[\boldsymbol{X}_i]} \prod_{k\in\partial i\setminus j}\int d\boldsymbol{X}_k\, c_{k\setminus i}[\boldsymbol{X}_k] e^{-S^{\mathrm{int}}_{ik}[\boldsymbol{X}_i,\boldsymbol{X}_k]} \tag{B.16a}$$

$$= e^{-\frac{1}{2}\boldsymbol{X}_i\left[\boldsymbol{G}^{\mathrm{loc}}_i\right]^{-1}\boldsymbol{X}_i^\top} \prod_{k\in\partial i\setminus j}\int d\boldsymbol{X}_k\, e^{-\frac{1}{2}(\boldsymbol{X}_k-\boldsymbol{M}_{k\setminus i})\boldsymbol{G}^{-1}_{k\setminus i}(\boldsymbol{X}_k-\boldsymbol{M}_{k\setminus i})^\top} e^{-\boldsymbol{X}_i\boldsymbol{J}_{ik}\boldsymbol{X}_k^\top} \tag{B.16b}$$

$$\propto e^{-\frac{1}{2}\boldsymbol{X}_i\left[\boldsymbol{G}^{\mathrm{loc}}_i\right]^{-1}\boldsymbol{X}_i^\top + \frac{1}{2}\sum_{k\in\partial i\setminus j}\boldsymbol{X}_i\boldsymbol{J}_{ik}\boldsymbol{G}_{k\setminus i}\boldsymbol{J}_{ik}^\top\boldsymbol{X}_i^\top + \sum_{k\in\partial i\setminus j}\boldsymbol{X}_i\boldsymbol{J}_{ik}\boldsymbol{M}^\top_{k\setminus i}}. \tag{B.16c}$$

Comparing the last expression with the Gaussian ansatz we obtain an identity for the exponents. Since the identity must hold for all $\boldsymbol{X}_i$, we can match the coefficients of the quadratic and linear

terms to obtain self-consistent equations for the cavity covariance and mean:

$$\boldsymbol{G}_{i\backslash j}^{-1} = \left[\boldsymbol{G}_i^{\mathrm{loc}}\right]^{-1} - \sum_{k\in\partial i\backslash j} \boldsymbol{J}_{ik}\boldsymbol{G}_{k\backslash i}\boldsymbol{J}_{ik}^{\top}, \tag{B.17}$$

$$\boldsymbol{G}_{i\backslash j}^{-1}\boldsymbol{M}_{i\backslash j}^{\top} = \sum_{k\in\partial i\backslash j} \boldsymbol{J}_{ik}\boldsymbol{M}_{k\backslash i}^{\top}. \tag{B.18}$$

Substituting Eq. (B.17) into Eq. (B.18) and rearranging we obtain

$$\left[\boldsymbol{G}_i^{\mathrm{loc}}\right]^{-1}\boldsymbol{G}_{i\backslash j} = \boldsymbol{I} + \sum_{k\in\partial i\backslash j} \boldsymbol{J}_{ik}\boldsymbol{G}_{k\backslash i}\boldsymbol{J}_{ik}^{\top}\boldsymbol{G}_{i\backslash j} \tag{B.19}$$

$$\left[\boldsymbol{G}_i^{\mathrm{loc}}\right]^{-1}\boldsymbol{M}_{i\backslash j}^{\top} = \sum_{k\in\partial i\backslash j} \boldsymbol{J}_{ik}\boldsymbol{M}_{k\backslash i}^{\top} + \sum_{k\in\partial i\backslash j} \boldsymbol{J}_{ik}\boldsymbol{G}_{k\backslash i}\boldsymbol{J}_{ik}^{\top}\boldsymbol{M}_{i\backslash j}^{\top}. \tag{B.20}$$

B.1.4 GECaM equations with non-vanishing response-field averages

The GECaM equations Eqs. (B.17) and (B.18) can be generalized to the case where the averages of the response fields do not vanish, i.e. $\langle \mathrm{i}\hat{x}_i(t)\rangle \neq 0$. This situation can arise when the system is locally constrained, i.e. when causality is broken. In this case, the cavity measures can still be assumed to be Gaussian, but with non-zero means for both the state and response fields. The propagator $\boldsymbol{G}_{i\backslash j}$ now includes non-zero blocks for the response-response correlations, and the cavity mean $\boldsymbol{M}_{i\backslash j}$ has non-zero components for the response fields

$$\boldsymbol{G}_{i\backslash j} = \begin{pmatrix} \boldsymbol{C}_{i\backslash j} & \boldsymbol{R}_{i\backslash j} \\ \boldsymbol{R}_{i\backslash j}^{\top} & \boldsymbol{B}_{i\backslash j} \end{pmatrix}, \quad \boldsymbol{M}_{i\backslash j} = \begin{pmatrix} \boldsymbol{\mu}_{i\backslash j} & \boldsymbol{P}_{i\backslash j} \end{pmatrix}, \tag{B.21}$$

where $\boldsymbol{B}_{i\backslash j}$ is the response-response correlation matrix and $\boldsymbol{P}_{i\backslash j}$ is the mean of the response fields. Since the $\boldsymbol{P}_{i\backslash j}$ is non-zero, the responses are taken as connected, i.e.

$$\mu_{i\backslash j}^{n} = \langle x_i^n\rangle_{i\backslash j}, \qquad P_{i\backslash j}^{n} = \langle \mathrm{i}\hat{x}_i^n\rangle_{i\backslash j}, \tag{B.22}$$

$$C_{i\backslash j}^{nm} = \langle x_i^n x_i^m\rangle_{i\backslash j} - \mu_{i\backslash j}^{n}\mu_{i\backslash j}^{m}, \qquad B_{i\backslash j}^{nm} = \langle \mathrm{i}\hat{x}_i^n \mathrm{i}\hat{x}_i^m\rangle_{i\backslash j} - P_{i\backslash j}^{n}P_{i\backslash j}^{m}, \tag{B.23}$$

$$R_{i\backslash j}^{nm} = \langle x_i^n \mathrm{i}\hat{x}_i^m\rangle_{i\backslash j} - \mu_{i\backslash j}^{n}P_{i\backslash j}^{m}. \tag{B.24}$$

With these definitions, the GECaM equations can be obtained by plugging the Gaussian ansatz into the dynamic cavity equations and matching the coefficients of the quadratic and linear terms, as done in section B.1.3. The resulting equations are the same as Eqs. (B.17) and (B.18), but with the generalized definitions of the propagator and mean.

By expanding the matrix products in Eqs. (B.17) and (B.18), we obtain a set of cavity equations for the responses and correlations

$$(\boldsymbol{R}_i^{\mathrm{loc}})^{-1}\boldsymbol{R}_{i\backslash j} = \boldsymbol{I} + 2T\Delta\boldsymbol{B}_{i\backslash j} + \Delta^2 \sum_{k\in\partial i\backslash j} J_{ik}J_{ki}\boldsymbol{R}_{k\backslash i}\boldsymbol{R}_{i\backslash j} + \Delta^2 \sum_{k\in\partial i\backslash j} J_{ik}^2\boldsymbol{C}_{k\backslash i}\boldsymbol{B}_{i\backslash j} \tag{B.25}$$

$$(\boldsymbol{R}_i^{\mathrm{loc}})^{-1}\boldsymbol{C}_{i\backslash j} = 2\Delta T\boldsymbol{R}_{i\backslash j}^{\top} + \Delta^2 \sum_{k\in\partial i\backslash j} J_{ik}^2\boldsymbol{C}_{k\backslash i}\boldsymbol{R}_{i\backslash j}^{\top} + \Delta^2 \sum_{k\in\partial i\backslash j} J_{ik}J_{ki}\boldsymbol{R}_{k\backslash i}\boldsymbol{C}_{i\backslash j} \tag{B.26}$$

$$\left[(\boldsymbol{R}_i^{\mathrm{loc}})^{-1}\right]^{\top}\boldsymbol{B}_{i\backslash j} = \Delta^2 \sum_{k\in\partial i\backslash j} J_{ik}J_{ki}\boldsymbol{R}_{k\backslash i}^{\top}\boldsymbol{B}_{i\backslash j} + \Delta^2 \sum_{k\in\partial i\backslash j} J_{ki}^2\boldsymbol{B}_{k\backslash i}\boldsymbol{R}_{i\backslash j} \tag{B.27}$$

and for the means

$$\begin{aligned}(\boldsymbol{R}_i^{\mathrm{loc}})^{-1}\boldsymbol{\mu}_{i\backslash j}^{\top} = \Delta \sum_{k\in\partial i\backslash j} J_{ik}\boldsymbol{\mu}_{k\backslash i}^{\top} + \Delta^2 \sum_{k\in\partial i\backslash j} J_{ik}J_{ki}\boldsymbol{R}_{k\backslash i}\boldsymbol{\mu}_{i\backslash j}^{\top} \\ + 2T\Delta\boldsymbol{P}_{i\backslash j}^{\top} + \Delta^2 \sum_{k\in\partial i\backslash j} J_{ik}^2\boldsymbol{C}_{k\backslash i}\boldsymbol{P}_{i\backslash j}^{\top}\end{aligned} \tag{B.28}$$

$$\left[(\boldsymbol{R}_i^{\mathrm{loc}})^{-1}\right]^{\top}\boldsymbol{P}_{i\backslash j}^{\top} = \Delta \sum_{k\in\partial i\backslash j} J_{ki}\boldsymbol{P}_{k\backslash i}^{\top} + \Delta^2 \sum_{k\in\partial i\backslash j} J_{ki}^2\boldsymbol{B}_{k\backslash i}\boldsymbol{\mu}_{i\backslash j}^{\top} + \Delta^2 \sum_{k\in\partial i\backslash j} J_{ik}J_{ki}\boldsymbol{R}_{k\backslash i}^{\top}\boldsymbol{P}_{i\backslash j}^{\top}. \tag{B.29}$$

In the continuous-time limit $\Delta \to 0$, the equations reduce to

$$\frac{d\mu_{i\setminus j}(t)}{dt} = -\lambda_i \mu_{i\setminus j}(t) + \sum_{k\in\partial i\setminus j} J_{ik}\mu_{k\setminus i}(t) + \sum_{k\in\partial i\setminus j}\int ds J_{ik} R_{k\setminus i}(t,s) J_{ki}\mu_{i\setminus j}(s) + 2TP_{i\setminus j}(t) + \sum_{k\in\partial i\setminus j}\int ds J_{ik}^2 C_{k\setminus i}(t,s) P_{i\setminus j}(s) \tag{B.30}$$

$$\frac{dP_{i\setminus j}(t)}{dt} = \lambda_i P_{i\setminus j}(t) - \sum_{k\in\partial i\setminus j} J_{ki}P_{k\setminus i}(t) - \sum_{k\in\partial i\setminus j}\int ds J_{ki}^2 B_{k\setminus i}(t,s)\mu_{i\setminus j}(s) - \sum_{k\in\partial i\setminus j}\int ds J_{ik}J_{ki} R_{k\setminus i}(s,t) P_{i\setminus j}(s) \tag{B.31}$$

$$\frac{\partial R_{i\setminus j}(t,t')}{\partial t} = -\lambda_i R_{i\setminus j}(t,t') + \sum_{k\in\partial i\setminus j}\int ds J_{ik} R_{k\setminus i}(t,s) J_{ki} R_{i\setminus j}(s,t') + \sum_{k\in\partial i\setminus j}\int ds J_{ik}^2 C_{k\setminus i}(t,s) B_{i\setminus j}(s,t') + \delta(t,t') \tag{B.32}$$

$$\frac{\partial C_{i\setminus j}(t,t')}{\partial t} = -\lambda_i C_{i\setminus j}(t,t') + \sum_{k\in\partial i\setminus j}\int ds J_{ik} R_{k\setminus i}(t,s) J_{ki} C_{i\setminus j}(s,t') + 2DR_{i\setminus j}(t',t) + \sum_{k\in\partial i\setminus j}\int ds R_{i\setminus j}(t',s) J_{ik}^2 C_{k\setminus i}(t,s) \tag{B.33}$$

$$\frac{\partial B_{i\setminus j}(t,t')}{\partial t} = \lambda_i B_{i\setminus j}(t,t') - \sum_{k\in\partial i\setminus j}\int ds J_{ik}J_{ki} R_{k\setminus i}(s,t) B_{i\setminus j}(s,t') - \sum_{k\in\partial i\setminus j}\int ds J_{ki}^2 B_{k\setminus i}(t,s) R_{i\setminus j}(s,t') \tag{B.34}$$

We can ascertain that if $P_{k\setminus i}$ and $B_{k\setminus i}$ vanish, the general GECaM equations reduce to the ones reported in the main text. In addition to that, $P_{i\setminus j}$ and $B_{i\setminus j}$ vanish consequently, and therefore the assumption is self-consistently verified.

B.1.5 Equilibrium GECaM equations

We consider the Time-Translationally Invariant (TTI) GECaM equations Eqs. (4.43) and (4.44)

$$\dot{R}_{i\backslash j}(\tau) = -\lambda_i R_{i\backslash j}(\tau) + \sum_{k\in\partial i\backslash j} J_{ik}J_{ki}\int_0^\tau du R_{k\backslash i}(u)R_{i\backslash j}(\tau-u) + \delta(\tau), \tag{B.35}$$

$$\begin{aligned}\dot{C}_{i\backslash j}(\tau) = &-\lambda_i C_{i\backslash j}(\tau) + \sum_{k\in\partial i\backslash j} J_{ik}J_{ki}\int_0^\infty du R_{k\backslash i}(u)C_{i\backslash j}(\tau-u)\\ &+ 2TR_{i\backslash j}(-\tau) + \sum_{k\in\partial i\backslash j} J_{ik}^2\int_\tau^\infty du C_{k\backslash i}(u)R_{i\backslash j}(u-\tau),\end{aligned} \tag{B.36}$$

where we have changed integration variable to $u = \tau - s$ in the integral terms.

If the interaction matrix $\underline{\underline{J}}$ is symmetric, i.e. $J_{ij} = J_{ji}$ for every i, j, the system satisfies detailed balance and it eventually reaches equilibrium after a sufficiently long time. Within this regime, applying the Fluctuation-Dissipation Theorem (FDT)

$$TR^{\mathrm{eq}}_{i\backslash j}(\tau) = -C^{\dot{e}q}_{i\backslash j}(\tau)\Theta(\tau), \tag{B.37}$$

we obtain an equation for the equilibrium correlations only,

$$\begin{aligned}C^{\dot{e}q}_{i\backslash j}(\tau)\,(1-2\Theta(-\tau)) = &-\lambda_i C^{eq}_{i\backslash j}(\tau) - \sum_{k\in\partial i\backslash j}\frac{J_{ik}^2}{T}\int_0^\infty du C^{\dot{e}q}_{k\backslash i}(u)C^{eq}_{i\backslash j}(\tau-u)\\ &- \sum_{k\in\partial i\backslash j}\frac{J_{ik}^2}{T}\int_\tau^\infty du C^{\mathrm{eq}}_{k\backslash i}(u)C^{\dot{e}q}_{i\backslash j}(u-\tau).\end{aligned} \tag{B.38}$$

Integrating by parts the last term

$$\begin{aligned}\int_\tau^\infty du C^{\mathrm{eq}}_{k\backslash i}(u)C^{\dot{e}q}_{i\backslash j}(u-\tau) = &\, C^{\mathrm{eq}}_{k\backslash i}(\infty)C^{eq}_{i\backslash j}(\infty) - C^{\mathrm{eq}}_{k\backslash i}(\tau)C^{eq}_{i\backslash j}(0)\\ &- \int_\tau^\infty du C^{\dot{\mathrm{eq}}}_{k\backslash i}(u)C^{eq}_{i\backslash j}(u-\tau).\end{aligned} \tag{B.39}$$

The equilibrium correlations decay to zero at infinite time differences $C^{\mathrm{eq}}_{k\backslash i}(\infty)C^{eq}_{i\backslash j}(\infty) = 0$. The first integral can be split into two parts

$$\int_0^\infty du = \int_0^\tau du + \int_\tau^\infty du,$$

so that the last term exactly cancels out with the integral in Eq. (B.39). The cavity equilibrium correlations are therefore obtained by solving the set of equations

$$\mathrm{sgn}(\tau)\dot{C}^{eq}_{i\backslash j}(\tau) = -\lambda_i C^{eq}_{i\backslash j}(\tau) + \sum_{k\in\partial i\backslash j}\frac{J_{ik}^2}{T}\left(C^{\mathrm{eq}}_{k\backslash i}(\tau)C^{eq}_{i\backslash j}(0) - \int_0^\tau du\dot{C}^{eq}_{k\backslash i}(u)C^{eq}_{i\backslash j}(\tau-u)\right). \tag{B.40}$$

The full equilibrium correlations are obtained from the cavity ones as

$$\mathrm{sgn}(\tau)\dot{C}^{eq}_{i}(\tau) = -\lambda_i C^{eq}_{i}(\tau) + \sum_{k\in\partial i}\frac{J_{ik}^2}{T}\left(C^{\mathrm{eq}}_{k\backslash i}(\tau)C^{eq}_{i}(0) - \int_0^\tau du\dot{C}^{eq}_{k\backslash i}(u)C^{eq}_{i}(\tau-u)\right). \tag{B.41}$$

Numerical solution A numerical solution of Eq. (B.40) can be found by discretizing time with a timestep Δ, i.e. $t = n\Delta$, $n = 0,\dots,T$ with $T = t_f/\Delta$. Within this discretization the equilibrium correlation function becomes a time vector with $T+1$ components $C^{eq,n}_{i\backslash j} = C^{eq}_{i\backslash j}(t = n\Delta)$. Then a discretized version of Eq. Eq. (B.40) is

$$\begin{aligned} C^{eq,n+1}_{i\backslash j} &= (1-\lambda_i\Delta)\,C^{eq,n}_{i\backslash j} + \Delta\sum_{k\in\partial i\backslash j}\frac{J_{ik}^2}{T}C^{\mathrm{eq,n}}_{k\backslash i}C^{eq,0}_{i\backslash j} \\ &\quad - \Delta\sum_{k\in\partial i\backslash j}\frac{J_{ik}^2}{T}\sum_{m=0}^{n-1}\left(C^{eq,m+1}_{k\backslash i} - C^{eq,m}_{k\backslash i}\right)C^{eq,n-m}_{i\backslash j}. \end{aligned} \tag{B.42}$$

$$\begin{aligned} &= (1-\lambda_i\Delta)\,C^{eq,n}_{i\backslash j} + \Delta\frac{C^{eq,0}_{i\backslash j}}{T}\sum_{k\in\partial i\backslash j}J_{ik}^2C^{\mathrm{eq,n}}_{k\backslash i} \\ &\quad - \Delta\sum_{m=0}^{n-1}\frac{C^{eq,n-m}_{i\backslash j}}{T}\sum_{k\in\partial i\backslash j}J_{ik}^2\left(C^{eq,m+1}_{k\backslash i} - C^{eq,m}_{k\backslash i}\right). \end{aligned} \tag{B.43}$$

while for the full correlation we obtain

$$\begin{aligned} C^{eq,n+1}_{i} &= (1-\lambda_i\Delta)\,C^{eq,n}_{i} + \Delta\frac{C^{eq,0}_{i}}{T}\sum_{k\in\partial i\backslash j}J_{ik}^2C^{\mathrm{eq,n}}_{k\backslash i} \\ &\quad - \Delta\sum_{m=0}^{n-1}\frac{C^{eq,n-m}_{i}}{T}\sum_{k\in\partial i\backslash j}J_{ik}^2\left(C^{eq,m+1}_{k\backslash i} - C^{eq,m}_{k\backslash i}\right). \end{aligned} \tag{B.44}$$

