

Robust Sequential Learning in Random Order Networks

William Guo
University of Pennsylvania
Philadelphia, PA, USA
willguo6@seas.upenn.edu

Edward Xiong
MIT
Cambridge, MA, USA
eyxiong@mit.edu

Jie Gao
Rutgers University
New Brunswick, NJ, USA
jg1555@rutgers.edu

ABSTRACT

In the sequential learning problem, agents in a network attempt to predict a binary ground truth, informed by both a noisy private signal and the predictions of neighboring agents before them. It is well known that social learning in this setting can be highly fragile: small changes to the action ordering, network topology, or even the strength of the agents’ private signals can prevent a network from converging to the truth. We study networks that achieve *random-order asymptotic truth learning*, in which almost all agents learn the ground truth when the decision ordering is selected uniformly at random. We analyze the robustness of these networks, showing that those achieving random-order asymptotic truth learning are resilient to a bounded number of adversarial modifications. We characterize necessary conditions for such networks to succeed in this setting and introduce several graph constructions that learn through different mechanisms. Finally, we present a randomized polynomial-time algorithm that transforms an arbitrary network into one achieving random-order learning using minimal edge or vertex modifications, with provable approximation guarantees. Our results reveal structural properties of networks that achieve random-order learning and provide algorithmic tools for designing robust social networks.

KEYWORDS

Social networks; Social Learning; Sequential learning; Bayesian learning; Robust networks; Approximation algorithms; Submodular Optimization

ACM Reference Format:

William Guo, Edward Xiong, and Jie Gao. 2026. Robust Sequential Learning in Random Order Networks. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 9 pages. <https://doi.org/10.65109/KNSC4490>

1 INTRODUCTION

The ability of a population to correctly aggregate dispersed information is fundamental to effective societal decision-making. Whether predicting a market crash, identifying a public health threat, or creating new social behaviors, success hinges on how weak signals can compound into stronger signals and how these signals propagate through the social network. It is in the interest of both individual agents and network designers to understand the circumstances that

enable society to learn as a collective, particularly when many network parameters — including the network topology and decision ordering — can be subject to change.

Social learning, or decision making in a social setting, is a well-studied topic in economics and social networks [7, 12, 20]. In the classical model of sequential social learning [6, 7, 9, 28, 31], agents take turns to make decisions, and each agent sees a private signal with uncertainty as well as the previous agent’s decisions. One major concern in this setting is *herding* or *information cascades*: with a constant probability, a few early agents make the wrong decision and subsequent agents ignore their own private signals, conforming to the actions of the agents they observe. This results in the majority of the network converging to the wrong value, even with perfectly rational Bayesian agents.

In this work, we consider sequential learning on a social network, following the model adopted in a few recent papers [4, 5, 19, 29]. A set of agents make predictions sequentially about an underlying binary ground truth. Each agent uses Bayesian inference to compute the truth value with maximum likelihood, using their independent bounded private signal as well as their neighbor’s predictions to construct their own belief. Recent work on this model demonstrates that the quality of information spread throughout a network is dependent on both its graph structure and the order in which individuals act. The ideal learning outcome is characterized by the notion of *asymptotic learning*, in which all but $o(n)$ agents in a network of size n correctly learn the ground truth with probability approaching 1 as n goes to infinity. However, the precise relationship between graph topology/decision ordering and whether a network is capable of learning asymptotically still remains elusive.

Of particular interest are networks that achieve asymptotic learning when the decision ordering is chosen *uniformly at random*. In practice, imposing a carefully crafted ordering on a network can be challenging when there is no global coordination. On the other hand, random orderings result in the “average case” or expected behavior, allowing us to isolate the role of topology in social learning.

Due to the inherent difficulty of preserving key graph structures under random orderings, networks that learn under random orderings are more difficult to design. In the literature, only two instances of network learning under random orders have been identified: the “celebrity network” in [5], and symmetric expander graphs satisfying the “local learning requirement” in [4]. Both networks seem to possess graph properties robust to adversarial deletion: [4] explicitly shows the existence of a network obtaining random-order learning, even after the removal of all but an arbitrary constant fraction of the network.

In this paper, we focus on characterizing and enabling asymptotic learning with random ordering. We characterize the robustness properties of random-order learning networks to adversarial modifications. Furthermore, we offer constructions for networks



This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). <https://doi.org/10.65109/KNSC4490>

achieving random-order learning through different means, and provide an algorithmic approach to enable arbitrary networks to learn using a small number of modifications.

1.1 Our Contributions

To motivate our exposition into random-order learning networks, we first show that strategic-order learning can be quite fragile. We exhibit a graph family that achieves asymptotic learning under a strategic ordering only when the uncertainty of the agent’s private signals is in a specified range. This demonstrates the lack of stability in networks that rely on such strategic orderings, as discussed in [19] and [29]. In contrast, we show that asymptotic learning under random ordering is robust to a constant number of adversarial modifications, generalizing the robustness properties discussed in [4] to arbitrary random-order learning networks. Namely, a random-order learning network with learning rate $1 - \epsilon$ is robust to $o(1/\epsilon)$ many modifications. This bound is nearly tight in the worst-case, as removing $\Theta(1/\epsilon)$ many vertices from a learning celebrity network can result in a non-learning network.

Motivated by these robustness properties, we provide constructions for networks that support random-order learning. We start with the classic example of a complete network that suffers from an information cascade [6]. We first show that any learning network must have an independent set of size $\omega(1)$, a property that a complete graph does not have. Surprisingly, by introducing a superconstant number of “guinea pigs” (degree one vertices), we find that a complete network can be “boosted” to achieve random-order asymptotic learning. This boosting approach is not limited to just guinea pigs: any graph structure capable of achieving asymptotic learning under a strategic order can be embedded to enable the complete graph to enable random-order learning.

As these constructions are specific to the complete graph, one may wonder how to boost an arbitrary network family to achieve random-order learning using a minimum number of edge/vertex modifications. We devise a randomized algorithm that achieves this by greedily boosting vertices with high coverage of other non-learning vertices. Our approach is closely related to the influence maximization framework introduced in [18], and more broadly submodular maximization [25, 32], which we use to show that our algorithm is a $O(g(n) \log n)$ -approximation of the minimum number of edge modifications needed to achieve random-order learning, for any superconstant function $g(n) = \omega(1)$.

Our algorithm, however, assumes the existence of an efficient “learning oracle” which determines whether a vertex achieves a high learning rate. The construction of such an oracle appears to be related to many other learning problems that are computationally hard. We conjecture the hardness of computing such an oracle for general graphs. Meanwhile, we give a deterministic local linear-time heuristic to check whether a vertex achieves asymptotic learning.

