UNIVERSITY OF SOUTHAMPTON

Faculty of Engineering and Physical Sciences
School of Electronics and Computer Science

Influence Maximization and Inference of Opinion Dynamics on Complex Networks

by

Zhongqi Cai

ORCiD: 0009-0005-0854-6942

A thesis for the degree of
Doctor of Philosophy

January 2024

University of Southampton

Abstract

Faculty of Engineering and Physical Sciences
School of Electronics and Computer Science

Doctor of Philosophy

Influence Maximization and Inference of Opinion Dynamics on Complex Networks

by Zhongqi Cai

Influence maximization has commonly been studied in the context of strategically allocating resources to agents in a network to maximize the spread of an opinion. In the first part of this thesis, we relate influence maximization to network control and consider a version of an inter-temporal influence maximization problem. Specifically, we study the competition between two external controllers with fixed campaign budgets. In this competition, one or both of the controllers have the flexibility to determine when to start control in order to maximize the share of a desired opinion in a group of agents who exchange opinions on a social network subject to voting dynamics. We first investigate the inter-temporal influence maximization in a constant-opponent setting where an active controller maximizes its vote share against a known-strategy opponent. Starting with a simplified model where influence starts to be allocated to all agents at the same time, we find that, for short time horizons, maximum influence is achieved by starting relatively later on more heterogeneous networks than in more homogeneous networks, while the opposite holds for long time horizons. Furthermore, by comparing the vote shares achievable via the node-specific optimization where each agent has different starting times and budget allocations with the same-starting-time scenario, we show that strategic allocations of the node-specific optimization are more effective when fewer resources are available to the controller. Apart from the constant-opponent setting, we also explore the game-theoretical setting where both controllers compete to maximize their influence. We find that the controller with budget superiority will start earlier, while the other controller will concentrate its resources in a later stage in order to use limited resources more effectively and achieve some pull-back from its opponent's initial advantage.

In the above, we have assumed either complete or no information about the opponent's strategy in solving the influence maximization problem. However, while the strategy will not always be known in real-world settings, we can infer information by observing the dynamics of opinion exchanges. Furthermore, instead of only passively observing the dynamics for inference, we are interested in influencing the opinion dynamics in such a way that observations can be improved. For this purpose, we propose a framework of strategically influencing a dynamical process with the aim of making hidden parameters more easily inferable. More specifically, we consider a model of networked agents who exchange opinions subject to voting dynamics. Agent dynamics are subject to peer influence and to the influence of two controllers. One of these controllers is treated as passive and we presume its influence is unknown. We then consider a scenario in which the other active controller attempts to infer the passive controller's influence from observations. We explore how the active controller can strategically deploy its own influence to manipulate the dynamics with the aim of accelerating the convergence of its estimates of the opponent. Along with benchmark cases, we then propose two heuristic algorithms for designing optimal influence allocations. We first demonstrate that it is possible to accelerate the inference process by strategically interacting with the network dynamics. Investigating configurations in which optimal control is deployed, we find that agents with higher degrees and larger opponent allocations are harder to predict. Furthermore, even when applying the heuristic algorithms, opponent's influence is typically the harder to predict the more degree-heterogeneous the social network.

In the third part of the thesis, we extend the previous framework from the voter model to the Ising model. Different from the linearity of the voter model, which results in high levels of mathematical tractability, the non-linearity of the Ising model requires different techniques for analysis. By comparing to benchmark cases of equally targeting, we first demonstrate that it is also possible to accelerate the inference by strategically interacting with the non-linear Ising dynamics. We then apply series expansions to obtain an approximation of the optimized influence configurations in the high-temperature region. By using mean-field estimates, we demonstrate the applicability of the method in a more general scenario where real-time tracking of the system is infeasible. Last, by analyzing 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, so as to even out the inaccuracy for inferring larger values.

Contents

List of Figures

List of Tables

Declaration of Authorship

I declare that this thesis and the work presented in it is my own and has been generated by me as the result of my own original research.

I confirm that:

1. This work was done wholly or mainly while in candidature for a research degree at this University;
2. Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
3. Where I have consulted the published work of others, this is always clearly attributed;
4. Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
5. I have acknowledged all main sources of help;
6. Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
7. Parts of this work have been published as: Zhongqi Cai, Markus Brede, and Enrico Gerding. Influence maximization for dynamic allocation in voter dynamics. In Rosa M. Benito, Chantal Cherifi, Hocine Cherifi, Esteban Moro, Luis Mateus Rocha, and Marta Sales-Pardo, editors, *Complex Networks & Their Applications IX*, pages 382–394, Cham, 2021. Springer International Publishing. ISBN 978-3-030-65347-7
   Zhongqi Cai, Enrico Gerding, and Markus Brede. Accelerating opponent strategy inference for voting dynamics on complex networks. In Rosa Maria Benito, Chantal Cherifi, Hocine Cherifi, Esteban Moro, Luis M. Rocha, and Marta Sales-Pardo, editors, *Complex Networks & Their Applications X*, pages 844–856, Cham, 2022a. Springer International Publishing. ISBN 978-3-030-93409-5
   Zhongqi Cai, Enrico Gerding, and Markus Brede. Control meets inference: Using network control to uncover the behaviour of opponents. *Entropy*, 24(5), 2022b. ISSN 1099-4300. . URL https://www.mdpi.com/1099-4300/24/5/640

Signed:..... Date:.....

Acknowledgements

First and foremost, I would like to express my deepest gratitude to my supervisors, Dr. Markus Brede and Prof. Enrico Gerding, for their invaluable guidance, support, and encouragement throughout my Ph.D. research. I am tremendously grateful for the time spent discussing ideas, troubleshooting issues, and pushing me to grow as a researcher.

In addition, I wish to acknowledge the funding sources from the China Scholarship Council that made this research possible. Their financial support allowed me to fully dedicate time and energy toward this project.

Finally, my deepest thanks go to my family and friends for their unconditional love, patience, and belief in me. I am profoundly grateful to my parents for instilling in me a passion for learning and discovery. This thesis would not have been possible without all of you.

Chapter 1

Introduction

With hundreds of millions of users worldwide, the prevalence of online social media has facilitated information diffusion and led to a growing influence on opinion formation (Xiong and Liu, 2014). As a result of the small-world and scale-free structure of online social networks (Dorogovtsev and Mendes, 2002), opinions by any user in the networks can become influential through replicating and forwarding. This revolutionized mass opinion formation process further invokes concerns about the spreading of negative influences, since, while providing new channels to guide and influence people for public benefits (e.g., health (Korda and Itani, 2013), education (Tess, 2013)), the increasing use of social media has also led to a wider spread of fake news (van der Linden et al., 2020) and misinformation (Loomba et al., 2021).

Owing to the considerable social and commercial impacts of influencing people on a large scale, the intervention in influence propagation through social networks has already been shown to be powerful in practical applications (Guille et al., 2013). One of the applications of governmental intervention on influence propagation is proposed by Yadav et al. (2016), which strategically selects the most influential participants to maximally spread influence to raise awareness of HIV. More recently, issues of vaccine hesitancy fueled by anti-vaccination efforts towards COVID-19 vaccination programs on social media have found much attention (Puri et al., 2020; Chadwick et al., 2021). Outside of the public benefit field, there are also plenty of high-impact applications in the commercial and political domains (Azaouzi et al., 2021), where powerful agents can leverage their influence to shape and dominate prevailing views through communication channels. Typical examples include viral marketing (Razali et al., 2023; Goyal et al., 2019), political manipulation (McFaul and Kass, 2019; Badawy et al., 2018), and the propagation of radical opinions and formation of radicalization (Galam and Javarone, 2016). Consequently, it is of great importance to understand the intricate dynamics stemming from external influence exerted by influential agents, so as to prevent manipulation or guide public opinions.

A canonical setting for modelling opinion propagation intervention in social networks is known as *influence maximization* (IM) (Kempe et al., 2003). The concept of influence maximization was first proposed by Domingos and Richardson (Domingos and Richardson, 2001; Richardson and Domingos, 2002) and has since garnered significant attention. IM involves modeling one or more external influences, often referred to as *external controllers* (Masuda, 2015), that strategically allocate resources (e.g. money, information, incentives) to agents in the network. The objective is to maximize the expected number of agents adopting their opinions, also known as the *expected spread* (Goyal et al., 2013). For example, advertisers may offer free products or discounts to a number of selected agents to spread brand awareness or drive the adoption of products through word-of-mouth marketing. However, achieving this objective necessitates operating within a budget constraint, as resources available for allocations to agents are inherently limited (Kempe et al., 2003; Bharathi et al., 2007). By employing the framework of IM, it enables a comprehensive exploration of optimal implementation strategies, including the strategic manipulation of opinion dynamics as well as the identification of effective preventive measures against manipulation (Chen et al., 2022). Additionally, the framework provides ways to understand the effects arising from dynamic interactions between external controllers and internal agents (Alshamsi et al., 2018).

Usually, the IM problem is explored under the assumptions that the underlying networks over which the opinion spreads are fully observed or we have prior knowledge about the opponents' strategies. However, in many realistic situations, the network structures or opponents' strategies are unknown and can only be obtained indirectly from observations of the spreading process (Hoang et al., 2019). Therefore, how to infer the network topology or reconstruct opponents' strategies from observed system dynamics, is also a fundamental problem in network science (Casadiego et al., 2017), and is also part of our research topics. Moreover, the issue of inferring structural and modelling parameters of complex-networked systems from observed system dynamics has gained significant attention in a variety of research domains. Applications range from discovering neural connection networks from observed neural spike data (Cocco et al., 2009), uncovering genetic interaction networks from gene expression level data (Lezon et al., 2006), reconstructing three-dimensional protein structure from amino acids sequences (Morcos et al., 2011), revealing contact networks from infection data (Fajardo and Gardner, 2013), to inferring online social networks from information flows (Myers and Leskovec, 2010).

Despite a large amount of research devoted to tackling influence maximization and network inference, several challenges remain to be addressed in these fields. In the rest of this chapter, we detail the research challenges and our contributions towards addressing parts of these challenges, and present the organization of the rest of this thesis.

1.1 Research Challenges and Questions

As described in the preceding section, influence maximization serves as a prevailing framework for examining the interplay between external opinion control and the dynamics of opinions within a social system. However, the majority of research in this field has focused on maximizing opinion spreads in equilibrium (stationary) states (Brede et al., 2019a), after dynamics have fully unfolded. Nonetheless, the potentially prolonged dynamics of the opinion propagation process contrasts with the need for real-world influence maximization to obtain optimal outcomes in shorter time horizons (Hegselmann et al., 2014). In this context, the transient dynamics before equilibrium matter, as a strategy optimized for equilibrium may be suboptimal early on. However, the transient analysis of the system dynamics in influence maximization is less researched and is needed in order to reveal how to allocate resources at specific times to steer the trajectory most effectively over short timescales.

In addition to the applicational challenges of the IM problem, the network inference problem also leaves unanswered questions. Specifically, most previous research on network inference uses passive inference approaches, which rely on a fixed dataset that has already been collected. There is no capability to dynamically choose which data to collect next based on what has been learned so far. For example, imagine we have made some initial inferences about a model's parameters from limited data. Passive methods cannot then actively decide to collect data from certain experiments that would help refine ambiguous parts of the model. Instead, they are constrained to make inferences from whatever static dataset happens to be available, which may be uninformative about key parameters. In contrast, by actively influencing dynamics during inference, data collection can be optimized adaptively, such as shifting focus to parameters with high uncertainty. However, methods for actively manipulating dynamics during inference to expedite the convergence of inference remain largely unexplored (see literature review in Section 2.5 for more details).

In the following, we enumerate the open research problems tackled in this thesis in more detail, and we identify a number of research questions for each.

1. **Leveraging transient dynamics for influence maximization through inter-temporal network control.** Compared to time-invariant resource allocations, the inter-temporal network control is more realistic for maximizing influence by leveraging transient dynamics. It allows controllers to allocate different amounts of resources over time in order to maximize their influence during the transient period, which can more efficiently exploit transient effects and take advantage of

short-term opinion evolution before equilibrium. However, the effect of transient dynamics and inter-temporal network control are less researched (see Section 2.4.3 for more details), and a variety of questions remain to be addressed, specifically:

(a) How can inter-temporal control be integrated into the existing framework of the classic IM problem to leverage transient opinion dynamics?

(b) Depending on network structures and time horizons, what are the optimal inter-temporal allocation strategies during transient phases before equilibrium?

(c) How do controller strategies in IM differ when facing a constant adversary with a known strategy (constant-opponent setting) versus competing in an environment where strategies are unknown (game-theoretical setting), especially when considering inter-temporal network control? What methods can be used to compute these strategies within the game-theoretical framework, and under what conditions can a Nash equilibrium be established?

2. **Actively manipulating dynamics to expedite the convergence of inference.** In the typical setting investigated in the literature of network inference, observational data for reconstructing network structure and inferring parameters of dynamical processes are assumed to be given and fixed. Little attention has been put on the aspect of enhancing the inference performance by improving the dataset's quality to obtain more accurate estimators with less data for social network inference. However, this approach has shown promise in biology (Ud-Dean and Gunawan, 2016), where the optimal design of gene knockout experiments reduces uncertain gene interactions and improves gene regulatory network inference. This motivates reframing the inference problem from an experimental design perspective: Can we shift the focus from purely algorithmic methods of passively inferring from a given dataset to an experimental design perspective, where intervention in network dynamics produces a dataset that allows for the most accurate inference from a constrained set of observations? How will manipulating agent interactions shape the data generation process, thus rendering the network dynamics into states that more readily expose the underlying parameters? Addressing these questions yields several specific research inquiries:

   (a) How does interacting with system dynamics through network control contribute to accelerating the convergence of inference strategies?

   (b) Within this framework, what are the methods to optimize the distribution of budget allocations for one controller in order to hasten the inference of another controller's strategy?

(c) Considering the diversity of opinion diffusion models, is the accelerated inference strategy universally viable, or does its applicability vary across different models? More specifically, how does the transition from linear to nonlinear opinion diffusion models affect the success of the accelerated inference approach?

(d) What constitutes the optimal configuration of budget allocations to effectively manipulate network dynamics, thereby generating a dataset that more accurately reveals an opponent's strategy?

(e) How do these optimal budget allocations depend on the particular network topologies and the constraints of budget availability?

This thesis aims to tackle the aforementioned research questions. The subsequent section will outline the specific contributions that arise from our work.

1.2 Research Contributions

This thesis explores the effects of external influence on opinion dynamics in two key scenarios. One is the context of influence maximization, where we study the temporal aspects of IM with the aim of finding out the system behaviours and optimal allocation strategies subject to inter-temporal network control. We also examine how these outcomes depend on factors like agent heterogeneity, network topology, and budget availability. The other scenario is about how we can best utilize the external influence to interact with the opinion dynamics to generate a more informative dataset for inference. Through these investigations, our aim is to gain deeper insights into mathematical properties of social systems, as well as unravelling the intricate interplay that unfolds between these systems and external sources of influence.

More specifically, we make the following contributions:

1. We present the first study of utilizing inter-temporal network control to achieve influence maximization over complex networks. Specifically, in the context of inter-temporal influence maximization, we assume that one or both controllers have the flexibility to determine when to start control in order to achieve maximum expected spread subject to time and budget constraints. Then, we consider a more complicated scenario where controllers optimize starting times and budget allocations for individual agents in the network, which is also referred to as *individual optimization*. This addresses Research Challenge 1a, as detailed in Section 1.1.

2. We conduct a comprehensive analysis of optimal starting times for the active controller in the constant-opponent setting, taking into account network topologies,

budget availability, and time horizons. To achieve this, we employ the heterogeneous mean-field method (Dezső and Barabási, 2002) and Taylor expansion techniques to derive estimates for the network's influence propagation timescales. This addresses Research Challenge 1b and parts of Research Challenge 1c, as detailed in Section 1.1.

Contributions 1 and 2 are described in the following paper:

Cai, Z., Brede, M. and Gerding, E. Influence maximization for dynamic allocation in voter dynamics. In Complex Networks and Their Applications IX: Volume 1, Proceedings of the Ninth International Conference on Complex Networks and Their Applications COMPLEX NETWORKS 2020 (pp. 382-394). Springer International Publishing.

3. We present the first investigation of inter-temporal influence maximization in the game-theoretical setting where both controllers compete to maximize their influence without prior knowledge about their opponent's strategy. Specifically, to study the game-theoretic aspects of competitive influence maximization, we use the iterative searching algorithm to find Nash equilibria. By doing so, we show the differences in the strategies of controllers for inter-temporal influence maximization between considering a known-strategy opponent and an unknown-strategy opponent. Moreover, we provide a possible way to numerically calculate the strategies of controllers in the game-theoretical setting. This addresses Research Challenge 1c, as detailed in Section 1.1.

4. We present the first study of investigating the network information inference from the perspective of accelerating the estimation of the opponent's strategy by optimally deploying the active controller's budget allocations. More specifically, we model the opinion propagation process for an individual agent in the network as a non-homogeneous Markov chain and further derive estimators of the opponent's strategy via maximum likelihood estimation. We also provide uncertainty quantification of our estimators by using the variance deduced from the expectation of the second-order derivative of the likelihood function. This, in turn, is used to inform decisions on the optimal allocations and understand the process of inference acceleration. This addresses Research Challenges 2a and 2b, as detailed in Section 1.1.

5. We provide a systematic investigation of how to optimally deploy resources in order to maximally accelerate the opponent strategy inference in the voter model, which is a commonly-used linear model in opinion dynamics for characterizing interactions between individuals. A detailed introduction for the voter model can be found in Section 2.2.2. To this end, we propose several novel heuristic algorithms for speeding up opponent strategy inference via minimizing the variance of estimators. Moreover, we verify the effectiveness of our algorithms in a variety

of scenarios with varying network topologies and budget availability. This addresses Research Challenges 2d and 2e, as detailed in Section 1.1.

Contributions 4 and 5 are described in the following two papers:

Cai, Z., Gerding, E., and Brede, M. Accelerating Opponent Strategy Inference for Voting Dynamics on Complex Networks. In Complex Networks and Their Applications X: Volume 1, Proceedings of the Tenth International Conference on Complex Networks and Their Applications COMPLEX NETWORKS 2021 (pp. 844-856). Springer International Publishing.

Cai, Z., Gerding, E., and Brede, M. (2022). Control Meets Inference: Using Network Control to Uncover the Behaviour of Opponents. Entropy, 24(5), 640.

6. After investigating the problem of inference acceleration in the linear voter model, it is of interest if similar ideas can be applied in relatively complex settings of non-linear models such as the Ising model. Therefore, we extend this idea and present the first study of accelerating the inference of the opponent's strategy in the Ising model. Different from the linearity of the voter model, which results in high levels of mathematical tractability, the non-linearity of the Ising model requires different techniques for analysis. To achieve this, we combine numerical results obtained from newly-proposed heuristics with detailed analytical explanations obtained via the Taylor series approximation to gain a deeper understanding for the structure of optimal allocation strategies in the high-temperature region. Furthermore, we analyze the performance of the heuristics for generating the optimal allocations in the scenario where the real-time tracking of the system states is inaccessible via mean-field approximations. This addresses Research Challenge 2c, as detailed in Section 1.1. This contribution is described in the following paper:

   Cai, Z., Gerding, E., and Brede, M. Accelerating Convergence of Inference in the Inverse Ising Problem. Under review in Physica A: Statistical Mechanics and its Applications.

1.3 Report Outline

The structure of this thesis is as follows. Chapter 2 covers the background needed to understand work in the fields of influence maximization and network inference, as well as state-of-the-art frameworks and techniques in these areas. Chapter 3 formalizes the influence maximization problem under the framework of inter-temporal network control based on voter dynamics, and provides both analytical and numerical solutions to solve the optimization problem. Chapter 4 investigates the reconstruction of the opponent's strategy from the perspective of accelerating the estimation of the opponent's strategy by optimally deploying the active controller's budget allocations in the voter model. Chapter 5 extends the context of reconstructing the opponent's strategy in the

Ising model. The report concludes with a summary and discussion of future work in Chapter 6.

Chapter 2

Background Theory and Literature Review

In this chapter, we conduct a thorough literature review, specifically focusing on influence maximization and network inference. In more detail, we start with an introduction for networks and the main network models in Section 2.1. Then, in Section 2.2, we review the most commonly-used opinion formation models with a special emphasis on the models that are employed in the remainder of the manuscript. Next, in Section 2.3, we discuss ways of how external influences can interact with the internal network dynamics. Section 2.4 covers the frameworks and techniques on influence maximization and points out the gaps in current research, with a special focus on the time-critical influence maximization problem. Section 2.5 presents state-of-the-art methodologies for network inference in the domains of epidemiology and information flow, and then delves into the widely adopted inverse Ising model, renowned in the statistical physics community for studying network inference from opinion dynamics.

2.1 Network Properties and Models

Social systems are complex webs of interactions and relationships between individuals. The evolution of individuals' opinions, cultural and linguistic characteristics, social standing, and other dynamic attributes are profoundly shaped by their locations in a social structure (Borgatti et al., 2018). To understand and model the structure and dynamics of social systems, networks provide a powerful representational framework that goes beyond simply analyzing attributes of individuals to examining the relational ties and patterns that connect them (Knoke and Yang, 2019).

Adopting the network perspective, a society is represented as a graph $G = \{V, E\}$, where individuals are represented by vertices (nodes) $v_i \in V$, and the edge $e_{ij} \in E$ represents an existing social interaction between v_i and v_j. Moreover, *neighbors* of node i refer to nodes that are directly connected to node i. To account for interaction intensity, weights w_{ij} can be assigned to each existing edge e_{ij}, while, non-existing links are modelled as $w_{ij} = 0$. These weights populate the weighted adjacency matrix $\boldsymbol{W} = \{w_{ij}\}$, signifying linking strengths and rendering the graph weighted. Unweighted graphs can also be captured in the weighted adjacency matrix $\boldsymbol{W}$, with binary elements $w_{ij} = 1$ or $w_{ij} = 0$ reflecting the presence or absence of an edge between nodes i and j.

Furthermore, in undirected graphs, the weighted adjacency matrix is symmetric, indicating a bidirectional connection between nodes i and j. However, in directed graphs, connections can be directed to indicate asymmetry in influence between individuals and will result in asymmetry in the weighted adjacency matrix. Moreover, by summing up the weights that are directly connected to a particular node i, we obtain the degree of node i, denoted as

$$k_i = \sum_j w_{ji}. \tag{2.1}$$

The node degree reflects the number of immediate weighted neighbors a node has [1]. Beyond node degrees, several key metrics are used to quantify properties of network topology (Newman, 2018):

Degree Distribution: The degree distribution $p(k)$ gives the fraction of nodes with degree k. Networks with power-law degree distribution $p(k) \sim k^{-\lambda}$ are called scale-free networks (Barabási and Bonabeau, 2003), where λ represents the degree exponent. This implies the existence of hubs – nodes with significantly more connections than others. Many real-world networks have been reported to be scale-free (Clauset et al., 2009).

Degree Assortativity: The assortativity coefficient r is commonly applied to unweighted networks with discrete integer degree values. It measures the degree correlation between nodes on either end of an edge and quantifies the preference for high-degree nodes to attach to other high-degree nodes (assortative mixing, $r > 0$) versus low-degree nodes attaching to high-degree nodes (disassortative mixing, $r < 0$). The assortativity coefficient can be calculated using the following formula (Newman, 2003a):

$$r = \frac{\sum_i q_{ii} - \sum_i u_i y_i}{1 - \sum_i u_i y_i}, \tag{2.2}$$

where q_{ii} is the fraction of edges that connect nodes of degree i to other nodes of the same degree i. This represents the probability that an edge links two nodes with identical degrees. In the denominator, u_i is the fraction of edges that start from a node with

[1] Note that, some literature (Newman, 2003b; Lewis, 2011) distinguishes "degree" for unweighted networks and "strength" for weighted ones. Additionally, in directed networks, we can differentiate between out-strength $k_i^{out} = \sum_j w_{ij}$ (the sum of outgoing edge weights) and in-strength $k_i^{in} = \sum_j w_{ji}$ (the sum of incoming edge weights).

degree i, and y_i is the fraction of edges that end at a node with degree i. This encodes the expected frequency of edges between nodes of degree i and any other degree, if connections were made randomly.

To gain an insight into how the above-mentioned properties of network topology influence social dynamics, several algorithmic frameworks are proposed for generating synthetic networks that exhibit similar structural characteristics to real-world networks. These network models define how nodes and edges are added, connected, or modified, thereby shaping the resulting network structure. Utilizing these network models allows for controlled experiments and investigations into various network properties that closely resemble real networks in structure and characteristics. In the following, we list some of the most widely used network models and their distinct characteristics (Newman, 2018).

Erdős-Rényi (ER) Model: The ER model (Erdős et al., 1960) has two primary variants – the $G(N, p)$ model and the $G(N, m)$ model. In the $G(N, p)$ model, a network of N nodes is formed by connecting each pair of nodes randomly with probability p. This generates a random graph with approximately $pN(N-1)/2$ edges distributed randomly. The $G(N, m)$ model, on the other hand, fixes the number of edges to m and connects pairs of nodes uniformly at random until m edges have been placed. Due to its mathematical tractability, the ER model allows analytical calculations of network properties (Erdős and Rényi, 1959) and is often used as a baseline for comparing with more complex network structures. However, the ER model has limitations in reproducing important features of many real-world networks, including heavy-tailed power-law degree distributions (Barabási and Albert, 1999), highly clustered connections (Watts and Strogatz, 1998), and heterogeneous nodes or edges (Newman, 2003b).

Random Regular Graphs: To construct a random regular graph (Kim and Vu, 2003), nodes are first assigned k stubs (half-edges). These stubs are then randomly paired by connecting two unconnected stubs at a time until no stubs remain. This results in a randomized graph topology where all nodes have the same degree k. The controlled generation of random regular graphs allows systematic experiments isolating the impact of narrow degree distributions on network structure and function for modeling and analysis. However, many other topological features of real networks like heavy-tailed degree distributions, high clustering, and heterogeneity are still not captured in the random regular graph.

Barabási-Albert (BA) Model: The BA model (Albert and Barabási, 2002) generates networks with a power law degree distribution using a growth process with preferential attachment. Specifically, networks start with a small graph. Then, new nodes are iteratively added and connected to existing nodes with probability proportional to their degrees, resulting in high-degree hub nodes acquiring more links over time. By doing

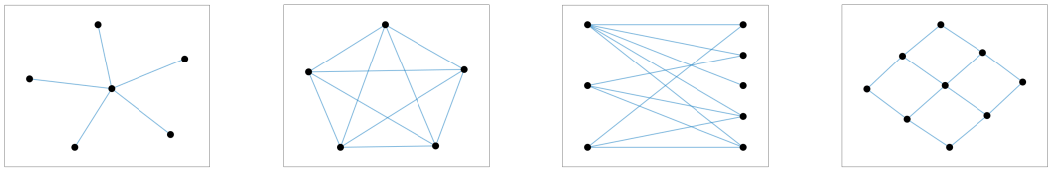

(a) (b) (c) (d)

FIGURE 2.1: The star graph (a); The complete graph (b); The bipartite graph (c); The two-dimensional lattice (d).

so, the BA model accurately captures the scale-free topology characteristic of numerous real-world networks, such as the Internet, social networks, and citation networks (Newman, 2005). Moreover, the BA model provides a way to systematically study the emergence and functions of hub nodes on network robustness and information spreading through the simple growth process (Liu et al., 2017). However, the BA model lacks structural properties like high clustering and community structure (Ravasz and Barabási, 2003). Variants of the BA model aim to address these limitations. For instance, the Holme-Kim model incorporates a triad formation step to produce clustered networks (Holme and Kim, 2002).

Configuration Model: The configuration model (Newman et al., 2001) generates networks with a given degree distribution by randomly connecting stubs of nodes. However, this random wiring process often creates invalid connections such as self-loops and multi-edges, which improperly distort network metrics. To address these issues, rewiring is needed to replace invalid connections while preserving the degree distribution. Moreover, the uncorrelated configuration model is an important variant that generates scale-free networks with a given power-law degree distribution (Catanzaro et al., 2005). In more detail, this model limits the maximum expected degree to avoid degree correlations. The resulting uncorrelated scale-free networks serve as useful null models for evaluating the analytical solutions of dynamic processes on complex networks (Váša and Mišić, 2022). Additionally, it enables controlled experiments on the impacts of varying power-law degree distributions versus random topology, and can be used to mimic scale-free structure observed in many real-world networks.

Watts-Strogatz (WS) Model: The WS model (Watts and Strogatz, 1998) generates small-world networks that exhibit both high local clustering and short average path lengths between nodes. The model starts with a regular lattice structure and randomly rewires a fraction of the edges to introduce randomness and create shortcuts. This rewiring allows for long-range connections and reduces path lengths while maintaining significant local clustering, which successfully replicates the small-world property observed in many real networks.

Apart from the aforementioned network models, in the following, we also present descriptions of commonly employed synthetic and simple network structures. These fundamental network topologies enable us to isolate and analyze the impact of simple network characteristics on opinion dynamics, influence maximization, and network inference.

Star Graphs: Star networks (see Fig. 2.1 (a)) are a specific type of network structure wherein a central node is directly connected to all other nodes, while the peripheral nodes are not connected to each other. This configuration forms a pattern resembling the shape of a star, hence the name. Star networks exhibit a high level of connectivity, with the central node serving as a focal point for communication and information flow. The central node plays a crucial role in coordinating and disseminating information to the peripheral nodes, making it an important hub within the network.

Complete Graphs: Complete graphs, as shown in Fig. 2.1 (b), are highly interconnected networks where every node is directly connected to all other nodes. Their distinct property is having the maximum possible number of edges, creating a dense structure. Complete graphs are commonly used as theoretical constructs and as a baseline for comparison in network analysis and algorithm design.

Bipartite Graphs: Bipartite graphs (see Fig. 2.1 (c)) consist of two sets of nodes, where edges only exist between nodes of different sets. These networks are generated by specifying the two sets of nodes and connecting nodes from one set to nodes in the other set. Bipartite graphs are useful for representing relationships between different types of entities, such as users and items in recommendation systems or actors and movies in a film network. The most distinct property of bipartite graphs is the bipartite structure, where nodes within the same set do not have direct connections, and all edges connect nodes from different sets.

Two-dimensional Lattices: Two-dimensional lattices (see Fig. 2.1 (d)) represent regular grid-like networks where nodes are arranged in a two-dimensional grid. Each node is connected to its neighboring nodes, resulting in a regular lattice structure. Two-dimensional lattices can be generated by specifying the grid size and connecting nodes to their adjacent neighbors. The most distinct property of two-dimensional lattices is their regularity and local connectivity. These networks are used to study spatially embedded systems, diffusion processes on grids, and lattice-based models in physics and biology (Kandasamy et al., 2014).

In summary, these network models and fundamental network structures provide valuable insights into various aspects of network theory, including connectivity, degree distributions, influence dynamics, and spatial relationships. By understanding the distinct properties of these models, researchers can analyze real-world networks, develop theoretical frameworks, and investigate the impact of network structures on the best strategies regarding influence maximization and network inference.

2.2 Diffusion Models

Building upon the understanding of network models, an essential component in studying opinion dynamics and influence maximization is the utilization of diffusion models. These models provide a framework for analyzing the spread of opinions, information, ideas, or behaviors within a network. By incorporating diffusion models into the context of network models, we can gain insights into the mechanisms through which opinions and influences propagate throughout a given network structure.

In more detail, diffusion models are designed to capture the dynamics of information dissemination and adoption across network nodes. In diffusion models, agents in the network interact with their directly connected neighbours according to a set of reasonable mathematical rules which represent social mechanisms responsible for opinion evolution. Resulting from the nontrivial collective effects of agent interactions, opinions propagate in the network and the opinion states of a population evolve accordingly (Li et al., 2018).

In the existing literature, the opinion diffusion models are divided into two categories (Chen et al., 2013): the progressive diffusion models (e.g., the independent cascade model, the linear threshold model) and the non-progressive diffusion models (e.g., the voter model, the susceptible-infected-susceptible (SIS) model, the Ising model). A major distinction between progressive and non-progressive models is whether the state of an agent can be changed back and forth due to agent interactions. In the following, we introduce the most commonly-used diffusion models for investigating the IM problem, classified within these two categories (Noorazar, 2020).

2.2.1 Progressive Diffusion Models

Progressive models assume that once a node adopts an opinion or behavior, it remains in that state indefinitely. Two commonly-used progressive models in the IM problem are the independent cascade model and the linear threshold model (Li et al., 2018).

The independent cascade model (Kempe et al., 2003) offers a way to model the influence propagation as a one-off activation. The continually-elapsing time is discretized as time steps in this model. Each connection between nodes has a chance of transmitting influence, represented by a probability, and each node can only be influenced once. The infected node will become inactive after one time step, which means that it can only propagate its influence to its neighbours during a single time slot. This one-hop dissemination of influence is *cascaded* throughout the entire network until all the nodes in the network are eventually infected, or the propagation stops when nodes fail to transform their neighbours due to the randomness of the transformed probability.

In the linear threshold model (Chen et al., 2010), an inactive node is also under the influence of opinions of its neighbours, but with a more granular rule. Specifically, in the linear threshold model, the probability on each directed edge is extended to a non-negative weight. For each node $v_i \in V$, the total incoming edge weights sum up to a maximum of 1. The rule specifically indicates that, at each time step, an inactive node v_i changes its state to "active" if and only if the total weight contributed by its incoming neighbours is at least θ_{v_i}. Here, thresholds θ_{v_i} characterize different levels of the tendency of the node v_i to adopt the opinion of its neighbours. Normally, the thresholds are determined randomly in classic settings (Kempe et al., 2003).

2.2.2 Non-progressive Diffusion Models

In contrast to progressive models, non-progressive models allow nodes to change states multiple times. Three seminal non-progressive models are the SIS model, the voter model, and the Ising model (Li et al., 2018).

The SIS model is originally an epidemiological framework used to study the spread of infectious diseases in complex networks (Xuan et al., 2020). In this model, individuals i within a network can exist in either a susceptible (denoted as $F_i(t) = 0$) or infected state (denoted as $F_i(t) = 1$) at time t. The infection process occurs when a susceptible individual becomes infected through contact with infected neighbors, akin to disease transmission. Specifically, the infection probability of node i is $\beta \sum_j w_{ij} F_j(t)$, where $\boldsymbol{W} = \{w_{ij}\}$ is the adjacency matrix, and β is the infection rate. Then infected individuals i can recover and revert to being susceptible again with recovery rate $\mu F_i(t)$, dependent on their own state only (Lee et al., 2013). While the SIS model is effective for studying disease spreads, it has an important difference compared to the typical opinion diffusion models like the voter model and Ising model in how state transitions occur. In opinion dynamics, transitions depend on the states of the node's neighbors in both directions, as opinions are influenced by peers in both directions – both in adopting and abandoning opinions. However, the SIS model lacks this realistic mutual influence capacity, since recoveries are spontaneous.

The voter model (Clifford and Sudbury, 1973) is a well-studied opinion dynamics model that has been prominent in the literature for decades (Redner, 2019). In the voter model, opinions are typically binary, e.g., an agent holding either opinion A or B. The underlying diffusion mechanism of the voter model is straightforward: At each time step, an agent is selected at random from the network, and then it copies the opinion of a randomly selected neighbor. Mathematically, the voter model can be described by tracking the evolution of the probability $x_i(t)$ of node i holding opinion A at time t (Mobilia et al., 2007):

$$\frac{dx_i}{dt} = (1 - x_i)\frac{\sum_j w_{ij} x_j}{k_i} - x_i \frac{\sum_j w_{ij}(1 - x_j)}{k_i}, \tag{2.3}$$

where w_{ij} represents the element in the i-th row and j-th column of the adjacency matrix, and k_i is the degree of node i. In Eq. (2.3), the first term stands for the probability of holding opinion B but selecting a neighbor with opinion A, while the second term stands for the probability of holding opinion A but choosing a neighbor with opinion B. Note that, in Eq. (2.3), the rates of opinion changes depend linearly on the sum of neighbors with the same states. This linear relationship enables tractable mathematical analysis of the model's properties (Redner, 2019). Additionally, on finite connected networks, the voting dynamics asymptotically reach a consensus where one opinion takes over the population. The specific outcome, whether converging towards opinion A or B, depends on the initial distribution of opinions and the underlying network structure (Even-Dar and Shapira, 2011). Notably, sparser connections dramatically slow consensus (Redner, 2019). Moreover, the voter model has been employed to model election outcomes. For example, empirical evidence of USA and UK election outcomes has supported the alignment of voting dynamics with actual voting data (Braha and De Aguiar, 2017; Vendeville et al., 2021; Fernández-Gracia et al., 2014).

The Ising model is one of the most popular models in the physics community used to study opinion dynamics and collective behavior in complex systems (Galam et al., 1982; Castellano et al., 2009). In the Ising model, each individual is represented as a spin (node) in a network, and their opinions are described as binary variables (e.g., $+1$ or -1). The model assumes that individuals are influenced by the opinions of their connected neighbors, and they tend to align their own opinions with those in their social vicinity. Following the commonly-used synchronous Glauber algorithm (Glauber, 1963; Galam, 2008) in simulating social dynamics, the Ising dynamics evolve according to the following probabilities:

$$\begin{aligned} Pr(s_i(t+1)=1) &= \frac{1}{1+e^{\left(-2\sum_j w_{ij}s_j(t)\right)/\tau}} \\ Pr(s_i(t+1)=-1) &= 1 - Pr(s_i(t+1)=1). \end{aligned} \tag{2.4}$$

Here, $Pr(s_i(t+1)=1)$ and $Pr(s_i(t+1)=-1)$ are the probabilities for spin i to have state 1 and -1 at time $t+1$, and τ is the temperature, controlling randomness. At low temperature τ, the tendency for individuals to align opinions with their neighbors typically overcomes random opinion fluctuations, leading to consensus and long-range order in the public opinion characterized by the emergence of majority opinion alignment across the population. However, above a critical value τ_c, individual opinion fluctuations dominate, preventing consensus and keeping the system in a disordered, fragmented state macroscopically (Castellano et al., 2009). Moreover, on two-dimensional lattice graphs, there is a sharp phase transition from an ordered phase to a disordered phase as the temperature passes the critical point τ_c. However, network heterogeneity changes the phase transition, which becomes less sharp, and the critical temperature

τ_c increases logarithmically with the network size (Aleksiejuk et al., 2002; Dorogovtsev et al., 2008).

Unlike the simple contagion observed in the voter model (Redner, 2019), where an agent can become "infected" with a certain probability after a single exposure to an infectious entity, the Ising model belongs to the class of complex contagion. In complex contagion, multiple exposures to infectious entities are necessary for an agent to change its state. In other words, the distinction between the voter model and the Ising model lies in how agents in the network respond to successive exposures (Min and San Miguel, 2018). Therefore, the Ising model provides a more accurate description of opinion adoption if the reinforcement of multiple sources of exposure is required for the adoption of a new idea (Centola, 2018; Vasconcelos et al., 2019). In Chapters 4 and 5, we address the network inference acceleration problem separately in the voter model and the Ising model. We first demonstrate that inference convergence can be accelerated by smartly targeting agents in the network using the voter model. As it is of interest if similar ideas can be applied in relatively complicated settings of complex contagion, we further investigate this problem in the Ising model. Different from the linearity of the voter model, which results in high levels of mathematical tractability, the non-linearity of the Ising model requires different techniques for analysis such as Taylor series approximation and mean-field approximation at certain temperature regions.

2.2.3 A Comparison between Progressive Models and Non-progressive Models

As mentioned in Section 2.2.1, the independent cascade model and the linear threshold model simulate the propagation of influence following a one-off activation, i.e., a node infected by an opinion remains infected by it. This assumption is appropriate for opinions that remain unchanged once decided, e.g., adopting a deeply-held belief. However, in many real-world settings, individuals may repeatedly flip their opinions back and forth due to peer and media influence, e.g., attitudes towards public or political issues. Since variants of the independent cascade and linear threshold models only allow a single activation for each node, they are unsuitable for modelling such opinion dynamics. Instead, such an opinion-changing process can be described by non-progressive models like the voter model (Redner, 2019) or Ising model (Lynn and Lee, 2016). Furthermore, the equilibrium or ground state in non-progressive models represents the self-optimizing tendency of social networks. As individuals interact, the system evolves toward configurations of minimal conflict and energy. This emergent convergence to stable states is absent from the independent cascade model. The self-organization in non-progressive models thus provides a more compelling analogy to opinion dynamics and consensus formation within social networks (Liu et al., 2010).

In Chapter 3, we choose the non-progressive voter model as the underlying opinion diffusion model for the IM problem for two main reasons. First, the voter model has been extensively studied in existing literature. For example, its basic properties, such as the exit probability (the likelihood that a given opinion takes over the entire network) and the consensus time (the expected time until a consensus is reached), have been analyzed on various network topologies (Redner, 2019; Castellano et al., 2009; Sood and Redner, 2005). The methodologies used in these analyses, like mean-field theory (Redner, 2019) and coalescing random walks (Yildiz et al., 2013), can provide insights into the inter-temporal IM problem. Second, the evolution of opinions in the voter model can be described by a linear system that can be solved analytically in some simple topologies (Masuda, 2015).

2.3 External Influence of Opinion Dynamics

In progressive models such as the classic independent cascade model and the linear threshold model, the diffusion process starts with a set of initially activated nodes in the network (known as the seed set) (Kempe et al., 2003). The role of external influence in progressive models is to identify the set of seed nodes in a network to trigger a cascading behavior or contagion process that maximizes diffusion.

