

Building Intuition for Dynamical Mean-Field Theory: A Simple Model and the Cavity Method

A tutorial oriented to the biophysics community

Emmy Blumenthal

Princeton University Department of Physics, Princeton, NJ 08540, USA

July 22, 2025

Dynamical Mean-Field Theory (DMFT) is a powerful theoretical framework for analyzing systems with many interacting degrees of freedom. This tutorial provides an accessible introduction to DMFT. We begin with a linear model where the DMFT equations can be derived exactly, allowing readers to develop clear intuition for the underlying principles. We then introduce the cavity method, a versatile approach for deriving DMFT equations for non-linear systems. The tutorial concludes with an application to the generalized Lotka–Volterra model of interacting species, demonstrating how DMFT reduces the complex dynamics of many-species communities to a tractable single-species stochastic process. Key insights include understanding how quenched disorder enables the reduction from many-body to effective single-particle dynamics, recognizing the role of self-averaging in simplifying complex systems, and seeing how collective interactions give rise to non-Markovian feedback effects.

Contents

1	Motivation	3
2	Exactly deriving the DMFT of a linear model	4
2.1	The model	4
2.2	Choosing a random particle: two types of randomness	4
2.3	Integrating out $N - 1$ degrees of freedom	5
2.4	Statistical analysis and self-averaging	7
2.5	Putting it together: the DMFT equation of motion	10
3	The cavity method	10
3.1	Step 1: add the cavity degree of freedom	11
3.2	Step 2: treat the cavity degree of freedom as a perturbation	11
3.3	Step 3: obtain the cavity equation of motion	13
3.4	Step 4: statistics under quenched disorder	13
3.5	Step 5: assert self-averaging to get self-consistent equations	14
3.6	Step 6: solve the self-consistent equations (hard!)	15
3.7	Why was it necessary to add the cavity degree of freedom?	17
4	A taste of possible analyses	17
4.1	Correlations	17
4.2	Response to perturbations	18
4.3	Fourier modes and power spectrum	19
5	A non-linear example: the generalized Lotka–Volterra model	19

5.1	The model	20
5.2	The cavity method	20
6	Conclusion	24
A	Appendix: Brief introduction to random matrix theory (RMT)	26
A.1	Observables	26
A.2	Wigner ensemble and the scaling of entries	27
A.3	Ginibre ensemble	27
A.4	Ginibre matrices as random rotations	27
A.5	Wick's theorem	28
A.6	The Green's function as the moment generating function	28
A.7	Derivation of the Wigner semicircle law	29
A.8	The skew-symmetric case	30
B	Appendix: Brief review of stationary Gaussian processes	30
B.1	Definition	31
B.2	Stationary Gaussian processes	31
B.3	Power spectral density and Wiener–Khinchin theorem .	32

A brief note of acknowledgment. Much of my understanding of this content comes from research studying ecological models. I learned many of these ideas from references [1–8]. Of course, all errors are my own. I am grateful to Akshit Goyal, Gautam Reddy, and Pankaj Mehta for their feedback on these notes. Please contact me if you identify any errors or have suggestions for improvement.¹

¹ eblu [ατ] princeton.edu

1 Motivation

Dynamical mean-field theory (DMFT) is a powerful theoretical framework to reduce the equations of motion for a dynamical system with many degrees of freedom to a single equation of motion for a single degree of freedom, sometimes called the ‘mean-field’ particle. DMFT was originally developed to study the dynamics of spin glasses but has since been applied to a wide variety of systems that involve dense and disordered interactions between many heterogeneous degrees of freedom. For example, in ecology, we might be interested in dynamics of the abundances of species in a community where each species interacts with all other species. A model for this system might look like:

$$\frac{dN_i}{dt} = N_i \left(1 - N_i - \sum_{j(\neq i)} A_{ij} N_j \right), \quad i = 1, \dots, S \quad (1)$$

where N_i is the abundance of species i and A_{ij} is a matrix of interaction parameters between species i and j . This model is called the generalized Lotka–Volterra (GLV) model of population dynamics. Making definitive statements about the dynamics of this entire system in the $S \gg 1$ limit is difficult because it is non-linear and involves $O(S^2)$ interaction parameters. To avoid this problem, in the $S \gg 1$ limit (the thermodynamic limit), we might consider parameters A_{ij} to be drawn from some distribution $P(\{A_{ij}\})$ and notice that species i is interacting with a sum over many other species (a ‘mean field’). We might replace such a sum with a stochastic process and get an equation of motion for just species i of the form:

$$\frac{dN_i}{dt} = \text{some expression only in terms of } N_i. \quad (2)$$

Such an equation of motion is called the DMFT equation of motion. Deriving this sort of equation is advantageous because it allows us to apply tools from the theory of stochastic process to study the dynamics of the (deterministic) system, potentially facilitating more tractable theoretical analysis. At the end of this note in section 5, we will derive the right-hand side of Eq. (2).

Another example application of DMFT is the study of neural networks which are some of the first biophysical subjects of study to which DMFT was applied [9, 10]. A model for neural dynamics might describe the activations $\{S_i\}$ of N neurons which each experience a local field h_i so that $S_i = \phi(h_i)$ with ϕ being a non-linear gain function (for example $\phi(x) = \tanh(gx)$). The local fields undergo dynamics:

$$\frac{dh_i}{dt} = -h_i + \sum_{j=1}^N J_{ij} S_j(t), \quad i = 1, \dots, N, \quad (3)$$

which models the local field h_i as a sum of contributions from all other neurons S_j weighted by some synaptic weights J_{ij} . A DMFT approach to this system would be to treat J_{ij} as a random matrix drawn from some distribution $P(\{J_{ij}\})$ and then derive an equation of motion for a single, randomly chosen neuron. This allows for analysis of many behaviors without having to specify $O(N^2)$ synaptic weights.

In appendix A, we will briefly review random matrix theory (RMT) which will serve as a useful tool for deriving the DMFT equations. In appendix B, we will briefly review stationary Gaussian processes which will appear in the DMFT equations.

2 Exactly deriving the DMFT of a linear model

In this section, we will exactly derive the DMFT equations for a linear model. The techniques we will apply here are typically not possible in any non-linear model because we can exactly solve the linear model. However, this example is useful because it allows us to precisely identify what assumptions we make when deriving a DMFT. It also clarifies why certain steps we take in a method like the cavity method are justified. It's important that a reader who has not been exposed to DMFT before knows that **the approach we take in this section is almost never possible for any setup except linear models**. An approach that works more generally is the cavity method, which we will discuss in section 3. Nonetheless, I believe this example gets to the heart of what DMFT is and why it is useful.

Figure 1 shows a schematic of the steps we will take to derive the DMFT equations for this linear model.

2.1 The model

We will consider a linear model of the form:

$$\frac{dx_i}{dt} = \sum_{j=1}^N A_{ij} x_j(t), \quad i = 1, \dots, N, \quad (4)$$

where A_{ij} is an $N \times N$ skew-symmetric ($A_{ij} = -A_{ji}$) matrix. We will assume that the entries of the matrix A_{ij} are drawn from a Gaussian distribution with statistics:²

$$\langle A_{ij} \rangle = 0, \quad \langle A_{ij} A_{kl} \rangle = \frac{\sigma^2}{N} (\delta_{ik} \delta_{jl} - \delta_{il} \delta_{jk}) \quad (5)$$

where σ^2 is a constant that sets the scale of the interactions and δ_{ij} is the Kronecker delta symbol (1 if $i = j$, 0 otherwise). This says that A_{ij} are independent up to the constraint that $A_{ij} = -A_{ji}$, have zero mean, and have variance σ^2/N . We scale the variance of A_{ij} by $1/N$ so that the time derivative of x_i is of order 1 in the thermodynamic limit $N \gg 1$. We will similarly assert that the initial conditions $x_1, \dots, x_N$ are drawn independently from a Gaussian distribution with zero mean and variance σ_0^2 :

$$\langle x_i(0) \rangle = 0, \quad \langle x_i(0) x_j(0) \rangle = \sigma_0^2 \delta_{ij}. \quad (6)$$

2.2 Choosing a random particle: two types of randomness

The goal of DMFT is to reduce the equations of motion for the N particles to a single equation of motion for a single particle. To do this, we choose a random particle which we will label $i = \alpha$ and then eliminate all other particles. In our simple example, we will eliminate all other

² It is not strictly necessary to assume that A_{ij} follow a normal distribution, but it simplifies the analysis.

particles by solving their equations of motion in terms of the random particle α and then substituting this solution into the equation of motion for α . Before we do this, we need to explicitly identify that there are two types of randomness in our model. There is the randomness due to sampling the entries of the matrix A_{ij} and initial conditions $x_i(0)$ which is called *quenched disorder*. There is also the randomness due to the choice of the random particle α , which I will call *shuffling randomness*. To be explicit, an average of a quantity X over quenched disorder is defined:

$$\langle X \rangle_{\mathcal{Q}} = \int \prod_{i,j=1}^N dA_{ij} \prod_{i=1}^N dx_i(0) P(\{A_{ij}\}, \{x_i(0)\}) X, \quad (7)$$

and an average over shuffling randomness is defined:

$$\langle X \rangle_{\text{sh}} = \frac{1}{N} \sum_{\alpha=1}^N (X \text{ with } \alpha \text{ being the chosen random particle}). \quad (8)$$

The goal of DMFT is to derive a stochastic differential equation which has some noise term in which sampling the noise is equivalent to shuffling randomness.³ That is to say: the dynamics of every particle are equivalent to the dynamics of a single particle with some noise term which is ‘resampled’ for each particle. Such an interpretation is useful because it allows us to use tools from the theory of stochastic processes even for a single, frozen realization of the quenched disorder (the matrix A_{ij} and initial conditions $x_i(0)$).

2.3 Integrating out $N - 1$ degrees of freedom

We will now proceed to derive the DMFT equations for the random particle α . We do this by solving the equations of motion for the other $N - 1$ particles in terms of the random particle α and then substituting this solution into the equation of motion for α . We split the equation of motion for the random particle α into two parts:

$$\begin{aligned} \frac{dx_i}{dt} &= \sum_{j(\neq\alpha)} A_{ij} x_j(t) + A_{i\alpha} x_{\alpha}(t), \quad (i \neq \alpha), \\ \frac{dx_{\alpha}}{dt} &= \sum_{i(\neq\alpha)} A_{\alpha i} x_i(t) + A_{\alpha\alpha} x_{\alpha}(t). \end{aligned} \quad (9)$$

We can exactly solve the equations of motion for the $N - 1$ particles $i \neq \alpha$ in terms of the random particle α :⁴

$$x_i(t) = \sum_{j(\neq\alpha)} [e^{A_{\setminus\alpha} t}]_{ij} x_j(0) + \int_0^t dt' \sum_{j(\neq\alpha)} [e^{A_{\setminus\alpha}(t-t')}]_{ij} A_{j\alpha} x_{\alpha}(t') \quad (10)$$

where we have introduced the notation $A_{\setminus\alpha}$ to denote the matrix A with the row and column corresponding to the random particle α removed. Here, $e^{A_{\setminus\alpha} t}$ is the matrix exponential. We then substitute this solution into the equation of motion for the random particle α :

$$\frac{dx_{\alpha}}{dt} = \Phi_{\alpha}(t) - \sigma^2 \int_0^t dt' \nu_{\alpha}(t-t') x_{\alpha}(t'), \quad (11)$$

³ The term ‘shuffling randomness’ is not standard, but I find it useful to distinguish the two types of randomness. It’s more standard to call the sort of average in Eq. (8) a single-realization system average.

⁴ This is accomplished by first writing the equations of motion for the $N - 1$ particles in matrix form:

$$\frac{d\mathbf{x}_{\setminus\alpha}}{dt} = A_{\setminus\alpha} \mathbf{x}_{\setminus\alpha}(t) + A_{\setminus\alpha, \alpha} x_{\alpha}(t),$$

where $\mathbf{x}_{\setminus\alpha}$ is the vector with the α th component removed. We then note that $\mathbf{x}_{\setminus\alpha}(t) = e^{A_{\setminus\alpha} t} \mathbf{x}_{\setminus\alpha}(0)$ solves the equation of motion if we neglect the term involving $x_{\alpha}(t)$. Finally, we use that the equation is linear to get Eq. (10).




