Physica A 632 (2023) 129348

Contents lists available at ScienceDirect

Physica A

journal homepage: www.elsevier.com/locate/physa

Accelerating convergence of inference in the inverse Ising problem

Zhongqi Cai *, Enrico Gerding, Markus Brede

School of Electronics and Computer Science, University of Southampton, UK

ARTICLE INFO

Keywords:
Network inference
Inverse kinetic Ising problem
Complex networks
Network control

ABSTRACT

Inferring modelling parameters of dynamical processes from observational data is an important inverse problem in statistical physics. In this paper, instead of passively observing the dynamics for inference, we focus on strategically manipulating dynamics to generate data that gives more accurate estimators within fewer observations. For this purpose, we consider the inference problem rooted in the Ising model with two opposite external fields, assuming that the strength distribution of one of the fields (labelled as passive) is unknown and needs to be inferred. In contrast, the other field (labelled active) is strategically deployed to interact with the Ising dynamics in such a way as to improve the accuracy of estimates of inferring the opposing passive field. By comparing to benchmark cases, we first demonstrate that it is possible to accelerate the inference by strategically interacting with the Ising dynamics. We then apply series expansions to obtain an approximation of the optimized influence configurations in the high-temperature region. Furthermore, by using mean-field estimates, we also demonstrate the applicability of the method in a more general scenario where real-time tracking of the system is infeasible. Last, analysing the optimized influence profiles, we describe heuristics for manipulating the Ising dynamics for faster inference. For example, we show that agents targeted more strongly by the passive field should also be strongly targeted by the active one.

1. Introduction

The inverse statistical problem [1] aims at inferring structural and modelling parameters of complex-networked systems from observed system dynamics. Due to the development of experimental techniques which allow for the accessibility of microscopic-level data, as well as advances in data storage in the last two decades, inverse statistical problems have gained increasing interest in a variety of research domains. Recent applications range from neuroscience [2], computational biology [3], epidemiology [4], financial economics [5], to social science [6]. Even though much effort has been devoted to inverse statistical problems, they remain challenging, especially when only a limited amount of observational data is available [7,8]. Therefore, in this paper, we focus on addressing the inverse statistical problem from the perspective of speeding up the convergence of inference with the aim of obtaining accurate estimations for the model parameters with less data. Our approach contrasts with the aim of most of the prior studies such as Yang et al. [7], Braunstein et al. [8] and Hoang et al. [9], which assume the existence of a given dataset, and focus on methods for inference. Instead, in this paper, we are interested in how data can best be generated. Our principal idea is that we allow strategic interference in the dynamics while observations are being gathered. Below we focus on how to best apply such interference with an aim to generate a more informative dataset that allows accurate inference with less data.

Specifically, in this work, we carry out our analysis of accelerating the convergence of inference based on the Ising model, which is one of the most popular pair-wise models in the physics community for studying opinion dynamics [10]. The problem of

* Corresponding author.
E-mail addresses: Zhongqi.Cai@soton.ac.uk (Z. Cai), eg@ecs.soton.ac.uk (E. Gerding), Markus.Brede@soton.ac.uk (M. Brede).

https://doi.org/10.1016/j.physa.2023.129348
Received 2 March 2023; Received in revised form 22 August 2023
Available online 31 October 2023

reconstructing parameters of the Ising-like model from extant time series data is named the inverse kinetic Ising problem [1]. As a subproblem of the more general problem of statistical inference, the core of most inference approaches for the inverse kinetic Ising problem is to maximize the likelihood of model parameters given time series of system states [1]. The majority of studies in this area concentrate on enhancing the inference performance by improving the accuracy or the scalability of proposed inference algorithms via utilizing varying approximations with various regimes of validity [9,11–19]. For instance, the mean-field method [11,12] and the Thouless-Anderson-Palmer approach [13–15] are utilized to provide approximated solutions for the inverse kinetic Ising problem in the weak and dense network connection region. Moreover, belief propagation [20] and replicas analysis [16,17] are used for inference in the strong and sparse network connection region. However, the above-mentioned methods only give close estimations for the inverse kinetic Ising problem for large sample sizes. Improving on this and considering the limitation of the dataset size obtained from experiments, the work of Hoang et al. [9] utilizes linear regression to provide accurate estimations for relatively small sample sizes. However, even though the inverse kinetic Ising problem has been extensively studied from an algorithmic perspective in the literature mentioned above, the data side has found relatively little attention [21]. In other words, little attention has been put in the aspect of enhancing the inference performance by improving the dataset's quality to obtain more accurate estimators with less data. This is, however, of great importance in many real-world applications when the measurements for the dynamical processes are costly or there are technical limitations in observing the whole process of networked dynamics such as the early stage of a rumour (or epidemic) spreading.

To the best of our knowledge, there are only two works that consider the inference problem from the data quality side [21,22]. Of these, Decelle et al. [21] measure the amount of information contained in a given dataset by applying a method that quantifies the effective rank of the correlation matrix. However, that work has also assumed data to be given and the authors do not investigate an active way of generating data for inference. The other work that is related to our modelling approach is our previous work [22]. That work investigates how to strategically influence networked dynamics with the aim of generating a higher-quality dataset for speeding up the convergence of inference. However, our previous study is based on the voter model, which belongs to the category of simple contagion [23]. It is of interest if similar ideas can be applied in more realistic settings of complex contagion [24,25] in the kinetic Ising model.

To bridge the gaps in the problem of inference acceleration from the data generating side, we extend our previous framework in [22], to the kinetic Ising model. Different from the linearity of the voter model, which results in high levels of mathematical tractability [26], the non-linearity of the Ising model requires different techniques for analysis. Inspired by the utilization of network control in [22], we treat the external magnetic fields as external controllers who exert influence in the network by building weighted and unidirectional links. Without loss of generality, we assume the existence of two opposing controllers. One is considered as passive, whose strategies are unknown and need to be inferred. The other controller is the focus of our investigation, and is devoted to minimizing the uncertainty of inference of its opponent's strategy by optimally manipulating the Ising dynamics with the aim of generating more informative data. Corresponding to the natural resource limits in real-world contexts, we also assume there are only limited resources available for the strategic controller.

Specifically, we make the following contributions. First, the inference of the inverse kinetic Ising problem of previous studies is constructed on a given and fixed dataset. Our work is the first to explore the inverse kinetic Ising problem from the perspective of manipulating the networked dynamics to generate a more informative dataset for inference. Second, our work provides a comprehensive characterization for the configuration of optimal budget allocations by the active controller. To achieve this, we combine numerical results obtained from heuristics with detailed analytical explanations by the Taylor series approximation [27] to gain a deeper understanding for the structure of optimal allocation strategies in the high-temperature region. Third, we analyse the performance of the heuristics for generating the optimal allocations in the scenario where the real-time tracking of the system states is inaccessible. More specifically, we substitute the real dataset with the mean-field state data generated by the mean-field approximation [28]. Our results establish that the convergence of inference can be accelerated by smartly targeting agents in the network with optimized control gains to improve the quality of generated dataset for inference. Moreover, we find two general patterns for the configuration of optimal budget allocations in the absence or presence of budget constraints. When not considering budget constraints, optimal allocations are determined only by the targeted node's neighbourhood and the influencing strength from the opponent. When budget constraints are taken into account, the optimal allocations are to re-weight the unconstrained optimal allocations by counting the uncertainty of all the other agents' estimations. More specifically, agents with larger uncertainty in estimation will have a larger chance to be allocated resources closer to their unconstrained optimal allocations. Further, by investigating the optimized influence profile, we find that agents targeted by the opponent with higher influences will be allocated more resources on average by the active controllers so as to even out the inaccuracy for inferring larger values.

The remainder of this paper is organized as follows. In Section 2, we give a formal description for the framework of accelerating the convergence of inference under the inverse Ising context and present heuristics of how to optimally interact with the Ising dynamics to generate more informative data. In Section 3, we provide numerical results for the optimal configuration of budget allocations by the active controller obtained by numerical optimization and analytical approximations. In Section 4, we summarize the main findings and contributions and discuss ideas for future work.

2. Model description and methods

Following the work by Galam [29], we interpret the Ising model in the context of opinion formation. In more detail, spins in the Ising model are considered as agents connected by a network. In the following, we consider a system consisting of N agents. Each agent is identical to a node embedded in the social network where the social links between agents are given by a weighted

adjacency matrix $W = \{w_{ij}\}_{i,j=1}^{N}$. Agents i and j are considered to be neighbours if there is a weighted link $w_{ij} > 0$ between them. We further assume that each agent holds one of two possible opinions at time step t, labelled as $s_i(t) = -1$ or $s_i(t) = 1$ for $i = 1, \dots, N$ and $t \geq 0$. Here, we assume the existence of two opposite external fields A and B, also referred to as controller A and B. To be more specific, controllers A and B are treated as external positive and negative fields that correspond to zealots who have unchanged opinions $s_A(t) = 1$ and $s_B(t) = -1$ for $\forall t \geq 0$. In addition, external controllers influence the Ising dynamics via unidirectional and positively-weighted connections to node i in the network at time t, denoted as $J_{A,i}(t) \in \mathbb{R}^+$ and $J_{B,i}(t) \in \mathbb{R}^+$, respectively. Note that, here we assume the resources b_A and b_B available for controllers A and B at each time step are limited, i.e., the sum of control gains $J_{A,i}(t)$ and $J_{B,i}(t)$ at each time step t are subject to the budget constraints: $\sum_{i=1}^{N} J_{A,i}(t) \leq b_A$, $\sum_{i=1}^{N} J_{B,i}(t) \leq b_B$.

Following the commonly-used Glauber algorithm [29,30] in simulating social dynamics, the parallel and discrete-time Ising dynamics consist of the following three steps: (i) Sum the weighted neighbouring states of node i ($i = 1, \dots, N$) at time t, which is denoted as $S_i(t) = \sum_j w_{ij} s_j(t)$. (iii) Compute the changes in the energy level of the Ising system if node i flips at time t, which is $\Delta E_i(t) = 2s_i(t)S_i(t) + 2J_{A,i}(t)s_i(t) - 2J_{B,i}(t)s_i(t)$. (iii) Flip the state of node i at time $t+1$ with probability $e^{-\Delta E_i(t)/T} / \left(1 + e^{-\Delta E_i(t)/T}\right)$ where T is the temperature. Correspondingly, the time-varying transition matrix $P_i(t)$ describing the probabilities of state changes of node i at time t is given by

$$P_i(t) = \left[P_i^{-1}(t) \quad P_i^{1}(t)\right] = \left[\frac{e^{\left(-2S_i(t)-2J_{A,i}(t)+2J_{B,i}(t)\right)/T}}{1+e^{\left(-2S_i(t)-2J_{A,i}(t)+2J_{B,i}(t)\right)/T}} \quad \frac{1}{1+e^{\left(-2S_i(t)-2J_{A,i}(t)+2J_{B,i}(t)\right)/T}}\right]. \tag{1}$$

Here, $P_i^{-1}(t)$ stands for the probability of node i to have state -1 at time $t+1$. From the third term of Eq. (1), we observe that the transition probabilities are independent of the updated node's current opinion state and are only determined by the node's neighbouring states and the control gains from external controllers.

As mentioned above, the external controllers interact with the internal Ising dynamics by targeting nodes with positively-continuous weighted links. In this paper, we are interested in addressing the problem of accelerating the reconstruction of the external controllers' strategy from agents' opinion changes. Specifically, we assume that controller A is an active controller who strategically distributes its control gains to alter the data generation with the aim of obtaining more accurate estimates of controller B's budget allocations within fewer observations. For this purpose, we assume that controller B is a constant opponent who has fixed budget allocations from time 0, i.e., $J_{B,i}(t) = J_{B,i}(0)$ for $\forall t \geq 0$. For simpler notation, we refer to $J_{B,i}(t)$ as $J_{B,i}$ in the following.