We leave the proofs of many of our results to the Appendix, which can be found in the full version [13].

1.2 Related Work

There is a large variation in model and parameter choices within social learning. We focus on the settings closest to ours, and refer the readers to some surveys in the field [2, 12, 20, 23].

Enabling Social Learning: To avoid herding in sequential learning, previous work suggests limiting social interactions, for example, removing the visibility between agents early in the sequence [26, 27], or introducing a stochastic social network among the agents [1]. A recent work also studied the timing effects of strategic behaviors of agents to delay actions [15]. When the private signals are unbounded (i.e., with probability arbitrarily close to 1), asymptotic learning occurs almost certainly [14, 28].

Sequential vs. Repeated learning: We study sequential learning where each agent makes one prediction following a linear order. In the repeated learning setting, agents broadcast a binary decision at every new time step t [21, 22]. For Bayesian agents with repeated learning, asymptotic truth learning is achieved if the graph is undirected and the distribution of private beliefs is not atomic¹ [22], or if agents repeatedly share with others their current beliefs [23].

Hardness of Bayesian Learning: It is well known that general Bayesian belief inference is difficult [10]. In the repeated social learning setting, it has been shown that computing the optimal decision is NP-hard even at $t = 2$ [17]. Generally, inferring the most likely ground truth value is PSPACE-hard [16]. In the sequential setting, determining whether a particular network can obtain a high learning rate under some strategic ordering is also NP-hard [29].

Influence Maximization: Our algorithm for enabling random-order learning relies on identifying a set of highly influential agents, a problem dating back to [11]. Prior work has established approximation algorithms for maximizing the expected influence under various diffusion models [8, 18]. Notably, so long as the influence objective is submodular, one can immediately obtain strong approximation guarantees [25, 32].

Robustness of Social Learning: The existing literature predominantly analyzes the robustness of social learning in light of faulty or misguided decision-making rules. For example, Bohren [3] considers a setting where some agents blindly predict based on their private signal and shows that significantly over-/underestimating the proportion of uninformed agents leads to undesirable behavior. Mueller-Frank [24] concerns the repeated learning setting, noting that the standard weighted average (DeGroot) decision model is not robust to arbitrarily small adjustments made to a single agent at each iteration. In the sequential learning setting, Arieli et. al [4] observed that networks learning under random orderings are robust to random vertex deletions, showing that robustness against random vertex deletions is embedded in the definition of uniform random orderings.

2 SETTING

We define a *network* \mathcal{F} to be an infinite family of undirected graphs with unbounded order $\sup_{G \in \mathcal{F}} |V(G)| = \infty$. For ease of notation, throughout this paper we describe networks as sequences $\mathcal{F} = \{G_n\}_{n \in \mathbb{N}}$ of undirected graphs $G_n = (V_n, E_n)$ where $|V_n| = n$. Each graph G_n represents n agents with perfect knowledge of the entire graph topology. Each graph is equipped with a decision ordering (equivalently, permutation) $\sigma_n : [n] \rightarrow [n]$, where σ_n is a bijection mapping each agent to a unique index. We distinguish between the *strategic ordering* setting, in which one may carefully choose

¹Note that a binary Bernoulli distribution for private signal as in our setting is not non-atomic.

the agent’s ordering for any graph G_n , and the *random ordering* setting, in which σ_n is chosen uniformly at random from the $n!$ possibilities. In either case, the decision ordering is known by all agents in advance.

Each agent seeks to learn the ground truth signal $\theta \in \{0, 1\}$, where $\Pr(\theta = 0) = \Pr(\theta = 1) = \frac{1}{2}$. Each agent is given a random private signal $p_v \in \{0, 1\}$, representing an independent noisy measurement of the ground truth. The agent’s private signals are correlated with θ with a common probability $q > 1/2$.

As per the decision ordering σ_n , agents will go in sequence to make their predictions. For any agent v , let $N(v)$ denote the set of neighbors of v in a graph. In sequential learning on a graph G , v can see the actions of their neighbors arriving before them in the ordering, namely in the set $N_\sigma(v) = \{u \in V_n : (u, v) \in E_n, \sigma(u) < \sigma(v)\}$. Each agent v then announces a prediction $a_v \in \{0, 1\}$ visible only to its neighbors succeeding it in the ordering. The ordering σ induces an acyclic directed graph \vec{G} which orients each edge in G to point from the lower index to the higher index in σ . An agent v ’s *subnetwork* $B(v)$ is defined on the set of all $u \in \vec{G}$ for which a u - v directed path exists in \vec{G} . Such vertices are called the ancestors of v . Informally, it is the set of all agents that may ultimately influence v ’s belief and action.

Following the literature in social learning [6], we assume all agents are perfectly Bayesian, and will predict the most probable ground truth value given their knowledge of G , σ , q , their private noisy signal, and the predictions of their neighbors arriving before them. For simplicity, if an agent finds that it is equally likely for $\theta = 0$ or 1 , the agent will side with their own private signal, although our results generalize to arbitrary tie-breaking rules.

For each agent $v \in V$, its *learning rate* $\ell_\sigma(v) := \ell_\sigma(v, q)$ given a fixed ordering σ is the probability that v correctly predicts the ground truth: $\ell_\sigma(v) = \Pr(a_v = \theta)$. The learning rate of a graph $G = \{V, E\}$ paired with a fixed ordering σ is the expected fraction of agents predicting the ground truth correctly, i.e.

$$L_\sigma(G) = \frac{1}{|V|} \sum_{v \in V} \ell_\sigma(v).$$

Similarly, we define the *random-order learning rate* of an agent v to be $\ell(v) = \mathbb{E}_\sigma[\ell_\sigma(v)]$, averaging the learning rate over all possible orderings σ weighted uniformly. Denote by S_n the set of all possible permutations of $[n]$. For an event $A \subseteq S_n$ dependent only on the random ordering of the vertices, the random-order learning rate conditioned on A is denoted as $\ell(v | A) = \Pr(a_v = \theta | \sigma \in A)$. The random-order learning rate of a graph G is the expected fraction of agents predicting the ground truth given a decision ordering selected uniformly at random, or

$$L(G) = \mathbb{E}_\sigma[L_\sigma(G)] = \frac{1}{|V|} \sum_{v \in V} \ell(v).$$

The learning rate of a network $\mathcal{F} = \{G_n\}_{n \in \mathbb{N}}$ is the *asymptotic learning rate* of its graph family, or $\lim_{n \rightarrow \infty} L(G_n)$. We say that a network obtains *random-order asymptotic truth learning* if $L(G_n) \rightarrow 1$ as $n \rightarrow \infty$. When \mathcal{F} is equipped with a sequence of fixed orderings $\{\sigma_n\}$, we say the network achieves *strategic order learning* if $\lim_{n \rightarrow \infty} L_{\sigma_n}(G_n) = 1$. We will abbreviate random order (resp. strategic order) asymptotic truth learning as simply random-order (resp. strategic order) learning, or even just “learning”. An

agent v with learning rate approaching 1 as n goes to infinity *learns* or *achieves learning*.

