

# Online Reputation and Polling Systems: Data Incest, Social Learning, and Revealed Preferences

Vikram Krishnamurthy, *Fellow, IEEE* and William Hoiles, *Member, IEEE*

(Invited Paper)

**Abstract**—This paper considers online reputation and polling systems where individuals make recommendations based on their private observations and recommendations of friends. Such interaction of individuals and their social influence is modeled as social learning on a directed acyclic graph. Data incest (misinformation propagation) occurs due to unintentional reuse of identical actions in the formation of public belief in social learning; the information gathered by each agent is mistakenly considered to be independent. This results in overconfidence and bias in estimates of the state. Necessary and sufficient conditions are given on the structure of information exchange graph to mitigate data incest. Incest removal algorithms are presented. Experimental results on human subjects are presented to illustrate the effect of social influence and data incest on decision-making. These experimental results indicate that social learning protocols require careful design to handle and mitigate data incest. The incest removal algorithms are illustrated in an expectation polling system where participants in a poll respond with a summary of their friends' beliefs. Finally, the principle of revealed preferences arising in microeconomics theory is used to parse Twitter datasets to determine if social sensors are utility maximizers and then determine their utility functions.

**Index Terms**—Afriat's theorem, Bayesian estimation, data incest, expectation polling, reputation systems, revealed preferences, social learning.

## I. INTRODUCTION

ONLINE reputation systems (Yelp, Tripadvisor, etc.) are of increasing importance in measuring social opinion. They can be viewed as sensors of social opinion—they go beyond physical sensors since user opinions/ratings (such as human evaluation of a restaurant or movie) are impossible to measure via physical sensors. Devising a fair online reputation system involves constructing a data fusion system that combines estimates of individuals to generate an unbiased estimate. This presents unique challenges from a statistical signal processing and data fusion point of view. First, humans interact with and influence other humans since ratings posted on online reputation systems strongly influence the behavior of individuals.<sup>1</sup> This can result in nonstandard information patterns due to correlations introduced by the structure of the

underlying social network. Second, due to privacy concerns, humans rarely reveal raw observations of the underlying state of nature. Instead, they reveal their decisions (ratings, recommendations, votes) which can be viewed as a low resolution (quantized) function of their raw measurements and interactions with other individuals.

### A. Motivation

This paper models how *data incest* propagates among individuals in online reputation and polling systems. Consider the following example comprising a multiagent system where agents seek to estimate an underlying state of nature. An agent visits a restaurant and obtains a noisy private measurement of the state (quality of food). She then rates the restaurant as excellent on an online reputation website. Another agent is influenced by this rating, visits the restaurant, and also gives a good rating on the online reputation website. The first agent visits the reputation site and notices that another agent has also given the restaurant a good rating—this double confirms her rating and she enters another good rating. In a fair reputation system, such “double counting” or “data incest” should have been prevented by making the first agent aware that the rating of the second agent was influenced by her own rating. The information exchange between the agents is represented by the directed graph of Fig. 1. The fact that there are two distinct paths (denoted in red) between Agent 1 at time 1 and Agent 1 at time 3 implies that the information of Agent 1 at time 1 is double counted thereby leading to a data incest event. Such data incest results in a bias in the estimate of the underlying state.

### B. Main Results and Organization

This paper has four parts. The first part, namely Section II, deals with the design of a fair online reputation system. Section II formulates the data incest problem in a multiagent system where individual agents perform social learning and exchange information on a sequence of directed acyclic graphs (DAGs). The aim is to develop a distributed data fusion protocol, which incorporates social influence constraints and provides an unbiased estimate of the state of nature at each node. Protocol 1 in Section II-D gives the complete design template of how incest can be avoided in the online reputation system. It is shown that by choosing the costs to satisfy reasonable conditions, the recommendations made by individuals are *ordinal* functions of their private observations and *monotone* in the prior information. This means that the Bayesian social learning follows simple intuitive rules and is, therefore, a useful

Manuscript received May 02, 2014; revised November 23, 2014; accepted November 23, 2014. Date of publication January 15, 2015; date of current version January 22, 2015. This work was presented in part in [1].

The authors are with the Department of Electrical and Computer Engineering, University of British Columbia, Vancouver, BC V6T 1Z4, Canada (e-mail: vikramk@ece.ubc.ca; whoiles@ece.ubc.ca).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TCSS.2014.2377891

<sup>1</sup>81% of hotel managers regularly check Tripadvisor reviews [2]. A one-star increase in the Yelp rating maps to 5%–9% revenue increase [3].

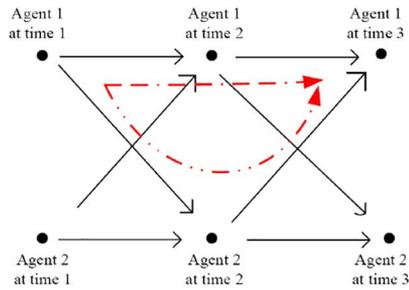


Fig. 1. Example of the information flow in a social network with two agents and over three event epochs. The arrows represent exchange of information.

idealization. Necessary and sufficient conditions for exact incest removal subject to a social influence constraint are given.

The second part of the paper, namely Section III, analyzes the data of an actual experiment that we performed on human subjects to determine how social influence affects decision-making. In particular, information flow patterns from the experimental data indicative of social learning and data incest are described. The experimental results illustrate the effect of social influence.

The third part of the paper, namely Section IV, describes how the data incest problem formulation and incest removal algorithms can be applied to an expectation polling system. Polls seek to estimate the fraction of a population that support a political party, executive decision, etc. In *intent* polling, individuals are sampled and asked who they intend to vote for. In *expectation* polling [4], individuals are sampled and asked who they believe will win the election. It is intuitive that expectation polling is more accurate than intent polling; since, in expectation, polling an individual considers its own intent together with the intents of its friends.<sup>2</sup> If the average degree of nodes in the network is  $d$ , then the savings in the number of samples is by a factor of  $d$ , since a randomly chosen node summarizes the results from  $d$  of its friends. However, the variance and bias of the estimate depend on the social network structure and data incest can strongly bias the estimate. We illustrate how the posterior distribution of the leading candidate in the poll can be estimated based on incestuous estimates.

Social learning assumes that agents (social sensors) choose their actions by maximizing a utility function. The final part of the paper, namely Section V, uses the principle of revealed preferences as a constructive test to determine: Are social sensors utility optimizers in their response to external influence? We present a remarkable result arising in microeconomics theory called Afriat's theorem [5], [6] which provides a necessary and sufficient condition for a finite dataset  $\mathcal{D}$  to have originated from a utility maximizer. The result is illustrated by tracking real-time tweets and retweets for reputation-based review Twitter accounts. In particular, Section V shows that social learning associated with reputation agencies has the structure illustrated in Fig. 2. That is, the sentiment of tweets and number

<sup>2</sup>Rothschild and Wolfers analyze all U.S. presidential electoral college results from 1952 to 2008 where both intention and expectation polling were conducted and shows a remarkable result: In 77 cases where expectation and intent polling pointed to different winners, expectation polling was accurate 78% of the time! The dataset from the American National Election Studies comprised voter responses to two questions:

*Intention:* Who will you vote for in the election for President?

*Expectation:* Who do you think will be elected President in November?

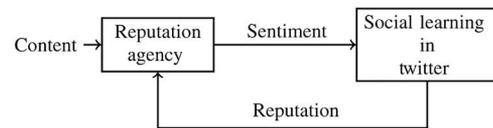


Fig. 2. Schematic of dynamics of the reputation agency and the Twitter network. Content (i.e., games, movies, and books) is provided to the reputation agency, which then publishes its review as a tweet. The Twitter network then responds to the sentiment of this tweet by retweeting. This may increase or decrease the reputation of the agency.

of followers (reputation) constitute the publicly available information that drives the social learning process. If the sentiment of the tweet published by the reputation agency improves its reputation, then this reinforces the reputation agency's belief that it adequately reviewed the content. This is an example of data incest.

### C. Related work

Reference [7] is a seminal paper in collective human behavior. The book [8] contains a complete treatment of social learning models. Social learning has been used widely in economics, marketing, political science, and sociology to model the behavior of financial markets, crowds, social groups, and social networks; see [8]–[11] and numerous references therein. Related models have been studied in sequential decision-making [12] and statistical signal processing [1], [13]. A tutorial exposition of social learning in and sensing is given in our recent paper [1]. Online reputation systems are reviewed and studied in [14]–[16]. Information caused by influential agents is investigated in [10] and [17]. In [18], examples are given that show that if just 10% of the population holds an unshakable belief, their belief will be adopted by the majority of society.

The role of social influence in decision-making (which we consider in Section II-C) is studied in [19]. The expectation polling [4] (which we consider in Section IV) is a form of *social sampling* [20] where participants in a poll respond with a summary of their friends responses. Dasgupta *et al.* [20] analyze the effect of the social network structure on the bias and variance of expectation polls. Social sampling has interesting parallels with the so-called Keynesian beauty contest, e.g., [https://en.wikipedia.org/wiki/Keynesian\\_beauty\\_contest](https://en.wikipedia.org/wiki/Keynesian_beauty_contest) for a discussion.

Data incest arises in other areas of electrical engineering. The so-called count to infinity problem in the distance vector routing protocol in packet switched networks [21] is a type of misinformation propagation. Data incest also arises in belief propagation (BP) algorithms in computer vision and error-correcting coding theory. BP algorithms require passing local messages over the graph (Bayesian network) at each iteration. For graphical models with loops, BP algorithms are only approximate due to the over-counting of local messages [22] which is similar to data incest. The algorithms presented in this paper can remove data incest in Bayesian social learning over nontree graphs that satisfy a topological constraint. In [17], data incest is considered in a network where agents exchange their private belief states (without social learning). Simpler versions were studied in [23] and [24].

Regarding revealed preferences, highly influential works in the economics literature include those by Afriat [5], [25] and Varian (chief economist at Google). In particular, Varian's work includes measuring the welfare effect of price discrimination [26], analyzing the relationship between prices of broadband Internet access and time of use service [27], and auctions for advertisement position placement on page search results from Google [27], [28].

#### D. Perspective

In Bayesian estimation, the twin effects of social learning (information aggregation with interaction amongst agents) and data incest (misinformation propagation) lead to nonstandard information patterns in estimating the underlying state of nature. Herding occurs when the public belief overrides the private observations and thus actions of agents are independent of their private observations. Data incest results in bias in the public belief as a consequence of the unintentional reuse of identical actions in the formation of public belief in social learning; the information gathered by each agent is mistakenly considered to be independent. This results in overconfidence and bias in estimates of the state.

Privacy and reputation pose conflicting requirements: privacy requirements result in noisier measurements or lower resolution actions (since individuals are not willing to disclose private observations), while a high degree of reputation requires accurate measurements. Utility functions, noisy private measurements and quantized actions are essential ingredients of the social learning models presented in this paper that facilitate modeling this tradeoff between reputation and privacy.

## II. REPUTATION DYNAMICS AND DATA INCEST

### A. Classical Social Learning

Consider a multiagent system that aims to estimate the state of an underlying finite state random variable  $x \in \mathbb{X} = \{1, 2, \dots, X\}$  with known prior distribution  $\pi_0$ . Each agent acts once in a predetermined sequential order indexed by  $k = 1, 2, \dots$ . Assume at the beginning of iteration  $k$ , all agents have access to the public belief  $\pi_{k-1}$  defined in Step 4) below. The social learning protocol proceeds as follows [8], [9].

Step 1) *Private observation*: At time  $k$ , agent  $k$  records a private observation  $y_k \in \mathbb{Y}$  from the observation distribution  $B_{iy} = P(y|x = i)$ ,  $i \in \mathbb{X}$ . Throughout this paper, we assume that  $\mathbb{Y} = \{1, 2, \dots, Y\}$  is finite.

Step 2) *Private belief*: Using the public belief  $\pi_{k-1}$  available at time  $k-1$  [Step 4) below], agent  $k$  updates its private posterior belief  $\eta_k(i) = P(x_k = i|a_1, \dots, a_{k-1}, y_k)$  using Bayes formula

$$\eta_k = \frac{B_{y_k} \pi}{\mathbf{1}'_X B_y \pi}, \quad B_{y_k} = \text{diag}(P(y_k|x = i), i \in \mathbb{X}). \quad (1)$$

Here,  $\mathbf{1}_X$  denotes the  $X$ -dimensional vector of ones and  $\eta_k$  is an  $X$ -dimensional probability mass function (pmf).