Figure 1 – Schematic of DMFT derivation presented in section 2. (a) Fully connected linear network with antisymmetric interactions ($A_{ij} = -A_{ji}$). Blue and red arrows represent positive and negative interactions, respectively. A randomly selected degree of freedom (star) is chosen and treated separately from the rest of the network. (b) Solve the dynamics for x_i without x_α . There are two linear contributions: the autonomous (self-driven) dynamics of the network without x_α and the dynamics driven by the dynamics of x_α . (c) The dynamics of x_α are driven by the dynamics of all other x_i . (d) After statistically analyzing the resulting equation of motion for x_α , one finds two linear contributions. One is due to the autonomous dynamics of all other x_i and converges to a Gaussian process in the large- N limit. The other is due to a back-reaction from x_α driving the rest of the system and the bulk driving x_α back; it becomes a memory term in the large- N limit with a deterministic kernel.

where we have defined:

$$\begin{aligned}\Phi_\alpha(t) &\equiv \sum_{i,j(\neq\alpha)} A_{\alpha i} [e^{A_{\setminus\alpha} t}]_{ij} x_j(0), \\ \nu_\alpha(t) &\equiv -\frac{1}{\sigma^2} \sum_{i,j(\neq\alpha)} A_{\alpha i} [e^{A_{\setminus\alpha} t}]_{ij} A_{j\alpha}.\end{aligned}\tag{12}$$

Note that we have dropped $A_{\alpha\alpha}$ because it is zero from the assumption that A_{ij} is skew-symmetric.⁵ This constitutes the first step in deriving the DMFT equation of motion. This result, however, is not yet in a form that is useful for analysis because we don't know much about $\Phi_\alpha(t)$ and $\nu_\alpha(t)$.

⁵ Generally, you can drop these sorts of terms in DMFT because they are of order $1/\sqrt{N}$ and thus negligible in the thermodynamic limit.

2.4 Statistical analysis and self-averaging

2.4.1 Φ_α fluctuations under quenched disorder at fixed α

To get a more useful result, we look to the statistics of the random variables $\Phi_\alpha(t)$ and $\nu_\alpha(t)$. Recall that $\Phi_\alpha(t)$ and $\nu_\alpha(t)$ may fluctuate over both quenched disorder and shuffling randomness. We will first treat the quenched disorder and then see how it interacts with the shuffling randomness. This means that we will first consider the case where α is fixed. Under this assumption, we might notice that $\Phi_\alpha(t)$ is a sum of $\{A_{\alpha i}\}$ which are independent Gaussian random variables which means that under fluctuations over quenched disorder, $\Phi_\alpha(t)$ is also a Gaussian process.⁶ A Gaussian process is completely characterized by its first two moments, the mean and covariance:

$$\begin{aligned}\langle \Phi_\alpha(t) \rangle_Q &= \sum_{i,j(\neq\alpha)} \langle A_{\alpha i} \rangle \langle [e^{A_{\setminus\alpha} t}]_{ij} \rangle \langle x_j(0) \rangle_Q = 0, \\ \langle \Phi_\alpha(t) \Phi_\alpha(s) \rangle_Q &= \sum_{i,j,k,l(\neq\alpha)} \langle A_{\alpha i} A_{\alpha k} \rangle \langle [e^{A_{\setminus\alpha} t}]_{ij} [e^{A_{\setminus\alpha} s}]_{kl} \rangle \langle x_j(0) x_l(0) \rangle \\ &= \sigma^2 \sigma_0^2 \frac{1}{N} \sum_{i,j(\neq\alpha)} \langle [e^{A_{\setminus\alpha} t}]_{ij} [e^{A_{\setminus\alpha} s}]_{ij} \rangle \\ &\approx \sigma^2 \sigma_0^2 \left\langle \frac{1}{N} \text{Tr} \left[e^{A_{\setminus\alpha} t + A_{\setminus\alpha}^T s} \right] \right\rangle \\ &= \sigma^2 \sigma_0^2 \left\langle \frac{1}{N} \sum_{i(\neq\alpha)} e^{\lambda_i t + \lambda_i^* s} \right\rangle.\end{aligned}\tag{13}$$

⁶ If you are uncomfortable with this argument, I've included a brief review of stationary Gaussian processes in appendix B.

Notice that in Eq. (14), we were able to factor the averages because the entries of the matrix $A_{\setminus\alpha}$ are (definitionally) independent of $A_{\alpha i}$. Here, λ_i are the eigenvalues of the matrix $A_{\setminus\alpha}$. In the $N \gg 1$ limit, we can replace the average over the eigenvalues with an integral over the spectral density of the matrix A .⁷

$$\frac{1}{N} \sum_{i(\neq\alpha)} e^{\lambda_i t + \lambda_i^* s} \rightarrow \int d^2\lambda \rho_A(\lambda) e^{\lambda t + \lambda^* s}\tag{15}$$

where:

$$\rho_A(\lambda) = \frac{1}{N} \sum_{i=1}^N \delta(x - \lambda_i)\tag{16}$$

⁷ We can use the spectral density of A instead of $A_{\setminus\alpha}$ because they are the same up to a small correction of order $1/N$ in the $N \gg 1$ limit.

is the empirical distribution of the eigenvalues of the matrix A , called the spectral density. In the thermodynamic limit this becomes:

$$\rho_A(x + iy) = \begin{cases} \delta(x) \cdot \frac{1}{2\pi\sigma^2} \sqrt{4\sigma^2 - y^2}, & \text{if } -2\sigma \leq y \leq 2\sigma, \\ 0, & \text{otherwise,} \end{cases} \quad (17)$$

which is the Wigner semicircle law but rotated onto the imaginary axis because A_{ij} is skew-symmetric and thus has purely imaginary eigenvalues.⁸ We can exploit these results to write:

$$\begin{aligned} \langle \Phi_\alpha(t) \Phi_\alpha(s) \rangle_{\mathcal{Q}} &= \frac{\sigma_0^2}{2\pi} \int_{-2\sigma}^{2\sigma} dy \sqrt{4\sigma^2 - y^2} e^{iy(t-s)} \\ &= \sigma_0^2 \sigma^2 \frac{J_1(2\sigma|t-s|)}{\sigma|t-s|}, \end{aligned} \quad (18)$$

where J_1 is the Bessel function of the first kind.⁹

2.4.2 Fluctuations of Φ_α under shuffling randomness

Now, we might want to consider how $\Phi_\alpha(t)$ fluctuates under shuffling randomness. To do this, „we might try to compute $\langle \Phi_\alpha(t) \rangle_{\text{sh}} = N^{-1} \sum_{\alpha=1}^N \Phi_\alpha(t)$ and $\langle \Phi_\alpha(t) \Phi_\alpha(s) \rangle_{\text{sh}} = N^{-1} \sum_{\alpha=1}^N \Phi_\alpha(t) \Phi_\alpha(s)$. However, there's not really much progress we can make because whatever expression we get will still involve a sum over α which is very difficult to compute. Instead, we will investigate what happens to these shuffled averages under fluctuations over quenched disorder. We are interested in the quenched averages of the shuffled averages:

$$\begin{aligned} \langle \langle \Phi_\alpha(t) \rangle_{\text{sh}} \rangle_{\mathcal{Q}} &= \frac{1}{N} \sum_{\alpha=1}^N \langle \Phi_\alpha(t) \rangle_{\mathcal{Q}} = \langle \Phi_\alpha(t) \rangle_{\mathcal{Q}} \\ \langle \langle \Phi_\alpha(t) \Phi_\alpha(s) \rangle_{\text{sh}} \rangle_{\mathcal{Q}} &= \frac{1}{N} \sum_{\alpha=1}^N \langle \Phi_\alpha(t) \Phi_\alpha(s) \rangle_{\mathcal{Q}} = \langle \Phi_\alpha(t) \Phi_\alpha(s) \rangle_{\mathcal{Q}} \end{aligned} \quad (19)$$

Because the shuffled average is linear in $\Phi_\alpha(t)$ and $\langle \Phi_\alpha(t) \rangle_{\mathcal{Q}}$ is independent of α , the quenched average of the shuffled average is just the quenched average of $\Phi_\alpha(t)$ at fixed α .¹⁰ The more interesting question is whether the shuffled averages fluctuate under quenched disorder. To

⁸ If you are unfamiliar with these concepts, I have included a brief review of random matrix theory in appendix A.

⁹ The Bessel function of the first kind is defined as:

$$J_1(x) = \frac{1}{\pi} \int_0^\pi \cos(\theta - x \sin \theta) d\theta.$$

It has the $x \rightarrow \infty$ asymptotic behavior:

$$J_1(x) \sim \sqrt{\frac{2}{\pi x}} \sin(x - \frac{\pi}{4}).$$

¹⁰ It is not a given that $\langle \Phi_\alpha(t) \rangle_{\mathcal{Q}}$ is independent of α in general. In our case, it is, and you can check this but substituting the explicit expression for $\Phi_\alpha(t)$ into the average and then computing the average, just as we did in Eqs. (13), (14). In some models, there are anomalous fluctuations in which case one would find $\text{var}_{\mathcal{Q}}[\langle \Phi_\alpha(t) \rangle_{\text{sh}}] = O(1)$ and proceed to compute quantities like $\langle \langle \Phi_\alpha(t) \rangle_{\text{sh}} \langle \Phi_\beta(t) \rangle_{\text{sh}} \rangle_{\mathcal{Q}}$ to understand the nature of these fluctuations. See Ref. [1] to see an example of this in the context of spin glasses.

answer this, we can compute the variance of the shuffled averages:

$$\begin{aligned}
\text{var}_{\mathcal{Q}} [\langle \Phi_{\alpha}(t) \rangle_{\text{sh}}] &= \frac{1}{N^2} \sum_{\alpha, \beta=1}^N \text{cov}_{\mathcal{Q}} [\Phi_{\alpha}(t), \Phi_{\beta}(t)] \\
&= \frac{1}{N^2} \sum_{\alpha, \beta=1}^N \text{cov}_{\mathcal{Q}} [\Phi_{\alpha}(t), \Phi_{\beta}(t)] \\
&= \frac{1}{N^2} \sum_{\alpha, \beta=1}^N \sum_{i, j (\neq \alpha)} \sum_{k, l (\neq \beta)} \left(\langle A_{\alpha i} A_{\beta k} \rangle \langle [e^{A_{\alpha} t}]_{ij} [e^{A_{\beta} t}]_{kl} \rangle \langle x_j(0) x_l(0) \rangle \right. \\
&\quad \left. - \langle A_{\alpha i} [e^{A_{\alpha} t}]_{ij} x_j(0) \rangle \langle A_{\beta k} [e^{A_{\beta} t}]_{kl} x_l(0) \rangle \right) \\
&= \frac{\sigma_0^2 \sigma^2}{N^3} \sum_{\alpha, \beta=1}^N \sum_{i, j (\neq \alpha)} \sum_{k, l (\neq \beta)} (\delta_{\alpha \beta} \delta_{ik} - \delta_{\alpha k} \delta_{i\beta}) \langle [e^{A_{\alpha} t}]_{ij} [e^{A_{\beta} t}]_{kl} \rangle \delta_{jl} \\
&= \frac{\sigma_0^2 \sigma^2}{N^3} \left(\sum_{\alpha} \sum_{i, j (\neq \alpha)} \langle [e^{A_{\alpha} t}]_{ij} [e^{A_{\alpha} t}]_{ij} \rangle \right. \\
&\quad \left. - \sum_{\alpha, \beta (\alpha \neq \beta)} \sum_{i (\neq \alpha, \beta)} \langle [e^{A_{\alpha} t}]_{\beta i} [e^{A_{\beta} t}]_{\alpha i} \rangle \right) \\
&= \sigma_0^2 \sigma^2 \frac{1}{N} \left(2 \cdot \frac{1}{N} \sum_{\alpha} \overbrace{\frac{1}{N} \text{Tr} \left[e^{A_{\alpha} t} (e^{A_{\alpha} t})^T \right]}^{O(1)} \right) \\
&= O\left(\frac{1}{N}\right).
\end{aligned} \tag{20}$$

In the final step of the calculation, we used that a trace of a random matrix like $\frac{1}{N} \text{Tr} M$ is $O(1)$ in the thermodynamic limit $N \gg 1$ (this is because it is an average of the $O(1)$ eigenvalues of the matrix M).¹¹ Through this computation, we've found that as $N \rightarrow \infty$, the shuffled average of $\Phi_{\alpha}(t)$ does not fluctuate under quenched disorder. This means that we can replace the quenched average of the shuffled average with the quenched average of $\Phi_{\alpha}(t)$ at fixed α . In fact, this is true for all shuffled averages involving $\Phi_{\alpha}(t)$:

$$\text{var}_{\mathcal{Q}} [\langle \Phi_{\alpha}(t_1) \Phi_{\alpha}(t_2) \cdots \Phi_{\alpha}(t_n) \rangle_{\text{sh}}] = O\left(\frac{1}{N}\right). \tag{21}$$

This means that we can always replace shuffled averages of $\Phi_{\alpha}(t)$ with quenched averages of $\Phi_{\alpha}(t)$ at fixed α :¹²

$$\langle \Phi_{\alpha}(t_1) \Phi_{\alpha}(t_2) \cdots \Phi_{\alpha}(t_n) \rangle_{\text{sh}} \approx \langle \Phi_{\alpha}(t_1) \Phi_{\alpha}(t_2) \cdots \Phi_{\alpha}(t_n) \rangle_{\mathcal{Q}}, \tag{22}$$

for $N \gg 1$.