We now describe a generalized improvement principle, which extends upon Proposition 1 in [19]:

PROPOSITION 2.1 (GENERALIZED IMPROVEMENT PRINCIPLE). *Given a graph $G = \{V, E\}$ and a particular vertex $v \in V$, consider a sub-graph $G' = \{V, E'\}$ of G such that $E' \subseteq E$, and for all $u, w \neq v$, $(u, w) \in E \implies (u, w) \in E'$. Then for any fixed ordering σ , the learning rate of v in graph G , denoted by $\ell_\sigma(v)$, is at least that of the learning rate of v in G' denoted by $\ell'_\sigma(v)$. Further, under random orderings $\ell(v) \geq \ell'(v)$.*

In essence, the generalized improvement principle ensures that each agent does at least as well as if it could only see a subset of its neighbors. This is true due to the rationality (Bayesian inference) of each agent, along with their knowledge of the graph G and ordering σ . As a corollary, one can see that each agent learns at least as well as its neighbors before it in the decision ordering; inductively, this implies that each agent does at least as well as all previous agents in its subnetwork.

Next, we establish a general property of random-order learning in a network, which holds independently of the decision rules.

LEMMA 2.2. *For a vertex v in G with a random learning rate at least $\ell(v) \geq 1 - \epsilon$, and an event $A \subseteq S_n$ which occurs with probability $\Pr(A) = \frac{|A|}{n!}$ on uniform random orderings, the learning rate of v conditioned on A is bounded by*

$$\ell(v | A) \geq 1 - \frac{\epsilon}{\Pr(A)}.$$

Finally, we note that any vertex with many learning neighbors itself must achieve learning.

LEMMA 2.3. *For any network $\mathcal{F} = \{G_n\}_{n \in \mathbb{N}}$ and vertex $v = \{v_n : v_n \in G_n, n \in \mathbb{N}\}$, if v has $d = \omega(1)$ neighbors that achieve learning, then v itself also achieves learning.*

One may expect that as long as the random ordering places v after one of its d learning neighbors, it too will learn. However, each neighbor has a lower learning rate conditioned on being early in the ordering, as, in expectation, they have less information to inform their beliefs (see Lemma 3.2). If v arrives too early, none of its neighbors before it may be in a position that enables learning. The decline in learning rate conditioned on being early in an ordering is a recurring obstacle throughout our analyses.

REMARK. *Many of our results invoke asymptotic notation despite referencing a single instance of a graph or vertex, which is done for brevity, as these results generalize to arbitrary values of n . When discussing guarantees related to asymptotic learning as in Lemma 2.3, we will clarify our results by referencing infinite families/sequences rather than fixed instances.*

3 ROBUSTNESS IN NETWORK LEARNING

3.1 Fragility of Strategic-Order Learning

To motivate our exposition into random-order learning networks, we first highlight the fragility of social learning under strategic orderings. Prior work (e.g., [19]) studied a myriad of networks that learn under strategic orderings. This type of learning is heavily




Figure 1: A network that obtains learning for $q < q_0$ but not for $q > q_0$, for some $q_0 \in [\frac{1}{2}, 1]$. The edges in the graph are oriented to indicate the direction of information flow.

dependent on the strategic ordering, yet even so, it can be fragile to other network parameters. Here, we introduce a network family that learns conditionally depending on q (the probability that private signals are correct), even after fixing the underlying graph topology and the strategic ordering on the vertices.

Our construction is shown in Figure 1, in which there is a special vertex v with four neighbors arriving before it, and a set of k neighbors $\{w_i\}$ arriving after it. Each w_i also has an independent neighbor that arrives before it. All the agents $\{w_i\}$ feed their predictions into a common neighbor u_0 arriving later. We let $k = \log n$, though any superconstant value of k suffices. All remaining vertices are arranged in a chain following the vertex u_0 .

In essence, the ability of the network to learn depends solely on whether u_0 is a learner. If u_0 could observe all (or even a constant fraction) of the independent private signals from the agents $\{w_i\}$, it would be able to correctly determine the ground truth with high probability by a simple Chernoff bound. However, all agents w_i observe the action of v . For $q > q_0$ for some $q_0 \approx 0.7887$, each w_i will default to the action of v , whereas for $q < q_0$, each w_i will act on its independent private signal with a constant probability. In the former case, all w_i 's will simply regurgitate the action of v , which u_0 will also adopt. But as v is not a learner, neither is u_0 . In the latter case, u_0 receives enough private signals to learn with high probability.

PROPOSITION 3.1. *There exists a strategic-order network \mathcal{F} that obtains asymptotic learning for $q < q_0$ and does not learn for $q > q_0$, for some fixed threshold q_0 .*

This example is rather counter-intuitive in that the network stops learning when the quality of private signals *improves*. However, this is not universal: a similar construction of a network that only learns for *high* values of q can be found in the full version of this paper. It can be seen that a small modification to the network, e.g., inserting or removing one of the early-arriving neighbors of v , can influence the entire network's learning behavior via a similar analysis.

A fundamental strength of random-order learning networks is the ability to learn almost independently of the decision ordering imposed on it. Such networks are able to learn seemingly due to inherent properties of the underlying topology. A natural question is whether such properties can be broken via minor perturbations to the network structure. In the next section, we investigate the robustness of random-order learning networks against *adversarial* edge/vertex modifications.

3.2 Robustness of Random-Order Learning

In this section, we argue that a network's robustness to edge/vertex insertion/deletions can be made explicit in terms of its random-order learning rate. We start with a rather intuitive observation: conditioned on being early in the decision ordering, an agent does worse than when conditioned on being later in the ordering. This can be attributed to the generalized improvement principle, as an agent arriving early and only seeing a few of its neighbors will be less informed than an agent arriving later and seeing more of its neighbors.

