

# Causal Influences over Social Learning Networks

Mert Kayaalp\*

Dalle Molle Institute for Artificial Intelligence (IDSIA USI-SUPSI)  
CH-6900, Lugano, Switzerland

{MERT.KAYAALP}@IDSIA.CH

Ali H. Sayed

École Polytechnique Fédérale de Lausanne (EPFL)  
CH-1015, Lausanne, Switzerland

{ALI.SAYED}@EPFL.CH

**Editor:** Christophe Giraud

## Abstract

This paper investigates causal influences between agents linked by a social graph and interacting over time. In particular, the work examines the dynamics of social learning models and distributed decision-making protocols, and derives expressions that reveal the causal relations between pairs of agents and explain the flow of influence over the network. The results turn out to be dependent on the graph topology and the level of information that each agent has about the inference problem they are trying to solve. Using these conclusions, the paper proposes an algorithm to rank the overall influence between agents to discover highly influential agents. It also provides a method to learn the necessary model parameters from raw observational data. The results and the proposed algorithm are illustrated by considering both synthetic data and real social media data.

**Keywords:** social influence, causal effect, diffusion of influence, spillover effects, opinion dynamics, causal ranking, distributed decision making

## 1. Introduction

Identifying influential agents in a network is a crucial problem with wide-ranging applications, such as detecting propaganda-sharing accounts (Smith et al., 2021) or selecting individuals to advertise to (Lagrée et al., 2018). However, over any social network, information usually spreads and mixes, leading to ripple effects that make the discovery of influence rather challenging. For example, the information leaving a source agent may be altered and combined with comments/data from other agents along the path until it reaches its destination. Most prior works measure influence through some network topology-based properties such as the eigenvector centrality of an agent (Dablander and Hinne, 2019), or through some descriptive importance factor depending on the problem at hand (Kempe et al., 2003; Banerjee et al., 2013; Shumovskaia et al., 2023). In comparison, this work treats influence as a *causal* quantity and approaches it from the perspective of structural causal models (Wright, 1934; Pearl, 2009). More specifically, influence will be defined as the change in behavior of the network when interventions occur at individual agents. This is a useful method to discard spurious and non-causal associations, unlike other methods based, for example, on the use of Granger causality (Granger, 1969). Obviously, conducting interventional experiments

---

\*. Corresponding author. Part of this work was conducted while the author was at EPFL.

may not be always feasible over real world social networks. However, with the help of appropriate representative models, one can rely on the use of raw observational data (Pearl and Mackenzie, 2018).

To that end, this paper considers some common mathematical models for social learning over graphs (DeGroot, 1974; Acemoglu et al., 2011; Jadbabaie et al., 2012; Chamley et al., 2013; Nedić et al., 2017). These models are widely used to model opinion formation from a behavioral perspective (Mossel and Tamuz, 2017; Krishnamurthy, 2022), and their generalization to real-world scenarios is verified by field experiments (Chandrasekhar et al., 2020). From an engineering perspective, the models can be utilized to design distributed decision-making systems where a group of agents (e.g., robots, sensors) collaboratively make decisions about the environment (Djurić and Wang, 2012). However, most of the research on social learning models has focused on analyzing their convergence and error behavior under different assumptions, such as verifying whether truth learning occurs in the presence of malicious agents supplying misinformation (Hare et al., 2019; Li et al., 2019), or whether the models can track drifts in the underlying hypothesis (Bordignon et al., 2021). In contrast, this study utilizes these opinion formation models to examine *causal effects* over social graphs. To the best of our knowledge, this appears to be the first study to do so by studying the diffusion of influence over space and time from a causal inference perspective.

## 1.1 Contributions

- Our work proposes a novel causal impact criterion for dynamic social networks in Sec. 4. It quantifies how much an agent  $m$  affects another agent  $k$  for each agent pair  $(m, k)$ .
- We derive closed-form expressions for these impact factors for two social network models: non-Bayesian social learning (Jadbabaie et al., 2012) and adaptive social learning (Bordignon et al., 2021) in Secs. 5.1 and 5.2, respectively. The non-Bayesian framework (Jadbabaie et al., 2012) models network dynamics through a (log)-linear autoregressive process, stated in future expression (4), and incorporates in-network interactions as well as possibly hidden out-of-network factors. The adaptive framework (Bordignon et al., 2021) additionally introduces temporal weightings that emphasize more recent interactions, as stated in (11). Beyond general causal influence expressions, we also analyze useful special cases to illustrate how the causal effects depend on the model parameters.
- Analyzing the influence for every pair  $(m, k)$  results in a network causality matrix, which offers various options to rank agents for their overall influence. In Sec. 6, we propose a particular algorithm, *CausalRank*, to rank agents based on their overall influence on the network, while accounting for the fact that more impactful agents should be weighted more. This algorithm leverages the eigenvector centrality of the bipartite causal relations matrix introduced in Sec. 4.
- Additionally, we introduce a graph causality learning (GCL) algorithm in Sec. 7, designed to estimate causal influences using observational data comprised of social interactions and the graph topology underlying the agents.




Figure 1: Some inherent challenges in estimating the causal influence of one agent on another, represented by the red and blue agents, respectively. *(Left)*: There can be confounding factors that influence both parties and induce non-causal correlation. *(Middle)*: Relationships within social networks are often bidirectional that are not instantaneous, but rather spread out over time. *(Right)*: Information transmitted from the source is processed and potentially changed along the way.

- In Sec. 8, we illustrate our results with synthetic data. In Sec. 9, we apply these findings to real social media data, demonstrating their practical usefulness.

## 1.2 Challenges for estimating influence over social networks

Influence estimation over social networks faces several challenges.

**Confounding factors.** Understanding influence requires disentangling correlation from causation. For example, over a social network, it is often observed that individuals who are connected tend to have similar (correlated) opinions. However, this does not necessarily imply a causal relationship and there can also exist confounding factors. For instance, individuals may obtain information from similar external sources (say, the same TV channels), or they may be connected to others who share similar preferences (a.k.a. homophily) — see left panel in Fig. 1. Similar issues arise in distributed decision-making systems, such as networks of wireless sensors or robots. Devices that communicate with each other are often in spatial proximity, leading to correlated observations. Therefore, accounting for these confounding factors is crucial for discovering *true causal* relationships.

**Temporal dynamics.** Social influence is not a one-time occurrence but rather a continuous process that unfolds over time. Therefore, we adopt a time-series approach to capture this dynamic nature. Unlike directed acyclic graph based models in the literature, we accommodate both cyclic networks and bidirectional links to capture feedback mechanisms — see middle panel in Fig. 1. Doing so is necessary in order to discover the propagation of influence over both space and time.

**Mixing and diffusion of information.** When examining the influence of an agent  $m$  on another agent  $k$ , it is essential to acknowledge that information leaving agent  $m$  can undergo alterations and become intertwined with the opinions of other agents in the network before reaching agent  $k$  — see right panel in Fig. 1. This phenomenon introduces additional

complexity to the study of influence propagation over graphs.

In this work, the social network models we consider allow us to treat these challenges together in some detail. Specifically, the expressions we derive within a rigorous causal theoretical framework quantify how the instantaneous and direct social effects diffuse over time and network. In doing so, as we seek to understand the total and overall causal effects, we take into account potential hidden confounding factors.

## 2. Related Work

**Influence quantification.** Most existing works in the literature quantify influence by relying on network topology-based measures (such as degree centrality) (Dablander and Hinne, 2019) or by examining task-specific importance factors (Kempe et al., 2003; Shah and Zaman, 2011; Banerjee et al., 2013; Shumovskaia et al., 2023). However, these measures often lack a causal interpretation and they may fail to eliminate non-causal correlation-inducing factors. Another line of work use randomized experiments in order to identify causal relations, which, in principle, can discard latent confounders (Aral and Walker, 2012; Bond et al., 2012). Unfortunately, conducting controlled experiments is impractical for many real-world scenarios. In this work, we study social influence from a structural causal model framework (Wright, 1934; Pearl, 2009), which will enable us to circumvent these issues with the help of useful models for interactions over social graphs.

**Causal influence.** While a few previous works have treated social influence as a causal effect, they have some limitations. For instance, (Sridhar et al., 2022) only considered a single time step of interaction and did not address the time series formulation, which is critical in capturing information diffusing over time. Likewise, (Soni et al., 2019) utilized Granger causality (Granger, 1969), which is a predictive tool that only depicts the natural behavior of the system and is not robust against latent confounders (Eichler and Didelez, 2010). In contrast, our approach is based on the causality principle that manipulating the causes, while keeping other variables constant, changes the effect (Woodward, 2004; Pearl, 2009). This idea was first applied to general time series in (Eichler and Didelez, 2010). Here, we apply it to a time series of variables over a graph consisting of interconnected agents.

**Networked time series.** Some works such as (Lee et al., 2023; Smith et al., 2018, 2021) relied on the Rubin’s potential outcomes framework (Rubin, 1974). Specifically, (Smith et al., 2018, 2021) analyzed the causal impact of each agent by *directly* modeling the response of other agents under a binary treatment (i.e., intervention) on a particular agent. Our work is complementary but different in the sense that we apply the procedure of hypothetical interventions on graphical models (Pearl, 2009) to the canonical models of social learning networks (Jadbabaie et al., 2012; Nedić et al., 2017; Bordignon et al., 2021). While doing so, we derive the expressions for the causal impact factors by analyzing the step-by-step dynamics of the social networks, and arrive at results that clarify how causal influences propagate over a graph.

**Social learning.** The literature on mathematical models for social learning is rich — see the surveys (Krishnamurthy, 2022; Mossel and Tamuz, 2017) and the papers (Jadbabaie et al., 2012; Bordignon et al., 2021; Nedić et al., 2017). In this study, we focus on the DeGroot

type of networked consensus modeling (DeGroot, 1974; Golub and Jackson, 2010), which has been widely applied in social network research. In particular, we examine the non-Bayesian extensions studied by (Jadbabaie et al., 2012; Molavi et al., 2018), which allow agents to interact with the environment and to exchange information within the network. These models and their extension to adaptive agents (Bordignon et al., 2021) are briefly reviewed in Sec. 3. Another related line of work in social learning considers stubborn agents (Acemoglu et al., 2010; Ghaderi and Srikant, 2014; Vial and Subramanian, 2022; Hunter and Zaman, 2022). This literature is primarily concerned with forward modeling. It considers questions like how agents with fixed opinions (e.g., misinformation spreading agents) affect collective behavior, including whether consensus is disrupted (Acemoglu et al., 2010), what equilibrium or disagreement patterns emerge (Ghaderi and Srikant, 2014), how quickly beliefs converge (Vial and Subramanian, 2022), and how outcomes depend on the agents’ locations over the network (Hunter and Zaman, 2022). Our objective is different in that we study hypothetical interventions in which an agent is assigned an arbitrary belief in order to identify that agent’s causal influence on the rest of the network. This has fundamentally different goals and leads to distinct formulas and conclusions.

**Network interference.** At the intersection of causal inference and networks, there also exists a body of work focused on network interference (Toulis and Kao, 2013; Sussman and Airolidi, 2017; Agarwal et al., 2022). These studies examine causal inference when interventions on certain agents can affect others within the network, and aim to design randomized experiments that take this into account. However, they typically assume that only immediate neighbors of an agent can impact an individual’s response. Although our problem setting differs, our analytical contributions may still prove useful for this area of research, since we discover how influence and interference propagate throughout the network. For example, our findings allow us to quantify how an agent’s influence diminishes with increasing distance from other agents in the graph.

**Notation.** Random variables are denoted in bold, e.g.,  $\mathbf{x}$ . A sequence of random variables  $\{\mathbf{x}_i\}$  over index  $i$  converging to a random variable  $\mathbf{x}$  in distribution is denoted by  $\mathbf{x}_i \xrightarrow{\text{dist.}} \mathbf{x}$ . Almost sure convergence is denoted by  $\mathbf{x}_i \xrightarrow{\text{a.s.}} \mathbf{x}$ . The  $k$ -th entry of a vector  $v$  is denoted by  $v_k$  or  $[v]_k$ . We denote the all-ones vector of dimension  $K \times 1$  by  $\mathbf{1}_K$ . For distributions  $p_1$  and  $p_2$ ,  $D_{\text{KL}}(p_1||p_2)$  denotes their KL divergence.

### 3. Social Learning

#### 3.1 Problem setting

We consider a network of  $K$  agents that are trying to infer the hypothesis that best explains their observations about the world. More formally, agents are trying to learn the true state of nature  $\theta^\circ \in \Theta$  from a finite set of  $H$  hypotheses,  $\Theta = \{\theta_1, \theta_2, \dots, \theta_H\}$ . At each time instant  $i$ , each agent  $k$  receives a personal and partially informative observation  $\xi_{k,i}$ , which encapsulates all out-of-network information available to  $k$  at time  $i$  and is distributed according to some marginal likelihood function  $L_k(\xi|\theta^\circ)$  dependent on the true state. The true state of nature can be viewed as an underlying environmental or contextual variable that governs the data-generating process of the observations. In engineering settings, it may also capture a distributional regime or shift. As a recent example, the state of nature may encode

whether there is an “AI bubble”. People then gather evidence from their local environments, for instance, news reportings, company revenues and spending, user adoption, or their own experience with AI tools, and update their beliefs through interactions with their peers.

The observations  $\xi_{k,i}$  are assumed i.i.d. over time, and the likelihood function  $L_k(\xi|\theta^\circ)$  is assumed to be time-invariant. From a behavioral perspective of social networks, the observations often represent external sources of information. Fixed distribution assumptions correspond to these sources providing relatively stable information over time. For instance, a TV channel may consistently support Candidate A over Candidate B for an election, justifying the assumption of time-invariant likelihood functions. Therefore, studying causal influences under the fixed distribution model is useful in many situations. As the analysis in the paper shows, even in this case, the arguments are not straightforward and there are many useful insights that can be gained from such studies. We acknowledge that in some applications the distributions can drift with time. We leave the study of this scenario to future work.

However, the observations are not necessarily independent across the agents. Since spatial independence is not required, our setting does not exclude possible latent confounders in the observation model. Each agent  $k$  knows the likelihood model  $L_k(\xi|\theta)$  for every possible hypothesis  $\theta \in \Theta$ . If the distribution  $L_k(\xi|\theta)$  for a hypothesis  $\theta \neq \theta^\circ$  is sufficiently different from  $L_k(\xi|\theta^\circ)$ , then it is easier for agent  $k$  to distinguish  $\theta^\circ$  from  $\theta$ . This “distinguishing power” can be quantified using the KL divergence between the likelihood models, namely,

$$d_k(\theta) \triangleq D_{\text{KL}}\left(L_k(\cdot|\theta^\circ)||L_k(\cdot|\theta)\right). \tag{1}$$

The larger this quantity is, the more *informative* agent  $k$ ’s observations are for distinguishing a wrong hypothesis  $\theta$  from the true hypothesis  $\theta^\circ$ . It can also be interpreted as the degree to which an agent is exposed to the true state of nature. In order to avoid pathological cases, we assume that  $d_k(\theta) < \infty$  for each agent  $k$  and hypothesis  $\theta$ . This condition makes sure that the likelihood functions for different hypotheses share the same support; and no observation alone is sufficient to refute any hypothesis.

**Definition 1 (Global identifiability)** *For each wrong hypothesis  $\theta \neq \theta^\circ$ , if there exists at least one clear-sighted agent  $k^*$  that can distinguish  $\theta$  from the true hypothesis  $\theta^\circ$ , i.e.,  $d_{k^*}(\theta) > 0$ , then the true state of the environment is said to be globally identifiable. ■*

Notice that global identifiability does not require *local* identifiability, which is the ability of each individual agent to distinguish  $\theta^\circ$  from any other hypothesis without cooperation with the other agents.

Based on the local observations and on interactions with other agents, each agent  $k$  forms an opinion (i.e., a belief vector), denoted by  $\mu_{k,i}$ , which is a probability mass function defined over the set of hypotheses  $\Theta$ . Here, the entry  $\mu_{k,i}(\theta)$  quantifies the confidence  $k$  has about  $\theta \in \Theta$  being the true hypothesis  $\theta^\circ$  at time  $i$ . We assume that all agents have positive initial beliefs and they do not reject any hypothesis at the start of the learning process, i.e.,  $\mu_{k,0}(\theta) > 0$  for each agent  $k$  and  $\theta \in \Theta$ . Agents exchange their beliefs with each other according to the communication topology described next.

### 3.2 Network topology




Figure 2: An illustration of a strongly connected network topology.

the matrix  $A$  is left-stochastic. Furthermore, since the underlying graph is strongly-connected,  $A$  is irreducible and aperiodic. Therefore, in view of the Perron-Frobenius theorem (Pillai et al., 2005), it satisfies

$$A^T \mathbf{1}_K = \mathbf{1}_K, \quad \mathbf{1}_K^T v = 1, \quad Av = v, \quad (2)$$

where  $v$  is called the Perron vector of  $A$  whose entries are all positive and add up to one. Here, the  $k$ -th entry of the vector  $v$  measures the relative centrality of agent  $k$  in the network.

### 3.3 Non-Bayesian social learning (NBSL)

In this section, we present the non-Bayesian social learning (NBSL) strategy from (Jadbabaie et al., 2012; Nedić et al., 2017; Lalitha et al., 2018). In this strategy, based on the observation  $\xi_{k,i}$  at time instant  $i$ , each agent first updates its belief in a *locally Bayesian* fashion to obtain its intermediate belief:

$$\psi_{k,i}(\theta) \propto L_k(\xi_{k,i}|\theta)\mu_{k,i-1}(\theta). \quad (3)$$

The proportionality sign  $\propto$  means that the entries of the resulting vector are normalized to add up to one, as befits a true probability mass function. The motivation for (3) is at least two-fold. From a behavioral point of view, Bayes’s rule is used to model human reasoning under uncertainty in neuroscience (Friston, 2005) and the social sciences (Oaksford and Chater, 2007; Easley and Kleinberg, 2010). From a system design perspective, Bayes’s rule is known to be an optimal information processing rule (Zellner, 1988). In the next step, the intermediate beliefs are shared with other agents, which may average them in a geometric manner to form the updated belief using the confidence scores they assign to their neighbors as follows:

$$\mu_{k,i}(\theta) \propto \prod_{\ell \in \mathcal{N}_k} (\psi_{\ell,i}(\theta))^{a_{\ell k}}. \quad (4)$$

The combination step (4) is a non-Bayesian way of combining beliefs and is inspired by the fact that humans are boundedly rational (Conlisk, 1996). In the above implementation, the agents are combining their neighbors’ *instantaneous* opinions, as opposed to behaving in a *fully* Bayesian manner (Acemoglu et al., 2011), which would require global information (such as the graph topology and access to all observations). This requirement makes the fully Bayesian solution NP-hard in general (Hazla et al., 2021). Although there are variations based on the arithmetic averaging (Jadbabaie et al., 2012), in this work we consider the geometric averaging form described above.