Quantities which have this property are called *self-averaging*. For our system, this means that under frozen quenched disorder, $\Phi_{\alpha}(t)$ behaves like a Gaussian process where the origin of the randomness is the shuffling randomness, not the quenched disorder (which is the origin from which we first said the fluctuations came).

¹¹ This is discussed in more detail in appendix A.

¹² One shows this by showing that the generating functional of $\Phi_{\alpha}(t)$ has $O(1/N)$ variance under quenched disorder:

$$\text{var}_{\mathcal{Q}} \left[\left\langle e^{\int dt \Phi_{\alpha}(t) j(t)} \right\rangle_{\text{sh}} \right] = O\left(\frac{1}{N}\right)$$

2.4.3 Fluctuations of ν_α

We can repeat the same sort of analysis for $\nu_\alpha(t)$. We can compute the quenched average of $\nu_\alpha(t)$ at fixed α :

$$\begin{aligned}\langle \nu_\alpha(t) \rangle_{\mathcal{Q}} &= -\frac{1}{\sigma^2} \sum_{i,j(\neq\alpha)} \langle A_{\alpha i} A_{j\alpha} \rangle \langle [e^{A_{\setminus\alpha} t}]_{ii} \rangle = \left\langle \frac{1}{N} \sum_{i(\neq\alpha)} [e^{A_{\setminus\alpha} t}]_{ij} \right\rangle \\ &= \int d^2\lambda \rho_A(\lambda) e^{\lambda t} = \frac{1}{2\pi\sigma^2} \int_{-2\sigma}^{2\sigma} dy \sqrt{4\sigma^2 - y^2} e^{iyt} \\ &= \frac{J_1(2\sigma t)}{\sigma t}.\end{aligned}\tag{23}$$

Next, compute its fluctuations under quenched disorder:

$$\begin{aligned}\text{var}_{\mathcal{Q}}[\nu_\alpha(t)] &= \frac{1}{\sigma^4} \sum_{i,j,k,l(\neq\alpha)} \text{cov}[A_{\alpha i}[e^{A_{\setminus\alpha} t}]_{ij} A_{j\alpha}, A_{\alpha k}[e^{A_{\setminus\alpha} t}]_{kl} A_{l\alpha}] \\ &= \frac{1}{N^2} \sum_{i,j,k,l(\neq\alpha)} \left((\delta_{ij}\delta_{kl} + \delta_{il}\delta_{jk} + \delta_{ik}\delta_{jl}) [e^{A_{\setminus\alpha} t}]_{ij} [e^{A_{\setminus\alpha} t}]_{kl} \right. \\ &\quad \left. - \delta_{ij}\delta_{kl} [e^{A_{\setminus\alpha} t}]_{ij} [e^{A_{\setminus\alpha} t}]_{kl} \right) \\ &= \frac{2}{N} \cdot \frac{1}{N} \text{Tr} \left[e^{A_{\setminus\alpha} t} (e^{A_{\setminus\alpha} t})^T \right] = O\left(\frac{1}{N}\right).\end{aligned}\tag{24}$$

This calculation shows that $\nu_\alpha(t)$ itself (not just its average) has $O(1/N)$ fluctuations under quenched disorder. This means that we can replace it (for all shuffling randomness and all quenched disorder) with its quenched average:

$$\nu_\alpha(t) \approx \langle \nu_\alpha(t) \rangle_{\mathcal{Q}}\tag{25}$$

2.5 Putting it together: the DMFT equation of motion

Returning to our definitions in Eqs. (11) and (12), we can put everything together to write the DMFT equation of motion for the random particle α (at frozen quenched disorder):

$$\begin{aligned}\frac{dx}{dt} &= \Phi(t) - \sigma^2 \int_0^t dt' \nu(t-t') x(t') \\ \Phi(t) &\sim \text{Gaussian Proc. (under shuffling randomness),} \\ \langle \Phi(t) \rangle_{\text{sh}} &= 0, \quad \langle \Phi(t) \Phi(s) \rangle_{\text{sh}} = \sigma^2 \sigma_0^2 \frac{J_1(2\sigma|t-s|)}{\sigma|t-s|}, \\ \nu(t) &= \frac{J_1(2\sigma t)}{\sigma t} \Theta(t).\end{aligned}\tag{26}$$

We have dropped the subscript α because the equation is valid for any random particle we choose. An equation of this form is the dynamical mean-field theory (DMFT) equation of motion. In section 4, we will discuss how this equation can provide insights into the dynamics of the system.

3 The cavity method

So far, we have derived the DMFT equation of motion for a single random particle in a system of N particles by exactly solving the equations of

motion for the other $N - 1$ particles in terms of the random particle and then substituting this solution into the equation of motion for the random particle. An advantage of this approach was that it allowed us to identify in an exact way how the quenched disorder and shuffling randomness can be treated interchangeably for this system. However, this approach is not very useful for non-linear systems. There are mainly two other approaches to DMFT: the cavity method and the replica method. In this note, we will only discuss the cavity method, which is often more direct and intuitive than the replica method.

The cavity method is a powerful technique for analyzing the dynamics of systems with quenched disorder. It is based on the observation that order parameters like $\frac{1}{N} \sum_{i=1}^N x_i(t)$ and $\frac{1}{N} \sum_{i=1}^N x_i(t)x_i(s)$ will be nearly the same in a system of size N and a system of size $N + 1$ in the thermodynamic limit $N \gg 1$. In the cavity method, one adds a new particle to the system and then analyzes how the dynamics of the system with $N + 1$ particles is related to the dynamics of the system with N particles. At the end of the cavity calculation, one finds that the equation of motion for the new particle can be written in terms of the order parameters of the system with N particles. Let's see how this works.

Before we proceed, we'll write our system of equation as:

$$\frac{dx_i}{dt} = \sum_{j=1}^N A_{ij}x_j(t) + h_i(t), \quad (27)$$

where $h_i(t)$ is some auxiliary field that we can perturb to analyze the system. At the end of the entire calculation, we will set $h_i(t) = 0$ and get the original equations of motion back.

Figure 2 shows a schematic of the cavity method.

3.1 Step 1: add the cavity degree of freedom

The first step in the cavity method is to add a new particle to the system which is called the cavity degree of freedom and is labeled $i = 0$. The new equations of motion for the system with $N + 1$ particles are:

$$\begin{aligned} \frac{dx_i}{dt} &= \sum_{j=1}^N A_{ij}x_j(t) + A_{i0}x_0(t) + h_i(t), \quad i = 1, \dots, N, \\ \frac{dx_0}{dt} &= \sum_{i=1}^N A_{0i}x_i(t) + A_{00}x_0(t) + h_0(t). \end{aligned} \quad (28)$$

Notice that we can drop the second term in the second equation because A_{ij} is skew-symmetric. In more general contexts, this term might not be zero, but we still drop it because it is of order $1/\sqrt{N}$.

3.2 Step 2: treat the cavity degree of freedom as a perturbation

Notice that for particles $i = 1, \dots, N$, the equations of motion are the same as before except for the new term $A_{i0}x_0(t)$. We can treat this as




Figure 2 – Schematic of the cavity method. (a) Add a new ‘cavity’ degree of freedom $x_0(t)$ to the system which adds a new, small term, $A_{i0}x_0(t)$, to each of the equations of motion for the other particles. This is detailed in step 1 (section 3.1). (b) Treat the addition of the cavity degree of freedom as a small perturbation to the dynamics of the other particles and apply linear response theory to relate the dynamics of the system with $N + 1$ particles to the dynamics of the system with N particles. This results in an expression for $x_i(t)$ in terms of the dynamics before the perturbation and a memory term involving a dynamical susceptibility $\nu_{ij}(t, t')$ which describes how the dynamics of particle i responds at time t to a perturbation to particle j at time t' . This is detailed in step 2 (section 3.2). (c) The cavity degree of freedom has dynamics driven by the dynamics of the other particles. Substitute the expression for $x_i(t)$ into the equation of motion for the cavity degree of freedom to obtain an equation of motion for the cavity degree of freedom in terms of the dynamics of the original system with N particles, the new coefficients A_{0i} and A_{j0} , the dynamical susceptibility $\nu_{ij}(t, t')$, and the cavity degree of freedom itself. This is detailed in step 3 (section 3.3). (d) Perform a statistical analysis of the equation of motion in the thermodynamic limit $N \gg 1$ to obtain the DMFT equation of motion for the cavity degree of freedom which includes two terms: a Gaussian process which describes direct influence from the other particles and a non-local term which describes the back-reaction the cavity degree of freedom experiences after it perturbs the other particles. This is detailed in section 3.4. The final step is to ‘close’ the DMFT by expressing the order parameters of the system in terms of the cavity degree of freedom which is discussed in section 3.5.

an effective perturbation to the auxiliary field $h_i(t)$:

$$h_i(t) \rightsquigarrow h_i(t) + A_{i0}x_0(t). \quad (29)$$

Because A_{i0} is $O(1/\sqrt{N})$, we can treat this as a small perturbation and use perturbation theory to relate the dynamics of the system with $N+1$ particles to the dynamics of the system with N particles. The linear response approximation is:¹³

$$x_i(t) = x_{i\setminus 0}(t) + \sum_{j=1}^N \int_0^t dt' \nu_{ij}(t, t') A_{j0} x_0(t') + O(N^{-1/2}), \quad (30)$$

where: $\nu_{ij}(t, t') \equiv \left. \frac{\delta x_i(t)}{\delta h_j(t')} \right|_{h=0}.$

3.3 Step 3: obtain the cavity equation of motion

Now, we can substitute this expression for $x_i(t)$ into the equation of motion for the cavity degree of freedom:

$$\frac{dx_0}{dt} = \sum_{i=1}^N A_{0i} x_{i\setminus 0}(t) + \sum_{i,j=1}^N \int_0^t dt' \nu_{ij}(t, t') A_{0i} A_{j0} x_0(t') + h_0(t). \quad (31)$$

What's notable about this equation is that it is an equation for $x_0(t)$ purely in terms of the dynamics of the system with N particles, the new coefficients A_{0i} and A_{j0} , and x_0 itself.¹⁴ To keep things tidy, we define:

$$\Phi_0(t) \equiv \sum_{i=1}^N A_{0i} x_{i\setminus 0}(t). \quad (32)$$

3.4 Step 4: statistics under quenched disorder

So far, we have assumed that the coefficients A_{ij} are fixed and do not fluctuate. Now, let's consider how each term in Eq. (31) fluctuates under quenched disorder. Notice that the first term, $\Phi_0(t)$, is a sum of many normal random variables $\{A_{0i}\}$, so it is a Gaussian process. This means that $\Phi_0(t)$ is fully characterized by its mean and covariance:

$$\langle \Phi_0(t) \rangle_{\mathcal{Q}} = \sum_{i=1}^N \langle A_{0i} \rangle \langle x_{i\setminus 0}(t) \rangle = 0 \quad (33)$$

$$\begin{aligned} \langle \Phi_0(t) \Phi_0(s) \rangle_{\mathcal{Q}} &= \sum_{i,j=1}^N \langle A_{0i} A_{0j} \rangle x_{i\setminus 0}(t) x_{j\setminus 0}(s) \\ &= \frac{\sigma^2}{N} \sum_{i,j=1}^N \delta_{ij} x_{i\setminus 0}(t) x_{j\setminus 0}(s) = \sigma^2 C(t, s), \end{aligned} \quad (34)$$

where we have defined:

$$C(t, s) \equiv \frac{1}{N} \sum_{i=1}^N x_{i\setminus 0}(t) x_{i\setminus 0}(s), \quad (35)$$

¹³ Here, $\frac{\delta F[h(t)]}{\delta h(t')}$ is the functional derivative of the functional $F[h(t)]$ with respect to the field $h(t')$. It measures how the function $F[h(t)]$ changes to linear order when we perturb the field $h(t)$ at time t' . It is defined so that $\frac{\delta h(t)}{\delta h(t')} = \delta(t - t')$ where δ is the Dirac delta function.