LEMMA 3.2 (MONOTONICITY). *Given a graph G , fix an agent $v \in G$. For all $i \in [n]$, let A_i be the event that v is at index i , or $\sigma(v) = i$. Then $\Pr(a_v = \theta \mid A_i) \leq \Pr(a_v = \theta \mid A_{i+1})$ for all $1 \leq i \leq n - 1$.*

Using monotonicity, we demonstrate that the random-order learning error of a graph only increases linearly in the number of modifications. The proof relies on the probability of any given agent arriving earlier in the random ordering before any of the vertex modifications.

THEOREM 3.3. *Given a graph G and a vertex $v \in G$ with random-order learning rate $1 - \epsilon$, after k vertex deletions, v obtains a learning rate at least $1 - (k + 1)\epsilon$.*

PROOF. Let D be the set of k deleted vertices, and let G' be the graph obtained by deleting D from G . For $n = |G|$, the size of G' is $n - k$. Consider the probability distribution \mathbb{P}_v of orderings on G' , obtained by first uniformly at random picking an ordering σ of G such that v precedes all $w \in D$, and then removing all the $w \in D$ to obtain an ordering σ' on G' . We use $\ell(v \mid \mathbb{P}_v)$ to denote the learning rate of $v \in G'$ under this distribution of orderings. Furthermore, let A_v denote the event that v precedes all vertices $w \in D$ in the original graph G .

The probability $\Pr(A_v)$ under a uniformly random ordering of vertices in G is $\frac{1}{k+1}$, so by Lemma 2.2 we have

$$\ell(v \mid A_v) \geq 1 - (k + 1)\epsilon.$$

For any ordering σ such that A_v occurs, we can trivially remove all vertices in D to obtain an ordering σ' of G' . The learning rate of v under σ and σ' is the same, since v is not affected by any vertex that arrives after it. Therefore,

$$\ell(v \mid \mathbb{P}_v) = \ell(v \mid A_v) \geq 1 - (k + 1)\epsilon.$$

Now observe that in the distribution \mathbb{P}_v , the probability that v is found at index i is

$$\Pr(\sigma(i) = v \mid \mathbb{P}_v) = \frac{\binom{n-i}{k}}{\binom{n}{k+1}}$$

and for the orderings such that $\sigma(i) = v$, all other vertices are uniformly random in the remaining indices. Importantly, note that $\Pr(\sigma(i) = v \mid \mathbb{P}_v)$ is monotonically decreasing in i . Thus by Lemma 3.2, the learning rate of v in G' is

$$\begin{aligned} \ell(v) &= \frac{1}{n-k} \sum_{i=1}^{n-k} \ell(v \mid \sigma(i) = v) \geq \sum_{i=1}^{n-k} \frac{\binom{n-i}{k}}{\binom{n}{k+1}} \ell(v \mid \sigma(i) = v) \\ &= \ell(v \mid \mathbb{P}_v) \geq 1 - (k+1)\epsilon \end{aligned}$$

as desired. \square

The same approach works to prove the following analogous statements on other modifications.

COROLLARY 3.3.1. *Given a graph G and a vertex $v \in G$ with random-order learning rate $1 - \epsilon$, after k vertex insertions, v obtains learning rate at least $1 - (k+1)\epsilon$.*

COROLLARY 3.3.2. *Given a graph G and a vertex $v \in G$ with random-order learning rate $1 - \epsilon$, after k edge deletions or insertions, v obtains learning rate at least $1 - \frac{2k+1}{2}\epsilon$.*

In this context, we note that a vertex insertion allows for the addition of an arbitrary set of edges from the new vertex to any existing vertices in the graph. Thus, we find that networks achieving random asymptotic learning are robust against a number of modifications dependent on the learning rate of the network:

THEOREM 3.4. *Given a network \mathcal{F} with random-order asymptotic learning rate $1 - \epsilon$, define the network \mathcal{F}' where each $G'_n \in \mathcal{F}'$ is obtained from some $G_m \in \mathcal{F}$ after at most k modifications, which may be vertex insertions/deletions or edge insertions/deletions. Then, \mathcal{F}' obtains random-order learning rate at least $1 - (k+1)\epsilon$. If $\epsilon = o(1)$ and $k = o(1/\epsilon)$, then \mathcal{F}' achieves asymptotic learning.*

We observe that for general random-order learning networks, this bound is nearly tight. As an example, consider the celebrity network [5], which consists of complete bipartite graphs $K_{n-k,k}$ with $k = o(n)$ and $k = \omega(1)$.² The larger partition of $n - k$ vertices can be thought of as “commoners”, while the smaller partition of k vertices represents “celebrities”. In a given random ordering, with high probability, the first celebrity in the ordering observes the independent signals of many commoners. Upon aggregating this information into a single high-quality signal, the celebrity can disseminate it to the rest of the network. When $k = \Omega(\log n)$, the learning rate of the celebrity network is at most $1 - \Theta(1/k)$, as at least $\Theta(n/k)$ commoners arrive before the first celebrity with probability $1 - o(1)$. Then, for $\epsilon = \Theta(1/k)$, one may adversarially delete the $k = \Omega(1/\epsilon)$ celebrities from the network, leaving $n - k$ lone vertices, which do not obtain asymptotic learning.

4 CONSTRUCTING ROBUST NETWORKS

4.1 Boosting a Complete Network

When given a network that does not support random-order learning, a natural engineering question is to ask how to achieve learning with minimal modifications, especially given the robustness guarantees proven in Section 3.2. Prior to this work, the understanding of

²When $k = \Omega(n)$ or $O(1)$, a poorly informed herd occurs with a constant probability. For example, when $k = \Omega(n)$, with a constant probability, the first celebrity sees only a constant number of commoners, so its probability of failure is at least a constant.




Figure 2: The celebrity graph as described in [5]

networks obtaining random-order learning has been quite limited, both in terms of necessary/sufficient properties and tangible examples of graph families. We first note a key necessary property of random-order learning networks, relying on a simple observation of Bayesian learning: the strength of a signal is dependent on the number of independent private signals backing it. It is known that each decision is a deterministic function of the announced private signals before it (see Lemma 4 in Appendix B of [19]). Then, for any agent v learning under random orderings, it is necessary that, with probability approaching 1, a superconstant number of agents preceding v act according to their own private signals. Conversely, if there is a constant probability that v 's subnetwork contains only $O(1)$ ancestors acting on their own private signals, with constant probability all of those ancestors receive incorrect private signals, implying that v has a learning rate bounded away from 1. We relate the number of agents predicting their own private signal to the size of the maximum independent set, and formalize a network-wide property in Proposition 4.1.