Step 3) *Myopic action*: Agent  $k$  takes action  $a_k \in \mathcal{A} = \{1, 2, \dots, A\}$  to minimize its expected cost

$$\begin{aligned} a_k &= \arg \min_{a \in \mathcal{A}} \mathbf{E}\{c(x, a)|a_1, \dots, a_{k-1}, y_k\} \\ &= \arg \min_{a \in \mathcal{A}} \{c'_a \eta_k\}. \end{aligned} \quad (2)$$

Here,  $c_a = (c(i, a), i \in \mathbb{X})$  denotes an  $X$ -dimensional cost vector, and  $c(i, a)$  denotes the cost incurred when the underlying state is  $i$  and the agent chooses action  $a$ .

Step 4) *Social learning filter*: Given the action  $a_k$  of agent  $k$ , and the public belief  $\pi_{k-1}$ , each subsequent agent  $k' > k$  performs social learning to update the public belief  $\pi_k$  according to the ‘‘social learning filter’’:

$$\pi_k = T(\pi_{k-1}, a_k), \quad T(\pi, a) = \frac{R' i_a \pi}{\sigma(\pi, a)} \quad (3)$$

where  $\sigma(\pi, a) = \mathbf{1}'_X R' i_a P' \pi$  is the normalization factor of the Bayesian update. In (3), the public belief  $\pi_k(i) = P(x_k = i|a_1, \dots, a_k)$  and  $R' i_a = \text{diag}(P(a|x = i, \pi), i \in \mathbb{X})$  has elements

$$\begin{aligned} P(a_k = a|x_k = i, \pi_{k-1} = \pi) &= \sum_{y \in \mathbb{Y}} P(a|y, \pi) P(y|x_k = i) \\ P(a_k = a|y, \pi) &= \begin{cases} 1, & \text{if } c'_a B_y P' \pi \leq c'_a B_y P' \pi, \tilde{a} \in \mathcal{A} \\ 0, & \text{otherwise.} \end{cases} \end{aligned}$$

The following result which is well known in the economics literature [8], [9].

*Theorem 2.1 ([9])*: The social learning protocol leads to an *information cascade*<sup>3</sup> in finite time with probability 1. That is, after some finite time  $\bar{k}$  social learning ceases and the public belief  $\pi_{k+1} = \pi_k$ ,  $k \geq \bar{k}$ , and all agents choose the same action  $a_{k+1} = a_k$ ,  $k \geq \bar{k}$ . ■

### B. Information Exchange Model

In comparison to Section II-A, we now consider social learning on a family of time-dependent DAGs—in such cases, apart from herding, the phenomenon of data incest arises.

Consider an online reputation system comprised agents  $\{1, 2, \dots, S\}$  that aim to estimate an underlying state of nature (a random variable). Let  $x \in \mathbb{X} = \{1, 2, \dots, X\}$  represent the state of nature (such as the quality of a restaurant/hotel) with known prior distribution  $\pi_0$ . Let  $k = 1, 2, 3, \dots$  depict epochs at which events occur. The index  $k$  marks the historical order of events and not absolute time. For simplicity, we refer to  $k$  as ‘‘time.’’

It is convenient also to reduce the coordinates of time  $k$  and agent  $s$  to a single integer index  $n$

$$n \triangleq s + S(k-1), \quad s \in \{1, \dots, S\}, \quad k = 1, 2, 3, \dots \quad (4)$$

<sup>3</sup>A *herd of agents* takes place at time  $\bar{k}$ , if the actions of all agents after time  $\bar{k}$  are identical, i.e.,  $a_k = a_{\bar{k}}$  for all time  $k > \bar{k}$ . An information cascade implies that a herd of agents occur. Trusov *et al.* quote the following anecdote of user influence and herding in a social network: ‘‘... when a popular blogger left his blogging site for a 2-week vacation, the site's visitor tally fell, and content produced by three invited substitute bloggers could not stem the decline.’’

We refer to  $n$  as a ‘‘node’’ of a time dependent information flow graph  $G_n$  which we now define. Let

$$G_n = (V_n, E_n), \quad n = 1, 2, \dots \quad (5)$$

denote a sequence of time-dependent DAGs<sup>4</sup> of information flow in the social network until and including time  $k$ . Each vertex in  $V_n$  represents an agent  $s'$  at time  $k'$  and each edge  $(n', n'')$  in  $E_n \subseteq V_n \times V_n$  shows that the information (action) of node  $n'$  (agent  $s'$  at time  $k'$ ) reaches node  $n''$  (agent  $s''$  at time  $k''$ ). It is clear that  $G_n$  is a subgraph of  $G_{n+1}$ .

The adjacency matrix  $A_n$  of  $G_n$  is an  $n \times n$  matrix with elements  $A_n(i, j)$  given by

$$A_n(i, j) = \begin{cases} 1, & \text{if } (v_j, v_i) \in E, \\ 0, & \text{otherwise,} \end{cases} \quad A_n(i, i) = 0. \quad (6)$$

The transitive closure matrix  $T_n$  is the  $n \times n$  matrix

$$T_n = \text{sgn}((\mathbf{I}_n - A_n)^{-1}) \quad (7)$$

where for matrix  $M$ , the matrix  $\text{sgn}(M)$  has elements

$$\text{sgn}(M)(i, j) = \begin{cases} 0, & \text{if } M(i, j) = 0 \\ 1, & \text{if } M(i, j) \neq 0. \end{cases}$$

Note that  $A_n(i, j) = 1$  if there is a single hop path between nodes  $i$  and  $j$ . In comparison,  $T_n(i, j) = 1$  if there exists a path (possible multihop) between  $i$  and  $j$ .

The information reaching node  $n$  depends on the information flow graph  $G_n$ . The following two sets will be used to specify the incest removal algorithms below:

$$\mathcal{H}_n = \{m : A_n(m, n) = 1\} \quad (8)$$

$$\mathcal{F}_n = \{m : T_n(m, n) = 1\}. \quad (9)$$

Thus,  $\mathcal{H}_n$  denotes the set of previous nodes  $m$  that communicate with node  $n$  in a single-hop. In comparison,  $\mathcal{F}_n$  denotes the set of previous nodes  $m$  whose information eventually arrives at node  $n$ . Thus,  $\mathcal{F}_n$  contains all possible multihop connections by which information from a node  $m$  eventually reaches node  $n$ .

*Properties of  $A_n$  and  $T_n$ :* Due to causality with respect to the time index  $k$  (information sent by an agent can only arrive at another agent at a later time instant), the following obvious properties hold (proof omitted).

*Lemma 1:* Consider the sequence of DAGs  $G_n$ ,  $n = 1, 2, \dots$

- 1) The adjacency matrices  $A_n$  are upper triangular.  $A_n$  is the upper left  $n \times n$  submatrix of  $A_{n+1}$ .
- 2) The transitive closure matrices  $T_n$  are upper triangular with ones on the diagonal. Hence,  $T_n$  is invertible.
- 3) Classical social learning of Section II-A is a trivial example with adjacency matrix  $A_n(i, j) = 1$  for  $j = i + 1$  and  $A_n(i, j) = 0$  elsewhere. ■

The appendix contains an example of data incest that illustrates the above notation.

<sup>4</sup>A DAG is a directed graph with no directed cycles. The ordering of nodes  $n = 1, 2, \dots$ , proposed here is a special case of the well-known result that the nodes of a DAG are partially orderable; see Section IV.

### C. Data Incest Model and Social Influence Constraint

Each node  $n$  receives recommendations from its immediate friends (one hop neighbors) according to the information flow graph defined above. That is, it receives actions  $\{a_m, m \in \mathcal{H}_n\}$  from nodes  $m \in \mathcal{H}_n$  and then seeks to compute the associated public beliefs  $\pi_m, m \in \mathcal{H}_n$ . If node  $n$  naively (incorrectly) assumes that the public beliefs  $\pi_m, m \in \mathcal{H}_n$  are independent, then it would fuse these as

$$\pi_{n-} = \frac{\prod_{m \in \mathcal{H}_n} \pi_m}{1'_X \prod_{m \in \mathcal{H}_n} \pi_m}. \quad (10)$$

This naive data fusion would result in data incest.

1) *Aim:* The aim is to provide each node  $n$  the true posterior distribution

$$\pi_{n-}^0(i) = P(x = i | \{a_m, m \in \mathcal{F}_n\}) \quad (11)$$

subject to the following *social influence constraint*: There exists a fusion algorithm  $\mathcal{A}$  such that

$$\pi_{n-}^0 = \mathcal{A}(\pi_m, m \in \mathcal{H}_n). \quad (12)$$

2) *Discussion: Fair Rating and Social Influence:* We briefly pause to discuss (11) and (12).

a) We call  $\pi_{n-}^0$  in (11) the *true or fair online rating* available to node  $n$  since  $\mathcal{F}_n$  defined in (9) denotes all information (multi-hop links) available to node  $n$ . By definition,  $\pi_{n-}^0$  is incest free and is the desired conditional probability that agent  $n$  needs.<sup>5</sup>

Indeed, if node  $n$  combines  $\pi_{n-}^0$  together with its own private observation via social learning, then clearly

$$\eta_n(i) = P(x = i | \{a_m, m \in \mathcal{F}_n\}, y_n), \quad i \in \mathbb{X}$$

$$\pi_n(i) = P(x = i | \{a_m, m \in \mathcal{F}_n\}, a_n), \quad i \in \mathbb{X}$$

are, respectively, the correct (incest free) private belief for node  $n$  and the correct after-action public belief. If agent  $n$  does not use  $\pi_{n-}^0$ , then incest can propagate; e.g., if agent  $n$  naively uses (10).

Why should an individual  $n$  agree to use  $\pi_{n-}^0$  to combine with its private message? It is here that the social influence constraint (12) is important.  $\mathcal{H}_n$  can be viewed as the *social message*, i.e., personal friends of node  $n$  since they directly communicate to node  $n$ , while the associated beliefs can be viewed as the *informational message*. As described in the remarkable recent paper [19], the social message from personal friends exerts a large social influence<sup>6</sup>—it provides significant incentive (peer pressure) for individual  $n$  to comply with the protocol of combining its estimate with  $\pi_{n-}^0$  and thereby prevent incest. Bond *et al.* [19] show that receiving messages

<sup>5</sup>For the reader unfamiliar with Bayesian state estimation, computing the posterior  $\pi_{n-}^0$  is crucial noisy observations. Then, the conditional mean estimate is evaluated as  $\mathbf{E}\{x | \{a_m, m \in \mathcal{F}_n\}\} = \sum_{x \in \mathbb{X}} x \pi_{n-}^0(x)$  and is the minimum variance estimate, i.e., optimal in the mean square error (MSE) sense and more generally in the Bregmann loss function sense. The conditional mean is a ‘‘soft’’ estimate and is unbiased by definition. Alternatively, the maximum *a posteriori* ‘‘hard’’ estimate is evaluated as  $\text{argmax}_x \pi_{n-}^0(x)$ .

<sup>6</sup>In a study conducted by social networking site *myYearbook*, 81% of respondents said they had received advice from friends and followers relating to a product purchase through a social site; 74% of those who received such advice found it to be influential in their decision (*Click Z*, January 2010).

from known friends has significantly more influence on an individual than the information in the messages. This study includes a comparison of information messages and social messages on Facebook and their direct effect on voting behavior. To quote [19], “The effect of social transmission on real-world voting was greater than the direct effect of the messages themselves. . .” In Section III, we provide results of an experiment on human subjects that also illustrates social influence in social learning.

#### D. Fair Online Reputation System: Protocol 1

The procedure specified in Protocol 1 evaluates the fair online rating by eliminating data incest in a social network. The aim is to achieve (11) subject to (12).

---

#### Protocol 1. Incest Removal in Online Reputation System.

---

##### (i) Information from Social Network:

- 1) *Social Message from Friends*: Node  $n$  receives social message  $\mathcal{H}_n$  comprising the names or photos of friends that have made recommendations.
- 2) *Informational Message from Friends*: The reputation system fuses recommendations  $\{a_m, m \in \mathcal{H}_n\}$  into the single informational message  $\pi_{n-}$  and presents this to node  $n$ .

The reputation system computes  $\pi_{n-}$  as follows:

- a)  $\{a_m, m \in \mathcal{H}_n\}$  are used to compute public beliefs  $\{\pi_m, m \in \mathcal{H}_n\}$  using Step 5) below.
- b)  $\{\pi_m, m \in \mathcal{H}_n\}$  are fused into  $\pi_{n-}$  as

$$\pi_{n-} = \mathcal{A}(\pi_m, m \in \mathcal{H}_n). \quad (13)$$

In Section II-G, fusion algorithm  $\mathcal{A}$  is designed as

$$l_{n-}(i) = \sum_{m \in \mathcal{H}_n} w_n(m) l_m(i), \quad i \in \mathbb{X}. \quad (14)$$

Here  $l_m(i) = \log \pi_m(i)$  and  $w_n(m)$  are weights.

(ii) *Observation*: Node  $n$  records private observation  $y_n$  from the distribution  $B_{iy} = P(y|x = i)$ ,  $i \in \mathbb{X}$ .