¹⁴ Note that up to this point, everything we have done is exact as $N \rightarrow \infty$. The only approximation we have made is that we have treated the new term $A_{i0}x_0(t)$ as a small perturbation, but the error is of order $O(1/\sqrt{N})$ which is negligible in the thermodynamic limit $N \gg 1$.

to be the empirical correlation function. As for the second term:

$$\begin{aligned} \left\langle \sum_{i,j=1}^N \nu_{ij}(t, t') A_{0i} A_{j0} \right\rangle &= \sum_{i,j=1}^N \nu_{ij}(t, t') \langle A_{0i} A_{j0} \rangle \\ &= -\frac{\sigma^2}{N} \sum_{i,j=1}^N \delta_{ij} \nu_{ij}(t, t') = -\sigma^2 \nu(t, t'), \end{aligned} \quad (36)$$

where we have defined:

$$\nu(t, t') \equiv \frac{1}{N} \sum_{i=1}^N \nu_{ii}(t, t') = \frac{1}{N} \sum_{i=1}^N \left. \frac{\delta x_i(t)}{\delta h_i(t')} \right|_{h=0}, \quad (37)$$

to be the average dynamical susceptibility. The variance is:¹⁵

$$\begin{aligned} \text{var}_{\mathcal{Q}} \left[\sum_{i,j=1}^N \nu_{ij}(t, t') A_{0i} A_{j0} \right] &= \sum_{i,j,k,l=1}^N \text{cov}[A_{0i} A_{j0}, A_{0k} A_{l0}] \nu_{ij}(t, t') \nu_{kl}(t, t') \\ &= \frac{\sigma^2}{N^2} \sum_{i,j,k,l=1}^N (\delta_{ik} \delta_{jl} + \delta_{il} \delta_{jk}) \nu_{ij}(t, t') \nu_{kl}(t, t') \\ &= \frac{\sigma^2}{N} \left(\underbrace{\frac{1}{N} \sum_{i,j=1}^N \nu_{ij}(t, t') \nu_{ij}(t, t')}_{O(1)} + \underbrace{\frac{1}{N} \sum_{i,j=1}^N \nu_{ij}(t, t') \nu_{ji}(t, t')}_{O(1)} \right) \\ &= O\left(\frac{1}{N}\right). \end{aligned} \quad (38)$$

In the final step of the calculation, we again used that the scaled trace of a random matrix like $\frac{1}{N} \text{Tr } M$ is $O(1)$ in the thermodynamic limit $N \gg 1$. This result means that the second term in Eq. (31) has $O(1/N)$ fluctuations under quenched disorder and can thus be replaced with its quenched average:

$$\sum_{i,j=1}^N \nu_{ij}(t, t') A_{0i} A_{j0} \approx -\sigma^2 \nu(t, t'). \quad (39)$$

Putting this together, we reduce the cavity equation of motion in Eq. (31) to effective dynamics (where fluctuations are due to quenched disorder):

$$\begin{aligned} \frac{dx_0}{dt} &= \Phi_0(t) - \sigma^2 \int_0^t dt' \nu(t, t') x_0(t') + h_0(t) \\ \Phi_0(t) &\sim \text{Gaussian Proc. (under quenched disorder),} \end{aligned}$$

$$\langle \Phi_0(t) \rangle_{\mathcal{Q}} = 0, \quad \langle \Phi_0(t) \Phi_0(s) \rangle_{\mathcal{Q}} = \sigma^2 C(t, s) = \sigma^2 \frac{1}{N} \sum_{i=1}^N x_{i \setminus 0}(t) x_{i \setminus 0}(s),$$

$$\nu(t, t') = \frac{1}{N} \sum_{i=1}^N \left. \frac{\delta x_i(t)}{\delta h_i(t')} \right|_{h=0}. \quad (40)$$

3.5 Step 5: assert self-averaging to get self-consistent equations

We now have an equation of motion for the cavity degree of freedom $x_0(t)$ in terms of the order parameters $C(t, s)$ and $\nu(t, t')$. However,

¹⁵ Here we used Wick's theorem to compute the expectation of the product of four Gaussian random variables. This is discussed in more detail in appendix section A.5.

based on the DMFT equation alone, we have no way to compute these order parameters. To compute them, we need to make an assumption about the self-averaging of the system. In particular, we assert that for the order parameters, their quenched averages are equal to their shuffled averages as $N \rightarrow \infty$:

$$\begin{aligned} C(t, s) &= \overbrace{\frac{1}{N} \sum_{i=1}^N x_{i \setminus 0}(t) x_{i \setminus 0}(s)}^{\text{"shuffled average"}} = \overbrace{\langle x_0(t) x_0(s) \rangle_{\mathcal{Q}}}^{\text{"quenched average"}} \\ \nu(t, t') &= \frac{1}{N} \sum_{i=1}^N \left. \frac{\delta x_i(t)}{\delta h_i(t')} \right|_{h=0} = \left\langle \left. \frac{\delta x_0(t)}{\delta h_0(t')} \right|_{h=0} \right\rangle_{\mathcal{Q}} \end{aligned} \quad (41)$$

These sorts of expressions are called *closure* or *self-consistency* conditions for the DMFT.¹⁶ We showed that these self-averaging properties are true for our model in the previous section by explicitly computing the fluctuations of the order parameters under quenched disorder. However, this is not possible for most systems, so the above ‘closing’ of the DMFT equations is usually done by *assuming* that the order parameters are self-averaging.

3.6 Step 6: solve the self-consistent equations (hard!)

We’ll first solve for the dynamical susceptibility $\nu(t, t')$. We do this by taking a functional derivative of the DMFT equation of motion with respect to $h_0(t')$:

$$\begin{aligned} \frac{\partial}{\partial t} \nu(t, t') &= \frac{\partial}{\partial t} \left\langle \frac{\delta x_0(t)}{\delta h_0(t')} \right\rangle = \left\langle \frac{\delta}{\delta h_0(t')} \frac{dx_0(t)}{dt} \right\rangle \\ &= -\sigma^2 \int_0^t dt'' \nu(t, t'') \nu(t'', t') + \delta(t - t'). \end{aligned} \quad (42)$$

This along with the initial condition $\nu(t, t') = 0$ for $t < t'$ is a self-consistent equation for $\nu(t, t')$. This equation implies that $\nu(t, t')$ is a function of the difference $t - t'$ alone, so we can write $\nu(t, t') = \nu(t - t')$. Let:

$$\tilde{\nu}(z) \equiv \int_0^\infty dt \nu(t) e^{-zt}, \quad (43)$$

be the Laplace transform of $\nu(t)$. Then Eq. (42) becomes:

$$z \tilde{\nu}(z) = -\sigma^2 [\tilde{\nu}(z)]^2 + 1 \implies \tilde{\nu}(z) = \frac{1}{2\sigma^2} \left(\sqrt{z^2 + 4\sigma^2} - z \right) \quad (44)$$

The inverse Laplace transform of this expression gives:

$$\nu(t) = \frac{J_1(2\sigma t)}{\sigma t} \Theta(t), \quad (45)$$

where $J_1(x)$ is the Bessel function of the first kind of order 1 and $\Theta(t)$ is the Heaviside step function (which is 1 for $t > 0$ and 0 for $t < 0$).¹⁷ Now that we have $\nu(t)$, we can obtain an explicit expression for $x_0(t)$ in terms of $\Phi(t)$. If we write:

$$x_0(t) = G(t) x_0(0) + \int_0^t dt' G(t - t') \Phi_0(t'), \quad (46)$$

¹⁶ Another order parameter,

$$m(t) \equiv \frac{1}{N} \sum_{i=1}^N x_{i \setminus 0}(t),$$

often appears in DMFT calculations, but it did not show up here because it is zero for our model. The closure/self-consistency condition for $m(t)$ is:

$$m(t) = \langle x_0(t) \rangle_{\mathcal{Q}}.$$

¹⁷ There are a few ways to compute this inverse Laplace transform, but the easiest way is to look at a table of Laplace transforms (on Wikipedia, for example) or use computer algebra software like Mathematica.

and substitute this back into the Eq. (40), we find that this is a general solution to the DMFT equation of motion when:

$$\frac{dG}{dt} = \delta(t) - \sigma^2 \int_0^t dt' \nu(t-t')G(t') \quad (47)$$

with initial condition $G(t) = 0$ for $t < 0$. This equation can be solved using the Laplace transform:

$$z\tilde{G}(z) = 1 - \sigma^2 \tilde{\nu}(z)\tilde{G}(z) \quad (48)$$

$$\Rightarrow \tilde{G}(z) = \frac{1}{z + \sigma^2 \tilde{\nu}(z)} = \frac{1}{2\sigma^2} (\sqrt{z^2 + 4\sigma^2} - z) \quad (49)$$

which we have already seen has the inverse Laplace transform: ¹⁸

$$G(t) = \frac{J_1(2\sigma t)}{\sigma t} \Theta(t). \quad (50)$$

To close the DMFT equations, we compute:

$$\begin{aligned} C(t, s) &= \langle x_0(t)x_0(s) \rangle \\ &= \sigma_0^2 G(t)G(s) + \sigma^2 \int_0^t dt' \int_0^s ds' G(t-t')G(s-s')C(t', s') \end{aligned} \quad (51)$$

If we define the two-variable Laplace transform:

$$\tilde{C}(z_t, z_s) = \int_0^\infty dt \int_0^\infty ds e^{-z_t t - z_s s} C(t, s), \quad (52)$$

we can apply the convolution theorem to write:

$$\tilde{C}(z_t, z_s) = \sigma_0^2 \tilde{G}(z_t)\tilde{G}(z_s) + \sigma^2 \tilde{G}(z_t)\tilde{G}(z_s)\tilde{C}(z_t, z_s) \quad (53)$$

which can be rearranged to give:

$$\tilde{C}(z_t, z_s) = \frac{\sigma_0^2 \tilde{G}(z_t)\tilde{G}(z_s)}{1 - \sigma^2 \tilde{G}(z_t)\tilde{G}(z_s)} \quad (54)$$

Using Eq. (48), we can write this as:

$$\tilde{C}(z_t, z_s) = \sigma_0^2 \frac{\tilde{G}(z_t)\tilde{G}(z_s)}{z_t + z_s} \quad (55)$$

In general:

$$\int_0^\infty dt \int_0^\infty ds e^{-z_t t - z_s s} f(t-s) = \frac{\tilde{f}(z_t)\tilde{f}(z_s)}{z_t + z_s} \quad (56)$$

for any well-behaved function $f(t)$ which tells us that $C(t, s)$ is a function of the difference $t - s$ alone and that:

$$C(t-s) = \sigma_0^2 G(|t-s|) = \sigma_0^2 \frac{J_1(2\sigma|t-s|)}{\sigma|t-s|}. \quad (57)$$

In the long time limit:

$$C(t) \sim \frac{\sigma_0^2}{\sqrt{\pi}\sigma^{3/2}t^{3/2}} \sin\left(2\sigma t - \frac{\pi}{4}\right), \quad \text{for } t \gg 1/\sigma. \quad (58)$$

¹⁸ The fact that $G(t) = \nu(t)$ appears in the solution is special to this model and is not true in general.




Figure 3 – Plot of the correlation function $C(t-s) = \langle x(t)x(s) \rangle$. The oscillations have period π/σ . Because $C(t) = \sigma_0^2 G(t) = \sigma_0^2 \nu(t)$, this plot also shows the shape of the dynamical susceptibility $\nu(t)$ and the Green's function $G(t)$.