B.2 Non-Hermitian random matrices

In what follows we will adopt the method proposed in [56] adapting it to our dynamical setting. To maintain generality, we consider complex-valued coupling matrices $\underline{\underline{\tilde{J}}} \in \mathbb{C}^{N,N}$. From a dynamical point of view, this is equivalent to duplicating the system introducing an auxiliary degree of freedom $y_i \in \mathbb{C}$ for each node of the graph. The duplicated system evolves according to the system of SDEs

$$\frac{dx_i(t)}{dt} = \sum_j \tilde{J}_{ij} y_j(t) - z y_i(t) + \xi_i(t) \tag{B.45}$$

$$\frac{dy_i(t)}{dt} = -\sum_j \tilde{J}^*_{ji} x_j(t) + z^* x_i(t) + \zeta_i(t) \tag{B.46}$$

where $z \in \mathbb{C}$ is a complex-valued parameter, corresponding to the eigenvalue of the coupling matrix. The variables ξ_i and ζ_i are complex Gaussian noise terms with mean and covariance given by

$$\langle \xi_i(t) \rangle = \langle \zeta_i(t) \rangle = 0 \tag{B.47}$$

$$\langle \xi_i^*(t) \xi_{i'}(t') \rangle = \langle \zeta_i^*(t) \zeta_{i'}(t') \rangle = 2T \delta_{ii'} \delta(t-t'). \tag{B.48}$$

The MSRJD functional integral formalism can now be applied, performing an integration over complex-valued trajectories using the complex Dirac delta function

$$\delta^{(2)}(x) = \delta(\mathrm{Re}\, x)\delta(\mathrm{Im}\, x) \propto \int_{\mathbb{C}} d\hat{x} d\hat{x}^*\, e^{-\hat{x}^* x + \hat{x} x^*}, \tag{B.49}$$

where $d\hat{x}d\hat{x}^* = d(\mathrm{Re}\, x)d(\mathrm{Im}\, x) = dx^2$ and the integral is taken over the whole complex plane. The dynamical partition function of the system after discretization of time is

$$\begin{aligned} Z \propto \Bigg\langle \int d[\underline{x}, \underline{\hat{x}}, \underline{y}, \underline{\hat{y}}]^2 \prod_{i,n} \exp \Bigg[& -(\hat{x}_i^n)^* \left(x_i^{n+1} - x_i^n - \Delta \sum_j \tilde{J}_{ij} y_j^n + \Delta z y_i^n - \Delta \xi_i^n \right) \\ & + \hat{x}_i^n \left(\left(x_i^{n+1}\right)^* - (x_i^n)^* - \Delta \sum_j \tilde{J}^*_{ij} \left(y_j^n\right)^* + \Delta z^* \left(y_i^n\right)^* - (\Delta \xi_i^n)^* \right) \\ & - (\hat{y}_i^n)^* \left(y_i^{n+1} - y_i^n + \Delta \sum_j \tilde{J}^*_{ji} x_j^n - \Delta z^* x_i^n - \Delta \zeta_i^n \right) \\ & + \hat{y}_i^n \left(\left(y_i^{n+1}\right)^* - (y_i^n)^* + \Delta \sum_j \tilde{J}_{ji} \left(x_j^n\right)^* - \Delta z (x_i^n)^n - (\Delta \zeta_i^n)^n \right) \Bigg] \Bigg\rangle, \end{aligned} \tag{B.50}$$

where we have introduced $\int d[\underline{\boldsymbol{x}},\underline{\hat{\boldsymbol{x}}},\underline{\boldsymbol{y}},\underline{\hat{\boldsymbol{y}}}]^2 = \prod_{i,n}\int (dx_i^n)^2\,(d\hat{x}_i^n)^2\,(dy_i^n)^2\,(d\hat{y}_i^n)^2$. For notational convenience, we set $\Delta = 1$ below. The average over the complex Gaussian noise can be performed by applying the identity

$$\left\langle e^{\underline{\hat{\boldsymbol{x}}}\underline{\boldsymbol{\Delta\xi}}^\dagger - \underline{\boldsymbol{\Delta\xi}}\,\underline{\hat{\boldsymbol{x}}}^\dagger + \underline{\hat{\boldsymbol{y}}}\,\underline{\boldsymbol{\Delta\zeta}}^\dagger - \underline{\boldsymbol{\Delta\zeta}}\,\underline{\hat{\boldsymbol{y}}}^\dagger}\right\rangle = e^{-\underline{\hat{\boldsymbol{x}}}\,\underline{\underline{g}}_x\,\underline{\hat{\boldsymbol{x}}}^\dagger - \underline{\hat{\boldsymbol{y}}}\,\underline{\underline{g}}_y\,\underline{\hat{\boldsymbol{y}}}^\dagger}, \tag{B.51}$$

where $\left[\underline{\underline{g}}_x\right]_{ii',nn'} = \langle(\Delta\xi_i^n)^*\,\Delta\xi_{i'}^{n'}\rangle = 2T\,\delta_{ii'}\delta_{nn'} = \left[\underline{\underline{g}}_y\right]_{ii',nn'} = \langle(\Delta\zeta_i^n)^*\,\Delta\zeta_{i'}^{n'}\rangle$ are the noise covariance matrices. The dynamical partition function can be factorized as

$$\begin{aligned}
Z \propto \int d[\underline{\boldsymbol{x}},\underline{\hat{\boldsymbol{x}}},\underline{\boldsymbol{y}},\underline{\hat{\boldsymbol{y}}}]^2 \prod_i \exp\Bigg\{\sum_n\Big[&-(\hat{x}_i^n)^*\left(x_i^{n+1} - x_i^n + z y_i^n\right)\\
&+\hat{x}_i^n\left(\left(x_i^{n+1}\right)^* - (x_i^n)^* + z^*\,(y_i^n)^*\right) - 2T\,|\hat{x}_i^n|^2\Big]\Bigg\}\\
\times\prod_{i<j}\exp\Bigg\{\sum_n\left[(\hat{x}_i^n)^*\,\tilde{J}_{ij}y_j^n + y_i^n\tilde{J}_{ji}\left(\hat{x}_j^n\right)^* - \hat{x}_i^n\tilde{J}_{ij}^*\left(y_j^n\right)^* - (y_i^n)^*\,\tilde{J}_{ji}^*\hat{x}_j^n\right]\Bigg\}\\
\times\prod_i\exp\Bigg\{\sum_n\Big[-(\hat{y}_i^n)^*\left(y_i^{n+1} - y_i^n - z^* x_i^n\right) + \hat{y}_i^n\left(\left(y_i^{n+1}\right)^* - (y_i^n)^* - z\,(x_i^n)^*\right) - 2T\,|\hat{y}_i^n|^2\Big]\Bigg\}\\
\times\prod_{i<j}\exp\Bigg\{\sum_n\left[-(\hat{y}_i^n)^*\,\tilde{J}_{ji}^* x_j^n - x_i^n\tilde{J}_{ij}^*\left(\hat{y}_j^n\right)^* + \hat{y}_i^n\tilde{J}_{ji}\left(x_j^n\right)^* + (x_i^n)^*\,\tilde{J}_{ij}\hat{y}_j^n\right]\Bigg\}.
\end{aligned} \tag{B.52}$$

It turns out that, in order to properly consider a dynamic cavity ansatz, the trajectories $\boldsymbol{x}_i$, $\hat{\boldsymbol{x}}_i$, $\boldsymbol{y}_i$ and $\hat{\boldsymbol{y}}_i$ have to be grouped together into a single variable $\mathscr{W}_i = (\boldsymbol{x}_i,\hat{\boldsymbol{x}}_i,\boldsymbol{y}_i,\hat{\boldsymbol{y}}_i)$. We can write the dynamical partition function as a quadratic form

$$Z \propto \int d\underline{\mathscr{W}}^2 \prod_i e^{-\mathscr{W}_i\mathscr{G}_{0,i}^{-1}\mathscr{W}_i^\dagger}\prod_{i<j}e^{-\mathscr{W}_i\tilde{\mathscr{J}}_{ij}\mathscr{W}_j^\dagger - \mathscr{W}_j\tilde{\mathscr{J}}_{ji}\mathscr{W}_i^\dagger}, \tag{B.53}$$

where we have introduced $\int D\underline{\mathscr{W}}^2 = \prod_i\int D\mathscr{W}_i^2 = \int D[\underline{\boldsymbol{x}},\underline{\hat{\boldsymbol{x}}},\underline{\boldsymbol{y}},\underline{\hat{\boldsymbol{y}}}]^2$ and the matrices

$$\mathscr{G}_{0,i}^{-1} = \begin{pmatrix} \mathsf{G}_{0,i}^{-1} & \mathsf{H} \\ -\mathsf{H}^\dagger & \mathsf{G}_{0,i}^{-1}\end{pmatrix}, \tag{B.54}$$

$$\tilde{\mathscr{J}}_{ij} = \begin{pmatrix} 0 & -\tilde{\mathsf{J}}_{ij} \\ \tilde{\mathsf{J}}_{ji}^\dagger & 0\end{pmatrix} = -\tilde{\mathscr{J}}_{ji}^\dagger, \tag{B.55}$$

whose elements are the matrices

$$\mathsf{G}_{0,i}^{-1} = \begin{pmatrix} 0 & \left[R_{0,i}^{-1}\right]^{\dagger} \\ R_{0,i}^{-1} & -2D\mathbb{I} \end{pmatrix}, \tag{B.56}$$

$$\mathsf{H} = \begin{pmatrix} 0 & z\mathbb{I} \\ z\mathbb{I} & 0 \end{pmatrix}, \tag{B.57}$$

$$\tilde{\mathsf{J}}_{ij} = \begin{pmatrix} 0 & \tilde{J}_{ij}\mathbb{I} \\ \tilde{J}_{ij}\mathbb{I} & 0 \end{pmatrix}. \tag{B.58}$$

The cavity messages are

$$c_{i\setminus j}(\mathscr{W}_i) \propto e^{-\mathscr{W}_i \mathscr{G}_{0,i}^{-1} \mathscr{W}_i^{\dagger}} \prod_{k\in\partial i\setminus j} \int d\mathscr{W}_k^2 \, c_{k\setminus i}(\mathscr{W}_k)\, e^{-\mathscr{W}_k \tilde{\mathscr{J}}_{ki} \mathscr{W}_i^{\dagger} - \mathscr{W}_i \tilde{\mathscr{J}}_{ik} \mathscr{W}_k^{\dagger}}. \tag{B.59}$$

We assume once again that the cavity messages have a Gaussian functional form, that is

$$c_{i\setminus j}(\mathscr{W}_i) \propto e^{-\mathscr{W}_i \mathscr{G}_{i\setminus j}^{-1} \mathscr{W}_i^{\dagger}}, \tag{B.60}$$

where the cavity propagator is given by

$$\mathscr{G}_{i\setminus j} = \begin{pmatrix} \mathsf{G}_{i\setminus j}^{\mathsf{XX}} & \mathsf{G}_{i\setminus j}^{\mathsf{XY}} \\ \mathsf{G}_{i\setminus j}^{\mathsf{YX}} & \mathsf{G}_{i\setminus j}^{YY} \end{pmatrix}. \tag{S12}$$

and $\mathsf{G}_{i\setminus j}^{\mathsf{AB}}$ is the usual propagator between variables A_i and B_i, with A and B being either $\mathsf{X} = (\boldsymbol{x}, \hat{\boldsymbol{x}})$ or $\mathsf{Y} = (\boldsymbol{y}, \hat{\boldsymbol{y}})$

$$\mathsf{G}_{i\setminus j}^{\mathsf{AB}} = \begin{pmatrix} \langle \boldsymbol{a}_i^{\dagger} \boldsymbol{b}_i \rangle_{i\setminus j} & \langle \boldsymbol{a}_i^{\dagger} \hat{\boldsymbol{b}}_i \rangle_{i\setminus j} \\ \langle \hat{\boldsymbol{a}}_i^{\dagger} \boldsymbol{b}_i \rangle_{i\setminus j} & \langle \hat{\boldsymbol{a}}_i^{\dagger} \hat{\boldsymbol{b}}_i \rangle_{i\setminus j} \end{pmatrix} = \begin{pmatrix} C_{i\setminus j}^{\mathsf{AB}} & R_{i\setminus j}^{\mathsf{AB}} \\ \left(R_{i\setminus j}^{\mathsf{BA}}\right)^{\dagger} & 0 \end{pmatrix}. \tag{B.61}$$

Integrating out the variables $\mathscr{W}_k$, we obtain

$$c_{i\setminus j}(\mathscr{W}_i) \propto e^{-\mathscr{W}_i \mathscr{G}_{0,i}^{-1} \mathscr{W}_i^{\dagger}} \prod_{k\in\partial i\setminus j} \int d\mathscr{W}_k^2 e^{-\mathscr{W}_k \mathscr{G}_{k\setminus i}^{-1} \mathscr{W}_k^{\dagger} - \mathscr{W}_k \tilde{\mathscr{J}}_{ki} \mathscr{W}_i^{\dagger} - \mathscr{W}_i \tilde{\mathscr{J}}_{ik} \mathscr{W}_k^{\dagger}} \tag{B.62a}$$

$$\propto e^{-\mathscr{W}_i \mathscr{G}_{0,i}^{-1} \mathscr{W}_i^{\dagger} + \sum_{k\in\partial i\setminus j} \mathscr{W}_i \tilde{\mathscr{J}}_{ik} \mathscr{G}_{k\setminus i} \tilde{\mathscr{J}}_{ki} \mathscr{W}_i^{\dagger}}, \tag{B.62b}$$

from which we get the extended GECaM matrix equation

$$\mathscr{G}_{i\backslash j}^{-1} = \mathscr{G}_{0,i}^{-1} - \sum_{k\in\partial i\backslash j} \tilde{\mathscr{J}}_{ik}\mathscr{G}_{k\backslash i}\tilde{\mathscr{J}}_{ki}. \tag{B.63}$$

To compute the spectral density, we only need the equations for the response functions. Conveniently expressed in matrix form and in the Laplace space, these equations are

$$\mathsf{R}_{i\backslash j}^{-1} = \mathsf{z} + \eta\mathsf{I} - \sum_{k\in\partial i\backslash j} \tilde{\mathsf{J}}_{ik}\mathsf{R}_{k\backslash i}\tilde{\mathsf{J}}_{ki}, \tag{B.64}$$

where we have introduced the matrices,

$$\mathsf{R}_{i\backslash j} = \begin{pmatrix} R_{i\backslash j}^{\mathsf{XX}} & R_{i\backslash j}^{\mathsf{XY}} \\ R_{i\backslash j}^{\mathsf{YX}} & R_{i\backslash j}^{\mathsf{YY}} \end{pmatrix} \tag{B.65}$$

$$\mathsf{z} = \begin{pmatrix} 0 & z \\ -z^* & 0 \end{pmatrix} \tag{B.66}$$

and redefined for convenience the coupling matrix $\tilde{\mathsf{J}}_{ij}$ as follows

$$\tilde{\mathsf{J}}_{ij} = \begin{pmatrix} 0 & A_{ij} \\ -A_{ji}^* & 0 \end{pmatrix}. \tag{B.67}$$

Equation (B.64) is equivalent to Eq. (64) in [55]. We can therefore compute the spectral distribution of $\boldsymbol{J}$ as

$$\rho_{\boldsymbol{J}}(z) = \lim_{\eta\to 0} \frac{1}{\pi N}\partial_{z^*} R_i^{\mathsf{YX}}(z-\lambda_i). \tag{B.68}$$

Appendix C

Detailed derivations for the Perturbative Closure

C.1 The Dyson equations

This section provides a compact derivation of the Dyson equations used in section 5.1, together with the identification of the self-energy components and their diagrammatic meaning within the MSRJD formalism.

Let $S[\mathrm{X}] = S^0[\mathrm{X}] + \varepsilon S^{\mathrm{p}}[\mathrm{X}]$ be the action with quadratic Gaussian part

$$S^0[\mathrm{X}] = \frac{1}{2}\mathrm{X}\big[\mathrm{G}^0\big]^{-1}\mathrm{X}^\top,$$

where $\mathrm{X} = (x, \mathrm{i}\hat{x})$ and all time products are understood as convolutions, e.g. $[\mathrm{AB}](t,t') = \int ds\, \mathrm{A}(t,s)\mathrm{B}(s,t')$. We have assumed for simplicity that the Gaussian measure has zero mean; the non-zero mean case is discussed at the end of this section. The Gaussian propagator G^0 is the two-point function of the unperturbed theory,

$$\mathrm{G}^0(t,t') = \langle \mathrm{X}(t)\,\mathrm{X}(t')^\top \rangle^0 = \begin{pmatrix} C^0(t,t') & R^0(t,t') \\ R^{0\top}(t,t') & 0 \end{pmatrix},$$

The full two-point function is instead evaluated with the full action S,

$$\mathrm{G}(t,t') = \langle \mathrm{X}(t)\,\mathrm{X}(t')^\top \rangle = \begin{pmatrix} C(t,t') & R(t,t') \\ R^\top(t,t') & 0 \end{pmatrix},$$

where $\langle\cdot\rangle^0$ and $\langle\cdot\rangle$ denote averages with respect to the Gaussian and full measures respectively. The perturbation action $S^{\mathrm{p}}[\mathrm{X}]$ is composed of two terms, one from the non-linearities in the drift $f(x)$, and one from the non-thermal noise covariance $g(x)$,

$$S^{\mathrm{p}}[\mathrm{X}] = -\int ds\, \mathrm{i}\hat{x}(s) f(x(s)) - \int ds\, g(x(s))\,(\mathrm{i}\hat{x}(s))^2 .$$

At first order in ε, the perturbation S^{p} produces two types of contributions: one from the nonlinear drift, proportional to $\langle f'(x)\rangle^0$, and one from the multiplicative noise, proportional to $\langle g(x)\rangle^0$. We start with the response function, whose first-order correction reads, in diagrammatic form,

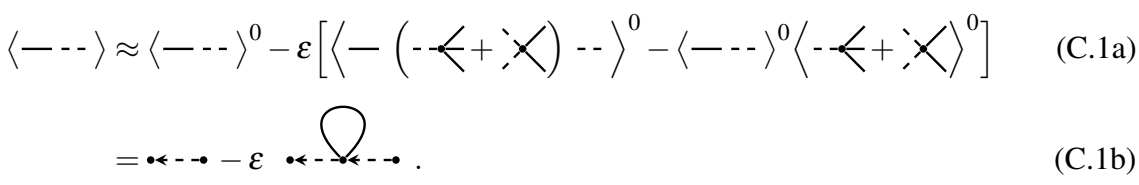

(C.1a)

(C.1b)

The first term is the bare response function $R^0(t,t')$, while the second term contains the self-energy contribution $\Sigma^R(t,t') = \delta(t-t')\langle f'(x(t))\rangle^0$. The correlation function is instead given by

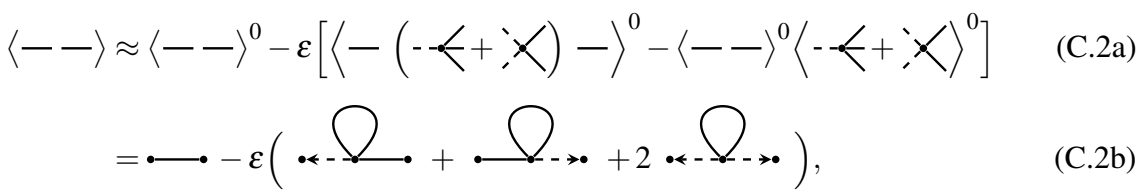

(C.2a)

