




Statistical mechanics of inference in epidemic spreading

Alfredo Braunstein ^{1,2,3}, Louise Budzynski ^{1,3,4} and Matteo Mariani ^{1,*}

¹*DISAT, Politecnico di Torino, Corso Duca Degli Abruzzi 24, 10129 Torino*

²*Collegio Carlo Alberto, P.zza Arbarello 8, 10122 Torino, Italy*

³*Italian Institute for Genomic Medicine, IRCCS Candiolo, SP-142, I-10060 Candiolo (TO), Italy*

⁴*Dipartimento di Fisica, Università “La Sapienza”, P.le A. Moro 5, 00185 Rome, Italy*



(Received 25 July 2023; accepted 21 November 2023; published 15 December 2023)

We investigate the information-theoretical limits of inference tasks in epidemic spreading on graphs in the thermodynamic limit. The typical inference tasks consist in computing observables of the posterior distribution of the epidemic model given observations taken from a ground-truth (sometimes called planted) random trajectory. We can identify two main sources of quenched disorder: the graph ensemble and the planted trajectory. The epidemic dynamics however induces nontrivial long-range correlations among individuals' states on the latter. This results in nonlocal correlated quenched disorder which unfortunately is typically hard to handle. To overcome this difficulty, we divide the dynamical process into two sets of variables: a set of stochastic independent variables (representing transmission delays), plus a set of correlated variables (the infection times) that depend deterministically on the first. Treating the former as quenched variables and the latter as dynamic ones, computing disorder average becomes feasible by means of the replica-symmetric cavity method. We give theoretical predictions on the posterior probability distribution of the trajectory of each individual, conditioned to observations on the state of individuals at given times, focusing on the susceptible infectious (SI) model. In the Bayes-optimal condition, i.e., when true dynamic parameters are known, the inference task is expected to fall in the replica-symmetric regime. We indeed provide predictions for the information theoretic limits of various inference tasks, in form of phase diagrams. We also identify a region, in the Bayes-optimal setting, with strong hints of replica-symmetry breaking. When true parameters are unknown, we show how a maximum-likelihood procedure is able to recover them with mostly unaffected performance.

DOI: [10.1103/PhysRevE.108.064302](https://doi.org/10.1103/PhysRevE.108.064302)

I. INTRODUCTION

Reconstructing information on epidemic spreading is crucial to develop advanced digital contact tracing strategies to mitigate the spreading of an epidemic. Based on partial information on the states of individuals at given times, the problem consists in reconstructing the posterior distribution on unobserved events, such as the initial state of the epidemic (the source), or undetected infected individuals. These inverse problems are known to be challenging, even for simple dynamics such as the susceptible infectious (SI) model. Several methods have been proposed to tackle inference problems in epidemics, including Monte Carlo [1–3], heuristic [4], belief propagation [5–8], mean-field [6], variational [9,10], and other [11,12] approaches. Although many of these methods have shown through extensive simulations to reconstruct efficiently some information on the posterior probability distribution in specific graphs sizes and ensembles, a study of the feasibility of inference in epidemic models is still generally lacking. A notable exception is given by preprint [8] (which appeared while we were finishing the present work). Its main aim is to provide a quantitative study of the feasibility of inference in epidemic spreading on random graphs, in the

form of phase diagrams, by means of extensive simulations on finite-size systems. The work focuses on the Bayes optimal setting, and uncovers interesting hints of failure of optimality, that are attributed to finite-size effects. In this work, we focus on the large size (thermodynamic) limit and use the replica-symmetric cavity method. Outside the Bayes optimal regime, we study the performances achieved when hyperparameters are inferred. We provide a theoretical analysis of inference tasks aiming at reconstructing individuals' trajectories from the partial knowledge of the state of a fraction of individuals at a given observation time *in the thermodynamic limit*, which we also show to be in good agreement with results on moderately large random graphs. To the best of our knowledge, this work is the first to study the information bounds of the risk assessment problem in the thermodynamic limit. Our method is semianalytical (we resort to numerics only to find fixed points) and it is based on the cavity method, which in the replica-symmetric regime (and for the graph ensembles studied) leads to the exact solution (up to numerical accuracy). We provide quantitative predictions on the information contained on the posterior probability given the observations, varying the characteristics of the epidemic, of the contact network, and of the observations. Our approach relies on a study of the properties of the posterior probability measure, for typical contact graphs and realization of the epidemic spreading, using the replica-symmetric (RS) cavity method. We focus

*matteo.mariani@polito.it

in this paper on the simple SI model [13], but the strategy is general and can be applied to other irreversible spreading process such as the SIR or SEIR models, see Sec. IID and Appendix C.

To perform this analysis, we need to compute averages of the inference task over realizations of a *planted* epidemic trajectory (the ground truth), from which observations are taken. These observations have thus to be treated as quenched disordered variables (along with the variables needed to describe the contact graph). However, their distribution is nonlocally correlated: the past history of the epidemic spreading induces long-range correlations between the state of individuals at the observation time. While the cavity method is well-suited for models in which variables used to describe the disorder are independent, applying it on a model with long-ranged correlated disorder is instead nontrivial. To circumvent this difficulty, we devised the following strategy. We separate the *planted* dynamical process in two sets of variables: (a) the transmission delays, which are independent, and (b) *infection times* and *observations*, which are a deterministic function of other infection times and of transmission delays through a set of local hard constraints. We treat the first set as quenched disorder, while the second set, together with the variables used to describe the inferred trajectory, are treated as dynamical variables. Although *planted* infection times are not truly dynamical variables, their deterministic dependence on the disorder allows us to consider them as such without modifying the probability distribution. This strategy thus effectively transfers correlations out of the quenched variables and into the dynamical ones, allowing a straightforward application of the cavity method.

The paper is organized as follows. In Sec. II we set up the problem, and present our strategy to adapt the RS formalism to inference in epidemic spreading. Results are presented in Sec. III. We start our analysis in the Bayes-optimal case (Sec. III A), where RS is expected to hold. We provide quantitative estimates of the feasibility of inference, including Bayes estimators, and the area under the receiver operating characteristic (ROC) curve (AUC). These RS predictions are in good agreement with the result of message-passing algorithm on large instances. We identify a region in the Bayes-optimal setting where belief-propagation algorithm fails to converge, both on finite-size instances (as already observed in Ref. [8]) and in the thermodynamic (large-size) limit. This observation is a strong hint for replica-symmetry breaking (RSB), and is also confirmed by a failure of Monte Carlo algorithm performing the inference task in this regime. This result is surprising, as it is often argued that being on the Nishimori line guarantees the absence of replica-symmetry breaking [14]. However, and although Nishimori's identities are always satisfied in the Bayes-optimal setting, our observations can be explained by the fact that the overlap between planted and inferred trajectories is not necessarily self-averaging in this problem (that is not gauge invariant). In Sec. III B, we explore regimes outside Bayes-optimal conditions. We identify a regime in which neither Belief Propagation nor the iterative numerical resolution of the RS cavity equations converge. This suggests the presence of an RSB transition. When the parameters of the model are unknown, one can rely on strategies such as expectation-maximization to infer them. These strategies are

equivalent to imposing the Nishimori conditions. We provide a quantitative study of an iterative strategy to infer the parameters of the prior in the thermodynamical limit. We show that, for a large range of the prior's parameters, it is possible to recover similar accuracy than the one of the Bayes optimal case, even when starting from initial conditions that are far from the prior's parameters. These results are in good agreement with simulations in finite systems.

II. ENSEMBLE STUDY FOR INFERENCE IN EPIDEMIC SPREADING

A. Epidemic inference

1. SI model on graphs

We consider the SI model of spreading, defined over a graph $\mathcal{G} = (V, E)$. At time t a node $i \in V$ can be in two states represented by a variable $x_i^t \in \{S, I\}$. At each time step, an infected node can independently infect each of its susceptible neighbors ∂i with probabilities $\lambda_{ij} \in [0, 1]$:

$$P(\underline{x}) = \prod_{i=1}^N \left[p(x_i^0) \prod_{t=0}^{T-1} p(x_i^{t+1} | x_i^t, x_{\partial i}^t) \right], \quad (1)$$

where $\underline{x} = \{x_i^t\}$ for $i = 1, \dots, N, t = 0, \dots, T$, and

$$p(x_i^{t+1} = I | x_i^t, x_{\partial i}^t) = 1 - \delta_{x_i^t, S} \prod_{j \in \partial i} (1 - \lambda_{ji} \delta_{x_j^t, I}). \quad (2)$$

The dynamics (2) is irreversible: A given node can only undergo the transition $S \rightarrow I$. Therefore, the trajectory in time of an individual can be parameterized by its infection time t_i . We assume that a subset of the nodes initiate with an infection time $t_i = 0$, i.e., $x_i^0 = I$. A realization of the SI process can be univocally expressed in terms of the independent transmission delays $s_{ij} \in \{1, 2, \dots, \infty\}$, following a geometrical distribution $w_{ij}(s) = \lambda_{ij}(1 - \lambda_{ij})^{s-1}$. Once the initial condition $\{x_i^0\}_{i \in V}$ and the set of transmission delays $\{s_{ij}, s_{ji}\}_{(ij) \in E}$ is fixed, the infection times can be uniquely determined from the set of equations:

$$t_i = \delta_{x_i^0, S} \min_{j \in \partial i} \{t_j + s_{ji}\}. \quad (3)$$

We assume that each individual has a probability γ to be infected at time $t = 0$, and we assume for simplicity that the transmission probabilities are site-independent: $\lambda_{ij} = \lambda$ for all $(ij) \in E$. The distribution of infection times conditioned on the realization of delays and on the initial condition can be written

$$P(\underline{t} | \{x_i^0\}, \{s_{ij}, s_{ji}\}) = \prod_{i \in V} \psi^*(t_i, \underline{t}_{\partial i}, x_i^0, \{s_{ji}\}_{j \in \partial i}), \quad (4)$$

with $\underline{t}_{\partial i} = \{t_j, j \in \partial i\}$, and where ψ^* enforces the above constraint on the infection times:

$$\psi^* = \mathbb{I}[t_i = \delta_{x_i^0, S} \min_{j \in \partial i} \{t_j + s_{ji}\}], \quad (5)$$

with $\mathbb{I}[A]$ the indicator function of the event A . Once averaged over the transmission delays and over the initial condition, we obtain the following distribution of times:

$$P(\underline{t}) = \prod_{i \in V} \psi(t_i, \underline{t}_{\partial i}), \quad (6)$$

where

$$\psi = \sum_{x_i^0} \gamma(x_i^0) \sum_{\{s_{ji}\}_{j \in \partial i}} \psi^*(t_i, \underline{t}_{\partial i}, x_i^0, \{s_{ji}\}_{j \in \partial i}) \prod_{j \in \partial i} w(s_{ji}), \quad (7)$$

and $\gamma(x) = \gamma \mathbb{I}[x = I] + (1 - \gamma) \mathbb{I}[x = S]$.

2. Inferring individual's trajectories from partial observations

In the inference problem, we assume that some information $\mathcal{O} = \{O_m\}_{m=1, \dots, M}$ is given on the trajectory by the result of M independent medical tests. The probability of observations $P(\mathcal{O}|\underline{t})$ factorizes over the set of tests:

$$P(\mathcal{O}|\underline{t}) = \prod_{m=1}^M \rho(O_m|t_{i_m}). \quad (8)$$

Each test gives information about the state $\sigma_m \in \{I, S\}$ of an individual i_m at a given time $\theta_m \in \{0, 1, \dots, T\}$ with a false rate $fr \in [0, 1]$. We can thus write each observation as $O_m = (i_m, \theta_m, \sigma_m, fr)$. The probability of observation conditioned to the infection time is therefore

$$\rho(O_m|t_{i_m}) = (1 - fr) \mathbb{I}[x_{i_m}^{\theta_m} = \sigma_m] + fr \mathbb{I}[x_{i_m}^{\theta_m} \neq \sigma_m]. \quad (9)$$

Here for simplicity we assume $fr_m \equiv fr$ constant which implies identical false-positive and negative rates. Using Bayes rule, the posterior probability of infection times is

$$P(\underline{t}|\mathcal{O}) = \frac{P(\underline{t})P(\mathcal{O}|\underline{t})}{P(\mathcal{O})}, \quad (10)$$

with $P(\underline{t})$ given in Eq. (6).

3. Bayes optimal setting

In the Bayes optimal setting, the parameters (λ, γ, fr) of the epidemic spreading process are known in the inference task. This means that the parameters (λ, γ, fr) used in the posterior probability (10) are the same than the true parameters used to generate the observations. However in many cases, values of the parameters are unknown, and need to be inferred. In such a case, we denote by $(\lambda^*, \gamma^*, fr^*)$ (respectively, $\lambda^l, \gamma^l, fr^l$) the parameters used to generate the observations (respectively, to infer the infection times).

B. Ensemble average

Our objective is to estimate how well observables on the true (or planted) infection trajectory $\underline{\tau}$ are approximated by those of the inferred trajectory \underline{t} , which follows the posterior distribution given the observations \mathcal{O} . We shall characterize the properties of the posterior distribution (10) on a random ensemble of contact graphs and realization of the epidemic spreading. An instance will be defined by a contact graph \mathcal{G} , a ground-truth trajectory $\underline{\tau}$ and a set of observation \mathcal{O} sampled from the distribution $P(\mathcal{O}|\underline{\tau})$. Three graph ensembles are considered: random regular (RR), Erdős-Rényi (ER) (defined for example in Ref. [15]), and graph ensemble with a truncated fat-tailed (FT) degree distribution. We will be interested in the large size limit $n \rightarrow \infty$, with n the number of individuals, at fixed degree distribution. In this limit, graph instances of the above-mentioned ensembles are locally treelike, allowing us to exploit the cavity method to determine the typical properties of the measure (10).

1. Correlated observations

A technical difficulty arises when one tries to apply the cavity method directly to the posterior distribution (10). This distribution is defined for a given realization of the observations $\mathcal{O} = \{O_m\}_{m=1, \dots, M}$. Observations \mathcal{O} have to be treated as quenched (disorder). While the cavity method is well-suited for local, independent random disorder, the past history of the epidemic spreading has introduced nontrivial long-range correlations between the observations $\{O_m\}$. To overcome this difficulty, we rely on the set of hard constraints (5) on the planted times that we recall here: $\tau_i = \delta_{x_i^0, S} \min_{j \in \partial i} \{\tau_j + s_{ji}\}$. These constraints are expressed in terms of independent random variables: the local delays s_{ij} and the initial-time state x_i^0 . In fact, knowing these two sets of variables, it is possible to determine each (planted) infection time: for fixed delays $\{s_{ij}, s_{ji}\}_{(ij) \in E}$ and seeds $\{x_i^0\}_{i \in V}$, the planted times are fixed to be the unique solution of the set of constraints (5). Our strategy is therefore to treat these local variables as disorder (quenched) variables, and consider instead the planted time as constrained dynamical (annealed) variable. To treat the noise in the observations, we also define the set of error bits $\{\varepsilon_m\}_{m=1, \dots, M}$, with $\varepsilon_m = \mathbb{I}[\sigma_m \neq x_{i_m}^{\theta_m}]$. We denote by $\mathcal{D} = \{\{x_i^0\}_{i \in V}, \{\varepsilon_m\}_{m=1}^M, \{s_{ij}, s_{ji}\}_{(ij) \in E}\}$ the set of all disordered variables. At fixed disorder, the set of planted times $\underline{\tau}$ and observations \mathcal{O} is uniquely determined. Averaging over the disordered variables is therefore equivalent to average over the set of observations: with this strategy we can perform the quenched average over correlated observations. Obviously, the price to pay with this approach is to treat planted times as annealed variables. As a result, the BP messages are over a couple of times (τ_i, t_i) instead of a single time t_i , increasing the complexity in the resolution of the RS equation with population dynamics (see Appendix A for further details).