For the voter model, with the inclusion of more than one external influence, mixed-equilibrium states where multiple opinions coexist will be reached. Specifically, some work aims at exploring the role of the external influence via either studying the scenario of strategically transforming agents into so-called *zealots*, where agents commit to a given opinion (Yildiz et al., 2013), or incorporating the IM problem with network control where the external influence is treated as an external controller who exerts influence by building unidirectional connections with agents in the network (Masuda, 2015). Contrary to the less realistic assumption that every agent can be converted into a zealot, in the network control model, the agent targeted by the external controller is also subject to peer influence since it has been realized that the individuals in the social network are exposed to various sources of social influence (Hu, 2017). Besides, it has been observed that it takes some time for the external influencer to influence individuals (Ali et al., 2018). Take the advertising campaign as an example. In order to achieve the purpose of product promotion, the advertisement needs to be broadcast repeatedly in a certain period of time to spread its influence. Therefore, we choose network control as our control model in the voter model when studying influence maximization and opponent strategy inference.

Originating from statistical physics and condensed matter physics, the Ising model is commonly used to study real-world systems that are not isolated and can be influenced

by external factors. To account for these external influences, the Ising model incorporates a concept known as "external fields" (Lynn and Lee, 2016). When introducing an external field in the Ising model, an additional term h_i is included in Eq. (2.4) to represent the interaction between the individual spin (node) i and the external field:

$$Pr(s_i(t+1)=1) = \frac{1}{1+e^{\left(-2\sum_j w_{ij}s_j(t)-2h_i\right)/\tau}} \tag{2.5}$$

By adjusting the strength and sign of the external field, one can effectively control the behavior of spins or nodes in the network, thereby manipulating network dynamics and guiding the system towards desired configurations. Similar to the voter model, the external fields can be viewed as building unidirectional connections from external controllers to agents in the network.

2.4 Existing Frameworks and Techniques for Influence Maximization

Since Domingos and Richardson first studied IM as an algorithmic problem (Domingos and Richardson, 2001; Richardson and Domingos, 2002), the IM problem has received extensive attention and been explored under different settings. We start our discussion of the related literature with the basic setting of IM where one controller in the network tries to maximize its influence in Section 2.4.1. Due to the limitation of the basic IM problem that it only considers one controller in a network, the problem has been extended into a competitive setting where multiple controllers are competing within a social network. Therefore, we proceed our discussion with the competitive IM problem in Section 2.4.2. As this work focuses on the inter-temporal allocations of budgets, we dedicate to introducing the related work and existing gaps in the time-critical IM field in Section 2.4.3.

2.4.1 Influence Maximization

Influence maximization has been first proposed for the independent cascade model (Kempe et al., 2003). This model has been formulated as a discrete optimization problem in which only one controller on a social network tries to maximize the spread of its opinion. Kempe et al. show that the IM problem is NP-hard and provide a greedy algorithm with a provable approximation guarantee (Kempe et al., 2003). As this proposed algorithm is computationally demanding for large-scale networks, extensive efforts have been made to improve its effectiveness in two directions. The first is introducing heuristics to accelerate the computation of the spread function for different seed sets (Leskovec et al., 2007; Chen et al., 2009). The second is restricting the search

to a small range of nodes as potential candidates for the next seed nodes (Goyal et al., 2011). As the evaluation function for the independent cascade model cannot be solved in polynomial time, the accuracy of these algorithms is guaranteed only by the submodularity of the evaluation function.

In addition to the vast use of heuristic strategies to identify the most influential seed nodes, some studies (Pan et al., 2016; Morone and Makse, 2015) map the diffusion process onto other well-studied problems. Specifically, Pan et al. (2016) apply a random walk to trace the most influential nodes and propose a linear-time complexity algorithm. Moreover, Morone and Makse (2015) translate the IM problem into optimal percolation, which aims to determine the minimal set of nodes whose removal from a network breaks down the system into unconnected clusters, and theoretically confirm that low-degree nodes surrounded by hub nodes are among the most influential nodes. Nevertheless, these methodologies are applied to progressive influence propagation models where the changes of states of agents in these models follow a progressive direction and cannot be reversed.

In contrast to progressive models that reach static states with no further cascades of influence propagation as mentioned in Section 2.2.1, the voter model allows nodes to change their opinions back and forth. As a result, in the voter model, all agents in a connected network will eventually align with the only existing external influencer in the asymptotic equilibrium (Even-Dar and Shapira, 2007). For this reason, the influence maximization for the voter model is usually explored against an opponent, which may be passive or actively maximizing their influence, and will be stated in detail in the next section.

2.4.2 Competitive Influence Maximization

Most studies of competitive IM either focus on the influence maximization problem for one of the competing controllers or study the game-theoretic aspects of competitive influence diffusion. For the former case, some research (Bharathi et al., 2007; Carnes et al., 2007; Budak et al., 2011) follows the idea of solving IM under the framework proposed by Kempe et al. (2003), and gives different heuristic algorithms for approximately computing the best response to an opponent's strategy based on independent-cascade-like models. Others apply the competitive IM to non-progressive diffusion models, utilizing greedy algorithms (Liu et al., 2010; Masuda, 2015; Brede et al., 2019b; Kuhlman et al., 2013; Yildiz et al., 2013) and gradient ascent algorithms (Lynn and Lee, 2016; Moreno et al., 2020; Romero Moreno et al., 2021b). Particularly, Masuda (2015) pioneers a linear-algebra-based approach for maximizing average vote shares in voter dynamics against a passive opponent. Moreover, Moreno et al. (2020) are the first to incorporate aspects of agent heterogeneity with influence maximization by modelling agents in the network with decreased probabilities of adopting a particular opinion. However, all of

these works only consider how to achieve optimal control in the stationary state. None of them takes into account the transient dynamics of the opinion propagation process and considers the non-stationary influence maximization.

For the game-theoretical setting, the competitive IM problem is mainly solved by finding the Nash equilibrium under a budget constraint in the stationary state (Goyal et al., 2019; Tzoumas et al., 2012; Fazeli and Jadbabaie, 2012; Kermani et al., 2017; Masucci and Silva, 2014; Chakraborty et al., 2019). Particularly, Goyal et al. (2019) study mixed Nash equilibria on the linear threshold model, and focus on identifying inefficiency of resource use at equilibrium. Kermani et al. (2017) obtain the pure Nash equilibrium by implementing numerical experiments in a real-world data set based on a variant of the independent cascade model. Although the game-theoretical IM problem based on the variant of the independent cascade model has been extensively investigated, it has received limited attention for non-progressive diffusion models. The most representative works applying game theory to solve IM on non-progressive models are Fazeli and Jadbabaie (2012); Masucci and Silva (2014); Chakraborty et al. (2019). Specifically, Fazeli and Jadbabaie (2012) represent the dynamics of influence spreading as a result of local coordination games in which each agent updates its state according to neighboring interaction. Based on this game-theoretic diffusion process, it describes the Nash equilibrium for the setting where two competitors simultaneously determine the number of their initial seeds with budget constraints to maximize the expected number of product adoption. Moreover, Masucci and Silva (2014) characterize the game-theoretical IM problem as optimal resource distribution in the voter model. However, both works treat the allocation of influence as a single time-step injection rather than continuous exertion of influence over a period of time. In other words, the controllers only affect the networks at the beginning of the campaign and only change the initial distribution of opinions. The exception considering continuous exertion of influence for non-progressive diffusion models is Chakraborty et al. (2019). They study the game-theoretical IM problem based on the voter model. However, their focus on optimal strategies is limited to a basic star topology, constraining its applicability to intricate real-world networks.

2.4.3 Time-critical Influence Maximization

Generally, the IM problem can be investigated from three orthogonal dimensions (Goyal et al., 2013): the number of seed nodes initially activated (known as the budget), the time required for the propagation, and the expected spread (i.e., the expected number of agents adopting a certain opinion). Recently, significant attention has been paid to the time-related influence maximization problem, which considers temporal factors, such as time minimization (Goyal et al., 2013; Alshamsi et al., 2017), and the time-constrained competition (Chen et al., 2012; Liu et al., 2012; Brede et al., 2019a; Ali et al.,

2018). Specifically, the time-minimization IM minimizes the time to achieve a preassigned expected spread under the budget constraint. While, for the time-constrained IM, it maximizes the expected spread under time and budget constraints.

For the time-minimization IM, Goyal et al. (2013) consider that the controller allocates all its budget at the beginning of the competition. Furthermore, Goyal et al. (2013) also show that the time-minimization IM problem is NP-hard and solve it by the greedy algorithm based on the independent cascade model. In contrast, Alshamsi et al. (2017) consider to target different individuals at different stages of influence propagation. They explore the dynamic strategy for time-minimization IM on an independent cascade model. Results show that the controller has to start targeting low-connectivity individuals and then switch to high-connectivity individuals at certain stages in order to minimize the time for achieving a predefined expected spread. On top of inter-temporally targeting, Tong et al. (2020) allow the controller to adjust its seeding strategy based on what it observes after a certain number of diffusion steps. However, all of the above work is addressed in a non-competitive setting where only a single external controller spreads its influence in the network. Given that competition for influence is also common in real-world contexts (e.g., political campaigns (Wilder and Vorobeychik, 2018) or radicalization prevention (Ramos et al., 2015)), the single-controller setting has a restricted range of applications.

For the time-constrained IM, Chen et al. (2012) incorporate the time delay aspect of influence diffusion. In more detail, they show that, if the influence delays follow the geometric distribution, the independent cascade model with delays maintains submodularity. Therefore, this problem can be solved by the greedy algorithm with a provable approximation guarantee. Furthermore, Liu et al. (2012) consider maximizing influence under the time constraint in the independent cascade model without influence delays and also prove monotonicity and submodularity of the influence spread function. However, these two works only consider the IM in the non-competitive setting where there is only one controller in the network, and assume that the controller allocates its budget only at the beginning of the campaign. To our best knowledge, only two works consider the time-constrained IM in the competitive setting. Most related to our modeling approach, Brede et al. (2019a) are the first to explore the effects of time horizons on IM in voter dynamics. However, their paper does not allow controllers to allocate different amounts of resources over time. It also does not consider the game-theoretical setting where competitors have no prior knowledge of the opponent's strategies. This contrasts with real-world scenarios like marketing, where the marketers can optimally distribute their budgets inter-temporally and only have little information of the competitor's strategy. The only directly related study that solves the time-constrained IM in the game-theoretical setting by considering when to initiate opinion propagation is by Ali et al. (2018). They use a technique called q-learning (Clifton and Laber, 2020) to model the inter-temporal seeding process, where an agent learns to select seed nodes by

iteratively updating its action values based on the rewards received in different states. However, to reduce the complexity of the q-learning framework, they only allow the controllers to select nodes from preassigned seed-combination strategies, which largely confines the flexibility of node selection. Besides, they focus on verifying the effectiveness of the q-learning framework from an algorithmic perspective and do not relate the mechanism behind the optimal strategies to transient dynamics of agents. Moreover, this work is based on the linear threshold model where, once an individual is influenced by a controller, its opinion will remain unchanged until the end of the diffusion process. In this sense, it is unsuitable for situations where opinions can be changed back and forth.

2.5 Network Structure Inference

As mentioned in Challenge 2 in Section 1.1, we are also interested in the question of how convergence of estimates can be accelerated through targeted interaction with the networked dynamics. To derive dynamical process parameters or reconstruct network topology from observational data, it is often necessary to draw on domain-specific expertise (Brugere et al., 2018). Here, we place the problem of speeding up inference in the context of opinion dynamics, with a specific focus on the rapid inference of unknown opponent's strategies. This focus is crucial for several reasons. First, it enhances models of real-world opinion dynamics, which involve undisclosed external influence attempts that conflict with common assumptions of opponent strategy transparency (Goyal et al., 2019). By quickly uncovering key aspects of such influence attempts, we can improve understanding of the actual drivers shaping public opinion formation (Myers et al., 2012). Additionally, with the swift propagation of manipulated messaging, quick identification of dissemination strategies is essential for effective counteraction before irreversible deception or polarization (Nguyen et al., 2012). Finally, vast social media datasets present opportunities for guiding influence if key signals can be quickly extracted. Inferencing faster from online dynamics would enable more responsiveness to emerging opinion trends (Willaert et al., 2020).

As we model the way of exerting influence by building unidirectional connections with agents in the network in our control setting (see Section 2.3), the connections from the external controllers can be viewed as edges that constitute part of the network topology. Given this, we present a comprehensive review in the highly relevant area of network structure inference. More specifically, in Section 2.5.1, we introduce the methodologies used for network structure inference. In Section 2.5.2, we detail the specific domain of reconstructing network structure from epidemiology and information networks as well as the gaps in current literature. In Section 2.5.3, we delve into the reviews of the widely adopted network inference model in opinion dynamics – the inverse Ising model.

2.5.1 Methodologies for Network Structure Inference

Existing literature on inferring network topologies from indirect data has been found in a wide range of domains and solved with domain-specialized knowledge. In the following, we numerate methodologies used for network structure inference categorized by different model representations. Generally, there are two types of models for network structure inference (Brugere et al., 2018): parametric models and non-parametric ones.

In more detail, parametric models are based on assumptions about the distributions of edges related to the input data, e.g., the probabilities of information flow between nodes in the existence of edges. A commonly-used method for solving parametric inference is the maximum likelihood estimation with major applications in epidemiology and information networks (Gomez-Rodriguez et al., 2012). We detail the formation of applying maximum likelihood estimation to network structure inference in Section 2.5.2. Another representative method in network structure inference is Bayesian inference which integrates Markov chain Monte Carlo to obtain estimations of the joint posterior distribution of network structure parameters (Young et al., 2020). To use this method, in addition to the input data, one also needs to provide a high-level description of how the data depends on the underlying network structure, normally represented by probability distributions. Instead of returning an exact inferred network, this method obtains a posterior distribution over possible network structures. Apart from the above two methods, causal models have been broadly used to infer climate networks and brain networks (Ebert-Uphoff and Deng, 2012; Dhamala et al., 2008), which generate causal networks represented as an acyclic digraph for the evolution of the system.

Non-parametric network structure inference refers to methods for learning the topology of networks without making assumptions about the form of the underlying model. One of the most straightforward methods for solving non-parametric network structure inference is called *interaction measures* (Kawale et al., 2013). It starts with measuring the correlations or similarities between two nodes and then chooses a threshold by hand-tuning or with domain-specific knowledge to determine if there exists a link between two nodes, which simply reduces the problem of inferring network topology to finding the optimal threshold. Depending on the statistical measure used to determine the existence of connections between nodes, we further categorize non-parametric models as correlation measures (Bialonski et al., 2011), entropy measures (Donges et al., 2009), frequency domain measures (Lachaux et al., 2002), and regression (Haury et al., 2012). For a more extensive review of methodologies of network structure inference, the readers can refer to Kolaczyk (2009) and Brugere et al. (2018).

2.5.2 Network Inference From Epidemiology and Information Flow

In this section, we focus on the specific domain of reconstructing network structure from the perspective of epidemiology and information networks. These fields align closely with our thesis topic: opinion dynamics inference – a field that has seen limited research. The manner in which diseases or information spreads in these fields offers a compelling parallel to the challenges of tracing opinion propagation in networks. By integrating these viewpoints, we seek to offer a comprehensive review of literature on opinion dynamics inference. As mentioned at the beginning of Section 2.2, the influence propagation usually takes place on a network. Even though it is possible to directly observe when nodes become infected, observing the source of influence, i.e., who infects whom, or the intensity of influence is often very challenging (Leskovec et al., 2009). For instance, considering the information propagation on social media, as the users discover new information, they post it without mentioning the source. Thus, we only observe the time when the user gets "infected", but not where it got infected from. In short, the aim of network structure inference in this field is to discover unobservable networks from data related to information or infection transmission over time.

A majority of works in this field are based on maximum likelihood estimation or explored under the expectation maximization frameworks embedded with a variety of optimization algorithms to obtain a good estimate of the network structure compatible with observations of information cascades. This method works as follows. Based on the assumption that the time of infection or discovering new information is observed, the likelihood of information or infection transmission time difference along an edge is defined by the underlying diffusion model, which further yields the likelihood that a node infects another node as well as the likelihood of information propagation sequences or trees of infection. Finally, by performing optimization based on the likelihood function, we obtain an estimator of the graph (Brugere et al., 2018).

Starting from the seminal work of Gomez-Rodriguez et al. (2012), inferring networks using maximum likelihood methods in this area has been extensively explored in a variety of scenarios. In Gomez-Rodriguez et al. (2012), the authors treat network structure inference as a binary optimization problem (representing whether or not there is an edge between two nodes) and propose the NetInf algorithm based on the maximization of the likelihood of the observed cascades in an independent cascade model. In more detail, they prove the monotonicity and submodularity of the likelihood function relative to edge selection and, therefore, select edges greedily with an accuracy guarantee. Furthermore, to improve the performance of the NetInf algorithm in the independent cascade model, Rodriguez and Schölkopf (2012) propose the MultiTree algorithm by including all directed trees in the optimization. In addition, algorithms have been developed to infer the intensity of connections in a network by Braunstein

et al. (2019) based on the susceptible–infected–recovered model, which is also a progressive model. Other works have incorporated prior knowledge about the network structure (e.g., sparsity (He and Liu, 2017), motif frequency (Tan et al., 2020), or community structure (Ramezani et al., 2017)) to improve the performance of network inference with limited amounts of data.

In order to incorporate uncertainty in inference, other works employ Bayesian inference using Markov chain Monte Carlo methods. Early works in the epidemiology domain (Britton and O'Neill, 2002; Meyers et al., 2005) treat the network model (e.g., ER graph, or scale-free networks) as known, and use Bayesian inference to discover the network model parameters as well as diffusion model parameters (e.g., infection rate). However, the assumption of knowing the network model is too restrictive and, in most cases, inference of structural information is necessary. The most representative work of using Bayesian inference to reconstruct network structure from information cascades is the work by Gray et al. (2020), which has improved estimates of network structure, especially in the presence of noise or missing data, and is also based on the independent cascade model. However, a major limitation arises from their assumption of a binary adjacency matrix for modeling the underlying graph. While simpler, this precludes inferring the intensities of connections between nodes, which limits applicability, as many real networks require modeling weights, not just topology.

Most of the above-mentioned works reconstruct network structures from observations of information cascades or infection trees and are based on progressive models. The exceptions that explore network structure based on non-progressive models are Barbillon et al. (2020), Li et al. (2017), Chen and Lai (2018) and Zhang et al. (2018). In more detail, Barbillon et al. (2020) apply the matrix-tree theorem to infer the network structure based on a SIS model. To maintain the information cascade as a directed acyclic graph as work based on progressive models, the information propagation has been encoded as a matrix with $T \times N$ dimensions where N represents the number of individuals and T is the length of time series. And as a node can flip its state back and forth, it would appear in the same propagation sequence several times. Unlike Barbillon et al. (2020) and all work based on progressive models which needs input sequences of nodes with infection times sorted from a root and monotonically increasing, the works by Li et al. (2017), Chen and Lai (2018) and Zhang et al. (2018) reconstruct network structure from observations of binary-state dynamics. In more detail, Li et al. (2017) translate the network structure inference into a sparse signal reconstruction problem by linearization and solve it via convex optimization. Moreover, Chen and Lai (2018) develop a model combining compressive sensing and clustering algorithms for network reconstruction. However, these above works only consider unidirectional infection (e.g., in the SIS model, if a susceptible node is in contact with an infected node, it will be infected according to a certain probability. Nevertheless, an infected node will not change to the susceptible state due to contact with another susceptible node

but according to the systematic recovery rate.). Instead, Zhang et al. (2018) solve the network inference problem by expectation maximization with a focus on the setting that two states are equivalent. This approach allows to utilize bidirectional dynamics to calculate transition probabilities to reduce the amount of data needed for accurate estimation. However, this work treats an edge as a binary variable (i.e., the existence or absence of a link between two nodes), and is not suitable for inferring the link weight between two nodes.

To sum up, most works in the field of epidemiology and information networks infer network structure from information cascades or infection trees which are identical to directed acyclic graphs, and are not applicable to situations where opinions can be changed back and forth. The work closest to our modelling approach is by Zhang et al. (2018), which uses both directions of transition probabilities to reconstruct networks from binary-state dynamics. However, it treats the edge as a binary variable. On the other hand, reconstructing the strategy of the opponent, which is presented as weighted links from binary-state dynamics is still largely unexplored in the fields of epidemiology and information flow. Most importantly, none of these works investigate the network inference problem from the perspective of manipulating the diffusion process to accelerate the convergence of estimations, which is an important lever if one wants to obtain an estimate with an accuracy guarantee within a short and limited observation time.

2.5.3 Inverse Ising Model

Alongside epidemiology and information flow, network inference has also been extensively studied in the statistical physics community. In the field of statistical physics, network inference is treated as an inverse problem, which 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, the inverse statistical problem has gained increasing interest in a variety of research domains. Recent applications are not limited to social dynamics but also extend to neuroscience (Cocco et al., 2009), computational biology (Lezon et al., 2006), epidemiology (Fajardo and Gardner, 2013), and financial economics (Sornette, 2014). One of the most typical and canonical settings for the inverse statistical problem is the inverse Ising model (Aurell and Ekeberg, 2012). Specifically, for the inverse Ising model, one aims at reconstructing parameters of the Ising model such as coupling strengths between spins from data like spins' states or magnetization.

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 (Nguyen et al., 2017). 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 (Roudi and Hertz, 2011; Mézard and Sakellariou, 2011; Kappen and Spanjers, 2000; Zeng et al., 2011; Dunn and Roudi, 2013; Bachschmid-Romano and Opper, 2014; Battistin et al., 2015; Campajola et al., 2019; Hoang et al., 2019; Lee et al., 2021). For instance, the mean-field method (Roudi and Hertz, 2011; Mézard and Sakellariou, 2011) and the Thouless-Anderson-Palmer approach (Kappen and Spanjers, 2000; Zeng et al., 2011; Dunn and Roudi, 2013) are utilized to provide approximated solutions for the inverse kinetic Ising problem in the weak and dense network connection region. Moreover, belief propagation (Zhang, 2012) and replicas analysis (Bachschmid-Romano and Opper, 2014; Battistin et al., 2015) 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. (2019) 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 (Decelle et al., 2016). 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 rumor (or epidemic) spreading.

To the best of our knowledge, there is only one work that considers the inference problem from the data quality perspective (Decelle et al., 2016). Specifically, the work by Decelle et al. (2016) measures the amount of information contained in a given dataset by applying a method that quantifies the effective rank of the correlation matrix. By removing the less informative rows in the dataset, they obtain a condensed matrix without loss of much information. By doing so, for the same length of the dataset, the data matrix obtained by removing the less informative parts will lead to more accurate inference. However, that work still assumes the data to be given and it does not investigate an active way of generating data for inference. Instead, we are interested in how to strategically influence networked dynamics with the aim of generating a higher-quality dataset for speeding up the convergence of inference in the Ising model.

2.6 Summary

While progressive models like the linear threshold and independent cascade models are widely used in influence maximization, they cannot model real-world opinion dynamics, which involve recurrent changes. Instead, non-progressive such as the voter model and the Ising model better capture empirical observations of evolving opinions through social interactions.

As for the frameworks proposed and techniques used to solve the influence maximization problem, it is a commonly-used method to first verify the submodularity and monotonicity of the evaluation function and then obtain solutions to the IM problem with heuristic algorithms by some accuracy guarantees. However, most works assume to initialize budget allocations from the beginning of the competition and intertemporal budget allocation in voter dynamics has not been considered (see challenge 1 in Section 1.1). Moreover, the aspects of heterogeneity in agent behaviour are rarely explored in influence maximization, especially for non-stationary influence maximization (see challenge 1 in Section 1.1).

We then delve into the problem of network inference, starting with its applications in epidemiology and information flow. Although substantial research has been dedicated to inferring network structures from observations of information cascades and infection trees, the exploration of uncovering weighted network connections from binary-state dynamics remains largely unexplored in the fields of epidemiology and information flow. Moreover, inferring parameters that describe complex systems from existing data is a crucial topic in statistical physics. In this regard, we provide a comprehensive review of the paradigmatic inverse Ising model, which serves as a widely-used approach for modeling opinion dynamics in statistical physics. Following our review, we observe that no existing work has combined aspects of network control with network structure inference where external controllers can interact with internal agents to elicit more information during inference in binary-state dynamics (see challenge 2 in Section 1.1).

Chapter 3

Inter-temporal Influence Maximization

In this chapter, we consider a version of an inter-temporal influence maximization problem. In this problem, one or both of the controllers have the flexibility to determine when to start control. The goal of the controller(s) is to maximize the share of a desired opinion in a group of agents who exchange opinions on a social network subject to voting dynamics.

This chapter is organized as follows: Section 3.1 gives an introduction for studying the inter-temporal influence maximization problem. Section 3.2 introduces detailed formulations of inter-temporal network control in constant-opponent and game-theoretical settings. Then, Section 3.3 outlines the algorithm for determining the optimal budget allocations in the game-theoretical setting. To proceed, Section 3.4.1 analyzes the influence of different starting times on the transient dynamics of vote shares. Following that, Section 3.4.2 discusses the optimal strategies against a known constant opponent, and an unknown opponent. Finally, in Section 3.4.3, we explore the effects of agent heterogeneity on optimized starting times and budget allocations within the inter-temporal influence maximization framework. We conclude in Section 3.5 with an overview of the main findings.

3.1 Introduction

As discussed in Section 2.4, the IM problem is typically explored without time constraints such as deadlines, and is only subject to a *budget* constraint where there are limited resources to allocate to agents in the network (Carnes et al., 2007; Budak et al., 2011; Kermani et al., 2017; Goyal et al., 2019; Romero Moreno et al., 2021b). However, the voting process towards equilibrium can take a long time (Brede et al., 2019a) and

many practical applications of IM have clear deadlines, e.g., seasonal promotions and political campaigns. Indeed, some researchers have incorporated temporal aspects in IM (Liu et al., 2012; Goyal et al., 2013; Alshamsi et al., 2017; Ali et al., 2018; Brede et al., 2019a; Tong et al., 2020). Related to our modeling approach, Brede et al. (2019a) are the first to explore the IM under time constraints in voter dynamics. However, their paper does not allow controllers to allocate different amounts of resources over time and does not consider the game-theoretical setting where competitors have no prior knowledge of the opponent's strategies. In other words, their assumptions do not properly reflect real-world scenarios such as marketing, where the marketers can optimally distribute their budgets inter-temporally and have no information of the competitor's strategy. Representative work considering effects of time scales and activating agents depending on stages of the diffusion process includes the papers by Alshamsi et al. (2017); Ali et al. (2018); Tong et al. (2020). Specifically, Alshamsi et al. (2017) concentrate on minimizing the diffusion time by targeting agents with different levels of connectivity at different stages of the contagion process. In addition to inter-temporal targeting, Tong et al. (2020) allow the controller to adjust its seeding strategy based on what it observes after a certain number of diffusion steps. However, the above two works are addressed in a non-competitive setting where only a single external controller spreads its influence in the network. Given that competition for influence is also common in real-world contexts (e.g., political campaigns (Wilder and Vorobeychik, 2018) or radicalization prevention (Ramos et al., 2015)), the single-controller setting has a restricted range of applications. The only directly related study that solves the time-constrained IM in the game-theoretical setting by considering when to initiate opinion propagation is the work by Ali et al. (2018). However, it focuses on verifying the effectiveness of a Q-learning framework from an algorithmic perspective and does not relate the mechanism behind the optimal strategies to transients of system dynamics. Furthermore, like other models discussed above, the diffusion model of Ali et al. (2018) is the independent cascade model and not appropriate for modeling fast-changing opinions. In addition to the lack of consideration of game-theoretical setting under the framework of time-constraint IM in most works, *agent heterogeneity* — in which agents have different levels of intrinsic bias towards adopting a specific opinion — has not found much attention (Aral and Dhillon, 2018). As agent heterogeneity is often used to model partial radicalism of agents, solving IM considering agent heterogeneity will help us find solutions to reducing the risk of extremism and radicalization. To the best of our knowledge, the only work that explores the IM problem in the presence of agent heterogeneity based on voter dynamics is by Moreno et al. (2020). However, their study focuses on finding the influence-maximizing strategy in the stationary state and does not consider agent heterogeneity in the context of the time-constraint IM.

To bridge these gaps in research about inter-temporal influence allocations, we study the inter-temporal influence maximization problem in voter dynamics under time and budget constraints for the following two settings: First, we consider an active controller

competing against a known and fixed-strategy opponent (referred to as the *constant-opponent setting*); Following this, we then consider a game-theoretical setting in which both controllers compete to maximize their influence (referred to as the *game-theoretical setting*).

Intuitively, in the context of inter-temporal allocations, one or both controllers have to design strategies to make efficient use of their budgets over time. This results in the following trade-off: If the controller starts allocating later, it has more disposable budgets per unit of time but less time left for its influence to become effective. In contrast, if a controller starts allocating earlier, its influence has much time to become effective, but part of it may be wasted, because it is spent before the time when vote shares are evaluated.

Against this background, in this chapter we address the problem of optimal inter-temporal allocations both analytically and numerically. In order to obtain analytical solutions for optimal allocations, we first explore the dynamic allocations in a simplified scenario where controllers only have the flexibility to determine *when* to start control and, once started, they target all agents equally. Then, we conduct numerical experiments on a more complicated scenario where one of the controllers optimizes starting times and budget allocations for individual agents in the network, which is also referred to as *individual optimization*. Given that the individual optimization concentrates on assigning different starting times and allocations on individual nodes, we also take agent heterogeneity into consideration and study its influence on dynamic allocations.

Specifically, in this chapter, we make the following contributions to address research questions 1a, 1b, and 1c on inter-temporal influence maximization outlined in Section 1.1: 1) To address research question 1a on integrating inter-temporal control into the classic IM framework, we propose an advanced model in this chapter. Our model introduces flexibility in budget allocation over time, allowing for strategic distribution rather than a one-time allocation at the campaign's beginning. By adjusting budget distribution over time, this model effectively captures and leverages the transient nature of opinion dynamics, maximizing influence in both constant-opponent and game-theoretical environments. 2) To address research question 1b on relating optimal inter-temporal allocation strategies to transients, we leverage the heterogeneous mean-field method (Brede et al., 2019a) and Taylor expansions to analyze timescales of influence propagation on networks and how they are affected by optimal allocation strategies given time horizons and budgets. Our findings indicate that for shorter time horizons, an optimized controller tends to delay the start of control in more heterogeneous networks compared to less heterogeneous ones. Conversely, for longer time horizons, an early start is more effective in highly heterogeneous networks. 3) For question 1c of how strategies differ within constant-opponent and game-theoretical settings, we observe distinct optimal strategies. In a constant-opponent setting, where one controller starts

influencing from the beginning, the optimal strategy involves initially allowing the opponent's influence and then strategically utilizing the budget towards the campaign's end. Conversely, in a game-theoretical setting, a controller with a larger budget typically initiates control earlier to gain an advantage before their opponent starts. Regarding Nash equilibrium, due to the complexity of analytically computing equilibria in continuous games with infinite strategy sets (Bonomi et al., 2012), we apply iterative gradient ascent (Bonomi et al., 2012). This numerical method, following the minimax theorem (Rosário Grossinho and Tersian, 2001), iteratively maximizes and minimizes the vote shares to find each controller's best response, converging when both have optimized their strategies in response to each other.

3.2 Model Description

Below, we consider a social network given by a graph $G(V, E)$ where a set of N agents is identified with the vertices $v_i \in V$ and edges $e_{ij} \in E$ indicate the existence of social connections between agent i and agent j. In line with most studies in the field, we assume an undirected and positively weighted network without self-loops. The weight of the link from node j to i is denoted by $w_{ji} \geq 0$. Agents in the network can hold one of two opinions: opinion A or opinion B. Apart from the independent agents $i = 1, \cdots, N$, we assume the existence of two external controllers which either favour opinion A or B, referred to as controller A and controller B, respectively. By definition, external controllers never change their opinions and are given as elements external to the network of independent agents. External controllers aim to influence the network, such as to maximize the vote shares of their own opinions. To achieve this, subject to an overall budget constraint, both controllers can build up unidirectional connections with internal agents. In other words, the control gains $a_i(t), b_i(t)$ by controller A and controller B are time-varying unidirectional link weights which indicate the allocation of budgets by A or B to agent i at time t. As we consider the inter-temporal allocation of resources, the control gains $a_i(t), b_i(t)$ are functions of time and they must satisfy the budget constraints: $\sum_{i=1}^{N} \int a_i(t)dt \leq b_A$ and $\sum_{i=1}^{N} \int b_i(t)dt \leq b_B$ where b_A, b_B are the given budgets, i.e. the total amounts of resources available to the controllers. Apart from the budget constraint, $a_i(t), b_i(t)$ also need to be non-negative, i.e. $a_i(t) \geq 0, b_i(t) \geq 0$.

Specifically, the updating process of opinions in the voting dynamics with the inclusion of two opposing controllers is described as follows (Masuda, 2015). First, one of the agents in the network, e.g., agent i, is selected randomly. Then, agent i selects an in-neighbour or a controller at random with a probability proportional to the weight of the incoming link (including control gains from controllers). Correspondingly, the probability of copying a specific opinion is $\frac{w_{ki}}{\sum_{j=1}^{N} w_{ji}+a_i(t)+b_i(t)}$ (to copy the opinion of a network neighbour k), $\frac{a_i(t)}{\sum_{j=1}^{N} w_{ji}+a_i(t)+b_i(t)}$ (to copy the opinion of the A controller), and

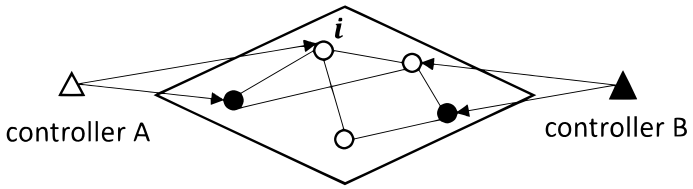

FIGURE 3.1: Schematic diagram of how controllers interact with the opinion dynamics and how agents update their opinions. Triangles stand for controllers and agents are represented by circles. Black and blank symbols indicate that the agents (or controllers) are holding opinions A or B, respectively. The lines between agents correspond to the social connections. External controllers A and B influence opinion dynamics by building unidirectional links to agents in the networks. Assuming unity link weights from the neighbours and controllers, in the next time step, the randomly picked agent i will have probability $3/4$ to stay in opinion A and probability $1/4$ to flip its opinion.

$\frac{b_i(t)}{\sum_{j=1}^{N} w_{ji}+a_i(t)+b_i(t)}$ (to copy the opinion of the B controller). For a better understanding of the updating process, an illustration is given in Fig. 3.1. Suppose agent i has been picked and assume unit-strength connections between agents and from the controllers. Agent i in Fig. 3.1 is linked with three other agents (one of which holds opinion B and two who hold opinion A), and is targeted by controller A. Therefore, in the next update, agent i will have probability $3/4$ to stay in opinion A and probability $1/4$ to flip its opinion to B.

Here, we follow the mean-field rate equation for probability flows (Masuda, 2015) by introducing x_i as the probability that agent i has opinion A. We then have:

$$\frac{dx_i}{dt} = (1-x_i)\frac{\sum_{j=1}^{N} w_{ji}x_j + a_i(t)}{\sum_{j=1}^{N} w_{ji} + a_i(t) + b_i(t)} - x_i\frac{\sum_{j=1}^{N}(1-x_j)w_{ji} + b_i(t)}{\sum_{j=1}^{N} w_{ji} + a_i(t) + b_i(t)}. \tag{3.1}$$

In Eq. (3.1) the first term corresponds to the probability of holding opinion B but being converted to opinion A, while the second term presents the probability of having opinion A but picking opinion B. Eq. (3.1) can be further rewritten in a matrix form as

$$\frac{d\vec{x}}{dt} = \left[\left(\vec{\Delta}\mathbb{1}^T\right) \odot \boldsymbol{W} + diag(-1)\right]\vec{x} + \vec{\Delta} \odot \vec{P}_A, \tag{3.2}$$

where $\vec{\Delta} = [\frac{1}{\sum_{j=1}^{N} w_{j1}+a_1(t)+b_1(t)}, ..., \frac{1}{\sum_{j=1}^{N} w_{jN}+a_N(t)+b_N(t)}]^T_{1\times N}$, $\vec{P}_A = [a_1(t), ..., a_N(t)]^T_{1\times N}$, $\vec{x} = [x_1(t), ..., x_N(t)]^T_{1\times N}$, $\mathbb{1}$ is an all-ones vector, $diag(-1)$ denotes an N-dimensional diagonal matrix with the elements -1 on the main diagonal, $\odot$ is the Hadamard product (Horn, 1990) which denotes pairwise multiplication and $\boldsymbol{W}$ is the weighted adjacency matrix of the social network. Specifically, Eq. (3.2) represents the opinion dynamics as a linear inhomogeneous non-autonomous system of first-order differential equations. Notably, Brede et al. (2019a) solve Eq. (3.2) in the time-invariant case by eigenvalue decomposition and find that the eigenvalues of the matrix $\left[\left(\vec{\Delta}\ \mathbb{1}^T\right) \odot \boldsymbol{W} + diag(-1)\right]$ determine the time scales of opinion changes. In more detail, eigenvalues for large-degree

nodes are smaller which indicates it would take longer for the high-degree nodes to reach the equilibrium.

Here, instead, we focus on the time-variant scenario of Eq. (3.1) regarding different starting times of controller A. To obtain an analytical estimate of vote-share trajectories of agents, we use the heterogeneous mean-field approximation (Dezső and Barabási, 2002) in Section 3.4.1. Additionally, for numerical results, we employ the Runge-Kutta method (Press et al., 1988) for numerical integration.

Consequently, the objective function for controller A is to maximize the average vote share obtained by solving Eq. (3.1) at the end of a given time T:

$$S_A(T) = \frac{\sum_{i=1}^{N} x_i(T)}{N}. \tag{3.3}$$

Correspondingly, since an agent in the network can either have opinion A or opinion B, for controller B, the objective function is $S_B(T) = 1 - S_A(T)$.

In the following, we present detailed formulations of inter-temporal network control in the context of constant-opponent and game-theoretical settings. Central to our discussion is the non-autonomous system, as represented by Eq. (3.1). In essence, a non-autonomous system, as defined in the field of dynamical systems, is one where its behavior is influenced by time-dependent factors. This means that unlike autonomous systems, whose behavior is governed by fixed rules, the governing equations of non-autonomous systems include variables or parameters that change over time. This temporal evolution adds layers of complexity, making it challenging to handle fully flexible influence allocations $a_i(t)$ and $b_i(t)$. Thus, we first consider a simplified model for inter-temporal influence allocation where the controller has the flexibility to determine the time of the start of the campaign in Sections 3.2.1 and 3.2.2. Furthermore, given the difference in transient dynamics of individual agents, we then refine the above assumption by assigning different allocations and starting times for individual agents in Section 3.4.2.2.

3.2.1 Inter-temporal Budget Allocations against a Constant Opponent

In the constant-opponent scenario, we consider maximizing the vote shares of opinion A at time T (see Eq. (3.3)) against a constant controller B who starts control from the beginning of the competition. To allow for analytical tractability, we start with a simplified campaign scheme in which all agents start at the same time with the same

allocation. Formally, the control gains $a_i(t), b_i(t)$ $(1 \leq i \leq N)$ by controller A and controller B under the above assumption are:

$$\begin{aligned} a_i(t) &= \begin{cases} 0 & 0 \leq t \leq t_a \\ \frac{b_A}{(T-t_a)N} & t_a < t \leq T, \end{cases} \\ b_i(t) &= \frac{b_B}{TN} \qquad\qquad 0 \leq t \leq T \end{aligned} \tag{3.4}$$

where t_a is the starting time of all nodes assigned by controller A. Consequently, this IM problem is equivalent to determining a single optimal t_a that maximizes $S_A(T)$, i.e.,

$$t_a^* = \arg\max_{t_a} S_A. \tag{3.5}$$

Note that, the optimized t_a^* should meet the time constraint, i.e., $0 \leq t_a^* \leq T$. Even though the simplicity of all nodes having the same starting time allows for analytical solutions for the optimal campaigning strategy, we are also interested in obtaining numerical results for a more general scenario in which the controller can assign different starting times and budget allocations for individual nodes by combining optimization schemes with numerical integration of Eq. (3.1). Specifically, let us consider a case in which controller A can choose specific campaigning starts $t_{a,i}$ for individual node i and has the option to split its budget unevenly between nodes by assigning individual nodes budgets $b_{A,i}$. To reduce the number of parameters, we still retain the assumption that once started campaigns proceed with constant intensity, i.e.,

$$a_i(t) = \begin{cases} 0 & 0 \leq t \leq t_{a,i} \\ \frac{b_{A,i}}{(T-t_{a,i})N} & t_{a,i} < t \leq T, \end{cases} \tag{3.6}$$

where $b_{A,i}$ and $b_{B,i}$ should satisfy $\sum_{i=1}^{N} b_{A,i} \leq b_A$ and $\sum_{i=1}^{N} b_{B,i} \leq b_B$ respectively. To obtain the optimal configurations for controller A in the constant opponent case, we use the interior-point algorithm (Press et al., 2007) for numerical optimization of Eq. (3.5) and stochastic hill climbing (Juels and Wattenberg, 1995) for individual optimization whose strategy space is defined by Eq. (3.6).