(iii) *Private Belief*: Node  $n$  uses  $y_n$  and informational message  $\pi_{n-}$  to update its private belief via Bayes rule

$$\eta_n = \frac{B_{y_n} \pi_{n-}}{\mathbf{1}'_X B_y \pi_{n-}}. \quad (15)$$

(iv) *Recommendation*: Node  $n$  makes recommendation

$$a_n = \arg \min_a c'_a \eta_n$$

and records this on the reputation system.

(v) *Public Belief Update by Network Administrator*: Based on recommendation  $a_n$ , the reputation system (automated algorithm) computes the public belief  $\pi_n$  using the social learning filter (3).

---

At this stage, the public rating  $\pi_{n-}$  computed in (13) of Protocol 1 is not necessarily the fair online rating  $\pi_n^0$  of

(11). Without careful design of algorithm  $\mathcal{A}$  in (13), due to inter-dependencies of actions on previous actions,  $\pi_{n-}$  can be substantially different from  $\pi_n^0$ . Then,  $\eta_n$  computed via (15) will not be the correct private belief and incest will propagate in the network. In other words,  $\eta_n$ ,  $\pi_{n-}$ , and  $\pi_n$  are defined purely in terms of their computational expressions in Protocol 1; they are not necessarily the desired conditional probabilities, unless algorithm  $\mathcal{A}$  is properly designed to remove incest. Note also the requirement that algorithm  $\mathcal{A}$  needs to satisfy the social influence constraint (12).

#### E. Ordinal Decision Making in Protocol 1

Protocol 1 assumes that each agent is a Bayesian utility optimizer. The following discussion shows that under reasonable conditions, such a Bayesian model is a useful idealization of agents' behaviors.

Humans typically make *monotone* decisions—the more favorable the private observation, the higher the recommendation. Humans make *ordinal* decisions<sup>7</sup> since humans tend to think in symbolic ordinal terms. Under what conditions is the recommendation  $a_n$  made by node  $n$  *monotone increasing* in its observation  $y_n$  and *ordinal*? Recall from Steps 3) and 4) of Protocol 1 that the recommendation of node  $n$  is

$$a_n(\pi_{n-}^0, y_n) = \arg \min_a c'_a B_{y_n} \pi_{n-}^0.$$

So, an equivalent question is: Under what conditions is the argmin increasing in observation  $y_n$ ? Note that an increasing argmin is an *ordinal* property—i.e.,  $\arg \min_a c'_a B_{y_n} \pi_{n-}^0$  increasing in  $y$  implies  $\arg \min_a \phi(c'_a B_{y_n} \pi_{n-}^0)$  is also increasing in  $y$  for any monotone function  $\phi(\cdot)$ .

The following result gives sufficient conditions for each agent to give a recommendation that is monotone and ordinal in its private observation.

**Theorem 2.2:** Suppose the observation probabilities and costs satisfy the following conditions.

- (A1)  $B_{iy}$  is TP2 (totally positive of order 2), i.e.,  $B_{i+1,y} B_{i,y+1} \leq B_{i,y} B_{i+1,y+1}$ .
- (A2)  $c(x, a)$  is submodular, i.e.,  $c(x, a+1) - c(x, a) \leq c(x+1, a+1) - c(x+1, a)$ .

Then,

- 1) Under (A1) and (A2), the recommendation  $a_n(\pi_{n-}^0, y_n)$  made by agent  $n$  is increasing and hence ordinal in observation  $y_n$ , for any  $\pi_{n-}^0$ .
- 2) Under (A2),  $a_n(\pi_{n-}^0, y_n)$  is increasing in belief  $\pi_{n-}^0$  with respect to the monotone likelihood ratio (MLR) stochastic order<sup>8</sup> for any observation  $y_n$ . ■

The proof is in the appendix. We can interpret the above theorem as follows. If agents make recommendations that are monotone and ordinal in the observations and monotone in the prior, then they mimic the Bayesian social learning model. Even

<sup>7</sup>Humans typically convert numerical attributes to ordinal scales before making decisions. For example, it does not matter if the cost of a meal at a restaurant is \$200 or \$205; an individual would classify this cost as “high.” Also, credit rating agencies use ordinal symbols such as AAA, AA, A.

<sup>8</sup>Given probability mass functions  $\{p_i\}$  and  $\{q_i\}$ ,  $i = 1, \dots, X$  then  $p$  MLR dominates  $q$  if  $\log p_i - \log p_{i+1} \leq \log q_i - \log q_{i+1}$ .

if the agent does not exactly follow a Bayesian social learning model, its monotone ordinal behavior implies that such a Bayesian model is a useful idealization.

Condition (A1) is widely studied in monotone decision-making; see the classical book by Karlin [30] and [31]; numerous examples of noise distributions are TP2. Indeed in the highly cited paper [32] in the economics literature, observation  $y + 1$  is said to be more “favorable news” than observation  $y$  if Condition (A1) holds.

Condition (A2) is the well-known submodularity condition [33]. Actually (A2) is a stronger version of the more general single-crossing condition [34], [35] stemming from the economics literature (see appendix)

$$(c_{a+1} - c_a)'B_{y+1}\pi^0 \geq 0 \implies (c_{a+1} - c_a)'B_y\pi^0 \geq 0.$$

This single-crossing condition is ordinal, since for any monotone function  $\phi$ , it is equivalent to

$$\phi((c_{a+1} - c_a)'B_{y+1}\pi^0) \geq 0 \implies \phi((c_{a+1} - c_a)'B_y\pi^0) \geq 0.$$

(A2) also makes sense in a reputation system for the costs to be well posed. Suppose, the recommendations in action set  $\mathcal{A}$  are arranged in increasing order and also the states in  $\mathbb{X}$  for the underlying state are arranged in ascending order. Then, (A2) says: if recommendation  $a + 1$  is more accurate than recommendation  $a$  for state  $x$ ; then recommendation  $a + 1$  is also more accurate than recommendation  $a$  for state  $x + 1$  (which is a higher quality state than  $x$ ).

In the experiment results reported in Section III, we found that (A1) and (A2) of Theorem 2.2 are justified.

#### F. Discussion and Properties of Protocol 1

For the reader’s convenience, we provide an illustrative example of data incest in the appendix. We now discuss several other properties of Protocol 1.

1) *Individuals Have Selective Memory:* Protocol 1 allows for cases where each node can remember some (or all) of its past actions or none. This models cases where people forget most of the past except for specific highlights. For example, in the information flow graph of the illustrative example in the appendix (Fig. 13), if nodes 1, 3, 4, and 7 are assumed to be the same individual, then at node 7, the individual remembers what happened at nodes 5 and 1, but not node 3.

2) *Security: Network and Data Administrator:* The social influence constraint (12) can be viewed as a separation of privilege requirement for network security. For example, the National Institute of Standards and Technology (NIST) of the U.S. Department of Commerce [36, Sec. 2.4] recommends separating the roles of data and systems administrator. Protocol 1 can then be interpreted as follows: A network administrator has access to the social network graph (but not the data) and can compute weights  $w_n(m)$  (13) by which the estimates are weighed. A data administrator has access to the recommendations of friends (but not to the social network). Combining the weights with the recommendations yields the informational message  $\pi_{n-}^0$  as in (14).

3) *Automated Recommender System:* Steps 1) and 5) of Protocol 1 can be combined into an automated recommender system that maps previous actions of agents in the social group to a single recommendation (rating)  $\pi_{n-}$  of (13). This recommender system can operate completely opaquely to the actual user (node  $n$ ). Node  $n$  simply uses the automated rating  $\pi_{n-}$  as the current best available rating from the reputation system. Actually, Algorithm  $\mathcal{A}$  presented below fuses the beliefs in a linear fashion. A human node  $n$  receiving an informational message comprised a linear combination of recommendation of friends, along with the social message has incentive to follow the protocol as described in Section II-C.

4) *Agent Reputation:* The cost function minimization in Step 4) can be interpreted in terms of the reputation of agents in online reputation systems. If an agent continues to write bad reviews for high-quality restaurants on Yelp, her reputation becomes lower among the users. Consequently, other people ignore reviews of that (low-reputation) agent in evaluating their opinions about the social unit under study (restaurant). Therefore, agents minimize the penalty of writing inaccurate reviews.

#### G. Incest Removal Algorithm $\mathcal{A}$

It only remains to describe the construction of algorithm  $\mathcal{A}$  in Step 2b) of Protocol 1 so that

$$\pi_{n-}(i) = \pi_{n-}^0(i), \quad i \in \mathbb{X}$$

where  $\pi_{n-}^0(i) = P(x = i | \{a_m, m \in \mathcal{F}_n\})$ . (16)

To describe algorithm  $\mathcal{A}$ , we make the following definitions: Recall  $\pi_{n-}^0$  in (11) is the *fair online rating* available to node  $n$ . It is convenient to work with the logarithm of the un-normalized belief: accordingly define

$$l_n(i) \propto \log \pi_n(i), \quad l_{n-}(i) \propto \log \pi_{n-}(i), \quad i \in \mathbb{X}.$$

Define the  $n - 1$  dimensional weight vector

$$w_n = T_{n-1}^{-1}t_n. \quad (17)$$

Recall that  $t_n$  denotes the first  $n - 1$  elements of the  $n$ th column of the transitive closure matrix  $T_n$ . Thus, the weights are purely a function of the adjacency matrix of the graph and do not depend on the observed data.

We present algorithm  $\mathcal{A}$  in two steps: 1) the actual computation is given in Theorem 2.3, 2) necessary and sufficient conditions on the information flow graph for the existence of such an algorithm to achieve the social influence constraint (12).

*Theorem 2.3 (Fair Rating Algorithm):* Consider the reputation system with Protocol 1. Suppose, the network administrator runs the following algorithm in (13):

$$l_{n-}(i) = \sum_{m=1}^{n-1} w_n(m) l_m(i) \quad (18)$$

where the weights  $w_n$  are chosen according to (17).

Then,  $l_{n-}(i) \propto \log \pi_{n-}^0(i)$ . That is, the fair rating  $\log \pi_{n-}^0(i)$  defined in (11) is obtained via (18). ■

Theorem 2.3 says that the fair rating  $\pi_{n-}^0$  can be expressed as a linear function of the action log-likelihoods in terms of the transitive closure matrix  $T_n$  of graph  $G_n$ .

### 1) Achievability of Fair Rating by Protocol 1:

- 1) Algorithm  $\mathcal{A}$  at node  $n$  specified by (13) needs to satisfy the social influence constraint (12)—i.e., it needs to operate on beliefs  $l_m, m \in \mathcal{H}_n$ .
- 2) On the other hand, to provide incest free estimates, algorithm  $\mathcal{A}$  specified in (18) requires all previous beliefs  $l_{1:n-1}(i)$  that are specified by the nonzero elements of the vector  $w_n$ .

The only way to reconcile points 1 and 2 is to ensure that  $A_n(j, n) = 0$  implies  $w_n(j) = 0$  for  $j = 1, \dots, n-1$ . This condition means that the single-hop past estimates  $l_m, m \in \mathcal{H}_n$  available at node  $n$  according to (13) in Protocol 1 provide all the information required to compute  $w'_n l_{1:n-1}$  in (18). We formalize this condition in the following theorem.

*Theorem 2.4 (Achievability of Fair Rating):* Consider the fair rating algorithm specified by (18). For Protocol 1 using the social influence constraint information  $(\pi_m, m \in \mathcal{H}_n)$  to achieve the estimates  $l_{n-}$  of algorithm (18), a necessary and sufficient condition on the information flow graph  $G_n$  is

$$A_n(j, n) = 0 \implies w_n(j) = 0. \quad (19)$$

Therefore, for Protocol 1 to generate incest free estimates for nodes  $n = 1, 2, \dots$ , condition (19) needs to hold for each  $n$ . [Recall  $w_n$  is specified in (18).] ■

*Summary:* Algorithm (18) together with the condition (19) ensure that incest free estimates are generated by Protocol 1 that satisfy social influence constraint (12).

## III. EXPERIMENTAL RESULTS ON HUMAN SUBJECTS

To illustrate social learning, data incest and social influence, this section presents an actual psychology experiment that was conducted by our colleagues at the Department of Psychology of University of British Columbia in September and October, 2013, see [37] for details. The participants comprised 36 undergraduate students who participated in the experiment for course credit.

### A. Experiment Setup

The experimental study involved 1658 individual trials. Each trial comprised two participants who were asked to perform a perceptual task interactively. The perceptual task was as follows: Two arrays of circles denoted left and right were given to each pair of participants. Each participant was asked to judge which array (left or right) had the larger average diameter. The participants answer (left or right) constituted their action. So, the action space is  $\mathcal{A} = \{0(\text{left}), 1(\text{right})\}$ .