2. A graphical model for the joint distribution over planted and inferred trajectories

The joint probability of the planted times $\underline{\tau}$, of the observations $\mathcal{O} = \{O_m\}$ and of the inferred times \underline{t} conditioned on the disorder \mathcal{D} is defined as

$$P(\underline{t}, \mathcal{O}, \underline{\tau}|\mathcal{D}) = P(\underline{\tau}|\mathcal{D})P(\mathcal{O}|\mathcal{D}, \underline{\tau})P(\underline{t}|\mathcal{O}, \mathcal{D}, \underline{\tau}) \quad (11)$$

$$= P(\underline{\tau}|\mathcal{D})P(\mathcal{O}|\mathcal{D}, \underline{\tau})P(\underline{t}|\mathcal{O}) \quad (12)$$

$$= \frac{1}{P(\mathcal{O})} P(\underline{\tau}|\mathcal{D})P(\mathcal{O}|\mathcal{D}, \underline{\tau})P(\mathcal{O}|\underline{t})P(\underline{t}).$$

Notice that the passage from Eq. (11) to Eq. (12) is because the inferred trajectory \underline{t} is, conditionally to \mathcal{O} , independent from \mathcal{D} and $\underline{\tau}$. To interpret the factorization in this definition, we notice that the joint distribution above is the probability to have a planted trajectory $\underline{\tau}$, to observe a set \mathcal{O} from the planted and to infer a trajectory \underline{t} from the observations, all conditioned to the disorder \mathcal{D} . Therefore, this is the joint probability of:

(1) sampling the planted $\underline{\tau}$ from the prior $P(\underline{\tau}|\mathcal{D})$, conditioned on the disorder,

(2) sampling the observations \mathcal{O} at given planted from the likelihood $P(\mathcal{O}|\mathcal{D}, \underline{\tau})$,

(3) sampling the inferred trajectory \underline{t} from the posterior $P(\underline{t}|\mathcal{O})$, which can be rewritten by means of Bayes' law as $\frac{1}{P(\mathcal{O})}P(\underline{t})P(\mathcal{O}|\underline{t})$.

Let us now analyze term by term the factors in Eq. (12). The first term in the product is

$$P(\underline{\tau}|\mathcal{D}) = \prod_{i \in V} \psi^*(\tau_i, \underline{\tau}_{\partial i}; x_i^0, \{s_{ji}\}_{j \in \partial i}),$$

with ψ^* given in Eq. (5). The second term in the product is the probability of having observation $\mathcal{O} = \{O_m\}$ given the planted times $\underline{\tau}$ and the disorder \mathcal{D} :

$$P(\mathcal{O}|\mathcal{D}, \underline{\tau}) = \prod_{m=1}^M [(1 - \varepsilon_m) \delta_{\sigma_m, x_m^{\theta_m}} + \varepsilon_m (1 - \delta_{\sigma_m, x_m^{\theta_m}})].$$

The third and the fourth terms are respectively given in Eqs. (8) and (6). Finally, the denominator

$$P(\mathcal{O}) = \sum_{\underline{t}} P(\underline{t})P(\mathcal{O}|\underline{t})$$

can be seen as a complicated function of the observations \mathcal{O} , but since the observations are a deterministic function of the disorder, we will denote it as a function of the latter:

$$P(\mathcal{O}) = Z(\mathcal{D}).$$

The fact that the observations are deterministically fixed when fixed disorder, also implies that we can sum Eq. (12) over \mathcal{O} and since the disorder is fixed there is only one corresponding nonzero value of $\mathcal{O} = \mathcal{O}(\mathcal{D})$. As a consequence, we have that $P(\underline{\tau}, \underline{t}|\mathcal{D}) = P(\underline{t}, \mathcal{O}(\mathcal{D}), \underline{\tau}|\mathcal{D})$. The joint probability distribution of planted and inferred times, conditioned on the disorder, is thus obtained by substituting in Eq. (12) the functional form of each factor:

$$P(\underline{\tau}, \underline{t}|\mathcal{D}) = \frac{1}{Z(\mathcal{D})} \prod_{i \in V} \psi^*(\tau_i, \underline{\tau}_{\partial i}; x_i^0, \{s_{ji}\}_{j \in \partial i}) \times \psi(\tau_i, \underline{t}_{\partial i}) \xi(\tau_i, t_i; \{\varepsilon_m\}_{i_m=i}), \quad (13)$$

with

$$\xi(\tau_i, t_i; \{\varepsilon_m\}_{i_m=i}) = \prod_{m:i_m=i} \rho(O_m|t_m). \quad (14)$$

Note that when the error probability is zero: $fr = 0$ so a single observation reads $O_m = (i_m, \theta_m, \sigma_m, 0)$, the error variables are always $\varepsilon_m = 0$ (no corruption), and $\rho(O_m|t_m) = \mathbb{I}[\sigma_m = x_{i_m}^{\theta_m}]$. The coupling term $\xi(\tau_i, t_i; \{\varepsilon_m\}_{i_m=i})$ between inferred and planted times in the joint probability becomes

$$\xi(\tau_i, t_i) = \prod_{O_m:i_m=i} \mathbb{I}[x_{i_m}^{\theta_m} = \sigma_m].$$

C. Nishimori conditions and replica symmetry

There is a general argument that hints at the absence of replica-symmetry breaking in the Bayes optimal case [14, 16]. However, it is not clear that this property holds in general. The argument is a consequence of the Nishimori conditions, that we reformulate here in our setting. Consider a given realization of the observations \mathcal{O} , and two configurations $\underline{t}_1, \underline{t}_2$ sampled independently from the posterior distribution

$P_I(\underline{t}|\mathcal{O})$, where the subscript I refers to the set of hyperparameters $\lambda^I, \gamma^I, fr^I$ used in the inference process. Let $f(\underline{t}_1, \underline{t}_2)$ be an arbitrary function of two configurations and $\langle f(\underline{t}_1, \underline{t}_2) \rangle_{\underline{t}_1, \underline{t}_2} = \sum_{\underline{t}_1, \underline{t}_2} f(\underline{t}_1, \underline{t}_2) P_I(\underline{t}_1|\mathcal{O}) P_I(\underline{t}_2|\mathcal{O})$ its average over the posterior distribution. This average depends on the observation set \mathcal{O} . If we average, in the Bayes optimal conditions, $\langle f(\underline{t}_1, \underline{t}_2) \rangle_{\underline{t}_1, \underline{t}_2}$ over \mathcal{O} , then we get

$$\mathbb{E}_{\mathcal{O}}[\langle f(\underline{t}_1, \underline{t}_2) \rangle_{\underline{t}_1, \underline{t}_2}] = \mathbb{E}_{\tau, \mathcal{O}}[\langle f(\underline{\tau}, \underline{t}) \rangle_{\underline{t}}]. \quad (15)$$

This equation goes under the name of Nishimori condition [14] and it is crucial to discuss the replica symmetry of the posterior. Before going through the proof (which appears in Ref. [14]), we comment its meaning: on the l.h.s. we have that both the arguments of f are averaged over the posterior trajectories $(\underline{t}_1, \underline{t}_2)$, while on the right-hand side (r.h.s.) one argument ($\underline{\tau}$) is averaged over the prior and the other (\underline{t}) over the posterior. In other words, Eq. (15) states that on average (over \mathcal{O}) it is the same to average over the planted or over the inferred trajectories. In this sense the Nishimori condition means that under the Bayes optimal conditions, the planted is a *fair sample* of the posterior [14]. We now show the proof of Eq. (15) taken from Ref. [14] and then we discuss its connection with replica symmetry. The left-hand side (l.h.s.) is

$$\begin{aligned} \mathbb{E}_{\mathcal{O}}[\langle f(\underline{t}_1, \underline{t}_2) \rangle_{\underline{t}_1, \underline{t}_2}] &= \sum_{\mathcal{O}} P_*(\mathcal{O}) \langle f(\underline{t}_1, \underline{t}_2) \rangle_{\underline{t}_1, \underline{t}_2} \\ &= \sum_{\mathcal{O}} \sum_{\underline{t}_1, \underline{t}_2} P_*(\mathcal{O}) f(\underline{t}_1, \underline{t}_2) P_I(\underline{t}_1|\mathcal{O}) P_I(\underline{t}_2|\mathcal{O}) \\ &= \sum_{\underline{t}_1, \mathcal{O}, \underline{t}_2} P_*(\mathcal{O}) \frac{P_I(\underline{t}_1) P_I(\mathcal{O}|\underline{t}_1)}{P_I(\mathcal{O})} P_I(\underline{t}_2|\mathcal{O}) f(\underline{t}_1, \underline{t}_2), \quad (16) \end{aligned}$$

where in the third line we used the Bayes law for $P_I(\underline{t}_1|\mathcal{O})$, and where the subscript $*$ in the probability $P_*(\mathcal{O})$ refers to the set of planted hyperparameters $\lambda^*, \gamma^*, fr^*$ used in the generation of the observables. In the Bayes optimal setting, the two sets of hyperparameters coincide $(\lambda^*, \gamma^*, fr^*) = (\lambda^I, \gamma^I, fr^I)$, so we can interchange the subscripts I with $*$. We therefore obtain the equality:

$$\begin{aligned} &= \mathbb{E}_{\mathcal{O}}[\langle f(\underline{t}_1, \underline{t}_2) \rangle_{\underline{t}_1, \underline{t}_2}] \\ &= \sum_{\underline{t}_1, \mathcal{O}, \underline{t}_2} P_*(\mathcal{O}) \frac{P_*(\underline{t}_1) P(\mathcal{O}, \underline{t}_1)}{P_*(\mathcal{O})} P_*(\underline{t}_2|\mathcal{O}) f(\underline{t}_1, \underline{t}_2) \\ &= \sum_{\underline{\tau}, \mathcal{O}, \underline{t}} P_*(\underline{\tau}) P(\mathcal{O}, \underline{\tau}) P_*(\underline{t}|\mathcal{O}) f(\underline{\tau}, \underline{t}) \\ &= \mathbb{E}_{\tau, \mathcal{O}}[\langle f(\underline{\tau}, \underline{t}) \rangle_{\underline{t}}], \end{aligned}$$

where in the second passage we simply renamed $\underline{t}_1 \rightarrow \underline{\tau}$ and $\underline{t}_2 \rightarrow \underline{t}$. The quenched average $\mathbb{E}_{\mathcal{O}}$ is over the set of observations \mathcal{O} . Note that in our case, we average instead over the set of disordered variables \mathcal{D} : as explained above, this is equivalent to average over observations \mathcal{O} , since \mathcal{O} is a deterministic function of \mathcal{D} . We now discuss the connections with replica symmetry by observing that the equality (15) is true in particular when the function f is taken to be the overlap between two configurations. In words, it states that the average

of the overlap q between two configurations sampled independently from the posterior, is equal to the average of the overlap q^* between the planted configuration and a random sample from the posterior. Applying this equality to higher moments of the overlap, it is actually possible to show that the two distributions are equal. In models with gauge invariance, such as the planted spin glass studied in Refs. [14,16], the overlap q^* coincides with the magnetization. It can be argued that the magnetization is a self-averaging quantity, and therefore that the overlap q is also self-averaging on the Nishimori line. This argument allows to conclude that the probability distribution of the overlap q is trivial, and therefore that there is no replica-symmetry-breaking phase in the Bayes optimal setting. In the case of epidemics, it is less clear that the overlap q^* between the planted configuration and a random sample of the posterior is a self-averaging quantity. In fact, we observed a region (for small seed probability γ and large transmission rate λ), which presents signs of a failure of optimality, signalled by a lack of convergence of belief-propagation in the thermodynamic limit (see Sec. III, where we conjecture that this is due to replica-symmetry breaking). A similar observation is made in Ref. [8] on finite-size instances.

D. Belief-propagation equations for the joint probability

The factor graph associated with the probability distribution (13) contains short loops which compromise the use of BP. To remove these short loops, we introduce the auxiliary variables $(\tau_i^{(j)}, \tau_j^{(i)}, t_i^{(j)}, t_j^{(i)})$ on each edge $(ij) \in E$ of the factor graph, which are the copied times $\tau_i^{(j)} = \tau_i$, and $t_i^{(j)} = t_i$ for all $j \in \partial i$. Let $T_{ij} = (\tau_i^{(j)}, \tau_j^{(i)}, t_i^{(j)}, t_j^{(i)})$ be the tuple gathering the copied times on edge $(ij) \in E$. The probability distribution on these auxiliary variables is

$$P(\{T_{ij}\}_{(ij) \in E} | \mathcal{D}) = \frac{1}{\mathcal{Z}(\mathcal{D})} \prod_{i \in V} \Psi(\{T_{il}\}_{l \in \partial i}; \mathcal{D}_i), \quad (17)$$

where $\mathcal{D}_i = \{\{s_{li}\}_{l \in \partial i}, x_i^0, \{\varepsilon_m\}_{i_m=i}\}$ is the disorder associated with vertex $i \in V$, and with

$$\begin{aligned} \Psi(\{T_{il}\}_{l \in \partial i}; \mathcal{D}_i) &= \xi(\tau_i^{(j)}, t_i^{(j)}; c_i) \psi^*(\tau_i^{(j)}, \underline{\tau}_{\partial i}^{(j)}; \{s_{li}\}_{l \in \partial i}, x_i^0) \\ &\times \psi(t_i^{(j)}, \underline{t}_{\partial i}^{(j)}) \prod_{l \in \partial i \setminus j} \delta_{t_i^{(j)}, t_i^{(l)}} \delta_{\tau_i^{(j)}, \tau_i^{(l)}}, \end{aligned} \quad (18)$$

where $j \in \partial i$ is a given neighbor of i . The factor graph associated with this probability distribution now mirrors the original graph $\mathcal{G} = (V, E)$ of contact between individuals, as shown in Fig. 1. The variable vertices live on the edges $(ij) \in E$, and the factor vertices associated with the function Ψ live on the original vertex set V . We introduce the Belief Propagation (BP) message $\mu_{i \rightarrow \Psi_j}$ on each edge $(ij) \in E$ as the marginal probability law of T_{ij} in the amputated graph in which node j has been removed. The set of BP messages obey a set of self-consistent equations:

$$\mu_{i \rightarrow \Psi_j}(T_{ij}) = \frac{1}{\mathcal{Z}_{\Psi_i \rightarrow j}} \sum_{\{T_{ik}\}_{k \in \partial i \setminus j}} \Psi(\{T_{il}\}_{l \in \partial i}; \mathcal{D}_i) \prod_{k \in \partial i \setminus j} \mu_{k \rightarrow \Psi_i}(T_{ik}), \quad (19)$$

where $\mathcal{Z}_{\Psi_i \rightarrow j}$ is a normalization factor. These equations are exact when the contact graph $\mathcal{G} = (V, E)$ is a tree. In practice,




FIG. 1. The factor graph construction. On the left there is an example of contact network among individuals. We map this onto the factor graph on the right: for each individual we place a corresponding factor node and for each edge between two individuals we place the super-variable $T_{ij} = (\tau_i^{(j)}, \tau_j^{(i)}, t_i^{(j)}, t_j^{(i)})$ containing the planted and the inferred trajectories of both individuals. This construction increases the complexity of the BP messages, but allows to obtain a disentangled factor graph map (without short loops), which mirrors the contact network.

the BP method is also used as a heuristic on random sparse instance. Introducing a horizon time T , the random variable $T_{ij} = (\tau_i^{(j)}, \tau_j^{(i)}, t_i^{(j)}, t_j^{(i)})$ lives in a space of size $O(T^4)$. We see in Appendix A how to simplify the BP Eqs. (19), and obtain a set of equivalent equations defined over modified BP messages living in a smaller space.

1. Extension to SIR and SEIR models

The scheme described so far and the results discussed in the next section are for the SI case. However, extending everything to richer and possibly more realistic epidemic models is almost straightforward. In Appendix C we give the necessary ingredients to rewrite the factor graph for SIR and SEIR models. We also show a parametrization for the SIR model which leads to BP messages having the same complexity of the SI case. The role of Appendix C is twofold: to show that conceptually all the difficulties of the ensemble average are already contained in the SI model and to give details to any reader who wants to apply the method to more realistic scenarios.