3.2.2 Inter-temporal Budget Allocations in the Game-theoretical Setting

In the game-theoretical scenario, we consider a zero-sum game of competitive vote-share maximization on graph G. Players of the game are controller A and controller B who have complete knowledge of graph G. Besides, these two players have to simultaneously decide on an inter-temporal allocation protocol at time zero. In the simplified campaign scheme, both controllers have the flexibility to choose the starting time of allocations, defined as t_a and t_b. More precisely, the sets of actions available to controller A and controller B are $\phi_A = \{t_a \mid 0 \leq t_a \leq T\}$ and $\phi_B = \{t_b \mid 0 \leq t_b \leq T\}$ respectively.

Hence, allocations of budgets on agent i $(1 \leq i \leq N)$ per unit time are:

$$a_i(t) = \begin{cases} 0 & 0 \leq t \leq t_a \\ \frac{b_A}{(T-t_a)N} & t_a < t \leq T, \end{cases}$$
$$b_i(t) = \begin{cases} 0 & 0 \leq t \leq t_b \\ \frac{b_B}{(T-t_b)N} & t_b < t \leq T. \end{cases} \tag{3.7}$$

Moreover, the payoff function for controller A and B are $u_A = S_A(T)$ and $u_B = 1 - S_A(T)$ where the vote shares $S_A(T)$ are continuous and satisfy $0 \leq S_A \leq 1$.

While pure-strategy Nash equilibria are theoretically guaranteed to exist in two-player zero-sum games under continuity, convexity and boundedness assumptions (Lu, 2007), explicitly calculating these equilibria through analytical means remains challenging for games with infinite, continuous strategy spaces (Guo et al., 2021). This holds in our problem setting, where the unknown strategies used by the opponent yield continuous action spaces that are mathematically intractable for closed-form equilibrium analysis. Additionally, the absence of closed-form expressions for the payoff functions of either competing controller further hinders direct analytical tractability for deriving equilibria. Therefore, numerically approximating the Nash equilibria stands as the most viable alternative. Here, we calculate the Nash equilibria via iterative gradient-based techniques guided by the minimax theorem (Ferreira et al., 2012), which is the most commonly used approach to numerically estimate Nash equilibria strategies when explicit derivation is hindered (Ansari et al., 2019). Details about this algorithm are illustrated in section 3.3.

3.3 Algorithm for Computing Nash Equilibria

In this section, we present the iterative searching algorithm for computing Nash equilibria in the game-theoretical setting in algorithm 1. According to the minimax theorem (Ferreira et al., 2012), in a two-player zero-sum game, each player can achieve their optimal strategy by maximizing their objective function while taking into account the opponent's attempt to minimize their objective function. Therefore, we compute Nash equilibria by maximizing and minimizing $S_A(T)$ iteratively until convergence is achieved (Bonomi et al., 2012). To reduce oscillations during iterations, we introduce an adaptive learning rate μ $(0 \leq \mu \leq 1)$ to be inversely proportional to the number of iterations in the search process (Dixon, 1972).

Specifically, we initialize the algorithm with some initial starting times of controller A and B, $t_a^{(0)}$ and $t_b^{(0)}$ (see line 1). Next, we perform a loop (see lines 2-8) and the loop will end on the condition that if the sum of the absolute changes in the starting times is less than a small enough preassigned threshold $\theta = 10^{-9}$. We then assume that the iteration

input : b_A, b_B, N, T, adjacency matrix $\boldsymbol{W}$, threshold θ
output: approximations for t_a^{NE} and t_b^{NE} at the Nash equilibrium; vote shares at time T, $S_A^{NE}(T)$

1 *Initialization:* $t_a^{(0)} = 0$; $t_b^{(0)} = 0$; $\Delta t_a = 1$; $\Delta t_b = 1$; *i=0*;
2 **while** $|\Delta t_a| + |\Delta t_b| \geq \theta$ **do**
3 $\quad t_a^{(i+1)} = t_a^{(i)} + \mu \left(\arg\max_{t_a} \left\{ S_A(T) \middle| t_b = t_b^{(i)} \right\} - t_a^{(i)} \right)$;
4 $\quad t_b^{(i+1)} = t_b^{(i)} + \mu \left(\arg\min_{t_b} \left\{ S_A(T) \middle| t_a = t_a^{(i+1)} \right\} - t_b^{(i)} \right)$;
5 $\quad \Delta t_a = t_a^{(i+1)} - t_a^{(i)}$;
6 $\quad \Delta t_b = t_b^{(i+1)} - t_b^{(i)}$;
7 $\quad$ i=i+1;
8 **end**
9 $t_a^{NE} = t_a^{(i-1)}$;
10 $t_b^{NE} = t_b^{(i-1)}$;
11 $S_A^{NE}(T) = \{ S_A(T) \mid t_a = t_a^{NE}, t_b = t_b^{NE} \}$

Algorithm 1: Iterative searching for Nash equilibrium

has converged. For the main body of the loop (lines 3-7), we carry out the optimization of t_a and t_b in turn. Note that, as $\mu = 1/i$ is in the interval $[0, 1]$, the updated $t_a^{(i+1)}$ and $t_b^{(i+1)}$ are always in the constraint plane $[0, T]$. The algorithm ends if both controller A and controller B have nothing to gain by unilaterally changing their strategies, which is consistent with the definition of the Nash equilibrium. By doing so, we determine the starting times of controller A and B in the Nash equilibrium named t_a^{NE} and t_b^{NE}, and also the vote share in the Nash equilibrium $S_A^{NE}(T)$.

3.4 Results

The focus of our study is on influence-maximizing strategies for one or both controllers on heterogeneous networks. Accordingly, in Section 3.4.1, we start our analysis with a heterogeneous mean-field approximation of Eq. (3.1) to explore the influence of different starting times on the transient dynamics of vote shares. Next, to further investigate the dependence of optimal starting times on network heterogeneity, we proceed with an analytical exploration of the timescales for the system to approach its equilibrium. To support and supplement our theoretical analysis, we also carry out detailed numerical experiments based on uncorrelated random scale-free networks with power-law degree distribution $p(k) \sim k^{-\lambda}$ constructed by the configuration model (Catanzaro et al., 2005) (see Section 2.1 for more details). After gaining some insights about the dependence of inter-temporal influence allocation on network heterogeneity, we conduct optimization in Section 3.4.2 to obtain the optimal strategies for the inter-temporal

network control in both constant-opponent and game-theoretical settings, as well as presenting the resulting improvements if we optimize each agent individually.

3.4.1 Heterogeneous Mean-field Analysis

Due to the large number of degrees of freedom of the system described in Eq. (3.1) which is in the same order of magnitude as the network size and the time-variant term given by the control gains, it is hard to obtain fully analytical solutions for voter dynamics. Here, to obtain an analytical estimate of vote-share trajectories of nodes, we take the commonly used heuristics called heterogeneous mean-field theory to investigate the complex dynamical processes of Eq. (3.1) in the context of competing against a constant opponent. Generally, this theory reduces the degrees of freedom of the system based on the assumption that the dynamics of agents of the same degree are statistically equivalent. This assumption works well for degree uncorrelated networks where there is no assortative or dis-assortative mixing by degree (Huang and Chen, 2019). Therefore, we group nodes of the same degree k and assume they follow roughly similar dynamics $x_k(t)$. We then approximate

$$\sum_j w_{ji} x_j \approx k_i \langle x \rangle , \tag{3.8}$$

where $k_i = \sum_j w_{ji}$ is the sum of incoming links of node i and $\langle x \rangle$ represents the average behaviour of a neighbour. Specifically, $\langle x \rangle$ is equal to

$$\langle x \rangle = \sum_k \frac{k}{\langle k \rangle} p_k x_k, \tag{3.9}$$

where p_k is the fraction of nodes with degree k and $\langle k \rangle = \sum_k k p_k$ is the average degree of the network. Inserting Eq. (3.8) into Eq. (3.1) and rewriting for the dynamics of nodes with degree k, we have

$$\dot{x}_k = \frac{a_k}{k + a_k + b_k} + \frac{k \langle x \rangle}{k + a_k + b_k} - x_k, \tag{3.10}$$

where a_k and b_k are influence allocations to nodes of degree k. Multiplying Eq. (3.10) by $k p_k / \langle k \rangle$ and summing over k, we obtain a differential equation for $\langle x \rangle$

$$\frac{d}{dt} \langle x \rangle = \beta + \alpha \langle x \rangle \tag{3.11}$$

with coefficients

$$\begin{aligned} \alpha &= \sum_k \frac{k^2 p_k}{\langle k \rangle} \frac{1}{k + a_k + b_k} - 1, \\ \beta &= \sum_k \frac{k p_k a_k}{\langle k \rangle} \frac{1}{k + a_k + b_k}, \end{aligned} \tag{3.12}$$

which are constants for a given network. As Eq. (3.11) is a first-order differential equation, it can be solved by eigenvalue decomposition. Based on the assumption that controller B starts control at time 0 and controller A has the freedom to choose the starting time t_a, we obtain the probability that nodes of degree k have opinion A at time t ($t > t_a$) and the corresponding vote share as follows:

$$\begin{aligned} x_k(t) &= \frac{a_k\alpha - \beta k + \frac{ke^{\alpha t}(\beta+\alpha x_k(t_a))}{\alpha+1}}{\alpha(a_k+b_k+k)} \quad - e^{-t}\left(\frac{a_k\alpha - \beta k + \frac{k(\beta+\alpha x_k(t_a))}{\alpha+1}}{\alpha(a_k+b_k+k)} - x_k(t_a)\right), \\ S_A(t) &= \sum_k p_k x_k(t), \end{aligned} \tag{3.13}$$

where

$$\begin{aligned} \gamma &= \sum_k \frac{k^2 p_k}{\langle k\rangle}\frac{1}{k+b_k} - 1, \\ x_k(t_a) &= x_0 e^{-t_a} + \frac{k}{k+b_k}x_0(e^{\gamma t_a}(1-e^{-t_a})). \end{aligned} \tag{3.14}$$

Here, $x_k(t_a)$ is the state of nodes of degree k at time t_a. From Figs. 3.4 (b) and (e) in subsection 3.4.2.1 which show the dependence of vote shares on the starting times, it can be seen that the mean-field method is a good approximation for the total vote shares $S_A(T)$.

According to Brede et al. (2019a), the equilibration dynamics of a node depend on its degree, which in turn, influences the optimal strategy in transient control. For example, hub nodes would typically have slow equilibration dynamics, which results in the poor vote shares of hub control for short time horizons. Inspired by that, the network's natural timescales towards equilibrium will also have an impact on determining the optimal inter-temporal budget allocations. Therefore, in the following, we investigate the dependence of equilibration dynamics in networks with different degrees of heterogeneity by systematically analyzing relaxation times as defined in Son et al. (2008) based on mean-field results of Eq. (3.13). In more detail, the relaxation time calculates the time required for the system to reach its equilibrium or steady state. In other words, it measures the speed at which a system stabilizes. To quantify the relaxation times for nodes with degree k, we define a normalized order parameter as:

$$r_k(t) = \frac{x_k(t) - x_k(\infty)}{x(0) - x_k(\infty)}. \tag{3.15}$$

Here, we always use the setting that all nodes have the same initial state $x(0) = x_0$. Then we measure the average relaxation times for nodes with degree k via:

$$\begin{aligned} \tau_{relax,k} &= \int_0^\infty r_k(t)dt = \int_0^\infty \frac{x_k(t) - x_k(\infty)}{x(0) - x_k(\infty)}dt \\ &= \frac{\alpha(\alpha x_0(a_k+b_k+k) - a_k\alpha + \beta k - kx_0) - \beta k}{\alpha(\alpha x_0(a_k+b_k+k) - a_k\alpha + \beta k)}. \end{aligned} \tag{3.16}$$

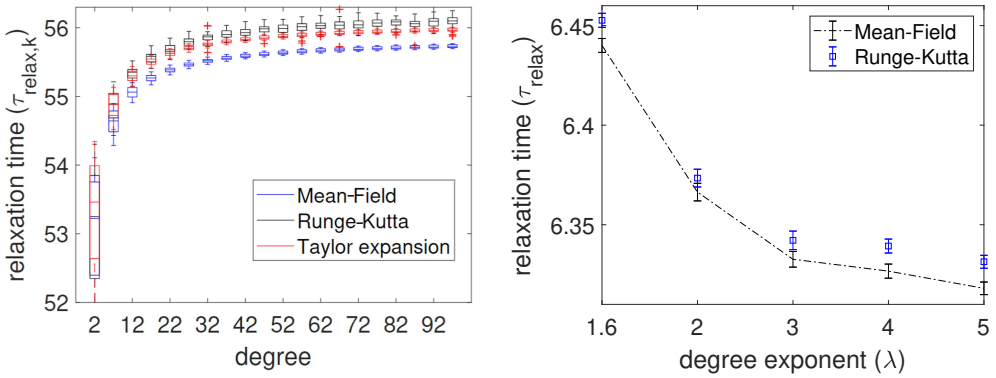

(a) (b)

FIGURE 3.2: Panel (a) shows dependence of relaxation time $\tau_{relax,k}$ on degree k calculated via direct integration using a Runge-Kutta method, the mean-field estimation of Eq. (3.16) and a Taylor expansion of Eq. (3.19) in k up to 2nd order on networks with exponent $\lambda = 1.6$. The data is represented in box plots with median, $25th$ and $75th$ percentiles and whiskers extending to the maximum or minimum values. Panel (b) shows the dependence of relaxation time τ_{relax} on network heterogeneity calculated numerically via integration and analytically by mean-field approximation. Error bars indicate 95% confidence intervals. Both results are based on networks with $N = 10000$, average degree $\langle k \rangle = 10.5$, averaged over 100 realizations. Both controller A and controller B start control at time 0 and allocate 0.1 or 1 resource on each node per unit time respectively for panels (a) and (b).

Inserting Eq. (3.12)'s α and β into Eq. (3.16) gives an involved expression. To still gain insight into the dependence of equilibration times on degree, we approximate Eq. (3.16) in the limit of $\frac{a_k+b_k}{k} < 1$ up to the second order and obtain

$$\begin{aligned}\tau_{relax,k} \simeq & \frac{\alpha-1}{\alpha} + \frac{a_k(x_0-1)+b_k x_0}{\beta+\alpha x_0} k^{-1} \\ & - \alpha \left(\frac{a_k(x_0-1)+b_k x_0}{\beta+\alpha x_0} \right)^2 k^{-2} + O\left(\frac{1}{k}\right)^3 .\end{aligned} \tag{3.17}$$

In particular, if controller A and controller B target all nodes equally, i.e. $a_k = a$ and $b_k = b$, we have:

$$\beta + \alpha = -\frac{b}{a}\beta. \tag{3.18}$$

As β is strictly positive and the sum of α and β is negative, we obtain that α is negative as well. Therefore:

$$\frac{a_k(x_0-1)+b_k x_0}{\beta+\alpha x_0} = \frac{x_0(a+b)-a}{\frac{x_0(a+b)-a}{a+b}\alpha} = \frac{a+b}{\alpha} < 0.$$

Eq. (3.17) can now be simplified to:

$$\tau_{relax,k} \simeq \frac{\alpha-1}{\alpha} + \frac{a+b}{\alpha} k^{-1} - \frac{(a+b)^2}{\alpha} k^{-2} + O\left(\frac{1}{k}\right)^3, \tag{3.19}$$

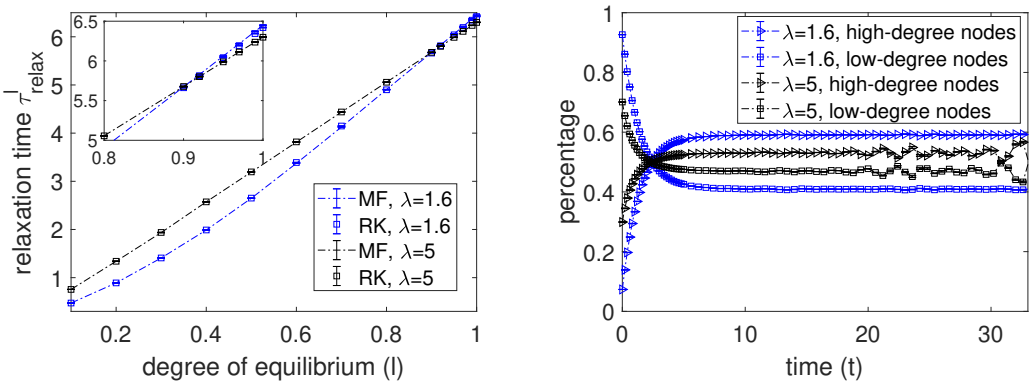

(a) (b)

FIGURE 3.3: Panel (a) shows the dependence of relaxation time on the degree of equilibrium l and network heterogeneity by Runge-Kutta and mean-field methods. Panel (b) shows the evolution of average vote share changes in the proportion. "low-degree nodes" and "high-degree nodes" refer to the first 80% low-degree nodes and top 20% high-degree nodes. The y-axis shows the proportion of the average state changes for high-degree nodes and low-degree nodes in the total changes. Both results are based on networks with $N = 10000$, average degree $\langle k \rangle = 10.5$, represented by error bars with 95% confidence intervals over 100 realizations. Both controller A and controller B start control at time 0 and allocate 1 resource on each node per unit time. The legends "$\lambda = 1.6$" and "$\lambda = 5$" are identical to power law distribution $p(k) \sim k^{-1.6}$ and $p(k) \sim k^{-5}$.

which indicates that the approximation of average relaxation times for nodes with degree k is independent of the initial state if both controllers target all nodes uniformly. Additionally, we find that the trend of $\tau_{relax,k}$ with degree k is mainly determined by the constant term $\frac{\alpha-1}{\alpha}$ and the first-order term $\frac{a+b}{\alpha}k^{-1}$. Specifically, the coefficient $\frac{a+b}{\alpha}$ is negative, which leads to the phenomenon that the larger the degree of a node is, the longer the relaxation time will be. Moreover, the second-order term $\frac{(a+b)^2}{\alpha}k^{-2}$ reduces the difference between relaxation time for nodes of different degrees.

To proceed, we compare $\tau_{relax,k}$ calculated via direct integration using a Runge-Kutta method, the mean-field estimation of Eq. (3.16) and a Taylor expansion of Eq. (3.17) in Fig. 3.2(a). From Fig. 3.2(a), it can be seen that x_k is monotonically increasing with degree k. This phenomenon is also consistent with the Gershgorin circle theorem (Chien et al., 2016). According to the Gershgorin theorem, eigenvalues for nodes with degree k of Eq. (3.2) lie within at least the discs with radii -1 + $\frac{1}{1+\frac{a_k+b_k}{k}}$ around zero. As we assume that the controllers target all nodes uniformly, the larger the degree of nodes, the smaller the absolute values of eigenvalues. In other words, the larger the node's degree, the longer its relaxation time scales towards equilibrium. Additionally, we find that the mean-field method and Taylor expansion are in reasonable agreement with numerical estimates for $\tau_{relax,k}$.

Furthermore, the overall average relaxation time (i.e., network's natural timescales towards equilibrium) for the equally targeting case is given by:

$$\begin{aligned}\tau_{relax} &= \sum_k p_k \tau_{relax,k} = \sum_k p_k \frac{k(\beta(a+b)+a)+\beta(a+b)^2}{\beta(a+b)(a+b+k)} \\ &\simeq \sum_k p_k \left(\frac{\langle k \rangle \left(k(a+b) - (a+b)^2 - k^2 \right)}{k^2(a+b)(\sum_k \frac{p_k}{k}(a+b)-1)} + 1 \right) + O\left(\frac{1}{k}\right)^3.\end{aligned} \tag{3.20}$$

As we investigate networks with different degrees of heterogeneity characterized by different degree exponents λ, Eq. (3.20) can be written as:

$$\begin{aligned}\tau_{relax} \simeq &- \frac{\langle k \rangle \, \lambda_1 (a+b) (\langle k \rangle \, (\lambda_1 - 2))^{\frac{1}{\lambda_1 - 2}}}{(\lambda_1 + 1) \left(a - \lambda_1 \left((\langle k \rangle \, (\lambda_1 - 2))^{\frac{1}{2-\lambda_1}} \right)^{\lambda_1} + b \right)} \\ &+ \frac{\langle k \rangle}{a+b} + 1 + O\left(\frac{1}{k}\right)^3.\end{aligned} \tag{3.21}$$

Our next aim is to investigate the dependence of overall relaxation times on network heterogeneity characterized by the degree exponent λ. For this purpose, we numerically calculate the average relaxation time τ_{relax} for different settings of λ. Fig. 3.2 (b) shows simulation results for τ_{relax} obtained numerically via Runge-Kutta integration and compares to mean-field results based on Eq. (3.20). The figure illustrates that the more heterogeneous the network, the larger its timescales towards equilibrium. Combined with the results in Fig. 3.2(a), we gain an intuition that the long timescales in highly heterogeneous networks are mainly caused by the higher degree nodes.

In the above, τ_{relax} only represents timescales towards equilibrium. However, we are also interested in time scales towards reaching non-equilibrium states. Therefore, we extend the notation of τ_{relax} by introducing the *degree of equilibrium* l. Here, l describes to which extent the termination state (depicted by vote share at time T, i.e., $S_A(T)$) approximates the equilibrium state (depicted by vote share at time ∞, i.e., $S_A(\infty)$). In more detail, $l = \frac{|S_A(T)-x(0)|}{|S_A(\infty)-x(0)|}$. Then, we define the average timescales towards $lS_A(\infty)$ as *l-percentage relaxation time*, given by:

$$\tau^l_{relax} = \int_0^{t'} r(t)dt = \int_0^{t'} \sum_k p_k r_k(t) dt, \tag{3.22}$$

where t' is determined implicitly by $S_A(t') = lS_A(\infty) + (1-l)x(0)$. This equation defines an average timescale at which the vote-share dynamics approach the desired l-percentage vote shares.

To explore the relationship between the l-percentage relaxation time and transient control, we plot the dependence of relaxation time on the degree of equilibrium l and

network heterogeneity in Fig. 3.3 (a). We clearly see a cross-over of τ^{l}_{relax} in Fig. 3.3 (a): relaxation times are larger for less heterogeneous networks than for more heterogeneous networks for low l, but this ordering is reversed for large degree of equilibrium (see inset in Fig. 3.3 (a)). We hypothesize that this is a consequence of the characteristic dynamics toward equilibrium in heterogeneous networks occurring via two stages. To illustrate this point, we visualize the evolution of vote shares for high-degree and low-degree nodes in Fig. 3.3 (b). In more detail, we sort nodes according to their degrees in ascending order. Then we assign the first 80% as low-degree nodes and the rest as high-degree nodes according to the Pareto principle (Price, 1965). To explore which role they play in the transient dynamics, we compute the state changes $\frac{dx_i}{dt}$ grouped by low-degree nodes (i.e. $\sum_{low} \frac{dx_i}{dt}$) and high-degree nodes (i.e. $\sum_{high} \frac{dx_i}{dt}$). Then the average contribution of low-degree nodes and high-degree nodes to the vote-share changes are:

$$\frac{0.2\sum_{low} \frac{dx_i}{dt}}{0.8\sum_{high} \frac{dx_i}{dt} + 0.2\sum_{low} \frac{dx_i}{dt}}, \text{and } \frac{0.8\sum_{high} \frac{dx_i}{dt}}{0.8\sum_{high} \frac{dx_i}{dt} + 0.2\sum_{low} \frac{dx_i}{dt}}. \tag{3.23}$$

In this way, we obtain Fig. 3.3 (b), where we also compare vote-share changes for networks constructed for different degree exponents. In Fig. 3.3 (b), we see that a large proportion of vote-share changes is caused by the low-degree nodes at the beginning of the evolution. As the evolution proceeds, the dynamics of high-degree nodes are increasingly becoming the leading cause of vote-share changes. Moreover, the degree of heterogeneity λ of the network will also make a difference in vote-share changes. For example, in the beginning, the state changes by low-degree nodes in highly heterogeneous networks make up a more significant proportion of total vote-share changes than that by low-degree nodes in less heterogeneous networks. We thus see that the state changes by high-degree nodes in highly heterogeneous networks account for a larger proportion in total vote-share changes than those by high-degree nodes in less heterogeneous networks.

Combining the results in Fig. 3.3 (a) and Fig. 3.3 (b), we obtain the following picture. For small l, as the state changes of vote shares are mainly driven by low-degree nodes (see the left front part of Fig. 3.3 (b)), the evolution of vote shares is dominated by the low-degree nodes. Since highly heterogeneous networks have many more low-degree nodes, they can approach the desired states faster. In contrast, for large l, the state changes of high-degree nodes play a crucial role in the total vote-share changes. Networks with high heterogeneity have many more high-degree nodes which thus delay the approach to equilibrium.

3.4.2 Analysis of Optimal Strategies

Above, we have applied the heterogeneous mean-field theory to investigate the voting dynamics and measure the equilibration timescales based on the mean-field approximation. After gaining some insights about the dynamics of networks with different heterogeneity, in this subsection, we carry out numerical experiments to find the optimal strategies for inter-temporal budget allocation in both constant-opponent and game-theoretical settings. Then we look back and explain the configurations of optimal inter-temporal allocation schemes by the results in Section 3.4.1. To make the explanation more explicit, in each setting, we start our analysis with the evolution of vote shares pertaining to networks with different degrees of heterogeneity. Based on that, we explore the effects of time horizons and relative budgets on optimal starting times. Additionally, to test the effectiveness of the simple scenario of assigning the same starting time for all nodes, we also compare the vote shares obtained by the simple scenario with the individual optimization setting.

3.4.2.1 Optimal Strategies in the Constant-opponent Setting

We start with finding the optimal strategies in the simple scenario where controller A only has the flexibility to determine when to start control. Specifically, we obtain the optimal starting time of controller A by interior-point optimization when controller B starts its control from time $t = 0$ subject to equal budgets.

Firstly, in order to have an intuition of how the optimal starting time affects the vote shares regarding network heterogeneity, we depict the evolution of vote shares for networks with degree exponents $\lambda = 1.6$ and $\lambda = 5$ within time horizons $T = 16$ and $T = 256$ in Figs. 3.4 (a) and (b) respectively. The choice of time horizons $T = 16$ and $T = 256$ allows us to investigate two scenarios whether there is enough time for the network to approach the equilibrium state or not, which corresponds to what we discuss in Section 3.4.1 about the degree of equilibrium. Further, the turning points where vote shares change dramatically in Figs. 3.4 (a) and (b) are the optimal times for controller A to start control.

By comparing the vote-share trajectories in Figs. 3.4 (a) and (b), we obtain the following observations. First, it is shown in Fig. 3.4 (a) that the optimized controller for networks with degree exponent $\lambda = 1.6$ will start later than for networks with degree exponent $\lambda = 5$. Additionally, at the very beginning of the competition, vote shares for networks with degree exponent $\lambda = 1.6$ decline faster, which roughly indicates that highly heterogeneous networks respond more quickly to the resource injection in the early stage. Second, in Fig. 3.4 (b), compared with networks with degree exponent $\lambda = 5$, the optimized controller for networks with degree exponent $\lambda = 1.6$ has to start earlier. Moreover, vote shares for networks with degree exponent $\lambda = 1.6$ have

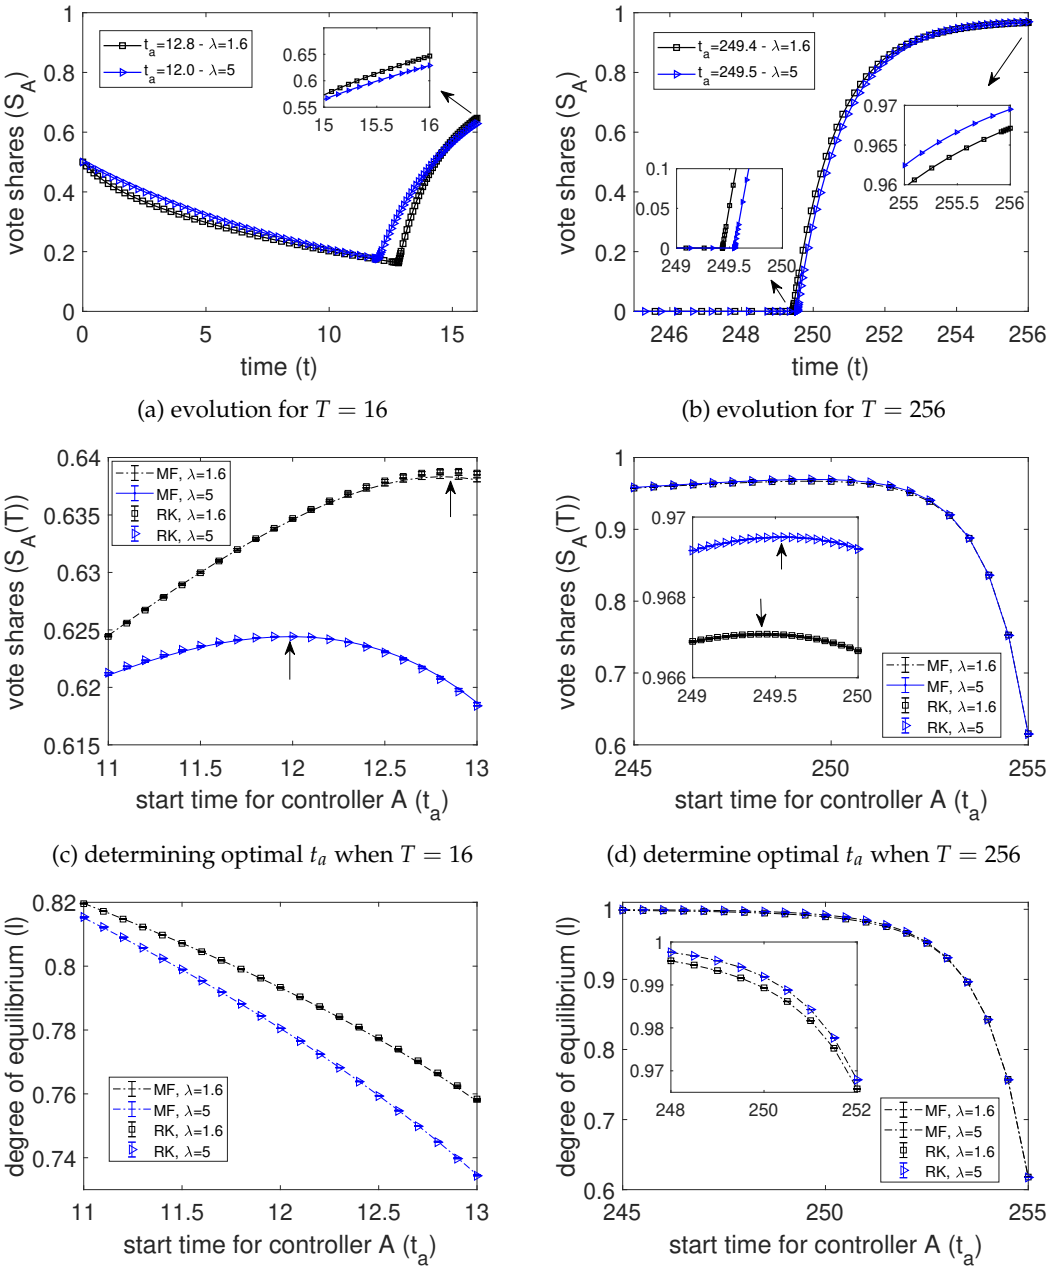

(a) evolution for $T = 16$

(b) evolution for $T = 256$

(c) determining optimal t_a when $T = 16$

(d) determine optimal t_a when $T = 256$

(e) dependence of l on t_a when T=16

(f) dependence of l on t_a when T=256

FIGURE 3.4: Panels (a) and (b) show the evolution of total vote shares when controller A takes optimal control for time horizons $T = 16$ and $T = 256$ respectively. The turning points are the time when controller A takes the optimal control. Panels (c) and (d) show the dependence of vote shares on controller A's starting time for time horizon $T = 16$ and $T = 256$ respectively. Panels (e) and (f) show the dependence of the degree of equilibrium l on starting time. To find the optimal control time, the networks have to strike a balance between budgets per node and degree of equilibrium. All the calculations are based on networks with $N = 10000$ and $\langle k \rangle = 10.5$ and tested in 100 realizations. Controller B always starts its control from time 0. The black squares and blue triangles stand for networks with degree of heterogeneity $\lambda = 1.6$ and $\lambda = 5$ respectively.

been exceeded at the end of the competition by networks with degree exponent $\lambda = 5$. Given that, for both networks, the vote shares when controller A starts control are the same (i.e., $S_A(t_a) = 0$), the changes in vote shares can only be a result of budgets and the network response speed.

Additionally, to prove the accuracy of interior-point optimization, we present data for the dependence of vote shares on the starting times of controller A in Figs. 3.4 (c) and (d). We note that the dependence is a convex shape with a maximum, which is marked with arrows in Figs. 3.4 (c) and (d). The peak values of curves are consistent with the optimal starting time in Figs. 3.4 (a) and (b) (see the turning points in Figs. 3.4 (a) and (b)).

In more detail, the maximum values of the vote shares are a result of a trade-off. On the one hand, if t_a is small, the controller will have more time left to influence the network but with small resource allocations on each node per unit time. In other words, the final vote shares are determined by $lS_A(\infty)$. Though an early start shifts the system closer to the equilibrium (i.e., l becomes larger), small resource allocations result in a small value of vote share in equilibrium (i.e., S_A becomes smaller). On the other hand, if controller A starts late, it will have more resource allocations on each node per unit time, which leads to a larger value of vote share in equilibrium, but there will be less time left for influence to become effective.

To proceed, we move on to determining optimized starting times for different time horizons T and relative budgets b_A/b_B. Fig. 3.5 (a) shows the dependence of the optimal control time of the targeting controller, i.e., $T - opt\{t_a\}$ where $opt\{t_a\} = t_a^*$ stands for the optimal starting time of controller A as shown in Eq. (3.5), on network heterogeneity and time horizons. Generally, the optimized controller only uses its budget near the end of the campaign. This means that the system is initially only subject to the influence of the opponent. Only when close to the end of the campaign T, the optimized controller exerts several times the allocations of its opponent on the network. In doing so, the system approaches equilibrium gradually, which can be seen from the monotonic rise of votes shares in Figs. 3.4 (a) and (b) for $t \geq t_a$. In addition, for short-time horizons, optimal control times for networks with large heterogeneity tend to start later, while for long-time horizons, optimal control on highly heterogeneous networks should start slightly earlier. This dependence of optimal starting times on network heterogeneity can be explained by our earlier observations in Fig. 3.3 (a).

For short time horizons, the network is still far from equilibrium at the end of the competition. In other words, the network's degree of equilibrium l is small, which corresponds to the lower-left corner of Fig. 3.3 (a). Therefore, the state changes of vote shares are dominated by the low-degree nodes, which have shorter timescales. As highly heterogeneous networks have more lower-degree nodes, they will respond much quicker

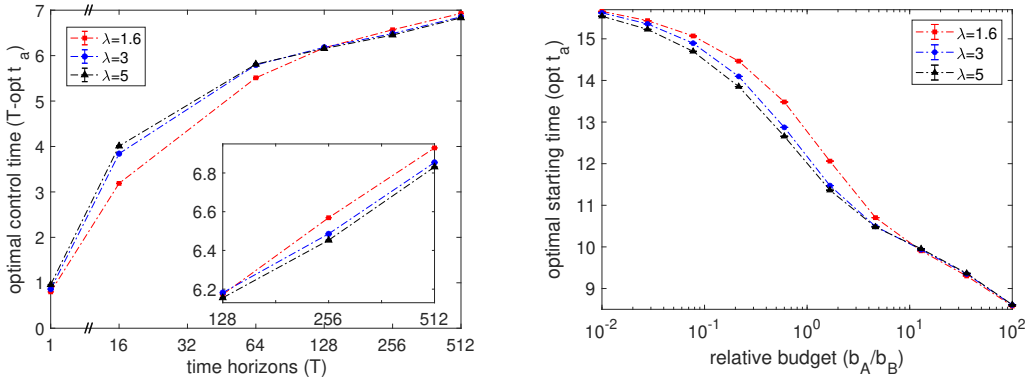

(a) (b)

FIGURE 3.5: Panel (a) shows the dependence of the optimal effective control time of controller A ($T - opt\{t_a\}$) on network heterogeneity and time horizons. The control gains of controller B pertaining to each node are all fixed as 1 per unit time from time 0. The total budgets of controller A are set to be the same as controller B's, e.g., for $T = 10$, $b_A = b_B = N \times T$. The y axis shows the difference between time horizon and optimized t_a. Panel (b) shows the dependence of the optimal starting time of controller A ($opt\{t_a\}$) on network heterogeneity and relative budgets. The time horizon is set as $T = 16$. The control gains of controller B pertaining to each node are all fixed as 1 per unit time from time 0. The total budgets of controller A are set to be a ratio to the budget of controller B as shown in the x axis. The calculations are based on networks with $N = 10000$ and $\langle k \rangle = 10.5$ and tested in 100 realizations. The legend "$\lambda = 1.6$" is identical to networks with power law distribution $p(k) \sim k^{-1.6}$. The initial states of opinion distribution are 0.5 for all nodes in all cases. The curves stand for networks with different power law exponents. Error bars indicate 95% confidence intervals.

to the resource allocations. Consequently, campaigns on highly heterogeneous networks should start slightly later than on less heterogeneous networks. In contrast, for long time horizons, the network is close to equilibrium at the end of the competition. In this case, the network's degree of equilibrium l approaches 1. From Fig. 3.3 (a), for a sufficiently larger l, the more heterogeneous the network, the larger the relaxation time. As a result, highly heterogeneous networks respond much more slowly to resource allocations, which explains an earlier start in optimized control.

To further confirm our conclusion, we also compare the degree of equilibrium for time horizon $T = 16$ and $T = 256$ in Figs. 3.4 (e) and (f). To this end, Fig. 3.4 (e) shows that, for short time horizon $T = 16$, the degree of equilibrium is always less than 0.82. Furthermore, as shown in Fig. 3.3 (a), when l is less than 0.9, the relaxation time τ^l_{relax} for highly heterogeneous networks is less than that for less heterogeneous networks. In this case, the highly heterogeneous networks respond faster to the exertion of control, so they can start later. In contrast, for long time horizons, the degree of equilibrium approaches 1 (see Fig. 3.4 (f)). In this case, less heterogeneous networks respond relatively faster to control. Therefore, optimal control for these networks can start later.

Fig. 3.5 (b) shows the dependence of the optimal starting times of the targeting controller (i.e., $opt\{t_a\} = t^*_a$) on network heterogeneity and relative budgets. In this figure,

we observe that when the optimized controller is in resource superiority, it will start earlier than when it is in budget disadvantage. Moreover, when relative budgets are not extremely small or large, there are larger differences in the optimal starting times for networks with varying levels of heterogeneity.

Network	Control settings		Opt t_a	Opt S_A	S_A for $t_a = 0$
Co-authorship network	$T = 16$	$b_A/b_B = 0.1$	14.588	0.279	0.098
		$b_A/b_B = 1$	12.658	0.746	0.500
		$b_A/b_B = 10$	10.652	0.961	0.909
	$T = 256$	$b_A/b_B = 0.1$	252.200	0.817	0.091
		$b_A/b_B = 1$	250.219	0.974	0.500
		$b_A/b_B = 10$	248.101	0.997	0.909
Email network	$T = 16$	$b_A/b_B = 0.1$	14.812	0.257	0.174
		$b_A/b_B = 1$	12.554	0.650	0.500
		$b_A/b_B = 10$	10.149	0.953	0.909
	$T = 256$	$b_A/b_B = 0.1$	251.651	0.713	0.091
		$b_A/b_B = 1$	249.578	0.969	0.500
		$b_A/b_B = 10$	247.965	0.997	0.909
Online social network	$T = 16$	$b_A/b_B = 0.1$	14.576	0.225	0.125
		$b_A/b_B = 1$	12.046	0.667	0.500
		$b_A/b_B = 10$	10.026	0.954	0.909
	$T = 256$	$b_A/b_B = 0.1$	251.334	0.748	0.091
		$b_A/b_B = 1$	249.603	0.970	0.500
		$b_A/b_B = 10$	247.949	0.996	0.909

TABLE 3.1: Comparison of optimized start times and shares of opinion A between two strategies of controller A: one using optimized inter-temporal allocations and the other one starting at time 0. Calculations are based on the co-authorship network, the email network, and the online social network, and are tested for varying time horizons $T = 16$ and $T = 256$, as well as different budget ratios $b_A/b_B = \{0.1, 1, 10\}$.

In addition to analyzing synthetic networks with controlled network heterogeneity, the idea of inter-temporal control can be further tested in real-world networks with different structures and link types. In Table 3.1, we present the optimized starting times (Opt t_a) and compare the resultant vote shares (Opt S_A) of controller A with the baseline scenario, where controller A begins targeting agents from time 0 (S_A for $t_a = 0$). We evaluate this across different time horizons ($T = \{16, 256\}$) and budget ratios ($b_A/b_B = \{0.1, 1, 10\}$).

