A Tutorial on Interactive Sensing in Social Networks

Vikram Krishnamurthy, *Fellow, IEEE* and H. Vincent Poor, *Fellow, IEEE*

(Invited Paper)

***Abstract*—This paper considers models and algorithms for interactive sensing in social networks in which individuals act as sensors and the information exchange between individuals is exploited to optimize sensing. Social learning is used to model the interaction between individuals that aim to estimate an underlying state of nature. In this context, the following questions are addressed: how can self-interested agents that interact via social learning achieve a tradeoff between individual privacy and reputation of the social group? How can protocols be designed to prevent data incest in online reputation blogs where individuals make recommendations? How can sensing by individuals that interact with each other be used by a global decision maker to detect changes in the underlying state of nature? When individual agents possess limited sensing, computation, and communication capabilities, can a network of agents achieve sophisticated global behavior? Social and game-theoretic learning are natural settings for addressing these questions. This article presents an overview, insights, and discussion of social learning models in the context of data incest propagation, change detection, and coordination of decision-making.**

***Index Terms*—Coordination, correlated equilibria, data incest, game-theoretic learning, information diffusion, reputation systems, social learning, social sampling.**

I. Introduction and Motivation

THE proliferation of social media such as real-time microblogging services (Twitter[1]), online reputation, and rating systems (Yelp) together with app-enabled smartphones, facilitate real-time sensing of social activities, social patterns, and behavior.

Social sensing, also called participatory sensing [1]–[5], is defined as a process by which physical sensors present in mobile devices such as GPS are used to infer social relationships and human activities. In this paper, we work at a higher level of abstraction. We use the term *social sensor* or *human-based sensor* to denote an agent that provides information about its environment (state of nature) on a social network after interaction with other agents. Examples of such social sensors include Twitter posts, Facebook status updates, and ratings on online reputation systems such as Yelp and Tripadvisor. Such social sensors go beyond physical sensors for social sensing. For example [6], user opinions/ratings (such as the quality of a restaurant) are available on Tripadvisor but are difficult to measure via physical sensors. Similarly, future situations revealed by the Facebook status of a user are impossible to predict using physical sensors.

Statistical inference using social sensors is relevant in a variety of applications including localizing special events for targeted advertising [7], [8], marketing [9], localization of natural disasters [10], and predicting sentiment of investors in financial markets [11], [12]. It is demonstrated in [13] that models built from the rate of tweets for particular products can outperform market-based predictors. However, social sensors present unique challenges from a statistical estimation viewpoint. First, social sensors interact with and influence other social sensors. For example, ratings posted on online reputation systems strongly influence the behavior of individuals.[2] Such interacting sensing can result in nonstandard information patterns due to correlations introduced by the structure of the underlying social network. Second, due to privacy concerns and time constraints, social sensors typically do not reveal raw observations of the underlying state of nature. Instead, they reveal their decisions (ratings, recommendations, and votes), which can be viewed as a low-resolution (quantized) function of their raw measurements and interactions with other social sensors.

As it is apparent from the above discussion, there is a strong motivation to construct mathematical models that capture the dynamics of interactive sensing involving social sensors. Such models facilitate understanding the dynamics of information flow in social networks and, therefore, the design of algorithms that can exploit these dynamics to estimate the underlying state of nature.

In this paper, *social learning* [16]–[18] serves as a useful mathematical abstraction for modeling the interaction of social sensors. Social learning in multiagent systems seeks to answer the following question:

A. How do Decisions Made by Agents Affect Decisions Made by Subsequent Agents?

In social learning, each agent chooses its action by optimizing its local utility function. Subsequent agents then use their private observations together with the actions of previous agents to estimate (learn) an underlying state. The setup is fundamentally different from classical signal processing in which sensors use

Manuscript received November 26, 2013; accepted January 29, 2014. Date of publication March 31, 2014; date of current version June 02, 2014. The work of V. Krishnamurthy was supported by the Canada Research Chairs program, Natural Sciences and Engineering Research Council of Canada, and Social Sciences and Humanities Research Council of Canada. The work of H. V. Poor was supported in part by the U.S. Army Research Office under MURI Grant W911NF-11-1-0036, and in part by the U.S. National Science Foundation under Grant CNS-09-05086.

V. Krishnamurthy is with the Department of Electrical and Computer Engineering, University of British Columbia, Vancouver V6T 1Z4, Canada (e-mail: vikramk@ece.ubc.ca).

H. V. Poor is with the Department of Electrical Engineering, Princeton University, Princeton, NJ08544 USA (e-mail: poor@princeton.edu).

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.2014.2307452

[1]On US Presidential election day in 2012, there were 15 000 tweets per second resulting in 500 million tweets in the day. Twitter can be considered as a real-time sensor.

[2]It is reported in [14] that 81% of hotel managers regularly check Tripadvisor reviews. Luca [15] reports that a one-star increase in the Yelp rating maps to 5–9% revenue increase.

noisy observations to compute estimates—in social learning, agents use noisy observations together with decisions made by previous agents to estimate the underlying state of nature.

In the last decade, 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 [16]–[21] and numerous references therein. Related models have been studied in the context of sequential decision-making in information theory [22], [23] and statistical signal processing [24], [25] in the electrical engineering literature.

Social learning models for interactive sensing can predict unusual behavior. Indeed, a key result in social learning of an underlying random variable is that rational agents eventually herd [17], i.e., they eventually end up choosing the same action irrespective of their private observations. As a result, the actions contain no information about the private observations and so the Bayesian estimate of the underlying random variable freezes. For a multiagent sensing system, such behavior can be undesirable, particularly if individuals herd and make incorrect decisions.

B. Main Results and Organization

In the context of social learning models for interactive sensing, the main ideas and organization of this paper are as follows:

1) Social Learning Protocol: Section II presents a tutorial formulation of the classical Bayesian social learning model, which forms the mathematical basis for modeling interactive sensing among humans. We illustrate the social-learning model in the context of Bayesian signal processing (for easy access to an electrical engineering audience). We then address how self-interested agents performing social learning can achieve useful behavior in terms of optimizing a social welfare function. Such problems are motivated by privacy issues in sensing. If an agent reveals less information in its decisions, it maintains its privacy; on the other hand, as part of a social group, it has an incentive to optimize a social welfare function that helps estimate the state of nature.

2) Data Incest in Online Reputation Systems: Section III deals with the question: how can data incest (misinformation propagation) be prevented in online reputation blogs where social sensors make recommendations?

In the classical social learning model, each agent acts once in a predetermined order. However, in online reputation systems such as Yelp or Tripadvisor, which maintain logs of votes (actions) by agents, social learning takes place with information exchange over a loopy graph (where the agents form the vertices of the graph). Due to the loops in the information exchange graph, *data incest* (misinformation) can propagate: suppose an agent wrote a poor rating of a restaurant on a social media site. Another agent is influenced by this rating, visits the restaurant, and then also gives a poor rating on the social media site. The first agent visits the social media site and notices that another agent has also given the restaurant a poor rating—this confirms his/her rating and he/she enters another poor 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 his/her own rating. Data incest results in a bias in the estimate of state of nature. How can automated protocols be designed to prevent data incest and thereby maintain a fair[3] online reputation system? Section III describes how the administrator of a social network can maintain an unbiased (fair) reputation system.

3) Interaction of Local and Global Decision Makers for Change Detection: Section IV deals with the question: in sensing where individual agents perform social learning to estimate an underlying state of nature, how can changes in the state of nature be detected? Section IV considers a sensing problem that involves change detection. Such sensing problems arise in a variety of applications such as financial trading where individuals react to financial shocks [26]; marketing and advertising [27], [28] where consumers react to a new product; and localization of natural disasters (earthquake and typhoons) [10].

For example, consider measurement of the adoption of a new product using a microblogging platform such as Twitter. The adoption of the technology diffuses through the market but its effects can only be observed through the tweets of select members of the population. These selected members act as sensors for the parameter of interest. Suppose the state of nature suddenly changes due to a sudden market shock or presence of a new competitor. Based on the local actions of the multiagent system that is performing social learning, a global decision maker (such as a market monitor or technology manufacturer) needs to decide whether or not to declare if a change has occurred. How can the global decision maker achieve such change detection to minimize a cost function comprising false alarm rate and delay penalty? The local and global decision makers interact, since the local decisions determine the posterior distribution of subsequent agents, which determines the global decision (stop or continue), which determines subsequent local decisions. We show that this social learning-based change detection problem leads to unusual behavior. The optimal decision policy of the stopping time problem has multiple thresholds. This is unusual: if it is optimal to declare that a change has occurred based on the posterior probability of change, it may not be optimal to declare a change when the posterior probability of change is higher.

4) Coordination of Decisions as a Noncooperative Game: Section V reviews game-theoretic learning in the context of social networks. A large body of research on social networks has been devoted to the diffusion of information (e.g., ideas, behaviors, and trends) [29], [30], and particularly on finding a set of target nodes so as to maximize the spread of a given product [31], [32]. Often customers end up choosing a specific product among several competitors. A natural approach to model this competitive process is via the use of noncooperative game theory [33], [34].

Game theory has traditionally been used in economics and social sciences with a focus on fully rational interactions where strong assumptions are made on the information patterns available to individual agents. In comparison, social sensors are agents with partial information, and it is the dynamic interactions between agents that are of interest. This motivates the need for game-theoretic learning models for agents interacting in social networks.

[3]Maintaining fair reputation systems has financial implications, as it is apparent from footnote 2.

Section V deals with the question: when individuals are self-interested and possess limited sensing, computation, and communication capabilities, can a network of such individuals achieve sophisticated global behavior? In Section V, we discuss a noncooperative game-theoretic learning approach for adaptive decision-making in social networks. This can be viewed as a non-Bayesian version of social learning, The aim is to ensure that all agents eventually choose actions from a common polytope of randomized strategies—namely, the set of correlated equilibria of a noncooperative game. Correlated equilibria are a generalization of Nash equilibria and were introduced by Aumann [35].[4]

5) Extensions: Section VI surveys briefly several extensions of the interactive sensing paradigm of this paper. In particular, the topics of global Bayesian games for coordinated sensing and sensing with information diffusion over large-scale social networks are discussed. These areas have witnessed much recent activity in the economics and computer science literature. Finally, a brief survey of how to obtain representative samples of a social network is given.

C. Perspective

The social learning and game-theoretic learning formalisms mentioned above can be used either as descriptive tools, to predict the outcome of complex interactions among agents in sensing, or as prescriptive tools, to design social networks and sensing systems around given interaction rules. Information aggregation, misinformation propagation, and privacy are important issues in sensing using social sensors. In this paper, we treat these issues in a highly stylized manner so as to provide easy accessibility to an electrical engineering audience. The underlying tools used in this paper are widely used by the electrical engineering research community in the areas of signal processing, control, information theory, and network communications.

In Bayesian estimation, the twin effects of social learning (information aggregation with interaction among 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 of nature.

Privacy issues impose important constraints on social sensors. Typically, individuals are not willing to disclose private observations. Optimizing interactive sensing with privacy constraints is an important problem. Privacy and trust pose conflicting requirements on human-based sensing: privacy requirements result in noisier measurements or lower resolution actions, while maintaining a high degree of trust (reputation) requires accurate measurements. Utility functions, noisy private measurements, and quantized actions are the essential ingredients of the social and game-theoretic learning models presented in this paper that facilitate modeling this tradeoff between reputation and privacy.