E. Estimators

To quantify the feasibility of inference tasks, some estimators are defined in this paragraph and studied in Sec. III. In view of that, it is useful to define $P_{i,t}(x_i^{*,t}, x_i^t)$ as the marginal probability of having the planted state $x_i^{*,t}$ and the inferred state x_i^t of one individual $i \in V$ at a given time $t \in \{0, 1, \dots, T\}$.

1. Maximum mean overlap

The overlap at a given time t between the planted configuration $\underline{x}^{*,t}$ and an estimator $\hat{\underline{x}}^t$ is $O_t(\underline{x}^{*,t}, \hat{\underline{x}}^t) = \frac{1}{N} \sum_{i=1}^N \delta_{x_i^{*,t}, \hat{x}_i^t}$. In the inference process, on a given instance, the planted configuration is not known, and the best Bayesian estimator is obtained by assuming that $\underline{x}^{*,t}$ is distributed according to the posterior distribution. The best estimator of the overlap $\hat{\underline{x}}^{t, \text{MMO}}$ is the one maximising the overlap averaged over the

posterior:

$$\text{MO}_t(\hat{x}^t) = \sum_{\underline{x}^t} P(\underline{x}^t | \mathcal{O}) O_t(\underline{x}^t, \hat{x}^t), \quad (20)$$

which is achieved for $\hat{x}_i^{t, \text{MMO}} = \text{argmax}_{x_i^t} (P_{i,t}(x_i^t | \mathcal{O}))$. The overlap $O_t(\underline{x}^{*,t}, \hat{x}^{t, \text{MMO}})$ provides a quantitative estimation of the accuracy of the maximum mean overlap estimator $\hat{x}^{t, \text{MMO}}$. We compute this quantity, averaged over the graph ensemble, and the realization of the planted configurations and of the observations:

$$\mathbb{E}_{\mathcal{G}, \mathcal{D}}[O_t(\hat{x}^{*,t}, \hat{x}^{t, \text{MMO}})].$$

Note that in our formalism, we have access the marginal probability over planted and inferred configurations, conditioned on the disorder \mathcal{D} : $P_{i,t}(x_i^{*,t}, x_i^t | \mathcal{D})$. However, as previously explained, fixing the disorder variables is sufficient to fix the planted configuration and the observations \mathcal{O} .

2. Minimum mean-squared error

We also consider the squared error (SE) at a given time t between the planted configuration $\underline{x}^{*,t}$ and an estimator \hat{x}^t :

$$\text{SE}_t(\underline{x}^{*,t}, \hat{x}) = \frac{1}{N} \sum_{i=1}^N (x_i^{*,t} - x_i)^2.$$

As for the overlap, the best Bayesian estimator for the squared error $\hat{x}^{t, \text{MMSE}}$ is the one minimizing the squared error averaged over the posterior:

$$\text{MSE}_t(\hat{x}^t) = \sum_{\underline{x}^t} P(\underline{x}^t | \mathcal{O}) \text{SE}(\underline{x}^t, \hat{x}^t), \quad (21)$$

which is achieved for $\hat{x}_i^{t, \text{MMSE}} = \sum_{x_i^t} P_{i,t}(x_i^t | \mathcal{O}) x_i^t$. We compute the average, over the graph ensemble and the disorder, of the squared error between the planted configuration and the MMSE estimator:

$$\mathbb{E}_{\mathcal{G}, \mathcal{D}}[\text{SE}_t(\hat{x}^{*,t}, \hat{x}^{t, \text{MMSE}})].$$

3. Area under the curve (AUC)

On a single instance, the ROC curve is computed as follows. At a fixed time t , one computes for each individual its marginal probability $P_i(x_i^t = I | \mathcal{O})$. For a given threshold $\rho \in [0, 1]$, the true positive rate $\text{TPR}(\rho)$ [respectively, false-positive rate $\text{FPR}(\rho)$] is the fraction of positive (respectively, negative) individuals with $P_i(x_i^t = I | \mathcal{O}) \geq \rho$. The ROC is the parametric plot of $\text{TPR}(\rho)$ versus $\text{FPR}(\rho)$, with ρ the varying parameter. Note that the FNR (and therefore the ROC) is undefined when all individuals are infected (all positive). The area under the curve (AUC) can be interpreted as the probability that, picking one positive individual i and one negative individual j , their marginal probabilities allows to tell which is positive and which one is negative, i.e., that $P_i(x_i^t = I | \mathcal{O}) > P_j(x_j^t = I | \mathcal{O})$. This allows us to compute the AUC under the replica-symmetric formalism.

F. Inferring the hyperparameters

The prior parameters, with which the planted epidemic is generated, might not be accessible in the inference process.

In those cases, we propose to infer them from the observations by approximately maximizing $P(\mathcal{O} | \lambda^l, \gamma^l)$ at fixed observations \mathcal{O} . In this section we provide an upper bound to the feasibility of parameters inference. It is an upper bound because the process is described for the ensemble case, i.e., for an infinite-sized contact graph. This means that also the number of observations is infinite. The idea is to find the most typical parameters γ^l, λ^l for the given set of observations. For inferring γ we use the expectation maximization (EM) method, an iterative scheme which consists in separating the optimization process in two steps:

(1) At fixed BP messages, the update for γ at k th iteration reads

$$\gamma_k = \arg \max_{\gamma} \langle \log P(\underline{t}, \mathcal{O} | \gamma) \rangle_{\{\mu\}_k}, \quad (22)$$

where $\{\mu\}_k$ is a shorthand for the set of all BP messages at k th iteration.

(2) At fixed γ_k , the messages are updated with BP equations.

To understand Eq. (22) we recall the definition of the variational free energy:

$$F[Q](\gamma) := -\langle \log P(\underline{t}, \mathcal{O} | \gamma) \rangle_Q + \langle \log Q(\underline{t}) \rangle_Q. \quad (23)$$

The posterior distribution $P(\underline{t} | \mathcal{O}, \gamma)$ can be shown to be the distribution Q which minimizes F (see, for example, Ref. [17]). If we evaluate averages with fixed BP messages, then the dependency of F on γ is only on the first addend of the right hand side. Then the optimization on γ reduces to Eq. (22). By setting to zero the first derivative of Eq. (22) w.r.t. γ we have, for the k th iteration,

$$\gamma_k = \frac{1}{N} \sum_{i \in V} p_i^{t,k}(t_i = 0 | \mathcal{O}), \quad (24)$$

where $p_i^{t,k}(t_i = 0 | \mathcal{O})$ is the posterior probability at k th iteration of EM for individual i to have infection time equal to 0 (i.e., to be the patient zero). The procedure we propose to include the inference of patient zero probability is therefore to simply update γ^l with Eq. (24) at every sweep of BP update on the population. We could write equations for EM in λ , but they would be more involved. We opted therefore to simply perform a gradient descent (GD) on the Bethe free energy. For the epidemic propagation on graph, the Bethe free energy can be related to the normalization of the BP messages and BP beliefs, respectively, addressed as $z_{\Psi_i \rightarrow j}$ and z_{Ψ_i} :

$$F_{\text{Bethe}} = \sum_{i \in V} \left[\left(\frac{1}{2} |\partial i| - 1 \right) \log z_{\Psi_i} - \frac{1}{2} \sum_{j \in \partial i} \log z_{\Psi_i \rightarrow j} \right].$$

Overall, the inference of prior hyperparameters is done by alternating one sweep of BP update of the messages at fixed parameters (γ^l, λ^l) with one step of EM for γ and one step of GD for λ .

III. RESULTS

In this section we explore the performance of inference tasks for the SI model. We consider two different regimes:

(1) when the parameters of the prior distribution are known in the inference process (Bayes optimal setting),




FIG. 2. Several measures (first row, MMSE; second row, MMO; third row, AUC) quantifying the hardness of epidemic inference, as a function of patient zero probability γ , and infection probability λ . Each column corresponds to three different times at which the quantities are computed (from left to right: initial time $t = 0$, intermediate time $t = 4$, and final time $t = T = 8$). The three measures display the same behavior, except for the initial time, when AUC is able to capture for high values of λ and γ that observations are not informative enough. Notice that MMSE quantifies the error in inferring individual's states, so it has a flipped behavior with respect to the other quantities (MMO and AUC are high when inference performance is good). These results were obtained for ER graph ensemble with average degree 3.

(2) when the parameters of the prior distribution are not known and are inferred.

The results are obtained by randomly initializing a population of BP messages and by iterating a population dynamics algorithm, as described in Appendix B. The algorithm stops when the BP messages satisfy a simple convergence criterion on the marginals, which must not fluctuate more than the square root of size of the population. If convergence criterion is not reached, then the algorithm stops after a fixed number of *maxiter* sweeps (typically we set the population size $N \sim 10^4$ and the total number of sweeps $maxiter = 100$). Convergence is (almost always) reached when the prior is known (or inferred), except for a rather interesting and unexpected regime which is discussed later on. The algorithm shows nonconvergence zones, as expected, also when the prior is not known.

A. Results in the Bayes-optimal case

In this paragraph, we study several measures and estimators that quantify the hardness of inference, varying epidemic parameters (transmission probability λ , patient zero probability γ), but also the fraction of observed individuals. We compare the minimum mean-squared error (MMSE), the maximum mean overlap (MMO), and the area under the ROC curve (see Sec. II E for their definition), and the Bethe free energy (Fe) associated with the posterior distribution $P(\underline{\ell}|\mathcal{O})$. In Fig. 2, we fix the fraction of unobserved individuals

(dilution) to $dil = 0.5$ (half the individuals are observed). We set the observation time at final time $T = 8$, and explore the space (γ, λ) . MMSE, MMO and AUC are computed at three different times (initial time $t = 0$, intermediate time $t = 4$, and final time $t = T = 8$). We can see that MMSE and MMO show the same behavior at all times. For very low infection probability λ , and patient zero probability γ , we see that MMSE is low (while MMO and AUC are high), meaning that the information contained in the inferred posterior distribution allows to recover the planted configuration with good accuracy. In this regime, typically seeds are surrounded by a small neighborhood of infected individuals, well-separated from the other seeds, making inference task easy. For high values of patient zero probability γ and infection λ , instead, the population becomes completely infectious in few time steps. Also in this regime, all the estimators show great performance for intermediate ($t = 4$) and final times, because the posterior marginals assign to every individual a probability 1 of being infectious. The hard regime is for intermediate values of γ, λ .

Note also that at $t = 0$, MMSE (respectively MMO) is low (respectively, high) for high values of γ . However, this does not mean that inference performance is good in this regime. Indeed for large γ , the majority of individuals are patients zero, and the other individuals are likely to be infected before the observation time T . Therefore, the observations are (almost) all positive, making impossible to distinguish the patients zero from the ones infected at later time. Thus, MMSE at time $t = 0$ is low because the marginal posteriors give high probability of being infected at $t = 0$, independently of the transmission rate λ . However, the (few) nonpatients zero will remain undetected. A quantity that is sensible to this problem is the AUC, which at time $t = 0$ has in fact a different behavior with respect to the other measures. Another (slightly) different quantity, for example, is the AUC evaluated only on non observed individuals. When many observations are done, the AUC is dominated by the observed individuals. Thus, evaluating AUC only on non observed individuals (AUCNO) can be a useful tool to understand the predicting power of the algorithm on individuals for which no information is given. To see the difference between AUC and AUCNO, we fix the patient zero probability $\gamma = 0.1$, and study these two measures as a function of the infection probability λ and the observations dilution dil (i.e., the fraction of unobserved individuals), see Fig. 3. We see that the two estimators behave differently, for example at the intermediate time $t = 4$, for low dilution (i.e., many observations) and low transmission rate λ . In this regime, there are only few infected individuals at the observation time (low γ and λ). While AUC is close to 1, AUCNO is low, indicating that it is actually hard to find who are the *unobserved* infected individuals.

1. Comparison with finite-size instances

It is natural to wonder whether our ensemble results obtained in the thermodynamic limit, with the RS cavity method, are consistent with large finite-size single instances. To check this point, we initialized large sized ($N = 30\,000$) graphs, and simulated discrete-time epidemic spreading and observation protocol. We used Sibyl [5], which is a Belief Propagation algorithm for calculating the posterior marginals in single

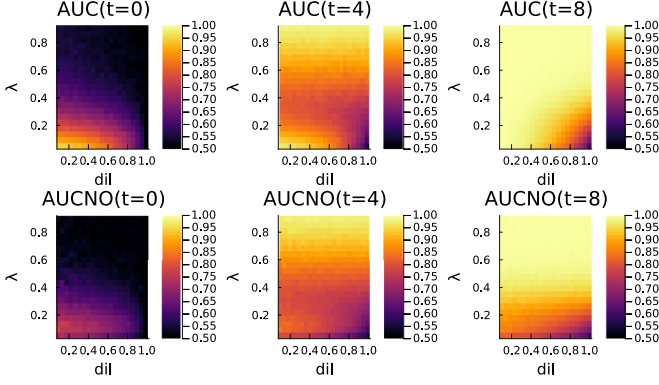


FIG. 3. A comparison between AUC evaluated on all individuals (AUC, first row) vs only unobserved individuals (AUCNO, second row). The two measures have a very similar but not identical behavior. In particular, at low dilution (many observations) the AUCNO is systematically smaller than the AUC as expected. This difference is more pronounced for low values of λ , and at intermediate and final time. These results are for ER graph ensemble, with average degree 3. We remark here that AUC is not 0.5 for dilution equal to 1. In fact, ER graphs are heterogeneous (with a Poisson law degree distribution). This implies that some information about the probability of infection of each node is contained in the graph itself. For example, the most connected nodes have highest probability of being infected. This allows to achieve some reconstruction also without any observation ($dil = 1$).

instance problems. We computed the MMSE, and we compared it with the RS predictions. An example is shown in Fig. 4, showing a good agreement between the RS predictions and the results on a single large instance.