(C.2b)

where the self-energy contributions are $\Sigma^R_{\mathrm{drift}}$ and $\Sigma^T(t,t') = 2\delta(t-t')\langle g(x(t))\rangle^0$. Collecting these results, we can write the first-order corrections to the response and correlation functions as

$$\delta R(t,t') = \int ds, R^0(t,s)\langle f'(x(s))\rangle^0 R^0(s,t'), \tag{C.3}$$

$$\begin{aligned}\delta C(t,t') = \int ds \Big[& R^0(t,s)\langle f'(x(s))\rangle^0 C^0(s,t') + C^0(s,t)\langle f'(x(s))\rangle^0 R^0(t',s) \\ & + 2R^0(t,s)\langle g(x(s))\rangle^0 R^0(t',s)\Big].\end{aligned} \tag{C.4}$$

These contributions correspond respectively to tadpole insertions on the $(x,\mathrm{i}\hat{x})$ and $(\hat{x},\hat{x})$ lines, whose values are weighted by the local averages $\langle f'(x)\rangle^0$ and $\langle g(x)\rangle^0$. At this order, the self-energy is local in time and renormalizes only the coefficients of the linear dynamics, corresponding to the Hartree-Fock approximation discussed in section 5.1. We can write the

first-order corrections in a compact form as

$$\mathrm{G} = \mathrm{G}^0 + \varepsilon \mathrm{G}^0 \Sigma^{(1)} \mathrm{G}^0, \tag{C.5}$$

where the first-order self-energy is

$$\Sigma^{(1)}(t,t') = \begin{pmatrix} 0 & \varsigma^R(t,t') \\ \varsigma^R(t',t) & 2\varsigma^T(t,t') \end{pmatrix} = \delta(t-t') \begin{pmatrix} 0 & \langle f'(x(t)) \rangle^0 \\ \langle f'(x(t)) \rangle^0 & 2\langle g(x(t)) \rangle^0 \end{pmatrix} \tag{C.6}$$

At second order in ε, we have to consider two types of contributions: the insertion of two tadpoles, and the insertion of a two-point loop. The latter contributions are non-local in time, and give rise to memory kernels in the Dyson equations. If we limit ourselves to the non-local contributions, we can write the second-order corrections to the response as

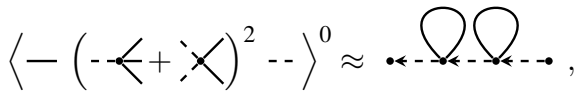

(C.7)

while the contributions to the correlation function are

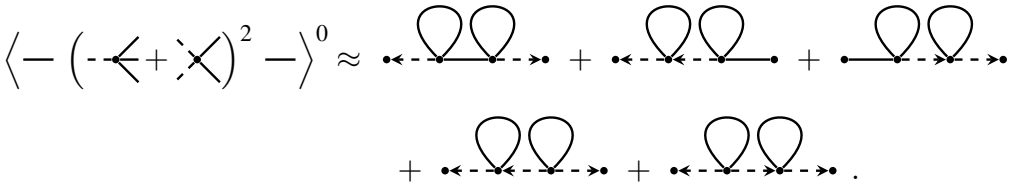

(C.8)

It is clear from the diagrams that these contributions can be factorized in terms of tadpole diagrams. These contributions can be written in a compact form as

$$\mathrm{G} = \mathrm{G}^0 + \varepsilon \mathrm{G}^0 \Sigma^{(1)} \mathrm{G}^0 + \frac{\varepsilon^2}{2} \mathrm{G}^0 \Sigma^{(1)} \mathrm{G}^0 \Sigma^{(1)} \mathrm{G}^0. \tag{C.9}$$

By summing the series to all orders in ε, we obtain the Dyson equation

$$\mathrm{G} = \mathrm{G}^0 + \varepsilon \mathrm{G}^0 \Sigma^{(1)} \mathrm{G}^0 + \frac{\varepsilon^2}{2} \mathrm{G}^0 \Sigma^{(1)} \mathrm{G}^0 \Sigma^{(1)} \mathrm{G}^0 + \cdots = \left(\mathrm{G}^0 - \Sigma^{(1)} \right)^{-1}, \tag{C.10}$$

which is exact if we limit ourselves to the local self-energy $\Sigma^{(1)}$. This corresponds to the Hartree-Fock approximation discussed in section 5.1. To go beyond this approximation, we need to include the non-local self-energy contributions arising at second order in ε. For example,

at second order in the response function we have, among other contributions,

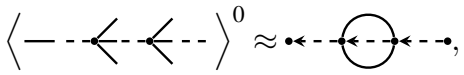

(C.11)

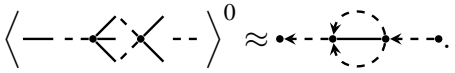

(C.12)

These contributions give rise to memory kernels in the Dyson equations, and can be resummed to all orders in ε using the same procedure as before. The full Dyson equation is then

$$\mathrm{G} = \mathrm{G}^0 + \mathrm{G}^0 \Sigma \mathrm{G}, \tag{C.13}$$

where the self-energy Σ now contains both local and non-local contributions.

C.2 The Bouchaud-Mézard model

In this appendix we derive in detail the Gaussian perturbative closure for the Bouchaud-Mézard (BM) dynamics Eq. (5.47) on random regular graphs, discuss the renormalization of the multiplicative noise amplitude, and compute the critical integral determining the noise-driven transition. Finally, we show how the adiabatic approximation used in [66] emerges as a limiting case of GECaM under homogeneous conditions.

C.2.1 Renormalization of the noise amplitude

We consider the BM dynamics on a Random Regular Graph (RRG) with connectivity K. Since the model is invariant under the scaling $x_i \to \alpha x_i$, we can set $\langle x_i(t) \rangle = 1$ without loss of generality and define fluctuations $\delta x_i(t) = x_i(t) - 1$.

The cavity measure can be written as

$$c_{i\setminus j}[\delta \mathrm{X}_i] \propto \exp\Big\{ - S^0_{i\setminus j}[\delta \mathrm{X}_i] - \varepsilon S^{\mathrm{p}}_{i\setminus j}[\delta \mathrm{X}_i] \Big\}, \tag{C.14}$$

where $S^0_{i\setminus j}[\delta \mathrm{X}_i] = \delta \mathrm{X}_i [\mathrm{G}^0_{i\setminus j}]^{-1} \delta \mathrm{X}_i^\top$ is the Gaussian part and $\delta \mathrm{X}_i = (\delta x_i, \mathrm{i}\hat{x}_i)$. The perturbation action originates from the multiplicative noise and reads

$$S^{\mathrm{p}}_{i\setminus j}[\mathrm{X}_i] = -\frac{\sigma^2}{2}\int_0^{t_f} ds\,(1+\delta x_i(s))^2 \big(\mathrm{i}\hat{x}_i(s)\big)^2 = -\frac{\sigma^2}{2}\int_0^{t_f} ds\,\big(\mathrm{i}\hat{x}_i(s)\big)^2 - \sigma^2 \int_0^{t_f} ds\,\delta x_i(s)\big(\mathrm{i}\hat{x}_i(s)\big)^2 - \frac{\sigma^2}{2}\int_0^{t_f} ds\,\delta x_i^2(s)\big(\mathrm{i}\hat{x}_i(s)\big)^2. \tag{C.15}$$

These three terms correspond respectively to the following diagrammatic vertices:

$$-\frac{\sigma^2}{2}\int_0^{t_f} ds\,\big(\mathrm{i}\hat{x}_i(s)\big)^2 = \tag{C.16}$$

$$-\sigma^2 \int_0^{t_f} ds\,\delta x_i(s)\big(\mathrm{i}\hat{x}_i(s)\big)^2 = \tag{C.17}$$

$$-\frac{\sigma^2}{2}\int_0^{t_f} ds\,\delta x_i^2(s)\big(\mathrm{i}\hat{x}_i(s)\big)^2 = \tag{C.18}$$

The response function remains unperturbed at the Hartree-Fock level, $R_{i\setminus j}(t,t') = R^0_{i\setminus j}(t,t')$, since there is no drift nonlinearity. The correlation function instead receives a local renormalization from the last vertex, while the second vertex contributes only at higher non-local orders. The first-order correction to the correlation function is

$\langle - - \rangle \approx \langle - - \rangle^0 - \varepsilon \Big[\Big\langle - \Big(\quad + \quad \Big) - \Big\rangle^0 - \langle - - \rangle^0 \Big\langle \quad + \quad \Big\rangle^0 \Big]$ (C.19a)

$= \bullet\!\!-\!\!-\!\!\bullet - 2\varepsilon$, (C.19b)

which reads explicitly

$$C_{i\setminus j}(t,t') \approx C^0_{i\setminus j}(t,t') + \varepsilon\sigma^2 \int ds R_{i\setminus j}(t,s) R_{i\setminus j}(t',s) C^0_{i\setminus j}(s,s). \tag{C.20}$$

The tadpole corrections can be resummed to all orders by substituting $C^0_{i\setminus j}(s,s)$ with its perturbed counterpart $C_{i\setminus j}(s,s)$ in the above equation, yielding at stationarity

$$C_{i\setminus j}(\tau) = C^0_{i\setminus j}(\tau) + \varepsilon\sigma^2(1+q^{\mathrm{st}}_{i\setminus j}) \int_\tau^\infty du R_{i\setminus j}(u) R_{i\setminus j}(u-\tau), \tag{C.21}$$

where $q^{\mathrm{st}}_{i\setminus j} = C_{i\setminus j}(0)$ is the equal time correlation. This result can be interpreted as a renormalization of the noise amplitude,

$$\left(\sigma^{\mathrm{r}}_{i\setminus j}\right)^2 = \sigma^2(1+q^{\mathrm{st}}_{i\setminus j}), \tag{C.22}$$

which increases with the amplitude of fluctuations. We notice that, since the unperturbed Gaussian process was noiseless ($T = 0$), the unperturbed correlation function is vanishing, $C^0_{i\setminus j}(\tau) = 0$, and the entire correlation function is generated by the multiplicative noise

$$C_{i\setminus j}(\tau) = \varepsilon\left(\sigma^{\mathrm{r}}_{i\setminus j}\right)^2 \int_\tau^\infty du R_{i\setminus j}(u)R_{i\setminus j}(u-\tau). \tag{C.23}$$

The equal time correlation function is then obtained from the integral of the correlation spectrum,

$$q^{\mathrm{st}}_{i\setminus j} = \varepsilon\left(\sigma^{\mathrm{r}}_{i\setminus j}\right)^2 \int_0^\infty d\omega R_{i\setminus j}(\omega)R_{i\setminus j}(-\omega). \tag{C.24}$$

Defining the integral

$$I_{i\setminus j} = \int_0^\infty d\omega R_{i\setminus j}(\omega)R_{i\setminus j}(-\omega), \tag{C.25}$$

we find

$$q^{\mathrm{st}}_{i\setminus j} = \varepsilon\left(\sigma^{\mathrm{r}}_{i\setminus j}\right)^2 I_{i\setminus j} = \varepsilon(1+q^{\mathrm{st}}_{i\setminus j})\sigma^2 I_{i\setminus j}, \tag{C.26}$$

which define a recursive geometric equation for $q^{\mathrm{st}}_{i\setminus j}$. Solving for $q^{\mathrm{st}}_{i\setminus j}$ and substituing into the renormalized noise amplitude Eq. (C.22), we obtain

$$\left(\sigma^{\mathrm{r}}_{i\setminus j}\right)^2 = \frac{\sigma^2}{1-\varepsilon\sigma^2 I_{i\setminus j}}. \tag{C.27}$$

C.2.2 Critical integral computation

On a K-regular graph, all cavity quantities are identical, and $\lambda = KJ$. The loop integral becomes

$$I = \int_{-\infty}^{\infty} \frac{d\omega}{2\pi} \frac{\mathrm{i}\omega + KJ - \sqrt{(\mathrm{i}\omega + KJ)^2 - 4(K-1)J^2}}{2(K-1)J^2} \times \frac{-\mathrm{i}\omega + KJ - \sqrt{(-\mathrm{i}\omega + KJ)^2 - 4(K-1)J^2}}{2(K-1)J^2} \tag{C.28a}$$

$$= \frac{2}{\pi\sqrt{K-1}J} \int_{-1}^{1} du \frac{\sqrt{1-u^2}}{\frac{K}{\sqrt{K-1}} - u + \sqrt{\left(\frac{K}{\sqrt{K-1}} - u\right)^2 - 1}} \tag{C.28b}$$

$$= -\frac{2\mathrm{i}}{\pi} \int_{\mathfrak{Im}} d(\mathrm{i}\omega) \frac{1}{\lambda + \mathrm{i}\omega + \sqrt{(\lambda + \mathrm{i}\omega)^2 - 4(K-1)J^2}} \times \frac{1}{\lambda - \mathrm{i}\omega + \sqrt{(\lambda - \mathrm{i}\omega)^2 - 4(K-1)J^2}} \tag{C.28c}$$

We define $I = -(2\mathrm{i}/\pi)\int_{\mathfrak{Im}} dz F(z)$ with $F(z) = f(z)f(-z)$ and

$$f(z) = \lambda + z + \sqrt{(\lambda + z)^2 - 4(K-1)J^2}.$$

The function $F(z)$ has two symmetric branch cuts $\mathscr{B}_{\pm} = [\pm\lambda - A, \pm\lambda + A]$ with $A = 2\sqrt{K-1}J$.

The contour is closed in the left half plane with a detour around $\mathscr{B}_-$ (see Fig. C.1). All large radius and infinitesimal arcs vanish, while the two integrals along the lips of $\mathscr{B}_-$ add because the square root changes sign across the cut. In particular, defining γ_3 and γ_4 as the upper and lower lips of $\mathscr{B}_-$ respectively, we have

$$\int_{\gamma_3} dz F(z) = \int_{-\lambda - A}^{-\lambda + A} dx \frac{1}{\lambda + x + \mathrm{i}\sqrt{A^2 - (\lambda + x)^2}} \frac{1}{\lambda - x + \sqrt{(\lambda - x)^2 - A^2}} \tag{C.29}$$

$$\int_{\gamma_4} dz F(z) = \int_{-\lambda + A}^{-\lambda - A} dx \frac{1}{\lambda + x - \mathrm{i}\sqrt{A^2 - (\lambda + x)^2}} \frac{1}{\lambda - x + \sqrt{(\lambda - x)^2 - A^2}} \tag{C.30}$$

where we used the parametrization $z = -xe^{\mathrm{i}\pi}$ for the first integral and $z = -xe^{-\mathrm{i}\pi}$ for the second one.

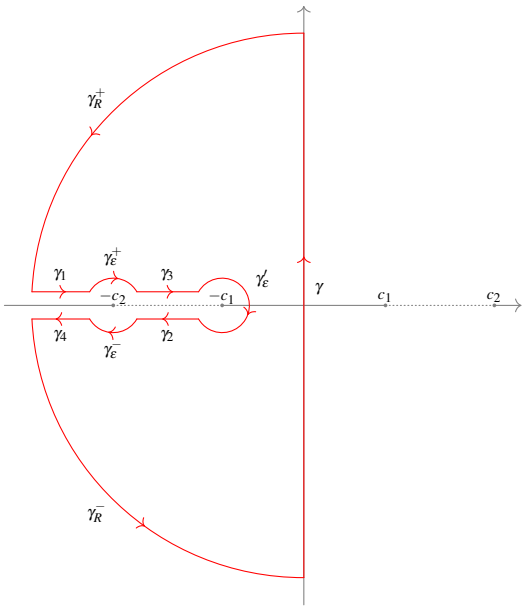

Fig. C.1 **Integration contour for the complex-plane calculation.** Contour used for the computation of Eq. (C.28c). The paths γ_R^+ and γ_R^- are circular arcs of radius $R \to \infty$, while γ_ε^+, γ_ε^- and γ'_ε are circular arcs of radius $\varepsilon \to 0$. It can be shown that all of them vanish in the proper limit.

The final expression is

$$I = \frac{4}{\pi A^2} \int_{-\lambda-A}^{-\lambda+A} dx \, \frac{\sqrt{A^2-(\lambda+x)^2}}{\lambda - x + \sqrt{(\lambda-x)^2 - A^2}}. \tag{C.31}$$

Introducing $u = (\lambda + x)/A \in [-1,1]$, we obtain

$$I = I_K = \frac{2}{\pi\sqrt{K-1}J} \int_{-1}^{1} du \, \frac{\sqrt{1-u^2}}{\frac{K}{\sqrt{K-1}} - u + \sqrt{\left(\frac{K}{\sqrt{K-1}} - u\right)^2 - 1}}. \tag{C.32}$$

In the fully connected limit ($J \to J/K$, $K \to \infty$) the integral reduces to $I_\infty = 1/(2J)$, and therefore

$$\sigma_c^2 = 2J(K-1). \tag{C.33}$$

C.2.3 Adiabatic approximation of GECaM

The approximation method proposed in [66, 108, 109] is based on the combination of two assumptions: (1) the "local field" $\sum_{j\in\partial i} x_j$ acting on i is assumed to change much more slowly than x_i (*adiabatic assumption*); (2) the neighbouring variables x_j for $j \in \partial i$ are assumed to

be independent, so that the local field satisfies the Central Limit Theorem and converges to a Gaussian distribution (*independence assumption*). In this Section, we show that this approximation naturally follows from the GECaM approximation under specific conditions. To simplify the calculations, we will only consider the case of a RRG with degree K.