The literature in the areas of social learning, sensing, and networking is extensive. Due to page restrictions, in each of the following sections, we provide only a brief review of relevant works. Seminal books in social networks include [36] and [37]. The book [18] contains a complete treatment of social learning models with several remarkable insights. For further references, we refer the reader to [38]–[42]. In [43], a nice description is given of how, if individual agents deploy simple heuristics, the global system behavior can achieve "rational" behavior. The related problem of achieving *coherence* (i.e., agents eventually choosing the same action or the same decision policy) among disparate sensors of decision agents without cooperation has also witnessed intense research; see [44] and [45]. Non-Bayesian social learning models are also studied in [46] and [47].

There is also a growing literature dealing with the interplay of technological networks and social networks [48]. For example, social networks overlaid on technological networks account for a significant fraction of Internet use. Indeed, as discussed in [48], three key aspects of that cut across social and technological networks are the emergence of global coordination through local actions, resource sharing models, and the wisdom of crowds (diversity and efficiency gain). These themes are addressed in the current paper in the context of social learning.

II. MULTIAGENT SOCIAL LEARNING

This section starts with a brief description of the classical social learning model. In this paper, we use social learning as the mathematical basis for modeling interaction of social sensors. A key result in social learning is that rational agents eventually herd, i.e., they choose the same action irrespective of their private observations, and social learning stops. To delay the effect of herding, and thereby enhance social learning, Chamley [18] (see also [49] for related work) has proposed a novel constrained optimal social learning protocol. We review this protocol, which is formulated as a sequential stopping time problem. We show that the constrained optimal social learning proposed by Chamley [18] has a threshold switching curve in the space of public belief states. Thus, the global decision to stop can be implemented efficiently in a social learning model.

A. Motivation: What is Social Learning?

We start with a brief description of the "vanilla"[5] social learning model. In social learning [18], agents estimate the

[4]Aumann's 2005 Nobel prize in economics press release reads: "Aumann also introduced a new equilibrium concept, correlated equilibrium, which is weaker than Nash equilibrium, the solution concept developed by John Nash, an economics laureate in 1994. Correlated equilibrium can explain why it may be advantageous for negotiating parties to allow an impartial mediator to speak to the parties either jointly or separately . . ."

[5]In typical formulations of social learning, the underlying state is assumed to be a random variable and not a Markov chain. Our description below is given in terms of a Markov chain since we wish to highlight the unusual structure of the social learning filter below to a signal processing reader who is familiar with basic ideas in Bayesian filtering. Further, we are interested in change detection problems in which the change time distribution can be modeled as the absorption time of a Markov chain.

underlying state of nature not only from their local measurements, but also from the actions of previous agents. (These previous actions were taken by agents in response to their local measurements; therefore, these actions convey information about the underlying state.) As we describe below, the state estimation update in social learning has a drastically different structure compared to the standard optimal filtering recursion and can result in unusual behavior.

Consider a countable number of agents performing social learning to estimate the state of an underlying finite state Markov chain x. Let $\mathbb{X} = \{1, 2, \dots, X\}$ denote a finite state space, $\boldsymbol{P}$ the transition matrix, and π_0 the initial distribution of the Markov chain.

Each agent acts once in a predetermined sequential order indexed by $k = 1, 2, \dots$ The index k can also be viewed as the discrete time instant when agent k acts. A multiagent system seeks to estimate x_0. Assume at the beginning of iteration k, all agents have access to the public belief π_{k-1} defined in Step iv) below. The social learning protocol proceeds as follows [17], [18]:

1) *Private Observation*: At time k, agent k records a private observation $y_k \in \mathbb{Y}$ from the observation distribution $B_{iy} = \boldsymbol{P}(y|x = i)$, $i \in \mathbb{X}$. Throughout this section, we assume that $\mathbb{Y} = \{1, 2, \dots, Y\}$ is finite.
2) *Private Belief*: Using the public belief π_{k-1} available at time $k-1$ (defined in Step iv) below), agent k updates its private posterior belief $\eta_k(i) = \boldsymbol{P}(x_k = i|a_1, \dots, a_{k-1}, y_k)$ as the following Bayesian update (this is the classical Hidden Markov Model filter [50]):

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

Here $\mathbf{1}_X$ denotes the X-dimensional vector of ones, η_k is an X-dimensional probability mass function (pmf) and $\boldsymbol{P}'$ denotes transpose of the matrix $\boldsymbol{P}$.
3) *Myopic Action*: Agent k takes action $a_k \in \mathcal{A} = \{1, 2, \dots, A\}$ to minimize its expected cost

$$a_k = \arg\min_{a\in\mathcal{A}} \boldsymbol{E}\{c(x,a)|a_1, \dots, a_{k-1}, y_k\} = \arg\min_{a\in\mathcal{A}}\{c_a'\eta_k\}. \quad (2)$$

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. Agent k then broadcasts its action a_k to subsequent agents.
4) *Social Learning Filter*: Given the action a_k of agent k, and the public belief π_{k-1}, each subsequent agent $k' > k$ performs social learning to compute the public belief π_k according to the following "social learning filter":

$$\pi_k = T(\pi_{k-1}, a_k), \quad \text{where } T(\pi, a) = \frac{R_a^\pi \boldsymbol{P}'\pi}{\sigma(\pi, a)} \quad (3)$$

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

$$\boldsymbol{P}(a_k = a|x_k = i, \pi_{k-1} = \pi) = \sum_{y\in\mathbb{Y}} \boldsymbol{P}(a|y,\pi)\boldsymbol{P}(y|x_k = i) \quad (4)$$

$$\boldsymbol{P}(a_k = a|y, \pi) = \begin{cases} 1, & \text{if } c_a' B_y \boldsymbol{P}'\pi \le c_{\tilde{a}}' B_y \boldsymbol{P}'\pi, \tilde{a} \in \mathcal{A} \\ 0, & \text{otherwise.} \end{cases}$$

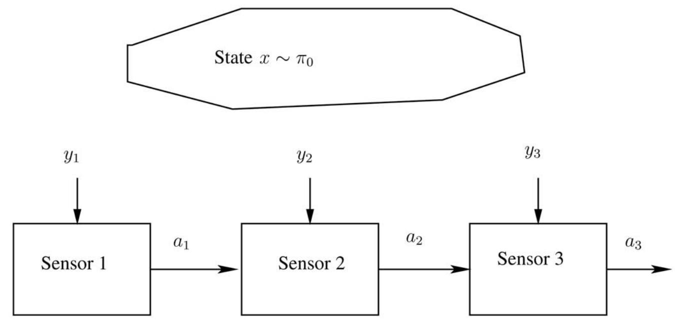

Fig. 1. Interaction of agents in social learning.

The derivation of the social learning filter (3) is given in the discussion as follows.

B. Discussion

Let us pause to give some intuition about the above-mentioned social learning protocol.

1) Information Exchange Structure: Fig. 1 illustrates the above-mentioned social learning protocol in which the information exchange is sequential. Agents send their hard decisions (actions) to subsequent agents. In the social learning protocol, we have assumed that each agent acts once. Another way of viewing the social learning protocol is that there are finitely many agents that act repeatedly in some predefined order. If each agent chooses its local decision using the current public belief, then the setting is identical to the social learning setup. We also refer the reader to [19] for several recent results in social learning over several types of network adjacency matrices.

2) Filtering with Hard Decisions: Social learning can be viewed as agents making *hard* decision estimates at each time and sending these estimates to subsequent agents. In conventional Bayesian state estimation, a *soft* decision is made, namely, the posterior distribution (or equivalently, observation) is sent to subsequent agents. For example, if $\mathcal{A} = \mathbb{X}$, and the costs are chosen as $c_a = -e_a$ where e_a denotes the unit indicator with 1 in the ath position, then $\arg\min_a c_a'\pi = \arg\max_a \pi(a)$, i.e., the maximum aposteriori probability (MAP) state estimate. For this example, social learning is equivalent to agents sending the hard MAP estimates to subsequent agents.

Note that rather than sending a hard decision estimate, if each agent chooses its action $a_k = y_k$ (i.e., agents send their private observations), then the right-hand side of (4) becomes $\sum_{y\in\mathbb{Y}} I(y = y_k)\boldsymbol{P}(y|x_k = i) = \boldsymbol{P}(y_k|x_k = i)$ and so the problem becomes a standard Bayesian filtering problem.

3) Dependence of Observation Likelihood on Prior: The most unusual feature of the above-mentioned protocol (to a signal processing audience) is the social learning filter (3). In standard state estimation via a Bayesian filter, the observation likelihood given the state is completely parametrized by the observation noise distribution and is functionally independent of the current prior distribution. In the social learning filter, the likelihood of the action given the state (which is denoted by R_a^π) is an explicit function of the prior π! Not only does the action likelihood depend on the prior, but it is also a discontinuous function, due to the presence of the $\arg\min$ in (2).

4) Derivation of Social Learning Filter: The derivation of the social learning filter (3) is as follows: define the posterior as $\pi_k(j) = \boldsymbol{P}(x_k = j|a_1, \ldots, a_k)$. Then

$$\begin{aligned}\pi_k(j) &= \frac{1}{\sigma(\pi_{k-1}, a_k)} \boldsymbol{P}(a_k|x_k = j, a_1, \ldots, a_{k-1}) \\ &\quad \cdot \sum_i \boldsymbol{P}(x_k = j|x_{k-1} = i)\boldsymbol{P}(x_{k-1} = i|a_1, \ldots, a_{k-1}) \\ &= \frac{1}{\sigma(\pi_{k-1}, a_k)} \sum_y \boldsymbol{P}(a_k|y_k = y, a_1, \ldots, a_{k-1}) \\ &\quad \cdot \boldsymbol{P}(y_k = y|x_k = j) \sum_i \boldsymbol{P}(x_k = j|x_{k-1} = i)\pi_{k-1}(i) \\ &= \frac{1}{\sigma(\pi_{k-1}, a_k)} \sum_y \boldsymbol{P}(a_k|y_k = y, \pi_{k-1}) \\ &\quad \cdot \boldsymbol{P}(y_k = y|x_k = j) \sum_i \boldsymbol{P}_{ij}\pi_{k-1}(i)\end{aligned}$$

where the normalization term is

$$\begin{aligned}\sigma(\pi_{k-1}, a_k) &= \sum_j \sum_y \boldsymbol{P}(a_k|y_k = y, \pi_{k-1}) \\ &\boldsymbol{P}(y_k = y|x_k = j) \sum_i \boldsymbol{P}_{ij}\pi_{k-1}(i).\end{aligned}$$

The above-mentioned social learning protocol and social learning filter (3) result in interesting dynamics in state estimation and decision-making. We will illustrate two interesting consequences that are unusual to an electrical engineering audience.

1) Rational Agents form herds and information cascades and blindly follow previous agents. This is discussed in Section II-C.
2) Making global decisions on change detection in a multiagent system performing social learning results in multithreshold behavior. This is discussed in Section IV.

C. Rational Agents form Information Cascades

The first consequence of the unusual nature of the social learning filter (3) is that social learning can result in multiple rational agents taking the same action independently of their observations. To illustrate this behavior, throughout this section, we assume that x is a finite state random variable (instead of a Markov chain) with prior distribution π_0.

We start with the following definitions; see also [18].

1) An individual agent k *herds* on the public belief π_{k-1} if it chooses its action $a_k = a(\pi_{k-1}, y_k)$ in (2) independently of its observation y_k.
2) 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 $k > \bar{k}$.
3) An *information cascade* occurs at time $\bar{k}$ if the public beliefs of all agents after time $\bar{k}$ are identical, i.e., $\pi_k = \pi_{\bar{k}}$, for all $k < \bar{k}$.

Note that if an information cascade occurs, then since the public belief freezes, social learning ceases. Moreover, from the above-mentioned definitions, it is clear that an information cascade implies a herd of agents, but the reverse is not true; see Section IV-C for an example.