An important quantity for the analysis of the strategy (3)–(4) is the log-belief ratio vector  $\boldsymbol{\lambda}_i(\theta) \triangleq [\boldsymbol{\lambda}_{1,i}(\theta), \dots, \boldsymbol{\lambda}_{K,i}(\theta)]^\top$  with individual entries defined as

$$\boldsymbol{\lambda}_{k,i}(\theta) \triangleq \log \frac{\boldsymbol{\mu}_{k,i}(\theta^\circ)}{\boldsymbol{\mu}_{k,i}(\theta)}. \quad (5)$$

For an agent  $k$ , if this quantity is positive for each  $\theta \neq \theta^\circ$ , then its belief vector is maximized at the true hypothesis and the agent can guess the correct hypothesis. Assuming no additional noise in the belief update and combination steps (3)–(4), these equations imply that the vector  $\boldsymbol{\lambda}_i(\theta)$  evolves according to the following linear recursion:

$$\boldsymbol{\lambda}_i(\theta) = A^\top (\boldsymbol{\lambda}_{i-1}(\theta) + \mathbf{x}_i(\theta)), \quad (6)$$

where  $\mathbf{x}_i(\theta) \triangleq [\mathbf{x}_{1,i}(\theta), \dots, \mathbf{x}_{K,i}(\theta)]^\top$  is the vector of log-likelihood ratios (LLRs) at time instant  $i$ :

$$\mathbf{x}_{k,i}(\theta) \triangleq \log \frac{L_k(\boldsymbol{\xi}_{k,i}|\theta^\circ)}{L_k(\boldsymbol{\xi}_{k,i}|\theta)}. \quad (7)$$

**Theorem 1 (Truth learning (Nedić et al., 2017; Lalitha et al., 2018))** *If agents employ the strategy (3)–(4), the asymptotic decay rate of a wrong hypothesis is equal for all agents and given by*

$$\frac{1}{i} \boldsymbol{\lambda}_{k,i}(\theta) \xrightarrow{\text{a.s.}} \sum_{\ell=1}^K v_\ell d_\ell(\theta). \quad (8)$$

Moreover, if global identifiability (Def. 1) holds, then all agents learn the truth with full confidence, i.e.,

$$\boldsymbol{\mu}_{k,i}(\theta^\circ) \xrightarrow{\text{a.s.}} 1. \quad (9)$$

■

### 3.4 Adaptive social learning (ASL)

The traditional non-Bayesian social learning (NBSL) strategy (3)–(4) described in the previous section has the drawback that agents do not prioritize new observations against their old observations. In addition to falling short in modelling human behavior, this strategy can be disadvantageous for engineering applications that require adaptation under non-stationary environments. To tackle this issue, the work (Bordignon et al., 2021) proposed changing the adaptation step (3) into<sup>1</sup>

$$\boldsymbol{\psi}_{k,i}(\theta) \propto L_k^\beta(\boldsymbol{\xi}_{k,i}|\theta) \boldsymbol{\mu}_{k,i-1}^{1-\delta}(\theta), \quad (10)$$

---

1. In fact, (Bordignon et al., 2021) only considers the special cases of  $\beta = \delta$  and  $\beta = 1$ . However, their results can be adapted to general  $\beta > 0$  straightforwardly.

where  $0 < \delta < 1$  and  $\beta > 0$  are design parameters. In particular, large values of  $\delta$  or  $\beta$  place more focus on new observations, whereas small values give importance to past beliefs. The modified adaptation step (10) alters the log-belief ratio recursion (6) to

$$\boldsymbol{\lambda}_i(\theta) = A^\top((1 - \delta)\boldsymbol{\lambda}_{i-1}(\theta) + \beta\mathbf{x}_i(\theta)). \quad (11)$$

In contrast to the NBSL case from the previous section where the beliefs converge to the truth almost surely (Theorem 1), in the *adaptive social learning* (ASL) strategy defined by steps (10) and (4), the beliefs will have everlasting random fluctuations that are necessary for keeping adaptation alive. The next result states that these random fluctuations have a regular behavior in the limit.

**Theorem 2 (Convergence in distribution (Bordignon et al., 2021))** *Under the ASL strategy (10) and (4), the log-belief ratios converge in distribution to the following absolutely convergent series:*

$$\boldsymbol{\lambda}_i(\theta) \xrightarrow{\text{dist.}} \beta \sum_{j=1}^{\infty} (1 - \delta)^{j-1} (A^\top)^j \mathbf{x}_j(\theta). \quad (12)$$

■

In the sequel, we need the following result to analyze the average causal effect between agents.


**Corollary 1 (Expected log-belief ratio in ASL)** *Theorem 2 implies that the log-belief ratios converge in the mean, i.e.,*

$$\lim_{i \rightarrow \infty} \mathbb{E}[\boldsymbol{\lambda}_i(\theta)] = \frac{\beta}{1 - \delta} \left( (I - (1 - \delta)A^\top)^{-1} - I \right) d(\theta) \quad (13)$$

where  $d(\theta) \triangleq [d_1(\theta), d_2(\theta), \dots, d_K(\theta)]^\top$  is the vector of network KL divergences.

**Proof** Since the series on the RHS of (12) is uniformly integrable, the expectation on  $\boldsymbol{\lambda}_i(\theta)$  converges to the expectation of the RHS of (12). The result then follows from the fact that  $\mathbb{E}[\mathbf{x}_j(\theta)] = d(\theta)$  for any time  $j$ , and from the closed-form expression for the series of absolutely convergent matrices. ■

## 4. Causal Effects in Social Learning



An intuitive and widely used assertion in defining causality is that manipulation of the causes should result in changes in the effect (Woodward, 2004). Based on this principle, interventions on a system, real or hypothetical, have been the primary tool for testing whether a variable causes another (Pearl, 2009, Chap. 1). In this work, in order to measure the causal influence strength between agents, we rely on the most basic intervention known as *atomic* and *persistent* intervention (Pearl, 2009, Chap. 3). Specifically,

Figure 3: An illustration of an atomic and persistent intervention on agent  $m = 1$  from Fig. 2. 9

in order to measure the social effect of an agent  $m$  on other agents, we analyze the belief evolution of these other agents if the belief of agent  $m$  is fixed to some particular constant belief vector, say,  $\boldsymbol{\mu}_{m,i} := \mu_m$  for all time instants — see Fig. 3 for a visual depiction<sup>2</sup>. In canonical causality notation, this intervention is denoted by  $do(\boldsymbol{\mu}_{m,i} := \mu_m)$  (Pearl, 2009). Since we consider only this intervention in this work and there is no room for ambiguity, we will use the notation that the post-intervention counterparts of the variables in Sec. 3 are topped with the symbol ‘ $\sim$ ’. For example, the log-belief ratio definition from (5) transforms into the following, under the intervention  $do(\boldsymbol{\mu}_{m,i} := \mu_m)$ :

$$\tilde{\lambda}_{k,i}(\theta) \triangleq \log \frac{\tilde{\boldsymbol{\mu}}_{k,i}(\theta^\circ)}{\tilde{\boldsymbol{\mu}}_{k,i}(\theta)}. \quad (14)$$

**Causal influence strength.** Intuitively, the amount of change in the effect following an intervention on the cause is expected to be related to the causal strength. Therefore, the difference between the post and pre-intervention distributions, or between appropriate functions of these distributions such as expectations, can be used to quantify the causal effect (Eichler, 2012; Pearl, 2009, Chap. 3). In this work, we employ the following definition in order to measure the causal influence of agent  $m$  on agent  $k$  (its connection to established definitions of causality in the literature will be discussed in the sequel):

$$C_{m \rightarrow k} \triangleq \mu_{k,\infty}(\theta^\circ) - \tilde{\mu}_{k,\infty}(\theta^\circ) \quad (15)$$

This formula measures the alteration of agent  $k$  as a consequence of an intervention on agent  $m$ . Specifically, it quantifies the magnitude of change of the expected asymptotic belief of agent  $k$  on the true state  $\theta^\circ$ . In general, its value depends on the belief  $\mu_m$  of the intervention  $do(\boldsymbol{\mu}_{m,i} := \mu_m)$ . Here, we use the following expression for the belief vector, which is explained in the sequel:

$$\mu_{k,\infty}(\theta^\circ) \triangleq \frac{1}{1 + \sum_{\theta \in \Theta \setminus \{\theta^\circ\}} \exp\{-\lambda_{k,\infty}(\theta)\}}. \quad (16)$$

This expression is defined in terms of the expected asymptotic log-belief ratio:

$$\lambda_{k,\infty}(\theta) \triangleq \lim_{i \rightarrow \infty} \mathbb{E}[\lambda_{k,i}(\theta)] \quad (17)$$

The variables topped with the symbol ‘ $\sim$ ’ for the intervention case are defined similarly (the existence of the limit for both NBSL and ASL under interventions will be discussed in the sequel). The transformation (16) is motivated by noting from (5) that:

$$\exp\{-\lambda_{k,i}(\theta)\} = \frac{\boldsymbol{\mu}_{k,i}(\theta)}{\boldsymbol{\mu}_{k,i}(\theta^\circ)}, \quad (18)$$

which implies

$$1 + \sum_{\theta \in \Theta \setminus \{\theta^\circ\}} \exp\{-\lambda_{k,i}(\theta)\} = \frac{1}{\boldsymbol{\mu}_{k,i}(\theta^\circ)} \sum_{\theta \in \Theta} \boldsymbol{\mu}_{k,i}(\theta) = \frac{1}{\boldsymbol{\mu}_{k,i}(\theta^\circ)}, \quad (19)$$

---

2. We use an atomic intervention since it isolates the causal contribution of a single agent by fixing only that agent’s belief and preserving the remaining interaction dynamics of the network. We also use it in a persistent manner since we are interested in the total long term effect rather than a transient one.

which, in turn, yields

$$\boldsymbol{\mu}_{k,i}(\theta^\circ) = \frac{1}{1 + \sum_{\theta \in \Theta \setminus \{\theta^\circ\}} \exp\{-\boldsymbol{\lambda}_{k,i}(\theta)\}}. \quad (20)$$

Here, if we replace log-belief ratio  $\boldsymbol{\lambda}_{k,i}(\theta)$  with the expected asymptotic log-belief ratio  $\lambda_{k,\infty}(\theta)$ , we arrive at the definition (16) for  $\mu_{k,\infty}(\theta^\circ)$ . Note that defining  $\mu_{k,\infty}(\theta^\circ)$  in terms of the expected log-belief ratios, as opposed to, say, expected beliefs (i.e.,  $\lim_{i \rightarrow \infty} \mathbb{E}[\boldsymbol{\mu}_{k,i}(\theta^\circ)]$ ), will enable us to obtain closed-form expressions for causality in terms of the informativeness of the agents, represented by the entries of  $d(\theta)$ , in Sec. 5. Next, we treat NBSL and ASL separately.

- **Non-Bayesian social learning.** In the idle case (i.e., no intervention) of NBSL, it is known from Theorem 1 that under global identifiability,  $\boldsymbol{\mu}_{k,i}(\theta^\circ) \xrightarrow{\text{a.s.}} 1$  (i.e.,  $\boldsymbol{\lambda}_{k,i}(\theta) \xrightarrow{\text{a.s.}} +\infty$  for each  $\theta \neq \theta^\circ$ ). Hence, for the NBSL case, the average (i.e., expected) causal influence (15) is given by

$$C_{m \rightarrow k}^{\text{NB}} = 1 - \tilde{\mu}_{k,\infty}(\theta^\circ). \quad (21)$$

This immediately implies that  $C_{m \rightarrow k}^{\text{NB}} \in [0, 1]$ , and it gets larger as the post-intervention belief diverges from the truth in expectation.

- **Adaptive social learning.** In a similar fashion, the causal effect strength for the ASL case is given by the general expression (15)

$$C_{m \rightarrow k}^{\text{ASL}} = \mu_{k,\infty}(\theta^\circ) - \tilde{\mu}_{k,\infty}(\theta^\circ), \quad (22)$$

where the pre-intervention asymptotic belief  $\mu_{k,\infty}(\theta^\circ)$  can be found by inserting  $\lambda_{k,\infty}(\theta)$  established in Corollary 1 to (16).

**Remark 1 (Controllability)** *Notably, the influence  $C_{m \rightarrow k}$  of agent  $m$  on agent  $k$  can also be interpreted as the controllability or manipulability of agent  $k$  by agent  $m$ .* ■

## 5. Theoretical Results

In this section, we derive closed-form expressions for  $\tilde{\lambda}_{k,\infty} \triangleq [\tilde{\lambda}_{k,\infty}(\theta_1), \dots, \tilde{\lambda}_{k,\infty}(\theta_H)]^\top$  in terms of the network topology and the informativeness of agents to obtain the causal strength measures  $C_{m \rightarrow k}^{\text{NB}}$  and  $C_{m \rightarrow k}^{\text{ASL}}$ . For ease of notation and without loss of generality, we set  $m = 1$ . One can obtain  $C_{m \rightarrow k}^{\text{NB}}$  by setting  $\beta = 1$  and  $\delta \rightarrow 0$  in  $C_{m \rightarrow k}^{\text{ASL}}$  due to (3) and (10). Nevertheless, we first present the analysis for NBSL since it is easier to derive and provides useful insights for ASL. Subsequently, we provide the results for ASL with proofs deferred to the appendix.

### 5.1 Non-Bayesian Social Learning

The intervention  $do(\boldsymbol{\mu}_{1,i} := \mu_1)$  ceases (or obstructs the use of) all incoming information at agent 1 from the neighbors  $\mathcal{N}_1$  and the use of the streaming observations  $\boldsymbol{\xi}_{1,i}$  from the

environment. Consequently, we can model this effect by redefining the combination matrix and the LLR vector counterparts under the intervention:

$$\tilde{A} \triangleq [\tilde{a}_{\ell k}], \quad \tilde{a}_{\ell k} \triangleq \begin{cases} 1, & \ell = k = 1 \\ 0, & \ell \neq k = 1 \\ a_{\ell k}, & \ell \neq 1, k \neq 1 \end{cases}, \quad \tilde{\mathbf{x}}_i(\theta) \triangleq [0, \mathbf{x}_{2,i}(\theta), \dots, \mathbf{x}_{K,i}(\theta)]^\top. \quad (23)$$

Observe that the effective combination matrix  $\tilde{A}$  can be obtained from  $A$  as follows:

$$A = \left[ \begin{array}{c|c} a_{11} & r^\top \\ a_{21} & \\ \vdots & R \\ a_{K1} & \end{array} \right] \implies \tilde{A} = \left[ \begin{array}{c|c} 1 & r^\top \\ 0 & R \end{array} \right], \quad (24)$$

for a  $(K-1) \times 1$  dimensional vector  $r$  and a  $(K-1) \times (K-1)$  dimensional matrix  $R$ :

$$r \triangleq [a_{12} \quad a_{13} \quad \dots \quad a_{1K}]^\top, \quad R \triangleq \begin{bmatrix} a_{22} & \dots & a_{2K} \\ \vdots & \ddots & \vdots \\ a_{K2} & \dots & a_{KK} \end{bmatrix}. \quad (25)$$

The matrix structure of  $\tilde{A}$  belongs to the class of *reducible* combination matrices (Resnick, 2013; Smith et al., 2014) that arise in the analysis of *weakly* connected networks (Molavi et al., 2013; Mossel et al., 2015; Ying and Sayed, 2016). As opposed to the strongly connected networks where information can flow thoroughly, in weakly connected networks information can flow only in one direction between certain parts of the network. In the current context, this corresponds to the fact that information continues to flow from agent  $m$  in the form of its belief vector fixed at  $\mu_m$ , but no information is flowing into it in the sense that agent  $m$  ignores all signals arriving from its neighbors and does not use them to update its local belief. However, in contrast to these prior works that analyze opinion dynamics under weakly connected networks, we are interested in the effect of the intervention on the original *strongly* connected network. This alters the LLR  $\mathbf{x}_i(\theta)$  as well — see (23). Similar to the original case in (6), we proceed by studying the log-belief ratio evolution that results from using  $\tilde{A}$ :

$$\tilde{\boldsymbol{\lambda}}_i(\theta) = \tilde{A}^\top (\tilde{\boldsymbol{\lambda}}_{i-1}(\theta) + \tilde{\mathbf{x}}_i(\theta)). \quad (26)$$

Recursive application of (26) across  $i$  iterations to the log-belief ratio  $\tilde{\boldsymbol{\lambda}}_i(\theta)$  yields

$$\tilde{\boldsymbol{\lambda}}_i(\theta) = \sum_{j=1}^i (\tilde{A}^{i-j+1})^\top \tilde{\mathbf{x}}_j(\theta) + (\tilde{A}^i)^\top \tilde{\boldsymbol{\lambda}}_0(\theta). \quad (27)$$

To study (27), we need to evaluate the powers of the effective combination matrix:

$$\tilde{A}^i = \left[ \begin{array}{c|c} 1 & r_i'^\top \\ 0 & R^i \end{array} \right], \quad \tilde{A}^\infty \stackrel{(a)}{=} \left[ \begin{array}{c|ccc} 1 & 1 & \dots & 1 \\ 0 & 0 & & \end{array} \right] \quad (28)$$

where (a) follows from the fact that the spectral radius of the matrix  $R$  is strictly smaller than 1, i.e.,  $R$  is a stable matrix (Ying and Sayed, 2016, Lemma 1). For each time  $i$  and  $0 < j \leq i$ , observe that

$$(\tilde{A}^{i-j+1})^\top \tilde{\mathbf{x}}_j(\theta) \stackrel{(23),(28)}{=} \left[ \begin{array}{c|c} 1 & 0 \\ \hline r'_{i-j+1} & R^{i-j+1\top} \end{array} \right] \begin{bmatrix} 0 \\ \mathbf{x}_{2,j}(\theta) \\ \vdots \\ \mathbf{x}_{K,j}(\theta) \end{bmatrix} \quad (29a)$$

$$\implies [(\tilde{A}^{i-j+1})^\top \tilde{\mathbf{x}}_j(\theta)]_1 = 0 \quad (29b)$$

where  $r'_i$  denotes the accumulated upper-right block in the power  $\tilde{A}^i$ , and

$$(\tilde{A}^i)^\top \tilde{\boldsymbol{\lambda}}_0(\theta) \stackrel{(23),(28)}{=} \left[ \begin{array}{c|c} 1 & 0 \\ \hline r'_i & R^{i\top} \end{array} \right] \begin{bmatrix} \log \frac{\mu_1(\theta^\circ)}{\mu_1(\theta)} \\ \log \frac{\boldsymbol{\mu}_{2,0}(\theta^\circ)}{\boldsymbol{\mu}_{2,0}(\theta)} \\ \vdots \\ \log \frac{\boldsymbol{\mu}_{K,0}(\theta^\circ)}{\boldsymbol{\mu}_{K,0}(\theta)} \end{bmatrix} \quad (30a)$$

$$\implies [(\tilde{A}^i)^\top \tilde{\boldsymbol{\lambda}}_0(\theta)]_1 = \log \frac{\mu_1(\theta^\circ)}{\mu_1(\theta)}. \quad (30b)$$

Inserting these into (27) verifies that the intervention is fixed as intended for all time instants, since

$$[\tilde{\boldsymbol{\lambda}}_i(\theta)]_1 \stackrel{(27)}{=} \sum_{j=1}^i [(\tilde{A}^{i-j+1})^\top \tilde{\mathbf{x}}_j(\theta)]_1 + [(\tilde{A}^i)^\top \tilde{\boldsymbol{\lambda}}_0(\theta)]_1. \quad (31a)$$

$$\tilde{\boldsymbol{\lambda}}_{1,i}(\theta) \stackrel{(29b),(30b)}{=} \log \frac{\mu_1(\theta^\circ)}{\mu_1(\theta)}, \quad \tilde{\boldsymbol{\mu}}_{1,i}(\theta) = \mu_1(\theta). \quad (31b)$$