This means that there are coherent oscillations in $x(t)$ which decay slowly as $t^{-3/2}$. Putting this all together Eq. (40) becomes:

$$\begin{aligned} \frac{dx}{dt} &= \Phi(t) - \sigma^2 \int_0^t dt' \nu(t, t') x(t') \\ \Phi(t) &\sim \text{Gaussian Proc. (under shuffling disorder),} \\ \langle \Phi(t) \rangle_{\mathcal{Q}} &= 0, \quad \langle \Phi(t) \Phi(s) \rangle_{\mathcal{Q}} = \sigma^2 \sigma_0^2 \frac{J_1(2\sigma|t-s|)}{\sigma|t-s|}, \\ \nu(t, t') &= \frac{1}{N} \sum_{i=1}^N \left. \frac{\delta x_i(t)}{\delta h_i(t')} \right|_{h=0} = \frac{J_1(2\sigma|t-t'|)}{\sigma|t-t'|} \Theta(t-t'), \end{aligned} \quad (59)$$

which precisely matches the DMFT equation of motion we derived in section 2.

3.7 Why was it necessary to add the cavity degree of freedom?

When deriving the DMFT equation of motion, we might have been tempted to look at the sum $\sum_{j(\neq i)} A_{ij} x_j(t)$, recognize that it is the sum of many normal random variables, and conclude that it is a Gaussian process with mean zero and variance $\sigma^2 C(t, s)$. Knowing what we know now, we can see that this would miss the back-reaction term, $-\sigma^2 \int_0^t dt' \nu(t, t') x_i(t')$, which is crucial for the DMFT. We miss this back-reaction term because $x_i(t)$ has $O(1/\sqrt{N})$ contributions to the dynamics of all other degrees of freedom $j \neq i$ which each then have $O(1/\sqrt{N})$ contributions to the dynamics of $x_i(t)$ which all add up to give an $O(1)$ contribution to the dynamics of $x_i(t)$. This feedback is not captured by simply taking averages of $\sum_{j(\neq i)} A_{ij} x_j(t)$ because we essentially assumed that x_j is independent of A_{ij} , which would be the case if $x_j(t)$ were a deterministic function, but we should actually treat it as a stochastic process that is correlated with A_{ij} . When we add the cavity degree of freedom, we instead get the sum: $\Phi_0(t) = \sum_{i=1}^N A_{0i} x_{i\backslash 0}(t)$. Because A_{0i} and $x_{i\backslash 0}(t)$ are independent (by definition), we are free to treat averages over A_{0i} and $x_{i\backslash 0}(t)$ separately.

4 A taste of possible analyses

Now that we have derived the DMFT equation of motion, let's explore how it might be used to analyze the dynamics of the system.

4.1 Correlations

The most immediate thing to ask is: if I look at the state of a random degree of freedom $x_i(t)$ at time t , how correlated is it with its own state at some later time $t + \tau$? This question is only well-defined in the case where we are looking at a random degree and integrating out the other degrees of freedom, otherwise everything is deterministic and perfectly correlated with itself. In solving the self-consistency equations in the previous section, we already determined $C(t, s) = \langle x_i(t) x_i(s) \rangle =$




Figure 4 – Plot of the dynamics of a single degree of freedom $x_\alpha(t)$ in a system $N = 64$.

$\frac{1}{N} \sum_{i=1}^N x_i(t)x_i(s)$. We can then determine the correlations:

$$\begin{aligned} \text{corr}[x_i(t), x_i(t+\tau)] &= \frac{\langle x_i(t)x_i(t+\tau) \rangle - \langle x_i(t) \rangle \langle x_i(t+\tau) \rangle}{\sqrt{\text{var}[x_i(t)] \text{var}[x_i(t+\tau)]}} \\ &= \frac{C(t, t+\tau)}{\sigma_0^2} = \frac{J_1(2\sigma\tau)}{\sigma\tau}. \end{aligned} \quad (60)$$

The shape of this function is shown in figure 5 along with a comparison to numerical simulations. We can see that there are oscillations between positive and negative correlation with approximate period π/σ which decay slowly as $\tau^{-3/2}$. This is remarkable because it says that there is a quasi-periodic structure to the dynamics of each degree of freedom which is coherent for a long time. This quasi-periodic structure can be seen in figure 4. Ultimately, this is a consequence of the fact that the matrix A_{ij} is skew-symmetric, meaning that its eigenvalues are purely imaginary, corresponding to oscillatory dynamics.

4.2 Response to perturbations

Another interesting question is: if I perturb a random degree of freedom x_i at time t_0 by a small amount, how much does it affect its own state at some later time $t_0 + t$? Because our equation is linear, we can interpret this as adding a small perturbation $h_i(t) = \epsilon\delta(t - t_0)$ to the right-hand-side of the equation of motion for $x_i(t)$. This response to perturbations is then captured by the dynamical susceptibility $\delta x_i(t_0 + t)/\delta h_i(t_0)$. Our result was that the average dynamical susceptibility is:

$$\nu(t_0 + t, t_0) = \frac{1}{N} \sum_{i=1}^N \left. \frac{\delta x_i(t_0 + t)}{\delta h_i(t_0)} \right|_{h=0} = \frac{J_1(2\sigma t)}{\sigma t}. \quad (61)$$

This functional form is identical to the correlation function, so you can see a plot of $\nu(t)$ in Fig. 3. This tells us that whether the effect of a perturbation is positive or negative depends on the time delay t in a quasi-periodic way and that the effect of the perturbation decays slowly as $t^{-3/2}$.

A common quantity of interest, especially in non-linear systems, is the Lyapunov exponent, which measures how quickly two nearby trajectories diverge from one another. Assume we perturb the system at time t_0 by some small amount ϵ in a direction ξ so that: $\mathbf{x}'(t_0) = \mathbf{x}(t_0) + \epsilon \xi$. The Lyapunov exponent is defined as:

$$\lambda = \lim_{\epsilon \rightarrow 0} \lim_{t \rightarrow \infty} \frac{1}{t} \ln \frac{\|\mathbf{x}'(t_0 + t) - \mathbf{x}(t_0 + t)\|}{\epsilon} \quad (62)$$

In the case where $\xi_i \sim \mathcal{N}(0, 1)$ are i.i.d. normal random variables, we can write the average Lyapunov exponent as:¹⁹

$$\langle \lambda \rangle_\xi = \lim_{t_0, t \rightarrow \infty} \frac{1}{2t} \ln \frac{1}{N} \sum_{i,j=1}^N \nu_{ij}(t_0 + t, t)^2, \quad (63)$$

where the average is taken over the random perturbation direction ξ [9]. The quantity $\frac{1}{N} \sum_{i,j=1}^N \nu_{ij}(t_0 + t, t)^2$ is self-averaging over quenched disorder and also happens to be the second moment of the eigenvalue




Figure 5 – Plot of the covariance function $C(t_0, t_0 + t) = \frac{1}{N} \sum_{i=1}^N x_i(t_0)x_i(t_0 + t)$ from the DMFT prediction (solid line) and from numerical simulations (dots). The numerical simulations were done with $N = 128$ and averaged over 32 random realizations of the dynamics. The error bars are the standard deviation of $C(t_0, t_0 + t)$ from replicate to replicate.

¹⁹ Taking the average over ξ is necessary mainly for formal mathematical reasons, but usually does not matter in practice because the Lyapunov exponent is self-averaging over ξ . Additionally, this sort of Lyapunov exponent ultimately converges numerically to the same value for almost all choices of ξ , so long as ξ is chosen ‘generically’ (i.e. not in a special direction).

spectrum of $\nu_{ij}(t_0 + t, t)$. To determine this quantity, one performs a two-particle cavity calculation in which one introduces two cavity degrees of freedom. This quantity has historically played an important role in the study of chaotic dynamics in neural networks [9–12]. In our simple linear model, this quantity is exactly zero for all choices of parameters because the eigenvalues of A_{ij} are purely imaginary.

4.3 Fourier modes and power spectrum

We can also ask: if I look at the Fourier modes of a random degree of freedom $x_i(t)$, how much power is in each mode? This measures how much of the dynamics is happening at different frequencies. The power spectral density of $x_i(t)$ at frequency ω is defined as:

$$S_x(\omega) = \left\langle \lim_{T \rightarrow \infty} \frac{1}{T} \left| \int_0^T dt x_i(t) e^{-i\omega t} \right|^2 \right\rangle. \quad (64)$$

It is the average total overlap of $x_i(t)$ with the Fourier mode $e^{i\omega t}$ per unit time. Because the DMFT for $x_i(t)$ is a Gaussian process²⁰, we can compute the power spectral density using the Wiener–Khinchin theorem:

$$S_x(\omega) = \int_{-\infty}^{\infty} dt C(t) e^{-i\omega t} = 2\sigma_0^2 \int_0^{\infty} dt \frac{J_1(2\sigma t)}{\sigma t} e^{-i\omega t} \quad (65)$$

To evaluate this integral, we use the Laplace identity:

$$\mathcal{L}_t \left[\frac{J_1(2\sigma t)}{\sigma t} \right] (z) = \frac{1}{2\sigma^2} \left(\sqrt{z^2 + 4\sigma^2} - z \right). \quad (66)$$

However, we are working in the complex plane, so we need to be careful about branch cuts of the square root function. Choosing the branch cut that ensures $S_x(\omega) \geq 0$ for all ω , we find:

$$S_x(\omega) = \begin{cases} \frac{\sigma_0^2}{\sigma^2} \sqrt{4\sigma^2 - \omega^2}, & |\omega| < 2\sigma \\ 0, & |\omega| > 2\sigma \end{cases}. \quad (67)$$

This is not too surprising because this is the same shape as the Wigner semicircle distribution, which is the eigenvalue spectrum of A_{ij} . This result means that the dynamics of $x_i(t)$ are dominated by frequencies in the range $[-2\sigma, 2\sigma]$ and that there is a sharp cutoff in power at these frequencies. Additionally, this means that if we were to drive the system with an external oscillatory input at frequency ω , the system would only respond if $|\omega| < 2\sigma$. Because the power spectral density peaks at $\omega = 0$ and increases as $|\omega| \rightarrow 0$, the system is most responsive to low-frequency inputs and has more slow dynamics than fast dynamics.

5 A non-linear example: the generalized Lotka–Volterra model

Finally, we briefly discuss how the cavity method can be applied to a non-linear model, specifically the generalized Lotka–Volterra (GLV) model of interacting species. This derivation roughly follows Ref. [3].

²⁰ Note that this sort of analysis might not work for all models because the application of the Wiener–Khinchin theorem requires that the process is stationary, ergodic, and Gaussian which is often not the case for non-linear models.




Figure 6 – Power spectral density $S_x(\omega)$ of the dynamics of a random degree of freedom $x_i(t)$ in the DMFT model. This quantity is the average total overlap of $x_i(t)$ with the Fourier mode $e^{i\omega t}$ per unit time.

5.1 The model

The GLV model describes the populations N_i of species $i = 1, \dots, S$. It is described by the coupled ordinary differential equations: [5]

$$\frac{dN_i}{dt} = N_i \left(1 - N_i - \sum_{j(\neq i)} A_{ij} N_j \right), \quad i = 1, \dots, S. \quad (68)$$

This is a phenomenological model of species interactions in which a species i grows logistically in isolation, but in the presence of another species j , its per-capita growth is decreased by an amount proportional to the population of species j . The amount of this decrease in growth is determined by the interaction matrix A_{ij} . In the $S \gg 1$ case, one assumes that the interaction matrix A_{ij} has random entries. We take A_{ij} to have the statistics:

$$\langle A_{ij} \rangle = \frac{\mu}{S}, \quad \text{cov}[A_{ij}, A_{kl}] = \frac{\sigma^2}{S} (\delta_{ik} \delta_{jl} + \gamma \delta_{il} \delta_{jk}) \quad (69)$$

This setup means that the variance of each entry is σ^2/S and the correlations between the entries are zero except for $\text{corr}[A_{ij}, A_{ji}] = \gamma$. When $\gamma = 1$, the model is symmetric, meaning that i effects j in the same way that j effects i . When $\gamma = -1$, the model is anti-symmetric, meaning that as i inhibits j , j promotes i .

The random GLV model has a very rich phenomenology including phase transitions, chaotic dynamics, marginal stability, aging, alternative stable states, and more [2, 3, 5, 13]. Properties of the GLV model and its relevance to real ecosystems are an active area of research which has led to the development of novel theoretical tools and fascinating results.