The circles were prepared for each trial as follows: two  $4 \times 4$  grids of circles were generated by uniformly sampling from the radii:  $\{20, 24, 29, 35, 42\}$  (in pixels). The average diameter of each grid was computed, and if the means differed by more than 8% or less than 4%, new grids were made. Thus in each trial, the left and right-array of circles differed in the average diameter by 4%–8%. Fig. 3 shows a typical example.

For each trial, one of the two participants was chosen randomly to start the experiment by choosing an action according to his/her observation. Thereafter, each participant was given

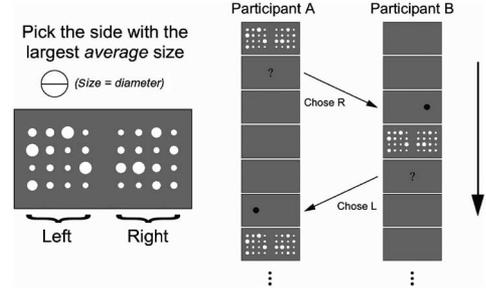


Fig. 3. Two arrays of circles were given to each pair of participants on a screen. Their task is to interactively determine which side (either left or right) had the larger average diameter. The partner’s previous decision was displayed on screen prior to the stimulus.

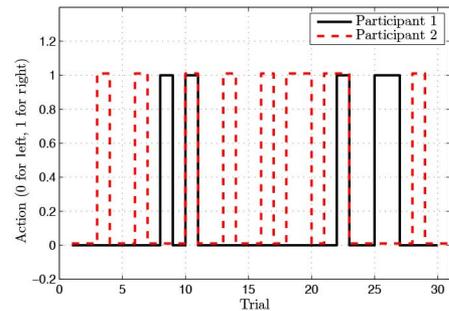


Fig. 4. Example of sample path of actions chosen by two participants in a single trial of the experiment. In this trial, both participants eventually chose the correct answer 0 (left).

access to their partner’s previous response (action) and the participants own previous action prior to making his/her judgment. This mimics the social learning Protocol 1 of Section II-D. The participants continued choosing actions according to this procedure until the experiment terminated. The trial terminated when the response of each of the two participants did not change for three successive iterations (the two participants did not necessarily have to agree for the trial to terminate).

In each trial, the actions of participants were recorded along with the time interval taken to choose their action. As an example, Fig. 4 illustrates the sample path of decisions made by the two participants in one of the 1658 trials. In this specific trial, the average diameter of the left array of circles was 32.1875 and the right array was 30.5625 (in pixels); so the ground truth was 0 (left).

### B. Experimental Results

The results of our experimental study are as follows:

- 1) *Social Learning Model:* As mentioned above, the experiment for each pair of participants was continued until both participants’ responses stabilized. *In what percentage of these experiments, did an agreement occur between the two participants?* The answer to this question reveals whether “herding” occurred in the experiments and whether the participants performed social learning (influenced by their partners). The experiments show that in 66% of trials (1102 among 1658), participants reached an agreement; that is herding occurred.

Further, in 32% of the trials, both participants converged to the correct decision after a few interactions.

To construct a social learning model for the experimental data, we consider the experiments where both participants reached an agreement. Define the social learning success rate as

$$\frac{\# \text{ expts where both participants chose correct answer}}{\# \text{ expts where both participants reached an agreement}}.$$

In the experimental study, the state space is  $\mathbb{X} = \{0, 1\}$  where  $x = 0$ , when the left array of circles has the larger diameter and  $x = 1$ , when the right array has the larger diameter. The initial belief for both participants is considered to be  $\pi_0 = [0.5, 0.5]$ . The observation space is assumed to be  $\mathbb{Y} = \{0, 1\}$ .

To estimate the social learning model parameters [observation probabilities  $B_{iy}$  and costs  $c(i, a)$ ], we determined the parameters that best fit the learning success rate of the experimental data. The best fit parameters obtained were<sup>9</sup>

$$B_{iy} = \begin{bmatrix} 0.61 & 0.39 \\ 0.41 & 0.59 \end{bmatrix}, \quad c(i, a) = \begin{bmatrix} 0 & 2 \\ 2 & 0 \end{bmatrix}.$$

Note that  $B_{iy}$  and  $c(i, a)$  satisfy both the conditions of the Theorem 2.2, namely TP2 observation probabilities and single-crossing cost. This implies that the subjects of this experiment made monotone and ordinal decisions.

2) *Data Incest*: Here, we study the effect of information patterns in the experimental study that can result in data incest. Since private observations are highly subjective and participants did not document these, we cannot claim with certainty if data incest changed the action of an individual. However, from the experimental data, we can localize specific information patterns that can result in incest. In particular, we focus on the two information flow graphs depicted in Fig. 5. In the two graphs of Fig. 5, the action of the first participant at time  $k$  influenced the action of the second participant at time  $k + 1$ , and thus, could have been double counted by the first participant at time  $k + 2$ . We found that in 79% of experiments, one of the information patterns shown in Fig. 5 occurred (1303 out of 1658 experiments). Further, in 21% of experiments, the information patterns shown in Fig. 5 occurred and at least one participant changed his/her decision, i.e., the judgment of participant at time  $k + 1$  differed from his/her judgments at time  $k + 2$  and  $k$ . These results show that even for experiments involving two participants, data incest information patterns occur frequently (79%) and causes individuals to modify their actions (21%). It shows that social learning protocols require careful design to handle and mitigate data incest.

#### IV. BELIEF-BASED EXPECTATION POLLING

We now move on to the third part of the paper, namely data incest in expectation polling. Unlike previous sections, the agents no longer are assumed to eliminate incest at each

<sup>9</sup>Parameter estimation in social learning is a challenging problem not addressed in this paper. Due to the formation of cascades in finite time, construction of an asymptotically consistent estimator is impossible, since actions after the formation of a cascade contain no information.

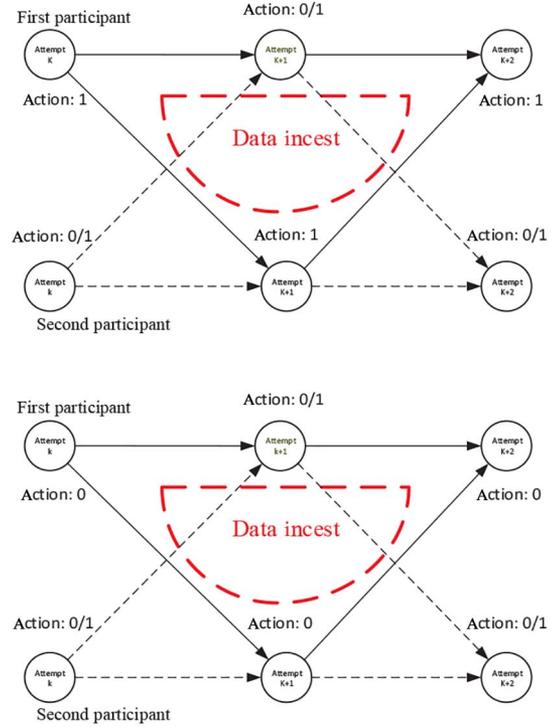


Fig. 5. Two information patterns from our experimental studies, which can result in data incest.

step. So, incest propagates in the network. Given the incestious beliefs, the aim is to compute the posterior of the state. We illustrate how the results of Section II can be used to eliminate data incest (and therefore bias) in expectation polling systems. Recall from Section I that in expectation polling [4], [20], individuals are sampled and asked who they believe will win the election; as opposed to intent polling where individuals are sampled and asked who they intend to vote for. The bias of the estimate from expectation polling depends strongly on the social network structure. Our approach below is Bayesian: we compute the posterior and therefore the conditional mean estimate (see Footnote 5).

We consider two formulations as follows.

- 1) An expectation polling system where in addition to specific polled voters, the minimal number of additional voters is recruited. The pollster then is able to use the incestious beliefs to compute the posterior conditioned on the observations and thereby eliminate incest (Section IV-A below).
- 2) An expectation polling system when it is not possible to recruit additional voters. The pollster then can only compute the posterior conditioned on the incestious beliefs. (Section IV-B below).

The first approach can be termed as “exact” since the estimate computed based on the incestious beliefs is equivalent to the estimate computed based on the private observations of the sampled voters. The second approach, although optimal given the available information, has a higher variance (see footnote 10).

Suppose  $X$  candidates contest an election. Let  $x \in \mathbb{X} = \{1, 2, \dots, X\}$  denote the candidate that is leading among the

voters, i.e.,  $x$  is the true state of nature. There are  $N$  voters. These voters communicate their expectations via a social network according to the steps listed in Protocol 2. We index the voters as nodes  $n \in \{1, \dots, N\}$  as follows:  $m < n$  if there exists a directed path from node  $m$  to node  $n$  in the DAG. It is well known that such a labeling of nodes in a DAG constitute a partially ordered set. The remaining nodes can be indexed arbitrarily, providing the above partial ordering holds.

Protocol 2 is similar to Protocol 1, except that voters exchange their expectations (beliefs) of who the leading candidate is, instead of recommendations. (So unlike the previous section, agents do not perform social learning.)

---

**Protocol 2.** Belief-Based Expectation Polling Protocol.

---

- 1) *Intention From Friends:* Node  $n$  receives the beliefs  $\{\pi_m, m \in \mathcal{H}_n\}$  of who is leading from its immediate friends, namely nodes  $\mathcal{H}_n$ .
- 2) *Naive Data Fusion:* Node  $n$  fuses the estimates of its friends naively (and therefore with incest) as

$$\pi_{n-} = \frac{\prod_{m \in \mathcal{H}_n} \pi_m}{\mathbf{1}'_X \prod_{m \in \mathcal{H}_n} \pi_m}.$$

- 3) *Observation:* Node  $n$  records its private observation  $y_n$  from the distribution  $B_{iy} = P(y|x = i)$ ,  $i \in \mathbb{X}$  of who the leading candidate is.
- 4) *Belief Update:* Node  $n$  uses  $y_n$  to update its belief via Bayes formula

$$\pi_n = \frac{B_{y_n} \pi_{n-}}{\mathbf{1}'_X B_{y_n} \pi_{n-}}. \quad (20)$$

- 5) Node  $n$  sends its belief  $\pi_n$  to subsequent nodes as specified by the social network graph.
- 

*Remark:* Similar to Theorem 2.2, the Bayesian update (20) in Protocol 2 can be viewed as an idealized model; under assumption (A1), the belief  $\pi_n$  increases with respect to the observation  $y_n$  (in terms of the MLR order). So even if agents are not actually Bayesian, if they choose their belief to be monotone in the observation, they mimic the Bayesian update.

### A. Exact Incest Removal in Expectation Polling

Assuming that the voters follow Protocol 2, the polling agency (pollster) seeks to estimate the leading candidate (state of nature  $x$ ) by sampling a subset of the  $N$  voters.

1) *Setup and Aim:* Let  $\mathcal{R} = \{\mathcal{R}_1, \mathcal{R}_2, \dots, \mathcal{R}_L\}$  denote the  $L - 1$  sampled *recruited* voters together with node  $\mathcal{R}_L$  that denotes the pollster. For example, the  $L - 1$  voters could be volunteers or paid recruits that have already signed up for the polling agency. Since, by Protocol 2, these voters have naively combined the intentions of their friends (by ignoring dependencies), the pollster needs to poll additional voters to eliminate incest. Let  $\mathcal{E}$  denote this set of *extra* (additional) voters to poll. Clearly, the choice of  $\mathcal{E}$  will depend on  $\mathcal{R}$  and the structure of the social network.

What is the smallest set  $\mathcal{E}$  of extra voters to poll in order to compute the posterior  $P(x|y_1, \dots, y_{\mathcal{R}_L})$  of state  $x$  (and therefore eliminate data incest)?

2) *Incest Removal:* Assume that all the recruited  $L - 1$  nodes in  $\mathcal{R}$  report to a central polling node  $\mathcal{R}_L$ . We assume that the pollster has complete knowledge of the network, e.g., in an online social network like Facebook.

Using the formulation of Section II-B, the pollster constructs the DAG  $G_n = (V_n, E_n)$ ,  $n \in \{1, \dots, \mathcal{R}_L\}$ . The methodology of Section II-G is straightforwardly used to determine the minimal set of extra (additional) voters  $\mathcal{E}$ . The procedure is as follows.

- 1) For each  $n = 1, \dots, \mathcal{R}_L$ , compute the weight vector  $w_n = T_{n-1}^{-1} t_n$ ; see (17).
- 2) The indices of the nonzero elements in  $w_n$  that are not in  $V_n$  constitute the minimal additional voters (nodes) that need to be polled. This is due to the necessity and sufficiency of (19) in Theorem 2.4.
- 3) Once the additional voters in  $\mathcal{E}$  are polled for their beliefs, the polling agency corrects belief  $\pi_n$  reported by node  $n \in \mathcal{R}$  to remove incest as

$$l_n^0(i) = l_n(i) + \sum_{m \in \mathcal{E}, m < n} w_n(m) l_m(i). \quad (21)$$

### B. Bayesian Expectation Polling Using Incestious Beliefs

Consider Protocol 2. However, unlike the previous subsection, assume that the pollster *cannot* sample additional voters  $\mathcal{E}$  to remove incest; e.g., the pollster is unable to reach marginalized sections of the population.