To obtain estimators for controller B's budget allocations, we use maximum likelihood estimation (MLE) [31] for parametric inference. More specifically, given the transition probabilities of opinion flips in Eq. (1), the logarithm of the likelihood function for observing time series of opinion changes during time span $[0, M]$ for node i is

$$L_i(M) = \sum_{t\in[0,M-1]} \left[\frac{1+s_i(t+1)}{2} \log\left(P_i^{1}(t)\right) + \frac{1-s_i(t+1)}{2} \log\left(P_i^{-1}(t)\right)\right]. \tag{2}$$

Inserting $P_i^{1}(t)$ and $P_i^{-1}(t)$ from Eq. (1) into Eq. (2) yields the full expression. By maximizing the likelihood function of Eq. (2) regarding the budget allocations of controller B, we obtain estimators of $J_{B,i}$, denoted as $\hat{J}_{B,i}$.

In accordance with the consistency of the MLE method, given sufficiently long time series of agents' opinion changes, the estimators $\hat{J}_{B,i}$ asymptotically approach the true value of $J_{B,i}$ [31]. Nevertheless, considering the cost of data collection in most real-world scenarios, it is always preferable to obtain more accurate estimators within a shorter observation period.

To assess the accuracy of the inference, we use the frequently-used Fisher information [32] as a metric to evaluate the quality of fit of the estimators derived by MLE. The Fisher information gives a measure for the dispersion between the deduced estimators from MLE and actual values. According to Efron and Hinkley [33], the Fisher information $I(J_{B,i}, M)$ pertaining to $J_{B,i}$ is defined as the expectation of the second-order partial derivative of Eq. (2) with respect to $J_{B,i}$, i.e.,

$$I(J_{B,i}, M) = E\left[\frac{\partial^2 L_i(M)}{\partial J_{B,i}^2}\right] = -\sum_{t\in[0,M-1]} \left[\frac{\operatorname{sech}^2\left(\frac{J_{A,i}(t)-J_{B,i}+S_i(t)}{T}\right)}{T^2}\right]. \tag{3}$$

As the function domain of hyperbolic secant ($\operatorname{sech}(\cdot)$) is limited within the range of $(0, 1]$, the Fisher information $I(J_{B,i}, M)$ is negative for all possible values of independent variables, i.e., $J_{A,i}(t) \in \mathbb{R}^+$, $J_{B,i} \in \mathbb{R}^+$ and $S_i(t) \in \mathbb{R}$. Moreover, as the length of observation M increases, the value of $I(J_{B,i}, M)$ decreases accordingly. Thereafter, by taking the negative reciprocal of Fisher information, one can generate confidence intervals for the MLE estimators. In more detail, for a large enough data sample, the estimators $\hat{J}_{B,i}$ obtained from MLE converge in a normal distribution to the actual value $J_{B,i}$. Therefore, we have

$$\left(\hat{J}_{B,i} - J_{B,i}\right) \xrightarrow{D} \mathcal{N}\left(0, \left[-I(J_{B,i}, M)\right]^{-1}\right). \tag{4}$$

Here, $\mathcal{N}(0, \left[-I(J_{B,i}, M)\right]^{-1})$ represents the normal distribution with standard deviation $\left[-I(J_{B,i}, M)\right]^{-1/2} > 0$, and mean value 0. Moreover, as $I(J_{B,i}, M)$ monotonically decreases with respect to an increase in the number of observations, we will obtain unbiased estimators $\hat{J}_{B,i}$ with lower standard deviation as we use more system updates. Note that, in Eq. (4), the true values of controller B's budget allocations are used to calculate the standard deviation. However, as the true values are what we want to infer and are normally unknown, in practical calculations, we substitute the true values $J_{B,i}$ ($1 \leq i \leq N$) with the estimated ones $\hat{J}_{B,i}$ [32].

In the following, to improve estimates of opponents, we minimize the standard deviation deduced by the Fisher information in Eq. (3). By doing so, we have transformed the problem of accelerating the convergence of inference to strategically choosing the

allocations of controller A's budget to minimize the standard deviation of the estimators. Moreover, from Eq. (3), we observe that the Fisher information can be calculated in a recursive way, i.e.,

$$I(J_{B,i}, t+1) = I(J_{B,i}, t) - \frac{\operatorname{sech}^2\left(\frac{J_{A,i}(t)-J_{B,i}+S_i(t)}{T}\right)}{T^2}. \tag{5}$$

Accordingly, for the estimated standard deviation of node i at current estimates $\hat{J}_{B,i}$, we have

$$\hat{\sigma}_i(\hat{J}_{B,i}, t+1) = \left(-I(\hat{J}_{B,i}, t) + \frac{\operatorname{sech}^2\left(\frac{J_{A,j}(t)-\hat{J}_{B,j}+S_j(t)}{T}\right)}{T^2} \right)^{-1/2}. \tag{6}$$

Extending our previous framework in [22], we apply the heuristics called the *one-step-ahead optimization* to the inverse kinetic Ising problem. In more detail, we optimize the configuration of controller A's budget allocations at the current step with the aim of minimizing the sum of the expected standard deviations in the next step. By doing so, we aim at a step-wise generation of a more informative dataset. Formally, the objective function for the inference acceleration problem is given by

$$\{ \overbrace{J^*_{A,1}(t), \ldots, J^*_{A,N}(t)}^{N \text{ agents in the network}} \} = \arg\min \sum_{i=1}^{N} \hat{\sigma}_i(\hat{J}_{B,i}, t+1) = \arg\min \sum_{i=1}^{N} \left[-I(\hat{J}_{B,i}, t) + \frac{\operatorname{sech}^2\left(\frac{J_{A,i}(t)-\hat{J}_{B,i}+S_i(t)}{T}\right)}{T^2} \right]^{-1/2} \tag{7}$$

subject to

$$\sum_{i=1}^{N} J^*_{A,i}(t) \le b_A$$
$$J^*_{A,i}(t) \ge 0 \qquad \text{for} \quad 1 \le i \le N.$$

Here, $J^*_{A,i}(t)$ represents the optimized budget allocation by controller A for node i, and the sum of optimized $J^*_{A,i}(t)$ (for $1 \le i \le N$) should satisfy the budget constraint b_A. Since the weighted sum of neighbouring states $S_i(t)$ can be observed and estimators $\hat{J}_{B,i}$ can be computed by maximizing Eq. (2), the only unknown parameters in Eq. (7) will be $J_{A,i}(t)$.

As Eq. (7) contains transcendental terms, it is challenging to obtain a closed-form solution for the optimization. However, by considering the expected standard deviation for each node separately in Eq. (6) and taking $J_{A,i}(t)$ as the only independent variable, we have the following analytical findings for each single node in the absence of budget constraints. First, the expected standard deviation for node i, $\hat{\sigma}_i(\hat{J}_{B,i}, t+1)$, is symmetric against the vertical line $J_{A,i}(t) = \hat{J}_{B,i} - S_i(t)$ if we relax the domain of $J_{A,i}(t)$ to $\mathbb{R}$. Second, $\hat{\sigma}_i(\hat{J}_{B,i}, t+1)$ has a global minimum (see Appendix A for a detailed proof), given by

$$J_{A,i}(t) = \begin{cases} \hat{J}_{B,i} - S_i(t), & \hat{J}_{B,i} \ge S_i(t) \\ 0, & \hat{J}_{B,i} < S_i(t) \end{cases}. \tag{8}$$

Additionally, by taking the second-order derivative of Eq. (6), we find that for lower temperature T, the function of $\hat{\sigma}_i(\hat{J}_{B,i}, t+1)$ has larger curvature at the minimum point $J_{A,i}(t) = \hat{J}_{B,i}(t) - S_i(t)$ compared with a setting of higher temperatures when curvatures are smaller (for the proof, see Appendix B).

2.1. High-temperature Taylor expansion

To investigate the optimal allocations in more depth, we use the Taylor expansion to come up with an analytical approximation for the optimal budget allocations. As the Taylor expansion for the standard deviation (see Eq. (6)) does not exist for low temperature, in the following, we focus on the approximation in the high temperature regime. By applying the Lagrange multiplier technique [34] based on the high-temperature Taylor approximation of standard deviation of Eq. (6), we obtain approximated solutions for the optimization of Eq. (7). A detailed proof is given in Appendix C. The approximated budget allocations obtained by the high-temperature Taylor expansion are denoted as $J^{approx}_{A,i}(t)$ for $1 \le i \le N$.

We find that there are two possible cases for $J^{approx}_{A,i}(t)$:

(i)

$$\text{If} \quad \sum_{i=1}^{N} J^{approx}_{A,i}(t) \le b_A, \quad \text{then} \quad J^{approx}_{A,i}(t) = \begin{cases} \hat{J}_{B,i} - S_i(t), & \hat{J}_{B,i} \ge S_i(t) \\ 0, & \hat{J}_{B,i} < S_i(t) \end{cases}. \tag{9}$$

That is, if controller A has a large enough budget, the approximated solution will be the same as the optimized one in the absence of a budget constraint (see Eq. (8)).

(ii) Otherwise, let

$$\phi_i(t) = \left(\hat{J}_{B,i} - S_i(t)\right) + \frac{\left[b_A - \sum_{J^{approx}_{A,j}(t)\neq 0} \left(\hat{J}_{B,j} - S_j(t)\right)\right] \hat{\sigma}_i^{-3}(\hat{J}_{B,i}, t)}{\sum_{J^{approx}_{A,j}(t)\neq 0} \hat{\sigma}_j^{-3}(\hat{J}_{B,j}, t)}. \tag{10}$$

Here, $\sum_{J_{A,j}^{approx}(t)\neq 0}$ stands for summing up all the j where the budget allocations from the controller A on nodes j are not 0. Thus, we have

$$J_{A,i}^{approx}(t) = \begin{cases} 0, & \hat{J}_{B,i} \leq S_i(t) \quad \text{or} \quad \phi_i(t) \leq 0 \\ \phi_i(t), & \phi_i(t) > 0 \end{cases}. \tag{11}$$

In this case, by summing up the approximated budget allocations over all nodes, we have $\sum_{i=1}^{N} J_{A,i}^{approx}(t) = b_A$. In practice, we compute $J_{A,i}^{approx}(t)$ as follows: (i) Calculate the approximation of optimal budget allocations $J_{A,i}^{approx}(t)$ according to Eq. (9) and then sum up $J_{A,i}^{approx}(t)$ for $1 \leq i \leq N$. If the sum exceeds the budget constraint b_A, continue with step (ii). Otherwise, the procedure is terminated. (ii) Let $J_{A,i}^{approx}(t) = 0$ for $\hat{J}_{B,i} \leq S_i(t)$. (iii) Set the rest of the non-zero approximated optimal budget allocations $J_{A,i}^{approx}(t) = \phi_i(t)$, and then proceed with determining whether all $J_{A,i}^{approx}(t)$ are non-negative. If so, the procedure is terminated. Otherwise, set the negative $J_{A,i}^{approx}(t)$ to be 0 and recalculate the rest of the non-zero $J_{A,i}^{approx}(t)$ according to Eq. (10). (iv) Repeat step (iii) until all $J_{A,i}^{approx}(t)$ are non-negative. Note that, the approximated solutions for the optimized budget configurations in Eq. (11) are to re-weight the unconstrained solutions in Eq. (8) by counting the standard deviation of other agents' estimations. Generally, an agent with larger standard deviation in estimations will be allocated with resources closer to its unconstrained optimized budget allocation.

2.2. High-temperature mean-field approximation

From the analytical solution of the high-temperature Taylor approximation shown in Eqs. (9) and (11), we find that the approximated optimal budget allocations are dependent on the neighbouring states of the targeted nodes. However, in some practical applications, it might be hard to keep track of the state changes at each step. Therefore, in the following we consider to replace the real states of nodes by mean-field approximated states. Specifically, assume that $\langle s_i \rangle$ denotes the average of the node i's states over the opinion dynamics. Using r_i to approximate $\langle s_i \rangle$ and employing the mean-field approximation (e.g., see [28]), we obtain

$$r_i = \tanh\left[\left(\sum_j w_{ij} r_j + J_{A,i} - J_{B,i}\right)/T\right] \tag{12}$$

for the opinion dynamics following the Glauber dynamics. Because of the hyperbolic term in Eq. (12), it is hard to obtain an explicit solution for r_i. In the following, we consider the mean-field approximation in high-temperature region. Therefore, by applying the Taylor expansion for high temperature with respect to Eq. (12), we have

$$r_i = \frac{\sum_{j=1}^{N} w_{ij} r_j + J_{A,i} - J_{B,i}}{T} + \mathcal{O}\left(T^{-3}\right) \tag{13}$$

To further simplify the solution of Eq. (13), we assume that all nodes are affected by the same mean field, in which we have:

$$r_i \approx \frac{d_i \langle r \rangle + J_{A,i} - J_{B,i}}{T} \tag{14}$$

where

$$\langle r \rangle = \frac{1}{N} \sum_i r_i. \tag{15}$$

Here, $d_i = \sum_j w_{ij}$ represents the weighted degree for node i. By summing Eq. (14) over all nodes, we obtain

$$\sum_i r_i = N \langle r \rangle = \frac{N \langle d \rangle \langle r \rangle + \sum_i J_{A,i} - b_B}{T} \tag{16}$$

i.e.,

$$\langle r \rangle = \frac{\sum_i J_{A,i} - b_B}{TN - \langle d \rangle N} \tag{17}$$

where $\langle d \rangle$ represents the averaged linking weights of the network. Inserting Eq. (17) into Eq. (14) yields the full expression for r_i. Note that $\langle r \rangle$ is dependent on the real budget used by controller A, which could be less or equal to the budget constraint b_A.