PROPOSITION 4.1. *Any network \mathcal{F} obtaining random-order learning must have maximum independent set size $\omega(1)$.*

Now, we introduce a new type of graph family obtaining asymptotic truth learning under random orderings that relies primarily on constant-degree vertices. We contrast this with the celebrity graph, in which *all* agents learn with high probability: there are no peripheral agents that do not learn yet still contribute significantly to the underlying learning scheme. In our construction, the constant-degree vertices do not learn themselves, but serve as guinea pigs for a select number of special vertices that propagate the signal to the rest of the network. Ignoring the constant-degree vertices, the rest of the network forms a complete graph, which is known not to achieve learning [6].

PROPOSITION 4.2. *There exists a network \mathcal{F} obtaining random-order learning that no longer achieves learning when all vertices of degree 1 are removed from each graph $G_n \in \mathcal{F}$.*

An instance of such a network is shown in Figure 3a. In essence, $g(n)$ special vertices in the complete graph K_n are each equipped with $h(n)$ guinea pigs with degree 1, where $g(n), h(n) = \omega(1)$ and $g(n) \cdot h(n) = o(n)$ so the number of added vertices is a negligible fraction of the entire network. With probability approaching 1, each special vertex sees $\Omega(h(n))$ of its guinea pigs, which all give their own private signals. By a Chernoff bound, we may deduce that each special vertex learns, and so by Lemma 2.3, the entire graph K_n achieves learning.

Despite the complete graph having poor learning capabilities on its own, Proposition 4.2 implies that it can achieve learning




Figure 3: (a): By equipping a small number of vertices in a K_n graph with many guinea pigs, the entire network can be boosted to achieve learning. (b): Example of K_n being boosted by \sqrt{n} complete binary trees of size $\log \log n$ to achieve random-order asymptotic learning.

upon “boosting” a small number of its vertices. One may note that equipping a vertex with a superconstant number of guinea pigs is not the only way to boost a vertex to achieve a high learning rate. Indeed, we show that *any* construction permitting a vertex to achieve high learning rates under fixed orderings can be embedded into a K_n network to boost the network at large. One example of this using complete binary trees is shown in Figure 3b.

Formally, suppose we are given a network $\mathcal{F}' = \{G'_n\}$ equipped with strategic orderings $\{\sigma_n\}$ such that for all $k \geq 1$, imposing the ordering $\sigma_k = \{v_1, v_2, \dots, v_k\}$ on G'_k permits v_k to learn, i.e. $\ell_{\sigma_k}(v_k) = 1 - o(1)$. Now, for sufficiently large n , let $k = \log \log n$. We begin with the complete graph K_n , and pick \sqrt{n} distinct vertices, denoted $S = \{v_k^{(i)} : i \in [\sqrt{n}]\}$. For each vertex $v_k^{(i)}$, attach a copy of $G'_k \setminus v_k$ so the induced graph on $\{G'_k \setminus v_k\} \cup v_k^{(i)} = G'_k$. For $m = n + \sqrt{n}(\log \log n - 1)$, we set G_m to be the final result, and allow $\mathcal{F} = \{G_m\}$ to be our final network.

PROPOSITION 4.3 (EMBEDDING ARBITRARY LEARNING STRUCTURES INTO RANDOM-ORDER LEARNING NETWORKS). *Given any network $\mathcal{F}' = \{G'_n\}$ with strategic orderings $\{\sigma_n\}$ for which there exists $v = \{v_n : v_n \in G'_n\}$ obtaining asymptotic learning, the construction described above yields a network \mathcal{F} that achieves random-order asymptotic learning.*

To summarize the proof idea, the probability over the random ordering that the agents in any G'_k copy are ordered identically to the sequence given by σ_k is

$$\frac{1}{k!} \geq \frac{1}{(\log \log n)^{\log \log n}} \geq \frac{1}{n^\delta}$$

for any $\delta > 0$. For any copy of the graph G'_k for which this occurs, the corresponding $v_k^{(i)}$ achieves learning by Proposition 2.1, i.e., its learning rate is at least that if it could only see its neighbors in the copy of G'_k . As the copies of G'_k are disjoint, we conclude that each $v_k^{(i)}$ has a probability $\geq 1/n^\delta$ of achieving learning. There are \sqrt{n} disjoint copies of $v_k^{(i)}$, so we expect many to achieve learning.

The key technical challenge arises from the fact that the earlier copies of $v_k^{(i)}$ have an exponentially lower probability of learning,

as necessarily all other vertices in its corresponding G'_k copy must appear before it. To circumvent this, we condition on the number of graphs whose entire vertex set appears before an index $j = n/\omega(1)$ for some small superconstant denominator. Since the internal ordering of each graph copy is independent, we proceed as above to show that at least one vertex $v_k^{(i)}$ before index j achieves learning. Thus, each vertex u after index j learns with probability approaching 1, and as each agent’s learning rate monotonically increases in its index, u achieves learning in aggregate.

We find this result rather surprising, as the inherent difficulty of preserving key graph properties under random orderings would suggest that complex structures such as binary trees cannot induce random-order learning. Notably, each boosted vertex v_i has a low probability of learning on its own, but it is almost guaranteed that in aggregate, at least one v_i learns.

4.2 Algorithmically Boosting Networks

We now discuss how to enable networks to achieve random-order learning while minimizing the number of modifications (limited to edge insertions/deletions and vertex insertions),³ which can be thought of as a generalization of the boosting approach depicted in Fig. 3a. Our algorithm runs in polynomial time, taking a fixed graph as input. To obtain the asymptotic guarantees for an entire network, one would need to execute our algorithm on each $G_n \in \mathcal{F}$.

A key assumption we make is that the algorithm has access to a learning oracle that can determine whether the learning rate of a vertex is above a certain threshold or not.

DEFINITION 4.1. *Let $f(v, t) : V \times (0, 1) \rightarrow \{0, 1\}$ be an oracle that outputs 1 if v learns at rate $\geq t$ under random orderings, and 0 otherwise.*

To the best of our knowledge, it is unclear whether such a function can be computed efficiently. We discuss a useful heuristic to simulate a learning oracle in Section 4.3.

When modifying a graph, the number of new vertices should be sublinear in the number of vertices in the original graph so as not to overrule the existing network. However, the number of vertices requiring boosting can be as large as $\Omega(n)$, as is the case in a sparse graph with constant average degree. To resolve this imbalance, we replicate the structure of the celebrity graph by connecting each non-learner with a small but superconstant number of learning vertices. This subroutine is described in BoostAgents.