2. Information contained in the observations

All the observables described so far are time-dependent quantities, giving an estimation on how easy/hard it is to infer the planted individual states $x_i^{*,t}$ at a given fixed time t . It is useful to define a time-independent observable, which gives a general overview of the inference process. We opted to study the Bethe free energy $F = -\log Z(\mathcal{D}) = -\log P(\mathcal{O})$, which can be expressed in terms of BP marginals (see its expression in Appendix A 6). It is the free energy associated with the posterior distribution:

$$P(\underline{t}|\mathcal{O}) = \frac{P(\underline{t}, \mathcal{O})}{P(\mathcal{O})}, \quad (25)$$

where

$$P(\mathcal{O}) = \sum_{\underline{t}} P(\underline{t}, \mathcal{O}) = \sum_{\underline{t}} P(\mathcal{O}|\underline{t})P(\underline{t}).$$

$F = -\log P(\mathcal{O})$ quantifies how informative the observations \mathcal{O} are: the quantity $P(\mathcal{O})$ is the sum of trajectories (weighted with their prior probability) which are compatible with the observation constraints \mathcal{O} . Observations reduce the space of the possible trajectories; as an extreme example, if we observed (noiselessly) every individual at every time, the space of possible trajectories would collapse on a single trajectory (the planted solution). The plots in Fig. 5 show the free energy in the two different regimes discussed above. The free energy is

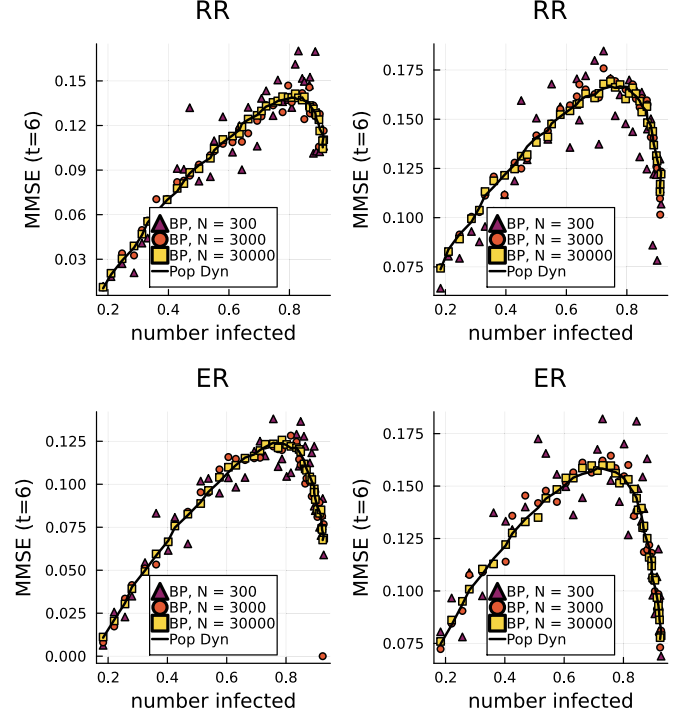


FIG. 4. The comparison between Sibyl algorithm (BP) for a single instance of $N = 300$ (triangles), $N = 3000$ (dots), and $N = 30000$ (squares) individuals, and the ensemble results obtained in the thermodynamic limit with population dynamics (black solid line). The plots show the MMSE at intermediate time $t = 6$, as function of the number of infected at final time $T = 8$, which is a function of infection parameters γ and λ . For this plot the patient zero probability is fixed at $\gamma = 0.15$. The first row represents the MMSE for random regular graphs (degree 3) while the second row is for Erdős-Rényi (ER) with average degree 3. Each column instead is associated with a value of observations dilution dil : the first column is for $dil = 0$ (all observed) while the second is for $dil = 0.5$. We see a very good agreement, that increases with the size of the the single instance contact graph, and that we checked to persist in the other observables (MMO and AUC).

obviously 0 for $dil = 1$ (no observation). However, it is close to 0 in other cases too, e.g., for high values of λ . For those values, the infection spreads very fast. As a consequence, at final time all the individuals are infected. Thus, since the observations are taken at final time, they do not bring valuable information on the planted trajectory: they will always register a positive (infected) result for all individuals at final time. In other words, all trajectories sampled from the prior are compatible with the observation that all individuals are infected at time T . Note however that inference can be easy in this regime, as it can be checked comparing Fig. 5 with Figs. 2 and 3 (for times $t = 4$ and $t = 8$). In this regime, although observations are not informative, the prior is concentrated on few trajectories (that are compatible with all individuals being infected at times $t = 4$ and $t = 8$), making inference task trivial. The interesting (and hard) regimes are at intermediate/low values of γ and λ and for nonzero dilution. In this regime, although observations are informative (F positive) the prior is not concentrated on few trajectories, making the inference a nontrivial task.




FIG. 5. Free-energy profile for two different regimes. Left panel: as a function of patient zero probability γ and infection probability λ , at fixed dilution $dil = 0.5$. Right panel: as a function of observations dilution dil and λ , at fixed patient zero probability $\gamma = 0.1$. The black part of the plot corresponds to the regimes in which observations do not bring any information, i.e., $F \simeq 0$. This happens obviously at $dil = 1$ because no observation is done. However, also for nonzero dilution, the free energy can be zero. When the infection transmission λ is high enough, in fact, all the individuals are infected at final time with probability almost equal to 1. Since for this plot the observations are performed at final time, then they all simply register the infectiousness of each individual, factually carrying no information with them. Only in the intermediate regimes, i.e., when the numbers of infectious and susceptible individuals are comparable to each others, observations carry information. In this regime the free energy is nonzero and inference is non trivial. The graph ensemble analyzed here is Erdős-Rényi with average degree 3.

3. More graph ensembles

The analysis shown so far has been performed on Erdős-Rényi graphs. To study if and how inference performance is affected by the graph structure, we compared the results on three families of graphs:

- (1) The random regular (RR) ensemble, where each node has the same degree d ;
- (2) The Erdős-Rényi (ER) ensemble. In the large-size limit, the degree distribution is a Poisson distribution of average d ;
- (3) A (truncated) *fat-tailed* (FT) ensemble of graphs, with a degree distribution $p(d) = \frac{1}{Z} \frac{1}{d^2+a}$ for $d \in [d_{\min}, d_{\max}]$ and $p(d) = 0$ if $d \notin [d_{\min}, d_{\max}]$. The quantity Z is the normalization of the distribution and the parameter a can be fixed by fixing the average degree.

The choice for the third graph ensemble allows for the existence of highly connected nodes, while still being handled by belief propagation (BP), since the distribution of the degree is truncated to a finite maximum value d_{\max} . In Fig. 6 (first row), we compare the minimum mean-squared error (MMSE) at time $t = 6$ for the three ensembles of graph. The average degree is fixed to 3 in all three graph ensembles.

4. Noise in observations

When noise affects individuals' observations, the inference results get typically worse. This can be seen in Fig. 6 (second row), where we studied the AUC as function of observations dilution and noise (false rate, fr). For false rate equal to 0.5, the observations carry no information, since they are wrong half of the time. This is identical to set dilution to $dil = 1$, i.e., not performing any observation. For intermediate values,

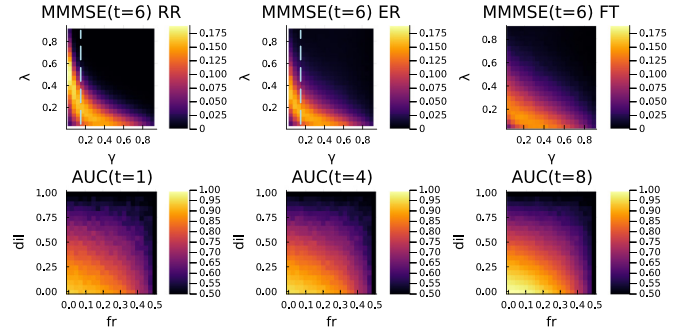


FIG. 6. Comparing feasibility of inference for different graph ensembles and for nonzero observation noise. First row: the plots show the MMSE at time $t = 6$, with observations made at final time $T = 8$, as functions of the patient zero probability γ and the infection probability λ . The three plots are (from left to right) for random regular (RR), Erdős-Rényi (ER) and fat-tailed (FT) graph ensembles. The average degree is fixed to 3 for the three ensembles examined. It can be seen that the profiles share similarities, but the more the degree distribution widens (from RR to ER to FT), the flatter is the MMSE. This is due to the presence (or absence) of high-degree nodes. In RR ensemble, all nodes have the same degree, so for example we see that inference is more difficult at low values of γ and values of λ . In this region, the presence of highly connected nodes simplifies inference because they (and their neighbours) will probably be infectious at time $t = 6$. The dashed lines correspond to the cases studied in Fig. 4 at dilution 0.5. The only difference is in the y axis, which is λ for this plot and the number of infected for the plots in Fig. 4. Second row: the AUC as function of observations' dilution (dil) and false rate (fr). The AUC decreases with fr and dil . The false-positive rate and false-negative rate are assumed to be the same. The patient zero probability is fixed to $\gamma = 0.03$ and the infection probability is $\lambda = 0.03$. The ensemble graph is random regular with degree 11.

we see that increasing false rate and/or dilution always leads to worse inference, as expected.

5. Application to a real network

We show here how our predictions on the infinite-size limit can be applied to more realistic random ensembles. Given a real finite network, we consider the configuration model ensemble with a degree distribution taken from the degree sequence of the graph. The cavity method allows to approximately find the expected value of each desired estimator in each regime for that ensemble, in the thermodynamic limit. We applied this procedure to a network of sexual contacts [18] from the data repository [19], and compared the results on the single instance with the cavity method approximation on the infinite-size limit (see Fig. 7). The first row shows the ensemble results, while the second shows single-instance results obtained with Sib. The phase diagrams show the AUC at initial, intermediate and final time as function of the infection and the observation dilution in the Bayes-optimal regime. The agreement between the two results is good, implying that the quality of inference performance in this setup depends mostly on the degree distribution, which is a global characteristic of the network.




FIG. 7. An application of the cavity method to a real network. First row: ensemble predictions for AUC at time 0,4,8. Second row: results obtained on the real network using Sib [5], the epidemic inference algorithm based on BP, the single instance method corresponding to the cavity method. The predictions are quite similar to the results obtained by Sib. Some differences can be seen for low values of the infection probability λ , where Sib has higher performance. A possible explanation could have been that the ensemble algorithm makes an incorrect estimate in the number of infectious individuals, however, we checked that the ensemble prediction for the number of infectious individuals in time is correct and consistent with Sib (not shown). The reason why Sib performs better than the cavity method should be related to the fact that at a single instance level the observations \mathcal{O} are more informative than in the ensemble case: if one individual is observed in the state S , then it is true that the infection cascade did not pass through it. Therefore, the correlations among its neighboring individuals drop down due to the S observation, which factually cuts the network. This reasoning does not subsist in the ensemble case, where the network is not fixed. Therefore, we expect lower AUC in the regime in which the S observations are more frequent (low infection λ). For this plot, the value of the patient zero probability is fixed at $\gamma = 0.2$ and there is no noise in observations, which are all made at time $T = 8$.

6. Convergence-breakdown for low seed probability

A surprising behavior of the Belief Propagation equations was observed in Ref. [8], for single instances at small values of γ , the patient zero probability. In fact, even in the Bayes optimal conditions, BP stops to converge. We checked that this breakdown of convergence is actually present even in the thermodynamic limit, using our population dynamic algorithm. This lack of convergence, therefore, seems to be related to a more profound reason. To understand what is happening, we simplified the framework by setting the infection probability λ to 1 and by observing all the individuals at final time. In this regime, in Fig. 8 the black dots represent the number of iterations needed for the population dynamics algorithm to converge. Around $\gamma = \tilde{\gamma} = 0.013$ the algorithm stops converging. An intuitive explanation to explain this behavior is the following: for γ around $\tilde{\gamma}$ at final time many individuals are observed infectious (I) and a small (but extensive) part are observed susceptible (S). As $\lambda = 1$, the sole nondeterministic part of the process is the initial state, so the inference problem reduces to guess the position of the patients zero. The S individuals not only signal that they were not infected in the epidemic process, but also that any patient zero must be at distance $>T$. For example, for a RR graph, this excludes a sphere centered in S -observed with $d(d - 1)^{T-1}$




FIG. 8. Study of population dynamics, Monte Carlo and number of clusters for zero-patience inference. We studied, for a RR graph with degree 3, the convergence of population dynamics (infinite graph) and Monte Carlo (finite-size graph). It breaks down at around $\tilde{\gamma} \simeq 0.013$. The black dots represent the number of iterations for population dynamics to reach convergence, normalized by the total number of iterations allowed. The continuous squared-marked lines represent the fraction of successful Monte Carlo runs, for several sizes ($N, 2N, 4N$, with $N = 5000$). We say that MC is successful every time it reaches a configuration which satisfies the observations (see main text). The failure of Monte Carlo coincides with BP failure. We conjecture that this is due to replica-symmetry breaking. We think that the main reason of this breakdown is because the space in which a patient zero can be placed in the posterior becomes clustered (see Fig. 9). To support this conjecture, we plot the fraction of connected components of the subgraph of all the individuals that could be the patients zero without violating the S observations. This number, as expected, grows sharply in the interval in which BP ceases to converge. The failure of convergence arises when the number of possible zones to place the patient zero (continue, blue line) becomes higher to the actual fraction of patients zero (dotted, blue line). This suggests that when the number of zones in which a patient zero might be becomes larger than the number of patients zero, then the problem becomes hard, as illustrated in Fig. 9.

individuals. For γ around $\tilde{\gamma}$ these spheres touch and intersect, so that the group of individuals eligible to be the patients zero gets separated in clusters. In Fig. 9 we give a plot in 2D of the phenomenon. When the number of observed S is sufficiently high, due to the fact that there are few patients zero, the space in which patients zero could physically be gets fragmented. To check that this is what actually happens in a random regular graph, we initialized a graph and counted the connected components in which a patient zero could be present without violating the S observations. This number actually sharply increases in the decreasing γ direction, around $\tilde{\gamma}$, i.e., when the algorithm stops converging. Further evidence of a phase transition is given by Monte Carlo dynamics. We implemented a simple Monte Carlo simulation on a graph. We first sampled the planted (ground-truth) configuration, from which we collected the observations. The observation protocol was set to observe all the individuals at the final time T (without observation noise). Then we started a Metropolis-Hasting Monte Carlo simulation to sample a configuration satisfying all the observations. To do so, we initialize a configuration by doing




FIG. 9. A 2D plot to visualize the geometric change undergone by the configuration space that could explain why population dynamics and Monte Carlo schemes stop to converge. In this plot, obtained by simulating epidemic spreading in a two-dimensional lattice, we compare two scenarios. To the right, γ is higher, namely there are more patients zero (green dots). This implies that the number of infected (green, red, and gray dots) is higher, so the number of S -observed individuals (no dots) is smaller. The patients zero can not be too close to the S -observed individuals because the infection probability is 1, so the observation constraint would be violated. The red dots represent all the individual which might be the patient zero according to the observations (i.e., individuals tested I and not too close to S -observed individuals). When the number of patients zero is lower, (left) the number of S -observed individuals increases. So the possible zones to accommodate patients zero (green plus red dots) reduce and get clustered. This could create several separated states of the posterior, each one corresponding to a possible combination of placements of patients zero.

a sample x of the prior distribution. The initialization configuration typically does not satisfy the observations. So we make the following move: we randomly select an individual and we change its $t = 0$ state by sampling the I state with probability γ [and the S state with probability $(1 - \gamma)$]. Subsequently, the initial state configuration is evolved (deterministically, since the infection probability is $\lambda = 1$). The configuration at final time is then checked to be consistent with the observations. In particular, we introduced the energy

$$U = - \sum_{i=1}^N \log p(o_i | x_i^T), \quad (26)$$

where x is the configuration and each o_i is the observation on the i th individual. In principle $p(o_i | x_i^T)$ should be either 0 (when the configuration does not satisfy the observations) or 1 (when the observation is satisfied). To avoid infinite energy barriers, we introduced a small noise in observations, which we reduced during the Monte Carlo by means of an annealing procedure. In other words, the energy is just a penalization for each broken constraint. At each step, the move in the space of initial states described above is made. The move is accepted by following a standard Metropolis scheme. The MC stops when the configuration satisfies all constraints. For each value of γ we repeated 60 times the MC scheme and computed the fraction of runs in which the algorithm was able to reach a configuration satisfying all observation constraints. We plot in Fig. 8 the fraction of Monte Carlo processes that reached a configuration of the posterior. We clearly see that this quantity drops down around $\tilde{\gamma}$. Due to the failure of BP equations (for finite and infinite graph), the explosion of possible patient




FIG. 10. Ensemble algorithm breakdown outside Nishimori condition. Keeping fixed $\lambda^* = 0.3$ and the parameters $\gamma^* = \gamma^I = 0.03$, while instead moving the inference parameter $0.0 < \lambda^I < 0.8$ and the dilution between 0 and 1 we see that for high values of λ^I the algorithm does not converge and provides nonphysical results for the free energy.

zero zones and the failure of the Monte Carlo scheme we conjecture replica-symmetry-breaking transition around $\tilde{\gamma}$.

B. Departing from Bayes-optimal conditions

It is well known that when inference is performed with imperfect knowledge of the prior distribution parameters, it is possible to observe a replica-symmetry-breaking (RSB) phase transition. This is due to the fact that outside the Bayes optimality regime, the Nishimori conditions are no longer guaranteed to be valid, see Sec. II C. An RSB phase can manifest itself with a convergence failure of the population dynamics algorithm, which is based on the replica-symmetry hypothesis. This is exactly what we see in Fig. 10. We recall that we use a star (*) to label the parameters with which the planted configuration is generated, e.g., λ^* is the true infection probability, while λ^I is the infection probability used by the algorithm in the inference process. For this plot we fixed $\gamma^* = \gamma^I$ (so we gave to the algorithm the exact value of patient zero probability) and we studied the free-energy landscape by varying λ^I and the observations dilution. There exists a zone in which the number of iterations reaches the maximum allowed number (which was set to 100). In this zone, the observables show an oscillating behavior. This suggests a breakdown of the algorithm validity, which may be caused by an RSB phase transition. In any case, when the prior is not known, some difficulties arise in epidemic inference. A good strategy to avoid them is to infer the prior parameters, as shown in the next paragraph.