By replacing the actual sum of weighted neighbouring states S_i with the sum of mean-field states $d_i \langle r \rangle$ in Eqs. (9) and (11), we have

$$J_{A,i}^{MF} = \begin{cases} \hat{J}_{B,i} - d_i \langle r \rangle, & \hat{J}_{B,i} \geq d_i \langle r \rangle \\ 0, & \hat{J}_{B,i} < d_i \langle r \rangle \end{cases}, \quad \text{If} \quad \sum_{i=1}^{N} J_{A,i}^{MF} \leq b_A. \tag{18}$$

and otherwise

$$J_{A,i}^{MF} = (\hat{J}_{B,i} - d_i \langle r \rangle) + \frac{(b_A - \sum_{J_{A,j}^{MF}\neq 0} \hat{J}_{B,j} + \sum_{J_{A,j}^{MF}\neq 0} d_j \langle r \rangle)\hat{\sigma}_i^{-3}(\hat{J}_{B,i}, t)}{\sum_{J_{A,j}^{MF}\neq 0} \hat{\sigma}_j^{-3}(\hat{J}_{B,j}, t)}. \tag{19}$$

Here, $J_{A,i}^{MF}$ represents the high-temperature mean-field approximation for the optimal budget allocation of node i. Eq. (19) exists only when $J_{A,i}^{MF}$ is non-negative. Importantly, for the special case of $b_A \geq b_B$, $J_{A,i}^{MF} = \hat{J}_{B,i}$ holds. This means that, even when the

active controller has more budget compared with its opponent, it will only spend the same amount of budget as its opponent for better inference in the mean-field scenario. According to Eq. (6), to compute $\hat{\sigma}_j^{-3}(\hat{J}_{B,j}, t)$, we also need the neighbouring states. To further simplify Eq. (19), we apply the Taylor expansion on $\hat{I}(\hat{J}_{B,i}, M)$ for high temperature, in which we have

$$I(\hat{J}_{B,i}, M) = \sum_{t=0}^{M-1}\left[-\frac{1}{T^2} + \frac{\left(J_{A,i}(t) - \hat{J}_{B,i} + S_i(t)\right)^2}{T^4} + \mathcal{O}\left(T^{-5}\right)\right] \approx \sum\left(-\frac{1}{T^2} + \frac{r_i^2}{T^2}\right). \tag{20}$$

Assume that node i is in equilibrium state. Then, for the mean-field approximation of the standard deviation $\hat{\sigma}_i(\hat{J}_{B,i}, M)$, we have

$$\hat{\sigma}_i(\hat{J}_{B,i}, M) \approx \left(\frac{M - r_i^2 M}{T^2}\right)^{-1/2} \tag{21}$$

where M is the length of the observation. Generally, for high temperature, r_i is small. Therefore, we have

$$\hat{\sigma}_i(\hat{J}_{B,i}, M)^{-3} = \left(\frac{M}{T^2}\right)^{3/2} - \frac{3}{2} r_i^2 \left(\frac{M}{T^2}\right)^{3/2} + \mathcal{O}\left(r_i^4\right). \tag{22}$$

If we only consider the first term in Eq. (22) and replace $\hat{\sigma}_i(\hat{J}_{B,i}, M)^{-3}$ with it in Eq. (19), we obtain

$$J_{A,i}^{MF} \approx \hat{J}_{B,i} + \frac{b_A - \sum_{J_{A,j}^{MF} \neq 0} \hat{J}_{B,j}}{Z} + \frac{b_A - b_B}{TN - \langle d \rangle N}\left(\frac{\sum_{J_{A,j}^{MF} \neq 0} d_j}{Z} - d_i\right), \tag{23}$$

where Z counts for the number of non-zero $J_{A,i}^{MF}$. Eq. (23) exists when $J_{A,i}^{MF}$ is non-negative.

2.3. Numerical one-step-ahead optimization

The approximated solutions for optimal budget allocations obtained in Sections 2.1 and 2.2 are deduced in the premise of high temperature. To obtain comprehensive results for the optimized budget allocations regardless of the temperature constraints, we use the interior-point method [35] for numerical optimization of Eq. (7) in all temperature regions. By doing this, we aim at getting step-wise optimal budget allocations $J_{A,i}^{opt}$ for $i = 1, \ldots, N$ that can be different at each time step t, and name this as the numerical one-step-ahead optimization. The procedure for the numerical one-step-ahead optimization is formalized in Algorithm 1.

Specifically, the detailed experimental setup for the numerical one-step-ahead optimization shown in Algorithm 1 is given as follows: (i) To meet the assumption of having enough samples before leveraging Fisher information to compute the standard deviation of MLE estimators given by Eq. (4), all nodes in the network are targeted by the same budget allocation $J_{A,f}$ by controller A for the first T_0 updates in the initialization part of Algorithm 1. We also keep track of the likelihood functions for the first T_0 updates for all nodes, denoted as $L_i(T_0)$ for $i = 1, \ldots, N$. (ii) Record the current observation time step t. If t is smaller than the total number of observations T_1, we compute the estimators $\hat{J}_{B,i}$ for all nodes based on current maximum likelihood functions $L_i(t)$ described in Eq. (2). Thereafter, we compute the Fisher information defined in Eq. (5) and insert them into the objective function in Eq. (7). By optimizing Eq. (7) with the interior point method, we obtain a set of optimized budget allocations for controller A, denoted as $J_{A,i}^{opt}(t)$. Finally, the network is updated with the new set of allocations $J_{A,i}^{opt}(t)$ following the Ising dynamics to obtain the next states at $t+1$ for all nodes. Lines 3–7 of Algorithm 1 codify the contents of this step. (iii) When the time step t exceeds the total number of observations T_1, the procedure is terminated. Note that the time complexity of the one-step-ahead optimization is $\mathcal{O}\left((T_1 - T_0)N^3\right)$.

Algorithm 1: Numerical one-step-ahead optimization

input : Adjacency matrix W, total number of observations T_1, fixed budget allocation $J_{A,f}$, length of updates before computing the standard deviation T_0

output: optimized budget allocation, $J_{A,i}^{opt}(t)$ for $T_0 \leq t \leq T_1$, $1 \leq i \leq N$

1 *Initialization: 50% of the initial opinions of agents $s_i(0)$ are -1 or 1; update the network with $J_{A,f}$ for the first T_0 steps; let $t = T_0$;*
2 **while** $t \leq T_1$ **do**
3 $\quad \hat{J}_{B,i} = \max_{J_{B,i}} L_i(t)$ for $1 \leq i \leq N$;
4 $\quad$ obtain $J_{A,i}^{opt}(t)$ by optimizing Eq. (7);
5 $\quad$ update the network following stochastic Ising dynamics with $\{J_{A,i}^{opt}(t)\}_{i=1}^{N}$;
6 $\quad t = t + 1$;
7 **end**

3. Results

In this section, we present the main results for the inference acceleration problem, with a special focus on exploring the configuration of the optimized budget allocations by the active controller A. First, in Section 3.1, we explore the inferrability of nodes in the case of no interference from controller A. Then, to gain some insights into how the inference is affected by the budget

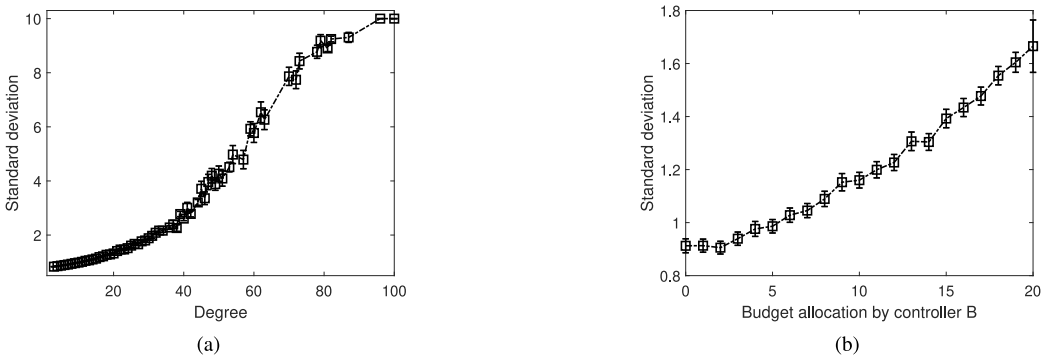

Fig. 1. (a) Dependence of the standard deviation of estimators for controller B's control gains on nodes' degrees in the absence of controller A. (b) Dependence of the standard deviation of estimators for controller B's control gains on controller B's budget allocations in the absence of controller A. The standard deviation is calculated at time step 1000 and temperature $T = 20$. We use a setting in which controller B targets all nodes with allocations randomly sampled from a uniform distribution $U(0, 20)$, in which the budget allocation by controller B per node on average is 10. We group the value of x axis into bins with width 1 and lower limits are inclusive, e.g., [0,1). Error bars indicate 95% confidence intervals.

allocations, we target all nodes with equal budget allocations. The results from the equally targeting scenario then work as the benchmark to verify the effectiveness of the proposed one-step-ahead heuristics. Second, we proceed with an investigation of the approximated configuration of the optimized budget allocations in the high-temperature region obtained by the high-temperature Taylor expansion in Section 3.2. Third, we further explore the profiles of optimized budget allocations in the scenario of not having priori knowledge about the neighbourhood by utilizing the high-temperature mean-field approximation in Section 3.3. Furthermore, in Section 3.4, we discuss the influence of budget availability on the optimal control.

The numerical experiments in this work are performed on uncorrelated random scale-free networks with network size $N = 1000$ and average degree $\langle d \rangle = 10$, generated according to the configuration model [36]. To ensure a large degree heterogeneity, the degree distribution of the constructed networks is $p_d \propto d^{-2}$ where d represents the node's degree. The results shown in the following context are all based on 50 repetitions of the corresponding experiments.

3.1. Benchmark — the equally targeting scenario

Before investigating the influence of budget allocations by the active controller on the inference, we are interested in the accuracy of inference without interference from controller A. Here, we use the standard deviation of estimators in Eq. (4) to quantify the accuracy of inference. For a preliminary investigation, we assume the passive controller B targets each node with random control gains drawn from a uniform distribution $U(0, 20)$ with the first moment of the distribution 10 equal to the average degree of the networks. To proceed, in Fig. 1, we present the dependence of the standard deviation of estimators for the opponent's strategies on nodes' degrees in panel (a) and on the allocations from controller B in panel (b) at temperature $T = 20$. In Panels (a) and (b) of Fig. 1, we observe clear patterns for the corresponding dependence. In the absence of interference from controller A, nodes are the harder to predict the larger their degrees and the larger the budget allocations from controller B. These results are consistent with the conclusion drawn for the voter model in [22]. A wider range of temperature values from $T = 1$ to $T = 100$ for inference without interference can be seen in Figs. 7 and 8 of Appendix D, showing similar patterns.

We then proceed with the qualitative exploration of the opponent strategy inference problem in the equally targeting scenario where the active controller A targets all nodes with the same control gain. By doing so, we aim at gaining some intuitions about how the inference is affected by the interference from the active controller A. Furthermore, we will later use the equally targeting strategy as a benchmark which will give insights into improvements in the estimations that can be obtained by optimization.