The following result, which is well known in the economics literature [17], [18], states that if agents follow the above-mentioned social learning protocol, then after some finite time $\bar{k}$, an *information cascade* occurs.[6] The proof follows via an elementary application of the martingale convergence theorem.

Theorem 2.1 [17]: The social learning protocol described in Section II-A leads to an information cascade in finite time with probability 1; i.e., there exists a finite time $\bar{k}$ after which social learning ceases, i.e., public belief $\pi_{k+1} = \pi_k$, $k \geq \bar{k}$, and all agents choose the same action, i.e., $a_{k+1} = a_k$, $k \geq \bar{k}$. □

Instead of reproducing the proof, let us give some insight as to why Theorem 2.1 holds. It can be shown using martingale methods that at some finite time $k = k^*$, the agent's probability $\boldsymbol{P}(a_k|y_k, \pi_{k-1})$ becomes independent of the private observation y_k. Then clearly from (4), $\boldsymbol{P}(a_k = a|x_k = i, \pi_{k-1}) = \boldsymbol{P}(a_k = a|\pi)$. Substituting this into the social learning filter (3), we see that $\pi_k = \pi_{k-1}$. Thus after some finite time k^*, the social learning filter hits a fixed point and social learning stops. As a result, all subsequent agents $k > k^*$ completely disregard their private observations and take the same action a_{k^*}, thereby forming an information cascade (and therefore a herd).

D. Constrained Interactive Sensing: Individual Privacy Versus Group Reputation

The above-mentioned social learning protocol can be interpreted as follows. Agents seek to estimate an underlying state of nature but reveal their actions by maximizing their privacy according to the optimization (2). This leads to an information cascade and social learning stops. In other words, agents are interested in optimizing their own costs (such as maximizing privacy) and ignore the information benefits that their action provides to others.

[6] A nice analogy is provided in [18]. If I see someone walking down the street with an umbrella, I assume (based on rationality) that he has checked the weather forecast and is carrying an umbrella since it might rain. Therefore, I also take an umbrella. So now there are two people walking down the street carrying umbrellas. A third person sees two people with umbrellas and based on the same inference logic, also takes an umbrella. Even though each individual is rational, such herding behavior might be irrational since the first person who took the umbrella, may not have checked the weather forecast.

Another example is that of patrons who decide to choose a restaurant. Despite their menu preferences, each patron chooses the restaurant with the most customers. So, eventually all patrons herd to one restaurant.

Trusov *et al.* [9] quote the following anecdote on user influence in a social network which can be interpreted as herding: "... when a popular blogger left his blogging site for a two-week vacation, the site's visitor tally fell, and content produced by three invited substitute bloggers could not stem the decline."

1) Partially Observed Markov Decision Process Formulation: We now describe an optimized social learning procedure that delays herding.[7] This approach, see [18] for an excellent discussion, is motivated by the following question: how can agents assist social learning by choosing their actions to trade off individual privacy (local costs) with optimizing the reputation[8] of the entire social group?

Suppose agents seek to maximize the reputation of their social group by minimizing the following social welfare cost involving all agents in the social group (compared to the myopic objective (2) used in standard social learning):

$$J_{\mu}(\pi_0) = \boldsymbol{E}^{\mu}_{\pi_0}\left\{\sum_{k=1}^{\infty} \rho^{k-1} c'_{a(\pi_{k-1}, y_k, \mu(\pi_{k-1}))}\eta_k\right\}. \tag{5}$$

In (5), $a(\pi, y, \mu(\pi)))$ denotes the decision rule that agents use to choose their actions, as explained below. Furtheremore, $\rho \in [0, 1)$ is an economic discount factor, and π_0 denotes the initial probability (prior) of the state x. $\boldsymbol{P}^{\mu}_{\pi_0}$ and $\boldsymbol{E}^{\mu}_{\pi_0}$ denote the probability measure and expectation of the evolution of the observations and underlying state, which are strategy dependent.

The key attribute of (5) is that each agent k chooses its action according to the privacy constrained rule

$$a_k = a(\pi_{k-1}, y_k, \mu(\pi_{k-1})). \tag{6}$$

Here, the policy

$$\mu : \pi_{k-1} \to \{1, 2 \ldots, L\}$$

maps the available public belief to the set of L privacy values. The higher the privacy value, the lesser the agent reveals through its action. This is in contrast to standard social learning (2) in which the action chosen is $a(\pi, y)$, namely a myopic function of the private observation and public belief.

The above-mentioned formulation can be interpreted as follows: individual agents seek to maximize their privacy according to social learning (6) but also seek to maximize the reputation of their entire social group (5).

Determining the policy μ^* that minimizes (5), and thereby maximizes the social group reputation, is equivalent to solving a stochastic control problem that is called a partially observed Markov decision process (POMDP) problem [41], [53]. A POMDP comprises a noisy observed Markov chain, such that the dynamics of the posterior distribution (belief state) are controlled by a policy (μ in our case).

2) Structure of Privacy Constrained Sensing Policy: In general, POMDPs are computationally intractable to solve and, therefore, one cannot say anything useful about the structure of the optimal policy μ^*. However, useful insight can be obtained by considering the following extreme case of the above-mentioned problem. Suppose there are two privacy values and each agent k chooses action

$$a_k = \begin{cases} y_k, & \text{if } \mu(\pi_k) = 1(\text{no privacy}) \\ \arg\min_a c'_a \pi_{k-1}, & \text{if } \mu(\pi_k) = 2(\text{full privacy}). \end{cases}$$

That is, an agent either reveals its raw observation (no privacy) or chooses its action by completely neglecting its observation (full privacy). Once an agent chooses the full privacy option, then all subsequent agents choose exactly the same option and therefore herd—this follows since each agent's action reveals nothing about the underlying state of nature. Therefore, for this extreme example, determining the optimal policy $\mu^*(\pi)$ is equivalent to solving a stopping time problem: determine the earliest time for agents to herd (maintain full privacy) subject to maximizing the social group reputation.

For such a quickest herding stopping time problem, one can say a lot about the structure of $\mu^*(\pi)$. Suppose the sensing system wishes to determine if the state of nature is a specific target state (say state 1). Then, Krishnamurthy [41] shows that under reasonable conditions on the observation distribution and supermodular conditions on the costs ([31] discusses supermodularity of influence in social networks), the dynamic programming recursion has a supermodular structure (see also [42], [54]–[57] for related results). This implies that the optimal policy μ^* has the following structure: there exists a threshold curve that partitions the belief space, such that when the belief state is on one side of the curve, it is optimal for agents to reveal full observations; if the belief state is on the other side of the curve, then it is optimal to herd. Moreover, the target state 1 belongs to the region in which it is optimal to herd.[9] This threshold structure of the optimal policy means that if individuals deploy the simple heuristic of "Choose increased privacy when belief is close to the target state,"

then the group behavior is sophisticated—herding is delayed and accurate estimates of the state of nature can be obtained.

III. Data Incest in Online Reputation Systems

This section generalizes the previous section by considering social learning in a social network. How can multiple social sensors interacting over a social network estimate an underlying state of nature? The state could be the position coordinates of an event [10] or the quality of a social parameter such as quality of a restaurant or a political party.

The motivation for this section can be understood in terms of the following sensing example. Consider the following interactions in a multiagent social network where agents seek to estimate an underlying state of nature. Each agent visits a restaurant based on reviews on an online reputation website. The agent then obtains a private measurement of the state (e.g., the quality of food in a restaurant) in noise. After that, he/she reviews the restaurant on the same online reputation website. The information exchange in the social network is modeled by a

[7]In the restaurant problem, an obvious approach to prevent herding is as follows. If a restaurant knew that patrons choose the restaurant with the most customers, then the restaurant could deliberately pay actors to sit in the restaurant, so that it appears popular thereby attracting customers. The methodology in this section where herding is delayed by benevolent agents is a different approach.

[8][51] and [52] contain lucid descriptions of quantitative models for trust, reputation, and privacy.

[9]In standard POMDPs where agents do not perform social learning, it is well known [58] that the set of beliefs for which it is optimal to stop is convex. Such convexity of the herding set does not hold in the current problem. But it is shown in [41] that the set of beliefs for which it is optimal to herd is connected and so is the set of beliefs for which it is optimal to reveal full observations.

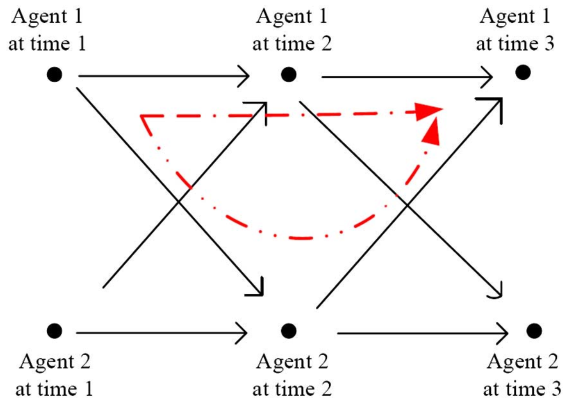

Fig. 2. Example of the information flow (communication graph) in a social network with two agents and over three event epochs. The arrows represent exchange of information.

directed graph. As mentioned in the introduction, data incest [59] arises due to loops in the information exchange graph. This is illustrated in the graph of Fig. 2. Agents 1 and 2 exchange beliefs (or actions) as depicted in Fig. 2. The fact that there are two distinct paths between Agent 1 at time 1 and Agent 1 at time 3 (these two paths are denoted in red) implies that the information of Agent 1 at time 1 is double counted, thereby leading to a data incest event.

How can data incest be removed, so that agents obtain a fair (unbiased) estimate of the underlying state? The methodology of this section can be interpreted in terms of the recent *Time* article [60], which provides interesting rules for online reputation systems. These include: 1) review the reviewers and 2) censor fake (malicious) reviewers. The data incest removal algorithm proposed in this paper can be viewed as "reviewing the reviews" of other agents to see if they are associated with data incest or not.

The rest of this section is organized as follows:

1) Section III-A describes the social learning model that is used to mimic the behavior of agents in online reputation systems. The information exchange between agents in the social network is formulated on a family of time-dependent directed acyclic graphs.
2) In Section III-B, a fair reputation protocol is presented and the criterion for achieving a fair rating is defined.
3) Section III-C presents an incest removal algorithm, so that the online reputation system achieves a fair rating. A necessary and sufficient condition is given on the graph structure of information exchange between agents, so that a fair rating is achievable.

Related work: Collaborative recommendation systems are reviewed and studied in [61] and [62]. In [63], a model of Bayesian social learning is considered in which agents receive private information about the state of nature and observe actions of their neighbors in a tree-based network. Another type of misinformation caused by influential agents (agents who heavily affect actions of other agents in social networks) is investigated in [19]. Misinformation in the context of this paper is motivated by sensor networks where the term "data incest" is used [59]. Data incest also arises in Belief Propagation (BP) algorithms [64], [65], which are used 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 [66], which is similar to data incest in social learning. With the algorithms presented in this section, data incest can be mitigated from Bayesian social learning over nontree graphs that satisfy a topological constraint. The closest work to the current paper is [59]. However, in [59], data incest is considered in a network where agents exchange their private belief states—i.e., no social learning is considered. Simpler versions of this information exchange process and estimation were investigated in [67]–[69]. We also refer the reader to [48] for a discussion of recommender systems.

A. Information Exchange Graph in Social Network

Consider an online reputation system comprising social sensors $\{1, 2, \dots, S\}$ 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 hotel) with known prior distribution π_0. Let $k = 1, 2, 3, \dots$ depict epochs at which events occur. These events involve taking observations, evaluating beliefs, and choosing actions as described below. The index k marks the historical order of events and not necessarily absolute time. However, for simplicity, we refer to k as "time."