1. Inferring epidemic prior's parameters

We infer the prior parameters by (approximately) minimizing the Bethe free energy. In particular, for the patient zero probability γ we use the expectation maximization (EM) method. For the infection probability λ , instead, we perform a gradient descent (GD) on the free energy. This mixed strategy (EM for inferring γ and GD for λ) was adopted due to its simplicity in terms of calculations. To check if the method works, we studied inference in the same conditions of Fig. 2. We therefore fixed observations dilution to 0.5, and we explored the space of patient zero probability γ^* against infection probability λ^* . Initializing the inferring parameters to $\gamma^I = \lambda^I = 0.5$, the results are shown in Fig. 11. The plot shows a comparison between the observables computed by




FIG. 11. A comparison between the inference feasibility (quantified here with MMMSE and free energy) in the case in which prior’s parameters are known (*first column*) and when instead they are learnt (*second column*). The quantities are represented as functions of the planted parameters γ^* and λ^* . In the first row the MMSE at intermediate time $t = 4$ is plotted: on the left there is the optimal Bayes result, already shown in Fig. 2, while on the right there is the result obtained when the infection and patient zero parameters λ^l and γ^l are learned. On the second row the same comparison (i.e., Bayes optimality on the left and hyperparameters’ learning on the right) is done for free energy. In both cases (MMSE and free energy) the initial conditions for the hyperparameters were set to $\lambda^l = 0.5$ and $\gamma^l = 0.5$. The results are for the Erdős-Rényi ensemble with average degree of 3. Observations are made at final time $T = 8$.

inferring the prior parameters and their respective quantities in the Bayes optimal case, i.e., the ones plotted in Fig. 2 (first row) and Fig. 5 (left). The prior parameters are learned by minimizing the free energy, which agrees almost perfectly with the optimal one. There is a strong agreement also for other observables, as the MMSE, which we plotted at time $t = 4$. To actually see how well the prior hyperparameters are inferred, we plot them in function of their planted respective quantity (see Fig. 12). It is again important to compare the results of prior parameters inference with the single instance results on finite graphs. Indeed, the inference results shown so far are for infinitely large graphs. The number of observations is therefore infinite too. It is then crucial to see whether for finite-size graphs (and finite information) it is possible to achieve comparable results to the ensemble. In Fig. 13 we see that this is the case. We compare the population dynamics and the single instance code by analyzing step-by-step their gradient descent on the patient zero and infection probabilities. The plot shows that, as expected, the ensemble inference is more precise due to infinite amount of information available. However, the values inferred by the single instance algorithm are very close to the true ones.




FIG. 12. The inferred prior parameters in function of their respective planted quantities. The plot is obtained at zero dilution (all individual observed) and for (uniformly) scattered observations in time. In the left panel, patient zero parameter γ^l is plotted in function of γ^* for different values of λ^* . The right panel’s lines are instead the values of the infection λ^l as function of λ^* for different values of γ^* .

2. Addressing biased observations

In realistic contexts, observations are not taken uniformly at random from the population. This is because infected people might manifest some symptoms, which push them to test themselves. The probability of being observed is therefore typically higher for infected people than for susceptible ones. A first consequence is that the fraction of infected individuals in the population is not equal to the fraction of infected ones in the set of observed individuals. If in the inference process this is not considered, then the risk is to achieve low performance. Not considering the bias of observations means to infer with an incorrect prior parameter, i.e., outside of




FIG. 13. Inference of prior parameters with the ensemble code (pop dynamics) compared to the single instance result, obtained running the belief propagation (BP) algorithm on a contact network of $N = 10\,000$ nodes. The plot shows the gradient descent in free energy with respect to the two parameters γ^l and λ^l which respectively represent the patient zero and the infection probabilities. For the infection λ^l the gradient descent is performed using the Sign Descender technique with learning rate 0.01, while for γ^l we used the expectation maximization method. The results are for Erdős-Rényi (ER) graphs with average degree 3. All the individuals are observed at final time.

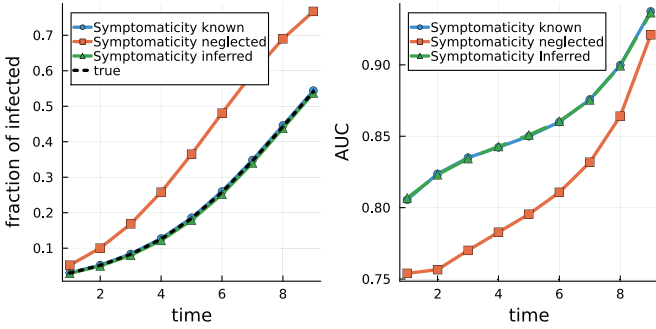


FIG. 14. Considering (and inferring) bias in observation allows to recover Nishimori conditions and improves inference performance. The bias is generated by symptomatic individuals, which are all assumed to be tested. The probability for an infected individual of being symptomatic was set to $p_+ = 0.5$. Asymptomatic individuals can also be tested. For this plot, the probability for an asymptomatic individual to be randomly selected for a test was set to $p_r = 0.04$. The left plot shows the estimated fraction of infected individuals over time. Considering the bias in the inference process allows to reconstruct this function. On the right plot we compared inference performance when the bias is considered vs when it is neglected. Considering the bias systematically improved the AUC. The patient zero probability was set to $\gamma = 0.03$ and the infection probability to $\lambda = 0.25$. The observations are all performed at time $T = 8$.

the Nishimori conditions. To quantify this bias, we introduce p_+ , the probability for an infected individual to be symptomatic. We assume that all infected symptomatic individuals are tested. Asymptomatic individuals are instead tested at random with probability p_r . From this, the probability for an infected individual to be tested, with a positive test result is

$$P(\text{tested, positive}|I) = (1 - \text{fr})[p_+ + p_r(1 - p_+)],$$

and similarly,

$$P(\text{tested, negative}|I) = \text{fr}[p_+ + p_r(1 - p_+)].$$

For susceptible states S :

$$P(\text{tested, positive}|S) = p_r \text{fr},$$

$$P(\text{tested, negative}|S) = p_r(1 - \text{fr}).$$

The unbiased case is recovered for $p_+ = 0$. We want to compare inference results when the bias p_+ is considered and when instead is neglected. In Fig. 14 (left panel), we see a substantial overestimation of the infection when ignoring the bias. In the right panel, we see that the AUC is systematically higher when the bias is included. We finally inferred the bias p_+ by minimizing the free energy (following exactly the same procedure of the transmission rate's inference). This process allows to include the unknown bias without affecting performances.

IV. CONCLUSION

In this paper, we study the feasibility of inference in epidemic spreading on a contact graph. Using the replica-symmetric cavity method, we give quantitative predictions of several estimators [minimum mean-square error (MMSE), maximum mean overlap (MMO), and area under

the ROC curve (AUC)], in different regimes depending on the characteristics of the epidemics, the observations and the contact graph.

In the Bayes-optimal setting, we show that for a large range of the model's parameters, the RS predictions are in good agreement with the results obtained on finite-size instances. It was noted in Ref. [8] that BP equations did not converge on large instances in a particular region of the parameters (at low seed probability and high rate transmission), a fact that is also confirmed by our simulations. Our simulations in that region show a lack of convergence of the cavity equations in the thermodynamic limit (answering thus negatively to the conjecture in Ref. [8] of it being to finite-size effects), and strongly hinting to replica-symmetry breaking.

In the non-Bayes optimal setting (i.e., when the parameters of the posterior differ from the parameters used in the prior), we observe a region where the iterative numerical resolution of the replica-symmetric cavity equations does not converge, suggesting the presence of a replica-symmetry-breaking transition. We show however that inferring parameters allowing to recover performance comparable to the one obtained in the Bayes optimal setting is possible with a simple iterative procedure, for a large range of the prior's parameters. There are however situations in which one is forced to work outside the Bayes optimal case (and for which replica-symmetry breaking is to be expected), e.g., when some parameters of the model are known only approximately but are too many to infer or when the knowledge of the contact network itself is imperfect.

Averaging over correlated disorder within the framework of the cavity method is the main technical issue addressed in this paper. The strategy developed here could be applied to more involved irreversible epidemic models, such as the SIR and SEIR model. The main limitation would be an increase of space size of the dynamical variables: each compartment added would come with a additional couple of transition times (one planted and one inferred time). The strategy could be applied more generally to any model in which disorder can be decomposed into a set of local (independent) random variables \underline{s} , and a set of correlated variables $\underline{\tau}$ that can be computed from the first set. Note however that each element of the correlated disorder $\underline{\tau}$ should be expressed only as a function of a local subset of $\underline{\tau}$ and \underline{s} . In other words, there must exist a function

$$\psi(\underline{\tau}|\underline{s}) = \prod_{i \in V} \psi_i(\tau_i | \underline{\tau}_{\partial i}, \underline{s}_{\partial i}), \quad (27)$$

with arbitrary factors ψ_i , which for fixed \underline{s} is nonzero only for a given value $\underline{\tau}$.

ACKNOWLEDGMENTS

This study was carried out within the FAIR—Future Artificial Intelligence Research—and received funding from the European Union Next-GenerationEU (Piano Nazionale Di Ripresa e Resilienza (PNRR)—Missione 4 Componente 2, Investimento 1.3—D.D. 1555 11/10/2022, PE00000013). This manuscript reflects only the authors' views and opinions, nei-

ther the European Union nor the European Commission can be considered responsible for them.

A public GitHub repository containing a Julia implementation of the algorithm and notebooks is available at [20].

APPENDIX A: BP EQUATIONS AND BETHE FREE ENERGY

In this Appendix we derive a simplified version of the BP Eqs. (19) introduced in Sec. IID. These simplified equations [given in Eqs. (A13) and (A14)] are over a set of modified messages represented in Fig. 15.

1. Clamping

In the numerical resolution of the cavity equations, it will be convenient to introduce a horizon time $T + 1$ above which the epidemic evolution is not observed. This results in a modification of the function ψ^* ensuring the constraints on infection times:

$$\psi^*(\tau_i, \underline{\tau}_{\partial i}, x_i^0, \{s_{ji}\}) = \mathbb{I}[\tau_i = \delta_{x_i^0, S} \times \min(T + 1, \min_{l \in \partial i}(\tau_l + s_{li}))]. \quad (\text{A1})$$

$$\psi^*(\tau_i^{(j)}, \underline{\tau}_{\partial i}^{(i)}, \{s_{li}\}_{l \in \partial i}, x_i^0) = \delta_{x_i^0, I} \delta_{\tau_i^{(j)}, 0} + \delta_{x_i^0, S} \prod_{l \in \partial i} \mathbb{I}[\tau_i^{(j)} \leq \tau_l^{(i)} + s_{li}] - \delta_{x_i^0, S} \mathbb{I}[\tau_i^{(j)} < T + 1] \prod_{l \in \partial i} \mathbb{I}[\tau_i^{(j)} < \tau_l^{(i)} + s_{li}] \quad (\text{A2})$$

and

$$\begin{aligned} \psi(t_i^{(j)}, \underline{t}_{\partial i}^{(i)}) &= \sum_{x_i^0} \gamma(x_i^0) \sum_{\{s_{li}\}_{l \in \partial i}} \prod_{l \in \partial i} w(s_{li}) \mathbb{I}[t_i^{(j)} = \delta_{x_i^0, S} \min(T + 1, t_l^{(i)} + s_{li})] \\ &= \gamma \delta_{\tau_i^{(j)}, 0} + (1 - \gamma) \left[\prod_{l \in \partial i} \left(\sum_{s=1}^{\infty} w(s) \mathbb{I}[t_i^{(j)} \leq t_l^{(i)} + s] \right) - \mathbb{I}[\tau_i^{(j)} < T + 1] \prod_{l \in \partial i} \left(\sum_{s=1}^{\infty} w(s) \mathbb{I}[t_i^{(j)} < t_l^{(i)} + s] \right) \right] \\ &= \gamma \delta_{\tau_i^{(j)}, 0} + (1 - \gamma) \left[\prod_{l \in \partial i} a(t_i^{(j)} - t_l^{(i)} - 1) - \mathbb{I}[\tau_i^{(j)} < T + 1] \prod_{l \in \partial i} a(t_i^{(j)} - t_l^{(i)}) \right] \\ &= \gamma(t_i^{(j)}) \left(\prod_{l \in \partial i} a(t_i^{(j)} - t_l^{(i)} - 1) - \phi(t_i^{(j)}) \prod_{l \in \partial i} a(t_i^{(j)} - t_l^{(i)}) \right), \end{aligned} \quad (\text{A3})$$

where we have defined

$$\begin{aligned} a(t) &= (1 - \lambda)^{tH(t)}, \\ \gamma(t) &= \begin{cases} \gamma & \text{if } t = 0 \\ 1 - \gamma & \text{if } t > 0 \end{cases}, \\ \phi(t) &= \begin{cases} 0 & \text{if } t = 0 \text{ or } t = T + 1 \\ 1 & \text{if } 0 < t < T + 1 \end{cases}, \end{aligned} \quad (\text{A4})$$

where $H(t)$ is the Heaviside step function, with $H(0) = 0$. We also notice that the function Ψ constrains the planted and inferred times of the incoming messages to the equality: $\tau_i^{(k)} = \tau_i^{(j)}$, and $t_i^{(k)} = t_i^{(j)}$ for all $k \in \partial i \setminus j$. We can now rewrite the first BP




FIG. 15. The compact optimized version of BP equations. The μ messages are functions of αT^2 values, since $\sigma \in \{0, 1, 2\}$ and $c \in \{0, 1\}$. For this reason we keep the population of μ messages. Each iteration of the optimized BP consists in computing ν message from a set of μ messages and after the extraction of disorder, according to Eq. (A13). Then from the ν message, by performing the summation on the last argument described in Eq. (A14), the new μ message is obtained.