Given the incest containing beliefs  $\pi_{\mathcal{R}_1}, \dots, \pi_{\mathcal{R}_L}$  of the  $L$  recruits and pollster, how can the pollster compute the posterior  $P(x|\pi_{\mathcal{R}_1}, \dots, \pi_{\mathcal{R}_L})$  and therefore, the unbiased conditional mean estimate  $\mathbf{E}\{x|\pi_{\mathcal{R}_1}, \dots, \pi_{\mathcal{R}_L}\}$  where  $x$  is the state of nature<sup>10</sup>?

1) *Optimal (Conditional Mean) Estimation Using Incestious Beliefs:* Define the following notation:

$$\begin{aligned} \pi_{\mathcal{R}} &= [\pi_{\mathcal{R}_1} \dots \pi_{\mathcal{R}_L}]', & l_{\mathcal{R}} &= [l_{\mathcal{R}_1} \dots l_{\mathcal{R}_L}]' \\ Y_{\mathcal{R}} &= [y_1 \dots y_{\mathcal{R}_L}]', & o(x) &= [\log b_{x y_1} \dots \log b_{x y_{\mathcal{R}_L}}]'. \end{aligned}$$

*Theorem 4.1:* Consider the beliefs  $\pi_{\mathcal{R}_1}, \dots, \pi_{\mathcal{R}_L}$  of the  $L$  recruits and pollster in an expectation poll operating according to Protocol 2. Then, for each  $x \in \mathbb{X}$ , the posterior is evaluated as

$$P(x|\pi_{\mathcal{R}}) \propto \sum_{Y_{\mathcal{R}} \in \mathcal{Y}} \prod_{m=1}^{\mathcal{R}_L} B_{x y_m} \pi_0(x). \quad (22)$$

Here,  $\mathcal{Y}$  denotes the set of sequences  $\{Y_{\mathcal{R}}\}$  that satisfy

$$\begin{aligned} \mathcal{O} o(x) &= l_{\mathcal{R}}(x) - \mathcal{O} e_1 l_0(x) \\ \text{where } \mathcal{O} &= [e_{\mathcal{R}_1} \ e_{\mathcal{R}_2} \ \dots \ e_{\mathcal{R}_L}]' (I - A')^{-1}. \end{aligned} \quad (23)$$

<sup>10</sup>The price to pay for not recruiting additional voters is an increase in variance of the Bayesian estimate. It is clear that  $E\{x|y_1, \dots, y_{\mathcal{R}_L}\}$  computed using additional voters in Section IV-A has a lower variance (and is thus a more accurate estimator) than  $\mathbf{E}\{x|\pi_{\mathcal{R}_1}, \dots, \pi_{\mathcal{R}_L}\}$  since the sigma-algebra generated by  $\pi_{\mathcal{R}_1}, \dots, \pi_{\mathcal{R}_L}$  is a subset of that generated by  $y_1, \dots, y_{\mathcal{R}_L}$ . Of course, any conditional mean estimator is unbiased by definition.

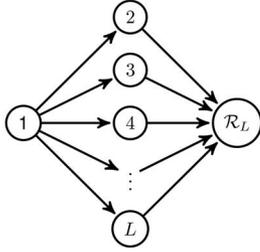


Fig. 6. Expectation polling where social network structure can result in significant bias. Node 1 has undue influence on the beliefs of other nodes. Node  $\mathcal{R}_L = L + 1$  represents the pollster. Only sampled voters are shown.

Recall  $A$  is the adjacency matrix and  $e_m$  denotes the unit  $\mathcal{R}_L$  dimension vector with 1 in the  $m$ th position. ■

The above theorem asserts that given the incest containing beliefs of the  $L$  recruits and pollster, the posterior distribution of the candidates can be computed via (22). The conditional mean estimate or maximum *a posteriori* estimate can then be computed as in footnote 5.

### C. An Extreme Example of Incest in Expectation Polling

The following example is a Bayesian version of polling in a social network described in [20]. We show that due to data incest, expectation polling can be significantly biased. Then, the methods of Sections IV-A and IV-B for eliminating incest are illustrated.

Consider the social network in Fig. 6 which represents an expectation polling system. The  $L - 1$  recruited nodes are denoted as  $\mathcal{R} = \{2, 3, \dots, L\}$ . These sampled nodes report their beliefs to the polling node  $L + 1$ . Since the poll only considers sampled voters, for notational convenience, we ignore labeling the remaining voters; therefore, Fig. 2 only shows  $L + 1$  nodes.

An important feature of the graph of Fig. 6 is that all recruited nodes are influenced by node 1. This unduly affects the estimates reported by every other node. For large  $L$ , even though node 1 constitutes a negligible fraction of the total number of voters, it significantly biases the estimate of  $x$  due to incest.<sup>11</sup>

To illustrate an extreme case of bias due to data incest, suppose that  $\mathbb{X} = \{1, 2\}$  (so there are  $X = 2$  candidates). Consider Fig. 6 with a total of  $L = 6$  recruited nodes; so  $\mathcal{R} = \{2, \dots, 8\}$ . Assume the private observations recorded at the nodes 1, 2,  $\dots$ , 8 are, respectively,  $[2, 1, 1, 1, 1, 1, 1, 2]$ . Suppose the true state of nature is  $x = 1$ , i.e., candidate 1 is leading the poll. The nodes exchange and compute their beliefs according to Protocol 2. For prior  $\pi_0 = [0.5, 0.5]'$  and observation matrix  $B = \begin{bmatrix} 0.8 & 0.2 \\ 0.2 & 0.8 \end{bmatrix}$ , it is easily verified that  $\pi_8(1) = 0.2$ ,  $\pi_8(2) = 0.8$ . That is, even though all six samples recorded candidate 1 as winning (based on private observations) and the

<sup>11</sup>In [20], a novel method is proposed to achieve unbiased expectation polling—weigh the estimate of each node by the reciprocal of its degree distribution, or alternatively sample nodes with probability inversely proportional to their degree distribution. Then, highly influential nodes such as node 1 in Fig. 6 cannot bias the estimate. Our paper is motivated by Bayesian considerations where we are interested in estimating the optimal (conditional mean) estimate which by definition is unbiased.

true state is 1, the belief  $\pi_8$  of the pollster is significantly biased toward candidate 2.

Next, we examine the two methods of incest removal proposed in Sections IV-A and IV-B, respectively. Due to the structure of adjacency matrix  $A$  of the network in Fig. 6, condition (19) does not hold and therefore exact incest removal is not possible unless node 1 is also sampled. Accordingly, suppose node 1 is sampled in addition to nodes  $\{2, \dots, 7\}$  and data incest is removed via algorithm (21). Then, the incest free estimate is  $P(x = 1|y_1, \dots, y_8) = 0.9961$ .

Finally, consider the case where node 1 cannot be sampled. Then using (22), the posterior is computed as  $P(x = 1|\pi_8) = 0.5544$ .

Comparing the three estimates for candidate 1, namely, 0.2 (naive implementation of Protocol 2 with incest), 0.9961 (with optimal incest removal), and 0.5544 (incest removal based on  $\pi_8$ ), one can see that naive expectation polling is significantly biased—recall that the ground truth is that candidate 1 is the winning candidate.

## V. REVEALED PREFERENCES AND SOCIAL LEARNING

A key assumption in the social learning models formulated in this paper is that agents are utility maximizers. In microeconomics, the principle of revealed preferences seeks to determine if an agent is an utility maximizer subject to budget constraints based on observing its choices over time. In this section, we will use the principle of revealed preferences on Twitter datasets to illustrate social learning dynamics in reputation agencies as a function of external influence (public belief).

### A. Afriat's Theorem

Given a time-series of data  $\mathcal{D} = \{(\pi_t, a_t), t \in \{1, 2, \dots, T\}\}$  where  $\pi_t \in \mathbb{R}^m$  denotes the public belief<sup>12</sup>,  $a_t$  denotes the response of agent, and  $t$  denotes the time index, is it possible to detect if the agent is a *utility maximizer*? An agent is a *utility maximizer* at each time  $t$  if for every public belief  $\pi_t$ , the chosen response  $a_t$  satisfies

$$a_t(\pi_t) \in \arg \max_{\{\pi'_t a_t \leq I_t\}} u(a) \quad (24)$$

with  $u(a)$  a nonsatiated utility function. Nonsatiated means that an increase in any element of response  $a$  results in the utility function increasing.<sup>13</sup> As shown by Diewert [38], without local nonsatiation the maximization problem (24) may have no solution.

In (24), the budget constraint  $\pi'_t a_t \leq I_t$  denotes the total amount of resources available to the social sensor for selecting the response  $x$  to the public belief  $\pi_t$ . In Section V-B, we will interpret this as a social impact budget.

Afriat's theorem [5], [25] provides a necessary and sufficient condition for a finite dataset  $\mathcal{D}$  to have originated from an utility maximizer.

<sup>12</sup>In this section, the public belief is no longer a probability mass function. Instead, it reflects the public's perception of the reputation agency based on the number of followers and tweet sentiment.

<sup>13</sup>The nonsatiated assumption rules out trivial cases such as a constant utility function, which can be optimized by any response.

*Theorem 5.1 (Afriat's Theorem):* Given a dataset  $\mathcal{D} = \{(\pi_t, a_t) : t \in \{1, 2, \dots, T\}\}$ , the following statements are equivalent.

- 1) The agent is a utility maximizer and there exists a nonsatiated and concave utility function that satisfies (24).
- 2) For scalars  $u_t$  and  $\lambda_t > 0$ , the following set of inequalities has a feasible solution:

$$u_{\tilde{a}}u - u_t - \lambda_t \pi'_t(a_{\tilde{a}}u - a_t) \leq 0$$

for  $t, \tilde{a}u \in \{1, 2, \dots, T\}$ . (25)

- 3) A nonsatiated and concave utility function that satisfies (24) is given by

$$u(a) = \min_{t \in T} \{u_t + \lambda_t \pi'_t(a - a_t)\}. \quad (26)$$

- 4) The dataset  $\mathcal{D}$  satisfies the generalized axiom of revealed preference (GARP), namely for any  $k \leq T$ ,  $\pi'_k a_k \geq \pi'_t a_{t+1} \quad \forall t \leq k-1 \implies \pi'_k a_k \leq \pi'_t a_1$ . ■

As pointed out in [6], an interesting feature of Afriat's theorem is that if the dataset can be rationalized by a nontrivial utility function, then it can be rationalized by a continuous, concave, and monotonic utility function. That is, violations of continuity, concavity, or monotonicity cannot be detected with a finite number of observations.

Verifying GARP (statement 4 of Theorem 5.1) on a dataset  $\mathcal{D}$  comprising  $T$  points can be done using Warshall's algorithm with  $O(T^3)$  [6], [39] computations. Alternatively, determining if Afriat's inequalities (25) are feasible that can be done via a LP feasibility test (using, e.g., interior point methods [40]). Note that the utility (26) is not unique and is ordinal by construction. Ordinal means that any monotone increasing transformation of the utility function will also satisfy Afriat's theorem. Therefore, the utility mimics the ordinal behavior of humans (see also Section II-E). Geometrically, the estimated utility (26) is the lower envelop of a finite number of hyperplanes that is consistent with the dataset  $\mathcal{D}$ .

### B. Example: Twitter Data of Online Reputation Agencies

In this section, we illustrate social learning associated with reputation agencies on Twitter. Content (games, movies, and books) is provided to the reputation agency which then publishes its review as a tweet. The sentiment of the tweet and number of followers (reputation) constitute the publicly available information that drives the social learning process. The Twitter network responds to the sentiment of the tweet by retweeting. Data incest information structures arise since if the sentiment of the tweet published by the reputation agency improves its reputation, then this reinforces the reputation agency's belief that it adequately reviewed the content.

The framework, which is illustrated in Fig. 7, involves utility maximization and dynamics of the public belief, which we will show (based on Twitter data) evolves according to an autoregressive process. Specifically, the goal is to investigate how the number of *followers* and *polarity*<sup>14</sup> affects the time before

<sup>14</sup>Here, polarity [41] is a real valued variable in the interval  $[-1, 1]$  and depends on the sentiment of the tweet; i.e., whether the tweet expresses a positive/negative/neutral statement.