To proceed, in Fig. 2(a), we show the dependence of the mean standard deviation $\bar{\sigma}_{eq}$ for the equally targeting scheme over all nodes on varying temperature T and different relative budget constraints $b_A = 0.5b_B$ (black circles), $b_A = b_B$ (blue triangles), $b_A = 2b_B$ (red squares). After a careful inspection of Fig. 2(a), we make the following observations. First, as shown in the upper-left corner of Fig. 2(a), it is hard to gain accurate estimations of the opponent's strategy in the low temperature regions. The main reason for the inaccuracy of the inference is that, at low temperature, the system falls into stabilization which leads to spontaneous magnetization [37]. Some nodes are keeping their states unchanged during the whole updating process. Therefore, no information would be obtained from the observation. Second, in the inset of Fig. 2(a), we find that, for extremely high temperature, the inference errors are roughly the same for different equally targeting budgets. Together with the probabilities of state flips in Eq. (1), we can thus understand why the Ising system is harder to control for inference acceleration as the temperature increases: In the context of high temperature, nodes have nearly equal probabilities to change to state 1 or −1. In this case, the temperature is the main factor for state flipping. As the control gains are divided by a large temperature, they have little influence in determining the agents' next states. Third, by combining the inset and the whole picture in Fig. 2(a), we find that the dependence of mean standard deviation on temperature is a convex shape. This means that there is a temperature at which predictions of opponent strategies are the most accurate.

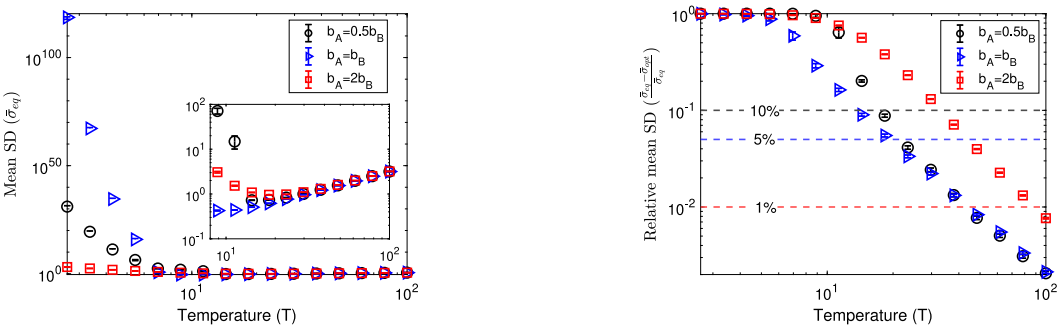

Fig. 2. (a) Dependence of the mean standard deviation $\bar{\sigma}_{eq}$ of estimators for controller B's control gains over all nodes at time step 1000 on varying temperature T when controller A targets all nodes equally. Different relative budget constraints $b_A/b_B = \{0.5, 1, 2\}$ are shown by different colours of curves. (b) Dependence of the relative mean standard deviation $\frac{\bar{\sigma}_{eq}-\bar{\sigma}_{opt}}{\bar{\sigma}_{eq}}$ on varying temperature T. $\bar{\sigma}_{opt}$ is calculated by averaging the standard deviation of estimators obtained via applying the numerical one-step-ahead optimization over all nodes at time step 1000. The three horizontal lines show the percentages of improvements in reducing the standard deviation by the optimization compared with the equally targeting strategy. We use a setting in which controller B targets all nodes with allocations randomly sampled from a uniform distribution $U(0, 20)$, in which the budget allocation by controller B per node on average is 10. Error bars indicate 95% confidence intervals.

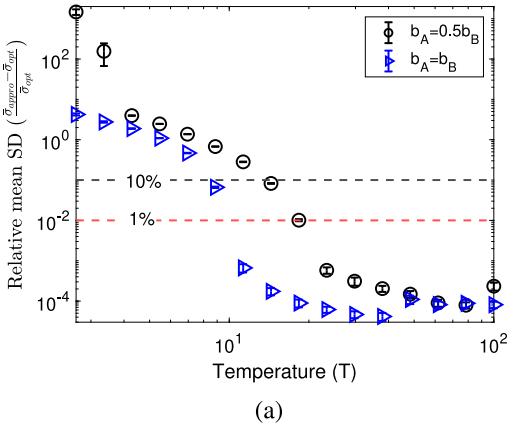

Fig. 3. (a) Dependence of the relative mean standard deviation $\frac{\bar{\sigma}_{approx}-\bar{\sigma}_{opt}}{\bar{\sigma}_{opt}}$ on varying temperature T. $\bar{\sigma}_{approx}$ stands for the mean standard deviation of estimators obtained via applying the high-temperature Taylor series approximation shown as Eqs. (9) and (11) over all nodes after the system initialization. Different relative budget constraints $b_A/b_B = \{0.5, 1\}$ are shown by different colours of curves. We use a setting in which controller B targets all nodes with allocations randomly sampled from a uniform distribution $U(0, 20)$, in which the budget allocation by controller B per node on average is 10. Error bars indicate 95% confidence intervals.

Next, to figure out if the inference of opponent budget allocations could be accelerated or not by the proposed one-step-ahead optimization algorithm, we compare the mean standard deviation of estimators $\bar{\sigma}_{eq}$ computed by the equally targeting strategy with the mean standard deviation $\bar{\sigma}_{opt}$ calculated by applying the optimized budget allocations of Algorithm 1 in Fig. 2(b). To focus on the improvements of inference in the high-temperature region, we consider the relative mean standard deviation $\frac{\bar{\sigma}_{eq}-\bar{\sigma}_{opt}}{\bar{\sigma}_{eq}}$ for different budget constraints $b_A/b_B = \{0.5, 1, 2\}$ in Fig. 2(b). It becomes clear that the optimization works best for extremely low temperature with a reduction of the relative standard deviation close to 100%. Moreover, with an increase of temperatures, the optimized standard deviation gets closer to the equally targeting one. However, more than 5% or about 1% improvements are achieved in the intermediate or high temperature regions compared with the equally targeted strategy. Additionally, we observe that the one-step-ahead optimization has a better performance in reducing the standard deviation of estimators if the active controller A has more budget than the passive controller B.

3.2. Results for high-temperature Taylor series approximation

As the approximated solutions in Eq. (11) for the optimal budget allocations are deduced under the assumption of high temperature, identifying the feasible region where the Taylor series approximation has close performance to the optimization is crucial. To proceed, in the panels of Fig. 3, we compare the mean standard deviation over all nodes calculated via the one-step-ahead optimization of Algorithm 1 with the Taylor series approximation in Eq. (11). In more detail, in Fig. 3(a), we show the

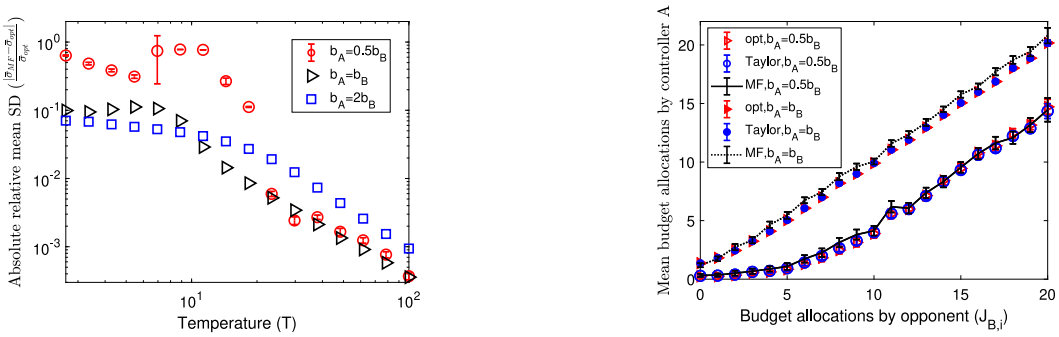

Fig. 4. (a) Dependence of the absolute value of the relative mean standard deviation $\frac{|\bar{\sigma}_{MF}-\bar{\sigma}_{opt}|}{\bar{\sigma}_{opt}}$ on varying temperature T. $\bar{\sigma}_{MF}$ is computed by averaging Eq. (21) over all nodes. Different relative budget constraints $b_A/b_B = \{0.5, 1, 2\}$ are shown by different shapes of symbols. (b) Dependence of the mean budget allocations by controller A over updates on the budget allocations by controller B. The triangles, squares and lines represent for the results calculated by the optimization, Taylor expansion and mean-field approximation, respectively. The blank symbols and solid line are for budget constraint $b_A = 0.5b_B$, while the filled symbols and dashed line are for $b_A = b_B$. The value of x axis are grouped into bins with width 1 and lower limits are inclusive, e.g., [0,1). Error bars indicate 95% confidence intervals.

dependence of relative mean standard deviation $\frac{\bar{\sigma}_{approx}-\bar{\sigma}_{opt}}{\bar{\sigma}_{opt}}$ on the temperature T at update 1000. Here, $\bar{\sigma}_{approx}$ stands for the mean standard deviation computed by the Taylor series approximation (see Eqs. (9) and (11)). To obtain $\bar{\sigma}_{approx}$, we consider the following procedure: After the first $T_0 = 100$ updates of system initialization as mentioned in Algorithm 1, we replace the optimization by the Taylor series approximation and obtain the approximated budget allocations at each update. Thereafter, we update the network with the approximated values following Ising dynamics. By doing this, the results of $\bar{\sigma}_{approx}$ are purely based on the approximated budget allocations but not the optimal budget allocations. Note that, unlike Fig. 2, we only present two budget cases $b_A = 0.5b_B$ and $b_A = b_B$ in Fig. 3(a). The reason for not presenting the results for $b_A = 2b_B$ is that: if controller A has sufficient budget, the optimized budget allocations will be the same as the approximated ones as indicated in Eqs. (8) and (9). As a consequence, the mean standard deviation $\bar{\sigma}_{opt}$ and $\bar{\sigma}_{approx}$ will be the same, and thus the relative mean standard deviation $\frac{\bar{\sigma}_{approx}-\bar{\sigma}_{opt}}{\bar{\sigma}_{opt}} = 0$ cannot be presented in a logarithmic figure.

Continuing with the results in Fig. 3(a), we find that even though the approximation is based on the assumption of high temperature, it also works well in the intermediate temperature region. For instance, for $b_A = b_B$, the approximation will result in a performance loss of less than 10% compared with the optimization at temperature above 10. Therefore, in the medium and high temperature region, one could substitute the optimized budget allocations with the approximated ones to reduce the computational complexity from $\mathcal{O}\left((T_1 - T_0)N^3\right)$ to $\mathcal{O}(T_1 - T_0)$ without sacrificing much of the performance. Moreover, consistent with the high-temperature assumption, the relative difference of standard deviation between the Taylor series approximation and the numerical optimization gets smaller with increasing temperatures. However, in the low-temperature region ($T < 6$), the high-temperature Taylor series approximation will lead to the performance loss of $\frac{\bar{\sigma}_{approx}-\bar{\sigma}_{opt}}{\bar{\sigma}_{opt}} \geq 1$. This is particularly pronounced for the insufficient budget scenario $b_A = 0.5b_B$, as demonstrated in the upper-left corner of Fig. 3(a), where the relative mean standard deviation will exceed 10^2. The significant performance loss in the case of $b_A = 0.5b_B$ is a result of budget insufficiency, which leads to bad initial estimations $\hat{J}_{B,i}$ for most nodes after the first $T_0 = 100$ updates after system initialization. The subsequent utilization of Eqs. (9) and (11) for approximated solutions allocates budgets only to nodes with extremely high standard deviations, leaving other nodes untargeted. The division of budgets across a large number of nodes with extremely high standard deviations prevents targeted nodes from receiving sufficient control gains to flip their states. Meanwhile, other untargeted nodes are hard to flip at low temperatures. Consequently, the approximated budget allocations and system states remain almost constant, leading to almost no information gain. This results in a persistently high mean standard deviation in the approximated solution. In contrast, for the $b_A = b_B$ scenario, the system initialization yields poor initial estimations for only a few nodes. Allocating resources to these nodes prompts state flips and thereafter leads to a relatively lower standard deviation compared to the $b_A = 0.5b_B$ case. A more detailed explanation of the difference in the relative mean standard deviation between $b_A = 0.5b_B$ and $b_A = b_B$ for the low-temperature region can be found in Appendix E.