2. Simplifications

To simplify the BP Eqs. (19), we will start by writing the functions ψ^* , ψ in a simplified way:

equation with the expression of ψ^* , ψ :

$$\begin{aligned}
 \nu_{\Psi_i \rightarrow j}(T_{ij}) = & \frac{\gamma(t_i^{(j)}) \xi(\tau_i^{(j)}, t_i^{(j)}, c_i)}{Z_{\Psi_i \rightarrow j}} \left\{ a(t_i^{(j)} - t_j^{(i)} - 1) \delta_{x_i^0, I} \delta_{\tau_i^{(j)}, 0} \prod_{k \in \partial i \setminus j} \left[\sum_{t_k^{(i)}} a(t_i^{(j)} - t_k^{(i)} - 1) \sum_{\tau_k^{(i)}} \mu_{k \rightarrow \Psi_i}(T_{ki}) \right] \right. \\
 & + a(t_i^{(j)} - t_j^{(i)} - 1) \delta_{x_i^0, S} \mathbb{I}[\tau_i^{(j)} \leq \tau_j^{(j)} + s_{ji}] \prod_{k \in \partial i \setminus j} \left[\sum_{t_k^{(i)}} a(t_i^{(j)} - t_k^{(i)} - 1) \sum_{\tau_k^{(i)}} \mu_{k \rightarrow \Psi_i}(T_{ki}) \mathbb{I}[\tau_i^{(j)} \leq \tau_k^{(i)} + s_{ki}] \right] \\
 & - a(t_i^{(j)} - t_j^{(i)} - 1) \delta_{x_i^0, S} \mathbb{I}[\tau_i^{(j)} < T + 1] \mathbb{I}[\tau_i^{(j)} < \tau_j^{(j)} + s_{ji}] \\
 & \times \prod_{k \in \partial i \setminus j} \left[\sum_{t_k^{(i)}} a(t_i^{(j)} - t_k^{(i)} - 1) \sum_{\tau_k^{(i)}} \mu_{k \rightarrow \Psi_i}(T_{ki}) \mathbb{I}[\tau_i^{(j)} < \tau_k^{(i)} + s_{ki}] \right] \\
 & - \phi(t_i^{(j)}) a(t_i^{(j)} - t_j^{(i)}) \delta_{x_i^0, I} \delta_{\tau_i^{(j)}, 0} \prod_{k \in \partial i \setminus j} \left[\sum_{t_k^{(i)}} a(t_i^{(j)} - t_k^{(i)}) \sum_{\tau_k^{(i)}} \mu_{k \rightarrow \Psi_i}(T_{ki}) \right] \\
 & - \phi(t_i^{(j)}) a(t_i^{(j)} - t_j^{(i)}) \delta_{x_i^0, S} \mathbb{I}[\tau_i^{(j)} \leq \tau_j^{(j)} + s_{ji}] \prod_{k \in \partial i \setminus j} \left[\sum_{t_k^{(i)}} a(t_i^{(j)} - t_k^{(i)}) \sum_{\tau_k^{(i)}} \mu_{k \rightarrow \Psi_i}(T_{ki}) \mathbb{I}[\tau_i^{(j)} \leq \tau_k^{(i)} + s_{ki}] \right] \\
 & + \phi(t_i^{(j)}) a(t_i^{(j)} - t_j^{(i)}) \delta_{x_i^0, S} \mathbb{I}[\tau_i^{(j)} < T + 1] \mathbb{I}[\tau_i^{(j)} < \tau_j^{(j)} + s_{ji}] \\
 & + \phi(t_i^{(j)}) a(t_i^{(j)} - t_j^{(i)}) \delta_{x_i^0, S} \mathbb{I}[\tau_i^{(j)} < T + 1] \mathbb{I}[\tau_i^{(j)} < \tau_j^{(j)} + s_{ji}] \\
 & \left. \times \prod_{k \in \partial i \setminus j} \left[\sum_{t_k^{(i)}} a(t_i^{(j)} - t_k^{(i)}) \sum_{\tau_k^{(i)}} \mu_{k \rightarrow \Psi_i}(T_{ki}) \mathbb{I}[\tau_i^{(j)} < \tau_k^{(i)} + s_{ki}] \right] \right\}, \tag{A5}
 \end{aligned}$$

where $T_{ki} = (\tau_k^{(i)}, \tau_i^{(k)} = \tau_i^{(j)}, t_k^{(i)}, t_i^{(k)} = t_i^{(j)})$ in the r.h.s., due to the constraint on the incoming times (and $T_{ij} = (\tau_i^{(j)}, \tau_j^{(i)}, t_i^{(j)}, t_j^{(i)})$ in the l.h.s.).

3. Summation over the planted times

We can see on the above equation that the r.h.s. depends on the planted time $\tau_j^{(i)}$ only through the sign:

$$\sigma_{ji} = 1 + \text{sgn}(\tau_j^{(i)} - \tau_i^{(j)} + s_{ji}), \tag{A6}$$

with the convention that $\text{sgn}(0) = 0$. We therefore introduce the notation

$$\tilde{\nu}_{\Psi_i \rightarrow j}(\tau_i^{(j)}, \sigma_{ji}, t_i^{(j)}, t_j^{(i)}) = \nu_{\Psi_i \rightarrow j}(\tau_i^{(j)}, \tau_j^{(i)}, t_i^{(j)}, t_j^{(i)}) \tag{A7}$$

for all $\tau_j^{(i)}$ such that $\sigma_{ji} = 1 + \text{sgn}(\tau_j^{(i)} - \tau_i^{(j)} + s_{ji})$. We also introduce the message

$$\tilde{\mu}_{i \rightarrow \Psi_j}(\sigma_{ij}, \tau_j^{(i)}, t_i^{(j)}, t_j^{(i)}) = \sum_{\tau_i^{(j)}} \mu_{i \rightarrow \Psi_j}(\tau_i^{(j)}, \tau_j^{(i)}, t_i^{(j)}, t_j^{(i)}) \mathbb{I}[\sigma_{ij} = 1 + \text{sgn}(\tau_i^{(j)} - \tau_j^{(i)} + s_{ij})]. \tag{A8}$$

With these definitions, the BP equation becomes

$$\begin{aligned}
 \tilde{\nu}_{\Psi_i \rightarrow j}(\tilde{T}_{ij}) = & \gamma(t_i^{(j)}) \xi(\tau_i^{(j)}, t_i^{(j)}, c_i) \left\{ a(t_i^{(j)} - t_j^{(i)} - 1) \delta_{x_i^0, I} \delta_{\tau_i^{(j)}, 0} \prod_{k \in \partial i \setminus j} \left[\sum_{t_k^{(i)}} a(t_i^{(j)} - t_k^{(i)} - 1) \sum_{\sigma_{ki}=0}^2 \tilde{\mu}_{k \rightarrow \Psi_i}(\tilde{T}_{ki}) \right] \right. \\
 & + a(t_i^{(j)} - t_j^{(i)} - 1) \delta_{x_i^0, S} \mathbb{I}[\sigma_{ji} \in \{1, 2\}] \prod_{k \in \partial i \setminus j} \left[\sum_{t_k^{(i)}} a(t_i^{(j)} - t_k^{(i)} - 1) \sum_{\sigma_{ki}=1}^2 \tilde{\mu}_{k \rightarrow \Psi_i}(\tilde{T}_{ki}) \right]
 \end{aligned}$$

$$\begin{aligned}
 & - a(t_i^{(j)} - t_j^{(i)} - 1) \delta_{x_i^0, S} \mathbb{I}[\tau_i^{(j)} < T + 1] \mathbb{I}[\sigma_{ji} = 2] \prod_{k \in \partial i \setminus j} \left[\sum_{t_k^{(i)}} a(t_i^{(j)} - t_k^{(i)} - 1) \tilde{\mu}_{k \rightarrow \Psi_i}(\sigma_{ki} = 2, \tau_i^{(j)}, t_k^{(i)}, t_i^{(j)}) \right] \\
 & - \phi(t_i^{(j)}) a(t_i^{(j)} - t_j^{(i)}) \delta_{x_i^0, I} \delta_{\tau_i^{(j)}, 0} \prod_{k \in \partial i \setminus j} \left[\sum_{t_k^{(i)}} a(t_i^{(j)} - t_k^{(i)}) \sum_{\sigma_{ki}=0}^2 \tilde{\mu}_{k \rightarrow \Psi_i}(\tilde{T}_{ki}) \right] \\
 & - \phi(t_i^{(j)}) a(t_i^{(j)} - t_j^{(i)}) \delta_{x_i^0, S} \mathbb{I}[\sigma_{ji} \in \{1, 2\}] \prod_{k \in \partial i \setminus j} \left[\sum_{t_k^{(i)}} a(t_i^{(j)} - t_k^{(i)}) \sum_{\sigma_{ki}=1}^2 \tilde{\mu}_{k \rightarrow \Psi_i}(\tilde{T}_{ki}) \right] \\
 & + \phi(t_i^{(j)}) a(t_i^{(j)} - t_j^{(i)}) \delta_{x_i^0, S} \mathbb{I}[\tau_i^{(j)} < T + 1] \mathbb{I}[\sigma_{ji} = 2] \prod_{k \in \partial i \setminus j} \left[\sum_{t_k^{(i)}} a(t_i^{(j)} - t_k^{(i)}) \tilde{\mu}_{k \rightarrow \Psi_i}(\sigma_{ki} = 2, \tau_i^{(j)}, t_k^{(i)}, t_i^{(j)}) \right] \Bigg\}, \tag{A9}
 \end{aligned}$$

where $\tilde{T}_{ij} = (\tau_i^{(j)}, \sigma_{ji}, t_i^{(j)}, t_j^{(i)})$ and $\tilde{T}_{ki} = (\sigma_{ki}, \tau_i^{(k)} = \tau_i^{(j)}, t_k^{(i)} = t_i^{(j)})$ for all $k \in \partial i \setminus j$. In the above equation we have dropped the normalization factor $z_{\Psi_i \rightarrow j}$, since the message $\tilde{v}_{\Psi_i \rightarrow j}$ is not a probability but the value taken by the (normalized) BP message $v_{\Psi_i \rightarrow j}$ for any $\tau_j^{(i)}$ achieving the equality (A6). The other BP equation becomes

$$\begin{aligned}
 \tilde{\mu}_{i \rightarrow \Psi_j}(\sigma_{ij}, \tau_j^{(i)}, t_i^{(j)}, t_j^{(i)}) &= \sum_{\tau_i^{(j)}=0}^{T+1} \mu_{i \rightarrow \Psi_j}(\tau_i^{(j)}, \tau_j^{(i)}, t_i^{(j)}, t_j^{(i)}) \mathbb{I}[\sigma_{ij} = 1 + \text{sgn}(\tau_i^{(j)} - \tau_j^{(i)} + s_{ij})] \\
 &= \sum_{\tau_i^{(j)}=0}^{T+1} v_{\Psi_i \rightarrow j}(\tau_i^{(j)}, \tau_j^{(i)}, t_i^{(j)}, t_j^{(i)}) \mathbb{I}[\sigma_{ij} = 1 + \text{sgn}(\tau_i^{(j)} - \tau_j^{(i)} + s_{ij})] \\
 &= \sum_{\tau_i^{(j)}=0}^{T+1} \tilde{v}_{\Psi_i \rightarrow j}(\tau_i^{(j)}, \sigma_{ji} = 1 + \text{sgn}(\tau_i^{(j)} - \tau_j^{(i)} + s_{ij}), t_i^{(j)}, t_j^{(i)}) \mathbb{I}[\sigma_{ij} = 1 + \text{sgn}(\tau_i^{(j)} - \tau_j^{(i)} + s_{ij})], \tag{A10}
 \end{aligned}$$

which gives for each value of σ_{ij}

$$\begin{aligned}
 \tilde{\mu}_{i \rightarrow \Psi_j}(0, \tau_j^{(i)}, t_i^{(j)}, t_j^{(i)}) &= \mathbb{I}[\tau_j - s_{ji} > 0] \sum_{\tau_i^{(j)}=0}^{\tau_j^{(i)} - s_{ji}} \tilde{v}_{\Psi_i \rightarrow j}(\tau_i^{(j)}, \sigma_{ji} = 2, t_i^{(j)}, t_j^{(i)}), \\
 \tilde{\mu}_{i \rightarrow \Psi_j}(1, \tau_j^{(i)}, t_i^{(j)}, t_j^{(i)}) &= \mathbb{I}[\tau_j - s_{ji} \geq 0] \tilde{v}_{\Psi_i \rightarrow j}(\tau_i^{(j)} = \tau_j^{(i)} - s_{ji}, \sigma_{ji} = 2, t_i^{(j)}, t_j^{(i)}), \\
 \tilde{\mu}_{i \rightarrow \Psi_j}(2, \tau_j^{(i)}, t_i^{(j)}, t_j^{(i)}) &= \sum_{\tau_i^{(j)}=\zeta_i^+}^{T+1} \tilde{v}_{\Psi_i \rightarrow j}(\tau_i^{(j)}, \sigma_{ji} = 1 + \text{sgn}(\tau_i^{(j)} - \tau_j^{(i)} + s_{ji}), t_i^{(j)}, t_j^{(i)}) \\
 &= \sum_{\tau_i^{(j)}=\zeta_i^+}^{\zeta_i^-} \tilde{v}_{\Psi_i \rightarrow j}(\tau_i^{(j)}, \sigma_{ji} = 2, t_i^{(j)}, t_j^{(i)}) \\
 &\quad + \mathbb{I}[\tau_j^{(i)} + s_{ji} \leq T + 1] \tilde{v}_{\Psi_i \rightarrow j}(\tau_i^{(j)} = \tau_j^{(i)} + s_{ji}, \sigma_{ji} = 1, t_i^{(j)}, t_j^{(i)}) \\
 &\quad + \mathbb{I}[\tau_j^{(i)} + s_{ji} < T + 1] \sum_{\tau_i^{(j)}=\tau_j^{(i)}+s_{ji}+1}^{T+1} \tilde{v}_{\Psi_i \rightarrow j}(\tau_i^{(j)}, \sigma_{ji} = 0, t_i^{(j)}, t_j^{(i)}), \tag{A11}
 \end{aligned}$$

where $\zeta_i^+ = \max(0, \tau_j^{(i)} - s_{ji} + 1)$ and $\zeta_i^- = \min(T + 1, \tau_j^{(i)} + s_{ji} - 1)$.

4. Summation over the inferred times

To reduce further the space of variables over which the BP messages are defined, we define the following message:

$$\mu'_{i \rightarrow \Psi_j}(\sigma_{ij}, \tau_j^{(i)}, c_{ij}, t_j^{(i)}) = \sum_{t_j^{(j)}} \tilde{\mu}_{i \rightarrow \Psi_j}(\sigma_{ij}, \tau_j^{(i)}, t_j^{(j)}, t_j^{(i)}) a(t_j^{(i)} - t_j^{(j)} - c_{ij}), \quad (\text{A12})$$

with $c_{ij} \in \{0, 1\}$. Using this definition, the first BP equation becomes

$$\begin{aligned} \tilde{\nu}_{\Psi_i \rightarrow j}(\tau_i^{(j)}, \sigma_{ji}, t_i^{(j)}, t_j^{(i)}) &= \gamma(t_i^{(j)}) \xi(\tau_i^{(j)}, t_i^{(j)}, c_i) \left\{ a(t_i^{(j)} - t_j^{(i)} - 1) \delta_{x_i^0, I} \delta_{\tau_i^{(j)}, 0} \prod_{k \in \partial i \setminus j} \left[\sum_{\sigma_{ki}=0}^2 \mu'_{k \rightarrow \Psi_i}(\sigma_{ki}, \tau_i^{(k)}, c_{ki} = 1, t_i^{(k)}) \right] \right. \\ &+ a(t_i^{(j)} - t_j^{(i)} - 1) \delta_{x_i^0, S} \mathbb{I}[\sigma_{ji} \in \{1, 2\}] \prod_{k \in \partial i \setminus j} \left[\sum_{\sigma_{ki}=1}^2 \mu'_{k \rightarrow \Psi_i}(\sigma_{ki}, \tau_i^{(k)}, c_{ki} = 1, t_i^{(k)}) \right] \\ &- a(t_i^{(j)} - t_j^{(i)} - 1) \delta_{x_i^0, S} \mathbb{I}[\tau_i^{(j)} < T + 1] \mathbb{I}[\sigma_{ji} = 2] \prod_{k \in \partial i \setminus j} \mu'_{k \rightarrow \Psi_i}(\sigma_{ki} = 2, \tau_i^{(k)}, c_{ki} = 1, t_i^{(k)}) \\ &- \phi(t_i^{(j)}) a(t_i^{(j)} - t_j^{(i)}) \delta_{x_i^0, I} \delta_{\tau_i^{(j)}, 0} \prod_{k \in \partial i \setminus j} \left[\sum_{\sigma_{ki}=0}^2 \mu'_{k \rightarrow \Psi_i}(\sigma_{ki}, \tau_i^{(k)}, c_{ki} = 0, t_i^{(k)}) \right] \\ &- \phi(t_i^{(j)}) a(t_i^{(j)} - t_j^{(i)}) \delta_{x_i^0, S} \mathbb{I}[\sigma_{ji} \in \{1, 2\}] \prod_{k \in \partial i \setminus j} \left[\sum_{\sigma_{ki}=1}^2 \mu'_{k \rightarrow \Psi_i}(\sigma_{ki}, \tau_i^{(k)}, c_{ki} = 0, t_i^{(k)}) \right] \\ &+ \left. \phi(t_i^{(j)}) a(t_i^{(j)} - t_j^{(i)}) \delta_{x_i^0, S} \mathbb{I}[\tau_i^{(j)} < T + 1] \mathbb{I}[\sigma_{ji} = 2] \prod_{k \in \partial i \setminus j} \mu'_{k \rightarrow \Psi_i}(\sigma_{ki} = 2, \tau_i^{(k)}, c_{ki} = 0, t_i^{(k)}) \right\}. \quad (\text{A13}) \end{aligned}$$