Employing as usual the Gaussian ansatz in Eq. (4.18), the local marginal of node i with a perturbation action $S_i^{\mathrm{p}}[\mathrm{X}_i]$ can be written as follows

$$c_i[\mathrm{X}_i] \propto e^{-\frac{1}{2}\mathrm{X}_i\mathrm{G}_{0,i}^{-1}\mathrm{X}_i^\top - S_i^{\mathrm{p}}[\mathrm{X}_i]} \prod_{k\in\partial i}\int \mathscr{D}[\mathrm{X}_k] e^{-\frac{1}{2}\left(\mathrm{X}_k-\mathrm{M}_{k\setminus i}\right)\mathrm{G}_{k\setminus i}^{-1}\left(\mathrm{X}_k-\mathrm{M}_{k\setminus i}\right)^\top} e^{-\mathrm{X}_i\mathrm{J}_{ik}\mathrm{X}_k^\top}. \tag{C.34}$$

We also assume that all couplings are the same, i.e. $\mathrm{J}_{ik}=\mathrm{J}$, and introduce a local field $K\mathrm{H}_i = \sum_{j\in\partial i}\mathrm{X}_j$. The previous expression becomes

$$\begin{aligned}
c_i[\mathrm{X}_i] &\propto e^{-\frac{1}{2}\mathrm{X}_i\mathrm{G}_{0,i}^{-1}\mathrm{X}_i^\top - S_i^{\mathrm{p}}[\mathrm{X}_i]} \prod_{k\in\partial i}\int \mathscr{D}[\mathrm{X}_k] e^{-\frac{1}{2}\left(\mathrm{X}_k-\mathrm{M}_{k\setminus i}\right)\mathrm{G}_{k\setminus i}^{-1}\left(\mathrm{X}_k-\mathrm{M}_{k\setminus i}\right)^\top} e^{-\mathrm{X}_i\mathrm{J}_{ik}\mathrm{X}_k^\top} \\
&\quad\times \int \mathscr{D}[\mathrm{H}_i]\delta\left(K\mathrm{H}_i - \sum_{j\in\partial i}\mathrm{X}_j\right) && \text{(C.35a)} \\
&\propto e^{-\frac{1}{2}\mathrm{X}_i\mathrm{G}_{0,i}^{-1}\mathrm{X}_i^\top - S_i^{\mathrm{p}}[\mathrm{X}_i]} \prod_{k\in\partial i}\int \mathscr{D}[\mathrm{X}_k] e^{-\frac{1}{2}\left(\mathrm{X}_k-\mathrm{M}_{k\setminus i}\right)\mathrm{G}_{k\setminus i}^{-1}\left(\mathrm{X}_k-\mathrm{M}_{k\setminus i}\right)^\top} e^{-\mathrm{X}_i\mathrm{J}\mathrm{X}_k^\top} \\
&\quad\times \int \mathscr{D}[\mathrm{H}_i,\hat{\mathrm{H}}_i] e^{-\mathrm{i}\hat{\mathrm{H}}_i\left(K\mathrm{H}_i-\sum_{k\in\partial i}\mathrm{X}_k\right)^\top} && \text{(C.35b)} \\
&\propto e^{-\frac{1}{2}\mathrm{X}_i\mathrm{G}_{0,i}^{-1}\mathrm{X}_i^\top - S_i^{\mathrm{p}}[\mathrm{X}_i]} \int \mathscr{D}[\mathrm{H}_i,\hat{\mathrm{H}}_i] e^{-K\mathrm{i}\hat{\mathrm{H}}_i\mathrm{H}_i^\top} \\
&\quad\times \prod_{k\in\partial i}\int \mathscr{D}[\mathrm{X}_k] e^{-\frac{1}{2}\left(\mathrm{X}_k-\mathrm{M}_{k\setminus i}\right)\mathrm{G}_{k\setminus i}^{-1}\left(\mathrm{X}_k-\mathrm{M}_{k\setminus i}\right)^\top + \left(\mathrm{i}\hat{\mathrm{H}}_i-\mathrm{J}^\top\mathrm{X}_i\right)\mathrm{X}_k^\top}. && \text{(C.35c)}
\end{aligned}$$

Performing the Gaussian integral and introducing $\mathrm{M}_i = \sum_{k\in\partial i}\mathrm{M}_{k\setminus i}$ and $\mathrm{G}_i = \sum_{k\in\partial i}\mathrm{G}_{k\setminus i}$, we get

$$\begin{aligned}
c_i[\mathrm{X}_i] \propto e^{-\frac{1}{2}\mathrm{X}_i\mathrm{G}_{0,i}^{-1}\mathrm{X}_i^\top - S_i^{\mathrm{p}}[\mathrm{X}_i]} \int \mathscr{D}[\mathrm{H}_i,\hat{\mathrm{H}}_i] e^{-K\mathrm{i}\hat{\mathrm{H}}_i\mathrm{H}_i^\top + \left(\mathrm{i}\hat{\mathrm{H}}_i-\mathrm{J}^\top\mathrm{X}_i\right)\mathrm{M}_i^\top} \\
\times e^{\frac{1}{2}\left(\mathrm{i}\hat{\mathrm{H}}_i-\mathrm{J}^\top\mathrm{X}_i\right)\mathrm{G}_i\left(\mathrm{i}\hat{\mathrm{H}}_i-\mathrm{J}^\top\mathrm{X}_i\right)^\top}.
\end{aligned} \tag{C.36}$$

The integral in $\hat{\mathrm{H}}$ is also Gaussian and it can be done to obtain

$$c_i[\mathrm{X}_i] \propto \int \mathscr{D}[\mathrm{H}_i] e^{-\frac{1}{2}\mathrm{X}_i\mathrm{G}_{0,i}^{-1}\mathrm{X}_i^\top - K\mathrm{X}_i\mathrm{J}\mathrm{H}_i^\top - S_i^{\mathrm{p}}[\mathrm{X}_i]} e^{-\frac{1}{2}(K\mathrm{H}_i-\mathrm{M}_i)\mathrm{G}_i^{-1}(K\mathrm{H}_i-\mathrm{M}_i)^\top}. \tag{C.37}$$

The local marginal can be written in a conditional form, which reminds of the previously discussed adiabatic assumption,

$$c_i[\mathrm{X}_i] = \int \mathscr{D}[\mathrm{H}_i] c_i[\mathrm{X}_i|\mathrm{H}_i] p[\mathrm{H}_i] \tag{C.38}$$

$$c_i[\mathrm{X}_i|\mathrm{H}_i] = \frac{1}{Z_{\mathrm{X}}[\mathrm{H}_i]} e^{-\frac{1}{2}\mathrm{X}_i \mathrm{G}_{0,i}^{-1} \mathrm{X}_i^\top - K\mathrm{X}_i \mathrm{J} \mathrm{H}_i^\top - S_i^{\mathrm{p}}[\mathrm{X}_i]} \tag{C.39}$$

$$p[\mathrm{H}_i] = \frac{1}{Z_{\mathrm{H}}} e^{-\frac{1}{2}(K\mathrm{H}_i - \mathrm{M}_i)\mathrm{G}_i^{-1}(K\mathrm{H}_i - \mathrm{M}_i)^\top}, \tag{C.40}$$

where the partition functions are

$$Z_{\mathrm{X}}[\mathrm{H}_i] = \int \mathscr{D}[\mathrm{X}_i] e^{-\frac{1}{2}\mathrm{X}_i \mathrm{G}_{0,i}^{-1} \mathrm{X}_i^\top - K\mathrm{X}_i \mathrm{J} \mathrm{H}_i^\top - S_i^{\mathrm{p}}[\mathrm{X}_i]} \tag{C.41}$$

$$Z_{\mathrm{H}} = \left((2\pi)^{2(T+1)} \det \mathrm{G}_i \right)^{-1/2}. \tag{C.42}$$

In the case of a random regular graph with homogeneous coupling constants, there is no distinction between nodes in the thermodynamic limit, therefore we drop all the corresponding indices. It follows that the contributions from the neighbors are the same

$$\mathrm{M}_i = \sum_{k \in \partial i} \mathrm{M}_{k \setminus i} = K \mathrm{M}_{\mathrm{cav}}, \qquad \mathrm{G}_i = \sum_{k \in \partial i} \mathrm{G}_{k \setminus i} = K \mathrm{G}_{\mathrm{cav}}, \tag{C.43}$$

therefore $\mathrm{H} = (h_1, h_2)$ is a local field with Gaussian statistics, with mean $\langle \mathrm{H} \rangle = \mathrm{M}_{\mathrm{cav}}$ and variance $\langle \mathrm{H}^2 \rangle - \langle \mathrm{H} \rangle^2 = \mathrm{G}_{\mathrm{cav}}/K$.

The first two moments of x can be computed as

$$\langle x(t) \rangle = \int \mathscr{D}[\mathrm{X}] c[\mathrm{X}] x(t) \tag{C.44a}$$

$$= \int \mathscr{D}[\mathrm{H}] p[\mathrm{H}] \int \mathscr{D}[\mathrm{X}] \frac{x(t)}{Z_{\mathrm{X}}[\mathrm{H}]} e^{-\frac{1}{2}\mathrm{X}\mathrm{G}_0^{-1}\mathrm{X}^\top - K\mathrm{X}\mathrm{J}\mathrm{H}^\top - S^{\mathrm{p}}[\mathrm{X}]} \tag{C.44b}$$

$$= -\frac{1}{KJ} \left\langle \frac{1}{Z_{\mathrm{X}}[\mathrm{H}]} \frac{\delta}{\delta h_2(t)} Z_{\mathrm{X}}[\mathrm{H}] \right\rangle_{\mathrm{H}} \tag{C.44c}$$

and

$$\langle x^2(t)\rangle = \int \mathscr{D}[\mathrm{X}]c[\mathrm{X}]x^2(t) \tag{C.45a}$$

$$= \int \mathscr{D}[\mathrm{H}]p[\mathrm{H}]\int \mathscr{D}[\mathrm{X}]\frac{x^2(t)}{Z_{\mathrm{X}}[\mathrm{H}]}e^{-\frac{1}{2}\mathrm{X}\mathrm{G}_0^{-1}\mathrm{X}^\top - K\mathrm{X}\mathrm{J}\mathrm{H}^\top - S^{\mathrm{p}}[\mathrm{X}]} \tag{C.45b}$$

$$= \frac{1}{(KJ)^2}\left\langle \frac{1}{Z_{\mathrm{X}}[\mathrm{H}]}\frac{\delta^2}{\delta h_2^2(t)}Z_{\mathrm{X}}[\mathrm{H}]\right\rangle_{\mathrm{H}}. \tag{C.45c}$$

Linear dynamics with additive noise

We now focus on the simplest case of a linear dynamics with additive noise ($S^{\mathrm{p}} = 0$), where $\lambda = KJ$ is the drift coefficient and T is the noise coefficient (temperature). The partition function is easily obtained as a Gaussian integral

$$Z_{\mathrm{X}}[\mathrm{H}] = \int \mathscr{D}[\mathrm{X}]e^{-\frac{1}{2}\mathrm{X}\mathrm{G}_0^{-1}\mathrm{X}^\top - K\mathrm{X}\mathrm{J}\mathrm{H}^\top} \propto e^{\frac{K^2}{2}(\mathrm{J}^\top\mathrm{H})\mathrm{G}_0(\mathrm{J}^\top\mathrm{H})^\top}, \tag{C.46}$$

where G_0 is the free particle propagator

$$\mathrm{G}_0(t,t') = \begin{pmatrix} C_{\mathrm{f}}(t,t') & R_{\mathrm{f}}(t,t') \\ R_{\mathrm{f}}(t',t) & 0 \end{pmatrix}, \tag{C.47}$$

with $C_{\mathrm{f}}(t,t') = T/\lambda\exp(-\lambda|t-t'|)$ being the free particle correlation function and $R_{\mathrm{f}}(t,t') = \Theta(t-t')\exp(-\lambda(t-t'))$ the free particle response function. It follows that the partition function can be written as

$$Z_{\mathrm{X}}[\mathrm{H}] \propto \exp\left\{\frac{(KJ)^2}{2}\int dt\int dt'\left(h_2(t)C_{\mathrm{f}}(t,t')h_2(t') + 2h_2(t)R_{\mathrm{f}}(t,t')h_1(t')\right)\right\}, \tag{C.48}$$

from which the first and second moments of x follow

$$\langle x(t)\rangle = -\frac{1}{KJ}\left\langle \frac{1}{Z_{\mathrm{X}}[\mathrm{H}]}\frac{\delta}{\delta h_2(t)}Z_{\mathrm{X}}[\mathrm{H}]\right\rangle_{\mathrm{H}} \tag{C.49a}$$

$$= -\frac{KJ}{2}\int dt'\left(2C_{\mathrm{f}}(t,t')\langle h_2(t')\rangle_{\mathrm{H}} + 2R_{\mathrm{f}}(t,t')\langle h_1(t')\rangle_{\mathrm{H}}\right) \tag{C.49b}$$

$$= -KJ\int dt'\left(C_{\mathrm{f}}(t,t')\mathrm{i}\hat{\mu}_{\mathrm{cav}}(t')\rangle_{\mathrm{H}} + R_{\mathrm{f}}(t,t')\mu_{\mathrm{cav}}(t')\right) \tag{C.49c}$$

$$= -KJ\int dt' R_{\mathrm{f}}(t,t')\mu_{\mathrm{cav}}(t') \tag{C.49d}$$

$$\langle x^2(t)\rangle = \frac{1}{(KJ)^2}\left\langle \frac{1}{Z_{\mathrm{X}}[\mathrm{H}]}\frac{\delta^2}{\delta h_2^2(t)}Z_{\mathrm{X}}[\mathrm{H}]\right\rangle_{\mathrm{H}} \tag{C.50a}$$

$$= C_{\mathrm{f}}(t,t) + \frac{(KJ)^2}{4}\left\langle \left[\int dt'\left(2C_{\mathrm{f}}(t,t')h_2(t') + 2R_{\mathrm{f}}(t,t')h_1(t')\right)\right]^2\right\rangle_{\mathrm{H}} \tag{C.50b}$$

$$\begin{aligned} = C_{\mathrm{f}}(t,t) + \frac{(KJ)^2}{4}\int dt'dt''\,\big(&4C_{\mathrm{f}}(t,t')C_{\mathrm{f}}(t,t'')\langle h_2(t')h_2(t'')\rangle_{\mathrm{H}} \\ &+4R_{\mathrm{f}}(t,t')R_{\mathrm{f}}(t,t'')\langle h_1(t')h_1(t'')\rangle_{\mathrm{H}} + 2C_{\mathrm{f}}(t,t')R_{\mathrm{f}}(t,t'')\langle h_2(t')h_1(t'')\rangle_{\mathrm{H}} \\ &+2R_{\mathrm{f}}(t,t')C_{\mathrm{f}}(t,t'')\langle h_1(t')h_2(t'')\rangle_{\mathrm{H}}\big) \end{aligned} \tag{C.50c}$$

$$\begin{aligned} = C_{\mathrm{f}}(t,t) + (KJ)^2\int dt'dt''\,\Bigg(&R_{\mathrm{f}}(t,t')R_{\mathrm{f}}(t,t'')\frac{\tilde{C}_{\mathrm{cav}}(t',t'')}{K} \\ &+R_{\mathrm{f}}(t,t')C_{\mathrm{f}}(t,t')\frac{R_{\mathrm{cav}}(t',t'')}{K}\Bigg), \end{aligned} \tag{C.50d}$$

where $C^{\mathrm{dc}}_{\mathrm{cav}}(t',t'') = K\langle h_1(t')h_1(t'')\rangle_{\mathrm{H}}$ is the cavity disconnected correlation. In summary, in the stationary state,

$$s^2 = \frac{T}{\lambda} + (KJ)^2\int_0^\infty d\tau\int_0^\infty ds\left(R_{\mathrm{f}}(\tau)R_{\mathrm{f}}(s)\frac{C_{\mathrm{cav}}(s-\tau)}{K} + R_{\mathrm{f}}(\tau)C_{\mathrm{f}}(s)\frac{R_{\mathrm{cav}}(s-\tau)}{K}\right). \tag{C.51}$$

It can be shown that the "adiabatic and independent" approximation of [66] corresponds to approximating the cavity quantities by the full quantities at zero time differences, i.e $C_{\mathrm{cav}}(s-\tau) \approx C(0) = s^2$ and $R_{\mathrm{cav}}(s-\tau) \approx R(0) = 0$, so that we obtain

$$s^2 = \frac{T}{\lambda} + \frac{s^2}{K}(KJ)^2\int_0^\infty d\tau e^{-\lambda\tau}\int_0^\infty ds e^{-\lambda s} \tag{C.52a}$$

$$= \frac{T}{\lambda} + \frac{s^2}{K}\left(\frac{KJ}{\lambda}\right)^2, \tag{C.52b}$$

and, since $\lambda = KJ$, we get the relation

$$s^2 = \frac{T}{J(K-1)}, \tag{C.53}$$

which can be obtained using the approximation in [66].

Bouchaud-Mézard model

We apply the previous approximation to the BM model, in which case $\lambda = KJ$ and $T = 0$ in the Gaussian part of the action. The interacting part of the action is non-Gaussian

$$S^{\mathrm{p}}[\mathrm{X}] = -\frac{1}{2}\sigma^2 \int dt x^2(t)(\mathrm{i}\hat{x}(t))^2 = -\frac{1}{8}\sigma^2 \left(\mathrm{X}\sigma_1 \mathrm{X}^\top\right)^2, \tag{C.54}$$

and it has to be computed perturbatively. In order to properly perturb the system, we have to rescale the variables as usual, defining $\mathrm{X} = (1+\delta x, \mathrm{i}\hat{x})$. The field $K\tilde{\mathrm{H}}_i = \sum_{j\in\partial i} \delta \mathrm{X}_j$ has now a different average. The site label will be dropped in the rest of the calculation. The field $\tilde{\mathrm{H}} = (h_1 - 1, h_2)$ is a Gaussian random variable with mean $\langle \tilde{\mathrm{H}} \rangle = \tilde{M}_{\mathrm{cav}}$ and variance $\langle \tilde{\mathrm{H}}^2 \rangle - \langle \tilde{\mathrm{H}} \rangle^2 = \mathrm{G}_{\mathrm{cav}}/K$, where $\tilde{\mathrm{M}}_{\mathrm{cav}} = (\mu_{\mathrm{cav}} - 1, 0)$. In summary the distribution of $\delta\mathrm{X}$ is

$$c[\delta \mathrm{X}] = \int \mathscr{D}[\tilde{\mathrm{H}}] c[\delta \mathrm{X}|\tilde{\mathrm{H}}] p[\tilde{\mathrm{H}}] \tag{C.55}$$

$$c[\delta \mathrm{X}|\tilde{\mathrm{H}}] = \frac{1}{Z_{\delta \mathrm{X}}[\tilde{\mathrm{H}}]} e^{-\frac{1}{2}\delta \mathrm{X} \mathrm{G}_0^{-1} \delta \mathrm{X}^\top - K \delta \mathrm{X} \mathrm{J} \tilde{\mathrm{H}}^\top - S^{\mathrm{p}}[\delta \mathrm{X}]} \tag{C.56}$$

$$p[\tilde{\mathrm{H}}] = \frac{1}{Z_{\tilde{\mathrm{H}}}} e^{-\frac{1}{2}\left(\tilde{\mathrm{H}} - \tilde{\mathrm{M}}_{\mathrm{cav}}\right) K \mathrm{G}_{\mathrm{cav}}^{-1} \left(\tilde{\mathrm{H}} - \tilde{\mathrm{M}}_{\mathrm{cav}}\right)^\top}, \tag{C.57}$$

where the partition functions are

$$Z_{\delta \mathrm{X}}[\tilde{\mathrm{H}}] = \int \mathscr{D}[\delta \mathrm{X}] e^{-\frac{1}{2}\delta \mathrm{X} \mathrm{G}_0^{-1} \delta \mathrm{X}^\top - K \delta \mathrm{X} \mathrm{J} \tilde{\mathrm{H}}^\top - S^{\mathrm{p}}[\delta \mathrm{X}]} \tag{C.58}$$

$$Z_{\tilde{\mathrm{H}}} = \left(\left(\frac{2\pi}{K}\right)^{2(T+1)} \det \mathrm{G}_{\mathrm{cav}} \right)^{-1/2}. \tag{C.59}$$

In terms of the field $\delta\mathrm{X}$, the interacting part of the action can be written as

$$\begin{aligned} S^{\mathrm{p}}[\delta \mathrm{X}] = -\frac{1}{2}\sigma^2 \int dt (1+\phi(t))^2 \hat{\phi}(t)^2 = -\frac{\sigma^2}{2} \int dt \hat{\phi}^2(t) \\ - \sigma^2 \int dt \hat{\phi}^2(t) \phi(t) - \frac{\sigma^2}{2} \int dt \hat{\phi}^2(t) \phi^2(t). \end{aligned} \tag{C.60}$$