5.2 The cavity method

Unlike our simple linear model, we cannot simply solve the equations of motion and integrate out all but one degree of freedom to get the DMFT equation of motion. Instead, we invoke the cavity method.

5.2.1 Step 1: Add a species

The first step is to add a cavity species $i = 0$ to the system:

$$\begin{aligned} \frac{dN_i}{dt} &= N_i \left(1 - N_i - \sum_{j(\neq i)} A_{ij} N_j + A_{i0} N_0 + h_i(t) \right), \quad i = 1, \dots, S, \\ \frac{dN_0}{dt} &= N_0 \left(1 - N_0 - \sum_{i=1}^S A_{0i} N_i + A_{00} N_0 + h_0(t) \right) \end{aligned} \quad (70)$$

In the above, we have added an auxiliary field $h_i(t)$ to the equations of motion for each degree of freedom which we will set to zero at the end.

5.2.2 Step 2: Treat the new species as a perturbation

The addition of the new species N_0 adds a term $A_{i0} N_0$ to the equations of motion for each species $i = 1, \dots, S$. This new term is $O(N^{-1/2})$, so

we can treat it as a small perturbation in the $S \gg 1$ limit. In particular, we can treat it as a shift in the auxiliary field $h_i(t)$:

$$h_i(t) \rightsquigarrow h_i(t) + A_{i0}N_0(t). \quad (71)$$

If we define the dynamical susceptibility matrix:

$$\nu_{ij}(t, t') \equiv \left. \frac{\delta N_i(t)}{\delta h_j(t')} \right|_{h=0}, \quad (72)$$

we can relate the dynamics of the system with $S + 1$ species to the dynamics of the system with S species:

$$N_i(t) \approx N_{i \setminus 0}(t) + \sum_{j=1}^S \int_0^t dt' \nu_{ij}(t, t') A_{j0} N_0(t'). \quad (73)$$

This linear response approximation becomes exact in the $S \gg 1$ limit.

5.2.3 Step 3: Obtain the cavity equation of motion

Substituting this expression into the equation of motion for $N_0(t)$, we get:

$$\frac{dN_0}{dt} = N_0 \left(1 - N_0 - \sum_{i=1}^S A_{0i} N_{i \setminus 0}(t) - \sum_{i,j=1}^S \int_0^t dt' \nu_{ij}(t, t') A_{0i} A_{j0} N_0(t') \right). \quad (74)$$

Note that we have dropped the A_{00} term because it is $O(1/\sqrt{S})$ and will not contribute in the $S \gg 1$ limit. This is the cavity equation of motion for the added species $N_0(t)$, and it is exact when we take $S \rightarrow \infty$.

5.2.4 Step 4: Statistics under quenched disorder

Next, we analyze the statistics of the terms in the cavity equation of motion under quenched disorder and discard $O(N^{-1/2})$ contributions. The statistics of the direct contribution term are:

$$\begin{aligned} \left\langle \sum_{i=1}^S A_{0i} N_{i \setminus 0}(t) \right\rangle &= \frac{\mu}{S} \sum_{i=1}^S N_{i \setminus 0}(t) = \mu m(t), \\ \left\langle \left(\sum_{i=1}^S A_{0i} N_{i \setminus 0}(t) \right) \left(\sum_{j=1}^S A_{0j} N_{j \setminus 0}(t') \right) \right\rangle &= \sum_{i,j=1}^S \langle A_{0i} A_{0j} \rangle N_{i \setminus 0}(t) N_{j \setminus 0}(t') \\ &= \frac{\sigma^2}{S} \sum_{i,j=1}^S \delta_{ij} N_{i \setminus 0}(t) N_{j \setminus 0}(t') = \sigma^2 C(t, t'), \end{aligned} \quad (75)$$

where we have defined:

$$m(t) = \frac{1}{S} \sum_{i=1}^S N_{i \setminus 0}(t), \quad C(t, t') = \frac{1}{S} \sum_{i=1}^S N_{i \setminus 0}(t) N_{i \setminus 0}(t'), \quad (76)$$

to be the empirical average and auto-covariance of the species abundances. The back-reaction term has statistics:

$$\begin{aligned}
\left\langle \sum_{i,j=1}^S \nu_{ij}(t, t') A_{0i} A_{j0} \right\rangle &= \gamma \sigma^2 \nu(t, t') \\
\text{var} \left[\sum_{i,j=1}^S \nu_{ij}(t, t') A_{0i} A_{j0} \right] &= \sum_{i,j,k,l=1}^S \nu_{ij}(t, t') \nu_{kl}(t, t') \text{cov}[A_{0i} A_{j0}, A_{0k} A_{l0}] \\
&= \frac{\sigma^4}{S^2} \sum_{i,j,k,l=1}^S \nu_{ij}(t, t') \nu_{kl}(t, t') (\delta_{ik} \delta_{jl} + \gamma^2 \delta_{il} \delta_{jk}) \\
&= \frac{\sigma^4}{S^2} \sum_{i,j=1}^S \nu_{ij}(t, t') \nu_{ij}(t, t') + \gamma^2 \frac{\sigma^4}{S^2} \sum_{i,j=1}^S \nu_{ij}(t, t') \nu_{ji}(t, t') \\
&= \frac{\sigma^4}{S} \underbrace{\left(\frac{1}{S} \text{Tr} [\nu(t, t') \nu(t, t')^T] + \gamma^2 \frac{1}{S} \text{Tr} [(\nu(t, t'))^2] \right)}_{O(1)} \\
&= O\left(\frac{1}{S}\right).
\end{aligned} \tag{77}$$

In the above, we have defined:

$$\nu(t, t') = \frac{1}{S} \sum_{i=1}^S \nu_{ii}(t, t') = \frac{1}{S} \sum_{i=1}^S \left. \frac{\delta N_i(t)}{\delta h_i(t')} \right|_{h=0}. \tag{78}$$

Thus, the back-reaction term is $O(1/\sqrt{S})$ and can be neglected in the $S \gg 1$ limit, meaning we can replace the back-reaction term with its mean. Putting it together:

$$\frac{dN_0}{dt} = N_0 \left(1 - N_0 - \mu m(t) - \sigma \eta(t) - \gamma \sigma^2 \int_0^t dt' \nu(t, t') N_0(t') \right), \tag{79}$$

where $\eta(t)$ is a zero-mean Gaussian process with covariance:

$$\langle \eta(t) \eta(t') \rangle = C(t, t'). \tag{80}$$

5.2.5 Step 5: Assert self-averaging

To close the cavity equation of motion, we need to assume that the order parameters $m(t)$, $C(t, t')$, and $\nu(t, t')$ are self-averaging:

$$\begin{aligned}
m(t) &= \frac{1}{S} \sum_{i=1}^S N_{i \setminus 0}(t) = \langle N_0(t) \rangle_{\mathcal{Q}}, \\
C(t, t') &= \frac{1}{S} \sum_{i=1}^S N_{i \setminus 0}(t) N_{i \setminus 0}(t') = \langle N_0(t) N_0(t') \rangle_{\mathcal{Q}}, \\
\nu(t, t') &= \frac{1}{S} \sum_{i=1}^S \left. \frac{\delta N_i(t)}{\delta h_i(t')} \right|_{h=0} = \left\langle \left. \frac{\delta N_0(t)}{\delta h_0(t')} \right|_{h=0} \right\rangle_{\mathcal{Q}}.
\end{aligned} \tag{81}$$

To be explicit, we can write an ODE for the dynamical susceptibility in which one writes $\tilde{\nu}(t, t') = \frac{\delta N_0(t)}{\delta h_0(t')}$ and obtains the differential equation:

$$\begin{aligned} \frac{\partial}{\partial t} \tilde{\nu}(t, t') &= \frac{\delta}{\delta h_0(t')} \frac{dN_0(t)}{dt} \\ &= N_0(t) \left(-\tilde{\nu}(t, t') - \gamma \sigma^2 \int_0^t dt'' \nu(t, t'') \tilde{\nu}(t'', t') + \delta(t - t') \right) \\ &\quad + \tilde{\nu}(t, t') \left(1 - N_0(t) - \mu m(t) - \sigma \eta(t) - \gamma \sigma^2 \int_0^t dt'' \nu(t, t'') N_0(t'') \right), \end{aligned} \quad (82)$$

and then obtains the dynamical susceptibility by averaging over resamplings of the DMFT stochastic process $N_0(t)$:²¹

$$\nu(t, t') = \langle \tilde{\nu}(t, t') \rangle_{\mathcal{Q}}. \quad (83)$$

It is not possible (to the best of my knowledge) to solve these equations analytically, but one can use numerical methods to solve them. To do this, one makes a guess for the order parameters $m(t)$, $C(t, t')$, and $\nu(t, t')$; solves the stochastic differential equation for $N_0(t)$ many times using the guessed order parameters; takes averages over the sampled trajectories to update the order parameters; and then repeats until convergence. Ref. [3] provides a detailed description of this procedure and its implementation in Python as well as suggestions to improve the convergence of the numerical method.

5.2.6 Analyzing the DMFT equation of motion

We can start ascribing meaning to the terms in the DMFT equation of motion. We see that the dynamics of a randomly selected species experiences a direct mean-field contribution from the average population of all species, $\mu m(t)$, which decreases the per-capita growth rate of the species. It additionally experiences fluctuations due to the random interactions with other species, $\sigma \eta(t)$, whose fluctuation size is set by the parameter σ and the typical abundance of species. The correlations of $\sigma \eta(t)$ decay the same way as the auto-correlation of the randomly selected species itself. Finally, the back-reaction term, $\gamma \sigma^2 \int_0^t dt' \nu(t, t') N_0(t')$, captures how the randomly selected species impacts the entire community and then feeds back into its own dynamics. Looking at Eq. (82), we can see that $\nu(t, t') \geq 0$ which is expected because if we increase the per-capita growth rate artificially (increase $h(t)$), we expect the species abundance to increase. We additionally see that $\nu(t, t) = \langle N_0(t) \rangle_{\mathcal{Q}} = m(t)$ and that $\nu(t, t')$ decreases as $|t - t'|$ increases, so the strength of the self-regulation due to back-reaction is set by $\gamma \sigma^2 m(t)$. Additionally, when $\gamma > 0$, we see that the back-reaction term indicates there is emergent self-limitation of a species' growth rate due to the interactions with other species. When $\gamma < 0$, the back-reaction term indicates that the species is promoted by its interactions with other species, which is a form of emergent cooperation.

5.2.7 Obtaining steady-state solutions

In the case where the system reaches a steady state, we can treat $N_0(t) = N_0$ as just a constant random variable, $m(t) = m$ a deterministic

²¹ Note that the notation $\tilde{\nu}(t, t')$ represents a single-realization dynamical susceptibility, while $\nu(t, t')$ represents the average over realizations. The former is not deterministic and is a stochastic process, while the latter is deterministic.

There is also another way to write a self-consistent equation for $\nu(t, t')$ which takes advantage of Novikov's theorem. See [3] appendix C for details or [14] for the original paper.

constant, and $\eta(t) = \eta$ a zero-mean Gaussian random variable with variance $q = C(t, t) = \langle N_0^2 \rangle_{\mathcal{Q}}$. For the back-reaction term, we perform the integral to define the integrated response: $\bar{\nu} = \int_0^\infty dt' \nu(t, t')$. This gives the static cavity equation:

$$0 = N_0 \left(1 - N_0 - \mu m - \sigma \sqrt{q} \xi - \gamma \sigma^2 \bar{\nu} N_0 \right), \quad (84)$$

where ξ is a zero-mean unit-variance Gaussian random variable [5]. We note that there are two possible solutions $N_0 > 0$ and $N_0 = 0$. We will select the $N_0 > 0$ solution whenever it exists because in ecology, we are generally interested in the uninhabitable steady-state (meaning the steady-state in which every species that could possibly stably coexist in the system is present). This means we can write the steady-state solution as:

$$N_0 = \max \left\{ 0, \frac{1 - \mu m - \sigma \sqrt{q} \xi}{1 + \gamma \sigma^2 \bar{\nu}} \right\}. \quad (85)$$

Because ξ is a zero-mean Gaussian random variable, this means that the steady-state distribution of species abundances is a truncated Gaussian distribution.