The second BP equation becomes

$$\begin{aligned} \mu'(0, \tau_j^{(i)}, c_{ij}, t_j^{(i)}) &= \mathbb{I}[\tau_j - s_{ji} > 0] \sum_{t_j^{(j)}} a(t_j^{(i)} - t_j^{(j)} - c_{ij}) \sum_{\tau_i^{(j)}=0}^{\tau_j^{(i)} - s_{ji}} \tilde{\nu}_{\Psi_i \rightarrow j}(\tau_i^{(j)}, \sigma_{ji} = 2, t_i^{(j)}, t_j^{(i)}), \\ \mu'(1, \tau_j^{(i)}, c_{ij}, t_j^{(i)}) &= \mathbb{I}[\tau_j - s_{ji} \geq 0] \sum_{t_j^{(j)}} a(t_j^{(i)} - t_j^{(j)} - c_{ij}) \tilde{\nu}_{\Psi_i \rightarrow j}(\tau_i^{(j)} = \tau_j^{(i)} - s_{ji}, \sigma_{ji} = 2, t_i^{(j)}, t_j^{(i)}), \\ \mu'(2, \tau_j^{(i)}, c_{ij}, t_j^{(i)}) &= \sum_{t_j^{(j)}} a(t_j^{(i)} - t_j^{(j)} - c_{ij}) \left[\sum_{\tau_i^{(j)}=\xi_i^+}^{\xi_i^-} \tilde{\nu}_{\Psi_i \rightarrow j}(\tau_i^{(j)}, \sigma_{ji} = 2, t_i^{(j)}, t_j^{(i)}) \right. \\ &+ \mathbb{I}[\tau_j^{(i)} + s_{ji} \leq T + 1] \tilde{\nu}_{\Psi_i \rightarrow j}(\tau_i^{(j)} = \tau_j^{(i)} + s_{ji}, \sigma_{ji} = 1, t_i^{(j)}, t_j^{(i)}) \\ &+ \left. \mathbb{I}[\tau_j^{(i)} + s_{ji} < T + 1] \sum_{\tau_i^{(j)}=\tau_j^{(i)}+s_{ji}+1}^{T+1} \tilde{\nu}_{\Psi_i \rightarrow j}(\tau_i^{(j)}, \sigma_{ji} = 0, t_i^{(j)}, t_j^{(i)}) \right]. \quad (\text{A14}) \end{aligned}$$

5. BP marginals

Once a fixed-point of the BP Eqs. (A13) and (A14) is found, the BP marginal can be expressed as

$$\begin{aligned} P_i(\tau_i, t_i) &= \sum_{\underline{\tau}_{\partial i}, \underline{t}_{\partial i}} P_{\Psi_i}(\tau_i, t_i, \underline{\tau}_{\partial i}, \underline{t}_{\partial i}) \\ &= \frac{1}{Z_{\Psi_i}} \sum_{\underline{\tau}_{\partial i}, \underline{t}_{\partial i}} \xi(\tau_i, t_i; c_i) \psi^*(\tau_i, \underline{\tau}_{\partial i}; \{s_{li}\}_{l \in \partial i}, x_i^0) \psi(\tau_i, \underline{t}_{\partial i}) \prod_{l \in \partial i} \mu_{l \rightarrow \Psi_i}(\tau_l, \tau_i, t_l, t_i) \\ &= \frac{1}{Z_{\Psi_i}} \gamma(t_i) \xi(\tau_i, t_i; c_i) \left(\delta_{x_i^0, I} \delta_{\tau_i, 0} \prod_{l \in \partial i} \left[\sum_{\sigma_{li}=0}^2 \mu'_{l \rightarrow \Psi_i}(\sigma_{li}, \tau_i, c_{li} = 1, t_i) \right] \right) \end{aligned}$$

$$\begin{aligned}
 & + \delta_{x_i^0, S} \prod_{l \in \partial i} \left[\sum_{\sigma_{li}=1}^2 \mu'_{l \rightarrow \psi_i}(\sigma_{li}, \tau_i, c_{li} = 1, t_i) \right] - \delta_{x_i^0, S} \mathbb{I}[\tau_i < T + 1] \prod_{l \in \partial i} \mu_{l \rightarrow \psi_i}(\sigma_{li} = 2, \tau_i, c_{li} = 1, t_i) \\
 & - \phi(t_i) \delta_{x_i^0, T} \delta_{\tau_i, 0} \prod_{l \in \partial i} \left[\sum_{\sigma_{li}=0}^2 \mu'_{l \rightarrow \psi_i}(\sigma_{li}, \tau_i, c_{li} = 0, t_i) \right] - \phi(t_i) \delta_{x_i^0, S} \prod_{l \in \partial i} \left[\sum_{\sigma_{li}=1}^2 \mu'_{l \rightarrow \psi_i}(\sigma_{li}, \tau_i, c_{li} = 0, t_i) \right] \\
 & + \phi(t_i) \delta_{x_i^0, S} \mathbb{I}[\tau_i < T + 1] \prod_{l \in \partial i} \mu'_{l \rightarrow \psi_i}(\sigma_{li} = 2, \tau_i, c_{li} = 0, t_i). \tag{A15}
 \end{aligned}$$

6. Bethe Free Energy

The Bethe free energy is written

$$F = - \sum_{i \in V} \log Z_{\psi_i} + \frac{1}{2} \sum_{i \in V} \sum_{j \in \partial i} \log Z_{ij}, \tag{A16}$$

where Z_{ψ_i} is the normalization of the BP marginal written above, and with

$$\begin{aligned}
 Z_{ij} & = \sum_{T_{ij}} \nu_{\psi_i \rightarrow j}(T_{ij}) \nu_{\psi_j \rightarrow i}(T_{ij}) \\
 & = \sum_{T_{ij}} \nu_{\psi_i \rightarrow j}(T_{ij}) \mu_{j \rightarrow \psi_i}(T_{ij}) \\
 & = \frac{1}{z_{\psi_i \rightarrow j}} \sum_{\{T_{il}\}_{l \in \partial i}} \Psi(\{T_{il}\}_{l \in \partial i}) \prod_{l \in \partial i} \mu_{l \rightarrow \psi_i}(T_{il}) \\
 & = \frac{Z_{\psi_i}}{z_{\psi_i \rightarrow j}}, \tag{A17}
 \end{aligned}$$

where $z_{\psi_i \rightarrow j}$ is the normalization of the BP message $\nu_{\psi_i \rightarrow j}$:

$$\begin{aligned}
 z_{\psi_i \rightarrow j} & = \sum_{\{T_{il}\}_{l \in \partial i}} \Psi(\{T_{il}\}_{l \in \partial i}) \prod_{k \in \partial i \setminus j} \mu_{k \rightarrow \psi_i}(T_{ik}) \\
 & = \sum_{\tau_i^{(j)}, \tau_j^{(i)}, t_i^{(j)}, t_j^{(i)}} \tilde{\nu}(\tau_i^{(j)}, \sigma_{ji} = 1 + \text{sgn}(\tau_j^{(i)} + s_{ji} - \tau_i^{(j)}), t_i^{(j)}, t_j^{(i)}), \tag{A18}
 \end{aligned}$$

where $\tilde{\nu}$ is the unnormalized message defined in Eq. (A7). We obtain an expression of the free-energy in terms of the normalizations Z_{ψ_i} , $z_{\psi_i \rightarrow j}$:

$$F = \frac{1}{N} \sum_{i \in V} \left(\frac{d_i}{2} - 1 \right) \log Z_{\psi_i} - \frac{1}{N} \frac{1}{2} \sum_{i \in V} \sum_{j \in \partial i} \log z_{\psi_i \rightarrow j}. \tag{A19}$$

7. Entropy and Energy

To compute the entropy it is sufficient to subtract energy and free energy:

$$S = U - F.$$

The energy is simply the average of the Hamiltonian:

$$\begin{aligned}
 H & = - \sum_i \log \psi - \sum_i \log \xi - \sum_i \log \psi^*, \\
 U & = - \sum_i \langle \log \psi \rangle - \sum_i \langle \log \xi \rangle - \sum_i \langle \log \psi^* \rangle \\
 & = - \sum_i \langle \log \psi \rangle.
 \end{aligned}$$

$$U = - \sum_i \frac{1}{Z_{\psi_i}} u_{\psi_i},$$

with

$$u_{\Psi_i} := \sum_{\{T_{ij}\}_{j \in \partial i}} \prod_{j \in \partial i} m_{j \rightarrow \Psi_i}(T_{ij}) \psi(t_i, t_{\partial i}) \psi^*(\tau_i, \tau_{\partial i}) \xi(t_i, \tau_i) \log \psi(t_i, t_{\partial i}).$$

Comparing this formula with the expression for Z_{Ψ_i} ,

$$Z_{\Psi_i} = \sum_{\{T_{ij}\}_{j \in \partial i}} \prod_{j \in \partial i} m_{j \rightarrow \Psi_i}(T_{ij}) \psi^*(\tau_i, \tau_{\partial i}) \xi(t_i, \tau_i) \psi(t_i, t_{\partial i}),$$

we see that the computation of the energy requires similar calculations to the ones already performed to compute free energy. The only additional difficulty is in the log ψ factor. Let us first trace over the planted variables. We keep τ_i fixed:

$$\begin{aligned} \sum_{\tau_{\partial i}} \prod_{j \in \partial i} m_{j \rightarrow \Psi_i}(T_{ij}) \psi^*(\tau_i, \tau_{\partial i}) &= \delta_{x_{i,0}^*, I} \delta_{\tau_i, 0} \prod_{j \in \partial i} \sum_{\tau_j} m_{\Psi_j \rightarrow i}(t_j, t_i, \tau_j, \tau_i) \\ &\quad + \delta_{x_{i,0}^*, S} \prod_{j \in \partial i} \sum_{\tau_j} m_{\Psi_j \rightarrow i}(t_j, t_i, \tau_j, \tau_i) \mathbb{I}[\tau_i \leq \tau_j + s_{ji}] \\ &\quad - \mathbb{I}[\tau_i \leq T] \delta_{x_{i,0}^*, S} \prod_{j \in \partial i} \sum_{\tau_j} m_{\Psi_j \rightarrow i}(t_j, t_i, \tau_j, \tau_i) \mathbb{I}[\tau_i < \tau_j + s_{ji}] \\ &=: \sum_{v=1}^3 \prod_{j \in \partial i} m_{j \rightarrow \Psi_i}^v(t_i, t_j, \tau_i). \end{aligned}$$

So we have

$$Z_{\Psi_i} = \sum_{t_i, \tau_i, \{t_j\}_{j \in \partial i}} \sum_{v=1}^3 f_v^{\tau_i, x_{i,0}^*} \prod_{j \in \partial i} m_{j \rightarrow \Psi_i}^v(t_i, t_j, \tau_i) \xi(t_i, \tau_i) \psi(t_i, t_{\partial i}).$$

Now remember that the original message m and the compressed factor to node message v are related by

$$\nu_{\Psi_i \rightarrow j}(t_i, t_j, \tau_i, 1 + \text{sign}(\tau_j - \tau_i + s_{ji})) = m_{\Psi_i \rightarrow j}(t_i, t_j, \tau_i, \tau_j).$$

So the sums we have to compute are

$$\begin{aligned} m_{j \rightarrow \Psi_i}^0(t_i, t_j, 0) &= \sum_{\tau_j} m_{\Psi_j \rightarrow i}(t_j, t_i, \tau_j, 0) \\ &= \sum_{\tau_j} \nu_{\Psi_j \rightarrow i}(t_j, t_i, \tau_j, 1 + \text{sign}(-\tau_j + s_{ij})), \\ m_{j \rightarrow \Psi_i}^1(t_i, t_j, \tau_i) &= \sum_{\tau_j \geq \tau_i - s_{ji}} m_{\Psi_j \rightarrow i}(t_j, t_i, \tau_j, \tau_i) \\ &= \sum_{\tau_j \geq \tau_i - s_{ji}} \nu_{\Psi_j \rightarrow i}(t_j, t_i, \tau_j, 1 + \text{sign}(\tau_i - \tau_j + s_{ij})), \\ m_{j \rightarrow \Psi_i}^2(t_i, t_j, \tau_i) &= \sum_{\tau_j > \tau_i - s_{ji}} m_{\Psi_j \rightarrow i}(t_j, t_i, \tau_j, \tau_i) \\ &= \sum_{\tau_j > \tau_i - s_{ji}} \nu_{\Psi_j \rightarrow i}(t_j, t_i, \tau_j, 1 + \text{sign}(\tau_i - \tau_j + s_{ij})). \end{aligned}$$

Notice that

$$m_{j \rightarrow \Psi_i}^1(t_i, t_j, \tau_i) = m_{j \rightarrow \Psi_i}^2(t_i, t_j, \tau_i) + \nu_{\Psi_j \rightarrow i}(t_j, t_i, \tau_i - s_{ji}, 2).$$

Now it is time to deal with planted times. Let us observe that

$$\psi(t_i, t_{\partial i}) = \gamma(t_i) (1 - \lambda)^{S_1} [1 - (1 \leq t_i \leq T)(1 - \lambda)^{S_2}],$$

where

$$\begin{aligned} \gamma(t_i) &= \gamma \delta_{t_i, 0} + (1 - \gamma)(1 - \delta_{t_i, 0}), \\ S_1 &:= \sum_{j \in \partial i} (t_i - t_j - 1)_+, \end{aligned}$$

$$S_2 := \sum_{j \in \partial i} \theta(t_i - t_j - 1).$$

We want to find the BP distribution of t_i, S_1, S_2 to average over $\psi \log \psi$. We define

$$F_k^{v, \tau_i}(t_i, S_1, S_2) := \sum_{\{t_j\}_{j \leq k}} \prod_{j \leq k} m_{j \rightarrow \psi_i}^v(t_i, t_j, \tau_i) \delta_{S_1, \sum_{j \leq k} (t_i - t_j - 1)_+} \delta_{S_2, \sum_{j \leq k} \theta(t_i - t_j - 1)}.$$

Therefore,

$$\begin{aligned} F_{k+1}^{v, \tau_i}(t_i, S_1, S_2) &= \sum_{\{t_j\}_{j \leq k+1}} \prod_{j \leq k+1} m_{j \rightarrow \psi_i}^v(t_i, t_j, \tau_i) \delta_{S_1, \sum_{j \leq k+1} (t_i - t_j - 1)_+} \delta_{S_2, \sum_{j \leq k+1} \theta(t_i - t_j - 1)} \\ &= \sum_{t_{k+1}} m_{k+1 \rightarrow \psi_i}^v(t_i, t_{k+1}, \tau_i) \sum_{\{t_j\}_{j \leq k}} \prod_{j \leq k} m_{j \rightarrow \psi_i}^v(t_i, t_j, \tau_i) \\ &\quad \times \delta_{S_1 - (t_i - t_{k+1} - 1)_+, \sum_{j \leq k} (t_i - t_j - 1)_+} \delta_{S_2 - \theta(t_i - t_{k+1} - 1), \sum_{j \leq k} \theta(t_i - t_j - 1)} \\ &= \sum_{t_{k+1}} m_{k+1 \rightarrow \psi_i}^v(t_i, t_{k+1}, \tau_i) F_k^{v, \tau_i}(t_i, S_1 - (t_i - t_{k+1} - 1)_+, S_2 - \theta(t_i - t_{k+1} - 1)) \end{aligned}$$

and

$$F_0^{v, \tau_i}(t_i, S_1, S_2) = \delta_{S_1, 0} \delta_{S_2, 0},$$

and for our purposes we want to find $F_{|\partial i|}^{v, \tau_i}(t_i, S_1, S_2)$. Now we have an iterative scheme to compute the measure. Once the function is found we simply have

$$\begin{aligned} u_{\psi_i} &= \sum_{t_i, \tau_i; \xi(t_i, \tau_i)=1} \sum_{v=1}^3 f_v^{\tau_i, \tau_i^*, 0} \sum_{S_1, S_2} F_{|\partial i|}^{v, \tau_i}(t_i, S_1, S_2) \psi(t_i, S_1, S_2) \\ &\quad \times \log \psi(t_i, S_1, S_2). \end{aligned}$$