The partition function $Z_{\delta \mathrm{X}}[\tilde{\mathrm{H}}]$ can be computed perturbatively. At first order in σ^2

$$
\begin{aligned}
Z_{\delta \mathrm{X}}[\tilde{\mathrm{H}}] \approx & \int \mathscr{D}[\delta \mathrm{X}] e^{-\frac{1}{2}\delta \mathrm{X} \mathrm{G}_0^{-1} \delta \mathrm{X}^{\top} - K \delta \mathrm{X} \mathrm{J} \tilde{\mathrm{H}}^{\top}} \left(1+\frac{\sigma^2}{2}\int dt \hat{\phi}^2(t)\right. \\
& \left. +\sigma^2 \int dt \hat{\phi}^2(t)\phi(t) + \frac{\sigma^2}{2}\int dt \hat{\phi}^2(t)\phi^2(t)\right) \qquad \text{(C.61a)} \\
= & \left\{1+\frac{\sigma^2}{2}\frac{1}{(KJ)^2}\int dt \frac{\delta^2}{\delta \tilde{h}_1^2(t)} - \sigma^2 \frac{1}{(KJ)^3}\int dt \frac{\delta^3}{\delta \tilde{h}_2(t)\delta \tilde{h}_1^2(t)}\right. \\
& \left. + \frac{\sigma^2}{2}\frac{1}{(KJ)^4}\int dt \frac{\delta^4}{\delta \tilde{h}_2^2(t)\delta \tilde{h}_1^2(t)}\right\} Z^0_{\delta \mathrm{X}}[\tilde{\mathrm{H}}] \qquad \text{(C.61b)} \\
= & \left\{1+\frac{\sigma^2}{2}\hat{T}_1 - \sigma^2 \hat{T}_2 + \frac{\sigma^2}{2}\hat{T}_3\right\} Z^0_{\delta \mathrm{X}}[\tilde{\mathrm{H}}], \qquad \text{(C.61c)}
\end{aligned}
$$

where $\hat{T}_1$, $\hat{T}_2$ and $\hat{T}_3$ are operators acting on the Gaussian single-particle partition function

$$
Z^0_{\delta \mathrm{X}}[\tilde{\mathrm{H}}] = \int \mathscr{D}[\delta \mathrm{X}] e^{-\frac{1}{2}\delta \mathrm{X} \mathrm{G}_0^{-1} \delta \mathrm{X}^{\top} - K \delta \mathrm{X} \mathrm{J} \tilde{\mathrm{H}}^{\top}} = e^{-\tilde{S}^0[\tilde{\mathrm{H}}]} \qquad \text{(C.62)}
$$

$$
\tilde{S}^0[\tilde{\mathrm{H}}] = -(KJ)^2 \int dt_1 \int dt_2 \tilde{h}_2(t_1) R_{\mathrm{f}}(t_1,t_2) \tilde{h}_1(t_2), \qquad \text{(C.63)}
$$

where we have used $C_{\mathrm{f}}(t,t') = 0$ for the BM model. The operators act on it as

$$
\begin{aligned}
\hat{T}_1 Z^0_{\delta \mathrm{X}}[\tilde{\mathrm{H}}] &= (KJ)^2 \int dt \int dt_1 R_{\mathrm{f}}(t_1,t)\tilde{h}_2(t_1) \int dt_2 R_{\mathrm{f}}(t_2,t)\tilde{h}_2(t_2) Z^0_{\delta \mathrm{X}}[\tilde{\mathrm{H}}] \\
&= T_1\left[\tilde{\mathrm{H}}\right] Z^0_{\delta \mathrm{X}}[\tilde{\mathrm{H}}] \qquad \text{(C.64)} \\
\hat{T}_2 Z^0_{\delta \mathrm{X}}[\tilde{\mathrm{H}}] &= (KJ)^3 \int dt \int dt_1 R_{\mathrm{f}}(t_1,t)\tilde{h}_2(t_1) \int dt_2 R_{\mathrm{f}}(t_2,t)\tilde{h}_2(t_2) \\
&\quad \times \int dt_3 R_{\mathrm{f}}(t,t_3)\tilde{h}_1(t_3) Z^0_{\delta \mathrm{X}}[\tilde{\mathrm{H}}] = T_2\left[\tilde{\mathrm{H}}\right] Z^0_{\delta \mathrm{X}}[\tilde{\mathrm{H}}] \qquad \text{(C.65)} \\
\hat{T}_3 Z^0_{\delta \mathrm{X}}[\tilde{\mathrm{H}}] &= (KJ)^4 \int dt \int dt_1 R_{\mathrm{f}}(t_1,t)\tilde{h}_2(t_1) \int dt_2 R_{\mathrm{f}}(t_2,t)\tilde{h}_2(t_2) \\
&\quad \times \int dt_3 \tilde{h}_1(t_3) R_{\mathrm{f}}(t,t_3) \int dt_4 R_{\mathrm{f}}(t,t_4)\tilde{h}_1(t_4) Z^0_{\delta \mathrm{X}}[\tilde{\mathrm{H}}] \\
&= T_3\left[\tilde{\mathrm{H}}\right] Z^0_{\delta \mathrm{X}}[\tilde{\mathrm{H}}], \qquad \text{(C.66)}
\end{aligned}
$$

where we have used $R_{\mathrm{f}}(t,t)=0$ in the Itô case. In summary,

$$Z_{\delta \mathrm{X}}[\tilde{\mathrm{H}}] \approx Z^0_{\delta \mathrm{X}}[\tilde{\mathrm{H}}]\left(1+\frac{\sigma^2}{2}T_1[\tilde{\mathrm{H}}]-\sigma^2 T_2[\tilde{\mathrm{H}}]+\frac{\sigma^2}{2}T_3[\tilde{\mathrm{H}}]\right) \tag{C.67a}$$

$$\approx e^{-\tilde{S}^0[\tilde{\mathrm{H}}]+\frac{\sigma^2}{2}T_1[\tilde{\mathrm{H}}]-\sigma^2 T_2[\tilde{\mathrm{H}}]+\frac{\sigma^2}{2}T_3[\tilde{\mathrm{H}}]} \tag{C.67b}$$

$$= e^{-\tilde{S}[\tilde{\mathrm{H}}]}. \tag{C.67c}$$

It is now possible to compute the first and second moments of the field ϕ. We neglect contributions in the local self-consistent field beyond second order. The mean is

$$\langle \phi(t)\rangle = -\frac{1}{KJ}\left\langle \frac{1}{Z_{\delta \mathrm{X}}[\tilde{\mathrm{H}}]}\frac{\delta}{\delta \tilde{h}_2(t)}Z_{\delta \mathrm{X}}[\tilde{\mathrm{H}}]\right\rangle_{\tilde{\mathrm{H}}} \tag{C.68a}$$

$$= -(KJ)\left(\int dt' R_{\mathrm{f}}(t,t')\langle \tilde{h}_1(t')\rangle_{\tilde{\mathrm{H}}}\right.$$
$$\left. -2\sigma^2\int dt_1 R_{\mathrm{f}}(t,t_1)\int dt_2 R_{\mathrm{f}}(t_2,t_1)\int dt_3 R_{\mathrm{f}}(t_1,t_3)\langle \tilde{h}_2(t_2)\tilde{h}_1(t_3)\rangle_{\tilde{\mathrm{H}}}\right) \tag{C.68b}$$

$$= 2\sigma^2(KJ)\int dt_1 R_{\mathrm{f}}(t,t_1)\int dt_2 R_{\mathrm{f}}(t_2,t_1)\int dt_3 R_{\mathrm{f}}(t_1,t_3)\frac{R_{\mathrm{cav}}(t_3,t_2)}{K}=0 \tag{C.68c}$$

due to causality ($R_{\mathrm{f}}(t,t')=R_{\mathrm{cav}}(t,t')=0$ if $t'>t$). Computing the second moment is more complicated. We use that

$$\langle \phi(t)^2\rangle = \frac{1}{(KJ)^2}\left\langle \frac{1}{Z_{\delta \mathrm{X}}[\tilde{\mathrm{H}}]}\frac{\delta^2}{\delta \tilde{h}_2^2(t)}Z_{\delta \mathrm{X}}[\tilde{\mathrm{H}}]\right\rangle_{\tilde{\mathrm{H}}} \tag{C.69a}$$

$$= \frac{1}{(KJ)^2}\left[\left\langle \left(\frac{\delta}{\delta \tilde{h}_2(t)}\tilde{S}[\tilde{\mathrm{H}}]\right)^2\right\rangle_{\tilde{\mathrm{H}}} - \left\langle \frac{\delta^2}{\delta \tilde{h}_2^2(t)}\tilde{S}[\tilde{\mathrm{H}}]\right\rangle_{\tilde{\mathrm{H}}}\right], \tag{C.69b}$$

where the first average is given by

$$\left\langle \left(\frac{\delta}{\delta \tilde{h}_2(t)} \tilde{S}[\tilde{\mathrm{H}}] \right)^2 \right\rangle_{\tilde{\mathrm{H}}} = (KJ)^4 \Big\{ \int dt_1 \int dt_2 R_{\mathrm{f}}(t,t_1) R_{\mathrm{f}}(t,t_2) \langle \tilde{h}_1(t_1) \tilde{h}_1(t_2) \rangle_{\tilde{\mathrm{H}}} + 2\sigma^2 \int dt_1 \int dt_2 R_{\mathrm{f}}(t,t_2) \int dt_3 R_{\mathrm{f}}(t,t_3) R_{\mathrm{f}}(t_1,t_3) \langle \tilde{h}_2(t_1) \tilde{h}_1(t_2) \rangle_{\tilde{\mathrm{H}}} \Big\} \tag{C.70a}$$

$$= (KJ)^4 \left\{ \int dt_1 \int dt_2 R_{\mathrm{f}}(t,t_1) R_{\mathrm{f}}(t,t_2) \frac{C_{\mathrm{cav}}(t_1,t_2)}{K} + 2\sigma^2 \int dt_1 \int dt_2 R_{\mathrm{f}}(t,t_2) \int dt_3 R_{\mathrm{f}}(t,t_3) R_{\mathrm{f}}(t_1,t_3) \frac{R_{\mathrm{cav}}(t_2,t_1)}{K} \right\}, \tag{C.70b}$$

and the second one is

$$\left\langle \frac{\delta^2}{\delta \tilde{h}_2^2(t)} \tilde{S}[\tilde{\mathrm{H}}] \right\rangle_{\tilde{\mathrm{H}}} = (KJ)^2 \sigma^2 \int dt_1 R_{\mathrm{f}}^2(t,t_1) \Big(1 + (KJ)^2 \int dt_2 \int dt_3 R_{\mathrm{f}}(t_1,t_2) R_{\mathrm{f}}(t_1,t_3) \langle \tilde{h}_1(t_2) \tilde{h}_1(t_3) \rangle_{\tilde{\mathrm{H}}} \Big) \tag{C.71a}$$

$$= (KJ)^2 \sigma^2 \int dt_1 R_{\mathrm{f}}^2(t,t_1) \Big(1 + (KJ)^2 \int dt_2 \int dt_3 R_{\mathrm{f}}(t_1,t_2) R_{\mathrm{f}}(t_1,t_3) \frac{C_{\mathrm{cav}}(t_1,t_2)}{K} \Big). \tag{C.71b}$$

In summary, at stationarity,

$$s^2 = \langle x^2 \rangle - \langle x \rangle^2 = \langle \phi^2 \rangle \tag{C.72a}$$

$$= \sigma^2 \int_0^\infty d\tau R_{\mathrm{f}}^2(\tau) + (KJ)^2 \int_0^\infty d\tau R_{\mathrm{f}}(\tau) \int_0^\infty ds R_{\mathrm{f}}(s) \frac{C_{\mathrm{cav}}(s-\tau)}{K} + \sigma^2 (KJ)^2 \int_0^\infty d\tau \int_0^\infty ds \int_0^{\min(\tau,s)} d\theta R_{\mathrm{f}}^2(\theta) R_{\mathrm{f}}(\tau-\theta) R_{\mathrm{f}}(s-\theta) \frac{C_{\mathrm{cav}}(\tau-\theta)}{K} + 2\sigma^2 (KJ)^2 \int_0^\infty d\tau R_{\mathrm{f}}(\tau) \int_0^\tau ds R_{\mathrm{f}}(\tau-s) \int_0^s d\theta R_{\mathrm{f}}(\theta) \frac{R_{\mathrm{cav}}(s-\theta)}{K}. \tag{C.72b}$$

Making the choice of response and correlation functions that corresponds to the adiabatic and independent approximation of [66], i.e setting $C_{\mathrm{cav}}(s-\tau) \approx C(0) = s^2$ and $R_{\mathrm{cav}}(s-\tau) \approx R(0) = 0$, we obtain

$$s^2 = \frac{\sigma^2}{2KJ} + \left(1 + \frac{\sigma^2}{2KJ} \right) \frac{s^2}{K}. \tag{C.73}$$

This formula leads to a wrong prediction of the critical point $\sigma_c^2 = 2KJ(K-1)$. This is due to the fact we neglected relevant higher-order contributions in the loop expansion generated by the noise term. Performing the resummation of all terms of the geometric series expansion containing chained one-loop diagrams only, the previous relation becomes

$$s^2 = \left(1 + \frac{\sigma^2}{2KJ} + \left(\frac{\sigma^2}{2KJ}\right)^2 + \dots\right)\left(1 + \frac{s^2}{K}\right) - 1 \tag{C.74a}$$

$$= \left(1 - \frac{\sigma^2}{2KJ}\right)^{-1}\left(1 + \frac{s^2}{K}\right) - 1, \tag{C.74b}$$

from which we get the same prediction $\sigma_c^2 = 2J(K-1)$ already obtained in [66].

Appendix D

Detailed derivations for Complex Ecosystems

D.1 Small-Coupling Expansion of the Cavity Equations

The small-coupling expansion provides a systematic approximation scheme to derive closed equations for local observables from the dynamic cavity formalism. In the weak-coupling regime, interactions between neighboring nodes can be treated perturbatively up to second order in the coupling strengths, leading to an effective single-species process driven by self-consistent Gaussian fields. In this section, we present the detailed steps of this derivation. First, in section D.1.1, we construct the MSRJD generating functional without assuming $x_i(t) > 0$, and we obtain the corresponding effective cavity process. Then, in section D.1.2, we derive explicitly the closed fixed-point equations governing the stationary properties of the system.

D.1.1 Correct derivation of the effective cavity process

The derivation presented here follows the general approach introduced in chapter 4, but avoids assuming that $x_i(t) > 0$ at all times, thereby providing a correct formulation of the MSRJD generating functional. We show that the resulting effective cavity process coincides with the one obtained through the simplified derivation in section 7.2.

The dynamical generating functional of the generalized Lotka-Volterra system can be written as usual through the MSRJD path integral

$$Z[\underline{\psi},\underline{\psi}] = \int \mathscr{D}^0[\underline{x},\underline{\hat{x}}] \prod_{i=1}^{N} e^{-\mathrm{i}\int dt\, \hat{x}_i(t)\left[\dot{x}_i(t)-x_i(t)\left(1-x_i(t)+\sum_{j\in\partial i} J_{ij}x_j(t)\right)\right]} \prod_{i=1}^{N} e^{\int dt\,(\psi_i(t)x_i(t)+\hat{\psi}_i(t)\mathrm{i}\hat{x}_i(t))}, \tag{D.1}$$

where we used the Fourier representation of the functional delta. We introduce auxiliary variables $f_i(t) = 1 - x_i(t) + \sum_{j\in\partial i} J_{ij}x_j(t)$, enforced via Dirac delta functionals expressed through their Fourier representations. The generating functional becomes

$$Z[\underline{\psi},\underline{\psi}] = \int \mathscr{D}^0[\underline{x},\underline{\hat{x}},\underline{f},\underline{\hat{f}}] \prod_{i=1}^{N} e^{-\mathrm{i}\int dt\left[\hat{x}_i(t)(\dot{x}_i(t)-x_i(t)f_i(t))+\hat{f}_i(t)(f_i(t)-(1-x_i(t)))\right]} \times \prod_{i=1}^{N}\prod_{j\in\partial i} e^{\mathrm{i}\int dt\,\hat{f}_i(t)J_{ij}x_j(t)} \prod_{i=1}^{N} e^{\int dt\,(\psi_i(t)x_i(t)+\hat{\psi}_i(t)\mathrm{i}\hat{x}_i(t))}, \tag{D.2}$$

which can be written in the compact form

$$Z[\underline{\psi},\underline{\psi}] = \int \mathscr{D}^0[\underline{x},\underline{\hat{x}},\underline{f},\underline{\hat{f}}] \exp\left(-\sum_{i=1}^{N} S_i^{\mathrm{loc}} - \sum_{(i,j)\in E} S_{ij}^{\mathrm{int}}\right), \tag{D.3}$$

with local and interaction actions

$$S_i^{\mathrm{loc}} = \mathrm{i}\int dt \left[\hat{x}_i(t)\Big(\dot{x}_i(t) - x_i(t)f_i(t)\Big) + \hat{f}_i(t)\Big(f_i(t) - \big(1 - x_i(t) + h_i(t)\big)\Big)\right], \tag{D.4}$$

$$S_{ij}^{\mathrm{int}} = -\mathrm{i}\int dt \left(J_{ij}\hat{f}_i(t)x_j(t) + J_{ji}\hat{f}_j(t)x_i(t)\right). \tag{D.5}$$

The factorized structure of the generating functional defines a graphical model over continuous trajectories, with edges corresponding to interaction terms. Applying the dynamic cavity formalism yields the recursion relations for the cavity marginals:

$$c_{i\setminus j}[\mathrm{X}_i,\mathrm{F}_i] = \frac{p_0(x_i(0))}{Z_{i\setminus j}} e^{S_i^{\mathrm{loc}}} \prod_{k\in\partial i\setminus j} \int \mathscr{D}[\mathrm{X}_k,\mathrm{F}_k]\, c_{k\setminus i}[\mathrm{X}_k,\mathrm{F}_k]\, e^{S_{ki}^{\mathrm{int}}}, \tag{D.6}$$

where we have introduced, for shortness of notation, the vectors $\mathrm{X}_i = (x_i, \hat{x}_i)$ and $\mathrm{F}_i = (f_i, \hat{f}_i$. Once the messages have converged, the full marginal of node i can be reconstructed as

$$c_i[\mathrm{X}_i, \mathrm{F}_i] = \frac{p_0(x_i(0))}{Z_{i\setminus j}} e^{S_i^{\mathrm{loc}}} \prod_{k\in\partial i} \int \mathscr{D}[\mathrm{X}_k, \mathrm{F}_k]\, c_{k\setminus i}[\mathrm{X}_k, \mathrm{F}_k]\, e^{S_{ki}^{\mathrm{int}}}, \tag{D.7}$$

The normalization factors $Z_{i\setminus j}$ and Z_i are fixed by imposing proper normalization of the cavity distributions.