One can proceed to use this expression to write down self-consistent expressions for m , q , and $\bar{\nu}$, and these self-consistent equations are much easier to handle than the full DMFT equation of motion. Analyzing the resulting equations additionally reveals the existence of phase transitions in the system, such as a transition from stable steady-state behavior to chaos, alternative-stable states, or unbounded growth [2, 3, 5, 13]. However, we will not reproduce the full analysis here, and we refer the reader to the original literature for more details.

6 Conclusion

In this tutorial, we've taken a journey through the core concepts of Dynamical Mean-Field Theory. Our goal was to demystify this powerful framework by applying it to a simple linear model where every step could be justified with exact calculations. We saw how the dynamics of a system with N interacting variables can, in the large N limit, be described by the stochastic dynamics of a single, representative variable. This simplification hinges on the self-averaging property, where macroscopic quantities for a single realization of the system become identical to an average over all possible realizations of the disorder.

We then introduced the cavity method, a more general and powerful technique that does not rely on having an exactly solvable model. By treating a newly added 'cavity' particle as a small perturbation, we re-derived the same DMFT equations for our linear model. This illustrated how the cavity method systematically accounts for the crucial 'back-reaction' of the system on itself, a feedback loop that is easy to miss with more naive approximations. The key insight is that the cavity particle's dynamics are governed by a stochastic process whose statistical properties, the correlation $C(t, s)$ and response function $\nu(t, s)$, are determined self-consistently by the dynamics of the cavity particle itself.

Finally, we saw how to apply the cavity method to the non-linear generalized Lotka–Volterra model from theoretical ecology. This led us to

a DMFT equation that, while not solvable with pen and paper, provides insights into the system's behavior and can be tractably analyzed to explore a rich phase diagram exhibiting stability, chaos, and unbounded growth.

The tools presented here are the foundation for studying a vast array of complex systems in biophysics and beyond, from the chaotic firing of neural networks to the complex energy landscapes of glassy materials. We hope this gentle introduction has provided you with the intuition to explore the world of disordered systems and to apply these methods in your own research.

Funding Acknowledgments: This work was supported by the Fannie and John Hertz Foundation Fellowship and the National Science Foundation Graduate Research Fellowship Program.

A Appendix: Brief introduction to random matrix theory (RMT)

Random matrix theory (RMT) is a branch of mathematics and physics that studies the properties of matrices with randomly chosen entries. It has its origins in nuclear physics, where it was used to model the energy levels of complex atomic nuclei, but has since found applications in many fields including condensed matter physics, number theory, statistics, and machine learning [2, 7]. In RMT, one often studies the properties of the matrices as they become large, referred to as the thermodynamic limit. This is because there are many universal properties that emerge in this limit, and if the size of the matrix is small, it might be more appropriate to study the specific matrix rather than its statistical properties.

A.1 Observables

When dealing with an $N \times N$ random matrix A_{ij} , it might be tempting to ask questions about the distribution of moments like $\langle \sum_{k=1}^N A_{ik} A_{kj} \rangle$ or $\langle \sum_{k,l=1}^N A_{ik} A_{kl} A_{lj} \rangle$, but this would result in answers that have size $O(N^2), O(N^3)$, etc. and are not very informative because they can fluctuate a lot from sample to sample. The analogous quantity to the moments of a random variable is the normalized trace:

$$\frac{1}{N} \text{Tr } A^k \quad (86)$$

For a well-behaved function f , one can also compute the normalized trace of $f(A)$ which is the average of f applied to the eigenvalues of A :

$$\frac{1}{N} \text{Tr } f(A) = \frac{1}{N} \sum_{i=1}^N f(\lambda_i). \quad (87)$$

Both of these quantities are self-averaging in the thermodynamic limit for many common random matrix ensembles, meaning that they converge to a deterministic value as $N \rightarrow \infty$. When $N \rightarrow \infty$, the eigenvalues λ_i of A are often treated as random samples from a deterministic distribution $\rho(\lambda)$ called the spectral density, which is defined as:

$$\rho(\lambda) = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N \delta(\lambda - \lambda_i), \quad (88)$$

where $\delta(x)$ is the Dirac delta function. This is interpreted as the empirical distribution of the eigenvalues of A . As $N \rightarrow \infty$, $\rho(\lambda)$ often converges to a continuous function, and the eigenvalues become dense in a certain interval. In the case where this is well-defined, we can write:

$$\frac{1}{N} \text{Tr } f(A) \rightarrow \int d\lambda \rho(\lambda) f(\lambda), \quad (89)$$

which reveals that the normalized trace is the expectation of $f(\lambda)$ when λ is drawn from the spectral density $\rho(\lambda)$.

A.2 Wigner ensemble and the scaling of entries

A common random matrix ensemble is the Wigner ensemble, which consists of $N \times N$ symmetric matrices A with entries A_{ij} that are i.i.d. normal random variables with mean zero and variance σ^2/N for $i < j$, and $A_{ii} = 0$. The Wigner ensemble is a prototypical example of a random matrix ensemble that exhibits universal properties in the thermodynamic limit and is analogous to the normal distribution for random variables because there are analogous theorems to the central limit theorem and law of large numbers.

The variance of the entries of the Wigner ensemble is scaled as $\sim 1/N$ because this ensures that the eigenvalues of A remain $O(1)$ as $N \rightarrow \infty$. To see this, consider the variance of the eigenvalue distribution which is given by the normalized trace of A^2 :

$$\text{var}[\lambda] \rightarrow \left\langle \frac{1}{N} \text{Tr} A^2 \right\rangle = \frac{1}{N} \sum_{i,j=1}^N \langle A_{ij} A_{ji} \rangle = \frac{2}{N} \sum_{i,j(i \neq j)} \text{var}[A_{ij}] \quad (90)$$

Unless $\text{var}[A_{ij}] \sim 1/N$, this quantity would either diverge or vanish as $N \rightarrow \infty$. The fact that with this scaling, quantities like $\frac{1}{N} \text{Tr} f(A)$ converge to $O(1)$ deterministic values as $N \rightarrow \infty$ plays an important role in the derivations presented in the main text.

The spectral density of the Wigner ensemble converges to the Wigner semicircle law in the thermodynamic limit:

$$\rho(\lambda) = \begin{cases} \frac{1}{2\pi\sigma^2} \sqrt{4\sigma^2 - \lambda^2}, & |\lambda| < 2\sigma, \\ 0, & |\lambda| \geq 2\sigma, \end{cases} \quad (91)$$

The spectral densities of random matrices often have compact support like this, meaning that there is a maximum and minimum eigenvalue in the thermodynamic limit. The exceptions to this are ensembles of matrices that become singular in the thermodynamic limit because this corresponds to having very large eigenvalues. In this appendix, we will derive the Wigner semicircle law, but it is not necessary for understanding the main text.

A.3 Ginibre ensemble

The Ginibre ensemble is another common random matrix ensemble which consists of $N \times N$ matrices Q with entries Q_{ij} that are i.i.d. normal random variables with mean zero and variance σ^2/N . The Ginibre ensemble is not symmetric, so its eigenvalues are generally complex numbers. The spectral density of the Ginibre ensemble converges to the circular law in the thermodynamic limit:

$$\rho(x + iy) = \begin{cases} \frac{1}{\pi\sigma^2}, & |x + iy| < \sigma, \\ 0, & |x + iy| \geq \sigma. \end{cases} \quad (92)$$

A.4 Ginibre matrices as random rotations

Recall that a matrix Q which has $Q^T Q = Q Q^T = I$ is called orthogonal, and that such matrices represent rotations and reflections in




Figure 7 – The Wigner semicircle law for the spectral density of a symmetric random matrix with i.i.d. normal entries with mean zero and variance σ^2/N .

N -dimensional space. If you have a vector $\mathbf{x} \in \mathbb{R}^N$ and you multiply it by an orthogonal matrix Q , the resulting vector $Q\mathbf{x}$ has the same length as $\mathbf{x}$. If you have some matrix M and form QMQ^T , this corresponds to rotating the basis in which M is expressed. A Ginibre matrix Q is not orthogonal, but it is approximately orthogonal in the thermodynamic limit in the sense that:

$$[QQ^T]_{ij} = \sum_{k=1}^N Q_{ik}Q_{jk} \approx \frac{\sigma^2}{N}\delta_{ij}, \quad (93)$$

where δ_{ij} is the Kronecker delta (1 if $i = j$ and 0 otherwise). This means that multiplying by a Ginibre matrix approximately preserves lengths and angles in high dimensions, so Ginibre matrices can be thought of as random rotations.

This is useful when diagonalizing a random symmetric matrix M . If M is a symmetric matrix, it can be diagonalized as $M = O\Lambda O^T$ where O is an orthogonal matrix and Λ is a diagonal matrix of eigenvalues. In many cases, the columns of O (which are the eigenvectors of M) are called ‘delocalized’ because the entries of each eigenvector are approximately independent normal random variables with variance $1/N$. It is often a good approximation to treat O as a Ginibre matrix in the thermodynamic limit, which is useful for computing averages over eigenvectors of random matrices. This means that one can think of many random symmetric matrices as being a random rotation of a diagonal matrix whose entries are sampled from the spectral density.

A.5 Wick’s theorem

A very useful tool for computing averages over Gaussian random variables is Wick’s theorem, which states that the average of a product of an even number of zero-mean Gaussian random variables can be expressed as a sum over all possible pairings of the variables, where you take the product of the averages of each pair. For example, if x_1, x_2, x_3, x_4 are zero-mean Gaussian random variables, then:

$$\langle x_1 x_2 x_3 x_4 \rangle = \langle x_1 x_2 \rangle \langle x_3 x_4 \rangle + \langle x_1 x_3 \rangle \langle x_2 x_4 \rangle + \langle x_1 x_4 \rangle \langle x_2 x_3 \rangle. \quad (94)$$

If there is an odd number of variables, the average is zero. In the math and statistics literature, this is often called Isserlis’ theorem.

A.6 The Green’s function as the moment generating function

A very useful tool for computing the spectral density of a random matrix is the Green’s function, which is defined as:

$$G(z) \equiv \lim_{N \rightarrow \infty} \frac{1}{N} \text{Tr} (zI - A)^{-1} = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N \frac{1}{z - \lambda_i}, \quad (95)$$

with z a complex number not in the spectrum of A . Let’s see what happens if we expand this in powers of $1/z$:

$$G(z) = \frac{1}{N} \sum_{i=1}^N \left(\frac{1}{z} + \frac{\lambda_i}{z^2} + \frac{\lambda_i^2}{z^3} + \dots \right) \quad (96)$$

This means that we can compute:

$$\frac{1}{N} \text{Tr} A^k = \left. \frac{d^{k+1}}{dw^{k+1}} G(w^{-1}) \right|_{w=0} \quad (97)$$

Thus, $G(z)$ is the moment generating function for the eigenvalue distribution of A . If we know $G(z)$, we can recover the spectral density using the Stieltjes inversion formula. To do this, first observe:

$$\begin{aligned} \lim_{\varepsilon \rightarrow 0^+} \frac{1}{x \pm i\varepsilon} &= \mp i\pi \lim_{\varepsilon \rightarrow 0^+} \frac{\varepsilon}{\pi(x^2 + \varepsilon^2)} + \lim_{\varepsilon \rightarrow 0^+} \frac{x}{x^2 + \varepsilon^2} \\ &= \mp i\pi \delta(x) + \text{P.V.} \frac{1}{x}, \\ \implies \frac{1}{\pi} \lim_{\varepsilon \rightarrow 0^+} \text{Im} \frac{1}{x - i\varepsilon} &= \delta(x), \end{aligned} \quad (98)$$

where P.V. denotes the Cauchy principal value and $\delta(x)$ is the Dirac delta function. We can use this to extract the spectral density from $G(z)$:

$$\begin{aligned} \rho(\lambda) &= \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N \delta(\lambda - \lambda_i) = \frac{1}{\pi} \lim_{N \rightarrow \infty} \lim_{\varepsilon \rightarrow 0^+} \text{Im} \frac{1}{N} \sum_{i=1}^N \frac{1}{\lambda - \lambda_i - i\varepsilon} \\ &= \frac{1}{\pi} \lim_{\varepsilon \rightarrow 0^+} \text{Im} G(\lambda - i\varepsilon). \end{aligned} \quad (99)$$

Thus, if we can compute $G(z)$ for complex z , we can recover the spectral density $\rho(\lambda)$.