Procedure 1 BoostAgents (G, S, k)

Construct $T = \{z_{11}, \dots, z_{kk}\} \cup \{w_1, \dots, w_k\}$
 Add edges (z_{ij}, w_i) for all $i, j \in [k]$
 For each $v \in S$ and $i \in [k]$, add edge (w_i, v)

As depicted in Figure 4, we introduce a set of k “celebrities” $\{w_i\}$, each equipped with k independent guinea pigs to guarantee that their learning rate is high. By Lemma 2.3, the added celebrities are capable of boosting the learning quality of the set S .

³Our analysis excludes vertex deletions, as events on the induced subgraph of the original vertices may change non-uniformly under deletions.




Figure 4: BoostAgents boosts the vertices $S \subseteq V$ by connecting them to an added set of celebrities $\{w_1, \dots, w_k\}$.

CLAIM 4.4. *Suppose $k = \omega(1)$. After BoostAgents on a set $S_n \subseteq V_n$, as $n \rightarrow \infty$, each $v \in S_n$ achieves random-order asymptotic truth learning.*

The number of vertices added by BoostAgents is $k^2 + k = \omega(1)$. As k can be any arbitrary slow-growing function of n , the number of new vertices can be made to be within the optimum by an additive value that is slightly superconstant. However, the number of edges added is $|S|k + k^2$. While the additive term k^2 is small, $|S|$ may be as large as $\Omega(n)$. Our goal is thus to S that is of *near-optimal* size, or within the optimum by a small multiplicative factor.

To lower-bound the number of modifications necessary for a non-learning network to achieve learning, we note that the set of modified vertices must be able to reach almost all non-learners with constant probability. It turns out that calling BoostAgents on a set of vertices maximizing coverage of non-learners is an appropriate method to enhance the learning rate of a network. We prove this in Lemma 4.5, where we argue that any non-learner reachable by many boosted vertices will also receive the effects of boosting.

LEMMA 4.5. *For a set $S \subseteq V$ on which BoostAgents has been called, suppose that each $u \in S$ has been equipped with $1/\epsilon = \omega(1)$ many guinea pigs. Then any vertex v that is reachable from S w.p. $\geq 1 - \delta$ has learning rate at least $1 - \delta - 2\sqrt[3]{81\epsilon}$. In particular, the probability v is reachable by some $u \in S$ for which $\sigma(u) \geq n\epsilon$ is at least $1 - \delta - \alpha\epsilon$, for any $\alpha = \omega(1)$, $\alpha = o(\frac{1}{\epsilon})$.*

The key challenge to proving Lemma 4.5 is, again, monotonicity. We would like to boost vertices with high coverage, but coverage increases the earlier a vertex is in the ordering. A reasonable worry is that our algorithm will choose to boost vertices that only have high coverage when placed earlier than any of the celebrities added by BoostAgents, thereby avoiding the effects of boosting. Fortunately, we show this is impossible in the second half of Lemma 4.5, which states that moving a vertex earlier by ϵn indices can only decrease the probability it reaches some other vertex by at most ϵ . Thus, having high coverage conditioned on being before all added celebrities implies high coverage more broadly.

Using Lemma 4.5, Lemma 4.6 follows immediately, suggesting that the minimum number of modifications needed through boosting alone, denoted $M_b(k, n)$, is within an arbitrarily slow-growing factor of the true optimum M . As networks achieving asymptotic

learning may leave $o(n)$ vertices uncovered, the exact optimum for our purposes may not be well defined. For precision, we fix a tolerance parameter $T := T(n) = o(n)$ and mandate that in expectation over the random ordering, at most T vertices are unaffected by the modifications. All our bounds hold for any threshold T .

LEMMA 4.6. *Given a network $\mathcal{F} = \{G_n\}_n$, consider any $M := M(n)$ such that it is possible for \mathcal{F} to obtain random-order learning by performing at most $M(n)$ vertex insertions or edge insertions/deletions to each G_n , where in expectation at most $T := T(n)$ non-learning vertices are unaffected by the modifications. Then there exists some $M_B := M_B(k, n)$ such that $M_B \leq O(k)M + 2k^2$ and it is possible to achieve random-order learning by adding at most $M_B(k, n)$ edges/vertices to each G_n solely through calls to BoostAgents(G, S, k), with at most T unaffected vertices in expectation.*

It remains to develop an algorithm to find a set S that maximizes its coverage. We leverage results from optimization problems with submodular functions [18, 25, 32]. Notably, fixing the size of the set $|S| = m$, the maximum coverage by the best m -element set can be approximated within a constant factor by simply greedily picking the agent with the largest marginal coverage increase at each iteration. Equivalently, fixing some target number of non-learners to be covered, the number of vertices to be boosted, chosen by the greedy approach is at most a $O(\log n)$ factor greater than the optimum. Our full algorithm is presented in Algorithm 1.

We first show that coverage as a function of a set S is submodular. Let $V' \subseteq V$ be the set of non-learners in G , and $C_\sigma(S)$ denote the number of agents in V' reachable from S under ordering σ . Then, $C = \frac{1}{n!} \sum_\sigma C_\sigma(S)$ is the expected number of agents reachable from S under a uniformly random ordering.

DEFINITION 4.2. [18] *Given some domain space U , a submodular function $f : 2^U \rightarrow \mathbb{R}$ satisfies*

$$f(S \cup \{x\}) - f(S) \geq f(T \cup \{x\}) - f(T)$$

for all $S \subseteq T \subseteq U$ and $x \in U$.

CLAIM 4.7. *For any $V' \subseteq V$, $C : 2^{V'} \rightarrow \mathbb{R}$ defined above is nonnegative, monotone, & submodular.*

While computing $C(S)$ exactly is not computationally feasible, we may estimate it using Monte Carlo sampling to obtain $\tilde{C}(S) = \frac{1}{N} \sum_{i=1}^N C_{\sigma_i}(S)$, for orderings $\sigma_1, \dots, \sigma_N$ sampled independently and uniformly at random. For a sufficiently large N , $\tilde{C}(S)$ can be made within a negligible additive error of $C(S)$ at all times. Our approach is the following: starting with $S = \emptyset$, we iteratively compute the vertex v maximizing $\tilde{C}(S \cup \{v\})$, adding it to S . We do this until the number of non-learners is less than T , at which point we boost S and terminate. Using classic results from the submodular optimization literature, we may show that our randomized greedy algorithm obtains a $O(\log n)$ -approximation.