To derive a tractable expression for the effective single-site dynamics, we perform a second-order expansion in the coupling parameters J_{ij} of the cavity recursion in Eq. (D.6), following the Gaussian-expansion method of chapter 4. Expanding to second order and neglecting averages involving only response fields, which vanish identically in the MSRJD formalism, we find

$$\begin{aligned} c_{i\setminus j}[\mathrm{X}_i, \mathrm{F}_i] \approx \frac{p_0(x_i(0))}{Z_{i\setminus j}} e^{S_i^{\mathrm{loc}}} \prod_{k\in\partial i\setminus j} \Bigg[1 + J_{ik} \int dt\, \mathrm{i}\hat{f}_i(t) \langle x_k(t)\rangle_{k\setminus i} \\ \int dt dt'\, J_{ik}J_{ki}\, \mathrm{i}\hat{f}_i(t) \langle x_k(t)\, \mathrm{i}\hat{f}_k(t')\rangle_{k\setminus i} x_i(t') \\ + \frac{1}{2} \int dt dt'\, \mathrm{i}\hat{f}_i(t) J_{ik}^2 \langle x_k(t) x_k(t')\rangle_{k\setminus i} \mathrm{i}\hat{f}_i(t') + \dots \Bigg]. \end{aligned} \tag{D.8}$$

Re-exponentiating the terms inside the product yields the approximate Gaussian form

$$c_{i\setminus j}[\mathrm{X}_i, \mathrm{F}_i] \approx \frac{p_0(x_i(0))}{Z_{i\setminus j}} \exp\left(-S_{i\setminus j}^{\mathrm{eff}}\right), \tag{D.9}$$

where the effective single-site action reads

$$\begin{aligned} S_{i\setminus j}^{\mathrm{eff}} = \mathrm{i} \int dt \Bigg[\hat{x}_i(t) \Big(\dot{x}_i(t) - x_i(t) f_i(t) \Big) + \hat{f}_i(t) \Big(f_i(t) - \big(1 - x_i(t) + h_i(t)\big) \\ + \sum_{k\in\partial i\setminus j} J_{ik} \mu_{k\setminus i}(t) + \sum_{k\in\partial i\setminus j} J_{ik}J_{ki} \int dt'\, R_{k\setminus i}(t,t') x_i(t') \Big) \\ + \int dt dt'\, \mathrm{i}\hat{f}_i(t) \sum_{k\in\partial i\setminus j} J_{ik}^2 C_{k\setminus i}(t,t')\, \mathrm{i}\hat{f}_i(t'). \end{aligned} \tag{D.10}$$

Here, the quantities

$$\mu_{i\setminus j}(t) = \langle x_i(t)\rangle_{i\setminus j}, \tag{D.11}$$

$$C_{i\setminus j}(t,t') = \langle x_i(t)x_i(t')\rangle_{i\setminus j} - \mu_{i\setminus j}(t)\mu_{i\setminus j}(t'), \tag{D.12}$$

$$R_{i\setminus j}(t,t') = \langle x_i(t)\,\mathrm{i}\hat{x}_i(t')\rangle_{i\setminus j} = \left.\frac{\delta\langle x_i(t)\rangle_{i\setminus j}}{\delta h_i(t')}\right|_{h=0}. \tag{D.13}$$

represent, respectively, the mean, response, and correlation functions of the neighboring cavity processes.

Reversing the steps of the path-integral construction and changing variables $x_i \to x_{i\setminus j}$ in order to make explicit the cavity dependence of the variable, the effective stochastic dynamics for node i on the cavity graph where node j is removed reads

$$\frac{dx_{i\setminus j}(t)}{dt} = x_{i\setminus j}(t)\Bigg[1 - x_{i\setminus j}(t) + \sum_{k\in\partial i\setminus j} J_{ik}\mu_{k\setminus i}(t) + h_{i\setminus j}(t) + \sum_{k\in\partial i\setminus j} J_{ik}J_{ki}\int_0^t dt'\, R_{k\setminus i}(t,t')\,x_{i\setminus j}(t') + \eta_{i\setminus j}(t)\Bigg], \tag{D.14}$$

where the effective noise $\eta_{i\setminus j}(t)$ is Gaussian with zero mean and covariance

$$\langle \eta_{i\setminus j}(t)\eta_{i\setminus j}(t')\rangle_{i\setminus j} = \sum_{k\in\partial i\setminus j} J_{ik}^2 C_{k\setminus i}(t,t'). \tag{D.15}$$

An analogous expression holds for the full dynamics of node i:

$$\frac{dx_i(t)}{dt} = x_i(t)\Bigg[1 - x_i(t) + \sum_{k\in\partial i} J_{ik}\mu_{k\setminus i}(t) + \sum_{k\in\partial i} J_{ik}J_{ki}\int_0^t dt'\, R_{k\setminus i}(t,t')\,x_i(t') + \eta_i(t)\Bigg], \tag{D.16}$$

with Gaussian noise satisfying

$$\langle \eta_i(t)\eta_i(t')\rangle_i = \sum_{k\in\partial i} J_{ik}^2 C_{k\setminus i}(t,t'). \tag{D.17}$$

The stochastic process derived above is thus fully equivalent to the heuristic derivation reported in section 7.2, confirming the internal consistency of the Gaussian-expansion and its validity even when the positivity of $x_i(t)$ is not explicitly enforced.

D.1.2 Derivation of the fixed point cavity equations

We provide here a detailed derivation of the cavity fixed point equations Eqs. (7.31) to (7.33). We start by assuming that the dynamics relaxes to a time-independent state, such that

$$\lim_{t\to\infty} x_i(t) = x_i^\star. \tag{D.18}$$

An analogous assumption is made for the cavity effective processes,

$$\lim_{t\to\infty} x_{i\backslash j}(t) = x_{i\backslash j}^\star, \qquad \lim_{t\to\infty} \xi_{i\backslash j}(t) = \xi_{i\backslash j}^\star. \tag{D.19}$$

Consequently, the mean cavity abundance and the two-point correlation function become stationary,

$$\lim_{t\to\infty} \mu_{i\backslash j}(t) = \mu_{i\backslash j}^\star = \langle x_{i\backslash j}^\star \rangle_{i\backslash j}^\star, \tag{D.20}$$

$$\lim_{t,t'\to\infty} C_{i\backslash j}(t,t') = q_{i\backslash j}^\star = \langle (x_{i\backslash j}^\star)^2 \rangle_{i\backslash j}^\star - (\mu_{i\backslash j}^\star)^2, \tag{D.21}$$

where $\langle\cdot\rangle_{i\backslash j}^\star$ denotes the average over the stationary cavity process. Under this assumption, the cavity response function becomes time-translationally invariant and depends only on the time difference $\tau = t - t'$. We define its time integral,

$$\chi_{i\backslash j}^\star = \int_0^\infty d\tau R_{i\backslash j}^\star(\tau), \tag{D.22}$$

which represents the integrated response (or susceptibility) of the cavity process.

At the fixed point, both $x_{i\backslash j}^\star$ and the effective field $\xi_{i\backslash j}^\star$ become static random variables. The variable $x_{i\backslash j}^\star$ depends functionally on $\xi_{i\backslash j}^\star$, which is Gaussian distributed with zero mean and variance

$$\langle (\xi_{i\backslash j}^\star)^2 \rangle_{i\backslash j}^\star = \sum_{k\in\partial i\backslash j} J_{ik}^2 q_{k\backslash i}^\star \equiv \Sigma_{i\backslash j}^2. \tag{D.23}$$

Since $x_{i\backslash j}^\star$ is constant in time, the fixed point abundances satisfy Eq. (7.14) with vanishing time derivative, i.e.

$$0 = x_{i\backslash j}^\star \left[1 - x_{i\backslash j}^\star + \sum_{k\in\partial i\backslash j} J_{ik}\mu_{k\backslash i}^\star + \sum_{k\in\partial i\backslash j} J_{ik}J_{ki}\chi_{k\backslash i}^\star x_{i\backslash j}^\star + \xi_{i\backslash j}^\star + h_{i\backslash j} \right]. \tag{D.24}$$

These equations admit two possible solutions: one corresponding to extinction, $x^\star_{i\backslash j} = 0$, and the other to a non-vanishing equilibrium abundance,

$$x^\star_{i\backslash j} = \frac{1 + \sum_{k\in\partial i\backslash j} J_{ik}\mu^\star_{k\backslash i} + \xi^\star_{i\backslash j} + h_{i\backslash j}}{1 - \sum_{k\in\partial i\backslash j} J_{ik}J_{ki}\chi^\star_{k\backslash i}}, \tag{D.25}$$

which is physically admissible only if the right-hand side is positive. The fixed-point solution can thus be written as

$$x^\star_{i\backslash j}(\xi^\star_{i\backslash j}) = \frac{1 + \sum_{k\in\partial i\backslash j} J_{ik}\mu^\star_{k\backslash i} + \xi^\star_{i\backslash j} + h_{i\backslash j}}{1 - \sum_{k\in\partial i\backslash j} J_{ik}J_{ki}\chi^\star_{k\backslash i}}\,\Theta\left(\frac{1 + \sum_{k\in\partial i\backslash j} J_{ik}\mu^\star_{k\backslash i} + \xi^\star_{i\backslash j} + h_{i\backslash j}}{1 - \sum_{k\in\partial i\backslash j} J_{ik}J_{ki}\chi^\star_{k\backslash i}}\right). \tag{D.26}$$

We now introduce the short-hand notations

$$\Delta_{i\backslash j} = \frac{1 + \sum_{k\in\partial i\backslash j} J_{ik}\mu^\star_{k\backslash i}}{\Sigma_{i\backslash j}}, \tag{D.27}$$

$$\Gamma_{i\backslash j} = 1 - \sum_{k\in\partial i\backslash j} J_{ik}J_{ki}\chi^\star_{k\backslash i}. \tag{D.28}$$

Since $\xi^\star_{i\backslash j}$ is Gaussian with mean zero and variance $\Sigma^2_{i\backslash j}$, we set $\xi^\star_{i\backslash j} = -\Sigma_{i\backslash j}\,s$, where s is a standard normal variable with zero mean and unit variance. In this way, the fixed point can be expressed as a function of s only:

$$x^\star_{i\backslash j}(s) = \frac{\Sigma_{i\backslash j}(\Delta_{i\backslash j} - s) + h_{i\backslash j}}{\Gamma_{i\backslash j}}\,\Theta\left(\frac{\Sigma_{i\backslash j}(\Delta_{i\backslash j} - s) + h_{i\backslash j}}{\Gamma_{i\backslash j}}\right). \tag{D.29}$$

The random variable $x^\star_{i\backslash j}$ has moments $\mu^\star_{i\backslash j}$, $q^\star_{i\backslash j}$, and susceptibility $\chi^\star_{i\backslash j}$, which are obtained self-consistently by averaging over the Gaussian variable s. The mean reads

$$\begin{aligned}\mu^\star_{i\backslash j} &= \langle x^\star_{i\backslash j}\rangle^\star_{i\backslash j}\Big|_{h_{i\backslash j}=0} = \int \frac{ds\,e^{-s^2/2}}{\sqrt{2\pi}}\,x^\star_{i\backslash j}(s)\Big|_{h_{i\backslash j}=0} \\ &= \frac{\Sigma_{i\backslash j}}{\Gamma_{i\backslash j}}\left[\Theta(\Gamma_{i\backslash j})\int_{-\infty}^{\Delta_{i\backslash j}} Ds\,(\Delta_{i\backslash j} - s) + \Theta(-\Gamma_{i\backslash j})\int_{\Delta_{i\backslash j}}^{\infty} Ds\,(\Delta_{i\backslash j} - s)\right],\end{aligned} \tag{D.30}$$

while the variance is

$$q^\star_{i\setminus j} = \langle (x^\star_{i\setminus j})^2 \rangle^\star_{i\setminus j}\Big|_{h_{i\setminus j}=0} - (\mu^\star_{i\setminus j})^2 = \int \frac{ds\, e^{-s^2/2}}{\sqrt{2\pi}} \left(x^\star_{i\setminus j}(s)\right)^2 \Big|_{h_{i\setminus j}=0} - (\mu^\star_{i\setminus j})^2$$
$$= \frac{\Sigma^2_{i\setminus j}}{\Gamma^2_{i\setminus j}} \left[\Theta(\Gamma_{i\setminus j}) \int_{-\infty}^{\Delta_{i\setminus j}} Ds\, (\Delta_{i\setminus j} - s)^2 + \Theta(-\Gamma_{i\setminus j}) \int_{\Delta_{i\setminus j}}^{\infty} Ds\, (\Delta_{i\setminus j} - s)^2 \right] - (\mu^\star_{i\setminus j})^2, \quad \text{(D.31)}$$

where $Ds = ds\, e^{-s^2/2}/\sqrt{2\pi}$. The susceptibility is obtained by differentiating the average fixed-point abundance with respect to a constant external field $h_{i\setminus j}$:

$$\chi^\star_{i\setminus j} = \frac{\partial \langle x^\star_{i\setminus j} \rangle^\star_{i\setminus j}}{\partial h_{i\setminus j}}\Big|_{h_{i\setminus j}=0} = \frac{1}{\Gamma_{i\setminus j}} \left[\Theta(\Gamma_{i\setminus j}) \int_{-\infty}^{\Delta_{i\setminus j}} Ds + \Theta(-\Gamma_{i\setminus j}) \int_{\Delta_{i\setminus j}}^{\infty} Ds \right]. \quad \text{(D.32)}$$

We observe that these integrals can be written in terms of the error function $\mathrm{erf}(x) = \int_0^x ds\, e^{-s^2} 2/\sqrt{\pi}$. Defining $\varphi(x) = e^{-x^2/2}/\sqrt{2\pi}$ as the standard normal density and $\Phi(x) = [1 + \mathrm{erf}(x/\sqrt{2})]/2$ its cumulative distribution function, the integrals can be expressed as moments of truncated normal distributions. In particular,

$$\int_{-\infty}^{\Delta} Ds = \Phi(\Delta), \qquad \int_{\Delta}^{\infty} Ds = \Phi(-\Delta),$$
$$\int_{-\infty}^{\Delta} Ds\, s = \varphi(\Delta), \qquad \int_{\Delta}^{\infty} Ds\, s = \varphi(\Delta),$$
$$\int_{-\infty}^{\Delta} Ds\, s^2 = \Phi(\Delta) - \Delta\varphi(\Delta), \qquad \int_{\Delta}^{\infty} Ds\, s^2 = \Phi(-\Delta) + \Delta\varphi(\Delta),$$

from which, substituting into Eqs. (D.30) to (D.32), we recover Eqs. (7.31) to (7.36).

D.2 Implementation details for the Population Dynamics algorithm

The population dynamics algorithm employed throughout this chapter is implemented in the open-source Julia package `RandomLotkaVolterraCavity.jl` [48]. Its structure follows the procedure outlined in section 7.3.1, with additional numerical optimizations introduced to improve stability and convergence.

A key modification concerns the way population elements are updated at each iteration. Instead of replacing a single randomly-chosen element after every update, as in the standard population dynamics scheme, we proceed by generating a random permutation of the population indices at the beginning of each iteration (or *sweep*). The algorithm then sequentially updates each element according to this shuffled order, ensuring that every population member is updated exactly once per sweep while minimizing correlations arising from a fixed update sequence. For each element in the shuffled population, the triplet $(\mu^{\mathrm{c}}, q^{\mathrm{c}}, \chi^{\mathrm{c}})$ is updated according to the self-consistent functions $\mathscr{F}_\mu$, $\mathscr{F}_q$, and $\mathscr{F}_\chi$ defined in section 7.3.1.

The update is performed softly by introducing a damping factor $d \in (0,1)$, which mixes the newly computed values with the previous ones:

$$\mu^{\mathrm{c}} \leftarrow d\,\mu^{\mathrm{c,new}} + (1-d)\,\mu^{\mathrm{c,old}}, \tag{D.33}$$

$$q^{\mathrm{c}} \leftarrow d\,q^{\mathrm{c,new}} + (1-d)\,q^{\mathrm{c,old}}, \tag{D.34}$$

$$\chi^{\mathrm{c}} \leftarrow d\,\chi^{\mathrm{c,new}} + (1-d)\,\chi^{\mathrm{c,old}}. \tag{D.35}$$

This under-relaxation procedure suppresses oscillations and improves numerical stability, although it typically slows down the convergence rate. In all the analyses presented in this work, the damping factor was chosen in the range $d \in [0.2, 0.6]$, depending on the convergence properties of the system.

After completing a full sweep over the shuffled population, convergence is checked by comparing the average values of μ, q, and χ before and after the sweep. If the relative variation of these averages falls below a fixed threshold, the algorithm is stopped and convergence is assumed. In our simulations, the tolerance threshold was set between 10^{-6} and 10^{-8}, depending on the parameter regime.

In principle, one should also verify the convergence of higher-order cumulants to ensure that the full population distribution has reached stationarity. However, this requirement substantially increases computational time. In practice, we found that monitoring the first moments provides reliable convergence for all parameter ranges explored in this chapter.

D.3 Phase diagrams for random regular graphs

This section details the numerical procedures employed to construct the stability phase diagrams presented in section 7.5. We first describe the algorithm used to identify critical points from population dynamics data, and then present the phase diagrams for the specific values of

connectivity K and interaction symmetry γ discussed in the Fig. 7.3, explaining how the critical boundaries were derived via interpolation of the numerical results.

D.3.1 Numerical determination of critical points

The phase boundaries were determined by performing systematic parameter sweeps. For the topological transition ($\sigma = 0$), we fixed the connectivity K and varied the interaction mean m. For the general phase diagrams ($\sigma > 0$), we fixed σ and scanned m around the critical region. For each parameter set, we computed the distribution of the stationary moments (μ, q, χ) with the Population Dynamics (PopDyn) algorithm, saving their average and standard deviation.

To identify the critical transition points m_c, we implemented an automated algorithm designed to separate the signal from the numerical noise floor. The procedure is illustrated in Figs. D.1 to D.3, and consists of the following steps:

1. **Baseline Characterization:** We define a baseline region using the first points of the scanned m-values, assumed to lie deep within the stable phase where the order parameter (e.g., the cavity variance q) should theoretically vanish. We compute the mean $\mathbb{E}_{\text{base}}[q]$ and standard deviation $\text{std}_{\text{base}}[q]$ of the observable q within this window to estimate the numerical noise floor.

2. **Noise Thresholds:** We define a detection threshold $T_q = \mathbb{E}_{\text{base}}[q] + 3\,\text{std}_{\text{base}}[q]$, corresponding to a 3-std deviation from the baseline noise. A similar threshold T_s is defined for the internal spread (standard deviation) of the population distribution.

3. **Transition Detection:** Scanning through the parameter space, a point m_i is flagged as a candidate transition if either of the following conditions is met for a persistence of 10 consecutive steps:

   - *Signal Shift:* The lower bound of the distribution departs from the baseline, $\mathbb{E}[q](m_i) - \text{std}[q](m_i) > T_q$, where $\mathbb{E}[q](m_i)$ and $\text{std}[q](m)$ are respectively the average and standard deviation of q among the population.
   - *Variance Explosion:* The standard deviation of q among the population exceeds the baseline fluctuation threshold, $\text{std}[q](m_i) > T_s$.

4. **Interpolation and Error Estimation:** Once the transition interval $[m_{i-1}, m_i]$ is identified, we approximate the behavior of q linearly between these points. The critical value m_c is defined as the exact point where this linear interpolation intersects the detection threshold

T. To estimate the uncertainty δm_c, we combine the grid resolution $\Delta m = m_i - m_{i-1}$ with the error propagated from the baseline noise. This definition ensures that sharp transitions have small errors dominated by the grid size, while slow, continuous transitions have larger uncertainties, reflecting the sensitivity of the critical point to numerical noise.

This methodology adapts naturally to the different physical regimes of the system. In the topological case ($\sigma = 0$), the transition is sharp: the cavity variance q exhibits a sudden jump from zero to a finite value. Our algorithm detects this discontinuity immediately, as shown in Fig. D.1.

Conversely, for the disorder-driven SE to ME transition ($\sigma > 0$), the change is slow and continuous. As illustrated in Fig. D.2 (for $K = 10, \gamma = -1$), the distribution of cavity variances spreads slowly from a delta function at zero to a broadened distribution. Here, the algorithm identifies the first value of m where the population statistics deviate sufficiently from the baseline, with the propagated error capturing the smoothness of the crossover.