To model the information exchange in the social network, we will use a family of directed acyclic graphs. It is convenient also to reduce the coordinates of time k and agent s to a single integer index n, which is used to represent agent s at time k:

$$n \triangleq s + S(k-1), \quad s \in \{1, \dots, S\}, k = 1, 2, 3, \dots \tag{7}$$

We refer to n as a "node" of a time-dependent information flow graph G_n, which we now define.

1) Some Graph Theoretic Definitions: Let

$$G_n = (V_n, E_n), \quad n = 1, 2, \dots \tag{8}$$

denote a sequence of time-dependent graphs of information flow in the social network until and including time k where $n = s + S(k-1)$. Each vertex in V_n represents an agent s' in the social network 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 the communication graph G_n is a subgraph of G_{n+1}. This means that the diffusion of actions can be modeled via a family of time-dependent directed acyclic graphs (a directed graph with no directed cycles).

The algorithms below will involve specific columns of the adjacency matrix and transitive closure matrix of the graph G_n. The Adjacency Matrix $\boldsymbol{A}_n$ of G_n is an $n \times n$ matrix with elements $\boldsymbol{A}_n(i,j)$ given by

$$\boldsymbol{A}_n(i,j) = \begin{cases} 1, & \text{if} (v_j, v_i) \in E, \\ 0, & \text{otherwise}, \end{cases} \quad \boldsymbol{A}_n(i,i) = 0. \tag{9}$$

The transitive closure matrix $\boldsymbol{T}_n$ is the $n \times n$ matrix

$$\boldsymbol{T}_n = \text{sgn}((\mathbf{I}_n - \boldsymbol{A}_n)^{-1}) \tag{10}$$

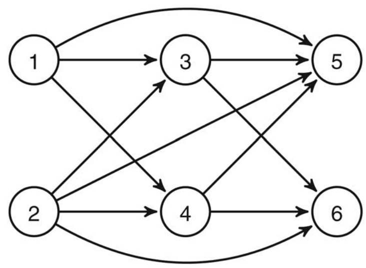

Fig. 3. Example of an information flow network with $S = 2$ two agents, namely $s \in \{1, 2\}$ and time points $k = 1, 2, 3$. Circles represent the nodes indexed by $n = s + S(k-1)$ in the social network, and each edge depicts a communication link between two nodes.

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

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

Note that $\boldsymbol{A}_n(i,j) = 1$ if there is a single hop path between nodes i and j, and in comparison, $\boldsymbol{T}_n(i,j) = 1$ if there exists a path (possible multi-hop) between node 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 following incest removal algorithms:

$$\mathcal{H}_n = \{m : \boldsymbol{A}_n(m,n) = 1\} \tag{11}$$

$$\mathcal{F}_n = \{m : \boldsymbol{T}_n(m,n) = 1\}. \tag{12}$$

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.

2) Example: To illustrate the above-mentioned notation, consider a social network consisting of $S = 2$ two groups with the following information flow graph for three time points $k = 1, 2, 3$.

Fig. 3 shows the nodes $n = 1, 2, \ldots, 6$, where $n = s + 2(k-1)$.

Note that in this example, as it is apparent from Fig. 2, each node remembers all its previous actions. The information flow is characterized by the family of directed acyclic graphs $\{G_1, G_2, G_3, G_4, G_5, G_6\}$ with adjacency matrices

$$\boldsymbol{A}_1 = [0], \quad \boldsymbol{A}_2 = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}, \quad \boldsymbol{A}_3 = \begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix},$$

$$\boldsymbol{A}_4 = \begin{bmatrix} 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}, \quad \boldsymbol{A}_5 = \begin{bmatrix} 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix}.$$

Since nodes 1 and 2 do not communicate, clearly $\boldsymbol{A}_1$ and $\boldsymbol{A}_2$ are zero matrices. Nodes 1 and 3 communicate as do nodes 2 and 3, hence $\boldsymbol{A}_3$ has two ones, etc. Finally, from (11) and (12),

$$\mathcal{H}_5 = \{1, 2, 3, 4\}, \quad \mathcal{F}_5 = \{1, 2, 3, 4\}$$

where $\mathcal{H}_5$ denotes all one hop links to node 5, whereas $\mathcal{F}_5$ denotes all multihop links to node 5.

Note that $\boldsymbol{A}_n$ is always the upper left $n \times n$ submatrix of $\boldsymbol{A}_{n+1}$. Moreover, due to causality with respect to the time index k, the adjacency matrices are always upper triangular.

B. Fair Online Reputation System

1) Protocol for Fair Online Reputation System: The procedure summarized in Protocol 1 aims to evaluate a fair reputation that uses social learning over a social network by eliminating incest.

Aim: Our aim is to design algorithm $\mathcal{A}$ in the automated recommender system (14) of Protocol 1, so that the following requirement is met:

$$\pi_{n-}(i) = \pi^0_{n-}(i), \quad i \in \mathbb{X}$$
$$\text{where } \pi^0_{n-}(i) = P(x = i | \{a_m, m \in \mathcal{F}_n\}). \tag{13}$$

We call π^0_{n-} in (13) the *true* or *fair online rating* available to node n, since $\mathcal{F}_n = \{m : T_n(m,n) = 1\}$ defined in (12) denotes all information (multihop links) available to node n. By definition, π^0_{n-} is incest-free since it is the desired conditional probability that we want. If algorithm $\mathcal{A}$ is designed, so that $\pi_{n-}(i)$ satisfies (13), then the computation (15) and Step v) yield

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

which are, respectively, the correct private belief for node n and the correct after-action public belief.

2) Discussion of Protocol 1: a) Data Incest: It is important to note that without careful design of algorithm $\mathcal{A}$, due to loops in the the public rating, π_{n-} computed using (14) can be substantially different from the fair online rating π^0_{n-} of (13). As a result, η_n computed via (15) will not be the correct private belief and incest will propagate in the network. In other words, η_n, π_{n-}, and π_n are defined purely in terms of their computational expressions in Protocol 1—at this stage, they are not necessarily the desired conditional probabilities, unless algorithm $\mathcal{A}$ is designed to remove incest.

Protocol 1 Incest Removal for Social Learning in an Online Reputation System

---

(i) *Information from Social Network:*

1) Recommendation from Friends: Node n receives past actions $\{a_m, m \in \mathcal{H}_n\}$ from previous nodes $m \in \mathcal{H}_n$ in the network. $\mathcal{H}_n$ is defined in (11).

2) Automated Recommender System: For these past actions $\{a_m, m \in \mathcal{H}_n\}$, the network administrator has already computed the public beliefs $(\pi_m, m \in \mathcal{H}_n)$ using Step v).

The automated recommender system fuses public beliefs $(\pi_m, m \in \mathcal{H}_n)$ into the single recommendation belief π_{n-} as

$$\pi_{n-} = \mathcal{A}(\pi_m, m \in \mathcal{H}_n). \tag{14}$$

The fusion algorithm $\mathcal{A}$ will be designed as follows.

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

(iii) *Private Belief*: Node n then uses y_n and public belief π_{n-} to update its private belief via Bayes' formula as

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

(iv) *Myopic Action*: Node n takes action

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

and inputs its action to the online reputation system.

(v) *Public Belief Update by Network Administrator*: Based on action a_n, the network administrator (automated algorithm) computes the public belief π_n using the social learning filter (3) with $\boldsymbol{P} = I$.

Note that instead of (14), node n could naively (and incorrectly) assume that the public beliefs $\pi_m, m \in \mathcal{H}_n$ that it received are independent. It would then fuse these public beliefs as

$$\pi_{n-} = \frac{\prod_{m\in\mathcal{H}_n} \pi_m}{\mathbf{1}'_X \prod_{m\in\mathcal{H}_n} \pi_m}. \tag{16}$$

This, of course, would result in data incest.

b) How much does an individual remember?: The above-mentioned protocol has the flexibility of modeling cases where each node remembers either some (or all) of its past actions or none of its past actions. This facilitates modeling cases in which people forget most of the past except for specific highlights.

c) Automated Recommender System: Steps i) and v) 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) π_{n-} of (14). This recommender system can operate completely opaquely to the actual user (node n). Node n simply uses the automated rating π_{n-} as the current best available rating from the reputation system.

d) Social Influence. Informational Message Versus Social Message: In Protocol 1, it is important that each individual n deploys Algorithm $\mathcal{A}$ to fuse the beliefs $\{\pi_m, m \in \mathcal{H}_n\}$; otherwise, incest can propagate. Here, $\mathcal{H}_n$ 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." The social message from personal friends exerts a large social influence—it provides significant incentive (peer pressure) for individual n to comply with Protocol 1 and thereby prevent incest. Indeed, a remarkable recent study described in [70] shows that social messages (votes) from known friends have significantly more influence on an individual than the information in the messages themselves. This study includes a comparison of information messages and social messages on Facebook and their direct effect on voting behavior. To quote [70], "The effect of social transmission on real-world voting was greater than the direct effect of the messages themselves . . ."

e) Agent Reputation: The cost function minimization in Step iv) 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, his/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 (or equivalently increase their reputations) by choosing proper actions.

f) Think and Act: Steps ii), iii) iv), and v) of Protocol 1 constitute standard social learning as described in Section II-A. The key difference with standard social learning is Steps i) performed by the network administrator. Agents receive public beliefs from the social network with arbitrary random delays. These delays reflect the time an agent takes between reading the publicly available reputation and making its decision. It is a typical behavior of people to read published ratings multiple times and then think for an arbitrary amount of time before acting.

C. Incest Removal Algorithm in Online Reputation System

We design algorithm $\mathcal{A}$ in Protocol 1, so that it yields the fair public rating π^0_{n-} of (13).

1) Fair Rating Algorithm: It is convenient to work with the logarithm of the un-normalized belief[10]; accordingly define

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

The following theorem shows that the logarithm of the fair rating π^0_{n-} defined in (13) can be obtained as a weighted linear combination of the logarithms of previous public beliefs.

Theorem 3.1 (Fair Rating Algorithm): Consider the online reputation system running Protocol 1. Suppose the following algorithm $\mathcal{A}(l_m, m \in \mathcal{H}_n)$ is implemented in (14) of Protocol 1 by the network administrator:

$$\begin{aligned} &l_{n-}(i) = w'_n l_{1:n-1}(i) \\ &\text{where } w_n = T^{-1}_{n-1} t_n. \end{aligned} \tag{17}$$

Then, $l_{n-}(i) \propto \log \pi^0_{n-}(i)$. That is, algorithm $\mathcal{A}$ computes the fair rating $\log \pi^0_{n-}(i)$ defined in (13).

In (17), $\mathbf{w}_n$ is an $n-1$ dimensional weight vector. Recall that t_n denotes the first $n-1$ elements of the nth column of the transitive closure matrix $\boldsymbol{T}_n$. □

Theorem 3.1 says that the fair rating π^0_{n-} can be expressed as a linear function of the action log-likelihoods in terms of the transitive closure matrix $\boldsymbol{T}_n$ of graph G_n. This is intuitive since π^0_{n-} can be viewed as the sum of information collected by the nodes, such that there are paths between all these nodes and n.

2) Achievability of Fair Rating by Protocol 1: We are not quite done.

1) On the one hand, algorithm $\mathcal{A}$ at node n specified by (14) has access only to beliefs $l_m, m \in \mathcal{H}_n$—equivalently, it

[10] The un-normalized belief proportional to $\pi_n(i)$ is the numerator of the social learning filter (3). The corresponding un-normalized fair rating corresponding to $\pi^0_{n-}(i)$ is the joint distribution $P(x = i, \{a_m, m \in \mathcal{F}_n\})$. By taking the logarithm of the un-normalized belief, Bayes formula merely becomes the sum of the log likelihood and log prior. This allows us to devise a data incest removal algorithm based on linear combinations of the log beliefs.

has access only to beliefs from previous nodes specified by $\boldsymbol{A}_n(:, n)$, which denotes the last column of the adjacency matrix $\boldsymbol{A}_n$.