Moreover, if we take the expectation of both sides of (27), we get

$$\begin{aligned} \mathbb{E}[\tilde{\boldsymbol{\lambda}}_i(\theta)] &= \sum_{j=1}^i (\tilde{A}^{i-j+1})^\top \mathbb{E}[\tilde{\mathbf{x}}_j(\theta)] + (\tilde{A}^i)^\top \mathbb{E}[\tilde{\boldsymbol{\lambda}}_0(\theta)] \\ &= \sum_{j=1}^i (\tilde{A}^{i-j+1})^\top \tilde{d}(\theta) + (\tilde{A}^i)^\top \mathbb{E}[\tilde{\boldsymbol{\lambda}}_0(\theta)]. \end{aligned} \quad (32)$$

where we use the definition  $\tilde{d}(\theta) \triangleq [0, d_2(\theta), \dots, d_K(\theta)]^\top$ . Hence, in the limit (the existence is guaranteed by the finiteness of LLRs and positive initial beliefs), it holds that

$$\begin{aligned} \lim_{i \rightarrow \infty} \mathbb{E}[\tilde{\boldsymbol{\lambda}}_i(\theta)] &= \lim_{i \rightarrow \infty} \sum_{j=1}^i (\tilde{A}^{i-j+1})^\top \tilde{d}(\theta) + (\tilde{A}^\infty)^\top \mathbb{E}[\tilde{\boldsymbol{\lambda}}_0(\theta)] \\ &= \sum_{j=1}^{\infty} (\tilde{A}^j)^\top \tilde{d}(\theta) + (\tilde{A}^\infty)^\top \mathbb{E}[\tilde{\boldsymbol{\lambda}}_0(\theta)]. \end{aligned} \quad (33)$$

If we incorporate (28) into (33), this implies for the log-belief ratios of all agents except agent  $m = 1$  that

$$\tilde{\lambda}_{-m,\infty}(\theta) \triangleq [\tilde{\lambda}_{2,\infty}(\theta), \dots, \tilde{\lambda}_{K,\infty}(\theta)]^\top = \sum_{j=1}^{\infty} (R^j)^\top d_{-m}(\theta) + \left( \log \frac{\mu_m(\theta^\circ)}{\mu_m(\theta)} \right) \mathbf{1}_{K-1} \quad (34)$$

where  $d_{-m}(\theta)$  is the  $(K - 1) \times 1$  dimensional vector of local KL divergences of the remaining agents, i.e.,  $d_{-m}(\theta) \triangleq \text{col}\{d_\ell(\theta)\}_{\ell=2}^K$ . Since  $R$  is a stable matrix (Ying and Sayed, 2016, Lemma 1), Eq. (34) can alternatively be written as

$$\tilde{\lambda}_{-m,\infty}(\theta) = \left( (I - R^\top)^{-1} - I \right) d_{-m}(\theta) + \left( \log \frac{\mu_m(\theta^\circ)}{\mu_m(\theta)} \right) \mathbf{1}_{K-1} \quad (35)$$

The causal influence of agent  $m = 1$  on agent  $k$  can now be calculated by inserting  $\tilde{\lambda}_{k,\infty}(\theta)$  into (16) to find  $\tilde{\mu}_{k,\infty}(\theta)$ , which is the input for (21) that yields  $C_{m \rightarrow k}^{\text{NB}}$ . More specifically, if we incorporate (35) into (16), we get

$$\tilde{\mu}_{k,\infty}(\theta^\circ) = \frac{1}{1 + \sum_{\theta \in \Theta \setminus \{\theta^\circ\}} \frac{\mu_m(\theta)}{\mu_m(\theta^\circ)} \exp \left\{ - \left[ \left( (I - R^\top)^{-1} - I \right) d_{-m}(\theta) \right]_k \right\}} \quad (36)$$

which, by (21), implies

$$C_{m \rightarrow k}^{\text{NB}} = 1 - \frac{1}{1 + \sum_{\theta \in \Theta \setminus \{\theta^\circ\}} \frac{\mu_m(\theta)}{\mu_m(\theta^\circ)} \exp \left\{ - \left[ \left( (I - R^\top)^{-1} - I \right) d_{-m}(\theta) \right]_k \right\}}. \quad (37)$$

Equation (35) is a general result which shows that  $C_{m \rightarrow k}^{\text{NB}}$  is a function of (i) the combination weights (via  $R$ ), and (ii) the individual informativeness of each agent (via  $d_{-m}(\theta)$ ).

**Remark 2 (Generalization to sub-networks)** *Note that expressions (35)-(37) can generalize to the influence of a sub-network with multiple agents rather than an individual agent  $m$ . In this case, upon intervening on all agents within the sub-network, the effective combination matrix becomes*

$$\tilde{A} = \left[ \begin{array}{c|c} I & r^\top \\ \hline 0 & R \end{array} \right] \quad (38)$$

where the first entry in (24) is replaced by an identity matrix. Namely, if the sub-network under consideration has size  $s$ , the identity,  $r$  and  $R$  matrices would be of dimensions  $s \times s$ ,  $(K - s) \times s$ , and  $(K - s) \times (K - s)$ , respectively. Similarly,  $d_{-m}(\theta)$  can be replaced with local KL divergences of the agents that do not belong to the treated sub-network. ■

**Remark 3 (Finite-time spread of influence)** *Expressions (35)-(37) reveal how the total overall effects depend on direct instantaneous effects in temporal networked interactions. For*

the spread of direct effects in a finite time instant  $n$ , we can modify (34) and (35) as

$$\tilde{\lambda}_{-m,n}(\theta) \triangleq [\tilde{\lambda}_{2,n}(\theta), \dots, \tilde{\lambda}_{K,n}(\theta)]^\top = \sum_{j=1}^n (R^j)^\top d_{-m}(\theta) + \left( \log \frac{\mu_m(\theta^\circ)}{\mu_m(\theta)} \right) \mathbf{1}_{K-1} \quad (39)$$

due to (27), and consequently,

$$\tilde{\lambda}_{-m,n}(\theta) = \left( (I - R^\top)^{-1} (I - R^{n+1\top}) - I \right) d_{-m}(\theta) + \left( \log \frac{\mu_m(\theta^\circ)}{\mu_m(\theta)} \right) \mathbf{1}_{K-1}. \quad (40)$$

■

For ease of the presentation, we continue to describe the total causal effects of a single agent, even though our results can be extended to sub-network influence and finite time analysis in a straightforward manner as discussed in the above remarks.

In (37),  $C_{m \rightarrow k}^{\text{NB}}$  represents a dose-response curve, assuming different values for different intervention strengths (i.e., dose)  $\mu_m$ . This is a typical situation in the context of continuous-valued interventions. In some applications, however, it proves beneficial to encapsulate the causal effect value with a single number. For this purpose, we may set  $\mu_m$  to be uniform across all hypotheses, i.e.,  $\mu_m(\theta) = 1/H, \forall \theta \in \Theta$ . This method of summarizing the causal effect is denoted as follows:

$$\bar{C}_{m \rightarrow k}^{\text{NB}} \triangleq C_{m \rightarrow k}^{\text{NB}} \Big|_{\mu_m(\theta)=1/H} \quad (41)$$

In Appendix A, we show that this choice effectively parallels the process of determining the average causal *derivative* effect (Peters et al., 2017, Chap. 6), which is a method commonly used in the literature for summarizing the causal effect. Basically, it quantifies the extent of change in agent  $k$  in response to an infinitesimal variation in the intervention strength  $\mu_m$ .

In the next section, we study two special network topologies that help illustrate the dependencies of the causal effect more explicitly.

### 5.1.1 SPECIAL CASES

**Fully-connected and federated architectures.** In this example, we consider a fully-connected network (see Fig. 4) with a rank-one combination matrix and Perron vector  $v$ , i.e.,

$$A = \begin{bmatrix} v_1 & v_1 & \cdots & v_1 \\ v_2 & v_2 & \cdots & v_2 \\ \vdots & \vdots & \ddots & \vdots \\ v_K & v_K & \cdots & v_K \end{bmatrix} = v \mathbf{1}_K^\top. \quad (42)$$

Note that in terms of performance, this system is equivalent to a federated architecture in which (i) agents send their beliefs to a fusion center after local adaptation, (ii) the center averages the received beliefs in a weighted manner, and (iii) then broadcasts the combined belief to all agents — see Fig. 5. Under intervention  $do(\mu_{1,i} := \mu_1)$ , we have




Figure 4: A fully connected network topology.




Figure 5: A federated architecture. The server broadcasts the weighted average of agents' beliefs back to them at each iteration. In terms of performance, this system is equivalent to the fully connected architecture of Fig. 4.

$$\tilde{A} = \begin{bmatrix} 1 & v_1 & \cdots & v_1 \\ 0 & v_2 & \cdots & v_2 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & v_K & \cdots & v_K \end{bmatrix} \implies R = \begin{bmatrix} v_2 & \cdots & v_2 \\ \vdots & \ddots & \vdots \\ v_K & \cdots & v_K \end{bmatrix} = v_{-m} \mathbf{1}_{K-1}^T \quad (43)$$

where  $v_{-m} \triangleq \text{col}\{v_\ell\}_{\ell=2}^K$  is a  $(K-1) \times 1$  dimensional vector consisting of all Perron entries except for agent  $m = 1$ . Observe that

$$R^2 = v_{-m} \mathbf{1}_{K-1}^T v_{-m} \mathbf{1}_{K-1}^T \stackrel{(a)}{=} (1 - v_1) v_{-m} \mathbf{1}_{K-1}^T, \quad (44)$$

where (a) follows from the fact that  $\mathbf{1}_K^T v = 1$  (Eq. (2)). Repeating the same arguments, it holds that

$$R^i = (1 - v_1)^{i-1} v_{-m} \mathbf{1}_{K-1}^T. \quad (45)$$

Therefore,

$$(I - R^T)^{-1} - I = \sum_{i=1}^{\infty} (R^T)^i = \mathbf{1}_{K-1} v_{-m}^T \sum_{i=1}^{\infty} (1 - v_1)^{i-1} = \frac{1}{v_1} \mathbf{1}_{K-1} v_{-m}^T. \quad (46)$$

Inserting this into (35), we arrive at the following expression for each agent  $k \neq m$ :

$$\tilde{\lambda}_{k,\infty}(\theta) = \frac{1}{v_m} \sum_{\ell=2}^K v_\ell d_\ell(\theta) + \log \frac{\mu_m(\theta^\circ)}{\mu_m(\theta)} \quad (47)$$

Combining (16) and (21) with (47) yields the causal effect:

$$\begin{aligned} C_{m \rightarrow k}^{\text{NB}} &\stackrel{(16),(21)}{=} 1 - \frac{1}{1 + \sum_{\theta \in \Theta \setminus \{\theta^\circ\}} \exp\{-\tilde{\lambda}_{k,\infty}(\theta)\}} \\ &\stackrel{(47)}{=} 1 - \frac{1}{1 + \sum_{\theta \in \Theta \setminus \{\theta^\circ\}} \frac{\mu_m(\theta)}{\mu_m(\theta^\circ)} \exp\left\{-\frac{1}{v_m} \sum_{\ell \neq m} v_\ell d_\ell(\theta)\right\}}. \end{aligned} \quad (48)$$

The effect of agent  $m$  on all other agents is the same, which is expected due to the symmetric nature of this special example. Furthermore, observe that the causal effect  $C_{m \rightarrow k}^{\text{NB}}$  decreases with increasing  $\tilde{\lambda}_{k,\infty}(\theta)$ . On that account, from (47) and (48) it can be seen that:

- Increasing the network centrality of agent  $m = 1$  (i.e., increasing  $v_m$ ) decreases  $\tilde{\lambda}_{k,\infty}(\theta)$ , and in turn increases the causal effect  $C_{m \rightarrow k}^{\text{NB}}$ . Therefore, an agent has more effect on other agents if it has a higher network centrality. In particular, if

$$v_m \rightarrow 0 \implies \tilde{\lambda}_{k,\infty}(\theta) \rightarrow +\infty \implies C_{m \rightarrow k}^{\text{NB}} \rightarrow 0, \quad (49)$$

which means that an agent with negligible network centrality has no causal effect on other agents.

- Increasing network centrality and informativeness of the other agents  $\ell \neq m$  (i.e., increasing  $v_\ell$  and  $d_\ell(\theta)$ ) increases  $\tilde{\lambda}_{k,\infty}(\theta)$ , and in turn decreases the causal effect  $C_{m \rightarrow k}^{\text{NB}}$ . In particular, if the most informative agents are equipped with the highest network centrality, then  $\sum_{\ell=2}^K v_\ell d_\ell(\theta)$  is large and it is harder for agent  $m = 1$  to control other agents.
- If the fixed belief on the true hypothesis  $\mu_m(\theta^\circ)$  decreases, then  $\tilde{\lambda}_{k,\infty}(\theta)$  decreases and the causal effect  $C_{m \rightarrow k}^{\text{NB}}$  increases. This suggests that the further from the truth the information an agent supplies, the more effect that agent will have on other agents. In other words, agents supplying misinformation have more effect on the rest of the network. Specifically, observe that if the rest of the agents have a low informativeness average, i.e., if

$$\sum_{\ell=2}^K v_\ell d_\ell(\theta) \approx 0 \implies C_{m \rightarrow k}^{\text{NB}} \approx 1 - \mu_m(\theta^\circ). \quad (50)$$

Therefore, the causal effect is proportional to the difference from the truth. It is maximized (i.e.,  $C_{m \rightarrow k}^{\text{NB}} = 1$ ) when the fixed belief assigns 0 to the true hypothesis.




Figure 6: A unidirectional ring.

**Ring architecture.** In this example, we consider a unidirectional ring network where each agent has a self-confidence of  $\alpha$ , and assigns a confidence of  $1 - \alpha$  to the preceding agent in the ring — see Fig. 6. Agents are indexed such that agent  $k + 1$  receives (or uses) information from agent  $k$  only. The combination matrix has the form:

$$A = \begin{bmatrix} \alpha & 1 - \alpha & 0 & \cdots & 0 \\ 0 & \alpha & 1 - \alpha & \cdots & 0 \\ \vdots & 0 & \alpha & \cdots & \vdots \\ 0 & \vdots & \vdots & \ddots & 1 - \alpha \\ 1 - \alpha & 0 & 0 & \cdots & \alpha \end{bmatrix}. \quad (51)$$

Under intervention  $do(\mu_{1,i} := \mu_1)$ , we have

$$\tilde{A} = \left[ \begin{array}{c|cccc} 1 & 1-\alpha & 0 & \cdots & 0 \\ 0 & \alpha & 1-\alpha & \cdots & 0 \\ 0 & 0 & \alpha & \cdots & \vdots \\ \vdots & \vdots & \vdots & \ddots & 1-\alpha \\ 0 & 0 & 0 & \cdots & \alpha \end{array} \right] \implies R = \begin{bmatrix} \alpha & 1-\alpha & \cdots & 0 \\ 0 & \alpha & \cdots & \vdots \\ \vdots & \vdots & \ddots & 1-\alpha \\ 0 & 0 & \cdots & \alpha \end{bmatrix}. \quad (52)$$

As a result,

$$(I - R^\top) = (1 - \alpha) \begin{bmatrix} 1 & 0 & \cdots & 0 \\ -1 & 1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & \cdots & -1 & 1 \end{bmatrix} \implies (I - R^\top)^{-1} = \frac{1}{1 - \alpha} \begin{bmatrix} 1 & 0 & \cdots & 0 \\ 1 & 1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 1 & 1 & \cdots & 1 \end{bmatrix}, \quad (53)$$

which implies

$$((I - R^\top)^{-1} - I)d_{-m}(\theta) = \frac{1}{1 - \alpha} \left[ \alpha d_2(\theta), d_2(\theta) + \alpha d_3(\theta), \dots, \sum_{\ell=2}^{K-1} d_\ell(\theta) + \alpha d_K(\theta) \right]^\top. \quad (54)$$

Consequently, for agent  $k$ ,

$$\tilde{\lambda}_{k,\infty}(\theta) = \frac{1}{1 - \alpha} \sum_{\ell=2}^{k-1} d_\ell(\theta) + \frac{\alpha}{1 - \alpha} d_k(\theta) + \log \frac{\mu_m(\theta^\circ)}{\mu_m(\theta)} \quad (55)$$

and, in addition, by definitions (16) and (21),

$$C_{m \rightarrow k}^{\text{NB}} = 1 - \frac{1}{1 + \sum_{\theta \in \Theta \setminus \{\theta^\circ\}} \frac{\mu_m(\theta)}{\mu_m(\theta^\circ)} \exp \left\{ -\frac{1}{1 - \alpha} \sum_{\ell=2}^{k-1} d_\ell(\theta) - \frac{\alpha}{1 - \alpha} d_k(\theta) \right\}}. \quad (56)$$

As stated before, the causal effect  $C_{m \rightarrow k}^{\text{NB}}$  decreases with increasing  $\tilde{\lambda}_{k,\infty}(\theta)$ . Therefore, the following remarks for (55) and (56) are in place:

- Since the KL divergence  $d_\ell(\theta)$  is non-negative,  $\tilde{\lambda}_{k,\infty}(\theta)$  is monotonically increasing along the path  $k = 2 \rightarrow \dots \rightarrow k = K$ . Therefore, the causal effect of agent  $m = 1$  is monotonically decreasing along the same path: the closer agent  $m$  is to an agent, the higher its effect on that agent. This is intuitive because the effect that agent  $m$  has on agent  $k + 1$  is transferred via agent  $k$  in the ring structure. The difference between the causal effects of agent  $m$  on agents  $k$  and  $k + 1$  is proportional to the increase in  $\tilde{\lambda}_{k,\infty}(\theta)$ , that is,

$$\tilde{\lambda}_{k+1,\infty}(\theta) - \tilde{\lambda}_{k,\infty}(\theta) = d_k(\theta) + \frac{\alpha}{1 - \alpha} d_{k+1}(\theta). \quad (57)$$

This means that informative agents with high KL divergence on the path between agent  $m$  and agent  $k$  reduce the causal effect  $C_{m \rightarrow k}^{\text{NB}}$ . In other words, the sphere of influence of an agent  $m$  is bigger if there are no other informative agents in the vicinity.

- For the immediate follower of agent  $m = 1$ , it follows that

$$\tilde{\lambda}_{2,\infty}(\theta) = \frac{\alpha}{1-\alpha}d_2(\theta) + \log \frac{\mu_1(\theta^\circ)}{\mu_1(\theta)}. \quad (58)$$

If agent 2 is not sufficiently informative itself, i.e.,  $d_2(\theta)$  is small, then  $\tilde{\lambda}_{2,\infty}(\theta)$  gets smaller and  $C_{1 \rightarrow 2}^{\text{NB}}$  gets higher. In other words, an agent is more controllable if it is not knowledgeable.

- The limiting average  $\tilde{\lambda}_{k,\infty}(\theta)$  increases with increasing  $\alpha$ . Therefore, if agents are more self-confident, the causal strength is smaller and agents are less controllable by other agents.