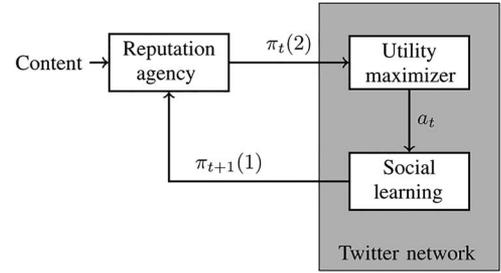


Fig. 7. Schematic of dynamics of the reputation agency and the Twitter network. The reputation agency receives *content* (books, games, and movies) which it reviews and publishes tweets at epochs  $t = 1, 2, \dots$ , with sentiment  $\pi_t(2)$ . The reputation agency index  $i$  has been omitted for clarity. The public belief  $\pi_t$  and response  $a_t$  are defined in Section V-B.

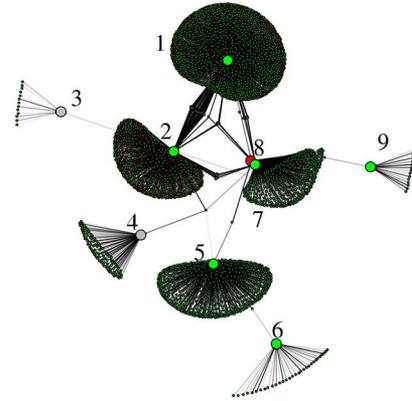


Fig. 8. Snapshot of the estimated retweet network obtained by tracking real-time tweets of reputation agencies. The labels 1, 2,  $\dots$ , 9 correspond to the Twitter accounts @IGN, @gamespot, @AmznMovieRevws, @creativereview, @HarvardBiz, @techreview, @pcgamer, @RottenTomatoes, and @LARewiewofBooks. The data are collected over a period of 24 h starting from November 17, 2014 at 9:00 P.M. The reputation agency nodes are denoted by large circles and retweeting nodes (followers) by small circles. The sentiment of the tweet published by the reputation agency is denoted by color: red is negative, green is positive, and gray is neutral. The time of the retweet is indicated by the shade intensity of the edges of the graph: the lighter the shade of an edge, the later the retweet was posted.

a retweet occurs and the total number of associated retweets. Apart from its relevance to social learning, the information provided by this analysis can be used in social media marketing strategies to improve a brand and for brand awareness. As discussed in [42], Twitter provides a significant amount of agent-generated data, which can be analyzed to provide novel personal advertising to agents.

1) *Twitter Datasets:* We consider nine well-known online reputation agencies: @IGN, @gamespot, @Amzn-MovieRevws, @creativereview, @HarvardBiz, @techreview, @pcgamer, @RottenTomatoes, @LARewiewofBooks. Fig. 8 provides the social retweet network, which includes the Twitter accounts of these nine reputation agencies. The Twitter data were collected on November 17, 2014 at 9:00 P.M. for a duration of 24 h. The data were obtained using the *Twitter Streaming API*<sup>15</sup> and a custom python script. The sentiment

<sup>15</sup>[Online]. Available: <https://dev.twitter.com/streaming/overview>.

of the tweets and retweets is computed using *TextBlob*.<sup>16</sup> The social network contains 10 656 nodes with 11 054 edges. As illustrated in Fig. 8, numerous nodes retweet based on tweets posted by these nine agencies. The edge intensity in Fig. 8 can be used to gain intuition on the dynamics of retweets. For the nodes with high *in degree* such as @IGN and @HarvardBiz, the retweets typically occur in a short period of time on the order of 1–12 h. This behavior has been observed in popularity dynamics of papers and youtube videos [43] and is associated with a decrease in the ability to attract new attention after ageing.

In the following analysis, the aim is to determine if the Twitter followers of an online reputation agency exhibit utility maximization behavior in response to the tweet published by the reputation agency. The index  $t = 1, 2, \dots$ , denotes epochs at which the reputation agency publishes its tweets.<sup>17</sup> To apply Afriat's theorem (25), the public belief is defined by

$$\pi_t^i = [\#\text{followers}, \text{neutrality}]$$

for each reputation agency  $i$ . The  $\#\text{followers}$  is the number of followers of a tweet published by the online reputation agency. The *neutrality* of the tweet published by the reputation agency is computed as  $1/|\text{polarity}|$  where the *polarity* of a tweet is computed using *TextBlob*. The associated response taken by Twitter users that retweet in the network is given by

$$a_t^i = [\Delta t, \#\text{retweets}].$$

$\Delta t$  denotes the time between the tweet (published by the agency) and the first retweet (of a follower).  $\#\text{retweets}$  denotes the total number of retweets generated by followers prior to the next tweet from the reputation agency  $i$ . Next, we need to justify the linear budget in (24) to use Afriat's theorem. It is clear that as the number of followers of a reputation agency increases, the number of retweets will increase. Consequently, one expects a decrease in time between the first retweet  $a_t(1)$  as the number of followers  $\pi_t(1)$  increases. The results in [44] suggest that the higher the polarity of a tweets published by the reputation agency, the larger the number of retweets. So, we expect that as the *neutrality* (i.e., the lower the polarity) of the tweet increases the resulting number of retweets  $a_t(2)$  will decrease. So, it is reasonable to assume the existence of a social impact budget  $I_t$  for the utility maximization test (25) which satisfies  $I_t = \pi_t' a_t$ . We construct the datasets  $\mathcal{D}_i$  for each reputation agency  $i \in \{1, 2, \dots, 9\}$  from the Twitter data collected on November 17, 2014 at 9:00 P.M. for a duration of 24 h. The dataset  $\mathcal{D}_i = \{(\pi_t^i, a_t^i) : t \in \{1, 2, \dots, T^i\}\}$  was constructed using the public belief  $\pi_t^i$ , response  $a_t^i$ , and total number of tweets  $T^i$  for the reputation agency  $i$ . Note that  $t \in \{1, \dots, T^i\}$  denotes the tweets published by the reputation agency.

2) *Results*: We found that each of the Twitter datasets  $\mathcal{D}_i$  satisfy the utility maximization test (25). Using (26), from Afriat's Theorem, the associated utility function for the

<sup>16</sup>TextBlob—Python-based Text Processing Tool Suit [Online]. Available: <https://dev.twitter.com/streaming/overview>. <http://textblob.readthedocs.org/en/dev/>.

<sup>17</sup>The average time interval between tweets for @IGN, @gamespot, and @HarvardBiz are, respectively, 31, 33, and 34 min.

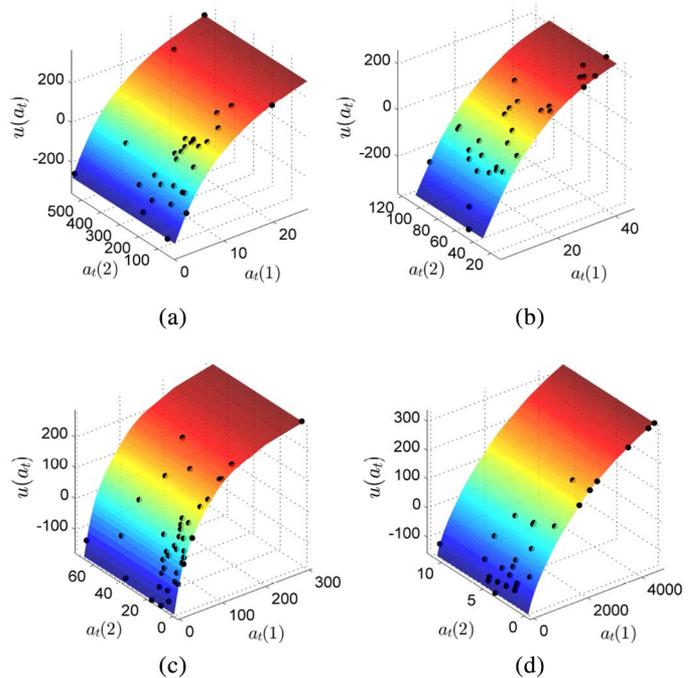


Fig. 9. Estimated utility function  $u(a_t)$  using the Twitter datasets  $\mathcal{D}_i$  for  $i \in \{1, 2, 5, 9\}$  defined in Section V-B constructed using (26) from Afriat's Theorem. Note that  $a_t(1)$  is the number of retweets, and  $a_t(2)$  has units of seconds. (a) @IGN twitter account. (b) @gamespot twitter account. (c) @HarvardBiz twitter account. (d) @LARReviewofBooks twitter account.

reputation agencies @IGN, @gamespot, @RottenTomatoes, and @LARReviewBooks is provided in Fig. 9(a)–(d). The utility function for the other five agencies is omitted as only minor differences are present compared with the utility functions provided in Fig. 9. Several key observations can be made from the results.

- Given that  $\mathcal{D}_i$  satisfies (25), this suggests that the number of followers  $\pi_t^i(1)$  and neutrality  $\pi_t^i(2)$  of the tweet contributes to the retweet dynamics [i.e., the delay before the first tweet  $a_t^i(1)$  and the total number of retweets  $a_t^i(2)$ ].
- The utility functions provided in Fig. 9 suggest that Twitter users prefer to increase the delay of retweeting  $a_t^i(1)$  compared with increasing the total number of retweets  $a_t^i(2)$ . The results also suggest that as the delay before the first retweet increases the associated number of retweets decreases. This effect is pronounced in Fig. 9(d) where the first and only retweet occurs approximately 2000 s after the original tweet.
- The revealed preferences of the reputation agencies, represented by the utility functions in Fig. 9(a)–(d), are not identical.

Observation a) is straightforward as a change in the number of followers will affect the time of a retweet. Additionally, as suggested in [44], as neutrality of the tweet increases the associated total number of retweets is expected to decrease. Observation b) illustrates an interesting characteristic of how users retweet to reputation agencies, it suggests that Twitter users prefer to increase the time prior to retweeting compared to increasing the total number of retweets. This result is caused by social features of users, which include the content of the tweet, and the

external context during which the tweet is posted [45]. In [45], over 250 million tweets are collected and analyzed and it was found that a high number of followers do not necessarily lead to shorter retweet times. A possible mechanism for this effect is that tweets that are exposed to a larger number of Twitter users then an individual user is less likely to retweet—this effect is known as *diffusion of responsibility*. For large retweet times  $a_t^i(1)$ , we observe that the total number of retweets is significantly reduced compared to short retweet times, refer to Fig. 9(a) and (d). This result has been observed in popularity dynamics [43] and is associated with an ageing effect—i.e., the interest of the tweet decreases with time. Observation c) is expected as different reputation agencies are expected to have different users retweeting. To quantify this result for the constructed ordinal (i.e., identical for any monotonic transformation) utility functions, the comparison of preferences is achieved using the *marginal rate of substitution* defined by

$$\text{MRS}_{12} = \frac{\partial u / \partial a^i(1)}{\partial u / \partial a^i(2)}. \quad (27)$$

In (27),  $\text{MRS}_{12}$  defines the amount of  $a^i(2)$  and the Twitter users are willing to give up for one additional unit of  $a^i(1)$ . In Fig. 9(a)–(d), it is clear that  $\text{MRS}_{12} > 1$  suggesting users prefer  $a^i(1)$  compared with  $a^i(2)$ . Additionally, notice that the  $\text{MRS}_{12}$  for each of the reputation agencies in Fig. 9(a)–(d) illustrates that the associated behavior of each is characteristically different.

3) *Social Learning*: To interpret the above results in terms of social learning, we next show that the response (action)  $a_t$  and public belief  $\pi_t$  at epoch  $t$  determine the public belief  $\pi_{t+1}$  at epoch  $t + 1$ . We found that the following autoregressive time series model

$$\pi_{t+1}(1) = \pi_t(1) + b a_t(2) + \epsilon_t \quad (28)$$

driven by a zero mean noise process  $\{\epsilon_t\}$  yields an excellent fit, where the parameter  $b$  is estimated from the Twitter data using least squares. Note that  $\pi_{t+1}(1) - \pi_t(1)$  in (28) models the total number of new followers resulting from the the total number of retweets  $a_t(2)$ . To test the accuracy of (28), we selected the reputation agencies @IGN, @gamespot, and @HarvardBiz. The experimentally measured and numerically predicted results are displayed in Fig. 10. As seen (28), accurately predicts the public opinion based on the response of the Twitter users. The *mean absolute percentage error* using the AR model for the reputation agencies @IGN, @gamespot, and @HarvardBiz are, respectively, 0.49%, 0.67%, and 0.41%.

## VI. NUMERICAL EXAMPLES

This section illustrates the incest removal algorithm of social learning Protocol 1 for two different types of social networks. Then, incest removal in expectation polling is illustrated for these two networks.