2) On the other hand, to provide incest-free estimates, algorithm $\mathcal{A}$ specified in (17) requires all previous beliefs $l_{1:n-1}(i)$ that are specified by the non-zero elements of the vector $\boldsymbol{w}_n$. The only way to reconcile points 1 and 2 is to ensure that $\boldsymbol{A}_n(j, n) = 0$ implies $\boldsymbol{w}_n(j) = 0$ for $j = 1, \ldots, n-1$. This condition means that the single hop past estimates $l_m, m \in \mathcal{H}_n$ available at node n according to (14) in Protocol 1 provide all the information required to compute $w_n' l_{1:n-1}$ in (17). This is a condition on the information flow graph G_n. We formalize this condition in the following theorem.

Theorem 3.2 (Achievability of Fair Rating): Consider the fair rating algorithm specified by (17). For Protocol 1 with available information $(\pi_m, m \in \mathcal{H}_n)$ to achieve the estimates l_{n-} of algorithm (17), a necessary and sufficient condition on the information flow graph G_n is

$$\boldsymbol{A}_n(j, n) = 0 \Rightarrow \boldsymbol{w}_n(j) = 0. \tag{18}$$

Therefore, for Protocol 1 to generate incest-free estimates for nodes $n = 1, 2, \ldots$, condition (18) needs to hold for each n. [Recall that $\boldsymbol{w}_n$ is specified in (17).] □

Note that the constraint (18) is purely in terms of the adjacency matrix $\boldsymbol{A}_n$, since the transitive closure matrix (10) is a function of the adjacency matrix.

Summary: Algorithm (17) together with the condition (18) ensure that incest-free estimates are generated by Protocol 1.

3) Illustrative Example (continued): Let us continue with the example of Fig. 2 where we already specified the adjacency matrices of the graphs G_1, G_2, G_3, G_4, and G_5. Using (10), the transitive closure matrices $\boldsymbol{T}_n$ obtained from the adjacency matrices are given by

$$\boldsymbol{T}_1 = [1], \quad \boldsymbol{T}_2 = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad \boldsymbol{T}_3 = \begin{bmatrix} 1 & 0 & 1 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{bmatrix},$$

$$\boldsymbol{T}_4 = \begin{bmatrix} 1 & 0 & 1 & 1 \\ 0 & 1 & 1 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad \boldsymbol{T}_5 = \begin{bmatrix} 1 & 0 & 1 & 1 & 1 \\ 0 & 1 & 1 & 1 & 1 \\ 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}.$$

Note that $\boldsymbol{T}_n(i, j)$ is non-zero only for $i \geq j$ due to causality, since information sent by a social group can only arrive at another social group at a later time instant. The weight vectors are then obtained from (17) as

$$\begin{aligned} \boldsymbol{w}_2 &= [0] \\ \boldsymbol{w}_3 &= [1 \quad 1]' \\ \boldsymbol{w}_4 &= [1 \quad 1 \quad 0]' \\ \boldsymbol{w}_5 &= [-1 \quad -1 \quad 1 \quad 1]'. \end{aligned}$$

Let us examine these weight vectors. $\boldsymbol{w}_2$ 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. 3. $\boldsymbol{w}_3$ means that node 3 uses estimates from nodes 1 and 2; $\boldsymbol{w}_4$ means that node 4 uses estimates only from nodes 1 and 2. The estimate from node 3 is not available at node 4. As shown in Fig. 3, the misinformation propagation occurs at node 5. The vector $\boldsymbol{w}_5$ says that node 5 adds estimates from nodes 3 and 4 and removes estimates from nodes 1 and 2 to avoid double counting of these estimates that is already integrated into estimates from nodes 3 and 4. Indeed, using the algorithm (17), incest is completely prevented in this example.

Let us now illustrate an example in which exact incest removal is impossible. Consider the information flow graph of Fig. 3 but with the edge between nodes 2 and 5 deleted. Then, $\boldsymbol{A}_5(2, 5) = 0$, while $\boldsymbol{w}_5(2) \neq 0$, and, therefore, the condition (18) does not hold. Hence, exact incest removal is not possible for this case.

D. Summary

In this section, we have outlined a controlled sensing problem over a social network in which the administrator controls (removes) data incest and thereby maintains an unbiased (fair) online reputation system. The state of nature could be geographical coordinates of an event (in a target localization problem) or quality of a social unit (in an online reputation system). As discussed earlier, data incest arises due to the recursive nature of Bayesian estimation and nondeterminism in the timing of the sensing by individuals. Details of proofs, extensions, and further numerical studies are presented in [59] and [71].

IV. Interactive Sensing for Quickest Change Detection

In this section, we consider interacting social sensors in the context of detecting a change in the underlying state of nature. Suppose a multiagent system performs social learning and makes local decisions as described in Section II. Given the public beliefs from the social learning protocol, how can quickest change detection be achieved? In other words, how can a global decision maker use the local decisions from individual agents to decide when a change has occurred? It is shown below that making a global decision (change or no change) based on local decisions of individual agents has an unusual structure resulting in a nonconvex stopping set.

A typical application of such social sensors arises in the measurement of the adoption of a new product using a microblogging platform such as Twitter. The adoption of the technology diffuses through the market, but its effects can only be observed through the tweets of select individuals of the population. These selected individuals act as sensors for estimating the diffusion. They interact and learn from the decisions (tweeted sentiments) of other members and, therefore, perform social learning. Suppose the state of nature suddenly changes due to a sudden market shock or presence of a new competitor. The goal for a market analyst or product manufacturer is to detect this change as quickly as possible by minimizing a cost function that involves the sum of the false alarm and decision delay.

Related work: López-Pintado [27], [28] models diffusion in networks over a random graph with arbitrary degree distribution. The resulting diffusion is approximated using deterministic dynamics via a mean field approach [72]. In the seminal paper [1], a sensing system for complex social systems is presented with data collected from cell phones. This data is used in [1] to recognize social patterns, identify socially significant locations, and infer relationships. In [10], people using a microblogging service such as Twitter are considered as sensors. In particular, Sakaki *et al.* [10] consider each Twitter user as a sensor and uses a particle filtering algorithm to estimate the centers of earthquakes and trajectories of typhoons. As pointed out in [10], an important characteristic of microblogging services such as Twitter is that they provide real-time sensing—Twitter users tweet several times a day, whereas standard blog users update information once every several days.

Apart from the above-mentioned applications in real-time sensing, change detection in social learning also arises in mathematical finance models. For example, in agent-based models for the microstructure of asset prices in high-frequency trading in financial systems [26], the state denotes the underlying asset value that changes at a random time τ^0. Agents observe local individual decisions of previous agents via an order book, combine these observed decisions with their noisy private signals about the asset, selfishly optimize their expected local utilities, and then make their own individual decisions (whether to buy, sell, or do nothing). The market evolves through the orders of trading agents. Given this order book information, the goal of the market maker (global decision maker) is to achieve quickest change point detection when a shock occurs to the value of the asset [73].

A. Classical Quickest Detection

The classical Bayesian quickest detection problem [74] is as follows: an underlying discrete time state process x jump-changes at a geometrically distributed random time τ^0. Consider a sequence of discrete time random measurements $\{y_k, k \geq 1\}$, such that conditioned on the event $\{\tau^0 = t\}$, y_k, $k \leq t$ are independent and identically distributed (i.i.d.) random variables with distribution B_1 and $y_k, k > t$ are i.i.d. random variables with distribution B_2. The quickest detection problem involves detecting the change time τ^0 with minimal cost. That is, at each time $k = 1, 2, \ldots$, a decision $u_k \in \{1(\text{stop and announce change}), 2(\text{continue})\}$ needs to be made to optimize a tradeoff between false alarm frequency and linear delay penalty.

To formalize this setup, let $\boldsymbol{P} = \begin{bmatrix} 1 & 0 \\ 1 - \boldsymbol{P}_{22} & \boldsymbol{P}_{22} \end{bmatrix}$ denote the transition matrix of a two state Markov chain x in which state 1 is absorbing. Then it is easily seen that the geometrically distributed change time τ^0 is equivalent to the time at which the Markov chain enters state 1. That is, $\tau^0 = \min\{k : x_k = 1\}$ and $\mathbf{E}\{\tau^0\} = 1/(1 - P_{22})$. Let τ be the time at which the decision $u_k = 1$ (announce change) is taken. The goal of quickest detection is to minimize the Kolmogorov–Shiryaev criterion for detection of a disorder [75]:

$$J_\mu(\pi_0) = d\mathbf{E}^\mu_{\pi_0}\{(\tau - \tau^0)^+\} + f\mathbf{P}^\mu_{\pi_0}(\tau < \tau^0). \tag{19}$$

Here $x^+ = x$ if $x \geq 0$ and 0 otherwise. The non-negative constants d and f denote the delay and false alarm penalties, respectively. So, waiting too long to announce a change incurs a delay penalty d at each time instant after the system has changed, while declaring a change before it happens incurs a false alarm penalty f. In (19), μ denotes the strategy of the decision maker. $\mathbf{P}^\mu_{\pi_0}$ and $\mathbf{E}^\mu_{\pi_0}$ are the probability measure and expectation of the evolution of the observations and Markov state, which are strategy-dependent. π_0 denotes the initial distribution of the Markov chain x.

In classical quickest detection, the decision policy μ is a function of the two-dimensional belief state (posterior pmf) $\pi_k(i) = \boldsymbol{P}(x_k = i | y_1, \ldots, y_k, u_1, \ldots, u_{k-1})$, $\quad i = 1, 2,$ with $\pi_k(1) + \pi_k(2) = 1$. So, it suffices to consider one element, say $\pi_k(2)$, of this pmf. Classical quickest change detection (see for example [74]) says that the policy $\mu^*(\pi)$, which optimizes (19), has the following threshold structure: there exists a threshold point $\pi^* \in [0, 1]$, such that

$$\mu^*(\pi_k) = \begin{cases} 2(\text{continue}), & \text{if } \pi_k(2) \geq \pi^* \\ 1(\text{announce change}), & \text{if } \pi_k(2) < \pi^*. \end{cases} \tag{20}$$

B. Multiagent Quickest Detection Problem

With the above-mentioned classical formulation in mind, consider now the following multiagent quickest change detection problem. Suppose that a multiagent system performs social learning to estimate an underlying state according to the social learning protocol of Section II-A. That is, each agent acts once in a predetermined sequential order indexed by $k = 1, 2, \ldots$ (Equivalently, as pointed out in the discussion in Section II-A, a finite number of agents act repeatedly in some predefined order and each agent chooses its local decision using the current public belief.) Given these local decisions (or equivalently the public belief), the goal of the global decision maker is to minimize the quickest detection objective (19). The problem now is a nontrivial generalization of classical quickest detection. The posterior π is now the public belief given by the social learning filter (3) instead of a standard Bayesian filter. There is now interaction between the local and the global decision makers. The local decision a_k from the social learning protocol determines the public belief state π_k via the social learning filter (3), which determines the global decision (stop or continue), which determines the local decision at the next time instant, and so on.