3.3. Results for high-temperature mean-field approximation

In the following, we consider a more practical scenario in the real-world context where we do not have access to real-time tracking of the system dynamics. In this scenario, we apply the high-temperature mean-field solutions presented in Section 2.2 to obtain a guess for nodes' neighbouring states and thereafter generate approximations for the optimized budget allocations in the high-temperature region. Similar to Section 3.2, we start with identifying the region where the mean-field approximation has close performance to the numerical one-step-ahead optimization. For this purpose, we evaluate the performance of the high-temperature mean-field approximation based on different temperatures. Fig. 4(a) shows the dependence of the absolute value of the relative

mean standard deviationn $\frac{\left|\bar{\sigma}_{MF}-\bar{\sigma}_{opt}\right|}{\bar{\sigma}_{opt}}$ on the temperature. Here $\bar{\sigma}_{MF}$ is computed by averaging Eq. (21) over all nodes. Notice that, on the y-axis, instead of calculating the relative values as Figs. 2 and 3, we compute the absolute difference between the numerical optimization and the mean-field approximation. The reason for this is that in some cases $\bar{\sigma}_{MF}$ will be smaller than $\bar{\sigma}_{opt}$. For example, consider the setting of $b_A \geq b_B$. In the scenario of $b_A \geq b_B$, the mean-field states r_i will be zero. Therefore, based on Eq. (21), $\bar{\sigma}_{MF} = \left(\frac{M}{T^2}\right)^{-1/2}$, which only considers the first term of the Taylor expansion of the Fisher information in Eq. (20) and ignores the non-negative term of $\frac{(J_{A,i}(t)-J_{B,i}+S_i(t))^2}{T^4}$ at each time step. This makes the high-temperature mean-field approximation of the standard deviation smaller than its actual value. To facilitate the analysis of the high-temperature mean-field approximation, we consider the relative standard deviation under three control settings, where controller A has less budget $b_A = 0.5b_B$, equal budget $b_A = b_B$, and more budget $b_A = 2b_B$ in Fig. 4(a). Consistent with the results of the high-temperature Taylor series approximation, the high-temperature mean-field approximation has the worst performance if the active controller has less budget compared with its opponent in the low temperature region (see the upper-left corner of Figs. 3(a) and 4(a)). However, the approximated standard deviation gradually converges to the optimized one as the temperature rises for these three budget settings. More importantly, even though the high-temperature mean-field approximation is deduced under the assumption of high temperature, it works well also in the intermediate temperature region where the difference in the standard deviation is less then 10% for temperature over 10 when $b_A \geq b_B$.

Moreover, in Fig. 4(a), for $b_A = 0.5b_B$, after $T = 5.5$, we observe an increase in the absolute value of the relative mean standard deviation for $6.5 \leq T \leq 13$. The jump is mainly caused by the quality of the approximations. To be more specific, $\bar{\sigma}_{MF}$ is computed by averaging Eq. (21), where M and T are given and r_i is dependent on $\langle r\rangle = \frac{\sum_i J_{A,i}-b_B}{TN-\langle d\rangle N}$ (see Eq. (17)). As the experiments are carried out on networks with average degree $\langle d\rangle = 10$, and given $\left|\sum_i J_{A,i} - b_B\right| \gg 0$ for $b_A = 0.5b_B$, we obtain high $|\langle r\rangle|$ for temperatures around 10 (i.e., $6.5 \leq T \leq 13$). This further leads to high r_i^2 and $\hat{\sigma}_i(\hat{J}_{B,i}, M) \approx 0$ (see Eq. (21)). As a result, $\frac{\left|\bar{\sigma}_{MF}-\bar{\sigma}_{opt}\right|}{\bar{\sigma}_{opt}}$ are close to 1 for $6.5 \leq T \leq 13$ in the setting of $b_A = 0.5b_B$. In contrast, for $b_A = b_B$ and $b_A = 2b_B$, according to Eq. (17), $\langle r\rangle$ is dependent on the real budget used by controller A, which could be less or equal to the budget constraint. In these cases, $\left|\sum_i J_{A,i} - b_B\right|$ is very close to 0, and will not result in significant increases in the mean-field standard deviations for temperatures around 10.

We next investigate the shape of the budget allocations obtained via the numerical optimization and approximations. As the approximations have better performance in the intermediate and high temperature regions, we first present the profile of budget allocations by controller A for an intermediate temperature $T = 20$. In more detail, in Fig. 4(b), we present the dependence of the mean budget allocations by controller A averaged over updates on the opponent's budget allocations at $T = 20$. Note that, for ease of observation, we group opponent's budget allocations in x-axis into bins with width 1 in Fig. 4(b). Results shown in this figure are obtained from three different algorithms where the triangles represent the numerical one-step-ahead optimization calculated as $\left\langle J_{A,i}^{opt}\right\rangle = \frac{\sum_{t=T_0}^{T_1} J_{A,i}^{opt}(t)}{T_1-T_0}$, circles are for Taylor series approximation $\left\langle J_{A,i}^{approx}\right\rangle = \frac{\sum_{t=T_0}^{T_1} J_{A,i}^{approx}(t)}{T_1-T_0}$ and lines stand for mean-field approximation $J_{A,i}^{MF}$. Moreover, the influence of budget constraints are also considered in Fig. 4(b), where marked symbol and dashed line represent $b_A = b_B$, and blank symbols and line are for $b_A = 0.5b_B$. By observing the trends of budget allocations by controller A for varying budget allocations from its opponent, we find that for agents targeted by the opponent with larger budget allocations, the active controller tends to allocate more resources on average as well. Moreover, we find that the high-temperature mean-field and Taylor series methods provide good approximations for the optimized budget allocations for both budget constraints. Note that, consistent with the analytical result obtained from the high-temperature mean-field method in which for $b_A \geq b_B$ $J_{A,i}^{MF} = \hat{J}_{B,i}$, we observe a nearly linear dependence in Fig. 4(b) for $b_A = b_B$.

To proceed, we investigate the accuracy of inference for agents after we performed the optimal control and compare with the results for the no interference scenario in Fig. 1. Similar to Fig. 1(a), we show the dependence of standard deviation of estimations on nodes with different degrees when applying the optimization of Algorithm 1 or using high-temperature Taylor and mean-field approximations in Fig. 5(a). By observing Fig. 5(a), we find that the high-temperature Taylor approximation provides close performance in depicting the influence of node degree heterogeneity on inference accuracy compared with the exact numerical optimization. Moreover, consistent with the no interference case in Fig. 1(a), we find that it is also harder to predict a higher degree node under optimization. In Fig. 5(b), we present the dependence of standard deviation of estimators on opponent budget allocations. We find that, compared with Fig. 1(b), the application of optimal control strongly alleviates the difference of standard deviation for nodes targeted by the opponent with various values (i.e., we observe that values vary by less than 0.01 as opposed to nearly 1 in Fig. 1(b)). Nevertheless, we still observe a clear negative correlation, i.e. the less the resource a node is targeted by the opponent, the harder it is to predict. This conclusion differs from the pattern observed in the no interference case in Fig. 1(b), where we observed that nodes targeted with higher control gains are harder to predict.

3.4. Influence of budget availability

To facilitate the analysis of the configurations of the budget allocations by the active controller, we further investigate the relationship between the budget allocations by the active controller and its opponent on the ratio of the budget constraints b_A/b_B. Similar to Fig. 4(b), we group budget allocations by the opponent into bins with width 1 and calculate the corresponding mean value of $\left\langle J_{A,i}^{opt}\right\rangle$, $\left\langle J_{A,i}^{approx}\right\rangle$ and $J_{A,i}^{MF}$ within the given bins, denoted as $\bar{J}_A^{opt,m}$, $\bar{J}_A^{approx,m}$ and $\bar{J}_A^{MF,m}$ for bins with higher limit m. Therefore,

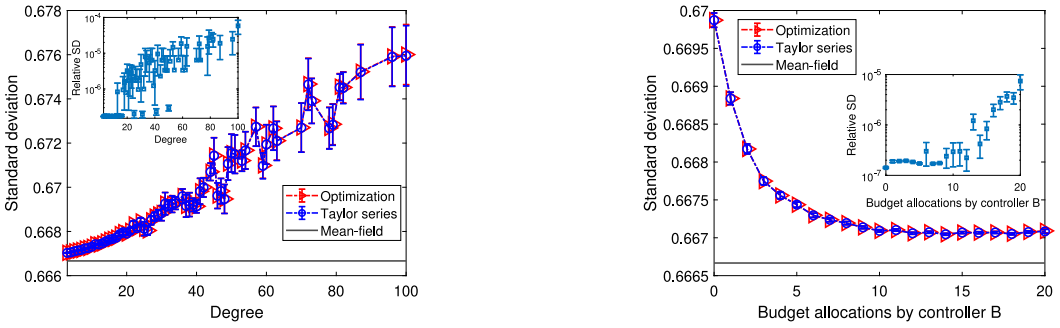

Fig. 5. (a) Dependence of the standard deviation of estimators for controller B's control gains on nodes' degrees. The inset shows the dependence of the absolute difference in standard deviations calculated by the optimization and by the Taylor expansion on degree. (b) Dependence of the standard deviation of estimators for controller B's control gains on controller B's budget allocations. The inset shows the dependence of the absolute difference in standard deviations calculated by the optimization and by the Taylor expansion on budget allocations by controller B. The standard deviation is calculated at time step 1000 and temperature $T = 20$. The red triangles, blue circles and lines are for results calculated by the one-step-optimization, Taylor expansion and mean-field approximation separately. We use a setting in which controller B targets all nodes with allocations randomly sampled from a uniform distribution $U(0, 20)$, in which the budget allocation by controller B per node on average is 10. Error bars indicate 95% confidence intervals.

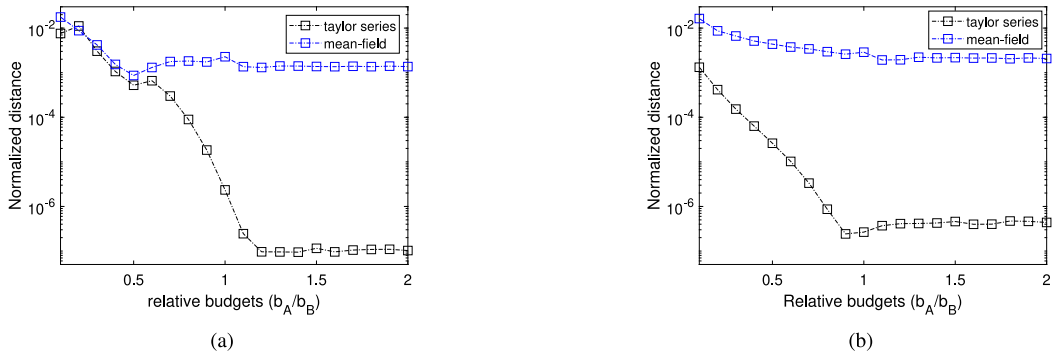

Fig. 6. Dependence of the fitness of the relationship of mean budget allocations by controller A on the allocations by controller B calculated by approximations compared with the optimization on relative budget constraints b_A/b_B. Here the fitness is calculated by taking square root of the sum of squares of difference between corresponding points calculated by optimization and by Taylor series or mean-field approximations calculated by Eqs. (24) and (25), respectively. The distance is further normalized by the actual budget used by controller A in the optimization. The blue squares show the normalized distance between Taylor expansion and the one-step-ahead optimization while the black squares are for the normalized distance between mean-field approximation and the one-step-ahead optimization. Panel (a) shows the results for temperature $T = 20$, and Panel (b) is $T = 50$.

the distance of the dependence of mean budget allocations by the active controller on $\hat{J}_{B,i}$ between the numerical optimization and the high-temperature Taylor approximation is calculated as

$$\sqrt{\sum_m \left(\bar{J}_A^{opt,m} - \bar{J}_A^{approx,m}\right)^2}/b_A^* \tag{24}$$

Similarly, the distance of the dependence between the mean-field approximation and numerical optimization is

$$\sqrt{\sum_m \left(\bar{J}_A^{opt,m} - \bar{J}_A^{MF,m}\right)^2}/b_A^* \tag{25}$$

Here b_A^* represents the actual total budget that the controller A has used during the optimization in Fig. 6.