### A. Social Learning and Incest in Corporate Network

Consider the network of Fig. 11 which depicts a corporate social network. Nodes 1 and 10 denote the same senior

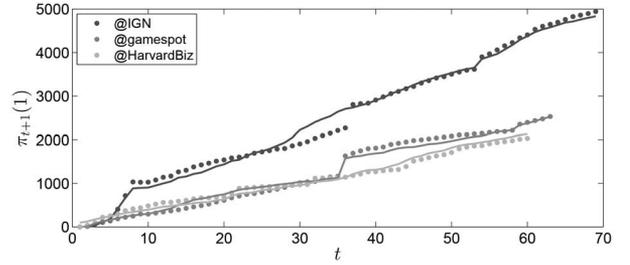


Fig. 10. Accuracy of the AR model (28) for the social learning dynamics of public belief  $\pi_{t+1}(1)$  of reputation agencies: @IGN, @gamespot, and @HarvardBiz. The dots represent the actual number of Twitter followers, while the solid line indicates the predicted number of followers using the AR model. The initial number of followers at  $t = 0$  is set to zero in the plot. The estimated parameter values of  $b$  in (28) for agencies @IGN, @gamespot, and @HarvardBiz are 0.4358, 0.8132, and 0.3825, respectively. The average time interval between tweet epochs  $t$  for @IGN, @gamespot, and @HarvardBiz are 31, 33, and 34 min, respectively.

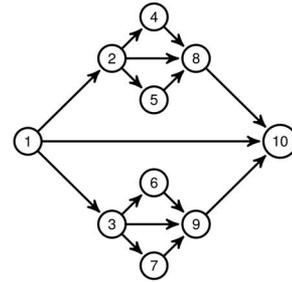


Fig. 11. Social learning in corporate network comprising higher level manager, mid-level manager, and workers. Data incest arises at nodes 8, 9, and 10.

level manager. Nodes 2 and 8 denote a mid-level manager; and nodes 3 and 9 denote another mid-level manager. The two mid-level managers attend a presentation (and therefore opinion) by a senior manager to determine (estimate) a specific parameter  $x$  about the company. Each mid-level manager then makes recommendations (decisions) and convey these to two of their workers. These workers eventually report back to their mid-level managers who in turn report back to the senior manager. The edge (1, 10) indicates that the senior manager recalls her decision at node 1 when making her decision at node 10. Similarly for edges (2, 8) and (3, 9).

We ran 1000 independent simulations with the following parameters in the social learning Protocol 1:  $\mathbb{X} = \{1, 2, \dots, X\}$ ,  $\mathcal{A} = \{1, 2, \dots, A\}$ ,  $\mathbb{Y} = \{1, 2, \dots, Y\}$ ,  $X = 10$ ,  $A = 10$ ,  $Y = 20$ , true state  $x$  uniform on  $\mathbb{X}$ , prior  $\pi_0$  uniform on  $\mathbb{X}$ , observation probabilities  $B_{x,y} \propto \exp -\frac{(y-x)^2}{2}$ , and costs  $c(i, a) = |\frac{A}{X}x - a|$ .

It is clear from Fig. 11 that incest arises at nodes 8, 9, and 10 due to multiple information paths. Table I(a) displays the MSE of the state estimates computed at these nodes. The incestious estimates were computed using naive data fusion (10). The incest free estimates were computed using the incest removal algorithm of Theorem 2.3. It is verified that the achievability condition of Theorem 2.4 holds; hence, incest can be removed completely given the information from single hop nodes (immediate friends). Table I shows that incest removal results in

TABLE I  
EFFECT OF INCEST REMOVAL IN SOCIAL LEARNING. THE MSE IN THE STATE ESTIMATE WERE OBTAINED BY AVERAGING OVER 1000 INDEPENDENT SIMULATIONS. (A) CORPORATE NETWORK OF FIG. 11 AND (B) MESH NETWORK OF FIG. 12

Node	8	9	10
MSE with incest	0.3666	0.3520	0.3119
MSE with incest removal	0.2782	0.2652	0.1376

(a)

Node	5	6	7	8	9
MSE with incest	0.3246	0.3420	0.3312	0.3149	0.3134
MSE with incest removal	0.2799	0.2542	0.2404	0.2267	0.2200

(b)

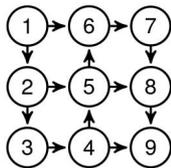


Fig. 12. Social learning in mesh network. Data incest occurs at nodes 5, 6, 8, and 9.

substantially more accurate estimates, particularly at node 10 which is the most informed node (since all nodes eventually report to node 10).

### B. Social Learning and Incest in Mesh Network

Consider the mesh network depicted in Fig. 12 with parameters identical to the previous example. Incest arises at nodes 5, 6, 8, and 9. Also, incest propagates to node 6 from node 5. It is verified that the condition of Theorem 2.4 holds for the mesh network and so incest can be completely removed. Table I(b) displays the MSE of naive data fusion and incest removal and shows that substantial improvements occur in the state estimate with incest removal.

### C. Expectation Polling Using Incestious Belief

We illustrate the results of Section IV for incest removal in expectation polling. Consider the network of Fig. 11 with the edge (1, 10) omitted. Node 10 denotes the pollster. Assume the pollster cannot sample node 1 so that exact data incest removal is impossible. [It can be verified that the necessary and sufficient condition (19) of Theorem 2.4 does not hold if node 1 is omitted.] Given the incestious beliefs (expectations)  $\pi_m$  from the sampled voters  $m \in \mathcal{R}$ , we use the estimator of Theorem 4.1 to compute the conditional mean estimate  $\hat{x} = E\{x|\pi_m, m \in \mathcal{R}\}$  of the leading candidate.

For the simulation, we chose  $\mathbb{X} = \{1, 2\}$  (two candidates), prior  $[0.4, 0.6]'$ ,  $\mathbb{Y} = \{1, 2\}$ ,  $B = \begin{bmatrix} 0.8 & 0.2 \\ 0.2 & 0.8 \end{bmatrix}$ . Table II(a) displays the MSE of  $\hat{x}$  for different recruited sample sets  $\mathcal{R}$ . The table also displays the MSE when all nodes are sampled in which case optimal incest removal is achieved. These are compared with the naive incestious estimator using  $\pi_{10}$  computed via Protocol 2.

TABLE II  
EFFECT OF INCEST REMOVAL IN EXPECTATION POLLING USING THE ESTIMATOR (22). THE MSEs IN THE STATE ESTIMATE WERE OBTAINED BY AVERAGING OVER 1000 INDEPENDENT SIMULATIONS. (A) CORPORATE NETWORK OF FIG. 11 WITH EDGE (1, 10) OMITTED AND (B) MESH NETWORK OF FIG. 12

Sampled voters $\mathcal{R}$	MSE
$\{8, 9, 10\}$	0.0250
$\{4, \dots, 10\}$	0.0230
$\{2, \dots, 10\}$	0.0161
$\{1, \dots, 10\}$	0.0118
Naive (incest)	0.0513

(a)

Sampled voters $\mathcal{R}$	MSE
$\{4, 8, 9\}$	0.0314
$\{2, 4, 8, 9\}$	0.0178
$\{1, \dots, 9\}$	0.0170
Naive (incest)	0.0431

(b)

Finally, consider expectation polling of voters in the mesh network of Fig. 12. Table II(b) displays the MSEs of the conditional mean estimates for different sampled nodes and the naive incestious estimate.

## VII. CONCLUSION AND EXTENSIONS

This paper considered data incest in reputation and expectation polling systems. In reputation systems, data incest arose in a multiagent social learning model. A necessary and sufficient condition on the adjacency matrix of the directed graph was given for exact incest removal at each stage. For expectation polling systems, it was shown that even if incest propagates in the network, the posterior of the state can be estimated based on the incestious beliefs. Finally, by analyzing Twitter datasets associated with several online reputation agencies, we used Afriat's theorem of revealed preferences to show utility maximization behavior and social learning. In future work, it is worthwhile extending the framework in this paper to active sensing and sequential decision-making. For the case of classical social learning, Spiro *et al.* [46] deals with sequential quickest detection and stochastic control. Extending this to data incest information patterns is challenging and nontrivial.

### ACKNOWLEDGMENT

The experimental data in Section III were obtained from Dr. G. Solman and Prof. A. Kingstone of the Department of Psychology, University of British Columbia. The experiments were conducted by them in the fall semester of 2013. The data analysis of Section III and Fig. 4 was prepared by Mr. M. Hamdi. A complete description of our results in the psychology experiment is in the preprint [37].

### APPENDIX

#### A. Illustrative Example

We provide here an example of the data incest problem setup of Section II. Consider  $S = 2$  two agents with information flow graph for three time points  $k = 1, 2, 3$  depicted in Fig. 13 characterized by the family of DAGs  $\{G_1, \dots, G_7\}$ . The

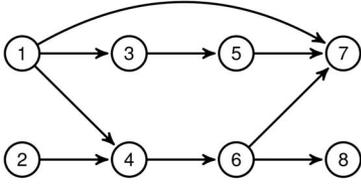


Fig. 13. Example of information flow network with  $S = 2$  two agents, namely  $s \in \{1, 2\}$  and time points  $k = 1, 2, 3, 4$ . Circles represent the nodes indexed by  $n = s + S(k - 1)$  in the social network and each edge depicts a communication link between two nodes.

adjacency matrices  $A_1, \dots, A_7$  are constructed as follows:  $A_n$  is the upper left  $n \times n$  submatrix of  $A_{n+1}$  and

$$A_7 = \begin{bmatrix} 0 & 0 & 1 & 1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}.$$

Let us explain these matrices. Since nodes 1 and 2 do not communicate, clearly  $A_1$  and  $A_2$  are zero matrices. Nodes 1 and 3 communicate, hence  $A_3$  has a single one, etc. Note that if nodes 1, 3, 4, and 7 are assumed to be the same individual, then at node 7, the individual remembers what happened at nodes 5 and 1, but not node 3. This models the case where the individual has selective memory and remembers certain highlights; we discuss this further in Section II-F. From (8) and (9)

$$\mathcal{H}_7 = \{1, 5, 6\}, \quad \mathcal{F}_7 = \{1, 2, 3, 4, 5, 6\}$$

where  $\mathcal{H}_7$  denotes all one hop links to node 7, whereas  $\mathcal{F}_7$  denotes all multihop links to node 7.

Using (7), the transitive closure matrices  $T_1, \dots, T_7$  are given by  $T_n$  is the upper left  $n \times n$  submatrix of

$$T_7 = \begin{bmatrix} 1 & 0 & 1 & 1 & 1 & 1 & 1 \\ 0 & 1 & 0 & 1 & 0 & 1 & 1 \\ 0 & 0 & 1 & 0 & 1 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}.$$

Note that  $T_n(i, j)$  is nonzero only for  $i \geq j$  due to causality since information sent by an agent can only arrive at another social group at a later time instant. The weight vectors are then obtained from (18) as

$$\begin{aligned} \mathbf{w}_2 &= [0], & \mathbf{w}_3 &= [1 \ 0]', & \mathbf{w}_4 &= [1 \ 1 \ 0]' \\ \mathbf{w}_5 &= [0 \ 0 \ 1 \ 0]', & \mathbf{w}_6 &= [0 \ 0 \ 0 \ 1 \ 0]' \\ \mathbf{w}_7 &= [-1 \ 0 \ 0 \ 0 \ 1 \ 1]'. \end{aligned}$$

$\mathbf{w}_2$  means that node 2 does not use the estimate from node 1. This formula is consistent with the constraints on information flow because the estimate from node 1 is not available to node 2 (see Fig. 13).  $\mathbf{w}_3$  means that node 3 uses estimates from

nodes 1;  $\mathbf{w}_4$  means that node 4 uses estimates only from node 1 and node 2. As shown in Fig. 13, the misinformation propagation occurs at node 7 since there are multiple paths from nodes 1 to 7. The vector  $\mathbf{w}_7$  says that node 7 adds estimates from nodes 5 and 6 and removes estimates from node 1 to avoid triple counting of these estimates already integrated into estimates from nodes 3 and 4. Using the algorithm (18), incest is completely prevented in this example.

Here is an example in which exact incest removal is impossible. Consider the information flow graph of Fig. 13 but with the edge between nodes 1 and 7 deleted. Then,  $A_7(1, 7) = 0$  while  $w_7(1) \neq 0$ , and therefore the condition (19) does not hold. Hence, exact incest removal is not possible for this case. In Section IV, we compute the Bayesian estimate of the underlying state when incest cannot be removed.

#### Proof of Theorem 2.2

The proof uses MLR stochastic dominance (defined in footnote 8) and the following single crossing condition.