The global decision maker's policy $\mu^* : \pi \rightarrow \{1, 2\}$ that optimizes the quickest detection objective (19) and the cost $J_{\mu^*}(\pi_0)$ of this optimal policy are the solution of "Bellman's dynamic programming equation"

$$\begin{aligned} \mu^*(\pi) &= \arg\min\{f\pi(2), d(1 - \pi(2)) \\ &\quad + \sum_{a \in \mathcal{A}} V(T(\pi, a))\sigma(\pi, a)\}, \quad J_{\mu^*}(\pi_0) = V(\pi_0) \\ V(\pi) &= \min\{f\pi(2), d(1 - \pi(2)) \\ &\quad + \sum_{a \in \mathcal{A}} V(T(\pi, a))\sigma(\pi, a)\}. \end{aligned} \tag{21}$$

Here $T(\pi, a)$ and $\sigma(\pi, a)$ are given by the social learning filter (3)—recall that a denotes the local decision. $V(\pi)$ is called

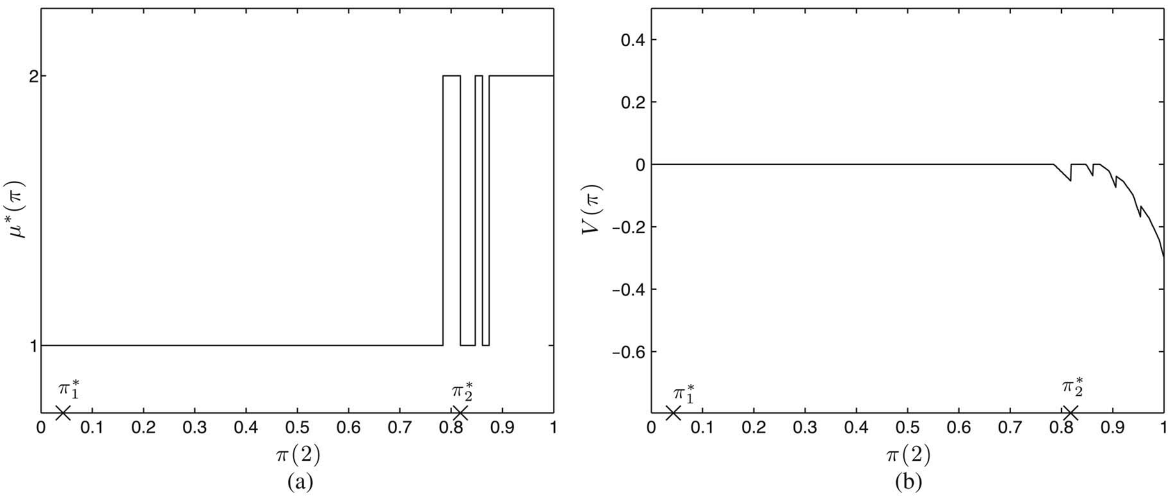

Fig. 4. Optimal global decision policy for social learning-based quickest change detection for a geometric distributed change time. The parameters are specified in Section IV-C. The optimal policy $\mu^*(\pi) \in \{1(\text{announce change}), 2(\text{continue})\}$ is characterized by a triple threshold—i.e., it switches from 1 to 2 three times as the posterior $\pi(2)$ increases. The value function $V(\pi)$ is nonconcave and discontinuous in π. As explained in the text, for $\pi(2) \in [0, \pi_1^*]$, all agents herd, while for $\pi(2) \in [\pi_2^*, 1]$, individual agents herd (see definitions in Section II-C): (a) optimal global decision policy $\mu^*(\pi)$ and (b) value function $V(\pi)$ for the global decision policy.

the "value function"—it is the cost incurred by the optimal policy when the initial belief state (prior) is π. As is shown in the numerical example below, the optimal policy $\mu^*(\pi)$ has a very different structure compared to classical quickest detection.

C. Numerical Example

We now illustrate the unusual multithreshold property of the global decision maker's optimal policy $\mu^*(\pi)$ in multiagent quickest detection with social learning. Consider the social learning model of Section II-A with the following parameters: The underlying state is a 2-state Markov chain x with state space $\mathbb{X} = \{1, 2\}$ and transition probability matrix $\boldsymbol{P} = \begin{bmatrix} 1 & 0 \\ 0.05 & 0.95 \end{bmatrix}$. So, the change time τ^0 (i.e., the time the Markov chain jumps from state 2 into absorbing state 1) is geometrically distributed with $E\{\tau^0\} = 1/0.05 = 20$.

Social Learning Parameters: Individual agents observe the Markov chain x in noise with the observation symbol set $\mathbb{Y} = \{1, 2\}$. Suppose the observation likelihood matrix with elements $B_{iy} = \boldsymbol{P}(y_k = y | x_k = i)$ is $\boldsymbol{B} = \begin{bmatrix} 0.9 & 0.1 \\ 0.1 & 0.9 \end{bmatrix}$. Agents can choose their local actions a from the action set $\mathcal{A} = \{1, 2\}$. The state-dependent cost matrix of these actions is $c = (c(i, a), i \in X, a \in \mathcal{A}) = \begin{bmatrix} 4.57 & 5.57 \\ 2.57 & 0 \end{bmatrix}$. Agents perform social learning with the above-mentioned parameters. The intervals $[0, \pi_1^*]$ and $[\pi_2^*, 1]$ in Fig. 4(a) are regions where the optimal local actions taken by agents are independent of their observations. For $\pi(2) \in [\pi_2^*, 1]$, the optimal local action is 2, and for $\pi(2) \in [0, \pi_1^*]$, the optimal local action is 1. Therefore, individual agents herd for belief states in these intervals (see the definition in Section II-C) and the local actions do not yield any information about the underlying state. Moreover, the interval $[0, \pi_1^*]$ depicts a region where all agents herd (again see the definition in Section II-C), meaning that once the belief state is in this region, it remains so indefinitely and all agents choose the same local action 1.[11]

Global Decision-Making: Based on the local actions of the agents performing social learning, the global decision maker needs to perform quickest change detection. The global decision maker uses the delay penalty $d = 1.05$ and false alarm penalty $f = 3$ in the objective function (19). The optimal policy $\mu^*(\pi)$ of the global decision maker where $\pi = [1 - \pi(2), \pi(2)]'$ is plotted versus $\pi(2)$ in Fig. 4(a). Note $\pi(2) = 1$ means that with certainty, no change has occurred; whereas, $\pi(2) = 0$ means that with certainty, a change has occurred. The policy $\mu^*(\pi)$ was computed by constructing a uniform grid of 1000 points for $\pi(2) \in [0, 1]$ and then implementing the dynamic programming equation (21) via a fixed point value iteration algorithm for 200 iterations. The horizontal axis $\pi(2)$ is the posterior probability of no change. The vertical axis denotes the optimal decision: $u = 1$ denotes stop and declare change, while $u = 2$ denotes continue.

The most remarkable feature of Fig. 4(a) is the multithreshold behavior of the global decision maker's optimal policy $\mu^*(\pi)$. Recall that $\pi(2)$ depicts the posterior probability of no change. So, consider the region where $\mu^*(\pi) = 2$ and sandwiched between two regions where $\mu^*(\pi) = 1$. Then as $\pi(2)$ (posterior probability of no change) increases, the optimal policy switches from $\mu^*(\pi) = 2$ to $\mu^*(\pi) = 1$. In other words, the optimal global decision policy "changes its mind"—it switches from no change to change as the posterior probability of a change decreases! Thus, the global decision (stop or continue) is a nonmonotone function of the posterior probability obtained from local decisions.

Fig. 4(b) shows the associated value function obtained via stochastic dynamic programming (21). Recall that $V(\pi)$ is the cost incurred by the optimal policy with initial belief state π. Unlike standard sequential detection problems in which the value

[11]Note that even if the agent k herds so that its action a_k provides no information about its private observation y_k, the public belief still evolves according to the predictor $\pi_{k+1} = P'\pi_k$. So an information cascade does not occur in this example.

function is concave, the figure shows that the value function is nonconcave and discontinuous. To summarize, Fig. 4 shows that social learning-based quickest detection results in fundamentally different decision policies compared to classical quickest detection (which has a single threshold). Thus, making global decisions (stop or continue) based on local decisions (from social learning) is nontrivial. In [42], a detailed analysis of this problem is given together with a characterization of this multithreshold behavior. More general phase-distributed change times are considered in [42].

V. Coordination of Decisions in Sensing—Noncooperative Game Approach

The discussion so far has dealt with Bayesian social learning models for sensing. In this section, we present a highly stylized non-Bayesian noncooperative game-theoretic learning approach for adaptive decision-making among agents.

Social and economic situations often involve interacting decision-making with diverging interests. Decision makers may act independently or form collaborative groups, wherein enforceable binding agreements ensure coordination of joint decisions. For instance, a person may choose the same cellphone carrier as the majority of family and friends to take advantage of the free talk times. Social networks diffuse information and hence facilitate coordination of such cooperative/self-interested units. This section examines how global coordination of decisions can be obtained when self-interested agents form a social network.

As mentioned in the Introduction, human-based sensing systems comprise agents with partial information and it is the dynamic interactions between agents that is of interest. This motivates the need for game-theoretic learning models for agents interacting in social networks. Learning dynamics in games typically can be classified into Bayesian learning [18], [21], adaptive learning [76], and evolutionary dynamics [77], [78]. We have already focussed on Bayesian social learning[12] in previous sections of the paper, and some further remarks are made in Section VI on global Bayesian games.

In this section, we focus on adaptive learning where individual agents deploy simple rule-of-thumb strategies. The aim is to determine if such simple individual behavior can result in sophisticated global behavior. We are interested in cases where the global behavior converges to the set of *correlated equilibria*.

A. Related Works

1) Correlated Equilibria: The set of correlated equilibria is a more natural construct in decentralized adaptive learning environments than the set of Nash equilibria,[13] since it allows for individual players to coordinate their actions. This coordination can lead to higher performance [35] than if each player chooses actions independently as required by a Nash equilibrium. As described in [79], it is unreasonable to expect in a learning environment that players act independently (as required by a Nash equilibrium), since the common history observed by all players acts as a natural coordination device.[14] The set of correlated equilibria is also structurally simpler than the set of Nash equilibria; the set of correlated equilibria is a convex polytope in the set of randomized strategies, whereas Nash equilibria are isolated points at the extrema of this set. Indeed, a feasible point in the set of correlated equilibria is obtained straightforwardly by a linear programming solver.

The works [80] and [81] report experimental results that explore the empirical validity of correlated equilibria. It is found in [81] that when private recommendations from a mediator are not available to the players, the global behavior is characterized by mixed-strategy Nash equilibria. Their main result is that players follow recommendations from a third party only if those recommendations are drawn from a correlated equilibrium that is "payoff-enhancing" relative to the available Nash equilibria. The work [82] addresses the problem of information release in social media via a game-theoretic formulation. The proposed model captures a user's willingness to release, withhold, or lie about information depending on the behavior of the user's circle of friends. The empirical study infers a relationship between the qualitative structure of the game equilibria and the automorphism group of the social network.

2) Game-Theoretic Learning and Regret-Based Algorithms: A comprehensive textbook in game-theoretic learning is [83]. Algorithms for game-theoretic learning are broadly classified into best response, fictitious play, and regret matching. In general, it is impossible to guarantee convergence to a Nash equilibrium without imposing conditions on the structure of the utility functions in the game. For supermodular games [84], best response algorithms can be designed to converge either to the smallest or the largest Nash equilibrium. Fictitious play is one of the oldest and the best-known models of learning in games; we refer the reader to [85] for convergence of stochastic fictitious play algorithms.