For comparison, Fig. 6(a) is computed at intermediate temperature $T = 20$ while (b) is for high-temperature at $T = 50$. By observing Fig. 6(a), we find that the distances of budget allocations between the numerical optimization and approximations calculated by Eqs. (24) and (25) are smaller for larger relative budgets b_A/b_B. This implies that the approximations work better if the active controller has more budget than its opponent. However, we also note thresholds at which if the b_A exceeds a certain value then increasing b_A does not result in improvements of the performance of the approximations. By comparing the normalized distance for different temperature, we find that, temperature will shift the "best" points of relative budgets where the high-temperature Taylor expansion will have the best performance in approaching the same budget allocations as the numerical optimization algorithm. For example, for $T = 20$, the performance of the high-temperature Taylor expansion is close to the numerical optimization when

$b_A/b_B \geq 1.2$. However, for higher temperature $T = 50$, high-temperature Taylor expansion has close performance to the optimization at $b_A/b_B \geq 0.9$.

4. Conclusion

Existing literature related to the inverse kinetic Ising problem has traditionally studied inference based on a given dataset of observations. In contrast, in this paper, we have investigated how to speed up the convergence of inference via strategically interfering with the Ising dynamics with the aim of generating a more informative dataset. To achieve this, we place the inference accelerating problem in the scenario of two opposing controllers. One is actively interacting with the opinion dynamics via optimizing its allocations step-wisely to minimize the uncertainty of its opponent's unknown strategies.

By comparing with the benchmark scenario in which the active controller targets all nodes equally, we establish that interacting with the Ising dynamics can substantially speed up the convergence of inference. Specifically, in the low-temperature region, the proposed heuristics decreases the uncertainty of inference by almost 100 percent compared with the benchmark case. This agrees with the finding of Decelle et al. [21], who demonstrate that out-of-equilibrium data allow for a much more accurate inference compared with equilibrium/stationary data. Moreover, even though the external controllers have limited power in manipulating the networked dynamics in higher temperature settings, the proposed heuristics will still make some improvements in speeding up the convergence of estimators of opponent strategies in the high-temperature region.

As a second contribution, we have provided a comprehensive exploration of the configuration of the optimal allocations both analytically and numerically. We find that, if the active controller has sufficient budget, it will allocate resources equal to the difference between the current estimators of opponent's allocations and the sum of the targeted node's weighted neighbouring states. In the case of insufficient budget, we utilize the high-temperature Taylor approximation and find that the optimized allocations will be re-weighted according to the standard deviation of current estimators based on the solutions for the sufficient budget case, in order to meet the budget constraints. By observing the profile of the mean optimized allocation over updates, we find a clear positive dependence of mean optimized allocation on opponent's budget allocations. Moreover, even by performing the optimal control, high degrees will impede nodes to be inferred accurately. The main reason for this is that, for nodes with large degree, when it changes its state, it is hard to distinguish whether the change is a result of control or its neighbours. This pattern is consistent with the findings presented in [38] and [22]. However, contrary to the results obtained from the voter model in [22], nodes targeted by the opponent with larger allocations will have higher inferrability.

Furthermore, we have extended our heuristics of optimally interacting with the networked dynamics to the case of not having the real-time tracking of the system dynamics. To cope with this situation, we substitute the real states with the mean-field states. This assumption makes our algorithm applicable to a wide range of scenarios when real-time tracking for the feedback from the population is infeasible. For a more extreme case, the results obtained via the mean-field approximation can be used as a guideline to have accurate estimators for opponent's strategies when only very limited data is accessible.

Our work illustrated above has been limited to inferring a fix-strategy opponent. Therefore, a possible line for future work is to investigate the inference acceleration problem when the opponent will change its allocations dynamically. Moreover, as indicated in the work of Romero Moreno et al. [39], knowing the strategy of the opponent will enable a better performance in influence maximization where controllers compete to maximize their influence in the network. Therefore, one could combine the opponent strategy inference problem with influence maximization where the active controller needs to find a trade-off between the resources putting to accelerate the inference and to maximize its influence.

CRediT authorship contribution statement

Zhongqi Cai: Conceptualization, Methodology, Software, Formal analysis, Investigation, Data curation, Writing – original draft, Writing – review & editing, Visualization. **Enrico Gerding:** Conceptualization, Methodology, Validation, Writing – review & editing, Supervision, Project administration. **Markus Brede:** Conceptualization, Methodology, Validation, Writing – review & editing, Supervision, Project administration.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

Acknowledgements

The authors acknowledge the use of the IRIDIS High Performance Computing Facility in the completion of this work. ZC acknowledges support from China Scholarships Council (No. 201906310134). MB acknowledges support from the Alan Turing Institute, United Kingdom (EPSRC grant EP/N510129/1, https://www.turing.ac.uk/) and the Royal Society, United Kingdom (grant IES\R2\192206, https://royalsociety.org/).

Appendix A. Global minima for a single node

Specifically, to find the minima for $\hat{\sigma}_i(\hat{J}_{B,i}, t+1)$ in Eq. (6), we take a first-order derivative of $\hat{\sigma}_i(\hat{J}_{B,i}, t+1)$ with respect to $J_{A,i}$. By doing so, we obtain

$$\frac{\partial \hat{\sigma}_i(\hat{J}_{B,i}, t+1)}{\partial J_{A,i}(t)} = \frac{T \tanh\left(\frac{J_{A,i}(t)-\hat{J}_{B,i}+S_i(t)}{T}\right) \operatorname{sech}^2\left(\frac{J_{A,i}(t)-\hat{J}_{B,i}+S_j(t)}{T}\right)}{\left(\operatorname{sech}^2\left(\frac{J_{A,i}(t)-\hat{J}_{B,i}+S_j(t)}{T}\right) - I(\hat{J}_{B,i}, t)T^2\right)^2 \hat{\sigma}_i(\hat{J}_{B,i}, t+1)}. \tag{26}$$

By letting Eq. (26) equal to 0, we find that, $J_{A_j} = \hat{J}_{B,i} - S_i(t)$ is always a solution for Eq. (26) and it is independent of the value of last-step Fisher information $I(\hat{J}_{B,i}, t)$. Furthermore, the sign of Eq. (26) is determined by the term of $\left(T \tanh\left(\frac{J_{A,i}(t)-\hat{J}_{B,i}+S_i(t)}{T}\right)\right)$, and we have

$$\begin{aligned} &\tanh\left(\frac{J_{A,i}(t) - \hat{J}_{B,i} + S_i(t)}{T}\right) > 0 \qquad \text{for} \qquad J_{A,i}(t) > \hat{J}_{B,i}(t) - S_i(t) \\ &\tanh\left(\frac{J_{A,i}(t) - \hat{J}_{B,i} + S_i(t)}{T}\right) < 0 \qquad \text{for} \qquad J_{A,i}(t) < \hat{J}_{B,i} - S_i(t). \end{aligned} \tag{27}$$

This means that the function of Eq. (6) increases monotonically when $J_{A,i}(t) > \hat{J}_{B,i}(t) - S_i(t)$ and decreases monotonically when $J_{A,i}(t) < \hat{J}_{B,i}(t) - S_i(t)$. Therefore, $J_{A,i}(t) = \hat{J}_{B,i}(t) - S_i(t)$ is a global minima for the standard deviation $\hat{\sigma}_i(\hat{J}_{B,i}, t+1)$ if we consider $J_{A,i}(t)$ in the domain of $\mathbb{R}$. However, as the budget allocation cannot be negative, when $\hat{J}_{B,i}(t) - S_i(t) < 0$, the global minima is $J_{A,i}(t) = 0$.

Appendix B. Curvature of expected standard deviation

The second-order derivative of Eq. (6) with respect to $J_{A,i}(t)$ is given by

$$\frac{\partial^2 \hat{\sigma}_i(\hat{J}_{B,i}, t+1)}{\partial J_{A,i}^2} = -\frac{\operatorname{sech}^4\left(\frac{J_{A,i}(t)-\hat{J}_{B,i}+S_i(t)}{T}\right)\left(I(\hat{J}_{B,i}, t)T^2 \cosh\left(\frac{2(J_{A,i}(t)-\hat{J}_{B,i}+S_i(t))}{T}\right) - 2I(\hat{J}_{B,i}, t)T^2 + 1\right)}{\left(I(\hat{J}_{B,i}, t)T^2 - \operatorname{sech}^2\left(\frac{J_{A,i}(t)-\hat{J}_{B,i}+S_i(t)}{T}\right)\right)^3 \hat{\sigma}_i(\hat{J}_{B,i}, t+1)}. \tag{28}$$

By letting $J_{A,i}(t) = \hat{J}_{B,i}(t) - S_i(t)$ in Eq. (28), we have

$$\left.\frac{\partial^2 \hat{\sigma}_i(\hat{J}_{B,i}, t+1)}{\partial J_{A,i}^2}\right|_{J_{A,i}(t)=\hat{J}_{B,i}(t)-S_i(t)} = \left(T^{-1} - I(\hat{J}_{B,i}, t)\right)^{1/2}. \tag{29}$$

For a larger temperature T, the value of Eq. (29) is smaller. Therefore, the curve of Eq. (6) is flatter for a high temperature compared with a lower temperature.

Appendix C. High-temperature Taylor expansion

For high temperature, we have

$$\hat{\sigma}_i(\hat{J}_{B,i}, t+1) = \hat{\sigma}_i(\hat{J}_{B,i}, t) - \frac{\hat{\sigma}_i^3(\hat{J}_{B,i}, t)}{2T^2} + \frac{3\hat{\sigma}_i^5(\hat{J}_{B,i}, t)}{8T^4} + \frac{(J_{A,i}(t) - J_{B,i} + S_i(t))^2 \hat{\sigma}_i^3(\hat{J}_{B,i}, t)}{2T^4} + \mathcal{O}\left(\left(\frac{1}{T}\right)^5\right). \tag{30}$$

Here, the $\mathcal{O}$ represents the Big O notation.

Next, to obtain an analytical solution for the minima of Eq. (7) subject to the inequality constraint $\sum_{i=1}^{N} J_{A,i}(t) \le b_A$, we apply the Lagrange multiplier technique [34] based on the approximation shown in Eq. (30). In order to handle the inequality constraint in the objective function, we introduce slack variables t and k_i $(1 \le i \le N)$. By letting

$$f\left(J_A(t)\right) = \sum_{i=1}^{N}\left[\hat{\sigma}_i(\hat{J}_{B,i}, t) - \frac{\hat{\sigma}_i^3(\hat{J}_{B,i}, t)}{2T^2} + \frac{3\hat{\sigma}_i^5(\hat{J}_{B,i}, t)}{8T^4} + \frac{\left(J_{A,i}(t) - \hat{J}_{B,i} + S_i(t)\right)^2 \hat{\sigma}_i^3\left(\hat{J}_{B,i}, t\right)}{2T^4}\right], \tag{31}$$

where $J_A(t) = \{J_{A,i}(t)\}_{i=1}^{N}$, Eq. (7) is converted into

$$\begin{aligned} &\{\overbrace{J_{A,1}^{approx}(t), \ldots, J_{A,N}^{approx}(t)}^{N \text{ agents in the network}}\} = \arg\min f\left(J_A(t)\right) \\ &\text{subject to} \\ &\qquad h\left(J_A(t)\right) = \sum_{i=1}^{N} J_{A,i}^{approx}(t) - b_A + t^2 = 0 \\ &\qquad g_i\left(J_A(t)\right) = -J_{A,i}^{approx}(t) + k_i^2 = 0 \qquad \text{for} \quad 1 \le i \le N, \end{aligned} \tag{32}$$

Here, $J_{A,i}^{approx}(t)$ for $1 \le i \le N$ denote the optimal budget allocations obtained from the Taylor expansion. Therefore, the Lagrangian is defined as:

$$L(J_A(t), \lambda, \{\beta_i\}_{i=1}^N) = f\left(J_A(t)\right) + \lambda h\left(J_A(t)\right) + \sum_{i=1}^{N}\left[\beta_i g_i\left(J_A(t)\right)\right]. \tag{33}$$