LEMMA 4.8. *Let $ALG := ALG(k, n)$ be the number of edges added by Algorithm 1. Then for any $k \ll \log n$ and $M := M(n)$ as defined in Lemma 4.6, $ALG \leq O(k \log n) \cdot M$ with probability at least $1 - \frac{1}{n^8}$.*

The runtime of Algorithm 1 is polynomial, and by Lemma 4.5 the modified network will obtain learning. Taking $k = g(n)$ for any $g(n) = \omega(1)$, we obtain Theorem 4.9.

Algorithm 1 BoostGraph-MonteCarlo(k)

- 1: $S \leftarrow \emptyset, \varepsilon = \frac{1}{n^2}, \delta = \frac{1}{n^9}, N = \frac{n^2}{2\varepsilon\delta} \ln \frac{2}{\delta}$
- 2: Using a learning oracle, determine the set of non-learners V'
- 3: **while** $\tilde{C}(S) < |V'| - T$ **do**
- 4: Sample orderings $\sigma_1, \dots, \sigma_N$ uniformly at random
- 5: For each $v \in V'$, compute

$$\tilde{C}(S \cup \{v\}) = \frac{1}{N} \sum_{i=1}^N C_{\sigma_i}(S \cup \{v\})$$

- 6: Let $v^* = \arg \max_{v \in V'} \tilde{C}(S \cup \{v\})$
 - 7: Set $S \leftarrow S \cup \{v^*\}$
 - 8: **end while**
 - 9: Run BOOSTAGENTS(G, S, k)
-

THEOREM 4.9. *For any $k = \omega(1)$, $T = o(n)$, and network $\mathcal{F} = \{G_n\}$, Algorithm 1 runs in polynomial time to modify each graph in \mathcal{F} to achieve random-order learning where with probability $\geq 1 - \frac{1}{n^8}$, at most T agents are unaffected in expectation. Furthermore, the number of modifications made $O(g(n) \log n)$ -approximates the optimum, for any function $g(n) = \omega(1)$.*

4.3 Efficiently Identifying Learning Vertices

We conjecture that it is NP-hard to compute, or even reasonably approximate, the output of any random-order learning oracle as defined in Definition 4.1. Even under a fixed ordering, computing the learning rate of a given vertex is difficult. In the repeated learning setting, in which agents have multiple attempts to predict the ground truth value, [17] gives an exponential time algorithm to simulate the decision of any agent, and proves NP-hardness of computing any individual agent’s decision when the private signals may be asymmetrically assigned (i.e. $\Pr(a_v = 1 \mid \theta = 0) \neq \Pr(a_v = 1 \mid \theta = 1)$). Although their results are intended for the repeated learning setting, they are directly applicable to a graph coupled with a strategic ordering in the sequential setting. The random ordering adds an additional layer of complexity.

CONJECTURE 4.10. *For general graphs G , $v \in V$, and $t \in [0, 1]$, computing the output of a learning oracle $f(v, t)$ is NP-hard.*

We observe that one may instead identify simple local structures that permit learning for particular vertices. Heuristically, many learning networks rely on agents that learn due to a specific graph structure independent of the rest of the network, which we refer to as “first learners”. For random-order learning, one indicator of whether a vertex is a first learner with high probability over the ordering is the number of independent neighbors it has. For example, any vertex boosted by BoostAgents becomes a first learner as it observes many independent private signals from its neighbors.

We describe a simple deterministic method to check if a vertex is a first learner. In particular, one may simply compute $\sum_{u \in N(v_n)} \frac{1}{\deg(u)}$ in $O(\deg(v)) = O(n)$ time. In essence, this tests whether an agent can directly see many independent signals from its neighbors.

CLAIM 4.11. *For $v = \{v_n : v_n \in G_n\}$, if*

$$\sum_{u \in N(v)} \frac{1}{\deg(u)} = \omega(1)$$

then v learns under random orderings.

The key insight is that for any $v \in G$, the expected number of neighbors arriving first within their own neighborhood under σ is $\sum_{u \in N(v)} \frac{1}{1 + \deg(u)} = \Theta(\sum_{u \in N(v)} \frac{1}{\deg(u)})$.⁴ Each such vertex is forced to predict its own private signal. If v can see the private signals of a superconstant number of neighbors, it will be able to learn with probability approaching 1.

We observe that a graph may learn despite no agent having a high probability of being a first learner. The construction in Proposition 4.3 shows that even arbitrarily complex structures with low probabilities of staying intact under random orderings can be made essential for a network’s learning. Under strategic orderings, there exist networks capable of learning without a clear set of first learners, such as the butterfly graph described in [19]. Unlike the heuristic in Claim 4.11, which can be efficiently used to test whether a vertex is a first learner, it is unclear how one may feasibly determine if a network learns via either of these alternative mechanisms.

5 CONCLUSION

In this work, we explored the structure of truth learning in social networks under uniformly random decision orderings. Our results highlight the fragility of networks learning under strategic orderings, while simultaneously proving the robustness of random-order learning networks. We also characterized various structural requirements, such as the necessity of large independent sets to achieve learning, and introduced constructions which demonstrate surprising capacity for learning under random orderings.

Our algorithm for transforming arbitrary networks into random-order learning networks uses a greedy approach to pick agents in the network to boost. With efficient access to an oracle that can calculate random-order learning rates for any vertex, this algorithm runs in polynomial time and succeeds with high probability. However [29] proves that the fixed-order learning decision problem is NP-hard; we conjecture that a similar result holds for the random-order learning rate of a network, but speculate that an approach reducing from 3SAT as done in both [29], [10] is unlikely to work due to their reliance on a strategic ordering.

Future work may explore the computational complexity of computing and/or approximating learning rates, and the design of alternative algorithms that do not rely on such oracles. One direction may be to derive stronger relationships between a network’s learning rate and properties of its underlying graph topology. To this end, it would be interesting to obtain stronger characterizations of learning networks, especially stronger sufficiency conditions on a network that would guarantee random-order asymptotic learning.

ACKNOWLEDGMENTS

This work was conducted when Guo and Xiong were participants of the 2025 DIMACS REU program (funded by CCF-2447342) at Rutgers University. In addition, Guo was also supported by NSF 2421503 supplement to DMS-2220271 and Xiong was supported by AI Institute (ACTION) under IIS-2229876. Gao would like to acknowledge NSF support through IIS-2229876, DMS-2220271, CNS-2515159, DMS-2311064, CCF-2208663 and CCF-2118953.

⁴This value is closely related to the β -measure in directed social networks - see [30].