Specifically, in Table 3.1, we select real-world networks that allow evaluation on both unweighted or weighted and undirected or directed graphs: The co-authorship network is an undirected and weighted graph, which depicts collaborations among scientists engaged in network theory (Newman, 2006). Our analysis is based on the largest connected component of this graph, encompassing 379 nodes and 914 links. This is a sparse graph with average degree $\langle k \rangle = 2.58$. Additionally, the email network in Table 3.1 represents the largest connected component of an email exchange system in University Rovira i Virgili (Guimera et al., 2003). The network is undirected and unweighted, consisting of 1133 nodes and 5451 links with average degree $\langle k \rangle = 9.62$. Moreover,

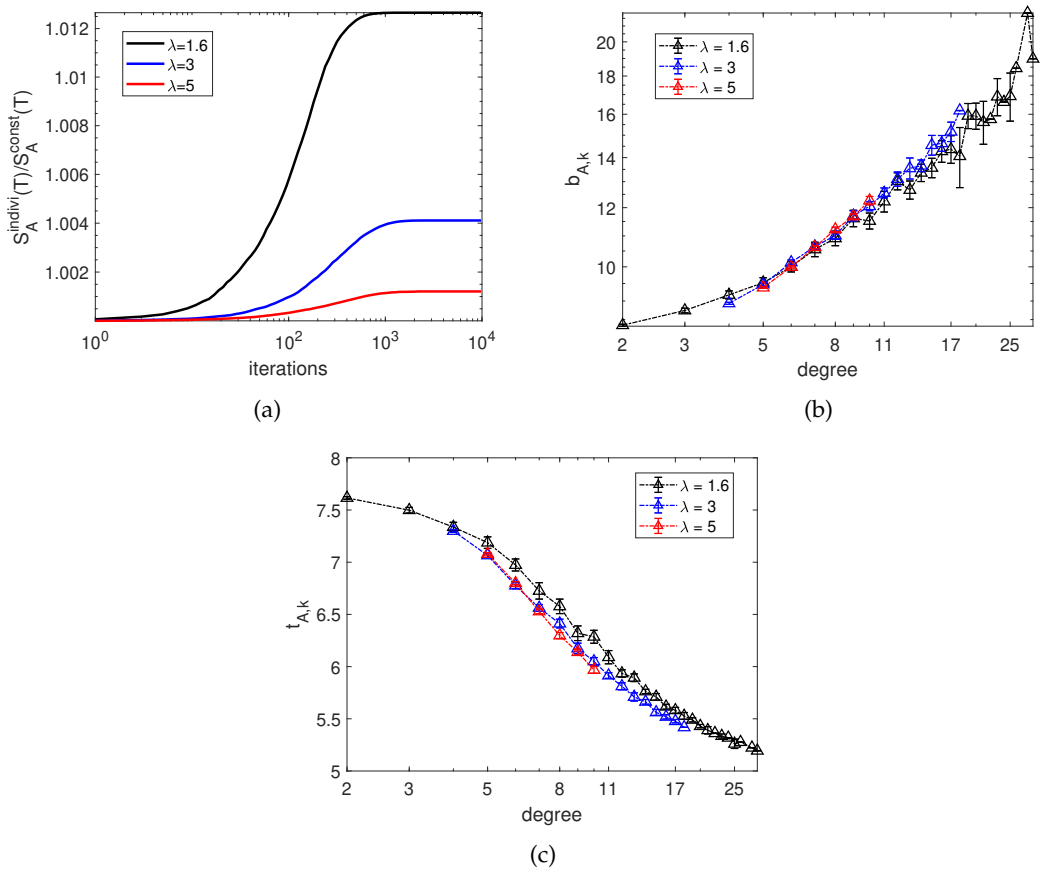

FIGURE 3.6: Panel (a) shows the optimization process characterized by the dependence of average relative vote share $S_A^{indivi}(T)/S_A^{const}(T)$ on iterations. $S_A^{indivi}(T)$ stands for vote shares at time T calculated by individual optimization while $S_A^{const}(T)$ represents vote shares calculated by assigning a single optimized starting time for the whole network. The black line, blue line, and red line stand for networks with degree of heterogeneity $\lambda = 1.6$, $\lambda = 3$ and $\lambda = 5$ respectively. Panels (b) and (c) show the corresponding distributions of budget allocations and starting times on nodes with degree k when the system depicted in panel (a) achieves maximum. All the calculations are based on networks with $N = 100$ and $\langle k \rangle = 6$ and tested in 10 realizations. The control gains of controller B pertaining to each node are all fixed as 1 per unit time from time 0. The time horizon T is set as $T = 10$. The total budgets of controller A are set to be the same as controller B's, i.e., $b_A = b_B = N \times T$. The initial states of opinion distribution are 0 for all nodes. Error bars indicate 95% confidence intervals.

we use the largest connected component of an online social network among students at University of California, Irvine (Panzarasa et al., 2009). This network contains 1294 nodes, and 19026 unweighted, directed links with average degree $\langle k \rangle = 10.93$. The results in Table 3.1 are consistent with the conclusions we draw from synthetic networks that when the optimized controller is in resource superiority, it will start earlier than it is in a budget disadvantage (see Fig. 3.5). This demonstrates the consistency of the optimization of starting times across topologies.

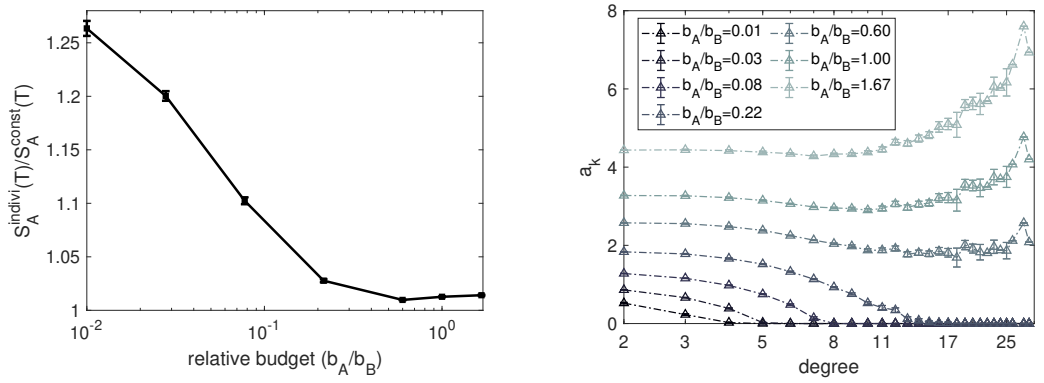

(a) (b)

FIGURE 3.7: Panel (a) shows the dependence of average relative vote share $S_A^{indivi}(T)/S_A^{const}(T)$ on relative budgets b_A/b_B. Panel (b) shows the dependence of control gains by individual optimization on nodes' degrees and relative budgets. All the calculations are based on networks with degree of heterogeneity $\lambda = 1.6$. The network size is $N = 100$ and $\langle k \rangle = 6$ and tested in 10 realizations. The control gains of controller B pertaining to each node are all fixed as 1 per unit time from time 0. The time horizon T is set as $T = 10$. The initial states of opinion distribution are 0 for all nodes. Error bars indicate 95% confidence intervals.

3.4.2.2 Optimal Strategies in the Individual Optimization Setting

Above, we have investigated the optimal strategy of controller A in a simplified scenario. As it is too restrictive that all nodes start at the same time with the same intensity, a question naturally arises: if the controller has more flexibility to deploy its resources, will it lead to a considerable improvement in the vote share? Regarding this, we have proposed the individual optimization scenario (see Eq. (3.6)) where the controller can choose different starting times and split its resources unevenly for individual nodes. Here, we use the stochastic hill climbing algorithm for the optimization. The procedure is described as follows: (i) Start with a given network configuration and an initialization of starting times $\{0 \leq t_{a,i} <= T\}$ and budget allocations $\{b_{A,i}\}$ for each agent $i = 1, \cdots, N$. Note that, the initial budget allocations $\{b_{A,i}\}$ should meet the budget constraint $\sum_i b_{A,i} = b_A$. In practice, we set the initial budget allocations as $\{b_{A,i}\} = b_A/N$. Let L_{max} be the maximum number of iterations. Then, compute the vote share of the initial configuration via Eqs. (3.1) and (3.3). (ii) If the current iteration is less than L_{max}, continue with step (iii). Otherwise, jump to step (iv). (iii) Generate a random number r according to the uniform distribution with lower bound 0 and upper bound 1. If $r < 0.5$, randomly pick two nodes i and j ($i \neq j$), and transfer a random fraction of budget allocations from node i to node j. It's important to note that transferring budgets between nodes ensures the total budget remains unchanged, thus satisfying the budget constraint. Otherwise, randomly select one node and randomly modify its starting time within the range of $[0, T]$. Update the vote share only if improvements are achieved via the above modification. Return to step (ii). (iv) The procedure is ended.

In the following, we present the numerical results for individual optimization according to the above procedure. In Fig. 3.6 (a), we evaluate the performance of individual optimization by comparing the vote shares achieved by node-specific starting times and budget allocations with the simple scenario where the optimized controller assigns a single starting time for the whole network. From our analysis of Fig. 3.6 (a), we observe that, roughly N^2 iterations yield near-optimal allocations, as each node is updated N times on average and no further improvements are observed beyond this point. Therefore, in the following experiments, we set $L_{max} = N^2$. Due to the limited improvements (not exceeding 2%, see Fig. 3.6 (a)) achievable via the node-specific optimization, we deviate from using the same initial distribution of opinion states 0.5 (i.e., half of the nodes start with opinion B) as shown in Figs. 3.4 and 3.5, which represent a neutral starting point. Instead, we adopt an initial distribution of opinion states 0 (i.e., all nodes start with opinion B) that will lead to slightly larger improvements compared to the 0.5 distribution (see Appendix A.1 for the quantitatively similar results of other initial state settings). Even though the improvements in vote shares achieved by individual optimization are very small, a clear difference in vote shares achievable for networks with different degrees of heterogeneity is identified. Specifically, individual optimization is more efficient for networks with larger heterogeneity. A possible explanation for this observation is that: We have confined the average degree of different networks to be the same. Therefore, the degrees of highly heterogeneous networks are distributed in a larger range compared with less heterogeneous networks. In the individual optimization case, the larger difference between degrees can be fully exploited in a highly heterogeneous network. For the same reason, bigger improvements can be achieved for larger networks.

To further confirm our hypothesis, in Figs. 3.6 (b) and (c), we present the dependence of optimized budget allocations and starting times of controller A for nodes with degree k (denoted as $b_{A,k}$ and $t_{A,k}$) on nodes' degrees. We observe an inverse relationship between the dependence of budget allocation on nodes' degrees and the dependence of optimal starting time on nodes' degrees if both controllers have the same budget constraints. In more detail, a node with a larger degree should be allocated more resources but start earlier in the equal budget case. Moreover, networks with different degrees of heterogeneity have a similar distribution of budget allocations and starting times regarding nodes' degrees.

To further explore if the limited improvements of vote shares in individual optimization are a result of the specific control configuration (i.e., $b_A = b_B$ and $T = 10$), we conduct more numerical experiments based on different relative budgets. Fig. 3.7 (a) shows the dependence of improvements of vote shares via individual optimization on relative budgets. We find that for a small relative budget, a large improvement can be achieved. In Fig. 3.7 (b) we further explain why the individual optimization is more efficient for small budget scenarios. In more detail, we clearly see two regimes

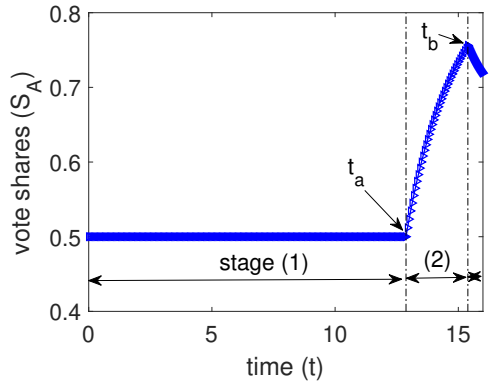

(a) evolution for networks with $\lambda = 1.6$

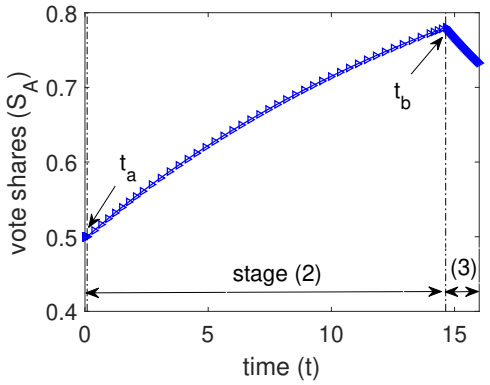

(b) evolution for networks with $\lambda = 5$

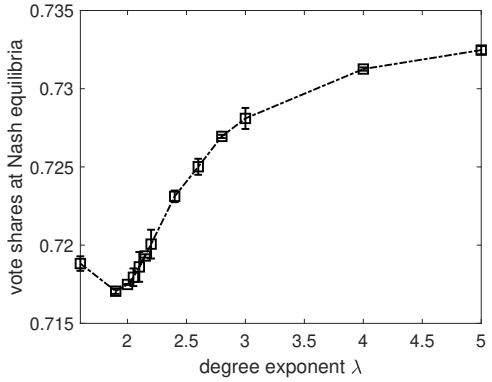

(c) vote shares in Nash equilibria

FIGURE 3.8: Panels (a) and (b) show the evolution of vote shares in the game-theoretical setting for networks with degree exponents $\lambda = 1.6$ (panel (a)) and $\lambda = 5$ (panel (b)) at Nash equilibria. Panel (c) shows the vote shares at Nash equilibria across degree exponent values. Calculations are based on networks with $N = 10000$ and $\langle k \rangle = 10.5$. The budgets of controller B are 10000 in total and 100000 for controller A. The initial states of opinion distribution are 0.5 for all nodes and the time horizon is set as $T = 16$. The dotted lines distinguish three stages. t_a and t_b shown by arrows are the solutions for Nash equilibria.

in Fig. 3.7 (b). For a small budget, the control gains decrease with the nodes' degree while for a large budget with the increase of degree, the optimized controller puts a larger allocation on it. Moreover, for small budgets, the optimized controller only targets nodes with small degrees and ignores large-degree nodes. By doing so, the optimized controller does not need to waste its resources on nodes with lower controllability, and can achieve significant improvements in the vote shares compared with the non-individual-optimization scenario where the controller has to target all nodes equally. In conclusion, by applying individual optimization, the vote shares can be improved, especially for highly heterogeneous networks, but the improvement is indeed very small for equal or large relative budget configurations.

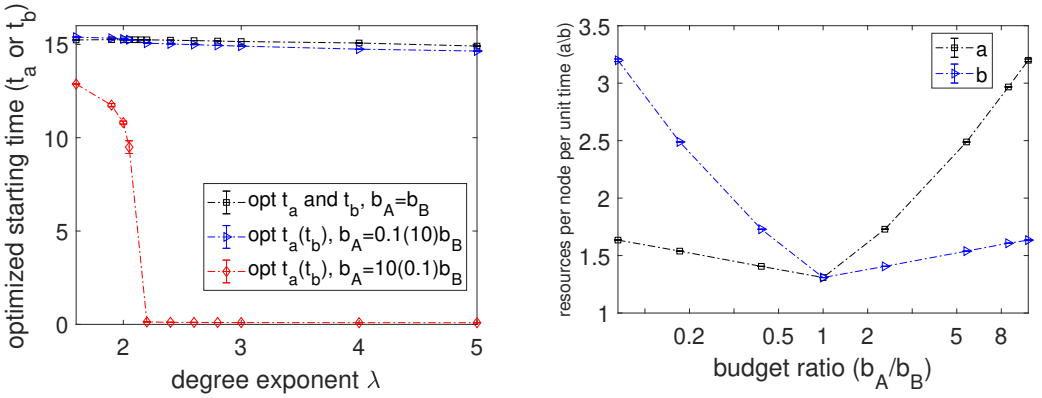

(a) (b)

FIGURE 3.9: Panel (a) shows the dependence of optimized starting times on network heterogeneity and budgets under the game-theoretical settings. Black squares show the best strategy for both controller A and controller B in the condition of equal budgets. Blue triangles show the best strategy for controller A if $b_A = 0.1b_B$ and the best strategy for controller B if $b_A = 10b_B$. Red rhombus show the best strategy for controller A if $b_A = 10b_B$ and the best strategy for controller B if $b_A = 0.1b_B$. Panel (b) shows the dependence of resource allocations on relative budgets. The black and blue lines represent the resources for each node per unit time allocated by the controller A and B respectively. Calculations are based on networks with $N = 10000$ and $\langle k \rangle = 10.5$ and tested in 10 realizations. The time horizon is set as $T = 16$ for all cases.

3.4.2.3 Optimal Strategies in the Game-theoretical Setting

Similarly, we start our analysis for optimal strategies in the game-theoretical setting by depicting the evolution of vote shares in the Nash equilibrium. For this purpose, the Nash equilibria are obtained by the iterative gradient ascent algorithm (see Algorithm 1 in Section 3.3). Due to the symmetry of the game for the equal budget, controllers A and B will take the same strategy in this scenario.

To fully explore the game-theoretical setting, we start with the context where one of the controllers is in resource superiority. In more detail, Figs. 3.8 (a) and (b) show the evolution of vote shares for a short time horizon $T = 16$ subject to control in a Nash equilibrium when controller A has a larger budget than controller B. We note that, in both Figs. 3.8 (a) and (b), controller A always starts earlier than controller B. Given that the initial setup for controller A and controller B is the same except for the budgets, we see that the controller with a budget advantage will start control earlier.

Furthermore, a closer inspection of Figs. 3.8 (a) and (b) gives us an insight into the game played by controller A and controller B to maximize their respective pay-off functions. In general, the evolution of vote shares can be divided into three stages (see the sequence numbers in Figs. 3.8). (i) No controller exerts influence. (ii) The first mover changes system states. (iii) The second mover starts control and the system is subject to both controllers. In contrast to the vote-share trajectories in the constant-opponent setting where the vote shares increase monotonically when the active controller starts

control, we note that there are turning points in the intersection of the second stage and the third stage. In the game-theoretical setting, the controller with budget advantages (i.e., controller A) always starts control first to seize the initiative. In contrast, the other controller will concentrate its resources in the final stage in order to use limited resources more effectively and achieve some pull-back from the other controller's initial advantage.

In addition, we observe that the starting times of the first and second movers in highly heterogeneous networks tend to be later than in less heterogeneous networks. This phenomenon can be explained by our earlier observations about network's timescales towards equilibrium in Section 3.4.1. Since for short time horizons (e.g., $T = 16$), highly heterogeneous networks respond to the resource allocation faster, the networks with degree exponent $\lambda = 1.6$ tend to start control later compared with networks with degree exponent $\lambda = 5$.

Building upon the prior analysis showing the evolution of vote shares over time for networks with specific degree exponents of $\lambda = 1.6$ (Fig. 3.8 (a)) and $\lambda = 5$ (Fig. 3.8 (b)), Fig. 3.8 (c) expands the examination across a range of degree exponent values from 1.6 to 5. While the previous figures focused on the dynamic changes in vote shares within a given network structure, this aggregated view captures the end-state vote shares obtained at Nash equilibria. The consistency of vote shares at equilibria, with values ranging from 0.717 to 0.733 despite differing degree heterogeneity and controller timing, highlights the robustness of the proposed approach.

Next, we consider the effect of relative budgets on optimal control. The optimal starting times of controller A and controller B for different budget ratios in networks with different degrees of heterogeneity characterized by different degree exponents are shown in Fig. 3.9 (a). It becomes clear that the controller with a budget advantage will always start earlier compared with the controller with a lower budget. Moreover, for a short time horizon (e.g., $T = 16$), controllers in highly heterogeneous networks will start later, which is shown by the monotonically descending curves. This is consistent with the results in the constant-opponent case.

We further explore the dependence of allocation per node on the budget ratio in Fig. 3.9 (b). It is clear that, although the controller with budget advantages tends to start earlier, it allocates more resources on each node per unit time as well. Therefore, for stage (3) in Figs. 3.8 (a) and (b), the second mover can only buffer the increase of S_A, it is impossible for the second mover to gain a larger vote share than the first mover.

3.4.3 Individual Optimization Considering Agent Heterogeneity

After investigating optimized but different starting times and budget allocations for individual agents, a natural extension is to combine agent heterogeneity with individual

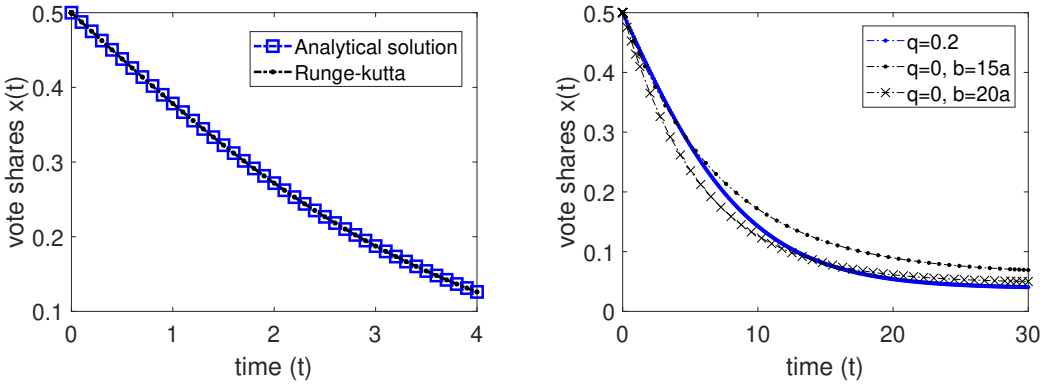

(a) (b)

FIGURE 3.10: Panel (a) compares vote shares calculated via analytical solution (see Eq. (3.27)) and direct integration using a Runge-Kutta method of Eq. (3.26). In panel (a), all nodes are assigned with zealotry $q = 0.2$ and the control gains of controller A pertaining to each node are all fixed as 1 per unit time from time 0 while controller B does not target nodes. Panel (b) compares the evolution of vote shares with and without zealots. The blue dotted line presents vote-share trajectory in the absence of controller B but considering the zealotry of all nodes as $q = 0.2$, while the black dotted line and the "x" line are for vote shares in the presence of controller B but not zealots with different relative budget $b = 15a$ and $b = 20a$. Calculations are based on complete networks with $N = 100$. The initial states of opinion distribution are 0.5 for all nodes.

optimization. Existing literature for the voter model has introduced so-called *zealots* (Mobilia et al., 2007) to represent agents that have biases towards adopting one of the opinions. Here, we model the level of zealotry with q_i, which describes the decreased chances of adopting opinion A for node i. Following Moreno et al. (2020), the rate equation for the probability of agent i holding opinion A changes to

$$\begin{aligned}\frac{dx_i}{dt} =&(1-q_i)(1-x_i)\frac{\sum_j w_{ji}x_j + a_i(t)}{\sum_j w_{ji} + a_i(t) + b_i(t)} - x_i\frac{\sum_j(1-x_j)w_{ji} + b_i(t)}{\sum_j w_{ji} + a_i(t) + b_i(t)} \\ =&(1-q_i+q_ix_i)\frac{a_i + \sum_j w_{ij}x_j}{k_i + a_i + b_i} - x_i.\end{aligned} \tag{3.24}$$

Before exploring the influence-maximizing strategies in the context of agent heterogeneity, we first show the impact of zealotry on the evolution of vote shares. It is intuitive from Eq. (3.24) that a larger q_i will result in a smaller probability for an agent to change from opinion B to opinion A. To have a quantitative understanding of the role of different levels of zealotry, we further investigate Eq. (3.24) in a simple scenario where there are two agents with zealotry q_1 and q_2 ($q_1 > q_2$) in the absence of external

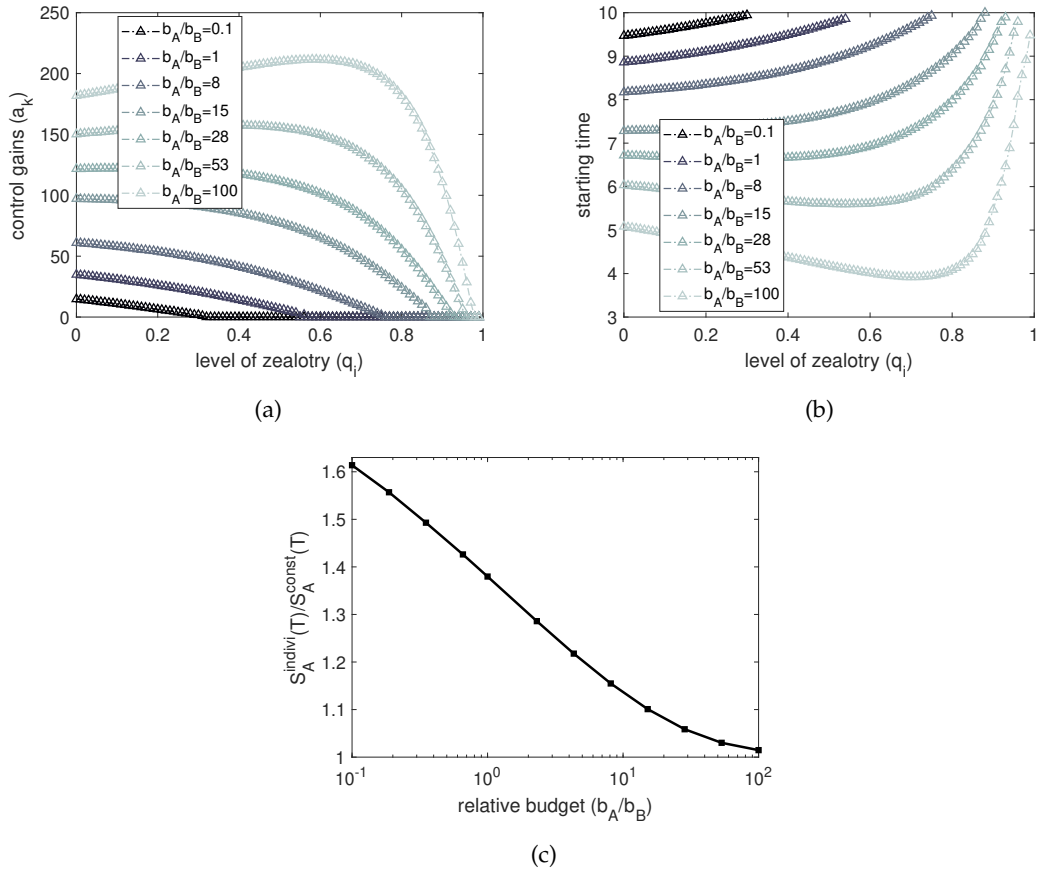

FIGURE 3.11: Panels (a) and (b) show the dependence of control gains and starting times on zealotry q_i and relative budget b_A/b_B obtained via the individual optimization method on a complete network with size $N = 100$. Panel (c) shows the resulting improvements in vote shares via individual optimization compared with assigning a single and optimal starting time for the whole network for varying relative budgets. The control gains of controller B pertaining to each node are all fixed as 1 per unit time from time 0. The time horizon T is set as $T = 10$. The initial states of opinion distribution are 0 for all nodes.

controllers. Assuming both nodes start from the same initial state x_0, then at the beginning of the evolution of the vote share, we have

$$\begin{aligned}\frac{dx_1}{dt} - \frac{dx_2}{dt} =& (1 - q_1 + q_1 x_0) - (1 - q_2 + q_2 x_0) \\ =& (q_2 - q_1)(1 - x_1) < 0.\end{aligned} \tag{3.25}$$

As $\frac{dx_1}{dt} - \frac{dx_2}{dt}$ is a negative value, a larger zealotry will lead to a faster decrease of vote shares x_i.

It is clear that the presence of an opposing controller will also lead to a decrease in vote shares. However, what is the difference between the influence of the opposing controller and zealotry towards the opponent on the evolution of vote-share trajectories? To isolate the effects of zealotry towards the opponent and the opposing controller, we solve Eq. (3.24) analytically in the following two scenarios. One is that all nodes have

the same level of zealotry q but only controller A exerts equal budget allocation a on every node. The other is controller B allocates equal budget b on every node in the absence of agent heterogeneity, i.e., $q = 0$. To further simplify the expression, we assume a fully connected network with size N.

For the first scenario, Eq. (3.24) can be written as

$$\frac{dx}{dt} = \frac{qN}{N+a}x^2 + \frac{aq - qN - a}{N+a}x + \frac{(1-q)a}{N+a}. \tag{3.26}$$

Solving Eq. (3.26), we obtain

$$x(t) = \frac{a(q-1) + Nq}{\frac{e^{\frac{ax}{a+N} - qx}(a(q-1)+Nqx_0)}{x_0 - 1} - Nq} + 1 \qquad (q \neq 0), \tag{3.27}$$

where x_0 is the initial state of the system. To verify the correctness of Eq. (3.27), we compare $x(t)$ calculated in two ways in Fig. 3.10 (a). One is by numerically integrating Eq. (3.26) using the Runge-Kutta method. The other is by using the analytical solution given in Eq. (3.27). The close agreement between the two curves in Fig. 3.10 (a) confirms that Eq. (3.27) provides the correct solution to the differential equation in Eq. (3.26).

For the second scenario, Eq. (3.24) can be written as

$$\frac{dx}{dt} = \frac{a + Nx}{N + a + b} - x. \tag{3.28}$$

Solving Eq. (3.28), we obtain

$$x(t) = \frac{(a(x_0 - 1) + bx_0)e^{-\frac{x(a+b)}{a+b+N}} + a}{a+b}. \tag{3.29}$$

To intuitively demonstrate the difference between the influence of the opposing controller and zealotry on network dynamics, we proceed with numerical results based on complete networks in Fig. 3.10 (b) where we squeeze the vote-share trajectory in the presence of zealots by tuning the relative budget b/a. In 3.10 (b), we find that zealotry has a significant effect on the evolution of vote shares as a small fraction of zealotry has a similar influence on vote shares as more than 10 times of the budget ratio b/a.

To further gain more intuition about how different levels of zealotry influence the optimal starting times and budget allocations of individual agents, we explore the influence-maximizing strategies in the complete network with a uniform distribution of zealotry in the range $[0, 1]$. From Figs. 3.11 (a) and (b), we find that there are two regimes of optimal control gains and starting times. For a small or not large enough budget ratio (e.g., $b_A/b_B = 0.1, 8$), controller A will only target nodes with low levels of zealotry. In these scenarios, the control gains decrease with the increase of zealotry,

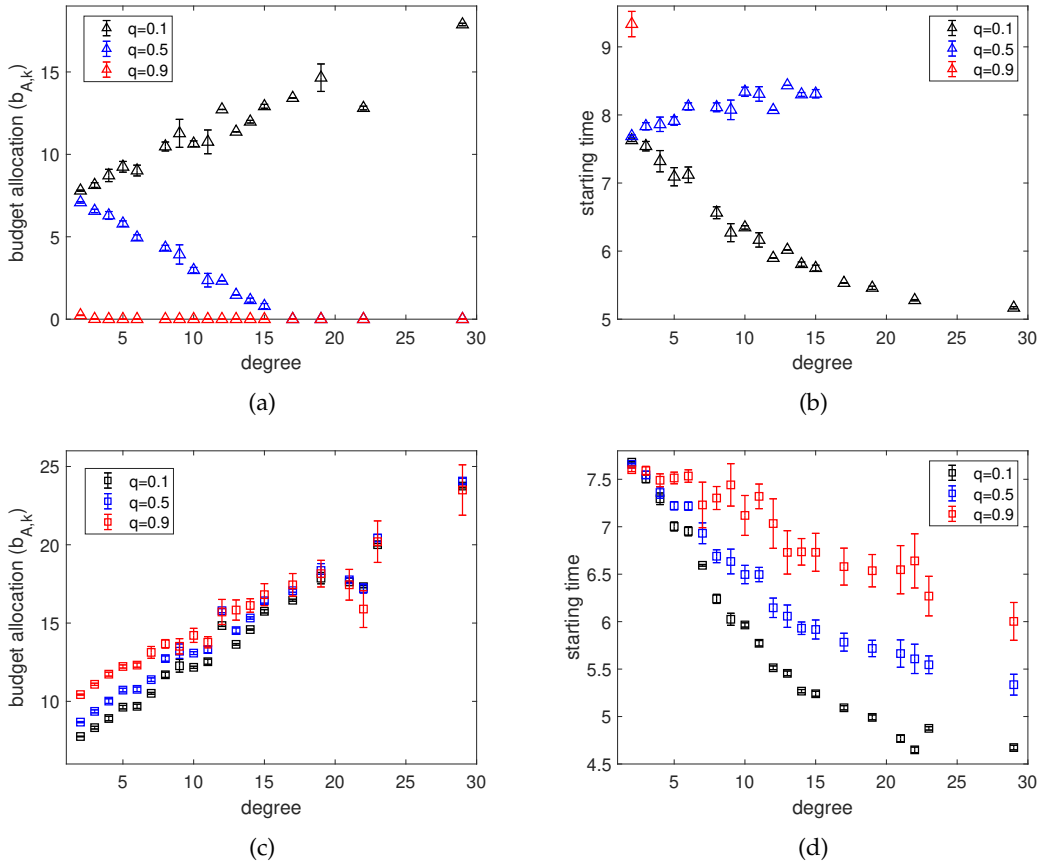

FIGURE 3.12: Panels (a) and (b) show the dependence of budget allocations and starting times on nodes' degrees for varying zealotry obtained by individual optimization for zealots. Panels (c) and (d) are the dependence of budget allocations and starting times on nodes' degree obtained by individual optimization for normal agents. Calculations are based on heterogeneous networks with $N = 100$, $\langle k \rangle = 6$, degree of heterogeneity $\lambda = 1.6$ and tested in 10 realizations. In each realization, we randomly choose 20% population of the network becoming zealots with the zealotry marked by q, i.e., $q = 0.1$, $q = 0.5$, $q = 0.9$. The control gains of controller B pertaining to each node are all fixed as 1 per unit time from time 0. The time horizon T is set as $T = 10$. The initial states of opinion distribution are 0 for all nodes. Error bars indicate 95% confidence intervals.

while the optimal starting times become later for highly zealous nodes. However, for sufficiently large budget (e.g., $b_A/b_B = 100$), within a certain scope of levels of zealotry, the optimized controller A will allocate more resources and start earlier for more zealous nodes. This pattern changes when the curve reaches a critical point of zealotry from where with the increase of zealot, the optimized controller will distribute fewer resources and starts later. Similar to Fig. 3.7 (a), the individual optimization is more efficient for small budgets as it can choose only to target a few nodes with high controllability.

Next, we focus on heterogeneous networks to further explore how the zealotry and nodes' degrees affect optimal starting times and budget allocations. For our experiments, we randomly select 20% of the population to become zealots with the same

zealotry q, and keep the remaining nodes as normal agents. In Figs. 3.12 (a)-(d) we show the relationship between optimal control gains and node degree and level of zealotry. We find that for medium and large zealotry, most resources will be allocated to normal agents. Moreover, zealots will behave in the opposite way of normal agents. In more detail, with the increase of degree, zealots will be allocated with fewer resources while normal agents will be targeted with larger resources. For the starting times, the optimized controller will start later for larger-degree zealots, however it will start earlier for larger-degree normal agents.

3.5 Summary

This chapter investigates the influence maximization problem from the perspective of inter-temporal budget allocations on the voter model. In doing so, we relate the problem of influence maximization to network control where external controllers exert influence via building unidirectional connections with agents in the network. We then integrate the time information into the opinion dynamics by exploring the optimal campaign starting times to achieve a maximum vote share within limited time horizons.

Intuitively, in the context of inter-temporal network control, one or both controllers have to design strategies to make efficient use of their budgets over time. This results in the following trade-off: If the controller starts allocating later, it has more disposable budget per unit of time but less time left for its influence to become effective. Here, in order to obtain analytical solutions for optimal allocations, we first explore the inter-temporal budget allocation in a simplified scenario where controllers have only the flexibility to determine when to start control, and once started, they have to target all agents equally. Then we conduct numerical experiments on a more complicated scenario where the controller optimizes starting times and budget allocations for individual agents in the network, which is referred to as individual optimization. Given that the individual optimization concentrates on assigning different starting times and allocations on individual nodes, we also take agent heterogeneity into consideration and study the influence of agent heterogeneity on inter-temporal budget allocations.

By doing so, our main findings are as follows:

1. In the constant-opponent setting, as we fix one controller to start control from the very beginning, the optimal strategy for the optimizing controller is to initially leave the system subject to the influence of the opposing controller and then only use its budget closer to the end of the campaign. Moreover, in the game-theoretical setting, the controller with a budget advantage will start control earlier than the controller with a smaller budget to ensure that, by the time its opponent begins to take control, it has gained an advantage in vote shares.

2. For short time horizons, the optimized controller tends to start control later in highly heterogeneous networks compared to less heterogeneous networks. In contrast, for long-time horizons, an earlier start is preferred for highly heterogeneous networks.

3. Compared with the simplified scenario of finding an equal and optimal starting time for all agents, the individual optimization only has limited superiority for equal or larger budget settings. However, if the optimized controller is at a budget disadvantage, by applying individual optimization, it will only focus on low-degree nodes and thus gain a large improvement in the vote share.

4. When considering agent heterogeneity in the context of individual optimization, we find that nodes with lower controllability (i.e., zealots with zealotry within a certain threshold) will receive more budget allocations and start earlier if the active controller has a large budget. Conversely, these nodes receive fewer budgets and activate later when budgets are limited. Further analysis of networks with heterogeneous degree distributions reveals the hierarchy of control difficulties: Zealots at hubs are the most difficult to control, followed by periphery zealots, normal hub nodes, and lastly normal periphery nodes. Optimal budget allocations and starting times appear to involve balancing budget availability with the difficulties of controlling each node.

Chapter 4

Opponent-strategy Reconstruction in the Voter Model

In this chapter, we explore the issue of accelerating the convergence of inference by discovering more informative observational data. More specifically, we develop and explore a framework of how convergence of estimates can be accelerated through targeted interaction with the networked dynamics. Our framework thus supposes that the observer can influence the dynamical process on the network and we explore how such influence can be optimally deployed to improve the inference of unknown parameters of the dynamics. To be more concrete, this chapter explores the problem of opponent strategy inference in the setting of competitive voting dynamics. The problem we are interested in is that one of the controllers can change its control allocations to accelerate its learning of the opposing controllers' targeting through observation of the voting dynamics.

This chapter is organized as follows: Section 4.1 sets the context and motivation for studying the inference acceleration problem in the voter model. Section 4.2 formalizes the problem of accelerating opponent strategy inference for the voter model and presents heuristics for solving the opponent strategy inference problem. Section 4.3 shows the corresponding results after applying the heuristics. Section 4.4 summarises the main findings.

4.1 Introduction

As mentioned in Section 2.5, most previous research has focused on the network reconstruction problem under the assumption that the entire time series of the network dynamics is accessible to ensure sufficient information is provided for accurate network inference. However, as investigated by the work of Braunstein et al. (2019), in

many real-world cases such as neuron cascades and epidemic spreading, the first stage of propagation is hard to measure and only a limited number of data points will be observed. Despite the experimental or technical limitations for data collection, obtaining high-precision estimations with less data is always desirable, especially when the measurement for the dynamical quantities is costly (Guo and Luk, 2013). Motivated by the dilemma between the availability of observational data and the accuracy of inference, in this chapter, we explore the issue of accelerating the convergence of inference by discovering more informative observational data. However, different from previous literature such as Braunstein et al. (2019), we develop and explore a framework of how the convergence of estimates can be accelerated through targeted interaction with the networked dynamics. Our framework thus supposes that the observer can influence the dynamical process on the network and we explore how such influence can be optimally deployed to improve the inference of unknown parameters of the dynamics.

To derive dynamical process parameters or reconstruct network topology from observational data, it is often necessary to draw on domain-specific expertise (Brugere et al., 2018). Here, we place the problem of speeding up inference in the context of opinion dynamics using the well-known competitive influence maximization framework (Li et al., 2018; Goyal et al., 2019), which studies the competition among external controllers who aim to maximally spread their opinions in the network through strategically distributing their influencing resources. Specifically, a common assumption while investigating the competitive influence maximization problem is that the external controllers are unaware of the strategy being used by their opponents during the competition. However, the need of inferring the opponent's behaviour in a short time frame is observed in many real-world contexts, such as finding out the source of fake news as soon as possible in the social network to stop it from spreading (Nguyen et al., 2012), analysing the provenance of extreme opinions to prevent radicalization (Galam and Javarone, 2016), and uncovering the strategy of the opposing political parties before a given deadline to gain advantages in the election (Hegselmann et al., 2014). Therefore, accelerating the inference to obtain better estimates of opponent's strategies from dynamical data within a short time frame is an important problem relevant to competitive influence maximization.

To be more concrete, in this chapter, we explore the problem of opponent strategy inference in the setting of the competitive voting dynamics as studied in Masuda (2015); Romero Moreno et al. (2021a); Brede et al. (2019a). This choice is motivated by the popularity of the voter model in opinion dynamics as well as its high levels of tractability (Redner, 2019). Specifically, in the voting dynamics, opinions are represented as binary variables, and each agent in the network holds one of two opinions. On top of the internal agents, following the work of Romero Moreno et al. (2020, 2021a); Brede et al. (2019a), the external controllers exert their influence on the network by building

unidirectional connections with agents, in which the intensity of their targeting is represented by link weights. The opinion propagates according to the rules that agents flip their opinion states with probabilities proportional to the number of agents with opposing opinions and link weights from opposing controllers (Masuda, 2015). The problem we are interested in is that one of the controllers can change its control allocations to accelerate its learning of the opposing controllers' targeting through observation of the voting dynamics.

Since we model the way of exerting influence from external controllers by building unidirectional connections with agents in the network, the connections from the external controller can also be viewed as edges that constitute part of the network topology. Therefore, our research problem of opponent strategy inference is closely related to the topic of network structure inference. There is a rich literature in the field of reconstructing network structure from information flows (Brugere et al., 2018), and a detailed review of the related work within the domains of epidemiology and information spreading is given in Section 2.5. Most relevant to our modelling approach, Li et al. (2017); Chen and Lai (2018); Zhang et al. (2018) infer the network topology from time series of binary-state data. More specifically, Li et al. (2017); Zhang et al. (2018) treat the connections between agents as binary variables, and transform the network inference problem to identify the existence of binary links between agents. Hence, these approaches are unsuitable for inferring continuous interaction intensity between agents and the external controllers. Further to the works of Li et al. (2017); Zhang et al. (2018), Chen and Lai (2018) remove the binary restriction and consider the network inference problem in a continuous space by developing a data-driven framework to predict link weights. Nevertheless, none of these works investigate the network inference problem from the perspective of manipulating the opinion diffusion process to accelerate the convergence of estimation, which is an important lever if one wants to obtain an estimate with an accuracy guarantee within a short and limited observation time.