To obtain the optimal solution of Eq. (32), we set the gradients of Eq. (33) to be 0:

$$\begin{aligned}
\frac{\partial L}{\partial J_{A,i}} &= \frac{\left(J_{A,i}(t) - \hat{J}_{B,i} + S_i(t)\right)\hat{\sigma}_i^3(\hat{J}_{B,i}, t)}{T^4} + \lambda - \beta_i = 0 \\
\frac{\partial L}{\partial \lambda} &= \sum_{i=1}^{N} J_{A,i}(t) - b_A + t^2 = 0 \\
\frac{\partial L}{\partial t} &= 2\lambda t = 0 \\
\frac{\partial L}{\partial \beta_i} &= -J_{A,i}(t) + k_i^2 = 0 \\
\frac{\partial L}{\partial k_i} &= 2\beta_i k_i = 0.
\end{aligned} \tag{34}$$

In accordance with the Karush–Kuhn–Tucker conditions [34], the inequality constraints should satisfy the complementary slackness condition, i.e., either the Lagrange multipliers are equal to zero or the inequality constraints are active. Moreover, for the minimization, the Lagrange multipliers should be non-negative, i.e.,

$$\begin{aligned}
\lambda &\ge 0 \\
\beta_i &\ge 0.
\end{aligned} \tag{35}$$

By solving the system of Eqs. (34) and (35), we have:

(i)

$$\text{If} \quad \sum_{i=1}^{N} J_{A,i}^{approx}(t) \le b_A, \quad J_{A,i}^{approx}(t) = \begin{cases} \hat{J}_{B,i} - S_i(t), & \hat{J}_{B,i} \ge S_i(t) \\ 0, & \hat{J}_{B,i} < S_i(t) \end{cases}. \tag{36}$$

(ii) Otherwise, let

$$\phi_i(t) = \left(\hat{J}_{B,i} - S_i(t)\right) + \frac{\left[b_A - \sum_{J_{A,j}^{approx}(t)\neq 0}\left(\hat{J}_{B,j} - S_j(t)\right)\right]\hat{\sigma}_i^{-3}(\hat{J}_{B,i}, t)}{\sum_{J_{A,j}^{approx}(t)\neq 0}\hat{\sigma}_j^{-3}(\hat{J}_{B,j}, t)}. \tag{37}$$

Here, $\sum_{J_{A,j}^{approx}(t)\neq 0}$ stands for summing up all the j where the budget allocations from the controller A on nodes j are not 0. Thus, we have

$$J_{A,i}^{approx}(t) = \begin{cases} 0, & \hat{J}_{B,i} \le S_i(t) \quad \text{or} \quad \phi_i(t) \le 0 \\ \phi_i(t), & \phi_i(t) > 0 \end{cases}. \tag{38}$$

Appendix D. Standard deviation of estimators for varying temperatures in the absence of controller A

In this section, we extend the experiments depicted in Fig. 1, where we initially evaluated the accuracy of inference without interference from controller A at temperature $T = 20$. In Figs. 7 and 8, we expand our investigation to include a wider range of temperature values from $T = 1$ to $T = 100$ with the aim of offering a more comprehensive understanding of the performance of inference without the inclusion of external interference.

In Figs. 7 and 8, we explore the dependence of the standard deviation of estimators for controller B's control gains on nodes' degrees and on the allocations from controller B. This exploration takes place in the absence of controller A, considering various temperatures. Specifically, we analyse $T = 1$ (7(a) and (d)), $T = 10$ (7(b) and (e)), $T = 30$ (7(c) and (f)), $T = 50$ (8(a) and (d)), $T = 70$ (8(b) and (e)), and $T = 100$ (8(c) and (f)). Similar to the patterns observed in Fig. 1, Figs. 7 and 8 also demonstrate that nodes are the harder to predict the larger their degrees and the larger the budget allocations from controller B. Furthermore, when comparing standard deviations at varying temperatures, we observe that, at low temperatures, obtaining precise estimations of the opponent's strategy is challenging, given that the states of certain nodes remain unchanged. As temperature increases, the differences in standard deviations among nodes of varying degrees and targeted by different control gains from controller B diminish. This is due to the higher temperature causing nodes to have a more similar probability of state flipping.

Appendix E. Mean standard deviation calculated by the high-temperature Taylor series approximation in low-temperature regions

In Fig. 3, we compare the mean standard deviation calculated via the one-step-ahead optimization of Algorithm 1 with the high-temperature Taylor series approximation in Eqs. (9) and (11) for two budget configurations: $b_A = 0.5b_B$ and $b_A = b_B$. Notably,

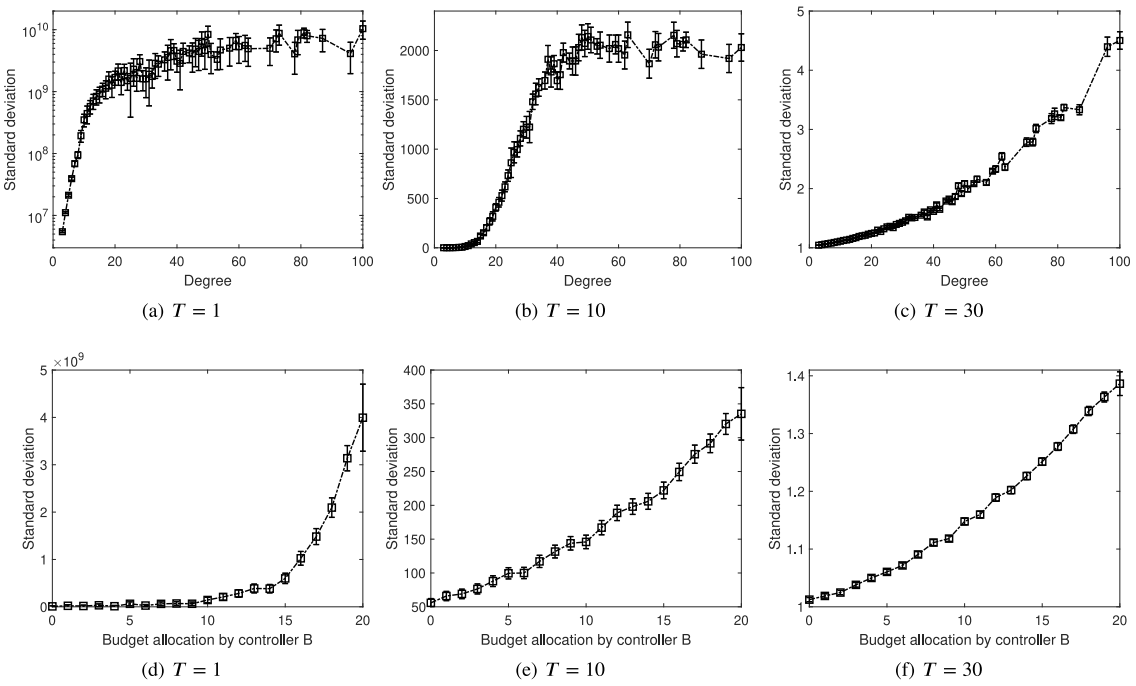

Fig. 7. (a–c) Dependence of the standard deviation of estimators for controller B's control gains on nodes' degrees in the absence of controller A. (d-f) Dependence of the standard deviation of estimators for controller B's control gains on controller B's budget allocations in the absence of controller A. The standard deviation is calculated at time step 1000 and temperature $T = 1$ (panels (a),(d)), $T = 10$ (panels (b),(e)), $T = 30$ (panels (c),(f)). We use a setting in which controller B targets all nodes with allocations randomly sampled from a uniform distribution $U(0, 20)$, in which the budget allocation by controller B per node on average is 10. We group the value of x axis into bins with width 1 and lower limits are inclusive, e.g., [0,1). Error bars indicate 95% confidence intervals.

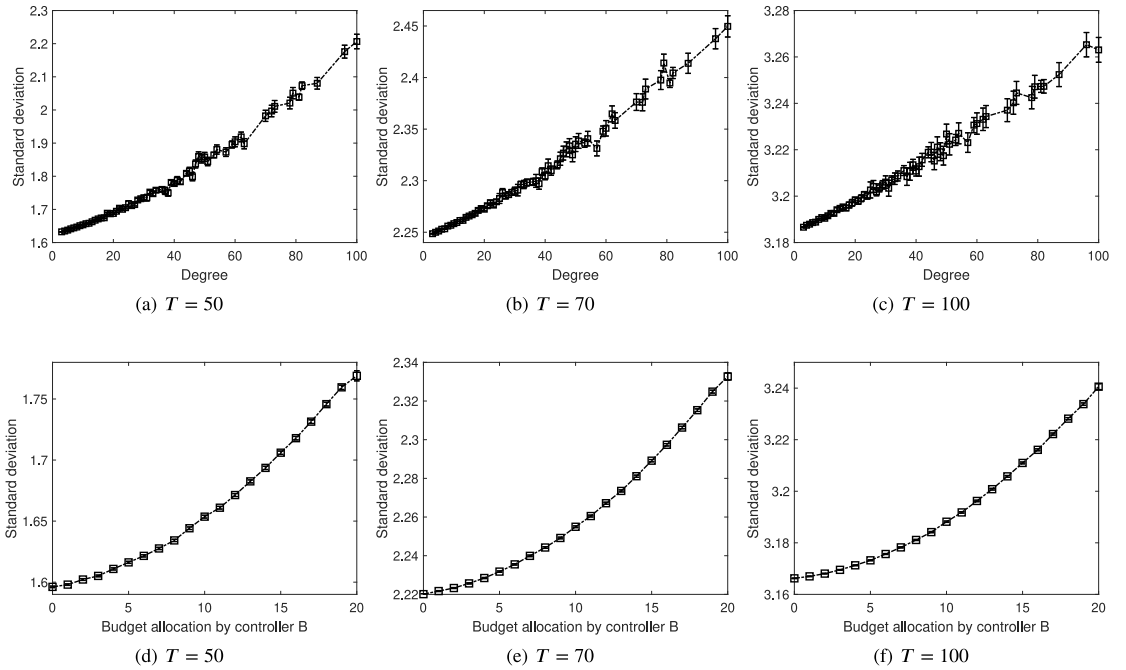

Fig. 8. (a–c) Dependence of the standard deviation of estimators for controller B's control gains on nodes' degrees in the absence of controller A. (d-f) Dependence of the standard deviation of estimators for controller B's control gains on controller B's budget allocations in the absence of controller A. The standard deviation is calculated at time step 1000 and temperature $T = 50$ (panels (a),(d)), $T = 70$ (panels (b),(e)), $T = 100$ (panels (c),(f)). We use a setting in which controller B targets all nodes with allocations randomly sampled from a uniform distribution $U(0, 20)$, in which the budget allocation by controller B per node on average is 10. We group the value of x axis into bins with width 1 and lower limits are inclusive, e.g., [0,1). Error bars indicate 95% confidence intervals.

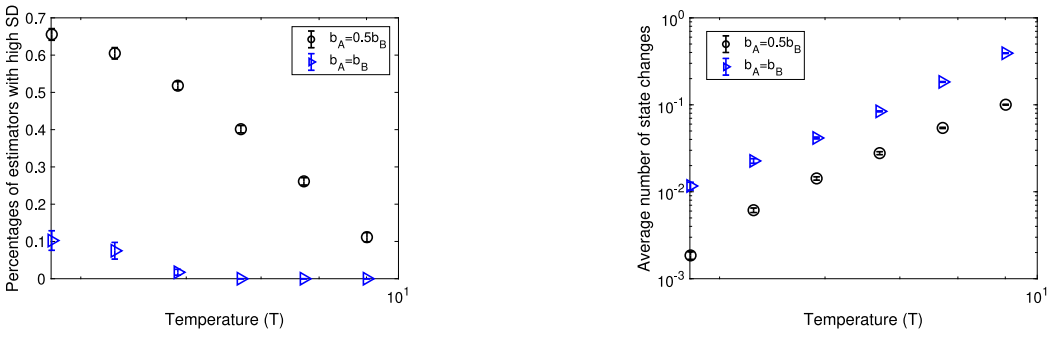

Fig. 9. (a) Percentages of estimators with standard deviations exceeding 10^2 among all nodes after system initialization $T_0 = 100$ in the low-temperature region. (b) Average number of state changes $\sum_{i=1}^{N}\sum_{t=T_0}^{T_1-1}\frac{|s_i(t+1)-s_i(t)|}{2N(T_1-T_0)}$ after applying the approximated budget allocations according to Eqs. (9) and (11) in the low-temperature region. Here, we set $T_1 = 1000$. Different relative budget constraints $b_A/b_B = \{0.5, 1\}$ are shown by different colours of curves. We use a setting in which controller B targets all nodes with allocations randomly sampled from a uniform distribution $U(0, 20)$, in which the budget allocation by controller B per node on average is 10. Error bars indicate 95% confidence intervals.

the high-temperature Taylor series approximation provides a close performance to the one-step-ahead optimization in the medium and high temperature regions for both $b_A = 0.5b_B$ and $b_A = b_B$. However, in the low-temperature region of Fig. 3, the relative mean standard deviation significantly increases as the temperature decreases for $b_A = 0.5b_B$, whereas no comparable rise is observed for $b_A = b_B$. An explanation for this phenomenon is as follows: After the first $T_0 = 100$ updates of the system initialization as illustrated in Algorithm 1, we obtain initial estimates of controller B's budget allocations $\hat{J}_{B,i}$ and corresponding standard deviations for each node. We record the percentages of initial estimates with standard deviations exceeding 10^2 among all nodes for both $b_A = 0.5b_B$ and $b_A = b_B$ in the low-temperature region ($T < 10$) in Fig. 9(a). As shown in Fig. 9(a), due to the budget insufficiency in $b_A = 0.5b_B$, we obtain a larger number of initial estimates with extremely high standard deviations compared with the case of $b_A = b_B$.