A.7 Derivation of the Wigner semicircle law

Here, we'll present one way to derive the Wigner semicircle law using the Green's function. There are many other ways to get this result that work for more general ensembles of random matrices, but this is one of the simplest.

Consider a Wigner matrix A with entries A_{ij} that are i.i.d. normal random variables with mean zero and variance σ^2/N for $i < j$, and $A_{ii} = 0$. Let $R(z) = (zI - A)^{-1}$ be the resolvent of A . If we start with $(zI - A)R(z) = I$ and take the normalized trace of both sides, we get:

$$zG(z) = 1 + \frac{1}{N} \text{Tr}[AR(z)]. \quad (100)$$

Let's turn to the second term on the right-hand side. We can invoke an identity called Stein's lemma: if Z is a zero-mean normal random variable with variance $\text{var}[Z]$, then for any well-behaved function f , we have:

$$\langle Zf(Z) \rangle = \text{var}[Z] \langle f'(Z) \rangle. \quad (101)$$

This can be derived by integration by parts. We can use this to compute:

$$\left\langle \frac{1}{N} \sum_{i,j=1}^N A_{ij} R_{ji}(z) \right\rangle = \frac{1}{N} \sum_{i,j=1}^N \text{var}[A_{ij}] \left\langle \frac{\partial R_{ji}(z)}{\partial A_{ij}} \right\rangle \quad (102)$$

Returning to $(zI - A)R(z) = I$, differentiating with respect to A_{ij} , using the product rule, and solving, we get:

$$\frac{\partial}{\partial A_{ij}} \left[\frac{1}{zI - A} \right]_{kl} = \left[\frac{1}{zI - A} \right]_{ki} \left[\frac{1}{zI - A} \right]_{jl} \quad (103)$$

Thus:

$$\begin{aligned} \frac{1}{N} \text{Tr}[AR(z)] &= \frac{\sigma^2}{N^2} \sum_{i,j=1}^N R_{ji}(z) R_{ji}(z) \\ &= \frac{\sigma^2}{N^2} \sum_{i,j(i \neq j)} R_{ji}(z) R_{ji}(z) + O\left(\frac{1}{N}\right) \\ &\approx \sigma^2 \left(\frac{1}{N} \text{Tr} R(z) \right)^2 \end{aligned} \quad (104)$$

where we have dropped the $O(1/N)$ terms in the thermodynamic limit. This gives us the equation:²²

$$zG(z) = 1 + \sigma^2 G(z)^2. \quad (105)$$

Solving this quadratic equation for $G(z)$, we get:

$$G(z) = \frac{z - \sqrt{z^2 - 4\sigma^2}}{2\sigma^2}, \quad (106)$$

where we choose the branch of the square root that ensures $G(z) \sim 1/z$ as $z \rightarrow \infty$. Finally, we can use the Stieltjes inversion formula to recover the spectral density:

$$\rho(\lambda) = \frac{1}{\pi} \lim_{\varepsilon \rightarrow 0^+} \text{Im} G(\lambda - i\varepsilon) = \begin{cases} \frac{1}{2\pi\sigma^2} \sqrt{4\sigma^2 - \lambda^2}, & |\lambda| < 2\sigma, \\ 0, & |\lambda| \geq 2\sigma. \end{cases} \quad (107)$$

This is the Wigner semicircle law.

A.8 The skew-symmetric case

In the main text, we deal with a random coupling matrix A that is skew-symmetric, meaning that $A^T = -A$. This means that the eigenvalues of A are purely imaginary. The matrix iA is Hermitian and has eigenvalues distributed according to the Wigner semicircle law with variance σ^2 . If μ is an eigenvalue of iA , then $\lambda = -i\mu$ is an eigenvalue of A . Thus, the spectral density of A is just the Wigner semicircle law rotated in the complex plane onto the imaginary axis:

$$\rho(x + iy) = \delta(x) \begin{cases} \frac{1}{2\pi\sigma^2} \sqrt{4\sigma^2 - y^2}, & |y| \leq 2\sigma, \\ 0, & |y| > 2\sigma, \end{cases} \quad (108)$$

where the Dirac delta $\delta(x)$ indicates that all eigenvalues lie on the imaginary axis.

B Appendix: Brief review of stationary Gaussian processes

A Gaussian process is a collection of random variables which are indexed by a continuous parameter (usually time or space) such that any finite

²² It is no coincidence that $\tilde{G}(z)$, the Laplace transform of the Green's function for the DMFT equation of motion, follows a very similar equation to $G(z)$ here. In particular (see Eq. (48) in the main text):

$$z\tilde{G}(z) = 1 - \sigma^2 \tilde{G}(z)^2.$$

There is a sign difference because the coupling matrix A_{ij} in the main text is skew-symmetric whereas the Wigner matrix here is symmetric. There is a very deep connection between the Green's functions of linear systems with random couplings and the spectra of random matrices. In fact, one can exploit this connection to quickly compute the spectral density of many random matrix ensembles using the cavity method. To learn more about this, see Ref. [4].

subset of these random variables has a multivariate normal distribution. Gaussian processes are useful for modeling functions and time series because they can represent a wide range of behaviors while still being mathematically tractable. For example, Gaussian processes can be used to model biological processes like neural activity or population dynamics. In this appendix, we'll provide a brief review of Gaussian processes and some special properties of those whose statistical properties are invariant under time translation, called stationary Gaussian processes.

B.1 Definition

Let's first consider some random variables $X_1, \dots, X_n$ that have a multivariate normal distribution with mean vector μ_i and covariance matrix Σ_{ij} . This means that their joint probability density function is given by:

$$p(X_1, \dots, X_n) \propto \exp \left(-\frac{1}{2} \sum_{i,j=1}^n (X_i - \mu_i) \Sigma_{ij}^{-1} (X_j - \mu_j) \right), \quad (109)$$

where Σ_{ij}^{-1} is the inverse of the covariance matrix. Now, let's consider $i = 1, 2, \dots, n$ indexing time points $t_1, t_2, \dots, t_n$ which are linearly spaced with spacing Δt . This means we could interpret X_i as the value of some random function $X(t)$ at time t_i . We might then ask what happens as we take the limit $\Delta t \rightarrow 0$ and $n \rightarrow \infty$ while keeping $n\Delta t = T$ fixed. In this limit, we can think of $X(t)$ as a random function defined for all $t \in [0, T]$. Eq. (109) then becomes:

$$p(\{X(t)\}) \propto \exp \left(-\frac{1}{2} \int_0^T dt \int_0^T ds (X(t) - \mu(t)) C^{-1}(t, s) (X(s) - \mu(s)) \right) \quad (110)$$

where $C(t, s)$ is the covariance function defined as:

$$C(t, s) = \langle (X(t) - \mu(t))(X(s) - \mu(s)) \rangle, \quad (111)$$

and $C^{-1}(t, s)$ is the unique function such that:

$$\int_0^T dt \int_0^T ds C^{-1}(t, s) C(s, t') = \delta(t - t'), \quad (112)$$

where $\delta(t - t')$ is the Dirac delta function. When $p(X_1, \dots, X_n)$ becomes $p(\{X(t)\})$ we call it a probability density functional (meaning that it takes a function as an input and outputs a probability density), and we say that $X(t)$ is a Gaussian process.

B.2 Stationary Gaussian processes

A stationary Gaussian process is a Gaussian process whose mean and covariance are invariant under time translation. This means that the mean function $\mu(t)$ is constant and the covariance function $C(t, s)$ depends only on the time difference $t - s$:

$$\mu(t) = \mu, \quad C(t, s) = C(t - s). \quad (113)$$

These sorts of processes appear in systems that are in equilibrium or steady state, where the statistical properties do not change over time. In the case where $\lim_{t \rightarrow \infty} C(t) = 0$, the process is not correlated with itself at long times. When this is the case, averages over long time intervals can be used to estimate ensemble averages.

B.3 Power spectral density and Wiener–Khinchin theorem

A useful tool for analyzing stationary Gaussian processes is the power spectral density (PSD), which describes how the variance of the process is distributed over different frequencies. The PSD of a function $X(t)$ is defined as the average squared magnitude of its Fourier transform:

$$S_X(\omega) = \left\langle \lim_{T \rightarrow \infty} \frac{1}{T} \left| \int_0^T dt X(t) e^{-i\omega t} \right|^2 \right\rangle, \quad (114)$$

where the expectation is over realizations of the process. The Wiener–Khinchin theorem states that the PSD is the Fourier transform of the auto-covariance function:

$$\begin{aligned} S_X(\omega) &= \lim_{T \rightarrow \infty} \frac{1}{T} \left| \int_0^T dt X(t) e^{-i\omega t} \right|^2 \\ &= \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T \int_0^T dt ds C(t-s) e^{-i\omega(t-s)} \\ &= \int_{-\infty}^{\infty} C(\tau) e^{-i\omega\tau} d\tau. \end{aligned} \quad (115)$$

References

- [1] Nishimori, H. *Statistical Physics of Spin Glasses and Information Processing: An Introduction* (Oxford University Press, 2001). URL <https://doi.org/10.1093/acprof:oso/9780198509417.001.0001>. <https://academic.oup.com/book/5185/book-pdf/54038185/acprof-9780198509400.pdf>.
- [2] Cui, W., III, R. M. & Mehta, P. Les houches lectures on community ecology: From niche theory to statistical mechanics (2024). URL <https://arxiv.org/abs/2403.05497>. 2403.05497.
- [3] Roy, F., Biroli, G., Bunin, G. & Cammarota, C. Numerical implementation of dynamical mean field theory for disordered systems: application to the lotka–volterra model of ecosystems. *Journal of Physics A: Mathematical and Theoretical* **52**, 484001 (2019). URL <https://dx.doi.org/10.1088/1751-8121/ab1f32>.
- [4] Cui, W., Rocks, J. W. & Mehta, P. An elementary mean-field approach to the spectral densities of random matrix ensembles. *Physica A: Statistical Mechanics and its Applications* **637**, 129608 (2024). URL <https://www.sciencedirect.com/science/article/pii/S037843712400116X>.
- [5] Bunin, G. Ecological communities with lotka-volterra dynamics. *Phys. Rev. E* **95**, 042414 (2017). URL <https://link.aps.org/doi/10.1103/PhysRevE.95.042414>.
- [6] Cui, W., Marsland, R. & Mehta, P. Effect of resource dynamics on species packing in diverse ecosystems. *Phys. Rev. Lett.* **125**, 048101 (2020). URL <https://link.aps.org/doi/10.1103/PhysRevLett.125.048101>.
- [7] Potters, M. & Bouchaud, J. *A First Course in Random Matrix Theory: for Physicists, Engineers and Data Scientists* (Cambridge University Press, 2020). URL <https://books.google.com/books?id=9K-WzQEACAAJ>.
- [8] Zou, W. & Huang, H. Introduction to dynamical mean-field theory of randomly connected neural networks with bidirectionally correlated couplings. *SciPost Phys. Lect. Notes* **79** (2024). URL <https://scipost.org/10.21468/SciPostPhysLectNotes.79>.

- [9] Crisanti, A. & Sompolinsky, H. Path integral approach to random neural networks. *Phys. Rev. E* **98**, 062120 (2018). URL <https://link.aps.org/doi/10.1103/PhysRevE.98.062120>.
- [10] Sompolinsky, H., Crisanti, A. & Sommers, H. J. Chaos in random neural networks. *Phys. Rev. Lett.* **61**, 259–262 (1988). URL <https://link.aps.org/doi/10.1103/PhysRevLett.61.259>.
- [11] Kadmon, J. & Sompolinsky, H. Transition to chaos in random neuronal networks. *Phys. Rev. X* **5**, 041030 (2015). URL <https://link.aps.org/doi/10.1103/PhysRevX.5.041030>.
- [12] Schuecker, J., Goedeke, S. & Helias, M. Optimal sequence memory in driven random networks. *Phys. Rev. X* **8**, 041029 (2018). URL <https://link.aps.org/doi/10.1103/PhysRevX.8.041029>.
- [13] Arnoulx de Pirey, T. & Bunin, G. Many-species ecological fluctuations as a jump process from the brink of extinction. *Phys. Rev. X* **14**, 011037 (2024). URL <https://link.aps.org/doi/10.1103/PhysRevX.14.011037>.
- [14] Novikov, E. A. Functionals and the random-force method in turbulence theory. *Sov. Phys. JETP* **20**, 1290–1294 (1965).