APPENDIX B: REPLICASYMMETRIC FORMALISM

The aim of the cavity method is to characterize the typical properties of the probability measure (17) that we recall here

$$P(\{T_{ij}\}_{(ij) \in E} | \mathcal{D}) = \frac{1}{Z(\mathcal{D})} \prod_{i \in V} \Psi(\{T_{il}\}_{l \in \partial i}; \mathcal{D}_i),$$

for typical random graphs and for typical realization of the disorder $\mathcal{D} = \{\mathcal{D}_i\}$, in the thermodynamic limit $N \rightarrow \infty$. In the simplest version of the cavity method, called replica symmetric (RS), one assumes a fast decay of the correlations between distant variables in the measure (17), in such a way that the BP Eqs. (19) converge to a unique fixed-point on a typical large instance, and that the measure (17) is well described by the locally treelike approximation. We consider a uniformly chosen edge $(ij) \in E$ in a random contact graph \mathcal{G} , and call \mathcal{P}^{RS} the probability law of the fixed-point BP message $\mu_{i \rightarrow \psi_j}$ thus observed. Within the decorrelation hypothesis of the RS cavity method, the incoming messages on a given factor node are *i.i.d.* with probability \mathcal{P}^{RS} . This implies that the probability law \mathcal{P}^{RS} must obey the following self-consistent equation:

$$\begin{aligned} \mathcal{P}^{\text{RS}}(\mu) &= \sum_{d=0}^{\infty} r_d \sum_{\mathcal{D}_i} P(\mathcal{D}_i) \int \prod_{i=1}^d d\mathcal{P}^{\text{RS}}(\mu_i) \\ &\quad \times \delta[\mu - f^{\text{bp}}(\mu_1, \dots, \mu_d; \mathcal{D}_i)], \end{aligned} \quad (\text{B1})$$

where $f^{\text{bp}}(\mu_1, \dots, \mu_d; \mathcal{D}_i)$ is a shorthand notation for the r.h.s. of Eq. (19), and $p(\mathcal{D}_i)$ is distribution of the local disorder $\mathcal{D}_i = \{\{s_{li}\}_{l \in \partial i}, x_i^0, \{\varepsilon_m\}_{i_m=i}\}$ associated with a given node i . We numerically solved these equations with population dynamics. Using the above simplifications, we are left with two types of BP messages: $\mu'_{i \rightarrow \psi_j}$ is defined over the variable $((\sigma_{ij}, \tau_j^{(i)}, c_{ij}, t_j^{(i)}))$ living in a space of size $6(T+1)^2$, and $\tilde{v}_{\psi_i \rightarrow j}$ is defined over the variable $(\tau_i^{(j)}, \sigma_{ji}, t_i^{(j)}, t_j^{(i)})$, living in a space of size $3(T+1)^3$. We store only a population of messages $\mu_{i \rightarrow \psi_j}$, this requires to keep in memory $O(\mathcal{N}T^2)$ numbers, with \mathcal{N} the population size. Computing a new element μ of the population requires in principle $O(T^4)$ operations, but can be reduced to $O(T^3)$ by computing the cumulants of the temporary message v .

1. Replica-symmetric Free Energy

Once averaged over the graph and disorder, the replica-symmetric prediction for the free energy is

$$\begin{aligned} F^{\text{RS}} &= \sum_d P_d \left(\frac{d}{2} - 1 \right) \sum_c p(c) \sum_x \gamma(x) \prod_{i=1}^d w(s_i) \int \prod_{l=1}^d \\ &\quad \times d\mathcal{P}^{\text{RS}}(\mu_l) \log Z_{\Psi_i}(\mu_1, \dots, \mu_d; x, c, s_1, \dots, s_d) \\ &\quad - \frac{d_{\text{av}}}{2} \sum_d r_d \sum_x \gamma(x) \sum_c p_c \sum_{s_1, \dots, s_d} \prod_{k=1}^d w(s_k) \int \prod_{k=1}^d \\ &\quad \times d\mathcal{P}^{\text{RS}}(\mu_k) \log z_{\Psi_i \rightarrow j}(\mu_1, \dots, \mu_d; x, c, s_1, \dots, s_d). \end{aligned} \quad (\text{B2})$$

APPENDIX C: GENERALIZATIONS TO SIR AND SEIR

We describe now, in the framework of the SIR model, a general strategy which can be straightforwardly applied also to the SEIR case. After that, we introduce a parametrization that, for the SIR model, allows to maintain the same number of variables of the SI case in the belief propagation algorithm.

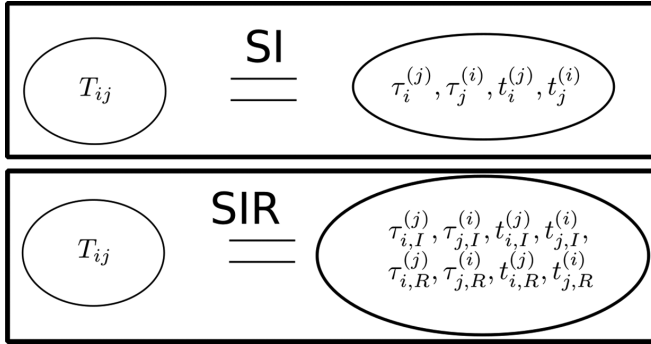


FIG. 16. The super variable for the SIR model includes copies of the planted and inferred recovery time.

1. General strategy

The idea is to simply increase the number of variables. While the SI model trajectory needs only one number per individual to be described, the SIR model needs two: the infection and the recovery times. The first step is therefore to introduce, for each individual i , its infection time t_i^I and its recovery time t_i^R . Then, we can rewrite Eq. (4), defining $\underline{t} = \{t_i^I, t_i^R\}_{i \in V}$:

$$P(\underline{t} | \{x_i^0\}, \{s_{ij}, s_{ji}\}, \{r_i\}) = \prod_{i \in V} \psi^*(t_i^I, t_i^R, \underline{t}_{\partial i}^I, \underline{t}_{\partial i}^R, x_i^0, \{s_{ji}\}_{j \in \partial i}, r_i),$$

where we introduced the new set of *recovery delays* $\{r_i\}_{i \in V}$, which are interpreted as the time interval in which the individual is in the I state. Each ψ^* has an expression which strongly resembles Eq. (5), except for the fact that now an individual $j \in \partial i$ can only infect the individual i before recovering:

$$\psi^* = \mathbb{I}[t_i^I = \delta_{x_i^0, S} \min_{j \in \partial i} f(t_j^I + s_{ji}, t_j^R)] \mathbb{I}[t_i^R = t_i^I + r_i],$$

where

$$f(t, t^R) = \begin{cases} t & \text{if } t < t^R \\ t^R & \text{if } t > t^R. \end{cases}$$

The meaning of f is to state that the tentative infection time from j to i can only be smaller than the recovery time, namely j can not infect i if j has already recovered. Once the ψ^* has




FIG. 17. The optimized factor graph for the SIR model.

been generalized, we can move to modifying Eq. (7),

$$\psi = \sum_{x_i^0, \{s_{ji}\}_{j \in \partial i}, r_i} \left[\gamma(x_i^0) \left(\prod_{j \in \partial i} \omega(s_{ji}) \right) \rho(r_i) \psi^* \times (t_i^I, t_i^R, \underline{t}_{\partial i}^I, \underline{t}_{\partial i}^R, x_i^0, \{s_{ji}\}_{j \in \partial i}, r_i) \right],$$

where ρ is the distribution of recovery delay, $\rho(r_i) = \nu(1 - \nu)^{r_i}$, $\nu \in [0, 1]$. Equation (13) becomes

$$P(\underline{t}, \underline{\tau} | \mathcal{D}) = \prod_i \psi^*(t_i^I, t_i^R, \underline{t}_{\partial i}^I, \underline{t}_{\partial i}^R, x_i^0, \{s_{ji}\}_{j \in \partial i}, r_i) \times \psi(t_i^I, t_i^R, \underline{t}_{\partial i}^I, \underline{t}_{\partial i}^R) \xi(t_i^I, t_i^R, \tau_i^I, \tau_i^R, \{\varepsilon_m\}_{m=i}). \quad (\text{C1})$$

The factor graph associated to this equation, as for Eq. (13), contains loops. Therefore, it is necessary to introduce copies of the infection and the recovery times, exactly as described in the main text for the SI model. To implement the cavity method it is therefore necessary to define each super-variable T_{ij} by including the recovery (planted and inferred) times, as shown in Fig. 16. The factor graph at this point is identical to the one in Fig. 1 and the BP equations are formally the same to Fig. 15, with the messages that however depend on more variables. This generalization can be done also for the SEIR model, simply introducing an exposure time t_i^E , an exposure delay e_i distributed according to $\zeta(e_i)$, and generalizing ψ^* by imposing that the infection can not happen during the exposure time interval.

2. Efficient parametrization for the SIR model

Switching from the infection times to the *transmission* times avoids to increase the number of variables in the BP messages. The idea is that, instead of using copies of the infection time $\{t_i^{(j)}\}_{i \in V}^{j \in \partial i}$, it is convenient to describe the epidemic trajectory with the infection transmission times $\{t_{ij}\}_{i \in V}^{j \in \partial i}$. Each t_{ij} is interpreted as the time at which i tries to infect its neighbor j . To actually know the infection time of j it is therefore necessary to take the minimum of the transmissions:

$$t_j^I = \min_{i \in \partial j} \{t_{ij}\}. \quad (\text{C2})$$

The scheme consists therefore to work with $\{t_{ij}\}_{i \in V}^{j \in \partial i}$, $\{t_i^R\}_{i \in V}$, and $\{t_i^I\}_{i \in V}$. After the new messages converge, the infection variables are found with Eq. (C2). We now show that with this parametrization no copies of the recovery time have to be introduced, differently from the previous paragraph. We define $\underline{t} = \{t_{ij}\}_{i \in V}^{j \in \partial i}$, $\{t_i^I\}_{i \in V}$, $\{t_i^R\}_{i \in V}$ and we have

$$P(\underline{t} | \{x_i^0\}, \{s_{ij}, s_{ji}\}, \{r_i\}) = \prod_{(i,j)} \psi^*(t_{ij}, t_i^I, t_i^R, \{t_{ki}\}_{k \in \partial i}, \underline{t}_{\partial i}^R, x_i^0, s_{ij}, r_i),$$

and each factor is

$$\psi^* = \mathbb{I}[t_{ij} = f(\delta_{x_i^0, S} \min_{k \in \partial i, j} \{t_{ki}\} + s_{ij}, t_i^R)] \mathbb{I}[t_i^R = t_i^I + r_i] \times \mathbb{I}[t_i^I = \min_{k \in \partial i} \{t_{ki}\}]. \quad (\text{C3})$$

The ψ function is always obtained by summing ψ^* over the delays, i.e., x_i^0, s_{ij}, r_i . The factor graph equation is therefore

$$P(\underline{t}, \underline{\tau} | \mathcal{D}) = \prod_i \xi(t_i^I, t_i^R, \tau_i^I, \tau_i^R, \{\varepsilon_m\}_{i_m=i}) \prod_{j \in \partial i} \\ \times (\psi^*(\tau_{ij}, \tau_i^I, \tau_i^R, \{\tau_{ki}\}_{k \in \partial i \setminus j}, \underline{t}_{\partial i}^R, x_i^0, s_{ij}, r_i) \\ \times \psi(t_{ij}, t_i^I, t_i^R, \{t_{ki}\}_{k \in \partial i \setminus j}, \underline{t}_{\partial i}^R)),$$

which, different from Eqs. (13) and (C1), does not contain any loops (see Fig. 17), so it is straightforward to implement BP equations on it: since the recovery times and the infection times are indeed in leaves (nodes attached only to one factor), they can be traced out, so that the BP equations only involve a number of variables which is the same of the SI case. Note that still the BP equations will be a bit slower than the SI case due to the different nature of the factors [Eq. (C3)], which are slower to compute.

-
- [1] N. Antulov-Fantulin, A. Lancic, H. Stefancic, M. Sikic, and T. Smuc, Statistical inference framework for source detection of contagion processes on arbitrary network structures, in *Proceedings of the IEEE 8th International Conference on Self-adaptive and Self-organizing Systems* (IEEE, Piscataway, NJ, 2014), pp. 78–83.
- [2] I. Bestvina and W. Thornton, Infection Model (2023), <https://github.com/ViraTrace/InfectionModel>.
- [3] R. Herbrich, R. Rastogi, and R. Vollgraf, CRISP: A probabilistic model for individual-level COVID-19 infection risk estimation based on contact data, [arXiv:2006.04942](https://arxiv.org/abs/2006.04942), <https://github.com/zalandoresearch/CRISP>.
- [4] P. Gupta, T. Maharaj, M. Weiss, N. Rahaman, H. Alsdurf, N. Minoyan, S. Harnois-Leblanc, J. Merckx, A. Williams, V. Schmidt, P.-L. St-Charles, A. Patel, Y. Zhang, D. L. Buckeridge, C. Pal, B. Schölkopf, and Y. Bengio, Proactive contact tracing, *PLOS Digital Health* **2**, e0000199 (2023).
- [5] F. Altarelli, A. Braunstein, L. Dall’Asta, A. Lage-Castellanos, and R. Zecchina, Bayesian inference of epidemics on networks via belief propagation, *Phys. Rev. Lett.* **112**, 118701 (2014).
- [6] A. Baker, I. Biazzo, A. Braunstein, G. Catania, L. Dall’Asta, A. Ingrassio, F. Krzakala, F. Mazza, M. Mézard, A. P. Muntoni, M. Refinetti, S. S. Mannelli, and L. Zdeborová, Epidemic mitigation by statistical inference from contact tracing data, *Proc. Natl. Acad. Sci. USA* **118**, e2106548118 (2021).
- [7] A. Braunstein and A. Ingrassio, Inference of causality in epidemics on temporal contact networks, *Sci. Rep.* **6**, 27538 (2016).
- [8] D. Ghio, A. L. M. Aragon, I. Biazzo, and L. Zdeborová, Bayes-optimal inference for spreading processes on random networks, *Phys. Rev. E* **108**, 044308 (2023).
- [9] A. Braunstein, G. Catania, L. Dall’Asta, M. Mariani, and A. P. Muntoni, Inference in conditioned dynamics through causality restoration, *Sci. Rep.* **13**, 7350 (2023).
- [10] I. Biazzo, A. Braunstein, L. Dall’Asta, and F. Mazza, A Bayesian generative neural network framework for epidemic inference problems, *Sci. Rep.* **12**, 19673 (2022).
- [11] D. Shah and T. Zaman, Rumors in a network: Who’s the culprit? *IEEE Trans. Inf. Theory* **57**, 5163 (2011).
- [12] P. C. Pinto, P. Thiran, and M. Vetterli, Locating the source of diffusion in large-scale networks, *Phys. Rev. Lett.* **109**, 068702 (2012).
- [13] L. J. S. Allen, Some discrete-time SI, SIR, and SIS epidemic models, *Math. Biosci.* **124**, 83 (1994).
- [14] L. Zdeborová and F. Krzakala, Statistical physics of inference: Thresholds and algorithms, *Adv. Phys.* **65**, 453 (2016).
- [15] M. Mézard and A. Montanari, *Information, Physics, and Computation* (Oxford University Press, Oxford, UK, 2009).
- [16] Y. Iba, The Nishimori line and Bayesian statistics, *J. Phys. A: Math. Gen.* **32**, 3875 (1999).
- [17] G. Parisi, *Statistical Field Theory* (Avalon Publishing, New York, NY, 1998).
- [18] L. E. C. Rocha, F. Liljeros, and P. Holme, Information dynamics shape the sexual networks of Internet-mediated prostitution, *Proc. Natl. Acad. Sci. USA* **107**, 5706 (2010).
- [19] R. A. Rossi and N. K. Ahmed, The network data repository with interactive graph analytics and visualization, in *Proceedings of the AAAI Conference on Artificial Intelligence* (AAAI, New York, 2015).
- [20] <https://github.com/matteomar26/Epidemblem.git>.