After the system initialization, we proceed with utilizing Eqs. (9) and (11) to calculate the approximated budget allocations of controller A. According to Eq. (11), the approximation prioritizes allocating budgets to agents with extremely high standard deviations to induce state flips. However, for $b_A = 0.5b_B$, as seen in Fig. 9(a), when $T < 4$ (corresponding to the first two points in Fig. 3(a) with relative mean standard deviations exceeding 10^2), the percentage of estimators with high standard deviations is larger than 0.5. In other words, dividing the limited budget $b_A = 0.5b_B$ among half of the network's agents yields little effect on spin state changes. Therefore, very limited information will be gained. This, in turn, results in very limited improvements in standard deviations, and approximated budget allocations stay constant. In contrast, for $b_A = b_B$, the approximated solutions initially allocate budgets to a smaller number of agents, inducing state flips. Unlike the case of $b_A = 0.5b_B$, this avoids getting stuck in a frozen state, leading to significantly lower relative mean standard deviation. To further illustrate the reasons for differences in standard deviations calculated by the high-temperature Taylor series approximation for $b_A = 0.5b_B$ and $b_A = b_B$ in the low-temperature region, we present the average number of state changes after applying the approximated budget allocations, i.e., $\sum_{i=1}^{N}\sum_{t=T_0}^{T_1-1}\frac{|s_i(t+1)-s_i(t)|}{2N(T_1-T_0)}$ in Fig. 9(b). The results in Fig. 9(b) are consistent with our previous analysis: For $b_A = 0.5b_B$ and $T < 4$, the average number of state changes are less than 10^{-2}. Therefore, little information can be obtained by observing the system, which results in large relative mean standard deviations exceeding 10^2 in Fig. 3.

References

[1] H.C. Nguyen, R. Zecchina, J. Berg, Inverse statistical problems: from the inverse Ising problem to data science, Adv. Phys. 66 (3) (2017) 197–261, http://dx.doi.org/10.1080/00018732.2017.1341604.

[2] S. Cocco, S. Leibler, R. Monasson, Neuronal couplings between retinal ganglion cells inferred by efficient inverse statistical physics methods, Proc. Natl. Acad. Sci. 106 (33) (2009) 14058–14062, http://dx.doi.org/10.1073/pnas.0906705106.

[3] T.R. Lezon, J.R. Banavar, M. Cieplak, A. Maritan, N.V. Fedoroff, Using the principle of entropy maximization to infer genetic interaction networks from gene expression patterns, Proc. Natl. Acad. Sci. 103 (50) (2006) 19033–19038, http://dx.doi.org/10.1073/pnas.0609152103.

[4] D. Fajardo, L.M. Gardner, Inferring contagion patterns in social contact networks with limited infection data, Netw. Spat. Econ. 13 (4) (2013) 399–426, http://dx.doi.org/10.1007/s11067-013-9186-6.

[5] D. Sornette, Physics and financial economics (1776–2014): puzzles, Ising and agent-based models, Rep. Progr. Phys. 77 (6) (2014) 062001, http://dx.doi.org/10.1088/0034-4885/77/6/062001.

[6] S. Myers, J. Leskovec, On the convexity of latent social network inference, Adv. Neural Inf. Process. Syst. 23 (2010) http://dx.doi.org/10.5555/2997046.2997090.

[7] Y. Yang, T. Luo, Z. Li, X. Zhang, P.S. Yu, A robust method for inferring network structures, Sci. Rep. 7 (1) (2017) 1–12, http://dx.doi.org/10.1038/s41598-017-04725-2.

[8] A. Braunstein, A. Ingrosso, A.P. Muntoni, Network reconstruction from infection cascades, J. R. Soc. Interface 16 (151) (2019) 20180844, http://dx.doi.org/10.1098/rsif.2018.0844.

[9] D.-T. Hoang, J. Song, V. Periwal, J. Jo, Network inference in stochastic systems from neurons to currencies: Improved performance at small sample size, Phys. Rev. E 99 (2) (2019) 023311, http://dx.doi.org/10.1103/PhysRevE.99.023311.

[10] C. Castellano, S. Fortunato, V. Loreto, Statistical physics of social dynamics, Rev. Modern Phys. 81 (2) (2009) 591–646, http://dx.doi.org/10.1103/RevModPhys.81.591.
[11] Y. Roudi, J. Hertz, Mean field theory for nonequilibrium network reconstruction, Phys. Rev. Lett. 106 (4) (2011) 048702, http://dx.doi.org/10.1103/PhysRevLett.106.048702.
[12] M. Mézard, J. Sakellariou, Exact mean-field inference in asymmetric kinetic Ising systems, J. Stat. Mech. Theory Exp. 2011 (07) (2011) L07001, http://dx.doi.org/10.1088/1742-5468/2011/07/L07001.
[13] H. Kappen, J. Spanjers, Mean field theory for asymmetric neural networks, Phys. Rev. E 61 (5) (2000) 5658–5663, http://dx.doi.org/10.1103/PhysRevE.61.5658.
[14] H.-L. Zeng, E. Aurell, M. Alava, H. Mahmoudi, Network inference using asynchronously updated kinetic Ising model, Phys. Rev. E 83 (4) (2011) 041135, http://dx.doi.org/10.1103/PhysRevE.83.041135.
[15] B. Dunn, Y. Roudi, Learning and inference in a nonequilibrium Ising model with hidden nodes, Phys. Rev. E 87 (2) (2013) 022127, http://dx.doi.org/10.1103/PhysRevE.87.022127.
[16] L. Bachschmid-Romano, M. Opper, Inferring hidden states in a random kinetic Ising model: replica analysis, J. Stat. Mech. Theory Exp. 2014 (6) (2014) P06013, http://dx.doi.org/10.1088/1742-5468/2014/06/P06013.
[17] C. Battistin, J. Hertz, J. Tyrcha, Y. Roudi, Belief propagation and replicas for inference and learning in a kinetic Ising model with hidden spins, J. Stat. Mech. Theory Exp. 2015 (5) (2015) P05021, http://dx.doi.org/10.1088/1742-5468/2015/05/P05021.
[18] C. Campajola, F. Lillo, D. Tantari, Inference of the kinetic Ising model with heterogeneous missing data, Phys. Rev. E 99 (6) (2019) 062138, http://dx.doi.org/10.1103/PhysRevE.99.062138.
[19] S. Lee, V. Periwal, J. Jo, Inference of stochastic time series with missing data, Phys. Rev. E 104 (2) (2021) 024119, http://dx.doi.org/10.1103/PhysRevE.104.024119.
[20] P. Zhang, Inference of kinetic Ising model on sparse graphs, J. Stat. Phys. 148 (3) (2012) 502–512, http://dx.doi.org/10.1007/s10955-012-0547-1.
[21] A. Decelle, F. Ricci-Tersenghi, P. Zhang, Data quality for the inverse lsing problem, J. Phys. A 49 (38) (2016) 384001, http://dx.doi.org/10.1088/1751-8113/49/38/384001.
[22] Z. Cai, E. Gerding, M. Brede, Control meets inference: Using network control to uncover the behaviour of opponents, Entropy 24 (5) (2022) 640, http://dx.doi.org/10.3390/e24050640.
[23] B. Min, M. San Miguel, Competing contagion processes: Complex contagion triggered by simple contagion, Sci. Rep. 8 (1) (2018) 1–8, http://dx.doi.org/10.1038/s41598-018-28615-3.
[24] D. Centola, How Behavior Spreads: The Science of Complex Contagions, in: Princeton Analytical Sociology Series, vol. 3, Princeton University Press, Princeton, NJ, 2018, pp. 1–296.
[25] V.V. Vasconcelos, S.A. Levin, F.L. Pinheiro, Consensus and polarization in competing complex contagion processes, J. R. Soc. Interface 16 (155) (2019) 20190196, http://dx.doi.org/10.1098/rsif.2019.0196.
[26] V. Sood, S. Redner, Voter model on heterogeneous graphs, Phys. Rev. Lett. 94 (17) (2005) 178701, http://dx.doi.org/10.1103/PhysRevLett.94.178701.
[27] T. Aoyama, T. Matsuo, Y. Shibusa, Improved Taylor expansion method in the Ising model, Progr. Theoret. Phys. 115 (3) (2006) 473–486, http://dx.doi.org/10.1143/PTP.115.473.
[28] C. Lynn, D.D. Lee, Maximizing influence in an ising network: A mean-field optimal solution, Adv. Neural Inf. Process. Syst. 29 (2016) http://dx.doi.org/10.5555/3157096.3157375.
[29] S. Galam, Sociophysics: A review of Galam models, Internat. J. Modern Phys. C 19 (03) (2008) 409–440, http://dx.doi.org/10.1142/S0129183108012297.
[30] R.J. Glauber, Time-dependent statistics of the Ising model, J. Math. Phys. 4 (2) (1963) 294–307, http://dx.doi.org/10.1063/1.1703954.
[31] I.J. Myung, Tutorial on maximum likelihood estimation, J. Math. Psych. 47 (1) (2003) 90–100, http://dx.doi.org/10.1016/S0022-2496(02)00028-7.
[32] A. Ly, M. Marsman, J. Verhagen, R.P. Grasman, E.-J. Wagenmakers, A tutorial on Fisher information, J. Math. Psych. 80 (2017) 40–55, http://dx.doi.org/10.1016/j.jmp.2017.05.006.
[33] B. Efron, D.V. Hinkley, Assessing the accuracy of the maximum likelihood estimator: Observed versus expected Fisher information, Biometrika 65 (3) (1978) 457–483, http://dx.doi.org/10.2307/2335893.
[34] D.P. Bertsekas, Constrained Optimization and Lagrange Multiplier Methods, Academic Press, 2014.
[35] S.-J. Kim, K. Koh, M. Lustig, S. Boyd, D. Gorinevsky, An interior-point method for large-scale ℓ_1-regularized least squares, IEEE J. Sel. Top. Sign. Proces. 1 (4) (2007) 606–617, http://dx.doi.org/10.1109/JSTSP.2007.910971.
[36] M. Catanzaro, M. Boguná, R. Pastor-Satorras, Generation of uncorrelated random scale-free networks, Phys. Rev. E 71 (2) (2005) 027103, http://dx.doi.org/10.1103/PhysRevE.71.027103.
[37] D. Chandler, Introduction to Modern Statistical Mechanics, Vol. 5, Oxford University Press, 1987, p. 449.
[38] Y.-Z. Chen, Y.-C. Lai, Sparse dynamical Boltzmann machine for reconstructing complex networks with binary dynamics, Phys. Rev. E 97 (3) (2018) 032317, http://dx.doi.org/10.1103/PhysRevE.97.032317.
[39] G. Romero Moreno, S. Chakraborty, M. Brede, Shadowing and shielding: Effective heuristics for continuous influence maximisation in the voting dynamics, Plos One 16 (6) (2021) e0252515, http://dx.doi.org/10.1371/journal.pone.0252515.