Finally, the nature of the instability is classified by monitoring the mean abundance μ. If μ remains finite across the transition, it is classified as a transition to the Multiple Equilibria (ME) phase (Fig. D.2). If instead μ diverges, it is classified as a transition to the Unbounded Growth (UG) phase, as shown in Fig. D.3. This fully automated procedure ensures that the phase diagrams are derived without manual bias.

D.3.2 Phase boundaries interpolation

The stability phase diagrams presented in Fig. 7.3 and expanded here in Fig. D.4 were constructed by mapping the critical boundaries in the (m, σ) plane for fixed connectivity ($K = 3, 10$) and interaction symmetry ($\gamma = -1, -0.8, -0.5, -0.2, 0, 1$).

For each configuration (K, γ), we performed a series of parameter sweeps to locate the discrete transition points (m_c, σ_c) separating the Single Equilibrium (SE), Multiple Equilibria (ME), and Unbounded Growth (UG) phases. Depending on the local slope of the boundary, we either fixed σ and scanned m (for horizontal-like boundaries) or fixed m and scanned σ (for vertical-like boundaries).

To obtain continuous phase boundaries from these discrete datasets, we performed a parametric interpolation using the Julia package `Dierckx.jl`. Specifically, we fitted a smooth parametric B-spline of order $k = 3$ to the sequence of detected critical coordinates $\{(m_c^{(i)}, \sigma_c^{(i)})\}_i$. This procedure ensures a robust reconstruction of the stability limits, smoothing out small numerical fluctuations while preserving the global topology of the phase diagram.

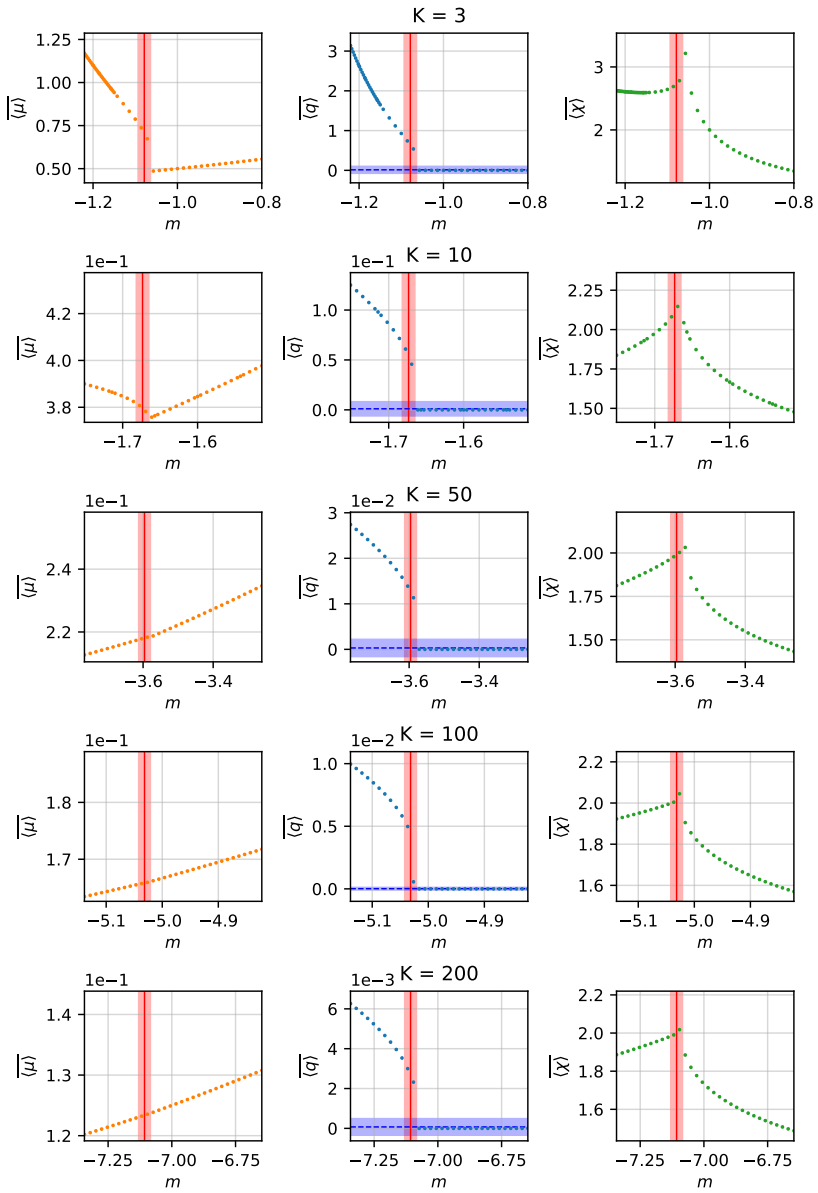

Fig. D.1 **Numerical determination of the topological phase transition** ($\sigma = 0$)**.** Each row corresponds to a different connectivity K. The columns display the population averages of the mean abundance $\overline{\langle \mu \rangle}$ (left), the cavity variance $\overline{\langle q \rangle}$ (center), and the susceptibility $\overline{\langle \chi \rangle}$ (right). The errors represent the standard deviation of the population distribution. In the central column (q), the blue horizontal line indicates the baseline noise floor (mean $\pm$ std), while the vertical red line marks the detected critical point m_c (with its estimated error width). Note the sharp discontinuity in q characteristic of the topological transition.

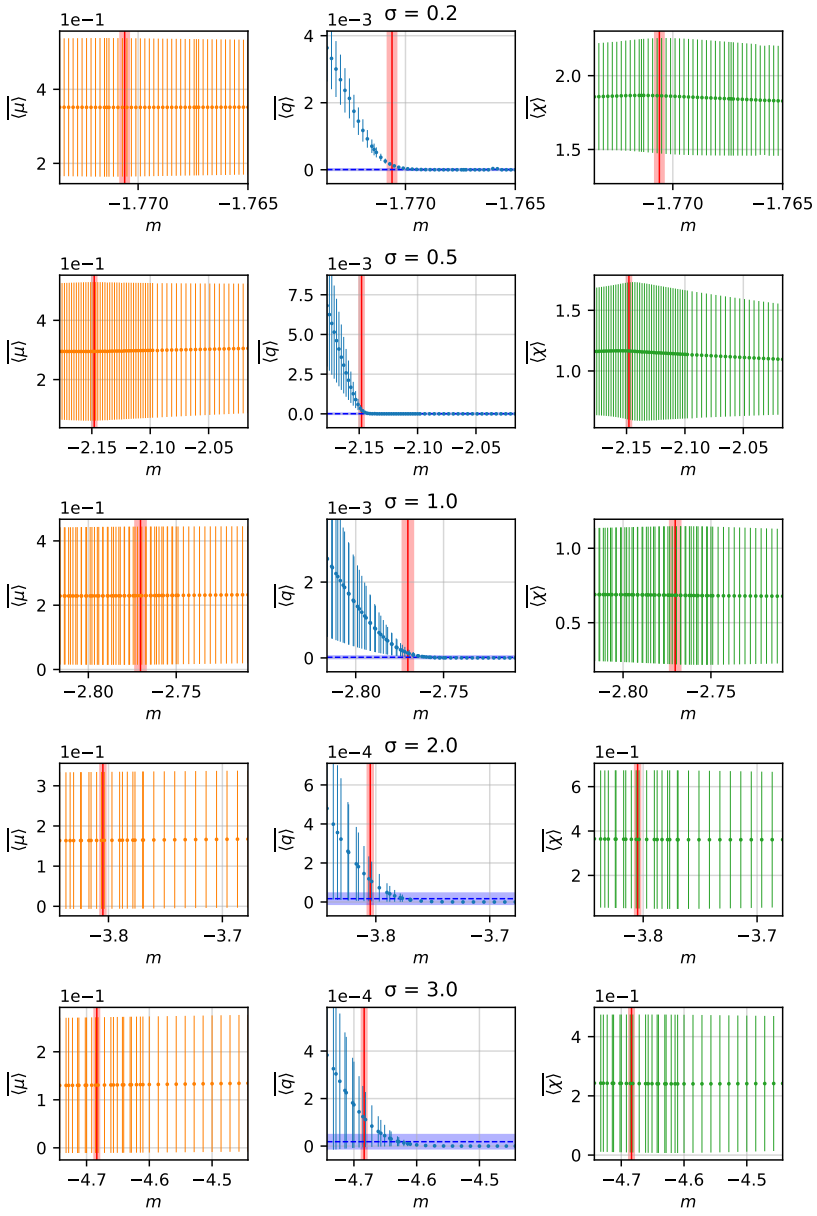

Fig. D.2 **Detection of the Single Equilibrium (SE) to Multiple Equilibria (ME) transition.** Parameters: $K = 10, \gamma = -1$. Each row corresponds to a different disorder strength σ. The columns display the population averages of the mean abundance $\overline{\langle \mu \rangle}$ (left), the cavity variance $\overline{\langle q \rangle}$ (center), and the susceptibility $\overline{\langle \chi \rangle}$ (right). The errors represent the standard deviation of the population distribution. In the central column (q), the blue horizontal line indicates the baseline noise floor (mean $\pm$ std), while the vertical red line marks the detected critical point m_c (with its estimated error width). The transition is identified by the departure of the cavity variance q from zero, while the mean abundance μ remains finite. Unlike the topological case, the order parameter grows continuously from zero, resulting in larger uncertainties for m_c as the distribution slowly broadens.

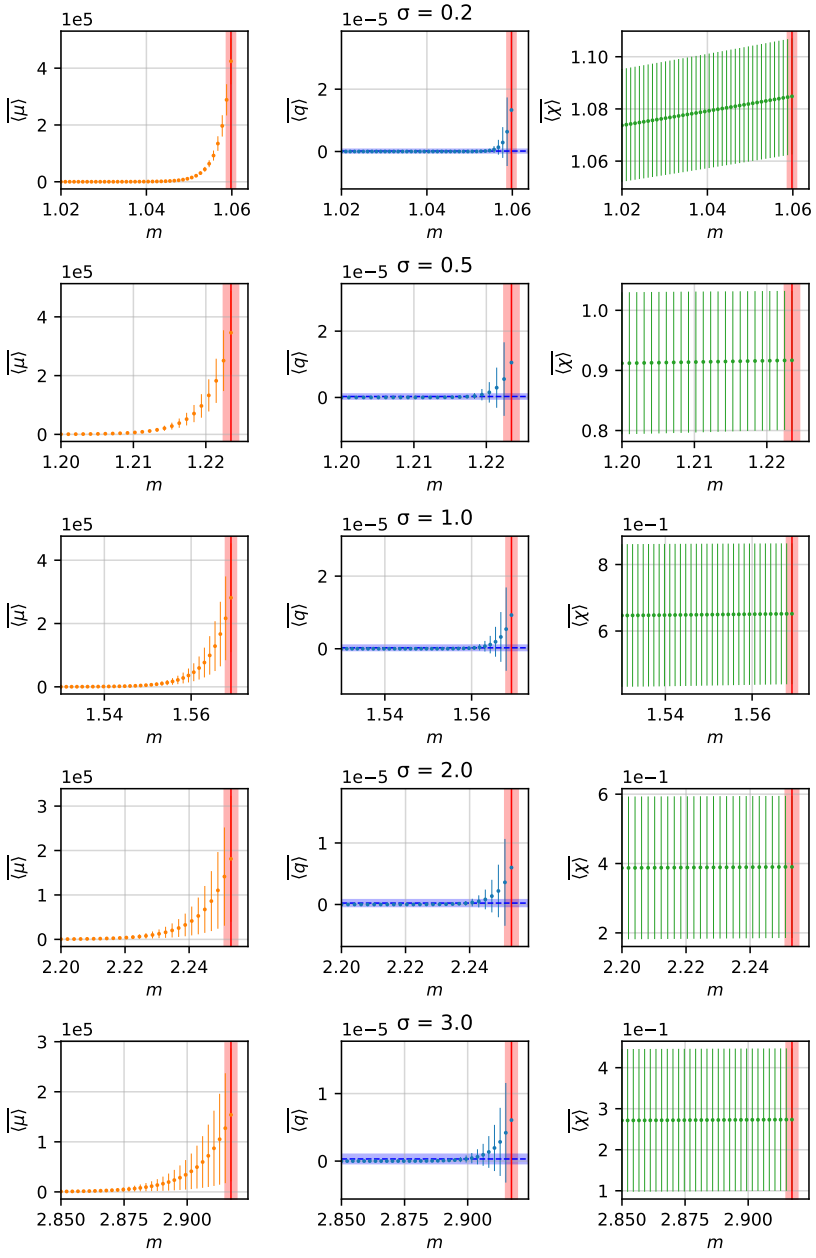

Fig. D.3 **Detection of the Single Equilibrium (SE) to Unbounded Growth (UG) transition.** Parameters: $K = 10, \gamma = -1$. Each row corresponds to a different disorder strength σ. The columns display the population averages of the mean abundance $\overline{\langle \mu \rangle}$ (left), the cavity variance $\overline{\langle q \rangle}$ (center), and the susceptibility $\overline{\langle \chi \rangle}$ (right). The errors represent the standard deviation of the population distribution. In the central column (q), the blue horizontal line indicates the baseline noise floor (mean $\pm$ std), while the vertical red line marks the detected critical point m_c (with its estimated error width). The critical point is associated with the divergence of the mean abundance μ. The vertical red line marks where the automated algorithm detects the instability, signaling the loss of linear stability and the onset of unbounded population growth.

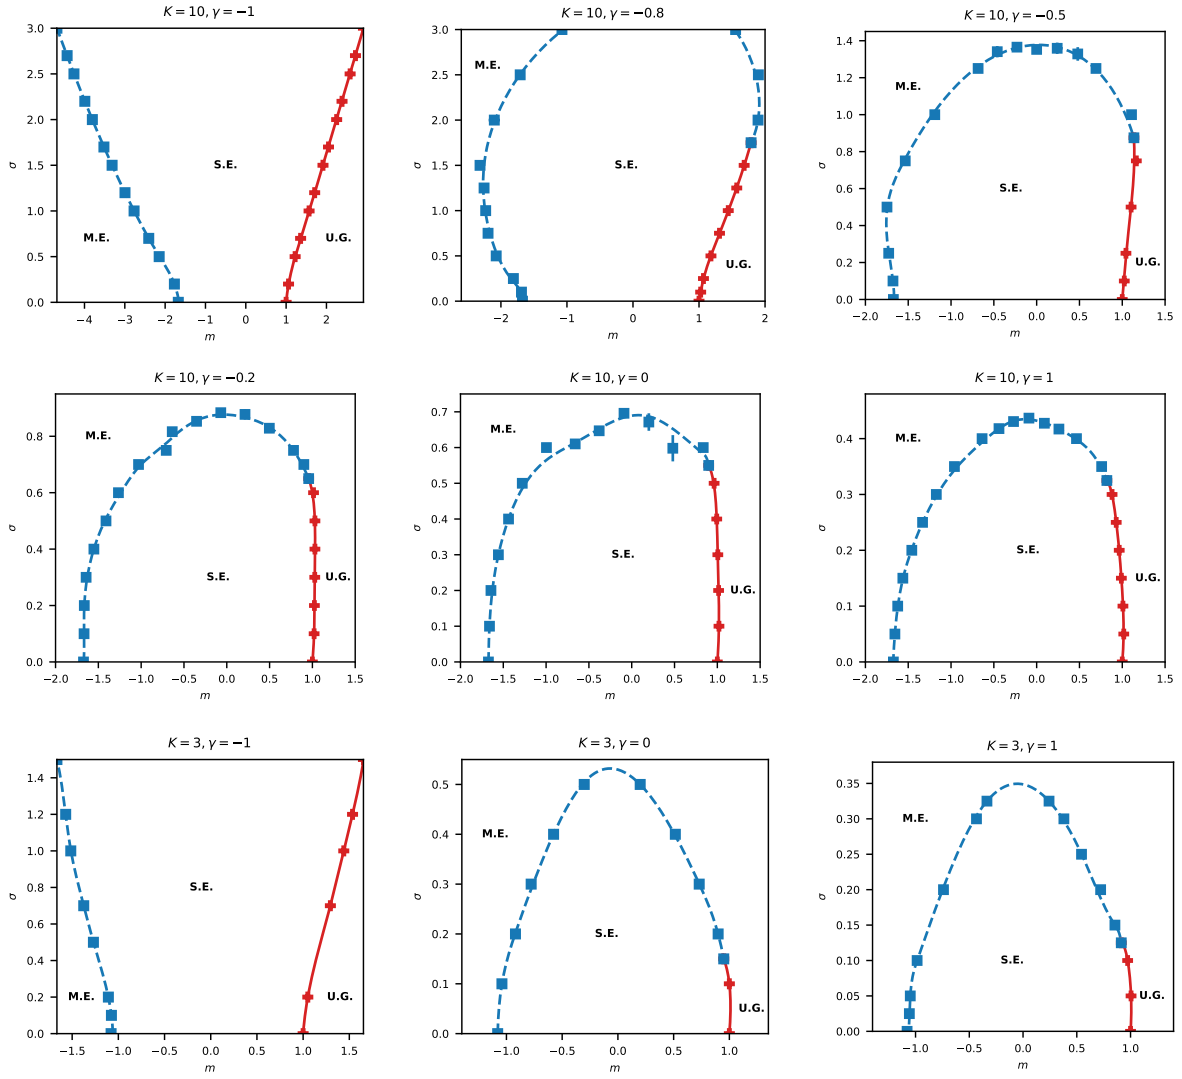

Fig. D.4 **Interpolated stability phase diagrams.** The grid displays the individual stability boundaries computed for different network configurations. **Top and Middle Rows ($K = 10$):** Phase diagrams for $\gamma = -1.0, -0.8, -0.5$ (top) and $\gamma = -0.2, 0.0, 1.0$ (middle). **Bottom Row ($K = 3$):** Phase diagrams for $\gamma = -1.0, 0.0, 1.0$. In each panel, the symbols represent the discrete critical points determined via the automated algorithm, while the lines represent the smooth parametric B-spline interpolation of order $k = 3$. The regions correspond to Single Equilibrium (SE), Multiple Equilibria (ME), and Unbounded Growth (UG) phases.

Appendix E

Detailed derivations for the Small-Coupling Dynamic Cavity Method

E.1 Path-integral formulation

This section presents an alternative derivation of the Small-Coupling Dynamic Cavity (SCDC) equations based on a path-integral representation of the epidemic dynamics. The approach closely follows the steps outlined in section 4.2 for continuous-state-space dynamics, adapted here to the discrete compartmental setting. We focus on the Susceptible-Infected (SI) model for clarity, but the derivation can be extended to other compartmental models with minor modifications.