To address the current gaps in accelerating the convergence of inference, in this chapter, we relate the problem of accelerating opponent strategy inference with network control. By doing so, we assume an active strategic controller who tries to minimize the uncertainty of inference of an opponent's strategy by optimally allocating its control resources to agents in the network based on the voter model. In other words, we explore how a controller can modify network dynamics such that the influence of opponents becomes easier to identify. Note that, we always assume only limited resources are available for the active controller to interfere with the network dynamics, since for most real-world applications (Goyal et al., 2019; Masucci and Silva, 2014), there are natural resource constraints.

Specifically, in this chapter, we make the following contributions to address research question 2 on network inference acceleration outlined in Section 1.1: 1) To address research questions 2a and 2b, we introduce a new perspective of strategically interacting

with opinion dynamics to quicken network inference. Here, we derive estimators of the opponent's strategy via maximum likelihood estimation and provide uncertainty quantification of estimators via the Fisher information (Ly et al., 2017). This, in turn, is used to inform decisions on the optimal allocations and understand the process of inference acceleration. Moreover, we've developed several heuristic algorithms to minimize estimator variance, thereby speeding up inference. 2) To address research questions 2d and 2e, we make the following findings. Optimal resource allocations vary with control scenarios: it's inversely proportional to neighboring opinion states in single-node control and depends on resource availability in multi-node control, focusing on the inferred node with limited resources and on neighbors with more resources. Strategic allocations gain importance with increased resources, especially in less heterogeneous networks. This highlights how both the network topology and budget availability critically influence the effectiveness of budget allocations in revealing an opponent's strategy.

4.2 Model Description

We consider a population of N agents exchanging their opinions through a social network. The undirected social connections between agents are represented by a symmetric adjacency matrix $\boldsymbol{W} = \{w_{ij}\}_{i,j=1}^{N}$, with $w_{ij} = 1$ indicating the existence of a social link between agent i and agent j and $w_{ij} = 0$ otherwise. Note that, agents i and j are called neighbours if they are directly linked. Moreover, we assume that each of the N agents holds a binary opinion at time t, denoted as $s_i(t) \in \{0, 1\}$ $(i = 1, ..., N)$. In addition to internal agents, we introduce two external controllers: A and B. These controllers are zealots who have fixed opinions: $s_A(t) = 1$ for controller A and $s_B(t) = 0$ for controller B for all $t \geq 0$.

This chapter departs from the framework in Chapter 3. There, we were engaged in a detailed exploration of how best to distribute different quantities of budget allocation across time, ensuring we neither surpassed the set budget nor the temporal constraints. This process was defined using the terms $a_i(t)$ and $b_i(t)$, and was bounded by constraints of the cumulative integration: $\sum_{i=1}^{N} \int a_i(t)dt \leq b_A$ and $\sum_{i=1}^{N} \int b_i(t)dt \leq b_B$. However, in this chapter, our primary interest lies in a stepwise optimization of the budget without the inter-temporal considerations of Chapter 3.

To avoid terminological confusion with Chapter 3, we introduce notation for the control gains from controllers A and B to agent i at time t as $J_{A,i}(t)$ and $J_{B,i}(t)$ respectively. By building unidirectional and non-negatively weighted links $J_{A,i}(t) \geq 0$ and $J_{B,i}(t) \geq 0$ to agent i at time t, the two external controllers A and B exert their influences on the social network and therefore interact with the intrinsic opinion dynamics. Here, the sum of the link weights are subject to budget constraints at each step, i.e., $\sum_{i=1}^{N} J_{A,i}(t) \leq b_A$ and

$\sum_{i=1}^{N} J_{B,i}(t) \leq b_B$, where b_A and b_B are the total resources available to controller A and B respectively.

In this study, we adopt the synchronous voter model (Gastner, 2015) for opinion propagation through the social network. Different from the asynchronous voting dynamics we described in Section 3.2 where only one agent is updated at each time step, the synchronous voter model considers a simultaneous update for the whole network, with the aim of encoding more information in shorter time series. Specifically, one significant drawback of the asynchronous approach is that nodes remain silent for the majority of the time, leading to vast periods of inactivity. This not only results in slower propagation of information but also provides less data for inference, potentially missing out on capturing rapid opinion changes. For instance, in a network with size $N = 1000$, on average it will take 1000 time steps for a node to update its state once in the asynchronous voting dynamics. Instead, the synchronous voter model assumes a parallel and discrete-time opinion updating for the whole population as follows: at time t, agent i ($i = 1, ..., N$) updates its opinion to $s_i(t+1) = 1$ with probability

$$Pr(s_i(t+1) = 1) = \frac{J_{A,i}(t) + \sum_{j=1}^{N} s_j(t) w_{ij}}{\sum_{j=1}^{N} w_{ij} + J_{A,i}(t) + J_{B,i}(t)}, \tag{4.1}$$

and to $s_i(t+1) = 0$ with probability

$$Pr(s_i(t+1) = 0) = \frac{J_{B,i}(t) + \sum_{j=1}^{N} \left(1 - s_j(t)\right) w_{ij}}{\sum_{j=1}^{N} w_{ij} + J_{A,i}(t) + J_{B,i}(t)}. \tag{4.2}$$

As shown in the equations of $Pr(s_i(t+1) = 1)$ and $Pr(s_i(t+1) = 0)$, the opinion transition probabilities are determined only by the neighbouring states of the updated agent and the weighted links from the controllers, and they are independent of the current opinion of the updated agent.

From the perspective of external controllers, they aim to maximize their influence by strategically allocating resources to agents in the network under the context of competitive influence maximization. According to Romero Moreno et al. (2021b), knowing the opponent's strategies allows for an efficient budget allocation to maximize influence. However, even though it may be possible to directly observe agents' opinions at each time step, observing the strategies of controllers, i.e., if an agent is targeted by the external controller, or even how strong the intensity of influence from the controllers is, are often very challenging (Leskovec et al., 2009).

To solve this problem of opponent-strategy reconstruction from observable data, we model the updating process of agent i ($i = 1, ..., N$) as a non-homogeneous Markov chain (Brémaud, 2020) where the Markov property is retained but the transition probabilities $Pr(s_i(t+1) = 1)$ and $Pr(s_i(t+1) = 0)$ depend on time. Further to this formalization, we assume an active controller A that infers the strategy of the passive and

constant controller B who has fixed budget allocations (i.e., $J_{B,i}(t) = J_{B,i}(0), i = 1, ..., N$, $\forall t \geq 0$) from the time series of agents' opinion changes. Here, the time series are given by a matrix $\boldsymbol{S} = [s_i(t)]_{N\times T}$ where T is the length of the observation period. In other words, while updating the voting dynamics, we obtain a data matrix $\boldsymbol{S}$ with N rows and T columns in which each row of $\boldsymbol{S}$ denotes the binary opinion dynamics of an agent over an observation period of length T. Taking the data matrix $\boldsymbol{S}$ as an input, we are interested in decoding the unknown parameters $J_{B,i}(t)$ (referred to as $J_{B,i}$ in the following) from the input. Given the transition probabilities $Pr(s_i(t+1) = 1)$ and $Pr(s_i(t+1) = 0)$ of the opinion flow between agents in the existence of controllers, a commonly-used method for solving such parametric inference is maximum-likelihood estimation (MLE) (Gomez-Rodriguez et al., 2012). Specifically, replacing $s_i(t+1)$ and $s_j(t)$ with data actually observed along time series from 0 to T yields the log-likelihood function of agent i

$$L_i(T) = \sum_{t\in[0,T-1]} \Bigg[s_i(t+1) \log \frac{J_{A,i}(t) + \sum_{j=1}^{N} w_{ij} s_j(t)}{J_{A,i}(t) + J_{B,i} + k_i} + (1 - s_i(t+1)) \log \Big(\frac{J_{B,i} + \sum_{j=1}^{N} w_{ij}(1 - s_j(t))}{J_{A,i}(t) + J_{B,i} + k_i} \Big), \Bigg] \tag{4.3}$$

where k_i is the degree of node i, i.e., $k_i = \sum_{j=1}^{N} w_{ij}$. This log-likelihood function gives the likelihood of observing an agent's time series, given the parameter $J_{B,i}$. Depending on the opinion states in the next step $s_i(t+1)$, either $Pr(s_i(t+1) = 1)$ or $Pr(s_i(t+1) = 0)$ is taken into account in the log-likelihood function of Eq. (4.3). We then estimate the budget allocations of controller B to be the values $J_{B,i}$ that are most likely to generate the given data matrix $\boldsymbol{S}$ after T observations. According to the approach of MLE, we maximize the log-likelihood function $L_i(T)$ in Eq. (4.3) with respect to the budget allocations of controller B to obtain an estimate of $J_{B,i}$, denoted as $\hat{J}_{B,i}$ in the following.

According to the consistency of maximum likelihood estimates (Myung, 2003), for a sufficiently large dataset, the estimator asymptotically converges to the true value. However, we are interested in the problem of whether the observations of opinion states can be improved by interfering with the opinion dynamics so that we will obtain good-fit estimates with fewer observations. To achieve this, instead of passively observing, we assume that the controller A is an active controller who strategically allocates its resources to accelerate the inference of the strategy of its opponent (i.e., $J_{B,i}$, $1 \leq i \leq N$). To evaluate the goodness of fit of the inference obtained from MLE, a commonly-used measurement is the Fisher information (Ly et al., 2017). Specifically, Fisher information is used to test if the maximum likelihood estimators are aligned with the dataset and to derive a measure of dispersion between the true value and the estimator. Following the work of Ly et al. (2017), the Fisher information $I(J_{B,i}, T)$ about

$J_{B,i}$ is given by the expectation of second-order partial derivative of Eq. (4.3) with respect to $J_{B,i}$, which is given by

$$I(J_{B,i},T) = E[\frac{\partial^2}{\partial J_{B,i}^2} L_i(T)] = -\sum_{t\in[0,T-1]} \frac{J_{A,i}(t)+\sum_{j=1}^{N} w_{ij}s_j(t)}{(J_{A,i}(t)+J_{B,i}+k_i)^2(k_i+J_{B,i}-\sum_{j=1}^{N} w_{ij}s_j(t))}$$
$$= \sum_{t\in[0,T-1]} \left[(J_{A,i}(t)+J_{B,i}+k_i)^{-2} - \left((J_{A,i}(t)+J_{B,i}+k_i)(k_i+J_{B,i}-\sum_{j=1}^{N} w_{ij}s_j(t)) \right)^{-1} \right]. \tag{4.4}$$

For ease of exposition, let

$$\beta_i(t) = \left[(J_{A,i}(t)+k_i+J_{B,i})(k_i+J_{B,i}-\sum_{j=1}^{N} w_{ij}s_j(t)) \right]^{-1},$$

and

$$\Psi_i(t) = (J_{A,i}(t)+k_i+J_{B,i})^{-2}.$$

Given this, Eq. (4.4) can be written as

$$I(J_{B,i},T) = \sum_{t\in[0,T-1]} (\Psi_i(t)-\beta_i(t)).$$

Moreover, in Eq. (4.4) we have,

$$\frac{J_{A,i}(t)+\sum_{j=1}^{N} w_{ij}s_j(t)}{(J_{A,i}(t)+J_{B,i}+k_i)^2(k_i+J_{B,i}-\sum_{j=1}^{N} w_{ij}s_j(t))} \geq 0.$$

Correspondingly, the negative sum of the above equation over t from 0 to $T-1$ is non-positive, and will decrease as the length of observation T increases. Hence, the Fisher information $I(J_{B,i},T)$ is also non-positive and monotonically decreasing as T increases.

As mentioned above, knowledge of the Fisher information can be used to determine whether the maximum likelihood estimator is close to the true value. Specifically, for a large enough sample (i.e., $T\to\infty$), the maximum likelihood estimator $\hat{J}_{B,i}$ converges in distribution of a normal distribution to the true value $J_{B,i}$ (Myung, 2003), i.e.

$$(\hat{J}_{B,i}-J_{B,i}) \xrightarrow{D} \mathcal{N}(0,-I(J_{B,i},T)^{-1}), \quad \text{as} \quad T\to\infty \tag{4.5}$$

where $\mathcal{N}(0,-I(J_{B,i},T)^{-1})$ stands for a normal distribution with mean $\mu=0$ and variance $\sigma^2(J_{B,i},T) = -I(J_{B,i},T)^{-1}$ for agent i. As the Fisher information is non-positive and monotonically decreasing with the number of observations, the variance is always positive and, after a long period of observations, we will obtain more information and produce an estimator $\hat{J}_{B,i}$ closer to the true value $J_{B,i}$. Moreover, by taking the first order

partial derivative of $\sigma^2(J_{B,i}, T)$ with respect to $J_{B,i}$, one obtains

$$\begin{aligned}
&\frac{\partial \sigma^2(J_{B,i}, T)}{\partial J_{B,i}} = \frac{\partial\{-I(J_{B,i}, T)^{-1}\}}{\partial I(J_{B,i}, T)} \frac{\partial I(J_{B,i}, T)}{\partial J_{B,i}} \\
&= I(J_{B,i}, T)^{-2} \sum_{t\in[0,T-1]} \left[\frac{\left(J_{A,i}(t) + \sum_{j=1}^{N} w_{ij} s_j(t)\right)(J_{A,i}(t) + 3(J_{B,i} + k_i) - 2\sum_{j=1}^{N} w_{ij} s_j(t)\right)}{(J_{A,i}(t) + J_{B,i} + k_i)^3 (J_{B,i} + k_i - \sum_{j=1}^{N} w_{ij} s_j(t))^2} \right] \\
&\geq 0,
\end{aligned} \tag{4.6}$$

and we find that the variance is monotonically increasing with the increase of $J_{B,i}$ regardless of the values of $J_{A,i}$ and s_i. Note that, the variance in Eq. (4.5) is calculated from Fisher information at the true value. As the true value of $J_{B,i}$ is unknown, in practical calculations, we later replace the true value of $J_{B,i}$ with $\hat{J}_{B,i}$ to calculate the estimated variance $\hat{\sigma}^2(\hat{J}_{B,i}, T)$.

By introducing the Fisher information, we transform the problem of accelerating opponent strategy inference by interacting with the opinion dynamics into strategically deploying the budget of controller A to maximally decrease the variance of estimates. As the Fisher information can be represented in a recursive way, where the Fisher information at time t is calculated by Fisher information at time $t-1$ plus two additional terms, the variance can also be calculated recursively via

$$\hat{\sigma}^2(\hat{J}_{B,i}, t) = -I(\hat{J}_{B,i}, t)^{-1} = -\left[I(\hat{J}_{B,i}, t-1) + \hat{\Psi}_i(t-1) - \hat{\beta}_i(t-1)\right]^{-1}, \tag{4.7}$$

where

$$\hat{\beta}_i(t) = \left[(k_i + \hat{J}_{B,i} - \sum_{j=1}^{N} w_{ij} s_j(t))(J_{A,i}(t) + k_i + \hat{J}_{B,i}) \right]^{-1},$$

$$\hat{\Psi}_i(t) = (J_{A,i}(t) + k_i + \hat{J}_{B,i})^{-2},$$

and $\hat{\sigma}^2(\hat{J}_{B,i}, t)$ represents the expected variance at time t.

Inspired by the recursive expression for the variance in Eq. (4.7), we propose two types of heuristics in which we explore configurations of the budget allocations of controller A at time t for node i (i.e., $J_{A,i}(t)$, $i = 1, ..., N$) to maximally decrease the expected variance of the estimators in future updates. Because of the combinatorics involved when dealing with arbitrary numbers of updates, we limit considerations to looking one or two steps ahead and correspondingly label the resulting heuristics *one-step-ahead optimization* and *two-step-ahead optimization*.

Our strategy here is as follows. At time t, controller A has an estimate of the influence of controller B and an estimate of the variance around it. It then allocates its influence in such a way as to minimize the expected variance of its next estimate either one or two updating steps in the future.

In the following, we delve into how the budget allocation on a single node i is optimized in order to minimize the variance of the corresponding estimator, $\hat{J}_{B,i}$. We present the formal expressions for this optimization in both one-step-ahead and two-step-ahead scenarios. The extensions of these two heuristics will be further discussed in Section 4.3 in which we consider optimizing the budget allocations over multiple nodes to minimize the sum of variance for the entire network.

4.2.1 One-step-ahead Optimization

Specifically, for the one-step-ahead optimization scenario, the argument of the objective function through which we aim to minimize the one-step-ahead variance of estimator $\hat{J}_{B,i}$ is

$$\begin{aligned} J^*_{A,i}(t) &= \arg\min \hat{\sigma}^2(\hat{J}_{B,i}, t+1) = \arg\min -I(\hat{J}_{B,i}, t+1)^{-1} \\ &= \arg\min -\left[I(\hat{J}_{B,i}, t) + \hat{\Psi}_i(t) - \hat{\beta}_i(t)\right]^{-1} \end{aligned} \tag{4.8}$$

where $J^*_{A,i}(t)$ is the optimized budget allocation for controller A at time t returned by the argument function "arg" where the expected variance at time $t+1$ takes its minimum value. Analogous to Eq. (4.7), we have

$$\hat{\beta}_i(t) = \left[(k_i + \hat{J}_{B,i} - \sum_{j=1}^{N} w_{ij} s_j(t))(J_{A,i}(t) + k_i + \hat{J}_{B,i})\right]^{-1},$$

and

$$\hat{\Psi}_i(t) = (J_{A,i}(t) + k_i + \hat{J}_{B,i})^{-2}.$$

By doing so, instead of relying on passive inference that uses a fixed dataset, our approach step-wisely generates time-series data based on what was previously learned. In further refining our methodology, we define an experimental setup that focuses on obtaining a step-wise optimized budget allocation $J^*_{A,i}(t)$ for node i at each time step t, while maintaining other nodes' budget allocations as J^f_A. This optimized, adaptive approach to data collection guides us in designing specific experiments to clarify model uncertainties and achieve highly accurate estimators.

More specifically, the one-step-ahead optimization algorithm then proceeds according to the following steps:

(i) To satisfy the premise of enough samples before using the Fisher information to calculate the variance of a maximum likelihood estimator, we let controller A target all nodes equally with fixed budget allocation $J^f_A = b_A/N$ for the first T_0 updates and record the likelihood at time T_0 as $L_i(T_0)$ (The determination of T_0 is further discussed in Section 4.3.2.).

(ii) If the current updating step t is less than the length of total time series T, we calculate the current estimator $\hat{J}_{B,i}$ by maximizing the likelihood function $L_i(t)$ with respect to $J_{B,i}$ and evaluate the Fisher information $I(\hat{J}_{B,i}, t)$. Then, we calculate the expectation of the variance defined in Eq. (4.8). Next, we obtain the optimized $J^*_{A,i}(t)$ by applying the interior point optimization algorithm (Press et al., 2007). Finally, we update the network with a new assignment of $J^*_{A,i}(t)$ and simulate the stochastic voting dynamics to gain the next-step states for all nodes.

(iii) The procedure is terminated when a fixed number of observations T have been made.

This procedure is more formally presented in Algorithm 2. The main body of Algorithm 2 (lines 3-7) corresponds to step (ii). After applying Algorithm 2, we obtain a sequence of $J^*_{A,i}(t)$ where $T_0 \leq t \leq T$. Note that the initial states of agents are generated randomly to ensure that 50% of the initial opinions of agents are 0 or 1.

input : targeted node i, length of observations T, fixed budget allocation J^f_A, number of updates before calculating the variance T_0, initial states of agents $s_j(0)$ for $1 \leq j \leq N$
output: optimized budget allocation for node i, $J^*_{A,i}(t)$ for $T_0 \leq t \leq T$
1 *Initialization: update the network with fixed J^f_A for the first T_0 steps; let $t = T_0$;*
2 **while** $t \leq T$ **do**
3 $\quad \max_{J_{B,i}} L_i(t) \Rightarrow \hat{J}_{B,i}$;
4 $\quad$ calculate $I(\hat{J}_{B,i}, t)$;
5 $\quad J^*_{A,i}(t) = \arg\min \hat{\sigma}^2(\hat{J}_{B,i}, t+1)$;
6 $\quad$ update the network by simulating the stochastic voting dynamics ;
7 $\quad t = t + 1$;
8 **end**

Algorithm 2: One-step-ahead optimization

4.2.2 Two-step-ahead Optimization

For the two-step-ahead optimization scenario, we label the optimized budget allocations for node i at time t and $t+1$ as $J^*_{A,i}(t)$ and $J^*_{A,i}(t+1)$. Then, the objective function for minimizing the two-step-ahead variance is calculated by the expected negatively

inverse Fisher information two steps ahead given by:

$$
\begin{aligned}
&\{J^*_{A,i}(t), J^*_{A,i}(t+1)\} = \arg\min \hat{\sigma}^2(\hat{J}_{B,i}, t+2) = \arg\min -I(\hat{J}_{B,i}, t+2)^{-1} \\
&= \arg\min -\Big[I(\hat{J}_{B,i}, t) + E[s_i(t+1)(\hat{\Psi}_i(t) + \hat{\Psi}_i(t+1)) + \bar{s}_i(t+1)(\hat{\Upsilon}_i(t) + \hat{\Upsilon}_i(t+1))]\Big]^{-1} \\
&= \arg\min -\Big[I(\hat{J}_{B,i}, t) + Pr(s_i(t+1)=1)Pr(s_i(t+2)=1)(\hat{\Psi}_i(t) + \hat{\Psi}_i(t+1)) \\
&\qquad + Pr(s_i(t+1)=0)Pr(s_i(t+2)=0)(\hat{\Upsilon}_i(t) + \hat{\Upsilon}_i(t+1)) \\
&\qquad + Pr(s_i(t+1)=1)Pr(s_i(t+2)=0)(\hat{\Psi}_i(t) + \hat{\Upsilon}_i(t+1)) \\
&\qquad + Pr(s_i(t+1)=0)Pr(s_i(t+2)=1)(\hat{\Upsilon}_i(t) + \hat{\Psi}_i(t+1))\Big]^{-1}
\end{aligned}
\tag{4.9}
$$

where

$$
\bar{s}_i(t+1) = 1 - s_i(t+1),
$$

$$
\hat{\Upsilon}_i(t) = -\frac{(J_{A,i}(t) + \sum_{j=1}^N w_{ij}s_j(t))(J_{A,i}(t) + 2k_i - \sum_{j=1}^N w_{ij}s_j(t) + 2\hat{J}_{B,i})}{(J_{A,i}(t) + k_i + \hat{J}_{B,i})^2(k_i - \sum_{j=1}^N w_{ij}s_j(t) + \hat{J}_{B,i})^2},
$$

$$
\hat{\Psi}_i(t) = (J_{A,i}(t) + k_i + \hat{J}_{B,i})^{-2}.
$$

Note that the probabilities of agent i having opinion 1 or 0 at the current time step are dependent on its neighbouring states at the previous time step. As in the one-step-ahead procedure, when performing the optimization of Eq. (4.9), $s_i(t)$ for $1 \leq i \leq N$ are known. Therefore, the expressions for $Pr(s_i(t+1) = 1)$ and $Pr(s_i(t+1) = 0)$ only contain one unknown parameter, which is $J_{A,i}(t)$. However, in the expressions of $Pr(s_i(t+2) = 1)$ and $Pr(s_i(t+2) = 0)$, the sum of their respective neighbouring opinions $\sum_{j=1}^N w_{ij}s_j(t+1)$ is unknown, and thus the full expressions for $Pr(s_i(t+2) = 1)$ and $Pr(s_i(t+2) = 0)$ are obtained via applying the law of total probability

$$
\begin{aligned}
Pr(s_i(t+2)=1) &= \sum_{m=0,..,k_i} Pr(s_i(t+2)=1 \mid \sum_{j=1}^N w_{ij}s_j(t+1) = m) Pr(\sum_{j=1}^N w_{ij}s_j(t+1) = m) \\
&= \sum_{m=0,..,k_i} \frac{J_{A,i}(t+1) + m}{J_{A,i}(t+1) + J_{B,i} + k_i} Pr(\sum_{j=1}^N w_{ij}s_j(t+1) = m) \\
Pr(s_i(t+2)=0) &= 1 - Pr(s_i(t+2)=1),
\end{aligned}
\tag{4.10}
$$

where

$$
Pr(\sum_j w_{ij}s_j(t+1) = m) = \sum_{\rho=1}^{l} \prod_{j \in c_\rho} Pr(s_j(t+1)=1) \prod_{j \in (Nei(i) \setminus c_\rho)} Pr(s_j(t+1)=0). \tag{4.11}
$$

In the above, l stands for the number of combinations leading to $\sum_{j=1}^N w_{ij}s_j(t+1) = m$ and the elements of $C = \{c_1, .., c_l\}$, represented as c_ρ for $1 \leq \rho \leq l$, indicate all possible combinations of the neighbourhood of node i adding up to m at time $t+1$. If we denote the neighbourhood of node i as $Nei(i)$, then $Nei(i) \setminus c_\rho$ returns the set of

elements in $Nei(i)$ but not in c_ρ. Inserting Eqs. (4.10) and (4.11) into Eq. (4.9) yields the full expression for the goal function.

The optimization procedure for the two-step-ahead scenario then follows along the lines of Algorithm 2 except for updating every two steps in step (ii) using Eq. (4.9), as we optimize $J_{A,i}(t)$ and $J_{A,i}(t+1)$ in one loop. As shown in Eqs. (4.10) and (4.11), to calculate the probability that node i has state 1 at time $t+2$, we have to list all combinations of nodes leading to having a sum of neighbouring states from 0 to k_i. Therefore, the time complexity for calculating Eq. (4.11) is $\mathcal{O}(k_i!)$ and will grow factorially if we look more than two steps ahead. As it will become infeasible to calculate the combinatorics for more than two steps ahead for large networks, in this thesis, we only consider looking one or two steps ahead.

4.3 Results

In this section, our focus is on exploring the best strategies of controller A who aims to accelerate the opponent-strategy reconstruction process by optimally allocating its budget to minimize the variance of estimators of controller B's targeting. In order to gain some first intuition about how the budget allocations influence the inference of the opponent's strategy, we start our analysis by exploring the dependence of variance of MLE on different budget allocations in the equally targeting scenario in Section 4.3.1. These results also provide a benchmark for later comparison to our optimization heuristics. Next, to investigate the efficiency of the one-step-ahead and two-step-ahead optimization algorithms, we proceed with a numerical exploration of the performance of these two algorithms in Sections 4.3.2 and 4.3.3, respectively. In more detail, we start with using the one-step-ahead and two-step-ahead algorithms to infer the opponent's control at a single node, and then extend the above setting to optimizing multiple nodes with the aim to minimize the sum of variance. To further investigate the dependency of the optimal budget allocations for inference acceleration on network heterogeneity, we carry out detailed numerical experiments based on uncorrelated random scale-free networks with power-law degree distribution $p(k) \sim k^{-\lambda}$ constructed according to the configuration model (Catanzaro et al., 2005). Here, k is the node's degree, and λ indicates the degree exponent. After that, in Section 4.3.4, we propose an algorithm called optimally equally targeting, which has reduced time complexity compared with the two-step-ahead algorithm at the cost of very little performance loss.

4.3.1 Opponent Strategy Inference in the Equally Targeting Scenario

We start with exploring the influence of budget allocations on the variance calculated from Eq. (4.7) in the equally targeting scenario where all nodes are targeted with the

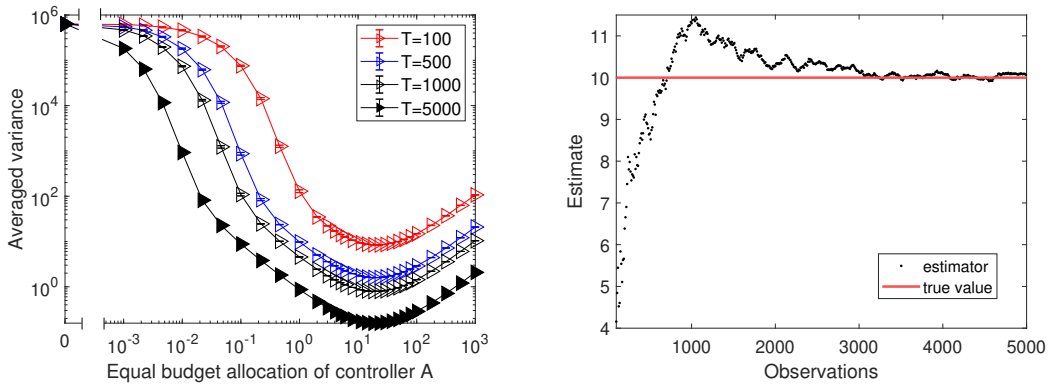

(a) (b)

FIGURE 4.1: Panel (a) shows the dependence of averaged variance of estimators for controller B's budget allocations over all agents on the budget allocations of controller A. Differently coloured curves correspond to different lengths of the observation periods T. The results are based on 100 repetitions of the experiment on random regular networks with $N = 1000$ nodes and average degree $\langle k \rangle = 10$, and we use a setting in which controller B targets all nodes equally with budget $J_{B,i} = 10$ for $1 \leq i \leq N$. Error bars indicate 95% confidence intervals. Panel (b) shows an example of one realization of the evolution of the estimator $\hat{J}_{B,1}$ over increasing numbers of observations. The true value of controller B's budget allocation is presented by the red line.

same budget allocation. To proceed, in Fig. 4.1 (a) we present numerical results for the dependence of the averaged variance over all agents in random regular networks on the varying budget allocations by the controller A for different observation periods T. In more detail, in panel (a) of Fig. 4.1, we observe a concave and asymmetric shape of the dependence of the averaged variance on the budget allocated by controller A, and clear minimum values of averaged variance can be identified for curves of different observation periods T. Moreover, the x-axis of Fig. 4.1 (a) starts from 0, which is identical to the scenario of no interference from controller A. In this scenario, agents will align with controller B after the first few updates and keep their opinions unchanged thereafter. As information is only gained in state flips, estimation in the scenario of no interference is almost impossible. Similarly, extremely small or large allocations (e.g., allocations less than 10^{-1} or bigger than 10^2) will cause difficulties in inferring the opponent's strategy as agents keep their opinions static in most updates. Further to the comparison of curves of different observation periods T in Fig. 4.1 (a), we find that, with the increase of the length of observation periods, the variance of the estimator decreases monotonically. In other words, a more accurate estimator will be obtained after a longer observation period, which is consistent with our analysis in Eqs. (4.4) and (4.5) where the variance will decrease monotonically with the increase of observational data. Additionally, we present the convergence of the maximum likelihood estimation for the opponent budget inference in Fig. 4.1 (b) by showing the dependence of the estimator of MLE on updates. With the increase of the number of observations, the estimator is approaching the true value.

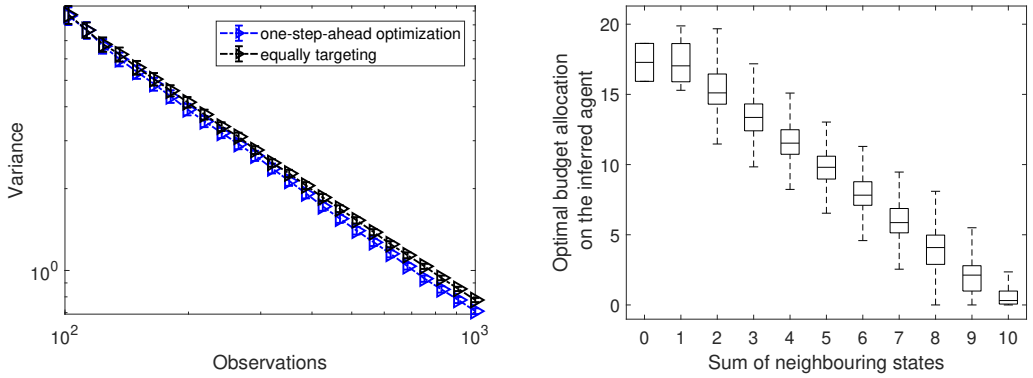

(a) (b)

FIGURE 4.2: Panel (a) compares the variance at a single inferred agent calculated by the one-step-ahead optimization with the variance calculated by the equally targeting strategy with increasing numbers of observations. Error bars in panel (a) present 95% confidence intervals. Panel (b) shows the dependence of the optimal budget allocations $f^*_{A,i}(t)$ over updates $t = 100$ to $t = 1000$ calculated by the one-step-ahead optimization in panel (a) on the sum of neighbouring states $\sum_{j=1}^{N} w_{ij} s_j(t)$, where i indicates the inferred agent i. Data in panel (b) is organized as box plots, where the central horizontal lines represent the median and the bottom and top box edges are for the $25th$ and $75th$ percentiles. The whiskers extend to the maximum or minimum data points. Results in both panels (a) and (b) are based on random regular networks with $N = 1000$ nodes and average degree $\langle k \rangle = 10$, and are averaged over 100 realizations of the experiment. Controller B targets all nodes equally with budget 10 and except for the inferred node, controller A targets all the other nodes with budget 20.

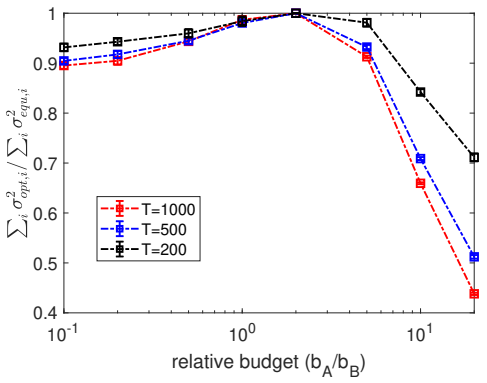

(a)

FIGURE 4.3: Relative sum of variance $\sum_i \sigma^2_{opt,i} / \sum_i \sigma^2_{equ,i}$ achievable by the one-step-ahead optimization compared to the equally targeting strategy for varying relative budgets b_A / b_B. Differently coloured curves correspond to different lengths of the observation periods T. Results are based on 100 repetitions of the experiment on scale-free networks with a power-law degree distribution of degree exponent $\lambda = 1.6$, network size $N = 1000$, average degree $\langle k \rangle = 10$ and are constructed according to the configuration model. We use a setting in which controller B targets all nodes with allocations randomly sampled from a uniform distribution, and the budget allocation by controller B per node on average is 10. Error bars indicate 95% confidence intervals.

4.3.2 Results for the One-step-ahead Optimization

To test the efficiency of the one-step-ahead optimization algorithm (see Algorithm 2), we start with exploring the optimal budget allocation for a single agent i according to Eq. (4.8) with the aim of minimizing the expected variance of the inferred agent step by step. In more detail, in Fig. 4.2 (a), we compare the variance of MLE calculated by the one-step-ahead optimization from Eq. (4.8) with the variance of the estimator obtained after applying the equally targeting strategy based on random regular networks. As indicated in Fig. 4.1 (a), for $J_{B,i} = 10$, the optimal minimum average variance points cluster around $J_{A,i} = 20$ across varying observation lengths T. Therefore, in Fig. 4.2 (a), we apply the same control parameters, i.e., $J_{B,i} = 10$ and $J_{A,i} = 20$ for $1 \leq i \leq N$, for the equally targeting scenario to assess the performance of the one-step-ahead optimization algorithm. In Fig. 4.2 (a), we find that, compared to the case of equally targeting, the one-step-ahead optimization algorithm achieves only a slight improvement in speeding up the convergence of the estimate (see the marginal difference in the dependence of variance on the number of observations in Fig. 4.2 (a)). Nevertheless, in order to shed light on the targeting strategy of controller A, in Fig. 4.2 (b), we further plot the dependence of the optimal budget allocations of controller A calculated by the one-step-ahead optimization averaged over updates $t = T_0$ to $t = T$ on the sum of neighbouring states $\sum_{j=1}^{N} w_{ij} s_j(t)$, where i represents the inferred agent. Note that, as depicted in Algorithm 2, T_0 represents the number of initialized updates before calculating the variance, and here we assign it as 100. As a result, we observe a clear pattern of the dependence of optimized budget allocations on the sum of neighbouring states: the larger the sum of neighbouring states, the lower the optimized budget allocation. In other words, to speed up estimates, controller A tends to target node i whenever all the node's neighbours differ from controller A.

In the following, we further generalize from the setting of attempting to infer the targeting of the B-controller at a single node to attempting to infer the targeting of the B-controller on all nodes. As a measure for the quality of estimates we use the sum of the variance of estimates at individual nodes and hence we aim at minimizing the sum of the variance of estimators for all agents. By extending Eq. (4.8), we have

$$\{\overbrace{J^*_{A,1}(t), ..., J^*_{A,N}(t)}^{N \text{ agents in the network}}\} = \arg\min \sum_{i=1}^{N} \hat{\sigma}^2(\hat{J}_{B,i}, t+1) = \arg\min \sum_{i=1}^{N} - \left[I(\hat{J}_{B,i}, t) + \hat{\Psi}_i(t) - \hat{\beta}_i(t)\right]^{-1}$$

subject to

$$J^*_{A,1}(t) + \cdots + J^*_{A,N}(t) \leq b_A, \tag{4.12}$$

where

$$\hat{\Psi}_i(t) = \left(J_{A,i}(t) + k_i + \hat{J}_{B,i}\right)^{-2},$$

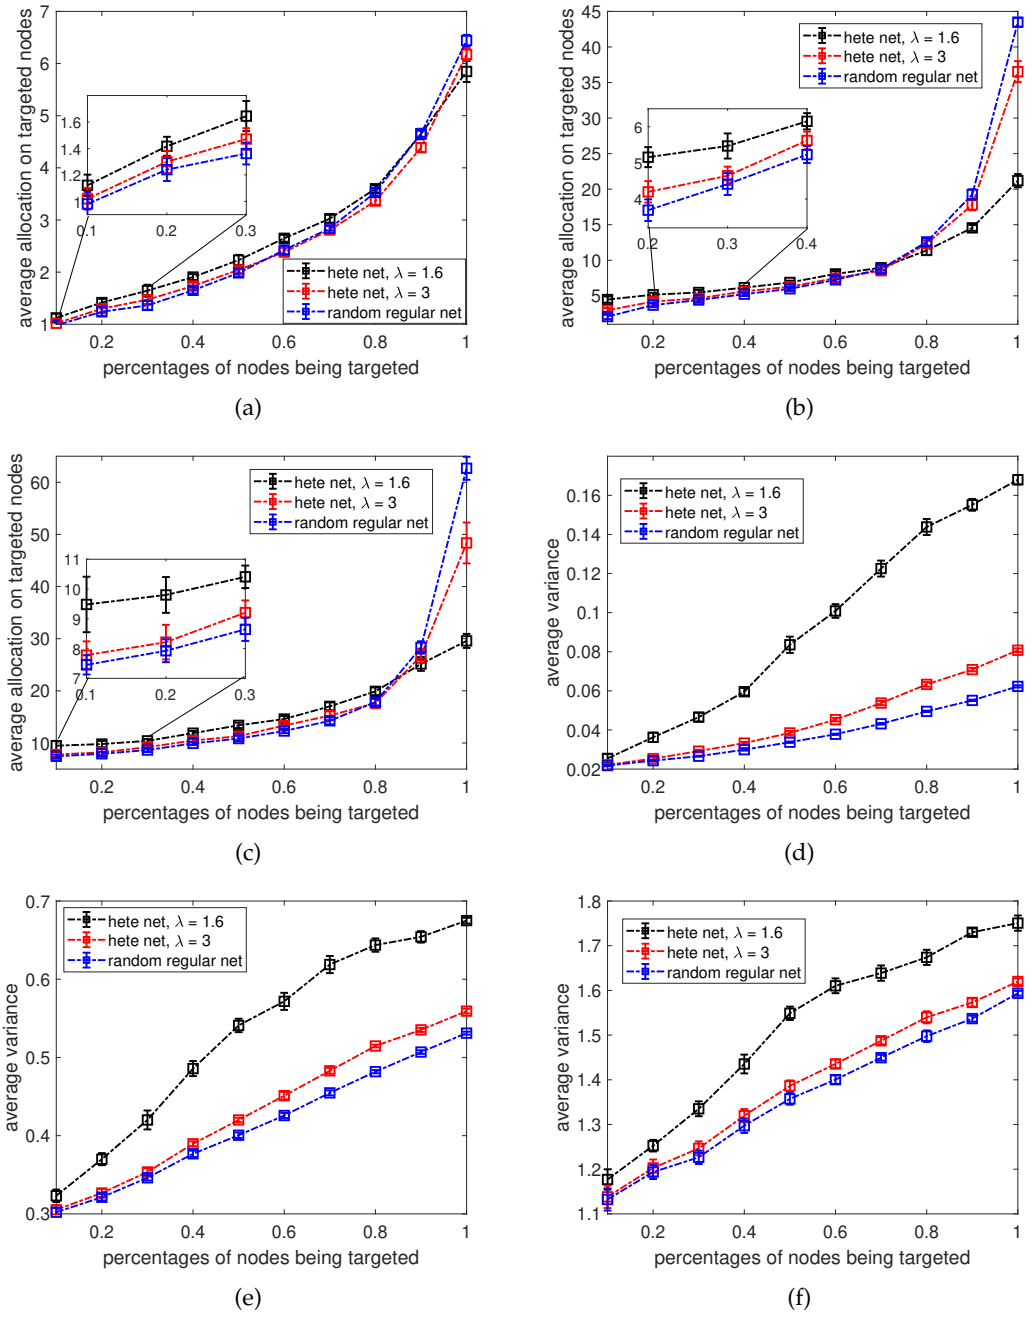

FIGURE 4.4: Panels (a-c) and (d-f) show the dependence of optimized average allocations for the one-step-ahead optimization and correspondingly normalized sum of variance of estimates on percentages of nodes being targeted. We use a setting in which controller B targets certain percentages of nodes with allocations randomly sampled from a uniform distribution, and the budget allocation per node on average by controller B is 1 (a,d), 5 (b,e), and 10 (c,f). The average allocations or variance are calculated by adding up the optimized allocations or variance for a certain percentage of targeted nodes and then divided by the number of nodes targeted. Black and red curves correspond to networks constructed according to the configuration model with a power-law degree distribution of exponent $\lambda = 1.6$ and $\lambda = 3$, respectively. Blue curves represent random regular networks. Results are based on 20 repetitions of the experiment on networks with size $N = 1000$, average degree $\langle k \rangle = 10$. Error bars indicate 95% confidence intervals.

$$\hat{\beta}_i(t) = \left[(k_i - \sum_{j=1}^{N} w_{ij} s_j(t) + \hat{J}_{B,i})(J_{A,i}(t) + k_i + \hat{J}_{B,i}) \right]^{-1},$$

and $J^*_{A,i}(t)$ stands for the optimized budget allocation for agent i by the controller A. Note that, the sum of the optimized budget allocations should be subject to the budget constraint at each time step, denoted as $J^*_{A,1}(t) + \cdots + J^*_{A,N}(t) \leq b_A$ in Eq. (4.12). Similar to Fig. 4.2 (a), in Fig. 4.3, we explore the one-step-ahead optimization algorithm in the scenario of minimizing the sum of variance of estimators in comparison with the equally targeting scenario for varying relative budgets b_A/b_B. In more detail, the improvement achieved by the one-step-ahead optimization is represented by the relative values of the ratios of the sum of variance $\sum_i \sigma^2_{opt,i} / \sum_i \sigma^2_{equ,i}$ at the end of iterations, where $\sum_i \sigma^2_{opt,i}$ denotes the sum of variance of estimators calculated by the one-step-ahead optimization and $\sum_i \sigma^2_{equ,i}$ stands for the sum of variance by the equally targeting strategy. After a careful inspection of Figure 4.3, we find that the one-step-ahead optimization can achieve a considerable improvement in reducing the variance compared with the equally targeting scheme if the active controller has much more budget than its opponent (i.e., $b_A \gg b_B$). In other settings, the one-step-ahead optimization only makes a slight improvement in minimizing the sum of variance, especially when the active controller has almost the same amount of available budget as its opponent. This indicates that the strategic allocation is more critical in accelerating the inference, if the active controller has more resources. In addition, we also find that with an increase in the length of the observation period, the relative values of sum of variance $\sum_i \sigma^2_{opt,i} / \sum_i \sigma^2_{equ,i}$ decreases. In other words, more benefits can accrue from the one-step-ahead optimization the longer the period of observation.