REFERENCES

- [1] Daron Acemoglu, Munther A Dahleh, Ilan Lobel, and Asuman Ozdaglar. 2011. Bayesian Learning in Social Networks. *Rev. Econ. Stud.* 78, 4 (March 2011), 1201–1236.
- [2] Daron Acemoglu and Asuman Ozdaglar. 2011. Opinion dynamics and learning in social networks. *Dyn. Games Appl.* 1, 1 (March 2011), 3–49.
- [3] J Aislinn Bohren. 2016. Informational Herding with Model Misspecification. *J. Econ. Theory* 163 (2016), 222–247.
- [4] Itai Arieli, Fedor Sandomirskiy, and Rann Smorodinsky. 2021. On Social Networks That Support Learning. In *Proceedings of the 22nd ACM Conference on Economics and Computation* (Budapest, Hungary) (EC '21). 95–96.
- [5] Gal Bahar, Itai Arieli, Rann Smorodinsky, and Moshe Tennenholtz. 2020. Multi-issue social learning. *Math. Soc. Sci.* 104 (March 2020), 29–39.
- [6] Abhijit V. Banerjee. 1992. A simple model of herd behavior. *Q. J. Econ.* 107, 3 (Aug. 1992), 797–817.
- [7] Sushil Bikhchandani, David Hirshleifer, and Ivo Welch. 1992. A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades. *J. Polit. Econ.* 100, 5 (Oct. 1992), 992–1026.
- [8] Christian Borgs, Michael Brautbar, Jennifer Chayes, and Brendan Lucier. 2014. Maximizing social influence in nearly optimal time. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms* (Portland, Oregon) (SODA '14). Society for Industrial and Applied Mathematics, USA, 946–957.
- [9] Christophe Chamley. 2004. *Rational Herds: Economic Models of Social Learning*. Cambridge University Press, Cambridge.
- [10] Gregory F Cooper. 1990. The computational complexity of probabilistic inference using bayesian belief networks. *Artif. Intell.* 42, 2-3 (March 1990), 393–405.
- [11] Pedro Domingos and Matt Richardson. 2001. Mining the network value of customers. In *Proceedings of the Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (San Francisco, California) (KDD '01). Association for Computing Machinery, New York, NY, USA, 57–66. <https://doi.org/10.1145/502512.502525>
- [12] Benjamin Golub and Evan Sadler. 2016. Learning in social networks. In *The Oxford Handbook of the Economics of Networks*. Oxford University Press, Oxford, 504–542.
- [13] William Guo, Edward Xiong, and Jie Gao. 2026. Robust Sequential Learning in Random Order Networks. <https://arxiv.org/abs/2602.08953>
- [14] Wade Hann-Caruthers, Vadim V Martynov, and Omer Tamuz. 2018. The speed of sequential asymptotic learning. *J. Econ. Theory* 173 (Jan. 2018), 383–409.
- [15] Wade Hann-Caruthers, Minghao Pan, and Omer Tamuz. 2025. Network and Timing Effects in Social Learning. In *Proceedings of the 26th ACM Conference on Economics and Computation*. Association for Computing Machinery, New York, NY, USA, 6. <https://doi.org/10.1145/3736252.3742486>
- [16] Jakub Házla, Ali Jadbabaie, Elchanan Mossel, and Mohammad Ali Rahimian. 2019. Reasoning in Bayesian opinion exchange networks is PSPACE-hard. In *Conference on Learning Theory*. PMLR, 1614–1648. <https://proceedings.mlr.press/v99/hazla19a.html>
- [17] Jan Házla, Ali Jadbabaie, Elchanan Mossel, and M Amin Rahimian. 2021. Bayesian decision making in groups is Hard. *Oper. Res.* 69, 2 (March 2021), 632–654.
- [18] David Kempe, Jon Kleinberg, and Eva Tardos. 2003. Maximizing the spread of influence through a social network. In *Conference on Learning Theory*. ACM, 137–146. <https://doi.org/10.1145/956750.956769>
- [19] Kevin Lu, Jordan Chong, Matt Lu, and Jie Gao. 2024. Enabling Asymptotic Truth Learning in a Social Network. In *Proceedings of the 20th Conference on Web and Internet Economics (WINE'24)*. 530–547.
- [20] Markus Mobius and Tanya Rosenblat. 2014. Social Learning in Economics. *Annu. Rev. Econom.* 6, 1 (Aug. 2014), 827–847.
- [21] Elchanan Mossel, Joe Neeman, and Omer Tamuz. 2013. Majority dynamics and aggregation of information in social networks. *Autonomous Agents and Multi-Agent Systems* 28, 3 (June 2013), 408–429. <https://doi.org/10.1007/s10458-013-9230-4>
- [22] Elchanan Mossel, Allan Sly, and Omer Tamuz. 2014. Asymptotic learning on Bayesian social networks. *Probab. Theory Relat. Fields* 158, 1-2 (Feb. 2014), 127–157.
- [23] Elchanan Mossel and Omer Tamuz. 2017. Opinion exchange dynamics. *Probab. Surv.* 14, none (Jan. 2017), 155–204.
- [24] Manuel Mueller-Frank. 2018. Manipulating Opinions in Social Networks. (Oct. 2018).
- [25] George L. Nemhauser, Laurence A. Wolsey, and Marshall L. Fisher. 1978. An analysis of approximations for maximizing submodular set functions—I. *Math. Program.* 14, 1 (Dec. 1978), 265–294.
- [26] Daniel Sgroi. 2002. Optimizing Information in the Herd: Guinea Pigs, Profits, and Welfare. *Games Econ. Behav.* 39, 1 (April 2002), 137–166.
- [27] Lones Smith. 1991. *Essays on dynamic models of equilibrium and learning*. Ph.D. Dissertation. University of Chicago.
- [28] Lones Smith and Peter Sørensen. 2000. Pathological Outcomes of Observational Learning. *Econometrica* 68, 2 (March 2000), 371–398.
- [29] Filip Úradník, Amanda Wang, and Jie Gao. 2025. Maximizing Truth Learning in a Social Network is NP-hard. In *Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025)*. 2078–2086.
- [30] René van den Brink and Robert P. Gilles. 2000. Measuring domination in directed networks. *Social Networks* 22, 2 (2000), 141–157.
- [31] Ivo Welch. 1992. Sequential sales, learning, and cascades. *J. Finance* 47, 2 (June 1992), 695–732.
- [32] Laurence A Wolsey. 1982. An analysis of the greedy algorithm for the submodular set covering problem. *Combinatorica* 2, 4 (Dec. 1982), 385–393.