E.1.1 Generating functional

As for the continuous case, we define the generating functional of the dynamics as

$$Z[\underline{\psi}] = \left\langle \exp\left(\sum_{i=1}^{N} \sum_{n=0}^{T} \sum_{x \in \mathscr{X}} \psi_i^n(x) \delta_{x_i^n, x} \right) \right\rangle \tag{E.1}$$

where the average is over all possible trajectories of the system, reweighted by their posterior probability Eq. (9.4b). Posterior averages of the epidemic trajectories can then be computed by differentiation of the generating functional, e.g.

$$\langle \delta_{x_i^n, x} \rangle = \left. \frac{\partial Z[\underline{\psi}]}{\partial \psi_i^n(x)} \right|_{\underline{\psi}=0}. \tag{E.2}$$

Explicitly, we have

$$Z[\underline{\boldsymbol{\psi}}] = \sum_{\underline{\boldsymbol{x}}} \prod_{i=1}^{N} p_i^0(x_i^0) p(O_i^T|x_i^T) e^{\sum_x \psi_i^T(x)\delta_{x_i^T,x}} \times \prod_{n=0}^{T-1} w_i^n(x_i^{n+1}|x_i^n, \underline{x}_{\partial i}^n) P(O_i^n|x_i^n) e^{\sum_x \psi_i^n(x)\delta_{x_i^n,x}}. \tag{E.3a}$$

$$= \sum_{\underline{\boldsymbol{x}}} \prod_{i=1}^{N} p_i^0(x_i^0) p(O_i^T|x_i^T) e^{\sum_x \psi_i^T(x)\delta_{x_i^T,x}} \times \prod_{n=0}^{T-1} \left[\sum_x \delta_{x_i^{n+1},x} w_i^n(x|x_i^n, \underline{x}_{\partial i}^n) \right] P(O_i^n|x_i^n) e^{\sum_x \psi_i^n(x)\delta_{x_i^n,x}}. \tag{E.3b}$$

Substituting the infection transition probabilities from Eqs. (8.7) and (8.8), we can rewrite the generating functional for the SI model as

$$Z[\underline{\boldsymbol{\psi}}] = \sum_{\underline{\boldsymbol{x}}} \prod_{i=1}^{N} p_i^0(x_i^0) \prod_{n=0}^{T} P(O_i^n|x_i^n) e^{\sum_x \psi_i^n(x)\delta_{x_i^n,x}} \prod_{n=0}^{T-1} \int dh_i^n \left[\delta_{x_i^{n+1},S} \delta_{x_i^n,S} \alpha_i^n e^{h_i^n} + \delta_{x_i^{n+1},I} \left(1 - \delta_{x_i^n,S} \alpha_i^n e^{h_i^n}\right) \right] \delta\left(h_i^n - \sum_{k\in\partial i} J_{ik} \delta_{x_k^n,I}\right), \tag{E.4}$$

where we introduced the infection pressure variable h_i^n defined in Eq. (8.5) and we used the shorthand notation $J_{ik} = \log(1-\lambda_{ik})$ for the infection couplings. The delta function enforcing the definition of h_i^n can be expressed in Fourier form as

$$\delta(h) = \int \frac{d\hat{h}}{2\pi} e^{-\mathrm{i}\hat{h}h}. \tag{E.5}$$

Substituting this expression into the generating functional, we obtain

$$Z[\underline{\boldsymbol{\psi}}] \propto \sum_{\underline{\boldsymbol{x}}} \prod_{i=1}^{N} p_i^0(x_i^0) \prod_{n=0}^{T} P(O_i^n|x_i^n) e^{\sum_x \psi_i^n(x)\delta_{x_i^n,x}} \times \prod_{n=0}^{T-1} \int d[h_i^n, \hat{h}_i^n] \left[(\delta_{x_i^{n+1},S} - \delta_{x_i^{n+1},I}) \delta_{x_i^n,S} \alpha_i^n e^{h_i^n} + \delta_{x_i^{n+1},I} \right] e^{-\mathrm{i}\hat{h}_i^n \left(h_i^n - \sum_{k\in\partial i} J_{ik} \delta_{x_k^n,I}\right)} \tag{E.6}$$

$$\propto \sum_{\underline{\boldsymbol{x}}} \int d\underline{\hat{\boldsymbol{h}}} \prod_{i=1}^{N} p_i^0(x_i^0) \prod_{n=0}^{T} P(O_i^n|x_i^n) e^{\sum_x \psi_i^n(x)\delta_{x_i^n,x}} \prod_{k\in\partial i} e^{\mathrm{i}\hat{h}_i^n J_{ik} \delta_{x_k^n,I}} \times \prod_{n=0}^{T-1} \left[(\delta_{x_i^{n+1},x_i^n} - \delta_{x_i^{n+1},I}) \alpha_i^n \delta(\hat{h}_i^n + \mathrm{i}) + \delta_{x_i^{n+1},I} \delta(\hat{h}_i^n) \right]. \tag{E.7}$$

We can extend the generating function to include source fields $\underline{\hat{\boldsymbol{\psi}}}$ coupled to the conjugate variables $\underline{\hat{\boldsymbol{h}}}$ as

$$Z[\underline{\boldsymbol{\psi}},\underline{\hat{\boldsymbol{\psi}}}] = \left\langle \exp\left(\sum_{i=1}^{N}\sum_{n=0}^{T}\sum_{x\in\mathscr{X}} \psi_i^n(x)\delta_{x_i^n,x} + \sum_{i=1}^{N}\sum_{n=0}^{T} \hat{\psi}_i^n \mathrm{i}\hat{h}_i^n \right) \right\rangle, \tag{E.8}$$

so that posterior averages of the conjugate variables can be computed as

$$\langle \mathrm{i}\hat{h}_i^n \rangle = \left. \frac{\partial Z[\underline{\boldsymbol{\psi}},\underline{\hat{\boldsymbol{\psi}}}]}{\partial \hat{\psi}_i^n} \right|_{\underline{\boldsymbol{\psi}}=0,\underline{\hat{\boldsymbol{\psi}}}=0}. \tag{E.9}$$

The generating functional with the additional source fields is factorized over nodes and edges of the contact network

$$Z[\underline{\boldsymbol{\psi}},\underline{\hat{\boldsymbol{\psi}}}] = \sum_{\underline{\boldsymbol{x}}} \int d\underline{\hat{\boldsymbol{h}}} \prod_{i=1}^{N} e^{S_i^{\mathrm{loc}}} \prod_{(i,j)\in E} e^{S_{ij}^{\mathrm{int}}} e^{\sum_{i=1}^{N}\sum_{n=0}^{T}\left(\sum_x \psi_i^n(x)\delta_{x_i^n,x} + \hat{\psi}_i^n \mathrm{i}\hat{h}_i^n\right)}, \tag{E.10}$$

where the local and pair actions are

$$S_i^{\mathrm{loc}} = \sum_{n=0}^{T-1} \log\left[\delta(\hat{h}_i^n + \mathrm{i})\alpha_i^n \left(\delta_{x_i^{n+1},x_i^n} - \delta_{x_i^{n+1},I} \right) + \delta(\hat{h}_i^n)\delta_{x_i^{n+1},I} \right], \tag{E.11}$$

$$S_{ik}^{\mathrm{int}} = \mathrm{i}\sum_{n=0}^{T-1} \left(\delta_{x_i^n,I} J_{ki}^n \hat{h}_k^n + \delta_{x_k^n,I} J_{ik}^n \hat{h}_i^n \right). \tag{E.12}$$

E.1.2 Cavity equations

We can apply the cavity method to the generating functional above, following the same steps outlined section 2.3. The main difference is that now, thanks to the introduction of the conjugate variables $\hat{h}_i^n$, there is no need to introduce the dual factor graph to decouple the interactions among neighboring nodes. The cavity equations for posterior marginals of the epidemic trajectories can be expressed as

$$c_{i\setminus j}(\boldsymbol{x}_i,\hat{\boldsymbol{h}}_i) = \frac{p_i^0(x_i^0)p(O_i^T|x_i^T)}{Z_{i\setminus j}} e^{S_i^{\mathrm{loc}}} \prod_{k\in\partial i\setminus j} \sum_{\boldsymbol{x}_k} \int d\hat{\boldsymbol{h}}_k c_{k\setminus i}(\boldsymbol{x}_k,\hat{\boldsymbol{h}}_k) e^{S_{ik}^{\mathrm{int}}}. \tag{E.13}$$

$$c_i(\boldsymbol{x}_i,\hat{\boldsymbol{h}}_i) = \frac{p_i^0(x_i^0)p(O_i^T|x_i^T)}{Z_i} e^{S_i^{\mathrm{loc}}} \prod_{k\in\partial i} \sum_{\boldsymbol{x}_k} \int d\hat{\boldsymbol{h}}_k c_{k\setminus i}(\boldsymbol{x}_k,\hat{\boldsymbol{h}}_k) e^{S_{ik}^{\mathrm{int}}}, \tag{E.14}$$

which are analogous to Eqs. (10.5) and (10.8).

E.2 Example of normalization issue for leaves of the contact graph

The small-coupling expansion requires assuming the normalization Eq. (10.11), which sums over all the possible trajectories of node i assuming $\boldsymbol{s}_i = 0$ at every time. The method thus considers all the trajectories in which the cavity node j is always susceptible, and therefore cannot infect node i. In particular, there are situations, such as the one considered in the example below, in which the normalization vanishes, meaning that it is not possible to explain an observed trajectory within the standard SI model. While this could seem pathological, it is worth stressing that the assumption made is necessary to obtain a message-passing algorithm that is independent of the trajectory of node j, a crucial condition to perform the expansion on which the present method is based. It is however possible to ensure that every trajectory of a node i remains feasible, the normalization constant being finite, by slightly modifying the epidemic model introducing a small self-infection probability. In addition to fixing the normalization issue, a small value of self-infection probability does not deteriorate the predictive power of the method.

To better illustrate this problem, we consider a leaf node i and its unique neighbor j. In the cavity graph corresponding to the message $c_{i\setminus j}(\boldsymbol{x}_i, \boldsymbol{s}_i)$, node i will appear as an isolated node. As a consequence, it is expected that the approximation behind the SCDC equations cannot explain, within the cavity graph, an infection actually transmitted from node j to node i. Indeed, because of the absence of further neighbors, the normalization term reads, after integration of the conjugate field $\boldsymbol{h}_i$,

$$Z_{i\setminus j} = \sum_{\boldsymbol{x}_i} p_0(x_i^0)\, p(\boldsymbol{O}_i|\boldsymbol{x}_i) \prod_{n=0}^{T-1} \left[\alpha_i^n (\delta_{x_i^{n+1}, x_i^n} - \delta_{x_i^{n+1}, I}) + \delta_{x_i^{n+1}, I} \right], \tag{E.15}$$

showing that an infection can only be explained by a self-infection event. When $\varepsilon_i^n = 0$ (i.e. $\alpha_i^n = 1$, no self-infections possible), the cavity message admits trajectories for which node i is always susceptible or infected. When a perfect observation (i.e. $\eta^+ = \eta^- = 0$) implies an infection event at some $n \neq 0$, the normalization vanishes, indicating an inconsistency in the model. This is prevented by the existence of a finite self-infection probability. Since it is recommended to operate in the limit of a vanishing self-infection, in the present case it is possible to analytically verify the limiting behavior for the cavity messages $\mu_{i\setminus j}^n$ and $\hat{\mu}_{i\setminus j}^n$.

As an example, we suppose that the leaf i is observed to be susceptible at time n_S and then is observed to be infected at time $n_I > n_S$. We consider a uniform self-infection probability $\varepsilon_i^n = \varepsilon$ ($\alpha_i^n = \alpha = 1 - \varepsilon$) for any time n and any node i, and a uniform prior probability

$p_0\left(x_i^0 = S\right) = 1-\gamma$, $p_0\left(x_i^0 = I\right) = \gamma$. The forward messages are

$$\rho_{\to n}^{i\backslash j}(S) = \begin{cases} (1-\gamma)\alpha^n & \text{if } n \leq n_I, \\ 0 & \text{if } n > n_I, \end{cases} \tag{E.16}$$

$$\rho_{\to n}^{i\backslash j}(I) = \begin{cases} \gamma + (1-\alpha)(1-\gamma)\sum_{l=0}^{n-1}\alpha^l & \text{if } n \leq n_S, \\ (1-\alpha)(1-\gamma)\sum_{l=n_S}^{n-1}\alpha^l & \text{if } n_S < n \leq n_I, \\ (1-\alpha)(1-\gamma)\sum_{l=n_S}^{n_I-1}\alpha^l & \text{if } n > n_I, \end{cases} \tag{E.17}$$

and the backward messages

$$\rho_{n\leftarrow}^{i\backslash j}(S) = \begin{cases} (1-\alpha)\alpha^{n_S-n}\sum_{l=0}^{n_I-1-n_S}\alpha^l & \text{if } n \leq n_S, \\ (1-\alpha)\sum_{l=0}^{n_I-1-n}\alpha^l & \text{if } n_S < n \leq n_I, \\ 1 & \text{if } n > n_I \end{cases} \tag{E.18}$$

$$\rho_{n\leftarrow}^{i\backslash j}(I) = \begin{cases} 0 & \text{if } n \leq n_S, \\ 1 & \text{if } n > n_S. \end{cases} \tag{E.19}$$

The normalization factor, taking into account the observations, is given by

$$\mathscr{Z}_{i\backslash j} = (1-\alpha)(1-\gamma)\sum_{l=n_S}^{n_I-1}\alpha^l, \tag{E.20}$$

so that the cavity marginal is given by

$$\mu_{i\backslash j}^n = \begin{cases} 0 & \text{if } n \leq n_S, \\ \frac{(1-\alpha)(1-\gamma)\sum_{l=n_S}^{n-1}\alpha^l}{(1-\alpha)(1-\gamma)\sum_{l=n_S}^{n_I-1}\alpha^l} & \text{if } n_S < n \leq n_I, \\ 1 & \text{if } n > n_I. \end{cases} \tag{E.21}$$

In the limit of vanishing self-infection $\varepsilon \to 0$ ($\alpha \to 1$), the cavity marginal takes the simple expression

$$\mu_{i\backslash j}^n = \begin{cases} 0 & \text{if } n \leq n_S, \\ \frac{n-n_S}{n_I-n_S} & \text{if } n_S < n \leq n_I, \\ 1 & \text{if } n > n_I, \end{cases} \tag{E.22}$$

which gives a reasonable probability profile for the node i to be infected in the absence of node i. It is worth stressing that this is not the full marginal m_i^n, which also depends on the messages coming from j to i. The cavity field instead diverges for times n between the two observation times

$$\hat{\mu}_{i\backslash j}^n = \begin{cases} 1 & \text{if } n \leq n_S, \\ -\infty & \text{if } n_S < n \leq n_I, \\ 0 & \text{if } n > n_I, \end{cases} \tag{E.23}$$

which is a clear consequence of having a vanishing normalization factor in the limit of zero self-infection. The divergence of the cavity field is thus the very non-physical effect of the inconsistency already discussed. In practice, in order to avoid divergences triggered by some peculiar combinations of observations, we then implement the algorithm using a cutoff $\hat{\mu}_c < 0$ on the values of $\hat{\mu}_{i\backslash j}$, such that the update rule Eq. (10.28) is implemented as follows

$$\hat{\mu}_{i\backslash j}^n = \max\left\{ \hat{\mu}_c, \frac{\rho_{\to n}^{i\backslash j}(S) M_{SS}^{i\backslash j} \left(\rho_{n+1\leftarrow}^{i\backslash j}(S) - \rho_{n+1\leftarrow}^{i\backslash j}(I) \right)}{\rho_{\to n}^{i\backslash j}(S)\, \rho_{n\leftarrow}^{i\backslash j}(S) + \rho_{\to n}^{i\backslash j}(I)\, \rho_{n\leftarrow}^{i\backslash j}(I)} \right\}. \tag{E.24}$$

E.3 Time-forward dynamics

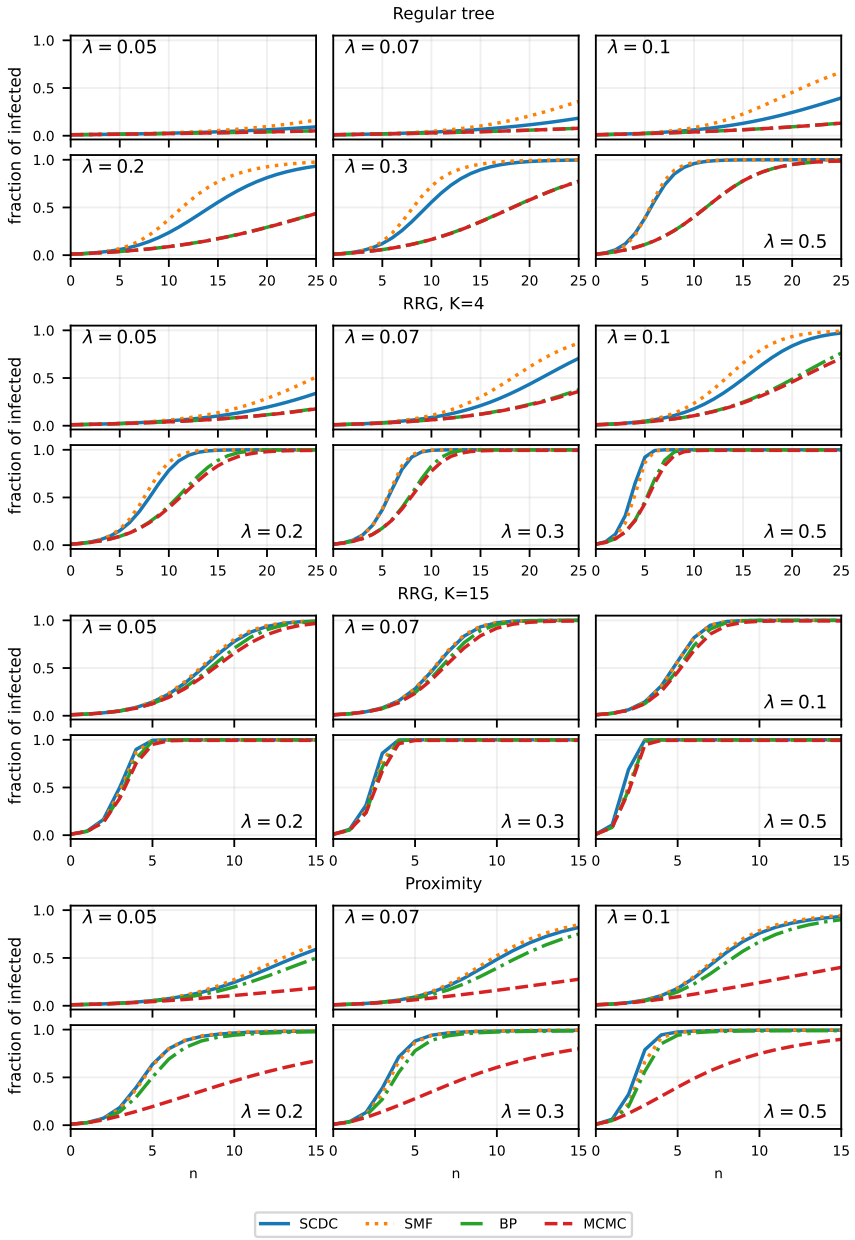

Fig. E.1 **Time-forward dynamics for varying infection probability.** Forward dynamics with varying λ. Fraction of infected individuals against epidemic time, with four different static contact graphs. From top to bottom: regular tree with degree $K = 4$ and $N = 485$, RRG with $N = 500$ and degree $K = 4$, RRG with $N = 500$ and $K = 15$, proximity graph with $N = 500$. Comparison is shown between SCDC, Individual-Based Mean-Field (IBMF), Belief Propagation (BP) and Monte Carlo simulations (with 10^4 samples). All the links have the same infection probability λ, whose value is reported inside each panel. Six different values of λ are used. In all cases, the probability of each individual being infected at time $n = 0$ is set to $p_0(x_i^0) = 5/N$, and the self-infection ε_i^n is set to 0 for every node i.