Above, in Fig. 4.3, we consider a scenario in which all agents in the network are subject to the control of the controller A, and the active controller wants to infer the budget allocations of its opponent over the entire network. However, in many real-world scenarios such as marketing, the controllers only focus on a subset of agents in the network, e.g. those who are most likely to buy their products. Inspired by this, we further assume that the controller A only distributes its budget among a certain fraction of agents targeted by controller B and tries to minimize the sum of variance among these agents. Additionally, we are also interested in the implications of network structure on the opponent strategy inference. Therefore, in Figs. 4.4, we show the dependence of average variance achievable by the one-step-ahead algorithm on the percentage of nodes being targeted by controller B.

The results in Fig. 4.4 are compared among regular random networks and scale-free networks with power-law degree exponents $\lambda = 1.6$ and $\lambda = 3$. Here, the average variance is calculated only within the targeted nodes, i.e., the average variance is equal to the sum of variance of the inferred agents divided by the number of nodes being targeted. More specifically, results for the dependence of averaged optimized budget allocation by the one-step-ahead optimization on varying percentages of nodes being

targeted are given in Figs. 4.4 (a) - (c), where panels correspond to controller B targeting nodes with allocations randomly sampled from a uniform distribution, and the budget allocation per node on average is 1, (Panel (a)), 5 (Panel (b)) and 10 (Panel (c)).

For the corresponding settings, the dependence for the optimized average variance are presented in Figs. 4.4 (d) - (f). From Figs. 4.4 (a) - (f), we obtain the following observations about the strategic allocations of the active controller. First, we see similar patterns in Figs. 4.4 (a) - (c) where with an increase in the number of agents being targeted, on average more resources are needed for the controller A to perform optimal inference. Depending on the budget availability of the controller B, the optimized controller A allocates more resources on each targeting node on average accordingly as b_B increases for the same amount of nodes being targeted, e.g., comparing the y-axis of the blue line in panel (a) to the blue line in panel (c). Meanwhile, in Figures 4.4 (d) - (f), as budget allocations from the opponent increase, the variance of the estimators rises. This is consistent with the analytical results in Eq. (4.6), which indicates that a higher value of budget allocation is harder to predict. Second, by comparing the curves of optimized budget allocations for different types of networks, we find that, only when a large portion of nodes are targeted then there is a significant difference in the optimized budget allocations among networks with different degree distributions. Otherwise, the optimized budget allocations are fairly close for networks with different degree heterogeneity. However, if we zoom in and compare the ordering of curves in Figs. 4.4 (a) - (c) for a small number of nodes being targeted with a large number of nodes being targeted, we find that there are two regimes for the strategy of the optimized controller depending on the network heterogeneity. In more detail, the optimized controller will allocate more resources on a more heterogeneous network than on a less heterogeneous network if only a small portion of nodes are targeted. The opposite holds if a large number of nodes are under control of the active controller. Third, in Figs. 4.4 (d) - (f), we find that more degree-heterogeneous networks always have higher average variance, i.e. opponent strategies are the more difficult to infer the more heterogeneous the network.

In Fig. 4.4, we always assume that the opponent targets nodes with allocations randomly sampled from a uniform distribution. However, we are also interested in the effects of the opponent's strategy on the predictability of the optimized controller. Therefore, in the following, we further explore the strategic allocations of the active controller based on different budget allocation strategies of its opponent. To proceed, we consider a scenario in which the opponent allocates resources as a function of the node's degree. More specifically, suppose the opponent generates random numbers r_i $(1 \leq i \leq N)$ from the interval of $[0, k_i^{\alpha}]$ for each of the N nodes, where k_i is the degree of node i and the exponent α indicates the varying strategies of the opponent. For instance, for $\alpha = 0$ an opponent would allocate independent of degree based on uniform random numbers, for $\alpha = 1$ the opponent would on average allocate proportional to degree,

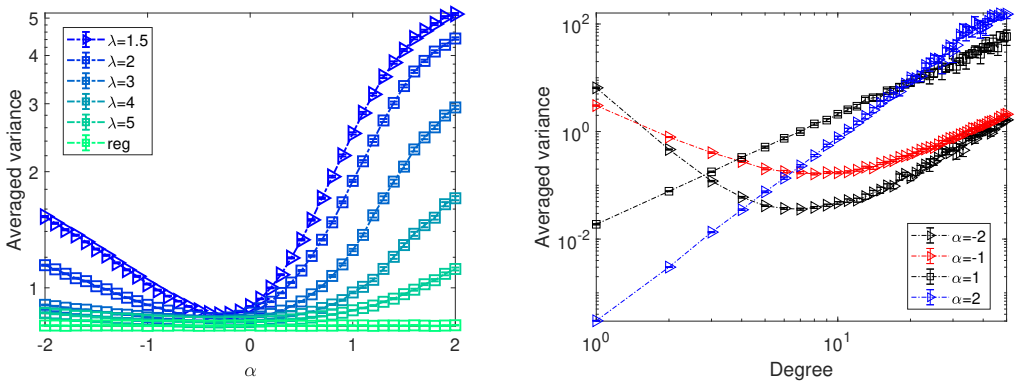

(a) (b)

FIGURE 4.5: Panel (a) shows the dependence of averaged variance obtained by the one-step-ahead optimization on the opponent strategy exponent α, where the budget allocation of the controller B is generated proportional to the random number within the interval $[0, k_i^\alpha]$. Different colours correspond to different degree exponents λ of the scale-free networks and "reg" corresponds to random regular graphs as indicated in the legend. Results are based on 100 repetitions of the experiment on networks with size $N = 1000$, average degree $\langle k \rangle = 10$. Error bars indicate 95% confidence intervals. For the setting corresponding to $\lambda = 1.5$ in panel (a), panel (b) shows the dependence of corresponding averaged variance of networks on the node's degrees for varying opponent strategy exponents α.

whereas for $\alpha = -1$ average allocations would be inversely proportional to degree. By then normalizing the random numbers r_i to satisfy the budget constraint of controller B, we obtain different budget allocations $J_{B,i}$.

In Fig. 4.5, we plot the dependence of optimized averaged variance obtained by the one-step-ahead optimization on the opponent strategies represented by the varying exponents α. We observe a concave shape of the averaged variance along with the changing exponents α, with minima near $\alpha = 0$. For all settings of α we generally also observe larger average variance the more heterogeneous the networks. To proceed, Fig. 4.5 (b) shows the dependence of averaged variance on nodes' degree. We find that, generally, nodes with a larger degree are more difficult to predict as the averaged variance of estimators is larger. In a similar vein, nodes being allocated larger budgets by the opposing controller are also harder to predict which can be seen from the curves for $\alpha = -1, -2$, as in this setting low degree nodes have larger averaged variance than the high degree nodes.

4.3.3 Results for the Two-step-ahead Optimization

In this section, we proceed with testing the efficiency of the two-step-ahead optimization algorithm. Similar to Section 4.3.2, we start by minimizing the variance of a single agent using the two-step-ahead optimization over random regular networks with network size $N = 1000$ and average degree $\langle k \rangle = 10$. In more detail, in Fig. 4.6, we

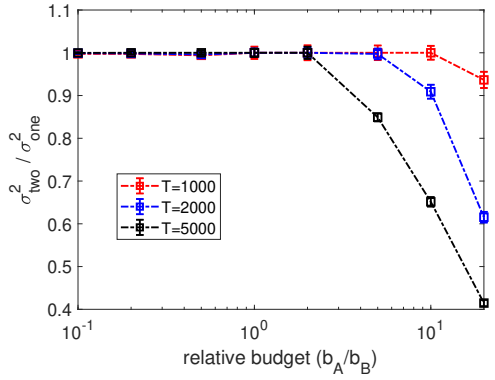

(a)

FIGURE 4.6: Relative variance $\sigma^2_{two}/\sigma^2_{one}$ of the estimate for a single inferred node achievable by the two-step-ahead optimization compared to the one-step-ahead scheme for varying relative budgets b_A/b_B. Differently coloured curves correspond to different lengths of the observation periods T. Results are based on 100 repetitions of the experiment on random regular networks with size $N = 1000$, average degree $\langle k \rangle = 10$. Controller B targets all nodes equally with allocation 10 and except the inferred node, controller A targets all the other nodes with budget 10. Error bars indicate 95% confidence intervals.

compare the variance of the estimator on a single node i calculated by two-step-ahead heuristics with the one-step-ahead optimization for varying relative budget constraints b_A/b_B based on different observation periods T. Note that, for the two-step-ahead optimization, we set the constraint b_A separately for time steps t and $t+1$. Therefore, assigning the optimized allocation obtained by the two-step-ahead algorithm at time steps t and $t+1$ as $J^*_{A,i}(t)$ and $J^*_{A,i}(t+1)$, we have $J^*_{A,i}(t) \leq b_A$ and $J^*_{A,i}(t+1) \leq b_A$. By observing the results in Fig. 4.6, we find that, similar to the results of Fig. 4.3, the two-step-ahead optimization can achieve a considerable improvement in reducing the variance compared with the one-step-ahead scheme only if the active controller has much more budget than its opponent.

Notice that, in the log-likelihood function of Eq. (4.3) composed of transition probabilities $Pr(s_i(t+1) = 1) = \frac{J_{A,i}(t)+\sum_j w_{ji}s_j(t)}{J_{A,i}(t)+J_{B,i}+k_i}$ and $Pr(s_i(t+1) = 0) = \frac{J_{B,i}+\sum_j w_{ji}(1-s_j(t))}{J_{A,i}(t)+J_{B,i}+k_i}$, the budget allocation from the controller A (i.e., $J_{A,i}(t)$) is not the only determinant that influences the inference of $J_{B,i}$. Instead, the sum of the neighbouring states $\sum_j w_{ji}s_j(t)$ of the inferred node i is also taken into consideration when inferring $J_{B,i}$. Therefore, a natural extension for the above scenario of minimizing the variance of a single node by only targeting that inferred node is to optimize the inference at the focus node by targeting it and its neighbours. For clarification, a schematic illustration of optimizing the budget allocations for the inferred node and its neighbourhood to minimize the variance of the central node is given in Fig. 4.7. In more detail, in Fig. 4.7, we have shown that to minimize the variance of the estimator $\hat{J}_{B,i}$ for node i at time step $t+2$ (marked as output), we have to optimize the budget allocation for the inferred node one step ahead and its neighbours two steps ahead (circled in red). A reason for optimizing the

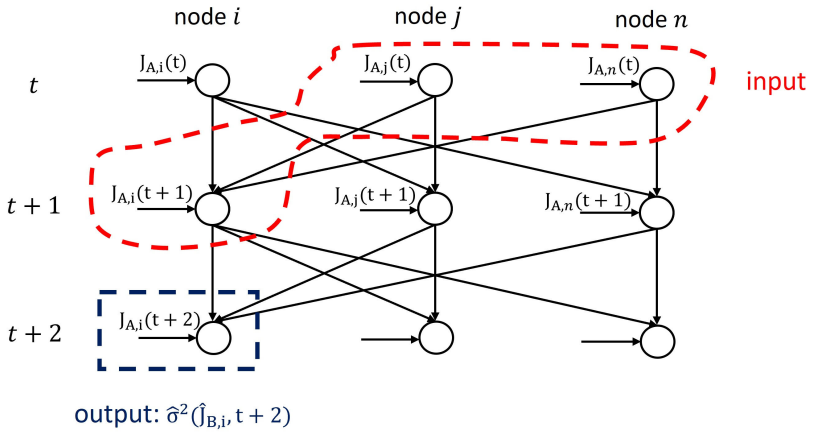

(a)

FIGURE 4.7: Schematic illustration of a variant of the two-step-ahead optimization in the context of optimizing the budget allocations for an inferred node i and its neighbourhood. Here, we assume that node j and node n are the two neighbours of node i. Each column presents the state dynamics of a node from time t to $t+2$. The arrows indicate interactions between nodes and controller A which determine the transition probabilities. For example, the state of node i at time $t+1$ depends on the states of node j and node n at time t, as well as the budget allocation $J_{A,i}(t)$. Therefore, there are arrows from node j and node n at time t point to node i at time $t+1$, as well as a horizontal arrow indicating the budget allocation from controller A at time t labelled by $J_{A,i}(t)$. The state of node i at time $t+2$ is determined by $J_{A,i}(t+1)$ and states of node j and n at time $t+1$. To influence states of node j and n at time $t+1$, we have to change the budget allocations at time t. Therefore, the inputs of the optimization of minimizing the variance of node i by optimizing the budget allocations for the inferred node and its neighbourhood are $J_{A,i}(t+1)$, $J_{A,j}(t)$ and $J_{A,n}(t)$ (see the variables circled by the red dashed line).

budget allocation of the neighbouring nodes two step ahead is that by doing so, we could influence the sum of neighbouring states at time $t+1$, and afterwards the variance of the inferred node at time $t+2$. Therefore, the optimization in this scenario can be viewed as a variant of the two-step-ahead optimization of Eq. (4.9), and the objective function is given by

$$\begin{aligned}
&\{J^*_{A,i}(t+1), \overbrace{J^*_{A,j}(t), ..., J^*_{A,n}(t)}^{\text{neighbours of node } i}\} = \arg\min \hat{\sigma}^2(\hat{J}_{B,i}, t+2) \\
&= \arg\min -\Big[Pr(s_i(t+2)=1) \times \hat{\Psi}_i(t+1) + Pr(s_i(t+2)=0) \times \hat{\Upsilon}_i(t+1)) + I(\hat{J}_{B,i}, t+1)\Big]^{-1}
\end{aligned}$$

subject to

$$J^*_{A,i}(t+1) + J^*_{A,j}(t) + \cdots + J^*_{A,n}(t) \leq b_A \tag{4.13}$$

where $Pr(s_i(t+2)=1)$ and $Pr(s_i(t+2)=0)$ represent the probabilities for node i to have opinion 1 and 0 at step $t+2$, respectively. Moreover, $\hat{\Psi}_i(t+1)$ and $\hat{\Upsilon}_i(t+1)$ are consistent with the definition in Eq. (4.9). Inserting Eq. (4.10) and Eq. (4.11) into Eq. (4.13) yields the full expression. Here, we use the interior-point method for the optimization of Eq. (4.13), and the corresponding time complexity to obtain $J^*_{A,i}(t+$

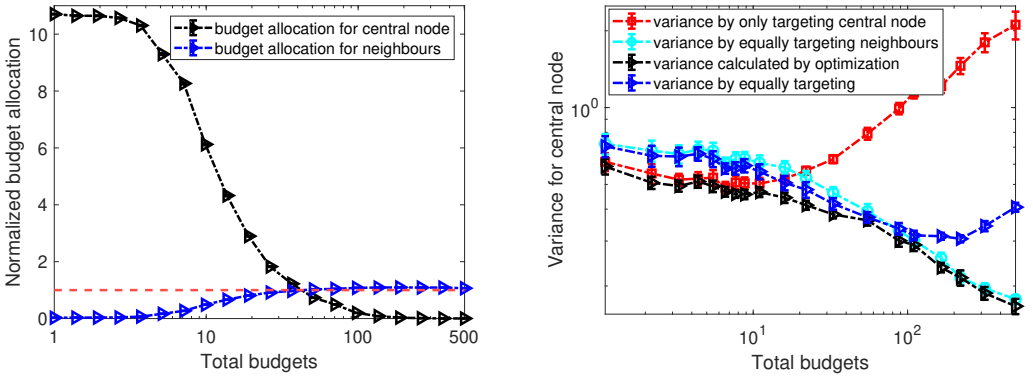

(a) (b)

FIGURE 4.8: Panel (a) shows the dependence of normalized budget allocations $\tilde{J}_{A,j} = \frac{J_{A,j}(k_i+1)}{b_A}$ after the first $T = 1000$ updates calculated by Eq. (4.13). The black triangles are the budget allocations for each neighbouring node where differences in allocations to different neighbours are characterized by error bars. Panel (b) shows the dependence of variance of MLE of the central node on varying total budgets at update $T = 1000$ based on four budget allocation strategies: only targeting the central node (red squares), equally targeting neighbours only (red circles), optimization described in Eq. (4.13) (black triangles), and equally targeting (blue triangles). The results are based on 20 realizations of random regular networks with 1000 nodes and average degree $\langle k \rangle = 10$. Controller B targets all nodes equally with budget 5, and except for the inferred node and its neighbours, controller A targets all the other nodes with budget 5. Error bars indicate 95% confidence intervals.

1), $J^*_{A,j}(t), \cdots, J^*_{A,n}(t)$ is $\mathcal{O}(k_i!T)$, where k_i is the degree of node i and T is the length of the observation period.

To distinguish differences in allocations made by the optimized controller to the central inferred node and on its neighbours, we partition the budget allocations for these two types of nodes into two groups and normalize by the average budget allocation to any node. We thus have $\tilde{J}_{A,i} = \frac{J_{A,i}(k_i+1)}{b_A}$ for the central node and for the neighbouring nodes $j \in Nei(i)$ we have $\tilde{J}_{A,j} = \frac{J_{A,j}(k_i+1)}{b_A}$. In Fig. 4.8 (a), we show the dependence of the normalized optimized budget allocations to the central node and its neighbours on varying budget availability b_A of controller A. We clearly observe two regimes of budget allocations for the central node and its neighbours. For small enough budget availability to controller A, all of the resources will be focused on only the central node. However, with an increase of the budgets available to the optimized controller, more and more resources will be diverted to its neighbours until a crossing point is reached. Finally, for large enough budget b_A, only the neighbouring nodes will be targeted.

Motivated by the optimized schemes of budget allocations for the central and neighbouring nodes in the context of extremely small and large budget constraints in Fig. 4.8, we propose two other heuristics. One is allocating all of the resources on the central node and leaving its neighbours un-targeted. The other is targeting the neighbouring

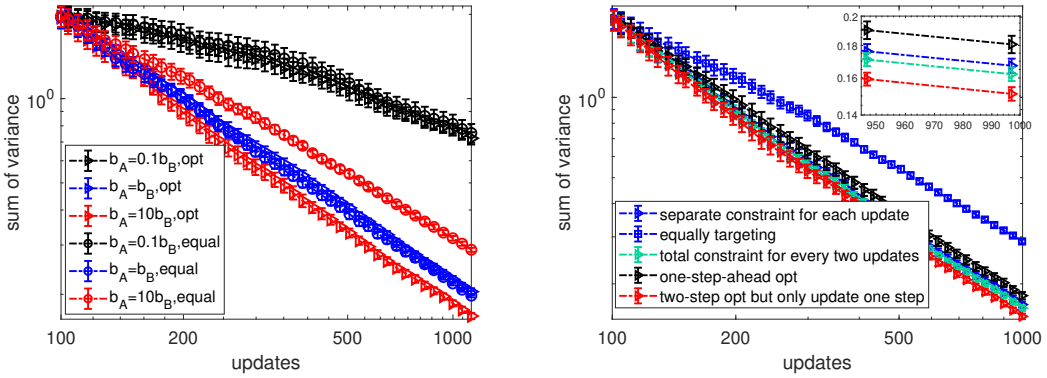

(a) (b)

FIGURE 4.9: Panel (a) compares the sum of expected variance of MLE calculated by the two-step-ahead optimization with the total constraint for every two updates (marked by "opt") and equally targeting strategy (marked by "equal") based on three different relative budget constraints $b_A/b_B = 0.1$, $b_A/b_B = 1$ and $b_A/b_B = 10$ with increasing numbers of observations. Panel (b) compares the sum of the expected variance of MLE calculated by 5 different methods under the control setting of $b_A = 10b_B$. Here, "two-step opt but only update one step" stands for using the two-step-ahead algorithm but only updating the budget allocation for the next step. "one-step-ahead opt" presents the one-step-ahead algorithm. Results are based on 20 repetitions of the experiment on ring networks with size $N = 10$. Controller B targets nodes with allocations randomly sampled from a uniform distribution, and the budget allocation per node on average by controller B is 2. Error bars indicate 95% confidence intervals.

nodes equally, but leaving the central node empty. In Fig. 4.8 (b), we compare the variance for the central node calculated by only targeting the central node (represented by red squares), by only equally targeting the neighbours (light blue circles), and by the optimization of Eq. (4.13) (marked as black triangles), and also the strategy of equally targeting all node (including the central node and its neighbours) as the benchmark. The results in Fig. 4.8 (b) are consistent with what we observe in Fig. 4.8 (a). Although the optimized strategy has the best performance in reducing the variance of the central node for all values of budget constraints compared with the three other strategies in Fig. 4.8 (b), for small total budgets, the variance calculated by the strategy of only targeting the central node is close to the variance by the optimized strategy. Meanwhile, for large total budgets, the strategy of only equally targeting the neighbours has almost the same performance as the optimized strategy. Our finding suggests that, instead of applying the optimization of Eq. (4.13) whose time complexity is $\mathcal{O}(k_i!T)$, we could substitute it with simple heuristics of targeting the central node or neighbours only without sacrificing much in performance.

In the following, we extend the scenario of minimizing the variance of a single node

to minimizing the sum of variance of estimators over the whole network with the two-step-ahead heuristics. In this context, we have

$$\begin{aligned}\begin{bmatrix} J^*_{A,1}(t) & \cdots & J^*_{A,N}(t) \\ J^*_{A,1}(t+1) & \cdots & J^*_{A,N}(t+1) \end{bmatrix} &= \arg\min \sum_{i=1}^{N} \hat{\sigma}^2(\hat{J}_{B,i}, t+2) \\ &= \arg\min \sum_{i=1}^{N} -I(\hat{J}_{B,i}, t+2)^{-1}.\end{aligned} \tag{4.14}$$

Generally, there are two options to set the constraint for the two-step-ahead optimization in the context of minimizing the sum of variance for the entire network. One is to set the budget constraint separately for each update, i.e., $J_{A,1}(t) + \cdots + J_{A,N}(t) \leq b_A$, referred to as *SCEU* in the following. The other is to set a total constraint for every two updates where $J_{A,1}(t) + \cdots + J_{A,N}(t) + J_{A,1}(t+1) + \cdots + J_{A,N}(t+1) \leq 2b_A$, referred to as *TCEW*. Due to the factorial time complexity of the two-step-ahead algorithm with respect to the network size N, we start our analysis of the optimal budget allocations for minimizing the sum of variance by numerical experiments conducted on a small ring graph of $N = 10$ nodes.

In more detail, Fig. 4.9 (a) compares the sum of variance calculated by the two-step-ahead method (marked as "opt") and the equally targeting strategy (marked as "equal") for varying budget constraints $b_A/b_B = 0.1, 1, 10$ with increasing numbers of observations. Moreover, we use a setting in which controller B targets all nodes with allocations randomly sampled from a uniform distribution, and the budget allocation per node on average is 2. After a careful inspection of Fig. 4.9 (a), we obtain a similar observation as in the case of optimizing one node by the two-step-ahead optimization in Fig. 4.6. The two-step-ahead optimization can make a significant improvement in reducing the sum of variance of estimators compared with the equally targeting strategy only when the active controller has much more budget than its opponent. This suggests the two-step-ahead optimization is more effective if the optimized controller has more available resources.

Since the substantial improvement in minimizing the sum of variance by the two-step-ahead optimization has only been observed when controller A is in budget superiority, we only investigate this scenario further. In the following, we proceed by comparing the sum of variance obtained by applying the one-step-ahead optimization, the two-step-ahead optimization and the equally targeting strategy in Fig. 4.9 (b) under the same network setting as Fig. 4.9 (a) and control setting of $b_A/b_B = 10$. From Fig. 4.9 (b), we find that the "two-step opt but only update one step" heuristics which uses a two-step-ahead algorithm but only updates the budget allocation for the next step has the best performance in the scenario of controller A having more budget than controller B. An explanation for this is: by optimizing two steps ahead, this heuristics accounts for the indirect influence between nodes, while, by only updating one step ahead, the controller can adjust its prediction of two-step-ahead variance after one update, as well

as adjust its budget allocation for the next step. Additionally, even though SCEU and TCEW are separated by different ways of imposing the budget constraint, there is no significant difference in the sum of variance obtained by these two methods. Moreover, as expected we also observe that the one-step-ahead method has the worst performance among all the heuristic methods, but is nevertheless still better than the equally targeting strategy.

4.3.4 Optimally Equally Targeting

As seen in Figs. 4.3 and 4.9, the one-step-ahead and two-step-ahead optimization algorithms have better performance in reducing the variance of estimators of the opponent's budget allocations compared with the equally targeting strategy. However, the cubic and factorial time complexity of the one-step-ahead and two-step-ahead optimization algorithms in terms of network size N make them unsuitable for application to large-size networks. To address this issue of scalability of the one-step-ahead and two-step-ahead optimizations, we propose a new heuristic algorithm named optimally-equally-targeting strategy (*OETS*), where we attempt to find an optimal equal allocation for all nodes in the network. Specifically, the heuristics of the optimally-equally-targeting strategy are motivated by the observations in Fig. 4.8 (b) which shows that putting too many resources on the inferred agent only will deteriorate the accuracy of the inference. Moreover, Figs. 4.3 and 4.9 also indicate that only limited improvement of variance reduction will be achieved by the one-step-ahead and two-step-ahead algorithms compared with the equally targeting strategy when the active controller has less or equal budgets compared to its opponent.

Formally, the objective function of the OETS is given by

$$J_A^* = \arg\min \sum_{i=1}^{N} \hat{\sigma}^2(\hat{J}_{B,i}, T) = \arg\min - \sum_{i=1}^{N} \left[\sum_{t=0}^{T-1} \left(\hat{\Psi}_i(t) - \hat{\beta}_i(t) \right) \right]^{-1}, \tag{4.15}$$

subject to

$$0 \leq J_A^* \leq b_A/N,$$

where J_A^* is the optimal budget allocation for all nodes to achieve a minimum sum of variance after T observations, b_A is the budget constraint for controller A, $\hat{\Psi}_i(t) = (J_A + k_i + \hat{J}_{B,i})^{-2}$, $\hat{\beta}_i(t) = \left[(k_i - \sum_{j=1}^{N} w_{ij} s_j(t) + \hat{J}_{B,i})(J_A + k_i + \hat{J}_{B,i}) \right]^{-1}$. Here, by proposing the optimally-equally-targeting strategy, we have reduced the parameter space from N (the one-step-ahead optimization) or $2N$ (the two-step-ahead optimization) to 1 and the time complexity to $\mathcal{O}(T)$ without sacrificing much of the performance.

To explore how the budget availability and network structures affect the optimally equally targeting strategy, in Fig. 4.10 (a), we show the corresponding sum of variance of MLE for varying equal budget allocations by the active controller on networks

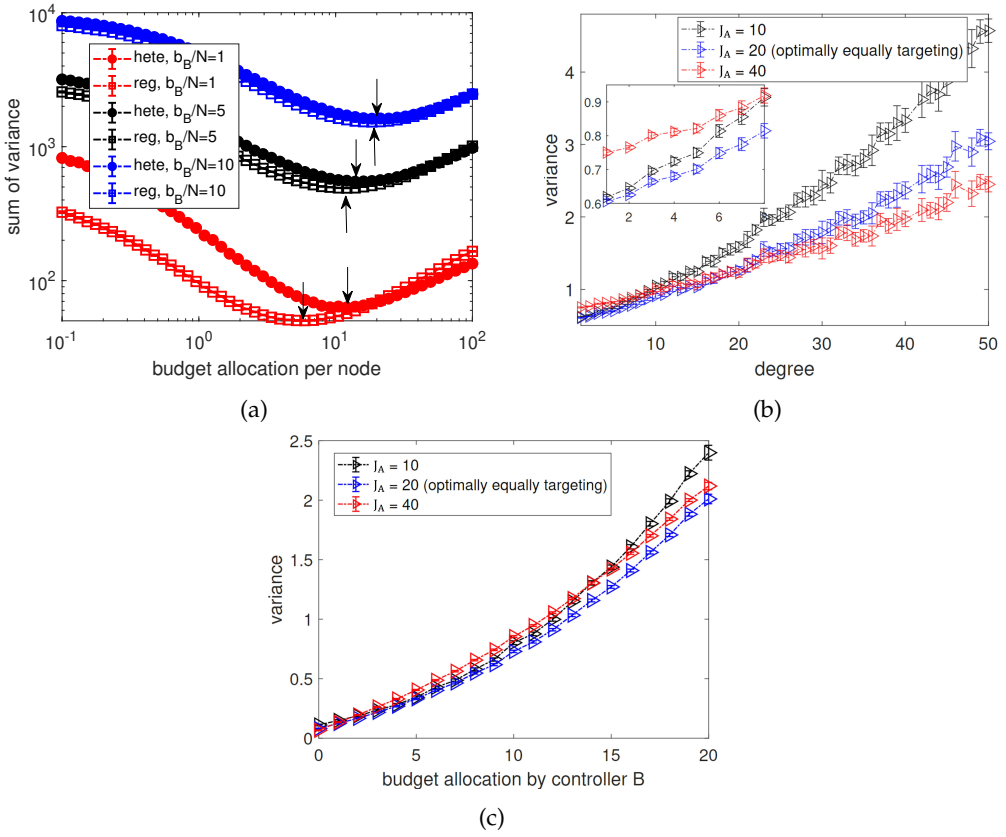

FIGURE 4.10: Panel **(a)** shows the dependence of sum of variance of estimators for controller B's budget allocations over all agents on the equal budget allocations of controller A at update $T = 1000$. Differently coloured curves correspond to varying budget constraints of controller B, e.g., the red lines marked with $b_B/N = 1$ indicate controller B targets each node with 1 on average. Circles and squares correspond to scale-free networks with degree exponent $\lambda = 1.6$ and random regular networks, respectively. The results are based on 20 repetitions of the experiment on networks with $N = 1000$ nodes and average degree $\langle k \rangle = 6$. Controller B targets nodes with allocations randomly sampled from a uniform distribution. Controller A targets all nodes equally. Panels **(b)** and **(c)** present the dependence of the corresponding variance of estimators achieved by equally targeting each node with allocations 10, 20, and 40 in Panel (a) on nodes' degrees (b) and budget allocations (c) by the opponent for the scale-free networks with degree exponent $\lambda = 1.6$ under the context of $b_B/N = 10$. Note that $J_A = 20$ is the optimal budget allocation for equally targeting obtained from Panel (a) (the minimum point) for $b_B/N = 10$. In Panel (c), we group the value of the x-axis into bins with width 1 and lower limits are inclusive, e.g., [0,1). Error bars indicate 95% confidence intervals.

with different heterogeneity in the context of the opponent targeting each node with average budget 1, 5 and 10. Additionally, controller B targets nodes with allocations randomly sampled from a uniform distribution. Note that, the optimally equally targeting strategy for each scenario in Fig. 4.10 (a) is marked by the arrows. By comparing curves for networks with different degree heterogeneity in Fig. 4.10 (a), we find that, similar to the results of Figs. 4.4, the variance of estimators for random regular networks is always smaller than that for the heterogeneous networks.

As the main difference of networks of different types is the degree distribution, to explore how the degree of nodes plays a role in OETS, we present the dependence of variance of estimators on nodes' degrees in Fig. 4.10 (b). Clearly, we observe a positive relationship between the variance and nodes' degree. This result has further explained the degree-based heuristics for the link weight prediction in Chen and Lai (2018) about why the solution obtained from a lower-degree node is preferred. Moreover, with a careful inspection of Fig. 4.10 (b), we observe two regimes. For low-degree nodes, a large equal allocation by controller A (e.g., $J_A = 40$) will result in a worse performance in predicting the budget allocations. However, for the hub nodes, a larger allocation is preferable in improving the accuracy of the prediction. Furthermore, by comparing the patterns of the dependence for equal budget allocations by controller A for $J_A = 10, 20, 40$ in Fig. 4.10 (b), we find that the OETS results from a trade-off. On the one hand, heterogeneous networks have more low-degree nodes, therefore relatively high budget allocations from the controller A should be avoided. On the other hand, as the hub nodes normally have much higher variance than low-degree nodes, low budget allocations from the controller A is inefficient in minimizing the sum of variance.

Another important factor in the strategy inference is the budget allocation by the opponent. Therefore, in Fig. 4.10 (c), we present the dependence of variance on opponent's budget allocations. Note that, as the budget allocations by the opponent are randomly sampled from a uniform distribution, for ease of observation, we group values into bins with width 1, i.e., $\{[0,1), [1,2), \cdots\}$ in Fig. 4.10 (c). Similar to Fig. 4.10 (b), with an increase of opponent's budget allocations, the variance of estimates rises monotonically. However, curves for different budget allocations J_A are fairly close and a larger J_A will not result in a lower variance for nodes which are allocated more resources by the opponent.

4.4 Summary

In this chapter, we have proposed an approach to apply network control in the context of a network inference problem. In our setting, an active controller interacts with a process of opinion dynamics on a network and aims to influence the resulting opinion dynamics in such a way that estimates of an opposing controller's strategy can be accelerated. Existing approaches related to such types of inference problems are often based on the assumption that the inference is performed using given data. In contrast, our approach aims to strategically interfere with the networked dynamics to generate more informative datasets.

By using the variance deduced from the Fisher information as a criterion of inference uncertainty, we have proposed several optimization heuristics. In the first step, in a

benchmark scenario in which an active controller can target nodes uniformly by an adjustable amount of influence, we have demonstrated that interference with the system's dynamics can substantially accelerate the convergence of estimates about opponents. We have then proceeded to develop more sophisticated optimization heuristics, based on step-wise updating of the interference with the dynamics and have shown that such approaches are typically effective if the active controller has a relatively large budget.

Next we have explored the one-step-ahead and two-step-ahead heuristics systematically in a variety of scenarios. First, in a scenario in which the active controller only aims at inference of a single node, we find that only very limited acceleration can be achieved by targeting only this node. However, far more substantial results can be achieved by also targeting the node's neighbours. For the latter setting we have demonstrated the effectiveness of simple heuristics, which relies on targeting only the focal node when the controller's budget is small and only conditionally influencing the focal node's neighbours when budget availability is large. Conditional targeting of neighbours should be carried out whenever a majority of them are not aligned with the active controller.

Furthermore, we have explored the effectiveness of inference acceleration for networks with varying amounts of degree heterogeneity for different settings of the opponent's influence allocations. As one might expect, we find that both, predicting opponent influence at nodes with large degrees, and precisely predicting large opponent influence nodes, are difficult. The first is essentially due to the presence of a large changing environment of the node which makes it difficult to distinguish the influence of control from the influence of neighbours. This finding is consistent with results presented in Chen and Lai (2018) in the context of link inference from static data. The second is due to the effect that large opponent control tends to fix a node in a static state, which makes it difficult to precisely predict the amount of opponent's influence.

As a consequence of the above, if an opponent targets uniformly at random, the inferrability of its influence is strongly related to the number of high-degree nodes on a network. Correspondingly, using our optimization schemes, we find that inference is the more difficult the larger the degree-heterogeneity of a network. The above finding also holds when opponent's influence strengths are drawn randomly with inverse proportionality to node degrees. In this case, networks with higher degree heterogeneity will also have larger average variance, since they have more low-degree nodes with large opponent influence, which also impedes inference.

Chapter 5

Opponent-strategy Reconstruction in the Ising Model

In this chapter, we extend the previous framework of inference acceleration in Chapter 4, which was initially applied to the voter model, to the inverse Ising model. The purpose is to evaluate the applicability and adaptability of this framework as we move from a linear model to a relatively intricate non-linear model. Similar to Chapter 4, instead of passively observing the dynamics for inference, we focus on strategically manipulating dynamics to generate data that give more accurate estimators with 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.

The structure of this chapter is organized as follows. In Section 5.1, we introduce the context and motivation for studying the inference acceleration problem in the Ising model. In Section 5.2, 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 5.3, we provide analytical and numerical results for the optimal configuration of budget allocations by the active controller. In Section 5.4, we summarise the main findings and contributions.

5.1 Introduction

As discussed in Section 2.5.3, even though much effort has been devoted to the inverse

Ising problem, it remain challenging, especially when only a limited amount of observational data is available (Yang et al., 2017; Braunstein et al., 2019). Therefore, in this chapter, we focus on addressing the inverse Ising 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. (2017); Braunstein et al. (2019); Hoang et al. (2019), which assume the existence of a given dataset, and focus on methods for inference. Instead, in this chapter, we are interested in how data can best be generated. Similar to Chapter 4, our principal idea is that we allow strategic interference in the dynamics while observations are being gathered, and 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 in the inverse Ising model.

Different from the linearity of the voter model, which results in high levels of mathematical tractability (Sood and Redner, 2005), the non-linearity of the Ising model requires different techniques for analysis. Inspired by the utilization of network control, 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, in this chapter, we make the following contributions to address research question 2 on network inference acceleration outlined in Section 1.1: 1) To address the research question (2c) on the applicability of inference acceleration across different opinion diffusion models, we contrast the approaches used in the linear voter model and the non-linear Ising model. In the voter model, transition probabilities are linear functions of neighboring states, and we employed greedy local optimization heuristics, such as a one-step-ahead algorithm, to minimize variance at individual nodes. However, in the non-linear Ising model, where transition probabilities are sigmoid functions of neighbor states, optimization becomes more complex. We adapted by using approximation methods, like high-temperature Taylor series expansion and mean-field approximation. These techniques help linearize the model's dynamics at high temperatures, offering a way to estimate optimal allocations in a non-linear system. Our results establish that the convergence of inference in non-linear Ising systems can also be accelerated by smartly targeting agents in the network with optimized control gains to

improve the quality of the generated dataset for inference. 2) To address research questions 2d and 2e, we make the following findings. Our study identifies two primary patterns for optimal budget allocation in network influence strategies, depending on the presence or absence of budget constraints. Without budget constraints, optimal allocations are governed solely by the neighborhood characteristics of the targeted node and the opponent's influence strength. However, when budget limitations are considered, the strategy shifts to recalibrating these optimal allocations based on the uncertainty levels in other agents' estimations. Specifically, agents with greater uncertainty in their estimations are more likely to receive resources aligned with their ideal, unconstrained allocations.

5.2 Model Description and Methods

Following the work by Galam (2008), 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 $\boldsymbol{W} = \{w_{ij}\}_{i,j=1}^{N}$. 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. In contrast to Chapters 3 and 4, where we denote the node states as 0 or 1 in the voter model, here, we follow the norm in the Ising model and label the node states as $s_i(t) = -1$ or $s_i(t) = 1$ for $i = 1, \cdots, N$ and $t \geq 0$. Here, we assume the existence of two opposite external fields A and B, also referred to as controllers 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 that 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: $\sum_i J_{A,i}(t) \leq b_A$, $\sum_i J_{B,i}(t) \leq b_B$.