## 5.2 Adaptive Social Learning

Similar to the modification in the NBSL case, the log-belief recursion (11) in ASL is modified as follows under intervention  $do(\mu_{1,i} := \mu_1)$ :

$$\tilde{\lambda}_i(\theta) = \tilde{A}^\top((1-\delta)\tilde{\lambda}_{i-1}(\theta) + \beta\tilde{\mathbf{x}}_i(\theta)). \quad (59)$$

The effective combination matrix  $\tilde{A}$  continues to be given by (23). However, the effective LLR is now given by

$$\tilde{\mathbf{x}}_i(\theta) \triangleq \left[ \frac{\delta}{\beta} \log \frac{\mu_1(\theta^\circ)}{\mu_1(\theta)}, \mathbf{x}_{2,i}(\theta), \dots, \mathbf{x}_{K,i}(\theta) \right]^\top, \quad (60)$$

where the first entry is different than the NBSL case. This is to compensate for the presence of the parameters  $\delta$  and  $\beta$ . Observe from (59) that  $\tilde{A}$  from (23) and  $\tilde{\mathbf{x}}_i(\theta)$  from (60) verify  $\tilde{\lambda}_{1,i}(\theta) = \log \frac{\mu_1(\theta^\circ)}{\mu_1(\theta)}$  for all time instants, i.e., the intervention is fixed, by similar arguments to (31b). In Appendix B, we derive the following expression for the limiting log-belief ratio expectations for the rest of the network:

$$\tilde{\lambda}_{-m,\infty}(\theta) = \left( \log \frac{\mu_m(\theta^\circ)}{\mu_m(\theta)} \right) (I - (1-\delta)R^\top)^{-1}r + \frac{\beta}{1-\delta} \left( (I - (1-\delta)R^\top)^{-1} - I \right) d_{-m}(\theta) \quad (61)$$

The causal effect  $C_{m \rightarrow k}^{\text{ASL}}$  can be calculated by inserting post-intervention expression (61) and pre-intervention expression (13) into the definitions (16) and (22). Notice from (61) that similar to the NBSL case, the causal effect depends on the informativeness of agents, the network topology, and the strength of intervention via  $d_{-m}(\theta)$ ,  $R$ , and  $\mu_m$ , respectively. In fact, if we set  $\delta = 0$  and  $\beta = 1$ , Eq. (61) reduces to the NBSL expression (35) as expected. This is because the left-stochastic nature of  $A$  implies that

$$r + R^\top \mathbf{1}_{K-1} = \mathbf{1}_{K-1} \iff (I - R^\top)^{-1}r = \mathbf{1}_{K-1}. \quad (62)$$

In addition, the causal effects in ASL are affected by the importance weighting parameters  $\delta$  and  $\beta$ , as well as by the vector  $r$  that represents the confidence weights other agents assign to agent  $m$ . In (35), the entries of  $r$  implicitly influence the causal effect via  $R$ : the column-wise summation of the entries of  $r$  and  $R$  results in 1 at all columns due to the left-stochastic nature of  $A$ . In comparison, in ASL, both  $r$  and  $R$  impact  $C_{m \rightarrow k}^{\text{ASL}}$  explicitly. If we take a closer look at the terms in (61), we can see that:

- The vector that scales the intervened log-belief ratio  $\log \frac{\mu_m(\theta^\circ)}{\mu_m(\theta)}$  can be expanded as

$$(I - (1 - \delta)R^\top)^{-1}r = r + (1 - \delta)R^\top r + (1 - \delta)^2 R^{2\top} r + (1 - \delta)^3 R^{3\top} r + \dots \quad (63)$$

On the RHS of this equation, the first  $r$  represents the scaling of the information transferred from agent  $m$  to the rest of the network *directly*. Namely, for an agent  $k$ , the scaling of the direct information is  $a_{mk}$  if  $m$  is an immediate neighbor ( $m \in \mathcal{N}_k$ ); 0 if it is not ( $m \notin \mathcal{N}_k$ ). On the other hand, the second term  $(1 - \delta)R^\top r$  describes the scaling of the information transferred from agent  $m$  to the rest of the network, which is then mixed with the other agents  $\forall k \neq m$  (via  $R^\top$ ) and “forgotten” (i.e., lose its recency) for one time instant by a factor of  $(1 - \delta)$ . The consecutive terms over time follow from the same scaling argument.

- In a similar manner, we can express the matrix that scales the vector of individual informativeness in the rest of the network  $d_{-m}(\theta)$  as:

$$\frac{1}{1 - \delta} \left( (I - (1 - \delta)R^\top)^{-1} - I \right) = R^\top + (1 - \delta)R^{2\top} + (1 - \delta)^2 R^{3\top} + \dots \quad (64)$$

Since there is no outgoing link from the rest of the network to agent  $m = 1$  under the intervention, the terms in this expression only depend on the combination matrix  $R$ . When new information arrives to the remaining agents, it is first mixed among these agents (corresponding to the first term  $R^\top$  on RHS), and then in the next iteration, it is mixed again but also gets forgotten due to the time discount factor  $\delta$  (corresponding to the second term  $(1 - \delta)R^{2\top}$  on RHS), and so on.

- Remember from (10) that  $\beta > 0$  scales the likelihood of observations, reflecting the weight agents place on their own observations originating from out-of-network sources. As a result, notice that in (61),  $\beta$  scales the individual informativeness  $d_\ell(\theta), \forall \ell \neq m$ . In other words, it amplifies the effect of self observations on the state of nature.

Next, we analyze the special cases introduced in the NBSL case under ASL framework.

### 5.2.1 SPECIAL CASES

**Fully-connected and federated architectures.** In Appendix C, we prove that the additional  $\delta$  and  $\beta$  parameters introduced for the ASL change the NBSL expression (47) to

$$\tilde{\lambda}_{k,\infty}(\theta) = \frac{1}{1 - (1 - \delta)(1 - v_m)} \left( \beta \sum_{\ell=2}^K v_\ell d_\ell(\theta) + v_m \log \frac{\mu_m(\theta^\circ)}{\mu_m(\theta)} \right) \quad (65)$$

Notice that as  $\delta \rightarrow 0$  and  $\beta \rightarrow 1$ , (65) recovers (47) as expected, and the following remarks from the NBSL case continue to hold here: (i) the influence of an agent  $m = 1$  is identical for each agent  $k \neq 1$  due to symmetry in the network topology, (ii) increasing the network centrality of agent  $m$  increases its causal influence, and (iii) increasing the network centrality and informativeness of the rest of the agents  $\ell \neq 1$  decreases the causal effect of  $m = 1$ . Moreover, since causal effect  $C_{m \rightarrow k}^{\text{ASL}}$  is a monotonic decreasing function of  $\tilde{\lambda}_{k,\infty}(\theta)$  by (22), Eq. (65) also implies the following conclusions:

- As stated after (61),  $\beta$  scales the informativeness of agents. Accordingly, (65) reveals that the causal effect is decreasing with increasing  $\beta$ . This is justifiable because as agents have greater reliance on their own observations about the state of nature, they are less influenced by other agents in the network. It is worth mentioning that in (65), the intervened log-belief ratio  $\log \frac{\mu_m(\theta^\circ)}{\mu_m(\theta)}$  behaves as “pseudo-informativeness”. It is scaled with the Perron entry  $v_m$  of agent  $m = 1$ , similar to how the rest of the agents’ informativeness is scaled with their own Perron entries. The difference is that the other agents’ informativeness are based on their log-likelihood ratios averaged with respect to the true distribution, whereas the intervened log-belief ratio can be arbitrary, possibly supplying misinformation.
- In the special case of the rest of the agents having no informativeness (i.e.,  $d_\ell(\theta) = 0, \forall \ell \neq m$ ), the limiting mean log-belief ratio vector in (65) turns into

$$\tilde{\lambda}_{k,\infty}(\theta) = \frac{1}{1 + \delta\left(\frac{1}{v_m} - 1\right)} \log \frac{\mu_m(\theta^\circ)}{\mu_m(\theta)}. \quad (66)$$

This is in contrast to the NBSL case, where,  $\tilde{\lambda}_{k,\infty}(\theta) = \log \frac{\mu_m(\theta^\circ)}{\mu_m(\theta)}$ . In other words, in steady-state of NBSL, the average beliefs of all agents become equal to the intervened fixed belief  $\mu_m$ , implying full controllability. In ASL, however, the controllability is reduced by a factor of  $(1 + \delta(1/v_m - 1)) \geq 1$  as shown in (66). In particular, increasing the forgetting factor  $\delta$  decreases controllability, especially when the network centrality of “controlling” agent  $m = 1$  is small. However, if agent  $m$  is highly central, i.e.,  $v_m \rightarrow 1$ , then the forgetting factor  $\delta$  has negligible effect on controllability.

- Considering the general case (65), note that unlike  $\beta$  which only affects non-intervened observations,  $\delta$  affects both intervened beliefs and non-intervened observations. Thus, to fully understand the impact of  $\delta$  on the overall causal effect, we must consider the exact values of the relevant parameters.

**Ring architecture.** In Appendix D, we prove that the additional  $\delta$  and  $\beta$  parameters introduced for the ASL change the NBSL expression (55) to

$$\tilde{\lambda}_{k,\infty}(\theta) = \left( \log \frac{\mu_m(\theta^\circ)}{\mu_m(\theta)} \right) \frac{(1 - \delta)^{k-2} (1 - \alpha)^{k-1}}{(1 - (1 - \delta)\alpha)^{k-1}} + \frac{\beta\alpha}{1 - (1 - \delta)\alpha} d_k(\theta) + \frac{\beta}{1 - \delta} \sum_{\ell=2}^{k-1} \frac{(1 - \delta)^{k-\ell} (1 - \alpha)^{k-\ell}}{(1 - (1 - \delta)\alpha)^{k-\ell+1}} d_\ell(\theta) \quad (67)$$

from which we can make the following observations:

- As  $\delta \rightarrow 0$  and  $\beta \rightarrow 1$ , the NBSL expression (55) for ring architectures is recovered. Similar to the earlier expressions, the causal effect decreases with increasing informativeness of agents along the path between  $m = 1$  and  $k$  ( $\ell = 2, \dots, k - 1$ ) (the last term on RHS), and also decreases with increasing informativeness of agent  $k$ . Informativeness is scaled by  $\beta$  as before.

- Recall that in NBSL, if the rest of the agents have no informativeness, it holds that  $\tilde{\lambda}_{k,\infty}(\theta) = \log \frac{\mu_m(\theta^\circ)}{\mu_m(\theta)}$ . In other words, agent  $m = 1$  can fully control other agents' beliefs. Instead, in ASL, if  $d_\ell(\theta) = 0, \forall \ell \neq 1$ , it holds that

$$\begin{aligned} \tilde{\lambda}_{k,\infty}(\theta) &= \frac{(1-\delta)^{k-2}(1-\alpha)^{k-1}}{(1-(1-\delta)\alpha)^{k-1}} \left( \log \frac{\mu_m(\theta^\circ)}{\mu_m(\theta)} \right) \\ &= \frac{1}{1-\delta} \left( 1 - \frac{1}{\frac{1-\alpha}{\delta} + \alpha} \right)^{k-1} \left( \log \frac{\mu_m(\theta^\circ)}{\mu_m(\theta)} \right). \end{aligned} \quad (68)$$

Observe that as the agent index  $k$  increases, the controllability decays at each hop by a factor of

$$\frac{\tilde{\lambda}_{k+1,\infty}(\theta)}{\tilde{\lambda}_{k,\infty}(\theta)} = 1 - \frac{1}{\frac{1-\alpha}{\delta} + \alpha} \in [0, 1]. \quad (69)$$

The decrease is higher when  $\delta$  is higher because the information from agent  $m = 1$  gets “partially forgotten” at each hop as  $k$  (i.e., the distance to agent 1) increases. However, in general, the informativeness of agents along the path is not 0, and they have a shadowing effect on agent  $m$ 's influence, as argued before. The forgetting factor  $\delta$  decreases this shadowing effect as well, particularly for agents far from agent  $k$ .

## 6. Causal Ranking of Agents

In the previous sections, we examined the bipartite influence between agents, that is, how much an agent  $m$  affects another agent  $k$  in the network. By calculating this influence for any pair of agents  $(m, k)$ , we can construct a  $K \times K$  influence matrix  $C$  with entries  $[C]_{mk} = C_{m \rightarrow k}$ . One is often interested in the overall influence of agent  $m$  on the network rather than its effect on individual agents. To that end, in this section, we describe a procedure to use  $C$  for ranking and quantifying the agents' cumulative effect over the network.

Since  $C$  is constructed from intervened belief dependent entries  $C_{m \rightarrow k}$ , an ordering based on  $C$  would be valid for a particular intervention. For an intervention dose independent ranking of agents, one can consider the matrix  $\bar{C}$ , which is formed with dose independent causal effects  $\bar{C}_{m \rightarrow k}$ :

$$\bar{C}_{m \rightarrow k} \triangleq C_{m \rightarrow k} \Big|_{\mu_m(\theta)=1/H} \quad (70)$$

where we extend the definition (41) for the NBSL case to the general case. For simplicity of the presentation, in the sequel, we focus on the NBSL case, even though our arguments keep holding for a general  $\bar{C}$  as well as an ordering based on intervention dose dependent matrix  $C$ , too. First, note that the causal effect for the NBSL case is given by

$$\begin{aligned} C_{m \rightarrow k}^{\text{NB}} &\stackrel{(16),(21)}{=} 1 - \frac{1}{1 + \sum_{\theta \in \Theta \setminus \{\theta^\circ\}} \exp\{-\tilde{\lambda}_{k,\infty}(\theta)\}} \\ &\stackrel{(35)}{=} 1 - \frac{1}{1 + \sum_{\theta \in \Theta \setminus \{\theta^\circ\}} \frac{\mu_m(\theta)}{\mu_m(\theta^\circ)} \exp\left\{-\left[(I - R^\top)^{-1} - I\right]d_{-m}(\theta)\right\}_k}. \end{aligned} \quad (71)$$

Here, setting  $\mu_m(\theta) = 1/H$  for any  $\theta \in \Theta$  based on (70) yields

$$\bar{C}_{m \rightarrow k}^{\text{NB}} = 1 - \frac{1}{1 + \sum_{\theta \in \Theta \setminus \{\theta^\circ\}} \exp \left\{ - \left[ ((I - R^\top)^{-1} - I) d_{-m}(\theta) \right]_k \right\}}. \quad (72)$$

Since all KL divergences are assumed to be finite ( $d_k(\theta) < \infty$ ) and the strongly connected graph assumption in Sec. 3 implies  $\rho(R) < 1$  (Ying and Sayed, 2016, Lemma 1),

$$\left\| ((I - R^\top)^{-1} - I) d_{-m}(\theta) \right\|_\infty < \infty. \quad (73)$$

Incorporating this into (72) implies that  $\bar{C}_{m \rightarrow k}^{\text{NB}} > 0, \forall m \neq k$ . Furthermore, regarding the diagonal elements of  $\bar{C}$ , it holds by definition that an intervention on agent  $m$  implies

$$\tilde{\lambda}_{m,\infty}(\theta) = \log \frac{\mu_m(\theta^\circ)}{\mu_m(\theta)} \quad (74)$$

As a result, if we set  $\mu_m(\theta) = 1/H, \forall \theta \in \Theta$

$$\begin{aligned} \bar{C}_{m \rightarrow m}^{\text{NB}} &\stackrel{(16),(21)}{=} 1 - \frac{1}{1 + \sum_{\theta \in \Theta \setminus \{\theta^\circ\}} \exp \{ -\tilde{\lambda}_{m,\infty}(\theta) \}} \\ &= 1 - \frac{1}{1 + \sum_{\theta \in \Theta \setminus \{\theta^\circ\}} \exp \{ 0 \}} \\ &= 1 - \frac{1}{H} \\ &> 0 \end{aligned} \quad (75)$$

Consequently, all entries of  $\bar{C}$  are positive, which implies that  $\bar{C}$  is a primitive matrix. Therefore, according to Perron's theorem (Pillai et al., 2005; Easley and Kleinberg, 2010),  $\bar{C}$  has a unique, real and positive eigenvalue  $\rho$  that dominates all other eigenvalues in magnitude. Moreover, the eigenvector  $q$  corresponding to  $\rho$  is unique up to a scaling and all its entries are positive, i.e.,

$$\bar{C}q = \rho q, \quad q_k > 0, \quad \forall k = 1, \dots, K. \quad (76)$$

The entry  $q_k$  is a measure of agent  $k$ 's overall influence over the network. The agents can be ranked with respect to these entries. We name the resulting algorithm **CausalRank** which is summarized in Algorithm 1. Importantly, the vector  $q$  — which is the output of Alg. 1 — differs from the network centrality eigenvector  $v$  in general. While  $v$  is determined solely by the combination matrix  $A$  (see (2)), as shown in previous sections, causal influences and hence  $q$  depend on the informativeness of agents as well. More specifically, (76) computes a *causal* eigenvector centrality that attributes higher importance to exerting influence on agents who are themselves influential. A possible alternative approach (which we call *average influence ranking* (AIR)) can treat all agents with equal regard in the averaging process by assigning the following ranking score to each agent  $m$ :

$$\text{AIR}(m) = \frac{1}{K-1} \sum_{k \neq m} \bar{C}_{m \rightarrow k} \quad (77)$$

---

**Algorithm 1** CausalRank Algorithm

---

- 1: **Input:** a network of  $K$  agents with indices  $\{1, 2, \dots, K\}$ , combination matrix  $A$ , set  $\Theta$  of  $H$  hypotheses, informativeness vector  $d(\theta)$
  - 2: **Initialize:**  $K \times K$ -dimensional influence matrix  $\bar{C}$
  - 3: **for** each agent  $m = 1, 2, \dots, K$  **do**
  - 4:     **for** each agent  $k = 1, 2, \dots, K$  **do**
  - 5:         **if**  $m = k$  **then**
  - 6:             set  $[\bar{C}]_{mm} := 1 - \frac{1}{H}$
  - 7:         **else**
  - 8:             compute the causal effect  $\bar{C}_{m \rightarrow k}$  with (72)
  - 9:             set  $[\bar{C}]_{mk} := \bar{C}_{m \rightarrow k}$
  - 10:         **end if**
  - 11:     **end for**
  - 12: **end for**
  - 13: find the largest eigenvalue  $\rho$  of  $\bar{C}$
  - 14: **Output:** the eigenvector  $q$  satisfying  $\bar{C}q = \rho q$
- 

In contrast, rather than employing a simple averaging, CausalRank seeks the equilibrium vector by assigning significant weights to those agents that have a higher influence on other influential agents. This concept bears resemblance to other methodologies based on eigenvector centrality, such as the PageRank algorithm (Brin and Page, 1998). While ranking websites, PageRank gives preferential treatment to links from more central websites.

It is also worth mentioning that CausalRank is distinct from the causal ordering methods for directed acyclic graphical models (Peters et al., 2017) since we are dealing with cyclic graphs with bidirectional links due to our time-series setting. Furthermore, CausalRank is not only useful for ranking, but also provides information on the strength of agents’ overall influence on others.

## 7. Causal Discovery from Observational Data

In Sec. 5, we derived the closed-form expressions (35) and (61) for the steady-state equilibrium of the network under interventions, which necessitate knowledge of the combination matrix  $A$  and the informativeness of agents  $d(\theta)$ . In practice, these parameters might not be readily available. The work (Shumovskaia et al., 2023) introduced the Graph Social Learning (GSL) algorithm, which can be used to recover  $A$  and  $d(\theta)$  using a sequence of publicly shared intermediate beliefs (a.k.a. *actions*)  $\{\psi_{k,i}\}$  in the *observational* setting of the ASL algorithm. Using observational data only can be especially useful in social network contexts where conducting experiments is not feasible. Nonetheless, (Shumovskaia et al., 2023) acknowledge that the algorithm may not perform well in real-world scenarios, mainly due to the limitations of the social learning model in accurately describing the real world. However, in many applications, some information about the underlying combination matrix  $A$  may already be available. For instance, in Twitter, the publicly available adjacency matrix can provide information about which user follows which other users.