*Definition 1 (Single Crossing [35]):*  $g: \mathbb{Y} \times \mathbb{A} \rightarrow \mathbb{R}$  satisfies a single crossing condition in  $(y, a)$  if  $g(y, a) - g(y, \bar{a}) \geq 0$  implies  $g(\bar{y}, a) - g(\bar{y}, \bar{a}) \geq 0$  for  $\bar{a} > a$  and  $\bar{y} > y$ . Then,  $a^*(y) = \operatorname{argmin}_a g(y, a)$  is increasing in  $y$ . ■

By (A1), it is verified that the Bayesian update satisfies

$$\frac{B_y \pi}{\mathbf{1}' B_y \pi} \leq_r \frac{B_{y+1} \pi}{\mathbf{1}' B_{y+1} \pi}$$

where  $\leq_r$  is the MLR stochastic order. (Indeed, the MLR order is closed under conditional expectation and this is the reason why it is widely used in Bayesian analysis.) By submodular assumption (A2),  $c_{a+1} - c_a$  is a vector with decreasing elements. Therefore,

$$(c_{a+1} - c_a)' \frac{B_y \pi}{\mathbf{1}' B_y \pi} \geq (c_{a+1} - c_a)' \frac{B_{y+1} \pi}{\mathbf{1}' B_{y+1} \pi}.$$

Since the denominator is nonnegative, it follows that  $(c_{a+1} - c_a)' B_{y+1} \pi \geq 0 \implies (c_{a+1} - c_a)' B_y \pi \geq 0$ . This implies that  $c'_a B_y \pi$  satisfies a single crossing condition in  $(y, a)$ . Therefore,  $a_n(\pi, y) = \operatorname{argmin}_a c'_a B_y \pi$  is increasing in  $y$  for any belief  $\pi$ .

#### Proof of Theorem 2.3

The local estimate at node  $n$  is given by (18), namely

$$l_{n-}(i) = w'_n l_{1:n-1}(i). \quad (29)$$

Define  $\bar{R}_{ia}^\pi = \log P(a|x = i, \pi)$  and the  $n - 1$ -dimensional vector  $\bar{R}_{1:n-1}^\pi(i) = [\bar{R}_{i,a_1}^{\pi_1} \ \bar{R}_{i,a_{n-1}}^{\pi_{n-1}}]$ . From the structure of transitive closure matrix  $T_n$

$$l_{1:n-1}(i) = T'_{n-1} \bar{R}_{1:n-1}^\pi(i), \quad l_{n-}(i) = t'_n \bar{R}_{1:n-1}^\pi(i). \quad (30)$$

Substituting the first equation in (30) into (29) yields

$$l_{n-}(i) = w'_n T'_{n-1} \bar{R}_{1:n-1}^\pi(i).$$

Equating this with the second equation in (30) yields  $w_n = T_{n-1}^{-1} t_n$ . (By Lemma 1,  $T_{n-1}$  is invertible.)

### Proof of Theorem 4.1

Given  $n$  nodes, it is clear from Bayes formula and the structure of the adjacency matrix  $A$  of the DAG that

$$l_{1:n}(i) = o_{1:n} + e_1 l_0(i) + A' l_{1:n}(i).$$

Since,  $I - A$  is invertible by construction

$$l_{1:n}(i) = (I - A')^{-1} o_{1:n} + (I - A')^{-1} e_1 l_0(i).$$

Then,  $l_{\mathcal{R}}(i) = [e_{\mathcal{R}_1}, e_{\mathcal{R}_2}, \dots, e_{\mathcal{R}_L}]' l_{1:n}(i)$  satisfies (23). Finally,  $P(x|\pi_{\mathcal{R}}) \propto \sum_{Y_{1:n} \in \mathcal{Y}} P(Y_{1:n}|x) \pi_0(x)$ . Here,  $\mathcal{Y}$  is the set of  $n$  dim. vectors satisfying (23).

### REFERENCES

- [1] V. Krishnamurthy and H. V. Poor, "A tutorial on interactive sensing in social networks," *IEEE Trans. Comput. Social Syst.*, vol. 1, no. 1, pp. 3–21, Mar. 2014.
- [2] B. Ifrach, C. Maglaras, and M. Scarsini, "Monopoly pricing in the presence of social learning," NET Institute Working Paper No. 12-01, 2011 [Online]. Available: <http://EconPapers.repec.org/RePEc:net:wpaper:1111>
- [3] M. Luca, "Reviews, Reputation, and Revenue: The Case of Yelp.com," Harvard Business School, Boston, MA, USA, Tech. Rep. 12-016, Sep. 2011.
- [4] D. Rothschild and J. Wolfers, "Forecasting elections: Voter intentions versus expectations," Working Paper, University of Pennsylvania, 2010.
- [5] S. Afriat, "The construction of utility functions from expenditure data," *Int. Econ. Rev.*, vol. 8, no. 1, pp. 67–77, 1967.
- [6] H. Varian, "The nonparametric approach to demand analysis," *Econometrica*, vol. 50, no. 1, pp. 945–973, 1982.
- [7] M. Granovetter, "Threshold models of collective behavior," *Amer. J. Sociol.*, vol. 83, no. 6, pp. 1420–1443, May 1978.
- [8] C. Chamley, *Rational Herds: Economic Models of Social Learning*. Cambridge, U.K.: Cambridge Univ. Press, 2004.
- [9] S. Bikchandani, D. Hirshleifer, and I. Welch, "A theory of fads, fashion, custom, and cultural change as information cascades," *J. Polit. Econ.*, vol. 100, no. 5, pp. 992–1026, Oct. 1992.
- [10] D. Acemoglu and A. Ozdaglar, "Opinion dynamics and learning in social networks," *Dyn. Games Appl.*, vol. 1, no. 1, pp. 3–49, 2011.
- [11] Y. Kanoria and O. Tamuz, "Tractable Bayesian social learning on trees," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Jul. 2012, pp. 2721–2725.
- [12] T. Cover and M. Hellman, "The two-armed-bandit problem with time-invariant finite memory," *IEEE Trans. Inf. Theory*, vol. IT-16, no. 2, pp. 185–195, Mar. 1970.
- [13] C. Chamley, A. Scaglione, and L. Li, "Models for the diffusion of beliefs in social networks: An overview," *IEEE Signal Process. Mag.*, vol. 30, no. 3, pp. 16–29, May 2013.
- [14] C. Dellarocas, "The digitization of word of mouth: Promise and challenges of online feedback mechanisms," *Manage. Sci.*, vol. 49, no. 10, pp. 1407–1424, 2003.
- [15] A. Jøsang, R. Ismail, and C. Boyd, "A survey of trust and reputation systems for online service provision," *Decis. Support Syst.*, vol. 43, no. 2, pp. 618–644, 2007.
- [16] E. Lim, V. Nguyen, N. Jindal, B. Liu, and H. Lauw, "Detecting product review spammers using rating behaviors," in *Proc. 19th ACM Int. Conf. Inf. Knowl. Manage.*, 2010, pp. 939–948.
- [17] V. Krishnamurthy and M. Hamdi, "MIS-information removal in social networks: Dynamic constrained estimation on directed acyclic graphs," *IEEE J. Sel. Topics Signal Process.*, vol. 7, no. 2, pp. 333–346, May 2013.
- [18] J. Xie et al., "Social consensus through the influence of committed minorities," *Phys. Rev. E*, vol. 84, no. 1, 2011, Art. ID. 011130.
- [19] R. Bond et al., "A 61-million-person experiment in social influence and political mobilization," *Nature*, vol. 489, pp. 295–298, Sep. 2012.
- [20] A. Dasgupta, R. Kumar, and D. Sivakumar, "Social sampling," in *Proc. 18th ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2012, pp. 235–243.
- [21] E. Royer and C. Toh, "A review of current routing protocols for ad hoc mobile wireless networks," *IEEE Personal Commun.*, vol. 6, no. 2, pp. 46–55, 1999.
- [22] J. Yedidia, W. Freeman, and Y. Weiss, "Constructing free-energy approximations and generalized belief propagation algorithms," *IEEE Trans. Inf. Theory*, vol. 51, no. 7, pp. 2282–2312, Jul. 2005.
- [23] R. J. Aumann, "Agreeing to disagree," *Ann. Statist.*, vol. 4, no. 6, pp. 1236–1239, Nov. 1976.
- [24] J. Geanakoplos and H. Polemarchakis, "We can't disagree forever," *J. Econ. Theory*, vol. 28, no. 1, pp. 192–200, 1982.
- [25] S. Afriat, *Logic of Choice and Economic Theory*. Oxford, U.K.: Clarendon, 1987.
- [26] H. Varian, "Price discrimination and social welfare," *Amer. Econ. Rev.*, vol. 75, no. 4, pp. 870–875, 1985.
- [27] H. Varian, "Revealed preference and its applications," *Econ. J.*, vol. 122, no. 560, pp. 332–338, 2012.
- [28] H. Varian, "Online ad auctions," *Amer. Econ. Rev.*, vol. 99, no. 2, pp. 430–434, 2009.
- [29] M. Trusov, A. V. Bodapati, and R. E. Bucklin, "Determining influential users in internet social networks," *J. Market. Res.*, vol. XLVII, pp. 643–658, Aug. 2010.
- [30] S. Karlin, *Total Positivity*. Stanford, CA, USA: Stanford Univ., 1968, vol. 1.
- [31] S. Karlin and Y. Rinott, "Classes of orderings of measures and related correlation inequalities. I. Multivariate totally positive distributions," *J. Multivariate Anal.*, vol. 10, no. 4, pp. 467–498, Dec. 1980.
- [32] P. Milgrom, "Good news and bad news: Representation theorems and applications," *Bell J. Econ.*, vol. 12, no. 2, pp. 380–391, 1981.
- [33] D. Topkis, *Supermodularity and Complementarity*. Princeton, NJ, USA: Princeton Univ. Press, 1998.
- [34] P. Milgrom and C. Shannon, "Monotone comparative statistics," *Econometrica*, vol. 62, no. 1, pp. 157–180, 1992.
- [35] R. Amir, "Supermodularity and complementarity in economics: An elementary survey," *Southern Econ. J.*, vol. 71, no. 3, pp. 636–660, 2005.
- [36] K. Scarfone, W. Jansen, and M. Tracy, *Guide to General Server Security*, NIST Special Publication, U.S. Department of Commerce, 2008, vol. 800, p. 123.
- [37] M. Hamdi, G. Solman, A. Kingstone, and V. Krishnamurthy, "Social learning in a human society: An experimental study," Arxiv preprint arXiv:1408.5378, 2014.
- [38] W. Diewert, "Afriat and revealed preference theory," *Rev Econ. Stud.*, vol. 40, pp. 419–425, 1973.
- [39] A. Fostel, H. Scarf, and M. Todd, "Two new proofs of Afriats theorem," *Econ. Theory*, vol. 24, no. 1, pp. 211–219, 2004.
- [40] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, U.K.: Cambridge Univ. Press, 2004.
- [41] A. Montejo-Rez, E. Martinez-Cmara, M. Martn-Valdivia, and L. Urea-Lpez, "Ranked wordnet graph for sentiment polarity classification in twitter," *Comput. Speech Lang.*, vol. 28, no. 1, pp. 93–107, 2014.
- [42] F. Giummolè, S. Orlando, and G. Tolomei, "A study on microblog and search engine user behaviors: How twitter trending topics help predict google hot queries," *HUMAN*, vol. 2, no. 3, p. 195, 2013.
- [43] H. Shen, D. Wang, C. Song, and A. Barabási, "Modeling and predicting popularity dynamics via reinforced poisson processes," Arxiv preprint arXiv:1401.0778, 2014.
- [44] S. Stieglitz and L. Dang-Xuan, "Emotions and information diffusion in social media-sentiment of microblogs and sharing behavior," *J. Manage. Inf. Syst.*, vol. 29, no. 4, pp. 217–248, 2013.
- [45] E. Spiro, C. Irvine, C. DuBois, and C. Butts, "Waiting for a retweet: modeling waiting times in information propagation," presented at the NIPS Workshop Social Netw. Social Media Conf., Lake Tahoe, NV, USA, Dec. 2012, vol. 12.
- [46] V. Krishnamurthy, "Quickest detection POMDPs with social learning: Interaction of local and global decision makers," *IEEE Trans. Inf. Theory*, vol. 58, no. 8, pp. 5563–5587, Aug. 2012.

**Vikram Krishnamurthy** (S'90–M91–SM'99–F'05) received the Ph.D. degree in electrical engineering from the Australian National University, Canberra, Australia, in 1992.

He is a Professor and Canada Research Chair with the Department of Electrical Engineering, University of British Columbia, Vancouver, BC, Canada. His research interests include statistical signal processing, computational game theory, and stochastic control in social networks.

Prof. Krishnamurthy served as a Distinguished Lecturer for the IEEE Signal Processing Society and Editor in Chief of THE IEEE JOURNAL SELECTED TOPICS IN SIGNAL PROCESSING. He received an honorary doctorate from KTH (Royal Institute of Technology), Sweden, in 2013.

**William Hoiles** (M'05) received the M.A.Sc. degree in electrical engineering from the Department of Engineering Science, Simon Fraser University, Vancouver, BC, Canada, in 2012. He is currently pursuing the Ph.D. degree at the Department of Electrical and Computer Engineering, University of British Columbia, Vancouver, Canada.

His research interests include social sensors and the bioelectronic interface.