Following the commonly-used Glauber algorithm (Glauber, 1963; Galam, 2008) in simulating social dynamics, the parallel and discrete-time Ising dynamics consist of the following three steps: (i) Sum the weighted neighboring states of node i ($i = 1, ..., N$) at time t, which is denoted as $S_i(t) = \sum_j w_{ij} s_j(t)$. (ii) 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 $e^{-\Delta E_i(t)/\tau} / \left(1 + e^{-\Delta E_i(t)/\tau}\right)$ where τ 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

$$\begin{aligned} P_i(t) &= \begin{bmatrix} Pr(s_i(t+1)=-1) & Pr(s_i(t+1)=1) \end{bmatrix} \\ &= \begin{bmatrix} \frac{e^{\left(-2S_i(t)-2J_{A,i}(t)+2J_{B,i}(t)\right)/\tau}}{1+e^{\left(-2S_i(t)-2J_{A,i}(t)+2J_{B,i}(t)\right)/\tau}} & \frac{1}{1+e^{\left(-2S_i(t)-2J_{A,i}(t)+2J_{B,i}(t)\right)/\tau}} \end{bmatrix}. \end{aligned} \tag{5.1}$$

Here, $Pr(s_i(t+1)=\pm 1)$ stands for the probability of node i to have state ± 1 at time $t+1$. From the third term of Eq. (5.1), 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 chapter, 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) (Myung, 2003) for parametric inference. More specifically, given the transition probabilities of opinion flips in Eq. (5.1), the logarithm of the likelihood function for observing time series of opinion changes during time span $[0, T]$ for node i is

$$L_i(T) = \sum_{t\in[0,T-1]} \left[\frac{1+s_i(t+1)}{2} \log\left(Pr(s_i(t+1)=1)\right) + \frac{1-s_i(t+1)}{2} \log\left(Pr(s_i(t+1)=-1)\right) \right]. \tag{5.2}$$

Inserting $Pr(s_i(t+1)=1)$ and $Pr(s_i(t+1)=-1)$ from Eq. (5.1) into Eq. (5.2) yields the full expression. By maximizing the likelihood function of Eq. (5.2) 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 $J_{B,i}$ (Myung, 2003). 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 (Ly et al., 2017) as a metric to evaluate the quality of fit of the estimators derived by MLE

as in Chapter 4. The Fisher information gives a measure for the dispersion between the deducted estimators from MLE and actual values. According to Efron and Hinkley (1978), the Fisher information $I(J_{B,i}, T)$ pertaining to $J_{B,i}$ is defined as the expectation of the second-order partial derivative of Eq. (5.2) with respect to $J_{B,i}$, i.e.,

$$I(J_{B,i}, T) = E\left[\frac{\partial^2 L_i(T)}{\partial J_{B,i}^2}\right] = -\sum_{t\in[0,T-1]}\left[\frac{\operatorname{sech}^2\left(\frac{J_{A,i}(t)-J_{B,i}+S_i(t)}{\tau}\right)}{\tau^2}\right]. \tag{5.3}$$

As the function domain of hyperbolic secant ($\operatorname{sech}(\cdot)$) is limited within the range of $(0, 1]$, the Fisher information $I(J_{B,i}, T)$ 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 T increases, the value of $I(J_{B,i}, T)$ decreases accordingly. Thereafter, by taking the negative reciprocal of Fisher information, one can generate confidence intervals for the MLE estimators (Gao et al., 2018). 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

$$\left(\hat{J}_{B,i} - J_{B,i}\right) \xrightarrow{D} \mathcal{N}\left(0, \left[-I(J_{B,i}, T)\right]^{-1}\right). \tag{5.4}$$

Here, $\mathcal{N}(0, \left[-I(J_{B,i}, T)\right]^{-1})$ represents the normal distribution with standard deviation $\left[-I(J_{B,i}, T)\right]^{-1/2} > 0$, and mean value 0. Moreover, as $I(J_{B,i}, T)$ 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. (5.4), 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 $\hat{J}_{B,i}$ (Ly et al., 2017).

In the following, to improve estimates of opponents, we minimize the standard deviation deduced by the Fisher information in Eq. (5.3). 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. (5.3), we observe that the Fisher information can be calculated in a recursive way, i.e.,

$$I(J_{B,i}, t+1) = I(J_{B,i}, t) - \frac{\operatorname{sech}^2\left(\frac{J_{A,i}(t)-J_{B,i}+S_i(t)}{\tau}\right)}{\tau^2}. \tag{5.5}$$

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

$$\hat{\sigma}_i(\hat{J}_{B,i}, t+1) = \left(-I(\hat{J}_{B,i}, t) + \frac{\operatorname{sech}^2\left(\frac{J_{A,j}(t)-\hat{J}_{B,j}+S_j(t)}{\tau}\right)}{\tau^2}\right)^{-1/2}. \tag{5.6}$$

Extending our previous framework in Chapter 4, 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

$$\begin{aligned}\{\overbrace{J^*_{A,1}(t),...,J^*_{A,N}(t)}^{N \text{ agents in the network}}\} &= \arg\min \sum_{i=1}^{N} \hat{\sigma}_i(\hat{J}_{B,i}, t+1) \\ &= \arg\min \sum_{i=1}^{N} \left[-I(\hat{J}_{B,i}, t) + \frac{\operatorname{sech}^2\left(\frac{J_{A,i}(t)-\hat{J}_{B,i}+S_i(t)}{\tau}\right)}{\tau^2} \right]^{-1/2} \end{aligned} \tag{5.7}$$

subject to

$$\begin{aligned} &\sum_{i=1}^{N} J^*_{A,i}(t) \leq b_A \\ &J^*_{A,i}(t) \geq 0 \qquad \text{for} \quad 1 \leq i \leq N. \end{aligned}$$

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 $1 \leq i \leq N$) 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. (5.2), the only unknown parameters in Eq. (5.7) will be $J_{A,i}(t)$.

As Eq. (5.7) 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. (5.6) 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.2 for a detailed proof), given by

$$J_{A,i}(t) = \begin{cases} \hat{J}_{B,i} - S_i(t), & \hat{J}_{B,i} \geq S_i(t) \\ 0, & \hat{J}_{B,i} < S_i(t) \end{cases}. \tag{5.8}$$

Additionally, by taking the second-order derivative of Eq. (5.6), we find that for lower temperatures τ, 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 A.3).

5.2.1 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. (5.6)) does not exist for low temperatures, in the following, we focus on the approximation in the high-temperature regime. Formally, for high temperatures, we have

$$\begin{aligned}\hat{\sigma}_i(\hat{J}_{B,i}, t+1) =& \hat{\sigma}_i(\hat{J}_{B,i}, t) - \frac{\hat{\sigma}_i^3(\hat{J}_{B,i}, t)}{2\tau^2} + \frac{3\hat{\sigma}_i^5(\hat{J}_{B,i}, t)}{8\tau^4} + \\ & \frac{(J_{A,i}(t) - J_{B,i} + S_i(t))^2 \hat{\sigma}_i^3(\hat{J}_{B,i}, t)}{2\tau^4} + \mathcal{O}\left(\left(\frac{1}{\tau}\right)^5\right).\end{aligned} \tag{5.9}$$

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

Next, to obtain an analytical solution for the minima of Eq. (5.7) subject to the inequality constraint $\sum_{i=1}^{N} J_{A,i}(t) \leq b_A$, we apply the Lagrange multiplier technique (Bertsekas, 2014) based on the approximation shown in Eq. (5.9). In order to handle the inequality constraint in the objective function, we introduce slack variables q and d_i ($1 \leq i \leq N$). By letting

$$f\left(J_A(t)\right) = \sum_{i=1}^{N}\left[\hat{\sigma}_i(\hat{J}_{B,i}, t) - \frac{\hat{\sigma}_i^3(\hat{J}_{B,i}, t)}{2\tau^2} + \frac{3\hat{\sigma}_i^5(\hat{J}_{B,i}, t)}{8\tau^4} + \frac{\left(J_{A,i}(t) - \hat{J}_{B,i} + S_i(t)\right)^2 \hat{\sigma}_i^3(\hat{J}_{B,i}, t)}{2\tau^4}\right], \tag{5.10}$$

where $J_A(t) = \{J_{A,i}(t)\}_{i=1}^{N}$, Eq. (5.7) is converted into

$$\begin{aligned}&\overbrace{\{J_{A,1}^{approx}(t), ..., J_{A,N}^{approx}(t)\}}^{N \text{ agents in the network}} = \arg\min f\left(J_A(t)\right) \\ &\text{subject to} \\ &\qquad h\left(J_A(t)\right) = \sum_{i=1}^{N} J_{A,i}^{approx}(t) - b_A + q^2 = 0 \\ &\qquad g_i\left(J_A(t)\right) = -J_{A,i}^{approx}(t) + d_i^2 = 0 \qquad \text{for} \quad 1 \leq i \leq N,\end{aligned} \tag{5.11}$$

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

$$\mathcal{L}(J_A(t), \lambda, \{\beta_i\}_{i=1}^{N}) = f\left(J_A(t)\right) + \lambda h\left(J_A(t)\right) + \sum_{i=1}^{N}\left[\beta_i g_i\left(J_A(t)\right)\right]. \tag{5.12}$$

To obtain the optimal solution of Eq. (5.11), we set the gradients of Eq. (5.12) to be 0:

$$\begin{aligned}
\frac{\partial \mathcal{L}}{\partial J_{A,i}} &= \frac{\left(J_{A,i}(t) - \hat{J}_{B,i} + S_i(t)\right)\hat{\sigma}_i^3(\hat{J}_{B,i}, t)}{\tau^4} + \lambda - \beta_i = 0 \\
\frac{\partial \mathcal{L}}{\partial \lambda} &= \sum_{i=1}^{N} J_{A,i}(t) - b_A + q^2 = 0 \\
\frac{\partial \mathcal{L}}{\partial q} &= 2\lambda q = 0 \\
\frac{\partial \mathcal{L}}{\partial \beta_i} &= -J_{A,i}(t) + d_i^2 = 0 \\
\frac{\partial \mathcal{L}}{\partial d_i} &= 2\beta_i d_i = 0.
\end{aligned} \tag{5.13}$$

In accordance with the Krush-Kuhn-Tucker conditions (Bertsekas, 2014), 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.,

$$\begin{aligned}
\lambda &\geq 0 \\
\beta_i &\geq 0.
\end{aligned} \tag{5.14}$$

By solving the system of equations (5.13) and (5.14), we find that there are two possible cases for $J_{A,i}^{approx}(t)$:

(i)

$$\text{If} \quad \sum_{i=1}^{N} J_{A,i}^{approx}(t) \leq b_A, \quad \text{then} \quad J_{A,i}^{approx}(t) = \begin{cases} \hat{J}_{B,i} - S_i(t), & \hat{J}_{B,i} \geq S_i(t) \\ 0, & \hat{J}_{B,i} < S_i(t) \end{cases}. \tag{5.15}$$

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. (5.8)).

(ii) Otherwise, let

$$\phi_i(t) = \left(\hat{J}_{B,i} - S_i(t)\right) + \frac{\left[b_A - \sum_{J_{A,j}^{approx}(t)\neq 0}\left(\hat{J}_{B,j} - S_j(t)\right)\right]\hat{\sigma}_i^{-3}(\hat{J}_{B,i}, t)}{\sum_{J_{A,j}^{approx}(t)\neq 0}\hat{\sigma}_j^{-3}(\hat{J}_{B,j}, t)}. \tag{5.16}$$

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

$$J_{A,i}^{approx}(t) = \begin{cases} 0, & \hat{J}_{B,i} \leq S_i(t) \quad \text{or} \quad \phi_i(t) \leq 0 \\ \phi_i(t), & \phi_i(t) > 0 \end{cases}. \tag{5.17}$$

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. (5.15) 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. (5.16). (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. (5.17) are to re-weight the unconstrained solutions in Eq. (5.8) 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.

5.2.2 High-temperature Mean-field Approximation

From the analytical solution of the high-temperature Taylor approximation shown in Eqs. (5.15) and (5.17), 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 replacing the real states of nodes with 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 (Lynn and Lee, 2016)), we obtain

$$r_i = \tanh\left[\left(\sum_j w_{ij} r_j + J_{A,i} - J_{B,i}\right)/\tau\right] \tag{5.18}$$

for the opinion dynamics following the Glauber dynamics. Because of the hyperbolic term in Eq. (5.18), it is hard to obtain an explicit solution for r_i. In the following, we consider the mean-field approximation in the high-temperature region. Therefore, by applying the Taylor expansion for high temperatures with respect to Eq. (5.18), we have

$$r_i = \frac{\sum_{j=1}^{N} w_{ij} r_j + J_{A,i} - J_{B,i}}{\tau} + \mathcal{O}\left(\tau^{-3}\right) \tag{5.19}$$

To further simplify the solution of Eq. (5.19), we assume that all nodes are affected by the same mean field, in which we have:

$$r_i \approx \frac{k_i \langle r \rangle + J_{A,i} - J_{B,i}}{\tau} \tag{5.20}$$

where

$$\langle r \rangle = \frac{1}{N} \sum_i r_i. \tag{5.21}$$

Here, $k_i = \sum_j w_{ij}$ represents the weighted degree for node i. By summing Eq. (5.20) over all nodes, we obtain

$$\sum_i r_i = N \langle r \rangle = \frac{N \langle k \rangle \langle r \rangle + \sum_i J_{A,i} - b_B}{\tau} \tag{5.22}$$

i.e.,

$$\langle r \rangle = \frac{\sum_i J_{A,i} - b_B}{TN - \langle k \rangle N} \tag{5.23}$$

where $\langle k \rangle$ represents the averaged linking weights of the network. Inserting Eq. (5.23) into Eq. (5.20) 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 $k_i \langle r \rangle$ in Eqs. (5.15) and (5.17), we have

$$J_{A,i}^{MF} = \begin{cases} \hat{J}_{B,i} - k_i \langle r \rangle, & \hat{J}_{B,i} \geq k_i \langle r \rangle \\ 0, & \hat{J}_{B,i} < k_i \langle r \rangle \end{cases}, \quad \text{If} \quad \sum_{i=1}^{N} J_{A,i}^{MF} \leq b_A. \tag{5.24}$$

and otherwise

$$J_{A,i}^{MF} = (\hat{J}_{B,i} - k_i \langle r \rangle) + \frac{(b_A - \sum_{J_{A,j}^{MF} \neq 0} \hat{J}_{B,j} + \sum_{J_{A,j}^{MF} \neq 0} k_j \langle r \rangle) \hat{\sigma}_i^{-3}(\hat{J}_{B,i}, t)}{\sum_{J_{A,j}^{MF} \neq 0} \hat{\sigma}_j^{-3}(\hat{J}_{B,j}, t)}. \tag{5.25}$$

Here, $J_{A,i}^{MF}$ represents the high-temperature mean-field approximation for the optimal budget allocation of node i. Eq. (5.25) 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. (5.6), to compute $\hat{\sigma}_j^{-3}(\hat{J}_{B,j}, t)$, we also need the neighbouring states. To further simplify Eq. (5.25), we apply the Taylor expansion on $\hat{I}(\hat{J}_{B,i}, T)$ for high temperatures, in which we have

$$I(\hat{J}_{B,i}, T) = \sum_{t=0}^{T-1} \left[-\frac{1}{\tau^2} + \frac{\left(J_{A,i}(t) - \hat{J}_{B,i} + S_i(t)\right)^2}{\tau^4} + \mathcal{O}\left(\tau^{-5}\right) \right] \approx \sum \left(-\frac{1}{\tau^2} + \frac{r_i^2}{\tau^2} \right). \tag{5.26}$$

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

$$\hat{\sigma}_i(\hat{J}_{B,i}, T) \approx \left(\frac{T - r_i^2 T}{\tau^2} \right)^{-1/2} \tag{5.27}$$

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

$$\hat{\sigma}_i(\hat{J}_{B,i}, T)^{-3} = \left(\frac{T}{\tau^2}\right)^{3/2} - \frac{3}{2} r_i^2 \left(\frac{T}{\tau^2}\right)^{3/2} + \mathcal{O}\left(r_i^4\right). \tag{5.28}$$

If we only consider the first term in Eq. (5.28) and replace $\hat{\sigma}_i(\hat{J}_{B,i}, T)^{-3}$ with it in Eq. (5.25), we obtain

$$J_{A,i}^{MF} \approx \hat{J}_{B,i} + \frac{b_A - \sum_{J_{A,j}^{MF} \neq 0} \hat{J}_{B,j}}{Z} + \frac{b_A - b_B}{\tau N - \langle k \rangle N} \left(\frac{\sum_{J_{A,j}^{MF} \neq 0} k_j}{Z} - k_i \right), \tag{5.29}$$

Where Z counts for the number of no-zero $J_{A,i}^{MF}$. Eq. (5.29) exists when $J_{A,i}^{MF}$ is non-negative.

5.2.3 Numerical One-step-ahead Optimization

The approximated solutions for optimal budget allocations obtained in Sections 5.2.1 and 5.2.2 are deduced in the premise of high temperatures. To obtain comprehensive results for the optimized budget allocations regardless of the temperature constraints, we use the interior-point method (Kim et al., 2007) for numerical optimization of Eq. (5.7) in all temperature regions. By doing this, we aim at getting step-wise optimal budget allocations $J_{A,i}^{opt}$ for $i = 1, \cdots, N$ 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 3.

Specifically, the detailed experimental setup for the numerical one-step-ahead optimization shown in Algorithm 3 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. (5.4), 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 3. We also keep track of the likelihood functions for the first T_0 updates for all nodes, denoted as $L_i(T_0)$ for $i = 1, \cdots, N$. (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. (5.2). Thereafter, we compute the Fisher information defined in Eq. (5.5) and insert them into the objective function in Eq. (5.7). By optimizing Eq. (5.7) 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 3 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 $\mathcal{O}\left((T_1 - T_0)N^3\right)$.

input : Adjacency matrix $\boldsymbol{W}$, 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 $T_0 \leq t \leq T_1$, $1 \leq i \leq N$

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* $t = T_0$;

while $t \leq T_1$ **do**

  $\hat{J}_{B,i} = \max_{J_{B,i}} L_i(t) \quad$ for $1 \leq i \leq N$;

  obtain $J_{A,i}^{opt}(t)$ by optimizing Eq. (5.7);

  update the network following stochastic Ising dynamics with $\{J_{A,i}^{opt}(t)\}_{i=1}^{N}$;

  $t = t + 1$;

end

Algorithm 3: Numerical one-step-ahead optimization

5.3 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 5.3.1, 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 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 5.3.2. 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 5.3.3. Furthermore, in Section 5.3.4, 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 $\langle k \rangle = 10$, generated according to the configuration model (Catanzaro et al., 2005). To ensure a large degree heterogeneity, the degree distribution of the constructed networks is $p_k \sim k^{-2}$ where k represents the node's degree. The results shown in the following context are all based on 50 repetitions of the corresponding experiments.

5.3.1 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. (5.4) 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 $\langle J_{B,i} \rangle = 10$ equal to the average degree of the networks. To proceed, in Fig. 5.1, 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). In Panels (a) and (b) of Fig. 5.1, 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 Chapter 4.

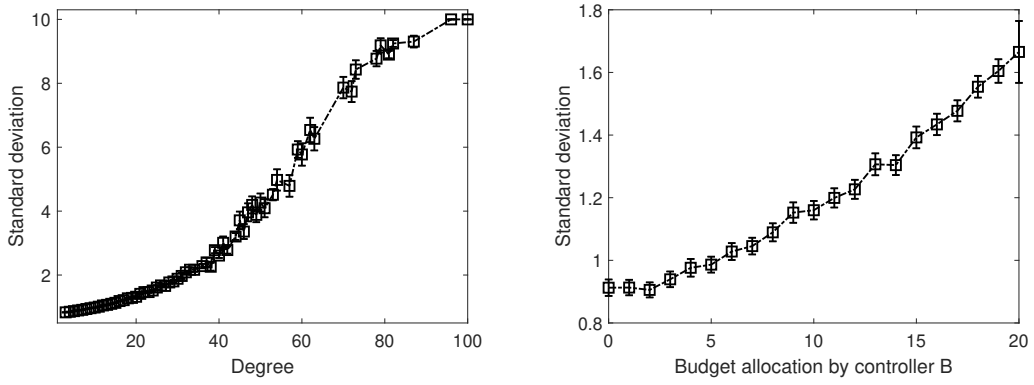

(a) (b)

FIGURE 5.1: **(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 $\tau = 20$. We use a setting in which controller B targets all nodes with allocations randomly sampled from a uniform distribution, and the budget allocation by controller B per node on average is 10. We group the value of the x-axis into bins with width 1 and lower limits are inclusive, e.g.,[0,1). Error bars indicate 95% confidence intervals.

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 optimisation.

To proceed, in Fig. 5.2 (a), we show the dependence of the mean standard deviation $\bar{\sigma}_{eq}$ for the equally targeting scheme over all nodes on varying temperatures τ 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. 5.2 (a), we make the following observations. First, as shown in the upper-left corner of Fig. 5.2 (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 temperatures, the system falls into stabilization which leads to spontaneous magnetization (Chandler, 1987). 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. 5.2 (a), we find that, for extremely high temperatures, the inference errors are roughly the same for different equally targeting budgets. Together with the probabilities of state flips in Eq. (5.1), we can thus understand why the Ising system is harder to control for inference acceleration as the temperature increases: In the context of high temperatures, 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. 5.2 (a), we find that the dependence of mean standard deviation on temperatures is a convex shape. This means that there is a temperature at which predictions of opponent strategies are the most accurate.

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 3 in Fig. 5.2 (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. 5.2 (b). It becomes clear that the optimization works best for extremely low temperatures with a reduction of the relative standard deviation close to 100%. Moreover, with an increase in 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.

5.3.2 Results for High-temperature Taylor Series Approximation

As the approximated solutions in Eq. (5.17) for the optimal budget allocations are deducted under the assumption of high temperatures, identifying the feasible region

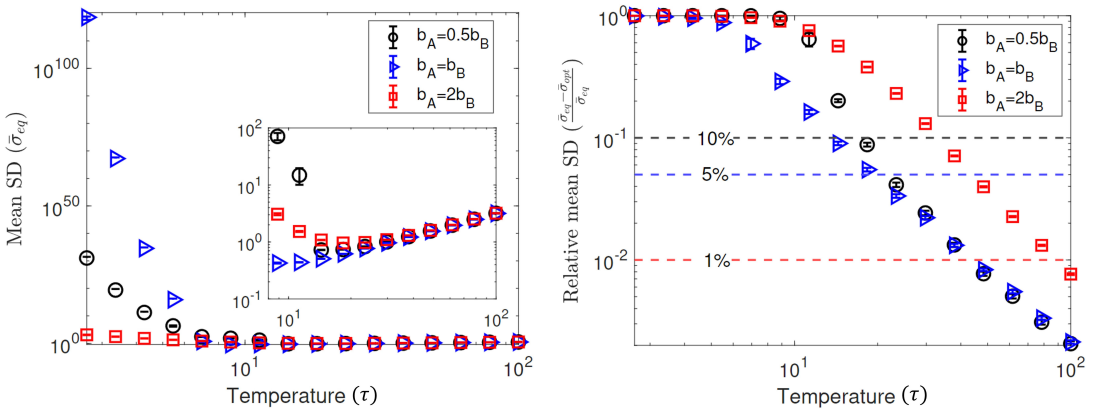

(a) (b)

FIGURE 5.2: **(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 temperatures τ 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 temperatures τ. $\bar{\sigma}_{opt}$ is calculated by averaging the standard deviation of estimators obtained via applying the 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, and the budget allocation by controller B per node on average is 10. Error bars indicate 95% confidence intervals.

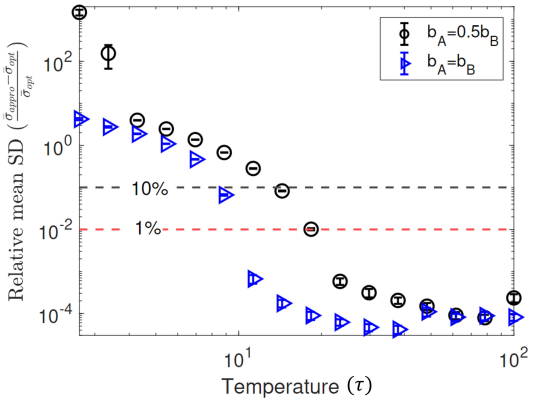

(a)

FIGURE 5.3: **(a)** Dependence of the relative mean standard deviation $\frac{\bar{\sigma}_{approx}-\bar{\sigma}_{opt}}{\bar{\sigma}_{opt}}$ on varying temperatures τ. $\bar{\sigma}_{approx}$ stands for the mean standard deviation of estimators obtained via applying the high-temperature Taylor series approximation shown as Eqs. (5.15) and (5.17) 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.

where the Taylor series approximation has close performance to the optimization is

crucial. To proceed, in the panels of Fig. 5.3, we compare the mean standard deviation over all nodes calculated via the one-step-ahead optimization of Algorithm 3 with the Taylor series approximation in Eq. (5.17). In more detail, in Fig. 5.3 (a), we show the dependence of relative mean standard deviation $\frac{\bar{\sigma}_{approx}-\bar{\sigma}_{opt}}{\bar{\sigma}_{opt}}$ on the temperature τ at update 1000. Here, $\bar{\sigma}_{approx}$ stands for the mean standard deviation computed by the Taylor series approximation (see Eqs. (5.15) and (5.17)). To obtain $\bar{\sigma}_{approx}$, we consider the following procedure: After the first $T_0 = 100$ updates of system initialization as mentioned in Algorithm 3, 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. 5.2, we only present two budget cases $b_A = 0.5b_B$ and $b_A = b_B$ in Fig. 5.3 (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. (5.8) and (5.15). As a consequence, the mean standard deviation $\bar{\sigma}_{opt}$ and $\bar{\sigma}_{approx}$ will also be the same, and thus the relative mean standard deviation $\frac{\bar{\sigma}_{approx}-\bar{\sigma}_{opt}}{\bar{\sigma}_{opt}} = 0$ can not be presented in a logarithmic figure.

Continuing with the results in Fig. 5.3 (a), we find that even though the approximation is based on the assumption of high temperatures, 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 temperatures 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 Figure 5.3 (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. (5.15) and (5.17) 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.

5.3.3 Results for High-temperature Mean-field Approximation

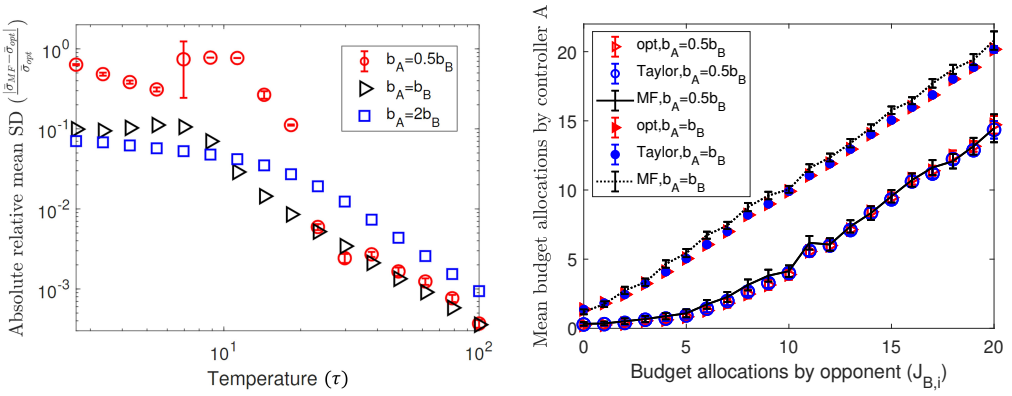

FIGURE 5.4: **(a)** Dependence of the absolute relative mean standard deviation $\frac{|\bar{\sigma}_{MF}-\bar{\sigma}_{opt}|}{\bar{\sigma}_{opt}}$ on varying temperatures τ. $\bar{\sigma}_{MF}$ is computed by averaging Eq. (5.27) over all nodes. Different relative budget constraints $b_A/b_B = \{0.5, 1, 2\}$ 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 values of the x-axis are grouped into bins with width 1 and lower limits are inclusive, e.g.,[0,1). Error bars indicate 95% confidence intervals.

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 5.2.2 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 5.3.2, 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. 5.4 (a) shows the dependence of the absolute relative mean standard deviation $\frac{|\bar{\sigma}_{MF}-\bar{\sigma}_{opt}|}{\bar{\sigma}_{opt}}$ on the temperature. Here $\bar{\sigma}_{MF}$ is computed by averaging Eq. (5.27) over all nodes. Notice that, on the y-axis, instead of calculating the relative values as Figs. 5.2 and 5.3, 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. (5.27), $\bar{\sigma}_{MF} = \left(\frac{T}{\tau^2}\right)^{-1/2}$, which only considers the first term of the Taylor expansion of the Fisher information in Eq. (5.26) and ignores the non-negative term of $\frac{(J_{A,i}(t)-J_{B,i}+S_i(t))^2}{\tau^4}$ 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. 5.4(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. 5.3(a) and 5.4(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 temperatures, it works well also in the intermediate temperature region where the difference in the standard deviation is less than 10% for temperatures over 10 when $b_A \geq b_B$.

Moreover, in Figure 5.4(a), for $b_A = 0.5b_B$, after $T = 5.5$, we observe an increase in the absolute values 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. (5.27), 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. (5.23)). As the experiments are carried out on networks with average degree $\langle d \rangle = 10$, and given $|\sum_i J_{A,i} - b_B| \gg 0$ 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. (5.27)). As a result, $\frac{|\bar{\sigma}_{MF} - \bar{\sigma}_{opt}|}{\bar{\sigma}_{opt}}$ 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. (5.23), $\langle r \rangle$ is dependent on the real budget used by controller A, which could be less or equal to the budget constraint. In these cases, $|\sum_i J_{A,i} - b_B|$ 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 pattern 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 $\tau = 20$. In more detail, in Fig. 5.4 (b), we present the dependence of the mean budget allocations by controller A averaged over updates on the opponent's budget allocations at $\tau = 20$. Note that, for ease of observation, we group opponent's budget allocations on the x-axis into bins with width 1 in Fig. 5.4 (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 is also considered in Fig. 5.4(b), where the marked symbols and the dashed line represent $b_A = b_B$, and blank symbols and the solid line are for $b_A = 0.5 b_B$. By observing the trends of budget allocations by controller A for varying budget allocations from the 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. 5.4 (b) for $b_A = b_B$.

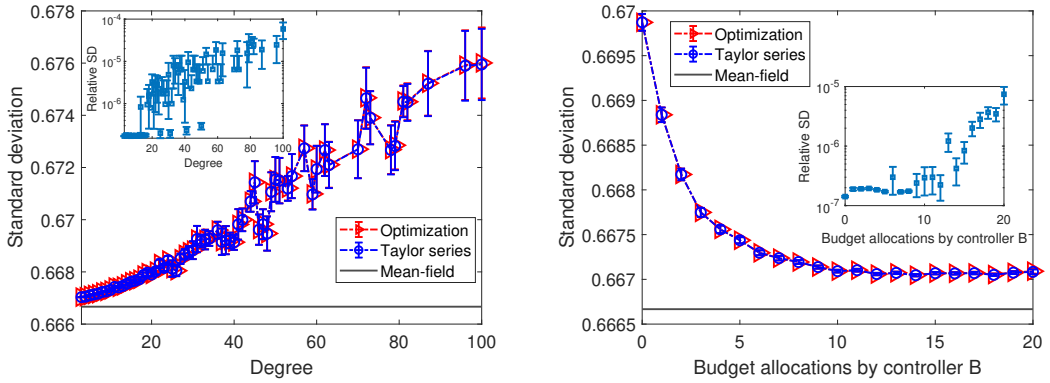

(a) (b)

FIGURE 5.5: **(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 degrees. **(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 $\tau = 20$. 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, and the budget allocation by controller B per node on average is 10. Error bars indicate 95% confidence intervals.

To proceed, we investigate the accuracy of inference for agents after we performed the optimal control and compare it with the results for the no interference scenario in Fig. 5.1. Similar to Fig. 5.1 (a), we show the dependence of standard deviation of estimations on nodes with different degrees when applying the optimization of Algorithm 3 or using high-temperature Taylor and mean-field approximations in Fig. 5.5 (a). By

observing Fig. 5.5 (a), we find that the high-temperature Taylor approximation provides close performance in explaining the influence of node degree heterogeneity on inference accuracy compared with the exact numerical optimization. Moreover, consistent with the no interference case in Fig. 5.1 (a), we find also when optimizing that it is harder to predict a higher degree node. In Fig. 5.5 (b), we present the dependence of the standard deviation of estimators on opponent budget allocations. We find that, compared with Fig. 5.1 (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. 5.1 (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. 5.1 (b), where we observed that nodes targeted with higher control gains are harder to predict.

5.3.4 Influence of Budget Availability

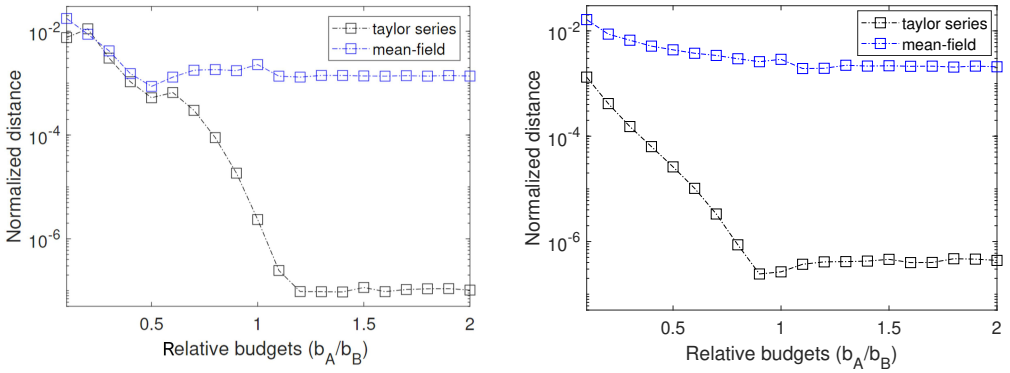

(a) (b)

FIGURE 5.6: 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. (5.30) and (5.31), 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 $\tau = 20$, and **Panel (b)** is $\tau = 50$.

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. 5.4 (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, 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

$$\sqrt{\sum_m \left(\bar{J}_A^{opt,m} - \bar{J}_A^{approx,m}\right)^2} / b_A^*. \tag{5.30}$$

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

$$\sqrt{\sum_m \left(\bar{J}_A^{opt,m} - \bar{J}_A^{MF,m}\right)^2} / b_A^*. \tag{5.31}$$

Here b_A^* represents the actual total budget that the controller A has used during the optimization in Figs. 5.6.

For comparison, Fig. 5.6 (a) is computed at intermediate temperature $\tau = 20$ while (b) is for high-temperature at $\tau = 50$. By observing Fig. 5.6 (a), we find that the distances of budget allocations between the numerical optimization and approximations calculated by Eqs. (5.30) and (5.31) 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 temperatures, we find that, temperatures 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 $\tau = 20$, 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 $\tau = 50$, high-temperature Taylor expansion has close performance to the optimization at $b_A/b_B \geq 0.9$.

5.4 Summary

Existing literature related to the inverse Ising problem has been restricted to conducting the inference of linking weights based on a given and fixed dataset. In contrast, in this chapter, 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 in a step-wise way. 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 strategically will substantially speed up the convergence of inference. Specifically, in the low-temperature region, the proposed heuristics will decrease the uncertainty of inference by almost 100 percent compared with the benchmark case. This agrees with the finding of Decelle et al. (2016), who demonstrate that out-of-equilibrium data will result in a much more accurate inference compared with equilibrium/stationary data. Moreover, even though the external controllers have limited power in manipulating the networked dynamics at higher temperatures, 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 a 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 neighbouring states. In the case of an insufficient budget, we utilize the 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 nodes' degree. Moreover, even by performing the optimal control, high degrees will impede nodes from being inferred accurately. The main reason for this is, for nodes with large degrees, 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 Chen and Lai (2018) and Cai et al. (2022b). However, contrary to the results obtained from the voter model in Chapter 4, 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 no access to tracking the system states at each time step by substituting 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 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.

Chapter 6

Conclusions and Future Work

This thesis focuses on the dual challenges of inter-temporal influence maximization and accelerating the convergence of inference. To this end, we have presented various studies targeting different facets of these challenges. The subsequent sections of this chapter summarize the findings derived from these studies and outline potential directions for future research.

6.1 Conclusions

In this thesis, we mainly tackle two research topics, including influence maximization, and its inverse problem named network inference, with a focus on the opponent-strategy reconstruction. As mentioned in Section 2.2, the solutions to influence maximization and opponent-strategy reconstruction are closely related to the underlying influence diffusion models. In this work, we choose the paradigmatic voter model and the Ising model as the opinion diffusion models, owing to their rich history. Furthermore, we assume the external influences interact with the internal opinion dynamics of nodes via strategically building connections to nodes in the network with the aim to maximally spread their opinions or expedite opponent strategy inference. In this approach, external influences effectively model targeted advertising or political campaigns in which nodes targeted by them are not assumed to be immediately committed to a given opinion but also subject to peer influence in the network.

In Chapter 3, we study the IM problem under the framework of inter-temporal network control in voter dynamics with time and budget constraints. This is explored under the constant-opponent setting where one active controller competes against a known and fixed-strategy opponent, and the game-theoretical setting in which both controllers simultaneously decide on an inter-temporal allocation protocol at time zero. By doing so,

we have obtained the following main findings as well as solved questions of challenge 1 in Section 1.1.

1. By assuming controllers have the flexibility to choose when to start control, we integrate the inter-temporal network control into the opinion dynamics (see question 1a). After an extensive investigation of the dependency of optimal starting times on the network configurations, we find that, for short time horizons, maximum influence is achieved by starting relatively later on more heterogeneous networks than in more homogeneous networks, while the opposite holds for long time horizons (see question 1b). Moreover, for the constant-opponent setting, the optimized controller will only use its budget closer to the end of the campaign. In contrast, for the game-theoretical setting, the controller with budget superiority will start earlier (see question 1c).

2. We extend the simple scenario in which all agents have the same starting time and budget allocation to a more generalized setting named individual optimization by assuming that each agent has individual starting times and budget allocations. In the individual optimization, we find that the strategic allocation is effective only when fewer resources are available to the controller, as it will only focus on low-degree nodes and leave high-degree nodes untargeted.

In Chapters 4 and 5, to address challenge 2 about reconstructing opponent's strategies from binary-state dynamics, we use maximum likelihood estimation to obtain estimates of opponent's budget allocations. Instead of passively observing the opinion dynamics, we integrate the network control into network reconstruction and aim to speed up the convergence of estimates of opponent's budget allocations via interference from the active controller. As the variance of a maximum likelihood estimate at each step can be derived from the Fisher information, the problem of accelerating the convergence of estimates is transformed into minimizing the variance of estimates by optimally distributing the budget allocations of the active controller over time. By doing so, we have obtained the following main findings as well as solved questions 2a, 2b, 2d, and 2e. Specifically, in Chapter 4, we have the following findings:

1. We find that it is possible to accelerate the inference process by strategically interacting with the network dynamics in the simple contagion model like the voter model.

2. We first consider accelerating the inference of the opponent strategy at a single node, when only the inferred node is controllable. In this setting, we find that the optimized resource allocation is inversely proportional to the sum of neighbouring opinion states.

3. We then consider minimizing the variance of the opponent influence at the inferred node when both the inferred node and also its neighbours are controllable. In this setting, we observe two regimes of the optimized resource allocations based on varying amounts of available resources for the active controller. If the active controller has very limited resources, then it should target the inferred node only. In contrast, if resources are large, a better strategy is to not target the inferred node, but instead focus only on neighbouring nodes.

4. In the scenario of inferring opponent strategies over entire networks, strategic allocations become increasingly important as more resources are available for the active controller. We also find that nodes with lower degrees and targeted with smaller amounts of resources by the opponent will generally have a smaller variance in inference.

Subsequently, in Chapters 5, we find:

1. Even in a setting with complex contagion, i.e., the Ising dynamics, the convergence of inference can be accelerated by selectively targeting agents in the network with optimized control gains to improve the quality of the generated dataset for inference.

2. There are 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 neighborhood 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.

3. When the real-time tracking of the system states is inaccessible, the mean-field approximation can provide a close performance to the optimal heuristics in the medium and high-temperature regions. Generally, 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.

Note that, the research presented in Chapters 4 and 5 focus specifically on developing and validating the framework of network inference acceleration. By testing the applicability and accuracy of the inference acceleration techniques in isolation, the aim is to comprehensively understand and optimize the inference process itself, without added complexities from the influence maximization scenario. However, integrating inference within influence maximization can be a potentially promising real-world application.

Nevertheless, this integration introduces many additional complexities that could obscure improvements directly attributable to the inference methodology. Therefore, a strategic phased approach was adopted here. The first phase, presented in Chapters 4 and 5, aim to establish a solid foundation in network inference before layering on the complexity of influence maximization algorithms. Time constraints also necessitated this staged prioritization. With the validation of the inference acceleration process, the future phase is to integrate it with influence maximization models as discussed in the future work in Section 6.2.

While this thesis aims to further understand influence and opinion dynamics in social networks, it is important to consider the ethical implications. The voter model and Ising model provide insights into how opinions spread, which could potentially be misused for propaganda or radicalization purposes. However, they only offer simplified representations of opinion spread in social networks, where real-world influences such as complex personal and social factors are not considered in these models. Therefore, this thesis represents an initial step toward mathematically characterizing influence processes and focuses on theoretical modeling and simulation. Any practical application would require extensive additional inputs from specific real-world systems. Moreover, further research could explore ways to mitigate potential harms, such as using influence modeling to promote diversity and reduce polarization.