Taking these aspects into account, in this section, we propose an algorithm that utilizes the adjacency matrix and a temporal sequence of publicly shared intermediate beliefs  $\{\boldsymbol{\psi}_{k,i}\}$  to estimate bipartite causal effects for both NBSL and ASL algorithms. Specifically, we leverage the graph of user connections to estimate combination matrix weights by using existing methods in the literature (e.g., averaging rule). Then, using the estimated combination matrix, we estimate the informativeness of agents by using belief update recursions. By inserting these to the closed-form expressions (35) and (61), we estimate the causal effects. To that end, observe that the intermediate log-belief ratios evolve based on a linear recursion due to (10):

$$\boldsymbol{\Lambda}_i = (1 - \delta)A^\top \boldsymbol{\Lambda}_{i-1} + \beta \mathbf{X}_i \quad (78)$$

where we are now defining the following  $K \times H$  matrices over all agents  $k \in \{1, \dots, K\}$  and hypotheses  $\theta_j \in \Theta$ :

$$[\boldsymbol{\Lambda}_i]_{kj} \triangleq \log \frac{\boldsymbol{\psi}_{k,i}(\hat{\boldsymbol{\theta}}^\circ)}{\boldsymbol{\psi}_{k,i}(\theta_j)}, \quad [\mathbf{X}_i]_{kj} \triangleq \log \frac{L_k(\boldsymbol{\xi}_{k,i}|\hat{\boldsymbol{\theta}}^\circ)}{L_k(\boldsymbol{\xi}_{k,i}|\theta_j)}. \quad (79)$$

Here,  $\hat{\boldsymbol{\theta}}^\circ$  is an estimate for the latent state of nature  $\boldsymbol{\theta}^\circ$  computed as follows after some time  $M$ :

$$\hat{\boldsymbol{\theta}}^\circ \triangleq \arg \max_{\theta \in \Theta} \sum_{k=1}^K \boldsymbol{\psi}_{k,M}(\theta). \quad (80)$$

The rationale behind (80) is that under proper assumptions we know from Theorems 1 and 2 that agents learn the true hypothesis with more confidence as  $M$  grows. Our goal is to infer the true combination matrix  $A$  and informativeness vector  $d(\theta)$  for each hypothesis from a sequence of  $M + 1$  matrices  $\{\boldsymbol{\Lambda}_M, \boldsymbol{\Lambda}_{M-1}, \dots, \boldsymbol{\Lambda}_0\}$  and the adjacency matrix of the agents.

We can estimate  $A$  by using existing procedures in the literature for forming combination matrices from adjacency matrices, e.g., by using the averaging or relative degree rules. For instance, the averaging rule assigns the same weight to all neighbors of an agent, i.e.,

$$[\hat{A}]_{\ell k} = \begin{cases} \frac{1}{|\mathcal{N}_k|}, & \text{if there is a link from } \ell \text{ to } k \text{ (i.e., } \ell \in \mathcal{N}_k) \\ 0, & \text{otherwise} \end{cases} \quad (81)$$

After forming the combination matrix estimate  $\hat{A}$ , we can insert it into (78) and average over available  $M$  samples to estimate the average log-likelihood ratios that correspond to the informativeness of agents using

$$\hat{\mathbf{D}} = \frac{1}{\beta M} \sum_{i=1}^M \left( \boldsymbol{\Lambda}_i - (1 - \delta) \hat{A}^\top \boldsymbol{\Lambda}_{i-1} \right). \quad (82)$$

Then, one can replace  $A$  and  $d_k(\theta_j)$  with  $\hat{A}$  and  $[\hat{\mathbf{D}}]_{kj}$  in Sec. 5 to obtain the causal effect estimate  $\hat{\mathbf{C}}_{m \rightarrow k}$ . The complete procedure is summarized in Alg. 2. Essentially, it combines our causality results in Sec. 5 with a straightforward adjustment to the GSL algorithm from (Shumovskaia et al., 2023).

The graph causality learning (GCL) algorithm (Alg. 2) only requires a sequence of shared intermediate beliefs (actions) and the knowledge of adjacency matrix. This enhances its

---

**Algorithm 2** Graph Causality Learning (GCL)

---

- 1: **Input:** a sequence of shared beliefs  $\{\psi_{k,i}\}$  for  $M + 1$  time instants, the graph topology of agents
- 2: **Parameters:** for NBSL  $\delta = 0, \beta = 1$ , for ASL  $\delta \in (0, 1), \beta > 0$
- 3: set true state of nature estimate:  $\hat{\theta}^\circ = \arg \max_{\theta \in \Theta} \sum_{k=1}^K \psi_{k,M}(\theta)$
- 4: form left-stochastic combination matrix estimate  $\hat{A}$  from input adjacency matrix, e.g., by (81)
- 5: **for**  $i = 0, 1, \dots, M$  **do**
- 6:   for each agent  $k$  and hypothesis  $\theta_j$ , set the entry:

$$[\mathbf{\Lambda}_i]_{kj} = \log \frac{\psi_{k,i}(\hat{\theta}^\circ)}{\psi_{k,i}(\theta_j)} \quad (83)$$

- 7: **end for**
- 8: estimate the informativeness:

$$\hat{D} = \frac{1}{\beta M} \sum_{i=1}^M \left( \mathbf{\Lambda}_i - (1 - \delta) \hat{A}^\top \mathbf{\Lambda}_{i-1} \right) \quad (84)$$

- 9: for any given agent pair  $(m, k)$ , compute approximate causal effect  $\hat{C}_{m \rightarrow k}$  by replacing  $A$  and  $d_k(\theta_j)$  with  $\hat{A}$  and  $[\hat{D}]_{kj}$  in the original expression (15) for  $C_{m \rightarrow k}$  (which also requires using (35) for NBSL and (61) for ASL)
  - 10: **Output:**  $\hat{C}_{m \rightarrow k}$
- 

practicality and makes it advantageous in terms of privacy for scenarios where only limited information is publicly accessible. For example, in a network of Twitter users, shared beliefs (opinions) in the form of tweets (posts) and the knowledge of who follows whom can usually be accessed by all users, while the external exposure to information (e.g., from mass media channels distinct from Twitter) may not be available. Therefore, the GCL algorithm can be useful for analyzing social media content while respecting privacy. In the next result, we provide a performance bound on the GCL algorithm.

**Theorem 3 (Causal influence estimation)** *For sufficiently small combination matrix estimation errors and  $\delta$  values, the error in causal influence estimation decreases with increasing number of samples  $M$  in expectation, namely,*

$$\mathbb{E} \left| C_{m \rightarrow k} - \hat{C}_{m \rightarrow k} \right| = O(1/\sqrt{M}) \quad (85)$$

for both NBSL and ASL under any intervention strength  $\mu_m$ .

**Proof** See Appendix E. ■

We show the practical usefulness of Alg. 2 by means of a real-world application to social media data in Sec. 9. Furthermore, a detailed analysis of the time complexity of the algorithms discussed in this paper is provided in Appendix F.




Figure 7: The strongly connected network architecture used in Sec. 8. Each agent also has a self-loop, which is omitted for visual simplicity.

Agent	$\nu_k$	$d_k(\theta')$
1	0.8	0.32
2	0.6	0.18
3	0.2	0.02
4	0.6	0.18
5	0	0
6	0	0
7	0.4	0.08
8	0.4	0.08
9	0.2	0.02
10	0.6	0.18
11	0.8	0.32

Table 1: Mean  $\nu_k$  of the observations under alternative hypothesis and the corresponding informativeness levels  $d_k(\theta')$  for each agent  $k$ .

## 8. Computer Simulations

For our numerical simulations, we first study a network of  $K = 11$  agents, interconnected with the strongly connected graph topology in Fig. 7. The agents observe data drawn from a Gaussian distribution and aim to distinguish the true state  $\theta^\circ$  from  $H = 2$  possible hypotheses. Under the true state, each agent  $k$  observes data that follows a Gaussian distribution with zero mean and unit variance, expressed as:

$$L_k(\xi|\theta^\circ) = \frac{1}{\sqrt{2\pi}} \exp\left\{-\frac{1}{2}\xi^2\right\}. \quad (86)$$

Under the alternative hypothesis  $\theta' \neq \theta^\circ$ , we assume that the data still has unit variance for all agents, but the mean vector  $\nu_k$  changes as shown in Table 1. Therefore, the informativeness of each agent, which is equal to the KL divergence between  $L_k(\xi|\theta^\circ)$  and  $L_k(\xi|\theta')$ , is given in Table 1 and is calculated as follows:

$$d_k(\theta') = D_{\text{KL}}(L_k(\xi|\theta^\circ)||L_k(\xi|\theta')) = \frac{1}{2}\nu_k^2. \quad (87)$$

Notably, agents 5 and 6 have no informativeness, that is, they are not able to learn the truth without cooperating with the other agents. Initially, we assume that the agents observe spatially independent data. In other words, the covariance matrix is an identity matrix.

We start with the NBSL case ( $\delta = 0, \beta = 1$ ). The right panel in Fig. 8 shows the combination matrix that is derived from the averaging rule applied to the graph topology in Fig. 7. Notice that the averaging rule generates a matrix whose entries are constant column-wise. The left panel in Fig. 8 shows the matrix of bipartite causal effects where the entry in  $m$ -th row  $k$ -th column represents  $\bar{C}_{m \rightarrow k}^{\text{NB}}$  (see (72) for the explicit formula).

Upon comparing the two heat maps in Fig. 8, it becomes apparent that the combination matrix entries do not reveal the causal relationships directly. For example, despite the




Figure 8: (*Left*): Bipartite causal influence matrix. (*Right*): Combination matrix corresponding to the network topology in Fig. 7 formed with averaging rule.

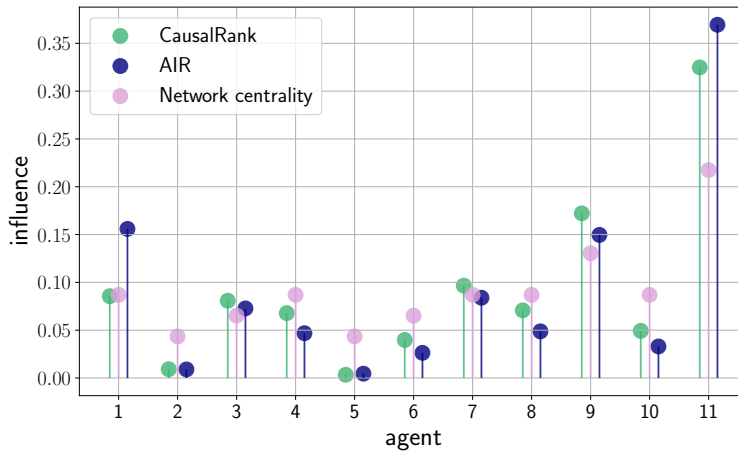


Figure 9: Ranking of agents based on their causal influence (for both CausalRank and AIR) and their network-topology based eigenvector centrality. All ranking scores are normalized to sum up to one.

absence of a direct connection in the combination matrix (as indicated by 0 entries), agent 11 exerts significant influence on agents 2 and 8. This phenomenon highlights the importance of taking the ripple effects over a network into account. Furthermore, the influence of agent 1 on agent 5 is notably high. Given the zero informativeness of agent 5, this finding aligns with our expectations, as low-informativeness agents are easier to control (remember the discussion in Sec. 5.1). Agent 5 being a low-informativeness agent also facilitates the propagation of influence from agent 1 to agent 6 via agent 5. Intriguingly, despite the absence of a direct connection between agents 1 and 6, this indirect influence is more substantial than the influence of agent 5 on agent 6. This shows that mixing of information over a network necessitates an understanding of causal influence beyond local interactions.




Figure 10: Average influence on agent  $k = 4$  from its neighborhood of 1,2 and 3 hop distances with respect to changing values of forgetting factor  $\delta$ .

Next, in Fig. 9, by using the matrices in Fig. 8, we compare the overall influences of agents using three methods: CausalRank, AIR, and network eigenvector centrality. Notably, the CausalRank and AIR metrics yield similar results as they both use the bipartite causal relations matrix for causal ranking. For instance, agents 2 and 5 possess relatively low rankings in both of these metrics. The network eigenvector centrality, on the other hand, only relies on the combination matrix, and often deviates from these two metrics. Specifically, it assigns relatively higher scores to agents 2 and 5, and a comparatively lower score to agent 11. Moreover, an interesting distinction between AIR and CausalRank becomes apparent when considering the case of agent 9. We can see from the causal influence matrix in Fig. 8 that agent 9 has a substantial impact on agent 11 — the most influential agent (see Fig. 9). Consequently, agent 9’s CausalRank score surpasses its AIR score. This can be attributed to CausalRank’s consideration of the significance of influencing agent 11. Unlike AIR, which assigns uniform weights, CausalRank assigns a higher weight to influences on more influential agents.

To gain insights into the influence of the forgetting factor  $\delta$  in the ASL case, we focus our attention on agent  $k = 4$ . In Fig. 10, we present the average influence exerted by agent 4’s neighbors that are 1, 2, and 3 hops away. It is clear from Fig. 10 that the influence of distant agents diminishes with increasing  $\delta$ . This is because increasing  $\delta$  increases the significance of recent observations, and since information from distant agents loses its recency by the time it arrives at agent 4, this implies assigning less importance to information from those distant agents.

In the simulations conducted so far, we have considered dose-independent causal effects  $\bar{C}_{m \rightarrow k}$ . Figure 12 demonstrates the dependency of  $C_{m \rightarrow k}$  on the intervened belief  $\mu_m(\theta)$  for  $\theta \neq \theta^0$ , as described in equation (37). We use  $m = 11$  and  $k = 6$  for this plot, which illustrates the theoretical result presented in equation (37).




Figure 11: Causal influence estimation error with respect to increasing number of time samples  $M$ . Averaged over 10 independent experiments.

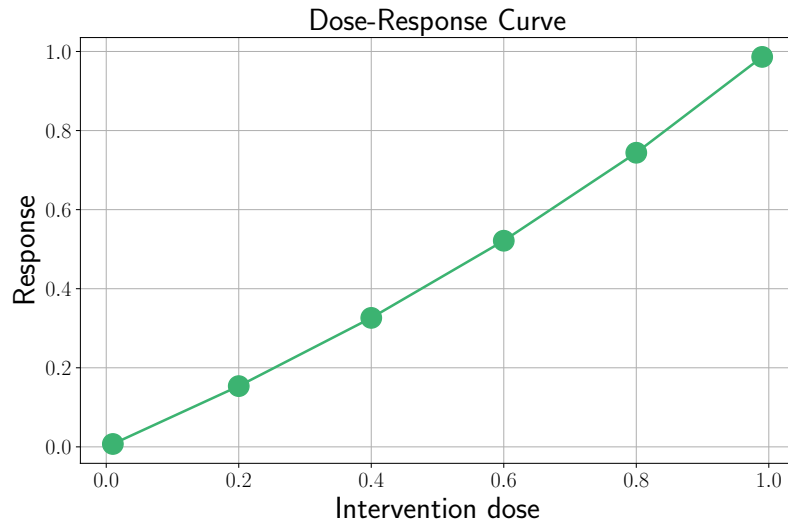


Figure 12:  $C_{m \rightarrow k}^{\text{NB}}$  value from (37) for  $m = 11$  and  $k = 6$  with respect to increasing intervention strength  $\mu_m(\theta)$  for  $\theta \neq \theta^\circ$ .




Figure 13: Correlation and causation between agents 6 and 11 with respect to varying dependence between the streaming observations they receive.

Then, we fix  $\delta = 0.1$ , and use the GCL algorithm (Alg. 2) in order to estimate the causal effects using observational data (shared beliefs) as described in Sec. 7. The norm disagreement of the causal influence matrix formed with estimates and the true causal influences, averaged over 10 Monte Carlo simulations, is given in Fig. 11. Observe that the error is decreasing as the number of samples  $M$  increases, which supports Theorem 3.

Next, we illustrate the distinction between causality and correlation by again considering agents  $m = 11$  and  $k = 6$ . The joint distribution of their data is changed by introducing varying levels of correlation to the observations that these agents are receiving. Fig. 13 shows that as the correlation in data increases, the correlation of the asymptotic beliefs of these agents also changes. However, the causal effects (both the effect of agent 6 on 11 and that of agent 11 on 6) remain constant. This shows that external observations can act as a correlation inducing confounding factor. Yet, our method maintains consistent results, which shows its robustness against non-causal factors.

In order to demonstrate the practical usefulness of CausalRank, we compare the robustness of the considered ranking methods against follower-inflation attacks. Such attacks are practically relevant in social networks because visible popularity metrics can be monetized through sponsorships, advertising and other reputation-based opportunities, which creates incentives to purchase fake indicators of social media influence. We use the same network topology and informativeness levels as in Fig. 7 and Table 1, but augment the network by introducing additional agents (bot accounts) that act as artificial followers of a randomly selected (attacked) agent. These bots are non-informative and are designed only to increase the apparent popularity of the attacked agent without contributing meaningful influence over the other agents. Fig. 14 shows the average rank of the attacked agent as the number of bots increases. The reported values are averaged over 5000 independent repetitions. As expected, the ranking based on the number of followers is highly sensitive to this type of manipulation, while AIR is also noticeably affected since it rewards influence over a larger number of agents irrespective of their importance. In contrast, the average rank produced by CausalRank remains close to the theoretical expectation of 6 for  $K = 11$  honest agents. We

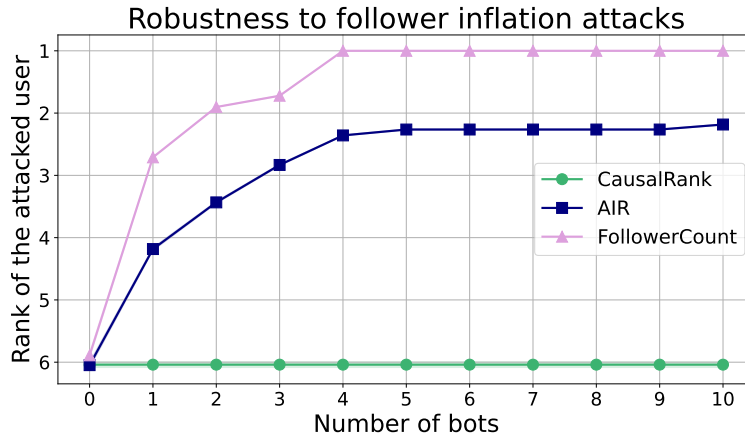


Figure 14: Influence rank of an attacked agent whose followers are artificially inflated with bot accounts. The attacked agents are randomly selected, and the experiment is repeated 5000 times.




Figure 15: CPU-time to compute influence matrix  $C$  and apply CausalRank for increasing network sizes.

provide additional experiments in Appendix G with networks of different architectures and sizes to show that this effect continues to hold across those settings.

Finally, we benchmark the computational time required to calculate the causal influence matrix  $C$  and CausalRank as a function of network size  $K$  for NBSL. For each value of  $K$  between 10 and 100, we generate 20 independent experiments, each on a distinct strongly-connected Erdős–Renyi graph. In Fig. 15, we report the total CPU time required for this computation in seconds, averaged over the independent experiments. As shown in Appendix F, the influence matrix calculation can be accelerated by reusing common matrix operations across different agents in the network. We provide additional plots illustrating this effect and also benchmark the CPU time for ASL in Appendix G.

## 9. Application to Social Media Data

In this section, we use the GCL algorithm (Alg. 2) on real-world data to assess the influence of Twitter users. Our approach distinguishes itself from prior works (Quercia et al., 2011; Jain and Sinha, 2020; Smith et al., 2021; Shumovskaia et al., 2023), which typically rely on some descriptive statistics to measure influence in Twitter. More specifically, in our approach,

- All input requirements are publicly available, i.e., publicly shared posts (tweets) by users and the information of who follows whom. This offers a significant advantage in terms of privacy, as we do not require any private feature about users.
- Going beyond providing a mere ranking of influential users, we also quantify the bipartite causal relations.
- We leverage natural language processing tools to extract meaningful information from the content of users’ posts to form belief inputs, rather than relying on traditional simpler metrics such as the posting frequency. In contrast to the binary treatment approach, adopting the continuous treatment metric in (15) for our causality definition enables us to allow for the use of continuous variables such as opinions and sentiment scores.