In this section, we focus on low *internal regret* algorithms as a strategy of play in game-theoretic learning. The internal regret[15] compares the loss of a strategy to the loss of a modified strategy, which consistently replaces one action by another—for example, "every time you bought Windows, you should have bought Apple instead." We refer the reader to [86] for an excellent discussion of internal and external regret-based algorithms. Two seminal papers in low internal regret algorithms for game-theoretic learning are [76] and [79] where the terminology "regret-matching" is used. In [79], it is proved that when all agents share stage actions and follow the proposed regret-based adaptive procedure, the collective behavior converges to the set of correlated equilibria. In [76], the authors assumed that agents

[12]The social learning protocol of Section II-A can be viewed as a Bayesian game comprising a countable number of agents, where each agent plays once in a specified order to minimize its cost; see [18] for further details on this game-theoretic interpretation.

[13]The set of correlated equilibria is defined in (25). Nash equilibria are a special case of correlated equilibria where the joint strategy is chosen as the product distribution for all players, i.e., all the agents choose their actions independently.

[14]Hart and Mas-Colell observe in [76] that for most simple adaptive procedures, "...there is a natural coordination device: the common history, observed by all players. It is thus reasonable to expect that, at the end, independence among players will not obtain."

[15]In comparison, the external regret compares the performance of a strategy selecting actions to the performance of the best of those actions in hindsight.

do *not* observe others' actions and proposed a reinforcement learning procedure that converges to the set of correlated equilibria. More recently, Krishnamurthy *et al.* [38], Maskery *et al.* [39], and Namvar *et al.* [87], [88] consider learning in a dynamic setting where a regret matching-type algorithm tracks a time varying set of correlated equilibria.

It should be noted that the concept of regret is well-known in the decision theory and random sampling literature [89]. Regret minimization methods are of interest in the multiarmed bandit literature, which is concerned with optimizing the cumulative objective function values realized over a period of time [90], which involves finding the best arm after a given number of arm pulls.

B. Regret-Based Decision-Making

Consider a noncooperative repeated game comprising L agents. Each agent l has a utility function $U^l(a^1,\dots,a^L)$. Here a^l denotes the action chosen by agent l, and a^{-l} denote the actions chosen by all agents excluding agent l. The utility function can be quite general. For example, Namvar *et al.* [91] consider the case in which the L agents are organized into M nonoverlapping social (friendship) groups, such that agents in a social group share the same utility function. The utility function could also reflect reputation or privacy using the models in [51] and [52].

Suppose each agent l chooses its actions according to the following adaptive algorithm running over time $k = 1, 2, \dots$

1) At time $k+1$, choose action a^l_{k+1} from pmf ψ^l_{k+1}, where

$$\psi^l_{k+1}(i) = \boldsymbol{P}\big(a^l_{k+1} = i | a^l_k\big) = \begin{cases} \frac{|r^l_k(i,j)|^+}{C}, & j \neq a^l_k \\ 1 - \sum_{j\neq i} \frac{|r^l_k(i,j)|^+}{C}, & j = a^l_k. \end{cases} \quad (22)$$

Here C is a sufficiently large positive constant, so that ψ^l_{k+1} is a valid pmf.

2) Update the internal regret matrix $\boldsymbol{r}^l_k$ that determines the pmf ψ^l_{k+1} via the stochastic approximation algorithm

$$\boldsymbol{r}^l_k(i,j) = \boldsymbol{r}^l_{k-1}(i,j) + \frac{1}{k}\Big(\big[U^l\big(j, a^{-l}_k\big) - U^l\big(a^l_k, a^{-l}_k\big)\big] I_{\{a^l_k = i\}} - \boldsymbol{r}^l_{k-1}(i,j)\Big). \quad (23)$$

Step 1) corresponds to each agent choosing its action randomly from a Markov chain with transition probability ψ^l_{k+1}. These transition probabilities are computed in Step 2) in terms of the internal regret matrix $\boldsymbol{r}^l_k$, which is the time-averaged regret agent l experiences for choosing action i instead of action j for each possible action $j \neq i$ (i.e., how much better off it would be if it had chosen action j instead of i):

$$\boldsymbol{r}^l_n(i,j) = \frac{1}{n}\sum_{k=1}^{n}\big[U^l\big(j, a^{-l}_k\big) - U^l\big(a^l_k, a^{-l}_k\big)\big] \cdot I_{\{a^l_k = i\}}. \quad (24)$$

The above-mentioned algorithm can be generalized to consider multiple social groups. If agents within each social group share their actions and have a common utility, then they can fuse their individual regrets into a regret for the social group. As shown in [91], this fusion of regrets can be achieved via a linear combination of the individual regrets where the weights of the linear combination depend on the reputation of the agents that constitute the social group.

C. Coordination in Sensing

We now address the following question:

If each agent chooses its action according to the above-mentioned regret-based algorithm, what can one say about the emergent global behavior?

By emergent global behavior, we mean the empirical frequency of actions taken over time by all agents. For each L-tuple of actions (a^l, a^{-l}), define the empirical frequency of actions taken up to time n as

$$z_n(a^l, a^{-l}) = \frac{1}{n}\sum_{k=1}^{n} I(a_k = a^l, a^{-l}_k = a^{-l}).$$

The seminal papers [79] and [43] show that the empirical frequency of actions z_n converges as $n \to \infty$ to the set of *correlated equilibria* of a noncooperative game. As noted previously, correlated equilibria constitute a generalization of Nash equilibria. The set of correlated equilibria $\mathcal{C}_e$ is the set of probability distributions on the joint action profile (a^l, a^{-l}) that satisfy

$$\mathcal{C}_e = \Big\{\mu : \sum_{a^{-l}} \mu^l(j, a^{-l})\big[U^l((i, a^{-l})) - U^l((j, a^{-l}))\big] \leq 0, \quad \forall l, j, i\Big\}. \quad (25)$$

Here $\mu^l(j, a^{-l}) = P^l(a^l = j, a^{-l})$ denotes the randomized strategy (joint probability) of player l choosing action j and the rest of the players choosing action a^{-l}. The correlated equilibrium condition (25) states that instead of taking action j (which is prescribed by the equilibrium strategy $\mu^l(j, a^{-l})$), if player l cheats and takes action i, it is worse off. So, there is no unilateral incentive for any player to cheat.

To summarize, the above-mentioned algorithm ensures that all agents eventually achieve *coordination (consensus) in decision-making*—the randomized strategies of all agents converge to a common convex polytope $\mathcal{C}_e$. Step 2) of the algorithm requires that each agent knows its own utility and the actions of other agents—but agents do not need to know the utility functions of other agents. In [76], a "blind" version of this regret-based algorithm is presented in which agents do not need to know the actions of other agents. These algorithms can be viewed as examples, in which simple heuristic behavior by individual agents (choosing actions according to the measured regret) results in sophisticated global outcomes [43], namely convergence to $\mathcal{C}_e$, thereby coordinating decisions.

We refer to [38], [39], and [87] for generalizations of the above-mentioned algorithm to the tracking case in which the step size for the regret matrix update is a constant. Such algorithms can track the correlated equilibria of games with time-varying parameters. Moreover, Namvar *et al.* [88] give sufficient conditions for algorithm to converge to the set of correlated equilibria when the regrets from one agent to other agents diffuse over a social network.

VI. Closing Remarks

In this paper, we have used social learning as a model for interactive sensing with social sensors. We summarize here some extensions of the social learning framework that are relevant to interactive sensing.

A. Crowd Behavior and Polling Agents

1) Wisdom of Crowds: Surowiecki's book [92] is an excellent popular piece that explains the wisdom-of-crowds hypothesis. The wisdom-of-crowds hypothesis predicts that the independent judgments of a crowd of individuals (as measured by any form of central tendency) will be relatively accurate, even when most of the individuals in the crowd are ignorant and error-prone. The book also studies situations (such as rational bubbles) in which crowds are not wiser than individuals. Collect enough people on a street corner staring at the sky, and everyone who walks past will look up. Such herding behavior is typical in social learning.

2) In Which Order Should Agents Act? In the social learning protocol, we assumed that the agents act sequentially in a predefined order. However, in many social networking applications, it is important to optimize the order in which agents act. For example, consider an online review site where individual reviewers with different reputations make their reviews publicly available. If a reviewer with high reputation publishes his/her review first, this review will unduly affect the decision of a reviewer with lower reputation. In other words, if the most senior agent "speaks" first, it would unduly affect the decisions of more junior agents. This could lead to an increase in bias of the underlying state estimate.[16] On the other hand, if the most junior agent is polled first, then since its variance is large, several agents would need to be polled in order to reduce the variance. We refer the reader to [94] for an interesting description of who should speak first in a public debate.[17] It turns out that for two agents, the seniority rule is always optimal for any prior—i.e., the senior agent speaks first followed by the junior agent; see [94] for the proof. However, for more than two agents, the optimal order depends on the prior and the observations in general.

B. Global Games for Coordinating Sensing

In the classical Bayesian social learning model of Section II, agents act sequentially in time. The global games model that has been studied in economics during the last two decades considers multiple agents that act simultaneously by predicting the behavior of other agents. The theory of global games was first introduced in [95] as a tool for refining equilibria in economic game theory; see [96] for an excellent exposition. Global games represent a useful method for decentralized coordination among agents; they have been used to model speculative currency attacks and regime change in social systems; see [96]–[98].

The most widely studied form of a global game is a one-shot Bayesian game, which proceeds as follows: consider a continuum of agents in which each agent i obtains noisy measurements y^i of an underlying state of nature x. Assume all agents have the same observation likelihood density $p(y|x)$ but the individual measurements obtained by agents are statistically independent of those obtained by other agents. Based on its observation y^i, each agent takes an action $a^i \in \{1, 2\}$ to optimize its expected utility $\boldsymbol{E}\{U(a^i, \alpha)|y^i\}$ where $\alpha \in [0, 1]$ denotes the fraction of all agents that take action 2. Typically, the utility $U(1, \alpha)$ is set to zero.

For example, suppose x (state of nature) denotes the quality of a social group and y^i denotes the measurement of this quality by agent i. The action $a^i = 1$ means that agent i decides not to join the social group, while $a^i = 2$ means that agent i joins the group. The utility function $U(a^i = 2, \alpha)$ for joining the social group depends on α, where α is the fraction of people who decide to join the group. In [97], the utility function is chosen as follows: if $\alpha \approx 1$, i.e., too many people join the group, then the utility to each agent is small, since the group is too congested and agents do not receive sufficient individual service. On the other hand, if $\alpha \approx 0$, i.e., too few people join the group, then the utility is also small, since there is not enough social interaction.

Since each agent is rational, it uses its observation y^i to predict α, i.e., the fraction of other agents that choose action 2. The main question is: what is the optimal strategy for each agent i to maximize its expected utility?

It has been shown that for a variety of measurement noise models (observation likelihoods $p(y|x)$) and utility functions U, the symmetric Bayesian Nash equilibrium of the global game is unique and has a threshold structure in the observation. This means that given its observation y^i, it is optimal for each agent i to choose its actions as follows:

$$a^i = \begin{cases} 1, & y^i < y^* \\ 2, & y^i \geq y^* \end{cases} \tag{26}$$

where the threshold y^* depends on the prior, noise distribution, and utility function.

In the above-mentioned example of joining a social group, the result means that if agent i receives a measurement y^i of the quality of the group, and y^i exceeds a threshold y^*, then it should join. This is yet another example of simple local behavior (act according to a threshold strategy) resulting in global sophisticated behavior (Bayesian Nash equilibrium). As can be seen, global games provide a decentralized way of achieving coordination among social sensors. In [98], the above-mentioned one-shot Bayesian game is generalized to a dynamic (multistage) game operating over a possibly infinite horizon. Such games facilitate modeling the dynamics of how people join, interact, and leave social groups.

The papers [40] and [99] use global games to model networks of sensors and cognitive radios. In [97], it has been shown that the

[16] To quote a recent paper from the Haas School of Business, U.C. Berkeley [93]: "In 94% of cases, groups (of people) used the first answer provided as their final answer... Groups tended to commit to the first answer provided by any group member." People with dominant personalities tend to speak first and most forcefully "even when they actually lack competence."