6.2 Future Work

We have enriched the extensive and evolving research realm of opinion dynamics, concentrating on influence maximization and opponent strategy inference. Given the complexity of real-world social systems, there are many potential extensions and variations of our frameworks, that could deepen our understanding of the studied opinion phenomena. In the following, we will outline the possible research directions in the future.

An intriguing avenue for future exploration involves integrating inter-temporal influence maximization with network inference acceleration. Although this thesis examines these two research problems separately, the inter-temporal influence maximization problem can be contextualized within scenarios of revealing unknown opponent strategies or network structures, by strategically allocating resources to interact with opinion dynamics. In more detail, we can assume an active controller that simultaneously attempts to acquire new knowledge about the network structure or opponent strategy and optimize its decisions based on existing knowledge. In other words, the agent attempts to balance the exploration and exploitation in order to maximize its total expected spreads over the period of time considered (Kuleshov and Precup, 2014). Note that, there is an inherent cost associated with not knowing the opponent's strategy and budget in influence maximization scenarios. This lack of knowledge can lead

to suboptimal decisions and strategies, potentially resulting in less effective influence campaigns. Quantifying this cost is crucial to understand the baseline against which the efficiency of the inference process should be measured. Therefore, the exploration should involve comparing the cost of performing the inference process with the cost of not knowing the opponent's strategy. This involves evaluating the resources (time, computational power, etc.) required for inference and weighing them against the potential losses incurred due to a lack of opponent strategy insights.

Furthermore, our analysis has been confined to a single layer of social connections. However, societies often involve simultaneous interactions of diverse types, and these distinctions could be integrated into the formulation of opinion dynamics. This can be elegantly captured through multiplex networks (Battiston et al., 2017), which encompass multiple graph layers sharing a common vertex set yet potentially differing edge sets. The dynamics across these layers may unfold on varying timescales or adopt different propagation models. A compelling challenge arises in indirectly shaping the dynamics of one layer through control exerted on the other layers.

Additionally, when addressing the problem of opponent-strategy reconstruction, we still assume that the controller has complete knowledge of the network topology. Therefore, another interesting direction for future work is to reconstruct the network topology as well as the continuous influence intensity between agents from binary-state dynamics. Consistent with the idea of obtaining improved performance of network inference from limited data, one can investigate the problem of inferring the network coupling strengths from partially observed time series data from the perspective of experimental design. In more detail, by assuming that only a certain fraction of observations for agent states are available, we could study how an observer, who wants to maximize the accuracy of the network inference, should distribute a limited number of observations to generate data that allow for the most accurate possible inference.

Appendix A

A.1 Improvements of Vote Shares Gained via Individual Optimization for Varying Initial Conditions

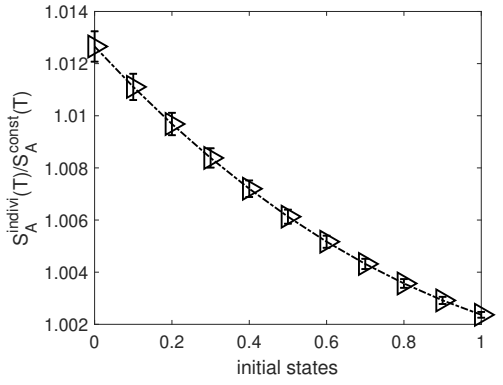

(a)

FIGURE A.1: Dependence of relative vote share $S_A^{indivi}(T)/S_A^{const}(T)$ on varying initial states. $S_A^{indivi}(T)$ stands for vote shares at time T calculated by individual optimization while $S_A^{const}(T)$ represents vote shares calculated by assigning a single optimized starting time for the whole network. All the calculations are based on networks with $N = 100$, degree of heterogeneity $\lambda = 1.6$, and $\langle k \rangle = 6$ and tested in 10 realizations. The control gains of controller B pertaining to each node are all fixed as 1 per unit time from time 0. The time horizon T is set as $T = 10$. The total budgets of controller A are set to be the same as controller B's, i.e., $b_A = b_B = N \times T$. Error bars indicate 95% confidence intervals.

A.2 Global Minima for a Single Node

Specifically, to find the minima for $\hat{\sigma}_i(\hat{J}_{B,i}, t+1)$ in Eq. (5.6), 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

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

By letting Eq. (A.1) equal to 0, we find that, $J_{A_j} = \hat{J}_{B,i} - S_i(t)$ is always a solution for Eq. (A.1) and it is independent of the value of last-step Fisher information $I(\hat{J}_{B,i}, t)$. Furthermore, the sign of Eq. (A.1) is determined by the term of $\left(\tau \tanh\left(\frac{J_{A,i}(t)-\hat{J}_{B,i}+S_i(t)}{\tau}\right)\right)$, and we have

$$\begin{aligned} &\tanh\left(\frac{J_{A,i}(t) - \hat{J}_{B,i} + S_i(t)}{\tau}\right) > 0 \qquad \text{for} \qquad J_{A,i}(t) > \hat{J}_{B,i}(t) - S_i(t) \\ &\tanh\left(\frac{J_{A,i}(t) - \hat{J}_{B,i} + S_i(t)}{\tau}\right) < 0 \qquad \text{for} \qquad J_{A,i}(t) < \hat{J}_{B,i} - S_i(t). \end{aligned} \tag{A.2}$$

This means that the function of Eq. (5.6) 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 can not be negative, when $\hat{J}_{B,i}(t) - S_i(t) < 0$, the global minima is $J_{A,i}(t) = 0$.

A.3 Curvature of Expected Standard Deviation

The second-order derivative of Eq. (5.6) with respect to $J_{A,i}(t)$ is given by

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

By letting $J_{A,i}(t) = \hat{J}_{B,i}(t) - S_i(t)$ in Eq. (A.3), we have

$$\left.\frac{\partial^2 \hat{\sigma}_i(\hat{J}_{B,i}, t+1)}{\partial J_{A,i}^2}\right|_{J_{A,i}(t)=\hat{J}_{B,i}(t)-S_i(t)} = \left(\tau^{-1} - I(\hat{J}_{B,i}, t)\right)^{1/2}. \tag{A.4}$$

For a larger temperature τ, the value of Eq. (A.4) is smaller. Therefore, the curve of Eq. (5.6) is flatter for a high temperature compared with a lower temperature.

References

Réka Albert and Albert-László Barabási. Statistical mechanics of complex networks. *Reviews of modern physics*, 74(1):47, 2002.

Agata Aleksiejuk, Janusz A Hołyst, and Dietrich Stauffer. Ferromagnetic phase transition in barabási–albert networks. *Physica A: Statistical Mechanics and its Applications*, 310(1-2):260–266, 2002.

Khurshed Ali, Chih-Yu Wang, and Yi-Shin Chen. A novel nested q-learning method to tackle time-constrained competitive influence maximization. *IEEE Access*, 7:6337–6352, 2018.

Aamena Alshamsi, Flávio L Pinheiro, and César A Hidalgo. When to target hubs? strategic diffusion in complex networks. *arXiv preprint arXiv:1705.00232*, 2017.

Aamena Alshamsi, Flávio L Pinheiro, and Cesar A Hidalgo. Optimal diversification strategies in the networks of related products and of related research areas. *Nature communications*, 9(1):1328, 2018.

Amirhossein Ansari, Masoud Dadgar, Ali Hamzeh, Jörg Schlötterer, and Michael Granitzer. Competitive influence maximization: integrating budget allocation and seed selection. *arXiv preprint arXiv:1912.12283*, 2019.

Sinan Aral and Paramveer S Dhillon. Social influence maximization under empirical influence models. *Nature human behaviour*, 2(6):375–382, 2018.

Erik Aurell and Magnus Ekeberg. Inverse ising inference using all the data. *Physical review letters*, 108(9):090201, 2012.

Mehdi Azaouzi, Wassim Mnasri, and Lotfi Ben Romdhane. New trends in influence maximization models. *Computer Science Review*, 40:100393, 2021.

Ludovica Bachschmid-Romano and Manfred Opper. Inferring hidden states in a random kinetic ising model: replica analysis. *J. Stat. Mech. Theory Exp.*, 2014(6):P06013, 2014.

Adam Badawy, Emilio Ferrara, and Kristina Lerman. Analyzing the digital traces of political manipulation: The 2016 russian interference twitter campaign. In *2018 IEEE/ACM international conference on advances in social networks analysis and mining (ASONAM)*, pages 258–265. IEEE, 2018.

Albert-László Barabási and Réka Albert. Emergence of scaling in random networks. *science*, 286(5439):509–512, 1999.

Albert-László Barabási and Eric Bonabeau. Scale-free networks. *Scientific american*, 288 (5):60–69, 2003.

Pierre Barbillon, Loïc Schwaller, Stéphane Robin, Andrew Flachs, and Glenn Davis Stone. Epidemiologic network inference. *Statistics and Computing*, 30(1):61–75, 2020.

Claudia Battistin, John Hertz, Joanna Tyrcha, and Yasser Roudi. Belief propagation and replicas for inference and learning in a kinetic ising model with hidden spins. *J. Stat. Mech. Theory Exp.*, 2015(5):P05021, 2015.

Federico Battiston, Vincenzo Nicosia, and Vito Latora. The new challenges of multiplex networks: Measures and models. *The European Physical Journal Special Topics*, 226:401–416, 2017.

Dimitri P Bertsekas. *Constrained optimization and Lagrange multiplier methods*. Academic press, 2014.

Shishir Bharathi, David Kempe, and Mahyar Salek. Competitive influence maximization in social networks. In *International workshop on web and internet economics*, pages 306–311. Springer, 2007.

Stephan Bialonski, Martin Wendler, and Klaus Lehnertz. Unraveling spurious properties of interaction networks with tailored random networks. *PloS one*, 6(8):e22826, 2011.

Guido Bonomi, Nicola Gatti, Fabio Panozzo, and Marcello Restelli. Computing equilibria with two-player zero-sum continuous stochastic games with switching controller. In *Twenty-Sixth AAAI Conference on Artificial Intelligence*, 2012.

Stephen P Borgatti, Martin G Everett, and Jeffrey C Johnson. *Analyzing social networks*. Sage, 2018.

Dan Braha and Marcus AM De Aguiar. Voting contagion: Modeling and analysis of a century of us presidential elections. *PloS one*, 12(5):e0177970, 2017.

Alfredo Braunstein, Alessandro Ingrosso, and Anna Paola Muntoni. Network reconstruction from infection cascades. *Journal of the Royal Society Interface*, 16(151): 20180844, 2019.

Markus Brede, Valerio Restocchi, and Sebastian Stein. Effects of time horizons on influence maximization in the voter dynamics. *Journal of Complex Networks*, 7(3):445–468, 2019a.

Markus Brede, Valerio Restocchi, and Sebastian Stein. Transmission errors and influence maximization in the voter model. *Journal of Statistical Mechanics: Theory and Experiment*, 2019(3):033401, 2019b.

Pierre Brémaud. Non-homogeneous markov chains. In *Markov Chains*, pages 399–422. Springer, 2020.

Tom Britton and Philip D. O'Neill. Bayesian inference for stochastic epidemics in populations with random social structure. *Scandinavian Journal of Statistics*, 29(3):375–390, 2002. ISSN 03036898, 14679469. URL http://www.jstor.org/stable/4616722.

Ivan Brugere, Brian Gallagher, and Tanya Y Berger-Wolf. Network structure inference, a survey: Motivations, methods, and applications. *ACM Computing Surveys (CSUR)*, 51(2):1–39, 2018.

Ceren Budak, Divyakant Agrawal, and Amr El Abbadi. Limiting the spread of misinformation in social networks. In *Proceedings of the 20th international conference on World wide web*, pages 665–674, 2011.

Zhongqi Cai, Markus Brede, and Enrico Gerding. Influence maximization for dynamic allocation in voter dynamics. In Rosa M. Benito, Chantal Cherifi, Hocine Cherifi, Esteban Moro, Luis Mateus Rocha, and Marta Sales-Pardo, editors, *Complex Networks & Their Applications IX*, pages 382–394, Cham, 2021. Springer International Publishing. ISBN 978-3-030-65347-7.

Zhongqi Cai, Enrico Gerding, and Markus Brede. Accelerating opponent strategy inference for voting dynamics on complex networks. In Rosa Maria Benito, Chantal Cherifi, Hocine Cherifi, Esteban Moro, Luis M. Rocha, and Marta Sales-Pardo, editors, *Complex Networks & Their Applications X*, pages 844–856, Cham, 2022a. Springer International Publishing. ISBN 978-3-030-93409-5.

Zhongqi Cai, Enrico Gerding, and Markus Brede. Control meets inference: Using network control to uncover the behaviour of opponents. *Entropy*, 24(5), 2022b. ISSN 1099-4300. . URL https://www.mdpi.com/1099-4300/24/5/640.

Carlo Campajola, Fabrizio Lillo, and Daniele Tantari. Inference of the kinetic ising model with heterogeneous missing data. *Phys. Rev. E*, 99(6):062138, 2019.

Tim Carnes, Chandrashekhar Nagarajan, Stefan M Wild, and Anke Van Zuylen. Maximizing influence in a competitive social network: a follower's perspective. In *Proceedings of the ninth international conference on Electronic commerce*, pages 351–360, 2007.

Jose Casadiego, Mor Nitzan, Sarah Hallerberg, and Marc Timme. Model-free inference of direct network interactions from nonlinear collective dynamics. *Nature communications*, 8(1):1–10, 2017.

Claudio Castellano, Santo Fortunato, and Vittorio Loreto. Statistical physics of social dynamics. *Reviews of modern physics*, 81(2):591, 2009.

Michele Catanzaro, Marián Boguná, and Romualdo Pastor-Satorras. Generation of uncorrelated random scale-free networks. *Physical review e*, 71(2):027103, 2005.

Damon Centola. *How behavior spreads: The science of complex contagions*, volume 3. Princeton University Press Princeton, NJ, 2018.

Andrew Chadwick, Johannes Kaiser, Cristian Vaccari, Daniel Freeman, Sinéad Lambe, Bao S Loe, Samantha Vanderslott, Stephan Lewandowsky, Meghan Conroy, Andrew RN Ross, et al. Online social endorsement and covid-19 vaccine hesitancy in the united kingdom. *Social Media+ Society*, 7(2):20563051211008817, 2021.

Sukankana Chakraborty, Sebastian Stein, Markus Brede, Ananthram Swami, Geeth de Mel, and Valerio Restocchi. Competitive influence maximisation using voting dynamics. In *Proceedings of the 2019 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, pages 978–985, 2019.

David Chandler. Introduction to modern statistical. *Mechanics. Oxford University Press, Oxford, UK*, 5:449, 1987.

Bo-Lun Chen, Wen-Xin Jiang, Yi-Xin Chen, Ling Chen, Rui-Jie Wang, Shuai Han, Jian-Hong Lin, and Yi-Cheng Zhang. Influence blocking maximization on networks: Models, methods and applications. *Physics Reports*, 976:1–54, 2022. ISSN 0370-1573. . URL https://www.sciencedirect.com/science/article/pii/S0370157322001600. Influence Blocking Maximization on Networks: Models, Methods and Applications.

Wei Chen, Yajun Wang, and Siyu Yang. Efficient influence maximization in social networks. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 199–208, 2009.

Wei Chen, Yifei Yuan, and Li Zhang. Scalable influence maximization in social networks under the linear threshold model. In *2010 IEEE international conference on data mining*, pages 88–97. IEEE, 2010.

Wei Chen, Wei Lu, and Ning Zhang. Time-critical influence maximization in social networks with time-delayed diffusion process. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence*, AAAI'12, page 592–598. AAAI Press, 2012.

Wei Chen, Laks VS Lakshmanan, and Carlos Castillo. Information and influence propagation in social networks. *Synthesis Lectures on Data Management*, 5(4):1–177, 2013.

Yu-Zhong Chen and Ying-Cheng Lai. Sparse dynamical boltzmann machine for reconstructing complex networks with binary dynamics. *Physical Review E*, 97(3):032317, 2018.

Mao-Ting Chien, Ssu-Ting Fang, and Yu-Xuan Su. A generalization of gershgorin circles. *Applied and Computational Mathematics*, 15(1):101–111, 2016.

Aaron Clauset, Cosma Rohilla Shalizi, and Mark EJ Newman. Power-law distributions in empirical data. *SIAM review*, 51(4):661–703, 2009.

Peter Clifford and Aidan Sudbury. A model for spatial conflict. *Biometrika*, 60(3):581–588, 1973.

Jesse Clifton and Eric Laber. Q-learning: Theory and applications. *Annual Review of Statistics and Its Application*, 7:279–301, 2020.

Simona Cocco, Stanislas Leibler, and Rémi Monasson. Neuronal couplings between retinal ganglion cells inferred by efficient inverse statistical physics methods. *Proceedings of the National Academy of Sciences*, 106(33):14058–14062, 2009.

Aurélien Decelle, Federico Ricci-Tersenghi, and Pan Zhang. Data quality for the inverse lsing problem. *Journal of Physics A: Mathematical and Theoretical*, 49(38):384001, 2016.

Zoltán Dezső and Albert-László Barabási. Halting viruses in scale-free networks. *Physical Review E*, 65(5):055103, 2002.

Mukeshwar Dhamala, Govindan Rangarajan, and Mingzhou Ding. Analyzing information flow in brain networks with nonparametric granger causality. *Neuroimage*, 41 (2):354–362, 2008.

Laurence Charles Ward Dixon. The choice of step length, a crucial factor in the performance of variable metric algorithms. *Numerical methods for non-linear optimization*, pages 149–170, 1972.

Pedro Domingos and Matt Richardson. Mining the network value of customers. In *Proceedings of the seventh ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 57–66, 2001.

Jonathan F Donges, Yong Zou, Norbert Marwan, and Jürgen Kurths. The backbone of the climate network. *EPL (Europhysics Letters)*, 87(4):48007, 2009.

Sergey N Dorogovtsev and Jose FF Mendes. Evolution of networks. *Advances in physics*, 51(4):1079–1187, 2002.

Sergey N Dorogovtsev, Alexander V Goltsev, and José FF Mendes. Critical phenomena in complex networks. *Reviews of Modern Physics*, 80(4):1275, 2008.

Benjamin Dunn and Yasser Roudi. Learning and inference in a nonequilibrium ising model with hidden nodes. *Phys. Rev. E*, 87(2):022127, 2013.

Imme Ebert-Uphoff and Yi Deng. Causal discovery for climate research using graphical models. *Journal of Climate*, 25(17):5648–5665, 2012.

Bradley Efron and David V Hinkley. Assessing the accuracy of the maximum likelihood estimator: Observed versus expected fisher information. *Biometrika*, 65(3):457–483, 1978. .

Paul Erdős, Alfréd Rényi, et al. On the evolution of random graphs. *Publ. math. inst. hung. acad. sci*, 5(1):17–60, 1960.

Paul Erdős and Alfréd Rényi. On random graphs i. *Publ. math. debrecen*, 6(290-297):18, 1959.

Eyal Even-Dar and Asaf Shapira. A note on maximizing the spread of influence in social networks. In *International Workshop on Web and Internet Economics*, pages 281–286. Springer, 2007.

Eyal Even-Dar and Asaf Shapira. A note on maximizing the spread of influence in social networks. *Information Processing Letters*, 111(4):184–187, 2011.

David Fajardo and Lauren M Gardner. Inferring contagion patterns in social contact networks with limited infection data. *Networks and spatial economics*, 13(4):399–426, 2013.

Arastoo Fazeli and Ali Jadbabaie. Game theoretic analysis of a strategic model of competitive contagion and product adoption in social networks. In *2012 IEEE 51st IEEE Conference on Decision and Control (CDC)*, pages 74–79. IEEE, 2012.

Juan Fernández-Gracia, Krzysztof Suchecki, José J Ramasco, Maxi San Miguel, and Víctor M Eguíluz. Is the voter model a model for voters? *Physical review letters*, 112(15):158701, 2014.

M. Ferreira, Marina Andrade, Maria Cristina Peixoto Matos, J. Filipe, and M. Coelho. Minimax theorem and nash equilibrium. *The International Journal of Latest Trends in Finance and Economic Sciences*, 2:36–40, 2012.

Serge Galam. Sociophysics: A review of galam models. *International Journal of Modern Physics C*, 19(03):409–440, 2008.

Serge Galam and Marco Alberto Javarone. Modeling radicalization phenomena in heterogeneous populations. *PloS one*, 11(5):e0155407, 2016.

Serge Galam, Yuval Gefen, and Yonathan Shapir. Sociophysics: A new approach of sociological collective behaviour. i. mean-behaviour description of a strike. *Journal of Mathematical Sociology*, 9(1):1–13, 1982.

Xu Gao, Daniel Gillen, and Hernando Ombao. Fisher information matrix of binary time series. *Metron*, 76:287–304, 2018.

Michael T Gastner. The ising chain constrained to an even or odd number of positive spins. *Journal of Statistical Mechanics: Theory and Experiment*, 2015(3):P03004, 2015.

Roy J Glauber. Time-dependent statistics of the ising model. *J. Math. Phys.*, 4(2):294–307, 1963.

Manuel Gomez-Rodriguez, Jure Leskovec, and Andreas Krause. Inferring networks of diffusion and influence. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 5(4):1–37, 2012.

Amit Goyal, Wei Lu, and Laks VS Lakshmanan. Celf++ optimizing the greedy algorithm for influence maximization in social networks. In *Proceedings of the 20th international conference companion on World wide web*, pages 47–48, 2011.

Amit Goyal, Francesco Bonchi, Laks VS Lakshmanan, and Suresh Venkatasubramanian. On minimizing budget and time in influence propagation over social networks. *Social network analysis and mining*, 3(2):179–192, 2013.

Sanjeev Goyal, Hoda Heidari, and Michael Kearns. Competitive contagion in networks. *Games and Economic Behavior*, 113:58–79, 2019.

Caitlin Gray, Lewis Mitchell, and Matthew Roughan. Bayesian inference of network structure from information cascades. *IEEE Transactions on Signal and Information Processing over Networks*, 6:371–381, 2020.

Adrien Guille, Hakim Hacid, Cecile Favre, and Djamel A Zighed. Information diffusion in online social networks: A survey. *ACM Sigmod Record*, 42(2):17–28, 2013.

Roger Guimera, Leon Danon, Albert Diaz-Guilera, Francesc Giralt, and Alex Arenas. Self-similar community structure in a network of human interactions. *Physical review E*, 68(6):065103, 2003.

Ce Guo and Wayne Luk. Accelerating maximum likelihood estimation for hawkes point processes. In *2013 23rd International Conference on Field programmable Logic and Applications*, pages 1–6. IEEE, 2013.

Shaoyan Guo, Huifu Xu, and Liwei Zhang. Existence and approximation of continuous bayesian nash equilibria in games with continuous type and action spaces. *SIAM Journal on Optimization*, 31(4):2481–2507, 2021.

Anne-Claire Haury, Fantine Mordelet, Paola Vera-Licona, and Jean-Philippe Vert. Tigress: trustful inference of gene regulation using stability selection. *BMC systems biology*, 6(1):1–17, 2012.

Xinran He and Yan Liu. Not enough data? joint inferring multiple diffusion networks via network generation priors. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*, pages 465–474, 2017.

Rainer Hegselmann, Stefan König, Sascha Kurz, Christoph Niemann, and Jörg Rambau. Optimal opinion control: The campaign problem. *arXiv preprint arXiv:1410.8419*, 2014.

Danh-Tai Hoang, Juyong Song, Vipul Periwal, and Junghyo Jo. Network inference in stochastic systems from neurons to currencies: Improved performance at small sample size. *Physical Review E*, 99(2):023311, 2019.

Petter Holme and Beom Jun Kim. Growing scale-free networks with tunable clustering. *Physical review E*, 65(2):026107, 2002.

Roger A Horn. The hadamard product. In *Proc. Symp. Appl. Math*, volume 40, pages 87–169, 1990.

Haibo Hu. Competing opinion diffusion on social networks. *Royal Society open science*, 4(11):171160, 2017.

Feng Huang and Han-Shuang Chen. An improved heterogeneous mean-field theory for the ising model on complex networks. *Communications in Theoretical Physics*, 71 (12):1475, 2019.

Ari Juels and Martin Wattenberg. Stochastic hillclimbing as a baseline method for evaluating genetic algorithms. *Advances in Neural Information Processing Systems*, 8, 1995.

WB Vasantha Kandasamy, Florentin Smaradache, and K Ilanthenral. *Pseudo Lattice graphs and their applications to fuzzy and neutrosophic models*. Infinite Study, 2014.

HJ Kappen and JJ Spanjers. Mean field theory for asymmetric neural networks. *Physical Review E*, 61(5):5658–5663, 2000. .

Jaya Kawale, Stefan Liess, Arjun Kumar, Michael Steinbach, Peter Snyder, Vipin Kumar, Auroop R Ganguly, Nagiza F Samatova, and Fredrick Semazzi. A graph-based approach to find teleconnections in climate data. *Statistical Analysis and Data Mining: The ASA Data Science Journal*, 6(3):158–179, 2013.

David Kempe, Jon Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 137–146, 2003.

Mehrdad Agha Mohammad Ali Kermani, Seyed Farshad Fatemi Ardestani, Alireza Aliahmadi, and Farnaz Barzinpour. A novel game theoretic approach for modeling competitive information diffusion in social networks with heterogeneous nodes. *Physica A: statistical mechanics and its applications*, 466:570–582, 2017.

Jeong Han Kim and Van H Vu. Generating random regular graphs. In *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, pages 213–222, 2003.

Seung-Jean Kim, Kwangmoo Koh, Michael Lustig, Stephen Boyd, and Dimitry Gorinevsky. An interior-point method for large-scale ℓ_1-regularized least squares. *IEEE journal of selected topics in signal processing*, 1(4):606–617, 2007.

David Knoke and Song Yang. *Social network analysis*. SAGE publications, 2019.

Eric D Kolaczyk. Network topology inference. In *Statistical Analysis of Network Data*, pages 1–48. Springer, 2009.

Holly Korda and Zena Itani. Harnessing social media for health promotion and behavior change. *Health promotion practice*, 14(1):15–23, 2013.

Chris J Kuhlman, VS Anil Kumar, and SS Ravi. Controlling opinion propagation in online networks. *Computer Networks*, 57(10):2121–2132, 2013.

Volodymyr Kuleshov and Doina Precup. Algorithms for multi-armed bandit problems. *arXiv preprint arXiv:1402.6028*, 2014.

Jean-Philippe Lachaux, Antoine Lutz, David Rudrauf, Diego Cosmelli, Michel Le Van Quyen, Jacques Martinerie, and Francisco Varela. Estimating the time-course of coherence between single-trial brain signals: an introduction to wavelet coherence. *Neurophysiologie Clinique/Clinical Neurophysiology*, 32(3):157–174, 2002.

Hyun Keun Lee, Pyoung-Seop Shim, and Jae Dong Noh. Epidemic threshold of the susceptible-infected-susceptible model on complex networks. *Physical Review E*, 87 (6):062812, 2013.

Sangwon Lee, Vipul Periwal, and Junghyo Jo. Inference of stochastic time series with missing data. *Phys. Rev. E*, 104(2):024119, 2021.

Jure Leskovec, Andreas Krause, Carlos Guestrin, Christos Faloutsos, Jeanne VanBriesen, and Natalie Glance. Cost-effective outbreak detection in networks. In *Proceedings of the 13th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 420–429, 2007.

Jure Leskovec, Lars Backstrom, and Jon Kleinberg. Meme-tracking and the dynamics of the news cycle. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 497–506, 2009.

Ted G Lewis. *Network science: Theory and applications*. John Wiley & Sons, 2011.

Timothy R Lezon, Jayanth R Banavar, Marek Cieplak, Amos Maritan, and Nina V Fedoroff. Using the principle of entropy maximization to infer genetic interaction networks from gene expression patterns. *Proceedings of the National Academy of Sciences*, 103(50):19033–19038, 2006.

Jingwen Li, Zhesi Shen, Wen-Xu Wang, Celso Grebogi, and Ying-Cheng Lai. Universal data-based method for reconstructing complex networks with binary-state dynamics. *Physical Review E*, 95(3):032303, 2017.

Yuchen Li, Ju Fan, Yanhao Wang, and Kian-Lee Tan. Influence maximization on social graphs: A survey. *IEEE Transactions on Knowledge and Data Engineering*, 30(10):1852–1872, 2018.

Bo Liu, Gao Cong, Dong Xu, and Yifeng Zeng. Time constrained influence maximization in social networks. In *2012 IEEE 12th international conference on data mining*, pages 439–448. IEEE, 2012.

Jing Liu, Mingxing Zhou, Shuai Wang, and Penghui Liu. A comparative study of network robustness measures. *Frontiers of Computer Science*, 11:568–584, 2017.

Shihuan Liu, Lei Ying, and Srinivas Shakkottai. Influence maximization in social networks: An ising-model-based approach. In *2010 48th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 570–576. IEEE, 2010.

Sahil Loomba, Alexandre de Figueiredo, Simon J Piatek, Kristen de Graaf, and Heidi J Larson. Measuring the impact of covid-19 vaccine misinformation on vaccination intent in the uk and usa. *Nature human behaviour*, 5(3):337–348, 2021.

Haishu Lu. On the existence of pure-strategy nash equilibrium. *Economics Letters*, 94 (3):459–462, 2007.

Alexander Ly, Maarten Marsman, Josine Verhagen, Raoul PPP Grasman, and Eric-Jan Wagenmakers. A tutorial on fisher information. *Journal of Mathematical Psychology*, 80:40–55, 2017.

Christopher Lynn and Daniel D Lee. Maximizing influence in an ising network: A mean-field optimal solution. In *Advances in Neural Information Processing Systems*, pages 2495–2503, 2016.

Antonia Maria Masucci and Alonso Silva. Strategic resource allocation for competitive influence in social networks. In *2014 52nd Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 951–958. IEEE, 2014.

Naoki Masuda. Opinion control in complex networks. *New Journal of Physics*, 17(3): 033031, 2015.

Michael McFaul and Bronte Kass. Understanding putin's intentions and actions in the 2016 us presidential election. *SECURING AMERICAN ELECTIONS*, page 1, 2019.

Lauren Ancel Meyers, Babak Pourbohloul, Mark EJ Newman, Danuta M Skowronski, and Robert C Brunham. Network theory and sars: predicting outbreak diversity. *Journal of theoretical biology*, 232(1):71–81, 2005.

Marc Mézard and J Sakellariou. Exact mean-field inference in asymmetric kinetic ising systems. *J. Stat. Mech. Theory Exp.*, 2011(07):L07001, 2011.

Byungjoon Min and Maxi San Miguel. Competing contagion processes: Complex contagion triggered by simple contagion. *Scientific reports*, 8(1):1–8, 2018.

Mauro Mobilia, Anna Petersen, and Sidney Redner. On the role of zealotry in the voter model. *Journal of Statistical Mechanics: Theory and Experiment*, 2007(08):P08029, 2007.

Faruck Morcos, Andrea Pagnani, Bryan Lunt, Arianna Bertolino, Debora S Marks, Chris Sander, Riccardo Zecchina, José N Onuchic, Terence Hwa, and Martin Weigt. Direct-coupling analysis of residue coevolution captures native contacts across many protein families. *Proceedings of the National Academy of Sciences*, 108(49):E1293–E1301, 2011.

Guillermo Romero Moreno, Edoardo Manino, Long Tran-Thanh, and Markus Brede. Zealotry and influence maximization in the voter model: when to target partial zealots? In *Complex Networks XI*, pages 107–118. Springer, 2020.

Flaviano Morone and Hernán A Makse. Influence maximization in complex networks through optimal percolation. *Nature*, 524(7563):65–68, 2015.

Seth A Myers and Jure Leskovec. On the convexity of latent social network inference. *arXiv preprint arXiv:1010.5504*, 2010.

Seth A Myers, Chenguang Zhu, and Jure Leskovec. Information diffusion and external influence in networks. In *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 33–41, 2012.

In Jae Myung. Tutorial on maximum likelihood estimation. *Journal of mathematical Psychology*, 47(1):90–100, 2003.

Mark Newman. *Networks*. Oxford university press, 2018.

Mark EJ Newman. Mixing patterns in networks. *Physical review E*, 67(2):026126, 2003a.

Mark EJ Newman. The structure and function of complex networks. *SIAM review*, 45 (2):167–256, 2003b.

Mark EJ Newman. Power laws, pareto distributions and zipf's law. *Contemporary physics*, 46(5):323–351, 2005.

Mark EJ Newman. Finding community structure in networks using the eigenvectors of matrices. *Physical review E*, 74(3):036104, 2006.

Mark EJ Newman, Steven H Strogatz, and Duncan J Watts. Random graphs with arbitrary degree distributions and their applications. *Physical review E*, 64(2):026118, 2001.

H Chau Nguyen, Riccardo Zecchina, and Johannes Berg. Inverse statistical problems: from the inverse ising problem to data science. *Advances in Physics*, 66(3):197–261, 2017.

Nam P Nguyen, Guanhua Yan, My T Thai, and Stephan Eidenbenz. Containment of misinformation spread in online social networks. In *Proceedings of the 4th Annual ACM Web Science Conference*, pages 213–222, 2012.

Hossein Noorazar. Recent advances in opinion propagation dynamics: A 2020 survey. *The European Physical Journal Plus*, 135:1–20, 2020.

Jingzhi Pan, Fei Jiang, and Jin Xu. Influence maximization in social networks based on non-backtracking random walk. In *2016 IEEE First International Conference on Data Science in Cyberspace (DSC)*, pages 260–267. IEEE, 2016.

Pietro Panzarasa, Tore Opsahl, and Kathleen M Carley. Patterns and dynamics of users' behavior and interaction: Network analysis of an online community. *Journal of the American Society for Information Science and Technology*, 60(5):911–932, 2009.

WH Press, SA Teukolsky, WT Vetterling, and BP Flannery. Section 10.11. linear programming: interior-point methods. In *Numerical Recipes: The Art of Scientific Computing*. Cambridge Univ. Press, 2007.

William H Press, Saul A Teukolsky, William T Vetterling, and Brian P Flannery. Numerical recipes in c, 1988.

Derek J De Solla Price. Networks of scientific papers. *Science*, pages 510–515, 1965.

Neha Puri, Eric A Coomes, Hourmazd Haghbayan, and Keith Gunaratne. Social media and vaccine hesitancy: new updates for the era of covid-19 and globalized infectious diseases. *Human Vaccines & Immunotherapeutics*, pages 1–8, 2020.

Maryam Ramezani, Hamid R Rabiee, Maryam Tahani, and Arezoo Rajabi. Dani: A fast diffusion aware network inference algorithm. *arXiv preprint arXiv:1706.00941*, 2017.

Marlon Ramos, Jia Shao, Saulo DS Reis, Celia Anteneodo, José S Andrade, Shlomo Havlin, and Hernán A Makse. How does public opinion become extreme? *Scientific reports*, 5:10032, 2015.

Erzsébet Ravasz and Albert-László Barabási. Hierarchical organization in complex networks. *Physical review E*, 67(2):026112, 2003.

Geofakta Razali, Masfiatun Nikmah, I Nyoman Tri Sutaguna, PA Andiena Nindya Putri, and Muhammad Yusuf. The influence of viral marketing and social media marketing on instagram adds purchase decisions. *CEMERLANG: Jurnal Manajemen dan Ekonomi Bisnis*, 3(2):75–86, 2023.

Sidney Redner. Reality-inspired voter models: A mini-review. *Comptes Rendus Physique*, 20(4):275–292, 2019.

Matthew Richardson and Pedro Domingos. Mining knowledge-sharing sites for viral marketing. In *Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 61–70, 2002.

Manuel Gomez Rodriguez and Bernhard Schölkopf. Submodular inference of diffusion networks from multiple trees. *arXiv preprint arXiv:1205.1671*, 2012.

Guillermo Romero Moreno, Edoardo Manino, Long Tran-Thanh, and Markus Brede. Zealotry and influence maximization in the voter model: when to target partial zealots? In *Complex Networks XI: Proceedings of the 11th Conference on Complex Networks CompleNet 2020*, pages 107–118. Springer, 2020.

Guillermo Romero Moreno, Sukankana Chakraborty, and Markus Brede. Shadowing and shielding: Effective heuristics for continuous influence maximisation in the voting dynamics. *PLOS ONE*, 16(6):1–21, 06 2021a. . URL https://doi.org/10.1371/journal.pone.0252515.

Guillermo Romero Moreno, Sukankana Chakraborty, and Markus Brede. Shadowing and shielding: Effective heuristics for continuous influence maximisation in the voting dynamics. *Plos one*, 16(6):e0252515, 2021b.

Maria Rosário Grossinho and Stepan Agop Tersian. *An Introduction to Minimax Theorems and Their Applications to Differential Equations*, volume 52 of *Nonconvex Optimization and Its Applications*. Springer US, 2001. ISBN 9781441948496. . URL http://link.springer.com/10.1007/978-1-4757-3308-2.

Yasser Roudi and John Hertz. Mean field theory for nonequilibrium network reconstruction. *Physical Review Letters*, 106(4):048702, 2011. .

Seung-Woo Son, Hawoong Jeong, and Hyunsuk Hong. Relaxation of synchronization on complex networks. *Physical Review E*, 78(1):016106, 2008.

Vishal Sood and Sidney Redner. Voter model on heterogeneous graphs. *Physical review letters*, 94(17):178701, 2005.

Didier Sornette. Physics and financial economics (1776–2014): puzzles, ising and agent-based models. *Reports on Progress in Physics*, 77(6):062001, 2014. .

Qi Tan, Yang Liu, and Jiming Liu. Motif-aware diffusion network inference. *International Journal of Data Science and Analytics*, 9(4):375–387, 2020.

Paul A Tess. The role of social media in higher education classes (real and virtual)–a literature review. *Computers in human behavior*, 29(5):A60–A68, 2013.

Guangmo Tong, Ruiqi Wang, Zheng Dong, and Xiang Li. Time-constrained adaptive influence maximization. *IEEE Transactions on Computational Social Systems*, 2020.

Vasileios Tzoumas, Christos Amanatidis, and Evangelos Markakis. A game-theoretic analysis of a competitive diffusion process over social networks. In *International Workshop on Internet and Network Economics*, pages 1–14. Springer, 2012.

SM Minhaz Ud-Dean and Rudiyanto Gunawan. Optimal design of gene knockout experiments for gene regulatory network inference. *Bioinformatics*, 32(6):875–883, 2016.

Sander van der Linden, Jon Roozenbeek, and Josh Compton. Inoculating against fake news about covid-19. *Frontiers in Psychology*, 11:2928, 2020.

František Váša and Bratislav Mišić. Null models in network neuroscience. *Nature Reviews Neuroscience*, 23(8):493–504, 2022.

Vítor V Vasconcelos, Simon A Levin, and Flávio L Pinheiro. Consensus and polarization in competing complex contagion processes. *Journal of the Royal Society Interface*, 16(155):20190196, 2019.

Antoine Vendeville, Benjamin Guedj, and Shi Zhou. Forecasting elections results via the voter model with stubborn nodes. *Applied Network Science*, 6:1–13, 2021.

Duncan J Watts and Steven H Strogatz. Collective dynamics of 'small-world'networks. *nature*, 393(6684):440–442, 1998.

Bryan Wilder and Yevgeniy Vorobeychik. Controlling elections through social influence. In *Proceedings of the 17th international conference on autonomous agents and multiagent systems*, pages 265–273. International Foundation for Autonomous Agents and Multiagent Systems, 2018.

Tom Willaert, Paul Van Eecke, Katrien Beuls, and Luc Steels. Building social media observatories for monitoring online opinion dynamics. *Social Media+ Society*, 6(2): 2056305119898778, 2020.

Fei Xiong and Yun Liu. Opinion formation on social media: an empirical approach. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 24(1):013130, 2014.

Weihao Xuan, Ruijie Ren, Philip E Paré, Mengbin Ye, Sebastian Ruf, and Ji Liu. On a network sis model with opinion dynamics. *IFAC-PapersOnLine*, 53(2):2582–2587, 2020.

Amulya Yadav, Hau Chan, Albert Xin Jiang, Haifeng Xu, Eric Rice, and Milind Tambe. Using social networks to aid homeless shelters: Dynamic influence maximization under uncertainty. In *AAMAS*, volume 16, pages 740–748, 2016.

Yang Yang, Tingjin Luo, Zhoujun Li, Xiaoming Zhang, and Philip S Yu. A robust method for inferring network structures. *Scientific Reports*, 7(1):1–12, 2017.

Ercan Yildiz, Asuman Ozdaglar, Daron Acemoglu, Amin Saberi, and Anna Scaglione. Binary opinion dynamics with stubborn agents. *ACM Transactions on Economics and Computation (TEAC)*, 1(4):1–30, 2013.

Jean-Gabriel Young, George T Cantwell, and MEJ Newman. Bayesian inference of network structure from unreliable data. *Journal of Complex Networks*, 8(6):cnaa046, 2020.

Hong-Li Zeng, Erik Aurell, Mikko Alava, and Hamed Mahmoudi. Network inference using asynchronously updated kinetic ising model. *Phys. Rev. E*, 83(4):041135, 2011.

Hai-Feng Zhang, Fang Xu, Zhong-Kui Bao, and Chuang Ma. Reconstructing of networks with binary-state dynamics via generalized statistical inference. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 66(4):1608–1619, 2018.

Pan Zhang. Inference of kinetic ising model on sparse graphs. *J. Stat. Phys.*, 148:502–512, 2012.