Note that we utilize Alg. 2 for the NBSL model ( $\delta = 0, \beta = 1$ ) as-is, without employing additional techniques to enhance its accuracy for real-world modeling. Our intention is to demonstrate the practical usefulness of our algorithm rather than striving to develop the most advanced practical algorithm available.

**Network structure.** Performance evaluation of the influence estimation algorithms in real-world social networks is challenging due to the absence of ground truth regarding influential users. There is also no ground truth reference for the confidence scores assigned by users to one another (i.e., combination weights) or for the information the users are obtaining from out-of-network resources (i.e., informativeness levels). Therefore, we utilize the framework from (Shumovskaia et al., 2023), namely, a sub-network consisting of  $K = 20$  Twitter users, as illustrated in Fig. 16. Notably, this sub-network incorporates Elon Musk (User 1), a public figure with 140 million followers across Twitter, who is reportedly influential on cryptocurrency prices (Ante, 2023). All users within the sub-network actively share posts related to cryptocurrencies and bitcoin-related topics. Furthermore, we highlight another user (User 2), who has 1,167 followers across Twitter, is notable for being followed by User 1, as depicted in Fig. 16. Importantly, the sub-network exhibits a strong connectivity among its members.

**Opinion processing.** The Twitter API is leveraged in order to collect the posts (tweets) of users between 01.01.2017 and 01.05.2022 relevant to crypto-currency discussions, using query keywords such as “coin”, “bitcoin”, or “crypto-currency”. To quantify the contextual information of these posts to form the input beliefs, sentiment analysis based on neural language models (Loureiro et al., 2022) is utilized. We refer to Fig. 17 for some illustrative examples. The sentiment scores obtained through natural language processing ranges from 0 to 1, signifying the degree of positive attitude towards Bitcoin. These scores correspond to the beliefs of the agents on the hypothesis of “Bitcoin is good/useful”. We consider two hypotheses, i.e.,  $H = 2$ , where the counter-hypothesis is “Bitcoin is bad/harmful”.

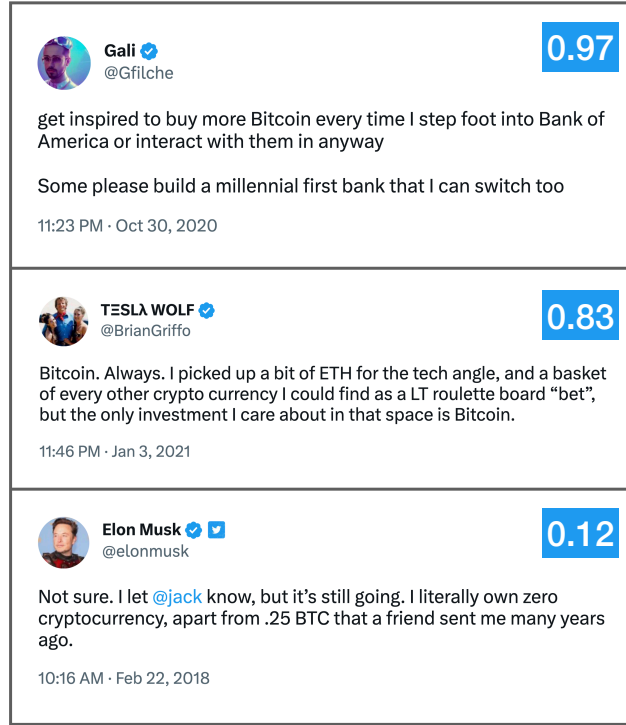


Figure 17: Sample tweets from the users in the sub-network. The number on the upper RHS quantifies the positive attitude towards Bitcoin.




Figure 16: Sampled Twitter sub-network with  $K = 20$  users. An arrow from a user  $k$  to  $\ell$  means that user  $\ell$  is following user  $k$ . The connectivity density of the graph is approximately 0.24.

We then integrated these beliefs obtained from users' tweets, along with the sub-network topology of who follows whom, into Alg. 2. Specifically, we employed NBSL modeling, i.e.,  $\delta = 0$  and  $\beta = 1$ . We binned the sequence of data to days, that is, each  $i$  in our algorithms corresponds to one day. Combination weights were estimated using an averaging rule on the sub-network topology.

**Bipartite causality.** Inserting the observational input into Alg. 2, the resulting average causal derivative effect matrix is shown in Fig. 18 in the form of a heat map. To facilitate comparison, we also include the adjacency matrix, which describes the connections between users. In these plots, the indices 1 and 2 correspond to User 1 and User 2, respectively.

Upon observing the heat map, it is evident that User 1 holds significant influence over all other users, as indicated by the high values in




Figure 18: (*Left*): Bipartite causal influence matrix. (*Right*): Adjacency matrix corresponding to the sub-network topology in Fig. 16.




Figure 19: Ranking of agents based on their causal influence and their corresponding network eigenvector centrality. Both ranking scores are summing up to one.

the 1th row, which aligns with our expectations. However, notice that the adjacency matrix does not precisely mirror the causal relationships. For instance, User 2 is followed by User 1, yet their influence on User 1, as depicted in the heat map, is relatively low. On the other hand, User 14 exerts a substantial impact on User 2, despite not being directly followed by User 2. This fact may arise from the fact that User 14 holds one of the highest influences on User 1 among all the users in this particular sub-network. These observations highlight the fact that the nature of influence dynamics within real-world social networks cannot simply be explained with direct follower relations.

**Causal impact ranking.** Once the influence matrix is determined, we apply the CausalRank algorithm (Alg. 1) to rank the agents based on their overall influence within the sub-network.

The resulting plot is depicted in Fig. 19. Notably, User 1 emerges as the most influential agent, aligning with our initial expectations.

However, an intriguing observation can be made regarding User 2. Despite having a high eigenvector centrality, their causal impact score appears relatively small. This phenomenon arises because the causal effect is not solely determined by network centrality but also takes into account the informativeness of the agents. For instance, if a user primarily retweets (reposts) what their neighbors are tweeting, such users tend to possess low informativeness, decreasing their causal impact score. Thus, even though User 2 may have a high centrality within the considered sub-network primarily due to being followed by User 1, their causal influence on their neighbors is low and does not propagate to other users, leading to a relatively small causal impact.

## 10. Conclusion

In this study, we analyzed causal influences among agents that are connected over a network and whose interactions occur over time. Using social learning models, we derived expressions for the causal relationships between pairs of agents. These expressions offer key insights into the diffusion of influences across a social network. We also proposed the CausalRank algorithm for ranking the overall influence of agents, which allows discovering highly influential actors within a network. Furthermore, to enhance the practical usage of our results, we proposed the graph causality learning algorithm (GCL) that learns the necessary model parameters from raw observational data in order to estimate the causal effects. We demonstrated how GCL can be applied in practice through an application to real social media data.

The social learning models we considered in this work are useful for both modeling opinion formation over social networks as well as for designing multi-agent decision-making systems. Therefore, potential applications range from the analysis of human social networks, such as those on social media platforms, to cooperative decision-making processes of socially intelligent machines like networks of drones or sensors. In addition to these, our results can be useful for applications that involve time-series networked interactions, since they provide insights on the diffusion of influence across graphs.

A key challenge in our experimental evaluation was the lack of definite ground truth for influential agents in real-world social networks. To address this issue, we conducted our social media experiment on data in which at least the most influential agent was readily identifiable. We also attempted to design controlled experiments using LLMs, whose capabilities have advanced rapidly since the initial submission of this work. However, our exploratory results indicated that their inherent biases remain too strong to support fully reliable evaluations of the algorithms at this stage. We therefore leave this direction for future work.

## Acknowledgments

The authors thank the anonymous reviewers and the action editor for their insightful and constructive feedback.

## Appendix A. Connection of (41) to Average Causal Derivative Effect

In this appendix, we demonstrate how definition (41) for the causal effect summary  $\bar{C}_{m \rightarrow k}^{\text{NB}}$  can also be interpreted as the increment in causal effect  $C_{m \rightarrow k}^{\text{NB}}$  after an infinitesimal change in the intervention strength  $\mu_m(\theta)$ . In the literature, this is referred to as the average causal *derivative* effect, as it computes the derivative in causal effect with respect to the intervention strength (Peters et al., 2017, Chap. 6). However, a simple derivative of  $C_{m \rightarrow k}^{\text{NB}}$  with respect to  $\mu_m$  fails to produce a dose-independent summary, given that  $C_{m \rightarrow k}^{\text{NB}}$  is not a linear function of  $\mu_m$ , which can be seen from (71). Nonetheless, we introduce the function

$$f(C_{m \rightarrow k}^{\text{NB}}) \triangleq \frac{C_{m \rightarrow k}^{\text{NB}}}{1 - C_{m \rightarrow k}^{\text{NB}}} \stackrel{(71)}{=} \sum_{\theta \in \Theta \setminus \{\theta^\circ\}} \frac{\mu_m(\theta)}{\mu_m(\theta^\circ)} \exp \left\{ - \left[ ((I - R^\top)^{-1} - I) d_{-m}(\theta) \right]_k \right\}, \quad (88)$$

and notice that  $f(C_{m \rightarrow k}^{\text{NB}}) \in [0, \infty)$  is a linear function of the intervened belief ratio vector  $s_m$ , defined as

$$s_m(\theta) \triangleq \frac{\mu_m(\theta)}{\mu_m(\theta^\circ)}, \quad s_m \triangleq [s_m(\theta_1), \dots, s_m(\theta_H)]^\top. \quad (89)$$

Here, the vector  $s_m$  quantifies the amount of relative misinformation the intervention  $\mu_m$  produces. Therefore, the gradient

$$\nabla_{s_m} f(C_{m \rightarrow k}^{\text{NB}}) = \sum_{\theta \in \Theta \setminus \{\theta^\circ\}} \frac{\partial f(C_{m \rightarrow k}^{\text{NB}})}{\partial s_m(\theta)} \stackrel{(88)}{=} \sum_{\theta \in \Theta \setminus \{\theta^\circ\}} \exp \left\{ - \left[ ((I - R^\top)^{-1} - I) d_{-m}(\theta) \right]_k \right\} \quad (90)$$

is independent of the intervention dose  $\mu_m$ , and satisfies

$$\nabla_{s_m} f(C_{m \rightarrow k}^{\text{NB}}) = f(\bar{C}_{m \rightarrow k}^{\text{NB}}). \quad (91)$$

In other words, to find the gradient of  $f(C_{m \rightarrow k}^{\text{NB}})$  with respect to  $s_m$ , setting

$$\frac{\mu_m(\theta)}{\mu_m(\theta^\circ)} = 1 \implies \mu_m(\theta) = \frac{1}{H}, \quad \forall \theta \in \Theta \quad (92)$$

in (88) as it was done in (41) for finding  $\bar{C}_{m \rightarrow k}^{\text{NB}}$ , is sufficient. The reason we are interested in  $\nabla_{s_m} f(C_{m \rightarrow k}^{\text{NB}})$  is the following: Notice from (88) that  $f(C_{m \rightarrow k}^{\text{NB}})$  is a monotonic increasing function of  $C_{m \rightarrow k}^{\text{NB}}$ . Also,  $s_m(\theta)$  is clearly a monotonic function of  $\mu_m(\theta)$  — see (89). Therefore,  $\nabla_{s_m} f(C_{m \rightarrow k}^{\text{NB}})$  can be considered as some proxy for the derivative of the causal effect  $C_{m \rightarrow k}^{\text{NB}}$  with respect to  $\mu_m(\theta)$ . Consequently, we conclude that setting  $\mu_m(\theta)$  as a uniform belief in  $C_{m \rightarrow k}^{\text{NB}}$  effectively parallels finding the average causal derivative effect. It is also important to note here that the average causal derivative effect is the analogue of the binary treatment causal effect. For binary interventions, the causal effect is given by the difference between the outcomes under the two possible interventions. When the intervention varies continuously, the corresponding quantity is the derivative with respect to intervention strength, which captures the infinitesimal change in the causal effect.

**Appendix B. Proof of (61)**

Theorem 2, which establishes convergence in distribution for the pre-intervention ASL case does not require the strongly connected graph assumption. In fact, it holds as long as the observations that the agents receive are i.i.d. over time. This condition is satisfied under the post-intervention case as well. Hence, the log-beliefs under intervention converge in distribution, i.e.,

$$\tilde{\lambda}_i(\theta) \xrightarrow{\text{dist.}} \beta \sum_{j=1}^{\infty} (1-\delta)^{j-1} (\tilde{A}^\top)^j \tilde{\mathbf{x}}_j(\theta). \quad (93)$$

Then, following the same arguments from Corollary 1, the limiting expectation becomes

$$\tilde{\lambda}_\infty(\theta) = \frac{\beta}{1-\delta} \left( (I - (1-\delta)\tilde{A}^\top)^{-1} - I \right) \tilde{d}(\theta) \quad (94)$$

where we define the vector of expected LLRs and the vector of limiting log-belief ratio expectations from across the network:

$$\tilde{d}(\theta) \triangleq \left[ \frac{\delta}{\beta} \log \frac{\mu_1(\theta^\circ)}{\mu_1(\theta)}, d_2(\theta), \dots, d_K(\theta) \right]^\top, \quad \tilde{\lambda}_\infty(\theta) \triangleq \left[ \log \frac{\mu_1(\theta^\circ)}{\mu_1(\theta)}, \lambda_{2,\infty}(\theta), \dots, \lambda_{K,\infty}(\theta) \right]^\top. \quad (95)$$

Using the block matrix form of  $\tilde{A}$  from (24), we have

$$I - (1-\delta)\tilde{A}^\top = \begin{bmatrix} \delta & 0 \\ -(1-\delta)r & I - (1-\delta)R^\top \end{bmatrix}, \quad (96)$$

which implies

$$(I - (1-\delta)\tilde{A}^\top)^{-1} = \begin{bmatrix} \frac{1}{\delta} & 0 \\ \frac{1-\delta}{\delta} (I - (1-\delta)R^\top)^{-1} r & (I - (1-\delta)R^\top)^{-1} \end{bmatrix}. \quad (97)$$

Inserting this into (94), we arrive at

$$\tilde{\lambda}_\infty(\theta) = \frac{\beta}{1-\delta} \begin{bmatrix} \frac{1-\delta}{\delta} & 0 \\ \frac{1-\delta}{\delta} (I - (1-\delta)R^\top)^{-1} r & (I - (1-\delta)R^\top)^{-1} - I \end{bmatrix} \begin{bmatrix} \frac{\delta}{\beta} \log \frac{\mu_1(\theta^\circ)}{\mu_1(\theta)} \\ d_{-m}(\theta) \end{bmatrix} \quad (98)$$

which again verifies that  $\tilde{\lambda}_{1,\infty}(\theta) = \log \frac{\mu_1(\theta^\circ)}{\mu_1(\theta)}$  and proves relation (61).

**Appendix C. Proof of (65)**

Recall from (42)–(43) that for this fully connected network:

$$A = v\mathbf{1}_K^\top, \quad R = v_{-m}\mathbf{1}_{K-1}^\top, \quad r = v_1\mathbf{1}_{K-1}. \quad (99)$$

As a result, it holds that

$$\begin{aligned} (I - (1 - \delta)R^\top)^{-1} &= \sum_{i=0}^{\infty} (1 - \delta)^i (R^\top)^i \\ &\stackrel{(45)}{=} I + (1 - \delta)\mathbf{1}_{K-1}v_{-m}^\top + (1 - \delta)^2(1 - v_1)\mathbf{1}_{K-1}v_{-m}^\top + \dots \\ &= I + (1 - \delta) \left( \sum_{i=0}^{\infty} (1 - \delta)^i (1 - v_1)^i \right) \mathbf{1}_{K-1}v_{-m}^\top \\ &= I + \frac{1 - \delta}{1 - (1 - \delta)(1 - v_1)} \mathbf{1}_{K-1}v_{-m}^\top. \end{aligned} \quad (100)$$

This implies that

$$\begin{aligned} (I - (1 - \delta)R^\top)^{-1}r &= \left( I + \frac{1 - \delta}{1 - (1 - \delta)(1 - v_1)} \mathbf{1}_{K-1}v_{-m}^\top \right) v_1\mathbf{1}_{K-1} \\ &\stackrel{(a)}{=} v_1 \left( \mathbf{1}_{K-1} + \frac{1 - \delta}{1 - (1 - \delta)(1 - v_1)} (1 - v_1)\mathbf{1}_{K-1} \right) \\ &= \frac{v_1}{1 - (1 - \delta)(1 - v_1)} \mathbf{1}_{K-1} \end{aligned} \quad (101)$$

where (a) follows from the fact that  $\sum_{\ell=2}^K v_\ell = 1 - v_1$ . Eq. (100) also implies that

$$\left( (I - (1 - \delta)R^\top)^{-1} - I \right) d_{-m}(\theta) = \frac{1 - \delta}{1 - (1 - \delta)(1 - v_1)} \left( \sum_{\ell=2}^K v_\ell d_\ell(\theta) \right) \mathbf{1}_{K-1} \quad (102)$$

Incorporating (101) and (102) into (61) concludes the proof.

**Appendix D. Proof of (67)**

Recall expressions (51) and (52) for  $A$  and  $R$  in the ring special case. Also observe that

$$r = \left[ 1 - \alpha, 0, \dots, 0 \right]^\top. \quad (103)$$

Accordingly, it holds that

$$(I - (1 - \delta)R^\top) = \begin{bmatrix} 1 - (1 - \delta)\alpha & 0 & \dots & 0 & 0 \\ -(1 - \delta)(1 - \alpha) & 1 - (1 - \delta)\alpha & \dots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & 1 - (1 - \delta)\alpha & 0 \\ 0 & 0 & \dots & -(1 - \delta)(1 - \alpha) & 1 - (1 - \delta)\alpha \end{bmatrix} \quad (104)$$

The inverse of this Toeplitz matrix is given by the following lower diagonal matrix:

$$(I - (1 - \delta)R^\top)^{-1} = \begin{bmatrix} \frac{1}{1 - (1 - \delta)\alpha} & 0 & \cdots & 0 & 0 \\ \frac{(1 - \delta)(1 - \alpha)}{(1 - (1 - \delta)\alpha)^2} & \frac{1}{1 - (1 - \delta)\alpha} & \cdots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \frac{(1 - \delta)^{K-3}(1 - \alpha)^{K-3}}{(1 - (1 - \delta)\alpha)^{K-2}} & \frac{(1 - \delta)^{K-4}(1 - \alpha)^{K-4}}{(1 - (1 - \delta)\alpha)^{K-3}} & \cdots & \frac{1}{1 - (1 - \delta)\alpha} & 0 \\ \frac{(1 - \delta)^{K-2}(1 - \alpha)^{K-2}}{(1 - (1 - \delta)\alpha)^{K-1}} & \frac{(1 - \delta)^{K-3}(1 - \alpha)^{K-3}}{(1 - (1 - \delta)\alpha)^{K-2}} & \cdots & \frac{(1 - \delta)(1 - \alpha)}{(1 - (1 - \delta)\alpha)^2} & \frac{1}{1 - (1 - \delta)\alpha} \end{bmatrix}. \quad (105)$$

Then, the matrix-vector products in the general formula (61) become

$$(I - (1 - \delta)R^\top)^{-1}r = \left[ \frac{1 - \alpha}{1 - (1 - \delta)\alpha}, \frac{(1 - \delta)(1 - \alpha)^2}{(1 - (1 - \delta)\alpha)^2}, \dots, \frac{(1 - \delta)^{K-2}(1 - \alpha)^{K-1}}{(1 - (1 - \delta)\alpha)^{K-1}} \right]^\top \quad (106)$$

and

$$\left( (I - (1 - \delta)R^\top)^{-1} - I \right) d_{-m}(\theta) = \begin{bmatrix} \frac{\alpha(1 - \delta)}{1 - (1 - \delta)\alpha} d_2(\theta) \\ \frac{\alpha(1 - \delta)}{1 - (1 - \delta)\alpha} d_3(\theta) + \frac{(1 - \alpha)(1 - \delta)}{(1 - (1 - \delta)\alpha)^2} d_2(\theta) \\ \vdots \\ \frac{\alpha(1 - \delta)}{1 - (1 - \delta)\alpha} d_K(\theta) + \cdots + \frac{(1 - \alpha)^{K-2}(1 - \delta)^{K-2}}{(1 - (1 - \delta)\alpha)^{K-1}} d_2(\theta) \end{bmatrix} \quad (107)$$

Inserting these into the general expression (61) concludes the proof.

## Appendix E. Proof of Theorem 3

If we denote the error in estimating the combination matrix by  $E \triangleq \hat{A} - A$ , then combining (78) and (82) yields:

$$\hat{D} = \frac{1}{\beta M} \sum_{i=1}^M \left( \beta \mathbf{X}_i - (1 - \delta) E^\top \mathbf{\Lambda}_{i-1} \right). \quad (108)$$

Therefore, the error in informativeness can be decomposed as

$$\begin{aligned}
 \mathbb{E}\|\widehat{\mathbf{D}} - D\|_{\text{F}}^2 &= \frac{1}{\beta^2 M^2} \mathbb{E}\left\| \sum_{i=1}^M \beta(\mathbf{X}_i - D) - (1 - \delta)E^\top \mathbf{\Lambda}_{i-1} \right\|_{\text{F}}^2 \\
 &= \frac{1}{M^2} \mathbb{E}\left\| \sum_{i=1}^M (\mathbf{X}_i - D) \right\|_{\text{F}}^2 + \frac{(1 - \delta)^2}{\beta^2 M^2} \mathbb{E}\left\| \sum_{i=1}^M E^\top \mathbf{\Lambda}_{i-1} \right\|_{\text{F}}^2 \\
 &\quad - \frac{2(1 - \delta)}{\beta M^2} \mathbb{E}\left[ \text{Tr}\left( \sum_{i=1}^M (\mathbf{X}_i - D) \right) \left( \sum_{i=1}^M (E^\top \mathbf{\Lambda}_{i-1}) \right)^\top \right] \tag{109}
 \end{aligned}$$

where  $\|\cdot\|_{\text{F}}$  denotes the Frobenius norm and  $D$  denotes the informativeness matrix with entries corresponding to  $[D]_{kj} \triangleq d_k(\theta_j)$ . Here, (i) because of the i.i.d. assumption on data over time, and (ii) under sufficiently small  $\delta$  values that keep the probability of the event ' $\widehat{\theta}^\circ \neq \theta^\circ$ ' sufficiently small with increasing  $M$  (see Theorem 2 and also refer to (Bordignon et al., 2021)), the first term satisfies

$$\begin{aligned}
 \frac{1}{M^2} \mathbb{E}\left\| \sum_{i=1}^M (\mathbf{X}_i - D) \right\|_{\text{F}}^2 &= \frac{1}{M} \mathbb{E}\|\mathbf{X}_i - D\|_{\text{F}}^2 + \frac{2}{M^2} \sum_{i=1}^M \sum_{j < i} \mathbb{E}\text{Tr}\left[ (\mathbf{X}_i - D)(\mathbf{X}_j - D)^\top \right] \\
 &= \frac{1}{M} \text{Tr}(\mathcal{R}), \tag{110}
 \end{aligned}$$

where we also have defined the covariance matrix of the log-likelihood ratios  $\mathcal{R}$ :

$$\mathcal{R} \triangleq \mathbb{E}\left( (\mathbf{X}_i - D)(\mathbf{X}_i - D)^\top \right). \tag{111}$$

Moreover, expanding recursion (78), it holds that

$$\sum_{i=1}^M \mathbf{\Lambda}_i = \sum_{i=1}^M (1 - \delta)^i (A^i)^\top \mathbf{\Lambda}_0 + \beta \sum_{i=1}^M \sum_{j=0}^{i-1} (1 - \delta)^j (A^j)^\top \mathbf{X}_{i-j}. \tag{112}$$

Furthermore, by the triangle inequality, it follows that

$$\begin{aligned}
 \left\| \sum_{i=1}^M \mathbf{\Lambda}_i \right\|_{\text{F}} &\leq \left\| \sum_{i=1}^M (1 - \delta)^i (A^i)^\top \mathbf{\Lambda}_0 \right\|_{\text{F}} + \beta \left\| \sum_{i=1}^M \sum_{j=0}^{i-1} (1 - \delta)^j (A^j)^\top \mathbf{X}_{i-j} \right\|_{\text{F}} \\
 &= O\left( \min\left\{ M^2, \frac{M}{\delta} \right\} \right). \tag{113}
 \end{aligned}$$

Hence, the second term in (109) satisfies

$$\frac{(1 - \delta)^2}{\beta^2 M^2} \mathbb{E}\left\| \sum_{i=1}^M E^\top \mathbf{\Lambda}_{i-1} \right\|_{\text{F}}^2 = O\left( \|E\|^2 \min\left\{ M^2, \frac{1}{\delta^2} \right\} \right) \tag{114}$$

and the third term in (109) satisfies

$$\frac{(1 - \delta)}{\beta M^2} \mathbb{E}\left[ \text{Tr}\left( \sum_{i=1}^M (\mathbf{X}_i - D) \right) \left( \sum_{i=1}^M (E^\top \mathbf{\Lambda}_{i-1}) \right)^\top \right] = O\left( \frac{\|E\|}{M} \min\left\{ M, \frac{1}{\delta} \right\} \right) \tag{115}$$

Consequently, for sufficiently small combination matrix errors, namely,

$$E = o\left(\max\left\{M^{-3/2}, \delta M^{-1/2}\right\}\right), \quad (116)$$

it holds that

$$\mathbb{E}\|\widehat{D} - D\|_{\mathbb{F}}^2 \leq \frac{1}{M}\text{Tr}(\mathcal{R}) + o(1/M). \quad (117)$$

Next, recall (61) which establishes log-belief ratios under interventions for ASL. For the first term in (61), the matrix estimation error  $E$  implies errors for the matrix components  $R$  and  $r$  that are also proportional to  $E$ , which means

$$\begin{aligned} & (I - (1 - \delta)(R + O(E))^{\top})^{-1}(r + O(E)) \\ &= (I - (1 - \delta)R^{\top})^{-1}\left(I - (I - (1 - \delta)R^{\top})^{-1}(1 - \delta)O(E)\right)^{-1}(r + O(E)) \\ &\stackrel{(a)}{=} (I - (1 - \delta)R^{\top})^{-1}\left(I + (1 - \delta)O(E)\right)(r + O(E)) \\ &= \left((I - (1 - \delta)R^{\top})^{-1} + (1 - \delta)O(E)\right)(r + O(E)) \\ &= (I - (1 - \delta)R^{\top})^{-1}r + O(E) \end{aligned} \quad (118)$$

where (a) holds as long as

$$\begin{aligned} E &= o\left((I - (1 - \delta)R^{\top})^{-1}\right) \\ &= o\left(1/(1 - (1 - \delta)\rho(R))\right) \\ &= o\left(1/(1 - \rho(R) + \delta\rho(R))\right) \\ &= o\left(1/(1 - \rho(R))\right) \end{aligned} \quad (119)$$

In (119) we assume that the graph topology is fixed (i.e., it is not changing with decreasing  $\delta$ ). By the same arguments, an error in the matrix expression of the second term of (61) implies

$$\begin{aligned} \left((I - (1 - \delta)(R + O(E))^{\top})^{-1} - I\right) &= \left((I - (1 - \delta)R^{\top})^{-1} + (1 - \delta)O(E) - I\right) \\ &= \left((I - (1 - \delta)R^{\top})^{-1} - I\right) + O(E), \end{aligned} \quad (120)$$

and an error in the matrix expression of the pre-intervention case of (13) implies

$$\begin{aligned} \left((I - (1 - \delta)\widehat{A}^{\top})^{-1} - I\right) &= \left((I - (1 - \delta)(A + E)^{\top})^{-1} - I\right) \\ &= \left((I - (1 - \delta)A^{\top})^{-1} - I\right) + O(E). \end{aligned} \quad (121)$$

Therefore, considering the fact that true causal effect can be written as

$$\begin{aligned}
 C_{m \rightarrow k}^{\text{ASL}} &= \mu_{k,\infty}(\theta^\circ) - \tilde{\mu}_{k,\infty}(\theta^\circ) \\
 &\stackrel{(16)}{=} \frac{1}{1 + \sum_{\theta \in \Theta \setminus \{\theta^\circ\}} \exp\{-\lambda_{k,\infty}(\theta)\}} - \frac{1}{1 + \sum_{\theta \in \Theta \setminus \{\theta^\circ\}} \exp\{-\tilde{\lambda}_{k,\infty}(\theta)\}} \\
 &\stackrel{(13),(61)}{=} \frac{1}{1 + \sum_{\theta \in \Theta \setminus \{\theta^\circ\}} \exp\left\{-\left[\frac{\beta}{1-\delta} \left((I - (1-\delta)A^\top)^{-1} - I\right) d(\theta)\right]_k\right\}} \\
 &\quad - \frac{1}{1 + \sum_{\theta \in \Theta \setminus \{\theta^\circ\}} \left(\frac{\mu_m(\theta)}{\mu_m(\theta^\circ)}\right)^{[(I - (1-\delta)R^\top)^{-1}r]_k} \exp\left\{-\left[\frac{\beta}{1-\delta} \left((I - (1-\delta)R^\top)^{-1} - I\right) d_{-m}(\theta)\right]_k\right\}},
 \end{aligned} \tag{122}$$

by (118), (120), and (121), the causal effect estimate satisfies

$$\begin{aligned}
 \hat{C}_{m \rightarrow k}^{\text{ASL}} &= \frac{1}{1 + \sum_{\theta \in \Theta \setminus \{\theta^\circ\}} \exp\left\{-\left[\frac{\beta}{1-\delta} \left((I - (1-\delta)A^\top)^{-1} - I + O(E)\right) \hat{d}(\theta)\right]_k\right\}} \\
 &\quad - \frac{1}{1 + \sum_{\theta \in \Theta \setminus \{\theta^\circ\}} \left(\frac{\mu_m(\theta)}{\mu_m(\theta^\circ)}\right)^{[(I - (1-\delta)R^\top)^{-1}r]_k + O(E)} \exp\left\{-\left[\frac{\beta}{1-\delta} \left((I - (1-\delta)R^\top)^{-1} - I + O(E)\right) \hat{d}_{-m}(\theta)\right]_k\right\}}.
 \end{aligned} \tag{123}$$

Here,  $\hat{d}(\theta)$  and  $\hat{d}_{-m}(\theta)$  are the estimates for  $d(\theta)$  and  $d_{-m}(\theta)$  respectively, which can be formed from informativeness estimation matrix  $\hat{\mathbf{D}}$ . In addition, the scalar function  $g$  with vector argument  $\lambda$  (over  $\Theta \setminus \{\theta^\circ\}$ ) defined as

$$g(\lambda) \triangleq \frac{1}{1 + \sum_{\theta \in \Theta \setminus \{\theta^\circ\}} \exp\{-\lambda(\theta)\}} \tag{124}$$

has the partial derivatives

$$\frac{\partial g(\lambda)}{\partial \lambda(\theta)} = \frac{\exp\{-\lambda(\theta)\}}{\left(1 + \sum_{\theta \in \Theta \setminus \{\theta^\circ\}} \exp\{-\lambda(\theta)\}\right)^2}, \tag{125}$$

which implies that

$$\begin{aligned}
 \|\nabla g(\lambda)\|_1 &= \sum_{\theta \in \Theta \setminus \{\theta^\circ\}} \left| \frac{\partial g(\lambda)}{\partial \lambda(\theta)} \right| \\
 &= \sum_{\theta \in \Theta \setminus \{\theta^\circ\}} \frac{\exp\{-\lambda(\theta)\}}{\left(1 + \sum_{\theta \in \Theta \setminus \{\theta^\circ\}} \exp\{-\lambda(\theta)\}\right)^2} \leq 1
 \end{aligned} \tag{126}$$

By the mean-value theorem, this further implies that for any two vectors  $\lambda_1$  and  $\lambda_2$ , there exists a vector  $\lambda_3$  such that

$$g(\lambda_1) - g(\lambda_2) = \nabla g(\lambda_3)^\top (\lambda_1 - \lambda_2). \quad (127)$$

Then, using Hölder's inequality on (127) yields that

$$\begin{aligned} |g(\lambda_1) - g(\lambda_2)| &\leq \|\nabla g(\lambda_3)\|_1 \|\lambda_1 - \lambda_2\|_\infty \\ &\stackrel{(126)}{\leq} \|\lambda_1 - \lambda_2\|_\infty. \end{aligned} \quad (128)$$

Note that the norm  $\|\cdot\|_\infty$  chooses the absolute maximum over all arguments induced from different hypotheses  $\theta \in \Theta \setminus \{\theta^\circ\}$ . If we apply (128) to the differences of first and second terms in (122) and (123), we get

$$\begin{aligned} &|C_{m \rightarrow k}^{\text{ASL}} - \widehat{C}_{m \rightarrow k}^{\text{ASL}}| \\ &\leq \frac{\beta}{1-\delta} \left\| \left[ \left( (I - (1-\delta)A^\top)^{-1} - I \right) d(\theta) - \left( (I - (1-\delta)A^\top)^{-1} - I + O(E) \right) \widehat{\mathbf{d}}(\theta) \right]_k \right\|_\infty \\ &+ \left\| O(E) \log \frac{\mu_m(\theta)}{\mu_m(\theta^\circ)} + \frac{\beta}{1-\delta} \left[ \left( (I - (1-\delta)R^\top)^{-1} - I \right) d_{-m}(\theta) - \left( (I - (1-\delta)R^\top)^{-1} - I + O(E) \right) \widehat{\mathbf{d}}_{-m}(\theta) \right]_k \right\|_\infty \\ &\leq \frac{\beta}{1-\delta} \left\| \left[ \left( (I - (1-\delta)A^\top)^{-1} - I \right) (d(\theta) - \widehat{\mathbf{d}}(\theta)) + O(E) \widehat{\mathbf{d}}(\theta) \right]_k \right\|_\infty \\ &+ \left\| O(E) + \frac{\beta}{1-\delta} \left[ \left( (I - (1-\delta)R^\top)^{-1} - I \right) (d_{-m}(\theta) - \widehat{\mathbf{d}}_{-m}(\theta)) - O(E) \widehat{\mathbf{d}}_{-m}(\theta) \right]_k \right\|_\infty \\ &\stackrel{(a)}{\leq} \frac{\beta}{1-\delta} \left\| \left[ \left( (I - (1-\delta)A^\top)^{-1} - I \right) (d(\theta) - \widehat{\mathbf{d}}(\theta)) \right]_k \right\|_\infty \\ &\quad + \frac{\beta}{1-\delta} \left\| \left( (I - (1-\delta)R^\top)^{-1} - I \right) (d_{-m}(\theta) - \widehat{\mathbf{d}}_{-m}(\theta)) \right\|_\infty + O(E) \\ &\stackrel{(b)}{\leq} \frac{\beta}{1-\delta} \left\| (I - (1-\delta)A^\top)^{-1} - I \right\|_\infty \left\| d(\theta) - \widehat{\mathbf{d}}(\theta) \right\|_\infty \\ &\quad + \frac{\beta}{1-\delta} \left\| (I - (1-\delta)R^\top)^{-1} - I \right\|_\infty \left\| d_{-m}(\theta) - \widehat{\mathbf{d}}_{-m}(\theta) \right\|_\infty + O(E) \end{aligned} \quad (129)$$

where (a) follows from the triangle inequality, and (b) follows from the sub-multiplicity of the matrix norms. Finally, since (117) implies

$$\mathbb{E} \left\| d(\theta) - \widehat{\mathbf{d}}(\theta) \right\|_\infty = O\left(\frac{1}{\sqrt{M}}\right) \quad (130)$$

and

$$\mathbb{E} \left\| d_{-m}(\theta) - \widehat{\mathbf{d}}_{-m}(\theta) \right\|_\infty = O\left(\frac{1}{\sqrt{M}}\right), \quad (131)$$

taking the expectation of both sides in (129) and incorporating (130) and (131) concludes that

$$\mathbb{E} \left| C_{m \rightarrow k}^{\text{ASL}} - \widehat{C}_{m \rightarrow k}^{\text{ASL}} \right| = O(1/\sqrt{M}). \quad (132)$$

In the case of NBSL, the only source of disagreement between  $C_{m \rightarrow k}^{\text{NB}}$  and  $\widehat{C}_{m \rightarrow k}^{\text{NB}}$  is at post-intervention, since pre-intervention steady state beliefs are fixed and identical (see (21)), thus their difference equals 0. Also, the proof for the post-intervention difference proceeds along the same lines of the ASL case above.

## Appendix F. Discussion of Computational Complexity

First, assuming that the matrix  $A$  and the informativeness vector  $d(\theta)$  for each hypothesis are known, the computational tasks for calculating the derived causal effect expressions (35) and (61) involve the following:

- Calculating  $(I - (1 - \delta)R^T)^{-1}$  for one agent can be achieved in  $O(K^3)$  time with naive matrix operations, where  $K$  is the size of the network. Performing this operation for all agents over the network *naively* results in  $O(K^4)$  worst case complexity. However, as we show in the sequel, doing such operations independently for each agent is highly inefficient, and in fact, all calculations can be done with  $O(K^3)$  instead of  $O(K^4)$ .
- For the matrix-vector multiplications, the complexity for each agent is  $O(K^2)$ , which leads to a complexity of  $O(K^3)$  when we consider all  $K$  agents.
- The computation to find the eigenvectors of the  $K \times K$  causal influence matrix also involves a computation time of  $O(K^3)$ .

Therefore, the total computational complexity with known  $A$  and  $d(\theta)$  is  $O(K^3)$ . If, however,  $A$  and  $d(\theta)$  need to be estimated first as in Alg. 2, we need to perform the following additional operations:

- Estimating  $A$ , based on the averaging rule, is  $O(K)$  or  $O(K^2)$  depending on the utilized graph data structure.
- Estimating  $d(\theta)$  involves learning from the average over  $M$  observations. Here, each computation involves a matrix multiplication between  $A$  and  $\Lambda$  (i.e., the observational data) which has  $O(K^2)$  complexity.

Consequently, considering causal discovery from observational data phase, the total computational complexity for causal ranking of the network is  $O(K^3 + M \cdot K^2)$ .

Now, we return to the calculation of  $(I - (1 - \delta)R^T)^{-1}$  for each agent and discuss how  $O(K^3)$  can be obtained by exploiting the fact that the  $K$  inverse computations are not independent. For each agent  $k = 1, \dots, K$ , let  $R_k$  denote the  $(K - 1) \times (K - 1)$  submatrix obtained by deleting the  $k$ -th row and the  $k$ -th column of  $A$ . We treat the ASL case  $\delta > 0$  and the NBSL case  $\delta = 0$  separately.

**1) ASL:** Since  $\delta > 0$ ,  $I - (1 - \delta)A^\top$  is invertible. Let  $G_\delta \triangleq (I - (1 - \delta)A^\top)^{-1}$ . The Schur complement identity hence implies that

$$(I - (1 - \delta)R_k^\top)^{-1} = (G_\delta)_{-k,-k} - \frac{(G_\delta)_{-k,k}(G_\delta)_{k,-k}}{[G_\delta]_{kk}}, \quad (133)$$

where  $(G_\delta)_{-k,-k}$  denotes the submatrix of  $G_\delta$  obtained by deleting the  $k$ -th row and column, and  $(G_\delta)_{-k,k}$  and  $(G_\delta)_{k,-k}$  denote the corresponding column and row blocks. Hence, it is sufficient to invert the full matrix  $I - (1 - \delta)A^\top$  once, which is  $O(K^3)$ . Recovering each individual inverse from the above formula costs  $O(K^2)$  (the complexity of outer product) and therefore recovering all  $K$  inverses costs  $O(K^3) + K \cdot O(K^2) = O(K^3)$ .

**2) NBSL:** In this case,  $\delta = 0$  and  $I - A^\top$  is singular. Therefore, Schur complement argument for ASL cannot be applied directly. However, all required inverses can still be recovered in total  $O(K^3)$  time as we show next.

For each agent  $k$ ,  $e_k$  denoting the  $k$ -th canonical basis vector,

$$B_k \triangleq I - (I - e_k e_k^\top)A^\top$$

replaces the  $k$ -th row of  $I - A^\top$  by  $e_k^\top$ . Furthermore, let  $P_k$  be a permutation matrix that moves the  $k$ -th coordinate to the first position. Then, it holds that

$$P_k^\top B_k P_k = \begin{bmatrix} 1 & 0 \\ -r_k^\top & I - R_k^\top \end{bmatrix},$$

where  $r_k$  is the  $k$ -th row of  $A$  with its  $k$ -th entry removed, written as a  $(K - 1) \times 1$  column vector. Therefore,

$$(P_k^\top B_k P_k)^{-1} = \begin{bmatrix} 1 & 0 \\ (I - R_k^\top)^{-1} r_k^\top & (I - R_k^\top)^{-1} \end{bmatrix},$$

and thus  $(I - R_k^\top)^{-1}$  is exactly the lower-right  $(K - 1) \times (K - 1)$  block of  $(P_k^\top B_k P_k)^{-1}$ . Now, if we fix a reference agent, say  $k = 1$  as before, computing  $B_1^{-1}$  once is  $O(K^3)$ . For any other agent  $k$ , by definition, we have

$$B_k = B_1 + e_k (Ae_k)^\top - e_1 (Ae_1)^\top = B_1 + U_k V_k^\top,$$

where

$$U_k \triangleq [e_k \quad -e_1], \quad V_k \triangleq [Ae_k \quad Ae_1].$$

Thus,  $B_k$  differs from  $B_1$  by a rank-2 update. By the Woodbury identity,

$$B_k^{-1} = B_1^{-1} - B_1^{-1} U_k (I + V_k^\top B_1^{-1} U_k)^{-1} V_k^\top B_1^{-1}.$$

Since  $U_k$  and  $V_k$  each have only two columns, computing  $B_k^{-1}$  from  $B_1^{-1}$  costs  $O(K^2)$  per agent. Hence, after one initial  $O(K^3)$  inversion for the reference agent, all  $K$  matrices  $\{(I - R_k^\top)^{-1}\}_{k=1}^K$  can be recovered in total

$$O(K^3) + K \cdot O(K^2) = O(K^3)$$

time.




Figure 20: Average rank of the attacked agent versus the number of added bots for Erdős–Rényi networks with  $K = 20$  and edge probability 0.3, averaged over 50 independent realizations.

## Appendix G. Additional Experiments

In this section, we report additional experiments that complement the results in Section 8.

**Robustness to follower-inflation attacks** To test whether the robustness observed in the main text persists beyond the already considered setting, we repeat the follower-inflation experiment on Erdős–Rényi networks of different sizes. For each  $K \in \{20, 40, 60, 80\}$ , we generate 50 independent Erdős–Rényi networks with edge probability 0.3. In each network, one agent is selected uniformly at random as the attacked user and we progressively add bot accounts that act as artificial followers of this agent.

The results are shown in Figs. 20–23. Across all considered network sizes, the behavior is consistent with that reported in the main text. In contrast to the other methods, *CausalRank* remains stable as bots are added. These experiments therefore prove that the robustness of *CausalRank* is not specific to a single network size, but persists across substantially larger Erdős–Rényi graphs as well.

**CPU-time comparisons** We next report additional CPU-timing results for the computation of the causal influence matrix and the rankings. In Fig. 24, we show the total CPU time for ASL with  $\delta = 0.1$  as a function of network size  $K$ , which shows that the behavior with NBSL is present also with ASL.

In Figs. 25 and 26, we compare the naive  $O(K^4)$  implementation with the accelerated  $O(K^3)$  formulation derived in Appendix F. As we discussed in Appendix F, the naive implementation computes  $K$  independent matrix inversions, whereas the accelerated version recovers all  $K$  required inverses from a single  $K \times K$  inversion by exploiting the Woodbury identity in the NBSL case ( $\delta = 0$ ) and the Schur’s complement in the ASL case ( $\delta > 0$ ). The empirical results verify that the speedup relative to the naive approach can be substantial (reaching to approximately  $10\times$  speedup by  $K = 80$ ).




Figure 21: Average rank of the attacked agent versus the number of added bots for Erdős–Rényi networks with  $K = 40$  and edge probability 0.3, averaged over 50 independent realizations.

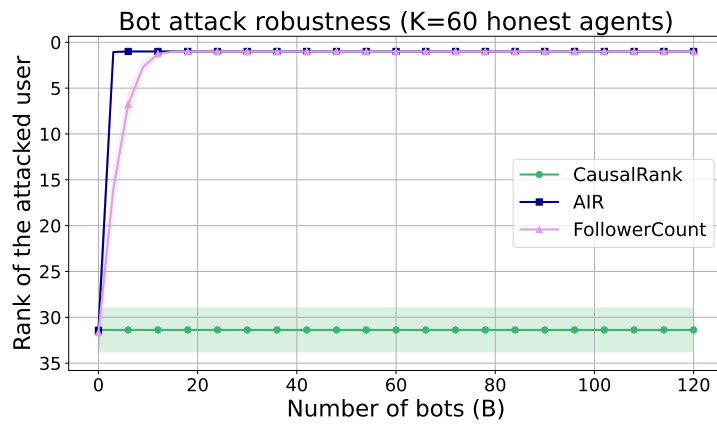


Figure 22: Average rank of the attacked agent versus the number of added bots for Erdős–Rényi networks with  $K = 60$  and edge probability 0.3, averaged over 50 independent realizations.




Figure 23: Average rank of the attacked agent versus the number of added bots for Erdős–Rényi networks with  $K = 80$  and edge probability 0.3, averaged over 50 independent realizations.

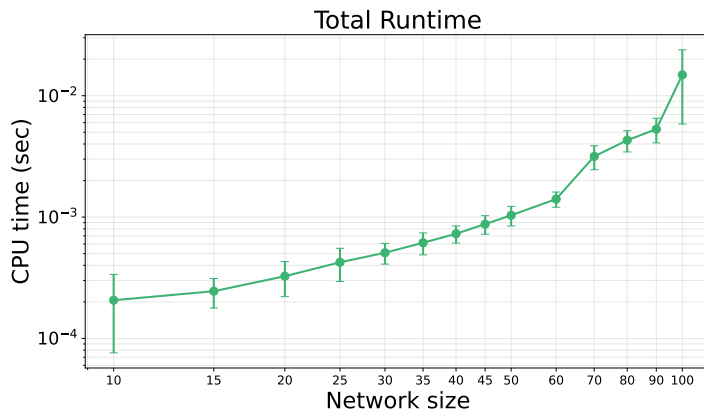


Figure 24: Total CPU time for ASL with  $\delta = 0.1$  as a function of network size  $K$ .

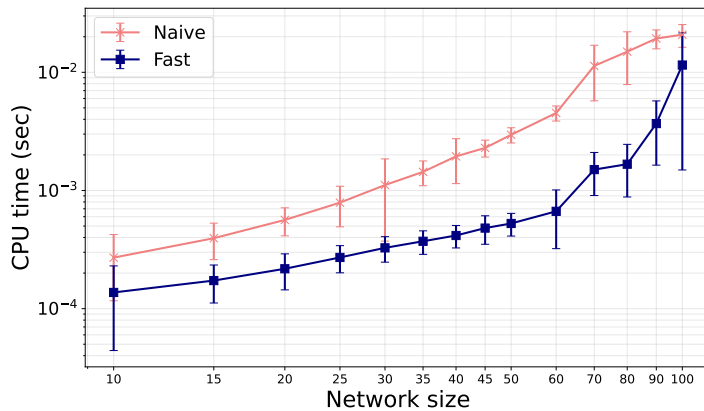


Figure 25: CPU-time comparison between the naive  $O(K^4)$  and accelerated  $O(K^3)$  implementations for NBSL.




Figure 26: CPU-time comparison between the naive  $O(K^4)$  and accelerated  $O(K^3)$  implementations for ASL.

## References

- D. Acemoglu, A. Ozdaglar, and A. ParandehGheibi. Spread of (mis) information in social networks. *Games and Economic Behavior*, 70(2):194–227, 2010.
- D. Acemoglu, M. A. Dahleh, I. Lobel, and A. Ozdaglar. Bayesian learning in social networks. *The Review of Economic Studies*, 78(4):1201–1236, 2011.
- A. Agarwal, S. Cen, D. Shah, and C. L. Yu. Network synthetic interventions: A framework for panel data with network interference. *arXiv:2210.11355*, 2022.
- L. Ante. How Elon Musk’s Twitter activity moves cryptocurrency markets. *Technological Forecasting and Social Change*, 186:1–14, 2023. doi: <https://doi.org/10.1016/j.techfore.2022.122112>.
- S. Aral and D. Walker. Identifying influential and susceptible members of social networks. *Science*, 337(6092):337–341, 2012.
- A. Banerjee, A. G. Chandrasekhar, E. Duflo, and M. O. Jackson. The diffusion of microfinance. *Science*, 341(6144):1236–1248, 2013.
- R. M. Bond, C. J. Fariss, J. J. Jones, A. D.I. Kramer, C. Marlow, J. E. Settle, and J. H. Fowler. A 61-million-person experiment in social influence and political mobilization. *Nature*, 489(7415):295–298, 2012.
- V. Bordinon, V. Matta, and A. H. Sayed. Adaptive social learning. *IEEE Transactions on Information Theory*, 67(9):6053–6081, 2021.
- S. Brin and L. Page. The anatomy of a large-scale hypertextual Web search engine. *Computer Networks and ISDN Systems*, 30(1):107–117, 1998.
- C. Chamley, A. Scaglione, and L. Li. Models for the diffusion of beliefs in social networks: An overview. *IEEE Signal Processing Magazine*, 30(3):16–29, 2013.

- A. G. Chandrasekhar, H. Larreguy, and J. P. Xandri. Testing models of social learning on networks: Evidence from two experiments. *Econometrica*, 88(1):1–32, 2020.
- J. Conlisk. Why bounded rationality? *Journal of Economic Literature*, 34(2):669–700, 1996.
- F. Dablander and M. Hinne. Node centrality measures are a poor substitute for causal inference. *Scientific Reports*, 9(1):6846, 2019.
- M. H. DeGroot. Reaching a consensus. *Journal of the American Statistical Association*, 69(345):118–121, 1974.
- P. M. Djurić and Y. Wang. Distributed Bayesian learning in multiagent systems: Improving our understanding of its capabilities and limitations. *IEEE Signal Processing Magazine*, 29(2):65–76, 2012.
- D. Easley and J. Kleinberg. *Networks, Crowds, and Markets: Reasoning About a Highly Connected World*. Cambridge University Press, 2010.
- M. Eichler. Causal inference in time series analysis. *Causality: Statistical Perspectives and Applications*, pages 327–354, 2012.
- M. Eichler and V. Didelez. On Granger causality and the effect of interventions in time series. *Lifetime data analysis*, 16:3–32, 2010.
- K. Friston. A theory of cortical responses. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 360(1456):815–836, 2005.
- J. Ghaderi and R. Srikant. Opinion dynamics in social networks with stubborn agents: Equilibrium and convergence rate. *Automatica*, 50(12):3209–3215, 2014.
- B. Golub and M. O. Jackson. Naive learning in social networks and the wisdom of crowds. *American Economic Journal: Microeconomics*, 2(1):112–149, 2010.
- C. W. J. Granger. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, pages 424–438, 1969.
- J. Z. Hare, C. A. Uribe, L. M. Kaplan, and A. Jadbabaie. On malicious agents in non-Bayesian social learning with uncertain models. In *Proc. Int. Conf. on Information Fusion*, pages 1–8, 2019.
- J. Hazla, A. Jadbabaie, E. Mossel, and M. A. Rahimian. Bayesian decision making in groups is hard. *Operations Research*, 69(2):632–654, 2021.
- D. Hunter and T. Zaman. Optimizing opinions with stubborn agents. *Operations Research*, 70(4):2119–2137, 2022.
- A. Jadbabaie, P. Molavi, A. Sandroni, and A. Tahbaz-Salehi. Non-Bayesian social learning. *Games and Economic Behavior*, 76(1):210–225, 2012.
- S. Jain and A. Sinha. Identification of influential users on Twitter: A novel weighted correlated influence measure for Covid-19. *Chaos, Solitons & Fractals*, 139:1–8, 2020.

- D. Kempe, J. Kleinberg, and É. Tardos. Maximizing the spread of influence through a social network. In *Proc. ACM SIGKDD*, pages 137–146, 2003.
- V. Krishnamurthy. Dynamics of social networks: Multi-agent information fusion, anticipatory decision making and polling. *arXiv:2212.13323*, Dec. 2022.
- P. Lagrée, O. Cappé, B. Cautis, and S. Maniu. Algorithms for online influencer marketing. *ACM Transactions on Knowledge Discovery from Data*, 13(1):1–30, 2018.
- A. Lalitha, T. Javidi, and A. D. Sarwate. Social learning and distributed hypothesis testing. *IEEE Transactions on Information Theory*, 64(9):6161–6179, 2018.
- Y. Lee, A. L. Buchanan, E. L. Ogburn, S. R. Friedman, M. E. Halloran, N. V. Katenka, J. Wu, and G. Nikolopoulos. Finding influential subjects in a network using a causal framework. *Biometrics*, pages 1–13, 2023.
- Y. Li, B. Qiu, Y. Chen, and H. V. Zhao. Analysis of information diffusion with irrational users: A graphical evolutionary game approach. In *Proc. IEEE ICASSP*, pages 2527–2531, 2019.
- D. Loureiro, F. Barbieri, L. Neves, L. E. Anke, and J. Camacho-Collados. TimeLMs: Diachronic language models from Twitter. In *Proc. Annual Meeting of the Association for Computational Linguistics*, pages 251–260, 2022.
- P. Molavi, A. Jadbabaie, K. R. Rad, and A. Tahbaz-Salehi. Reaching consensus with increasing information. *IEEE Journal of Selected Topics in Signal Processing*, 7(2):358–369, 2013.
- P. Molavi, A. Tahbaz-Salehi, and A. Jadbabaie. A theory of non-Bayesian social learning. *Econometrica*, 86(2):445–490, 2018.
- E. Mossel and O. Tamuz. Opinion exchange dynamics. *Probability Surveys*, 14:155–204, 2017.
- E. Mossel, A. Sly, and O. Tamuz. Strategic learning and the topology of social networks. *Econometrica*, 83(5):1755–1794, 2015.
- A. Nedić, A. Olshevsky, and C. A. Uribe. Fast convergence rates for distributed non-Bayesian learning. *IEEE Transactions on Automatic Control*, 62(11):5538–5553, 2017.
- M. Oaksford and N. Chater. *Bayesian Rationality: The Probabilistic Approach to Human Reasoning*. Oxford University Press, 2007.
- J. Pearl. *Causality*. Cambridge University Press, 2009.
- J. Pearl and D. Mackenzie. *The Book of Why: The New Science of Cause and Effect*. Basic Books, 2018.
- J. Peters, D. Janzing, and B. Schölkopf. *Elements of Causal Inference: Foundations and Learning Algorithms*. The MIT Press, 2017.

- S. U. Pillai, T. Suel, and S. Cha. The Perron-Frobenius theorem: some of its applications. *IEEE Signal Processing Magazine*, 22(2):62–75, 2005.
- D. Quercia, J. Ellis, L. Capra, and J. Crowcroft. In the mood for being influential on Twitter. In *Proc. IEEE International Conference on Social Computing*, pages 307–314, Boston, MA, 2011.
- S. I. Resnick. *Adventures in Stochastic Processes*. Springer, 2013.
- D. B. Rubin. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66(5):688, 1974.
- A. H. Sayed. Adaptation, learning, and optimization over networks. *Foundations and Trends in Machine Learning*, 7(4-5):311–801, July 2014.
- D. Shah and T. Zaman. Rumors in a network: Who’s the culprit? *IEEE Transactions on Information Theory*, 57(8):5163–5181, 2011.
- V. Shumovskaia, M. Kayaalp, M. Cemri, and A. H. Sayed. Discovering influencers in opinion formation over social graphs. *IEEE Open Journal of Signal Processing*, 4:188–207, 2023.
- S. T. Smith, E. K. Kao, K. D. Senne, G. Bernstein, and S. Philips. Bayesian discovery of threat networks. *IEEE Transactions on Signal Processing*, 62(20):5324–5338, 2014.
- S. T. Smith, E. K. Kao, D. C. Shah, O. Simek, and D. B. Rubin. Influence estimation on social media networks using causal inference. In *IEEE Statistical Signal Processing Workshop (SSP)*, pages 328–332, 2018.
- S. T. Smith, E. K. Kao, E. D. Mackin, D. C. Shah, O. Simek, and D. B. Rubin. Automatic detection of influential actors in disinformation networks. *Proc. National Academy of Sciences*, 118(4), 2021.
- S. Soni, S. L. Ramirez, and J. J. Eisenstein. Detecting social influence in event cascades by comparing discriminative rankers. In *Proc. Machine Learning Research*, volume 104, pages 78–99, 2019.
- D. Sridhar, C. D. Bacco, and D. Blei. Estimating social influence from observational data. In *Proc. Conference on Causal Learning and Reasoning*, volume 177, pages 712–733, 11–13 Apr 2022.
- D. L. Sussman and E. M. Airoldi. Elements of estimation theory for causal effects in the presence of network interference. *arXiv:1702.03578*, 2017.
- P. Toulis and E. Kao. Estimation of causal peer influence effects. In *Proc. ICML*, pages 1489–1497, 2013.
- D. Vial and V. Subramanian. Local non-Bayesian social learning with stubborn agents. *IEEE Transactions on Control of Network Systems*, 9(3):1178–1188, 2022.
- J. Woodward. *Making Things Happen: A Theory of Causal Explanation*. Oxford University Press, 2004.

- S. Wright. The method of path coefficients. *The Annals of Mathematical Statistics*, 5(3): 161–215, 1934.
- B. Ying and A. H. Sayed. Information exchange and learning dynamics over weakly connected adaptive networks. *IEEE Transactions on Information Theory*, 62(3):1396–1414, 2016.
- A. Zellner. Optimal information processing and Bayes’s theorem. *The American Statistician*, 42(4):278–280, 1988.