[17] As described in [94], seniority is considered in the rules of debate and voting in the U.S. Supreme Court. "In the past, a vote was taken after the newest justice to the Court spoke, with the justices voting in order of ascending seniority largely, it was said, to avoid the pressure from long-term members of the Court on their junior colleagues."

above-mentioned threshold structure (26) for the Bayesian Nash equilibrium breaks down if the utility function $U(2, \alpha)$ decreases too rapidly due to congestion. The equilibrium structure becomes much more complex and can be described by the following quotation [97]: Nobody goes there anymore. It's too crowded. —*Yogi Berra*

C. Sensing with Information Diffusion over Large-Scale Social Networks

In this paper, we have not considered the dynamics of large-scale random graphs. The fundamental theory of network science is well documented in seminal books such as [37] and [100], and involves the interplay of random graphs and game theory. In large-scale social networks, it is of significant interest to consider how information and, therefore, behavior of individuals diffuses over a social network comprising a population of interacting agents based on sampling the population. As described in [28], there is a wide range of social phenomena such as diffusion of technological innovations, cultural fads, and economic conventions [18] where individual decisions are influenced by the decisions of others.

Several recent papers investigate the diffusion of information in real-world social networks such as Facebook, Twitter, and blogs. Motivated by marketing applications, Sun *et al.* [101] study the diffusion (contagion) behavior in Facebook. Using data on around 260 000 Facebook *pages* (which advertise products, services, bands, and celebrities), Sun *et al.* [101] analyze how information diffuse on Facebook. In [102], the spread of *hashtags* on Twitter is studied.

Consider a social network where the states of individual nodes evolve over time as probabilistic functions of the states of their neighbors and an underlying target process. The evolution in the state of agents in the network can be viewed as diffusion of information in the network. Such *Susceptible-Infected-Susceptible* (SIS) models for diffusion of information in social networks have been extensively studied in [27], [28], [36], [37], and [103] to model, e.g., the adoption of a new technology in a consumer market.

Consistent with the interactive sensing paradigm of this paper, one can consider two extensions of the basic SIS model: first, the states of individual nodes evolve over time as a probabilistic function of the states of their neighbors *and* an underlying target process. The underlying target process can be viewed as the market conditions or competing technologies that evolve with time and affect the information diffusion. Second, the nodes in the social network are sampled randomly to determine their state. As the adoption of the new technology diffuses through the network, its effect is observed via sentiment (such as tweets) of these selected members of the population. These selected nodes act as social sensors. In signal processing terms, the underlying target process can be viewed as a signal, and the social network can be viewed as a sensor. The key difference compared to classical signal processing is that the social network (sensor) has dynamics due to the information diffusion.

The aim is to estimate the underlying target state and the state probabilities of the nodes by sampling measurements at nodes in the social network. In a Bayesian estimation context, this is equivalent to a filtering problem involving estimating the state of a prohibitively large-scale Markov chain in noise. A key idea is to use *mean field dynamics* as an approximation (with provable bounds) for the information diffusion and, thereby, obtain a tractable model. Such mean field dynamics have been studied in [72] and applied to social networks in [27], [28], and [36]. For an excellent recent exposition of interacting particle systems comprising agents each with a finite state space, see [104], where the more apt term "Finite Markov Information Exchange (FMIE) process" is used.

D. Sampling Social Networks

An important question regarding sensing in a social network is: how can one construct a small but representative sample of a large-scale social network? Leskovec and Faloutsos [105] study and compare several scale-down and back-in-time sampling procedures.

An obvious sampling scheme for a population is uniform sampling with or without replacement. Recently, there has been a significant progress in *social sampling* [106]. In social sampling, participants in a poll respond with a summary of their friends' responses. 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 form d of its friends. However, the variance and bias of the estimate depend strongly on the social network structure.[18]

Another important sampling methodology for social networks is *respondent-driven sampling* (RDS). RDS can be viewed as a Markov-chain Monte-Carlo sampling strategy and was introduced by Heckathorn [107]–[109] sampling from hidden populations in social networks. As mentioned in [110], the U.S. Centers for Disease Control and Prevention (CDC) recently selected RDS for a 25-city study of injection drug users that is part of the National HIV Behavioral Surveillance System [111]. RDS is a variant of the well-known method of snowball sampling where current sample members recruit future sample members. The RDS procedure is as follows: a small number of people in the target population serve as seeds. After participating in the study, the seeds recruit other people they know through the social network in the target population. The sampling continues according to this procedure with current sample members recruiting the next wave of sample members until the desired sampling size is reached.

VII. Summary

This paper has considered social learning models for interaction among sensors where agents use their private observations along with actions of other agents to estimate an underlying state of nature. We have considered extensions of the basic social learning paradigm to online reputation systems in which agents communicate over a social network. Despite the apparent

[18]Dasgupta *et al.* [106] also provide nice intuition in terms of intent polling and expectation polling. In intent polling, individual are sampled and asked who they intend to vote for. In expectation polling, individuals are sampled and asked who they think would win the election. For a given sample size, one would believe that expectation poling is more accurate than intent polling since in expectation polling, an individual would typically consider its own intent together with the intents of its friends.

simplicity in these information flows, the systems exhibit unusual behavior such as herding and data incest. Further, an example of social learning for change detection has been considered. Finally, we have discussed a non-Bayesian formulation, where agents seek to achieve coordination in decision-making by optimizing their own utility functions—this was formulated as a game-theoretic learning model.

The motivation for this paper stems from understanding how individuals interact in a social network and how simple local behavior can result in complex global behavior. The underlying tools used in this paper are widely used by the electrical engineering research community in the areas of signal processing, control, information theory, and network communications.

References

[1] N. Eagle and A. Pentland, "Reality mining: Sensing complex social systems," *Pers. Ubiquitous Comput.*, vol. 10, no. 4, pp. 255–268, 2006.
[2] C. C. Aggarwal and T. Abdelzaher, "Integrating sensors and social networks," in *Social Network Data Analytics*, C. C. Aggarwal, Ed., Springer-Verlag, 2011, ch. 14, pp. 397–412.
[3] A. Campbell, S. Eisenman, N. Lane, E. Miluzzo, R. Peterson, H. Lu et al., "The rise of people-centric sensing," *IEEE Internet Comput.*, vol. 12, no. 4, pp. 12–21, Jul./Aug. 2008.
[4] J. Burke, D. Estrin, M. Hansen, A. Parker, N. Ramanathan, S. Reddy, and M. Srivastava, "Participatory sensing," in *Proc. World Sensor Web Workshop ACM Sensys,* Boulder, CO, USA, 2006.
[5] A. Pentland, *Social Physics. How Good Ideas Spread–The Lessons from a New Science.* Baltimore, MD, USA: Penguin, 2014.
[6] A. Rosi, M. Mamei, F. Zambonelli, S. Dobson, G. Stevenson, and J. Ye, "Social sensors and pervasive services: Approaches and perspectives," in *Proc. 2011 IEEE Int. Conf. Pervasive Comput. Commun. Workshops* (PERCOM Workshops), IEEE, 2011, pp. 525–530.
[7] R. Lee and K. Sumiya, "Measuring geographical regularities of crowd behaviors for Twitter-based geo-social event detection," in *Proc. 2nd ACM SIGSPATIAL Int. Workshop Location Based Soc. Netw.*, ACM, 2010, pp. 1–10.
[8] Z. Cheng, J. Caverlee, and K. Lee, "You are where you tweet: A content-based approach to geo-locating Twitter users," in *Proc. 19th ACM Int. Conf. Inf. Knowl. Manag.*, ACM, 2010, pp. 759–768.
[9] M. Trusov, A. V. Bodapati, and R. E. Bucklin, "Determining influentual users in internet social networks," *J. Market. Res.*, vol. 47, pp. 643–658, Aug. 2010.
[10] T. Sakaki, M. Okazaki, and Y. Matsuo, "Earthquake shakes twitter users: Real-time event detection by social sensors," in *Proc. 19th Int. Conf. World Wide Web.* New York, NY, USA: ACM, 2010, pp. 851–860.
[11] J. Bollen, H. Mao, and X. Zeng, "Twitter mood predicts the stock market," *J. Comput. Sci.*, vol. 2, no. 1, pp. 1–8, 2011.
[12] B. Pang and L. Lee, "Opinion mining and sentiment analysis," *Found. Trends Inf. Retr.*, vol. 2, no. 1–2, pp. 1–135, 2008.
[13] S. Asur and B. A. Huberman, "Predicting the future with social media," in *Proc. IEEE/WIC/ACM Int. Conf. Web Intell. Intell. Agent Technol.* (WI-IAT), vol. 1, IEEE, 2010, pp. 492–499.
[14] B. Ifrach, C. Maglaras, and M. Scarsini, "Monopoly pricing in the presence of social learning," *NET Institute Working Paper No. 12-01,* 2011.
[15] M. Luca, "Reviews, reputation, and revenue: The case of yelp.com," Harvard Business School, Tech. Rep. 12-016, Sep. 2011.
[16] A. Banerjee, "A simple model of herd behavior," *Quaterly J. Econ.*, vol. 107, no. 3, pp. 797–817, Aug. 1992.
[17] S. Bikchandani, D. Hirshleifer, and I. Welch, "A theory of fads, fashion, custom, and cultural change as information cascades," *J. Political Econ.*, vol. 100, no. 5, pp. 992–1026, Oct. 1992.
[18] C. Chamley, *Rational herds: Economic Models of Social Learning.* Cambridge, U.K.: Cambridge Univ. Press, 2004.
[19] D. Acemoglu and A. Ozdaglar, "Opinion dynamics and learning in social networks," *Dynam Games Appl.*, vol. 1, no. 1, pp. 3–49, 2011.
[20] I. Lobel, D. Acemoglu, M. Dahleh, and A. Ozdaglar, "Preliminary results on social learning with partial observations," in *Proc. 2nd Int. Conf. Perform. Eval. Meth. Tools.* Nantes, France: ACM, 2007.
[21] D. Acemoglu, M. Dahleh, I. Lobel, and A. Ozdaglar, "Bayesian learning in social networks," *National Bureau of Economic Research*, Working Paper 14040, May 2008.
[22] T. Cover and M. Hellman, "The two-armed-bandit problem with time-invariant finite memory," *IEEE Trans. Inf. Theory*, vol. 16, no. 2, pp. 185–195, Mar. 1970.
[23] M. Hellman and T. Cover, "Learning with finite memory," *Ann. Math. Statist.*, vol. 41, no. 3, pp. 765–782, 1970.
[24] C. Chamley, A. Scaglione, and L. Li, "Models for the diffusion of beliefs in social networks: An overview," *IEEE Signal Proc. Mag.*, vol. 30, no. 3, pp. 16–29, May 2013.
[25] V. Krishnamurthy and H. V. Poor, "Social learning and Bayesian games in multiagent signal processing: How do local and global decision makers interact?" *IEEE Signal Process. Mag.*, vol. 30, no. 3, pp. 43–57, May 2013.
[26] M. Avellaneda and S. Stoikov, "High-frequency trading in a limit order book," *Quant. Finance*, vol. 8, no. 3, pp. 217–224, Apr. 2008.
[27] D. López-Pintado, "Contagion and coordination in random networks," *Int. J. Game Theory*, vol. 34, no. 3, pp. 371–381, 2006.
[28] D. López-Pintado, "Diffusion in complex social networks," *Games Econ. Behav.*, vol. 62, no. 2, pp. 573–590, 2008.
[29] M. Granovetter, "Threshold models of collective behavior," *Amer. J. Sociol.*, vol. 83, no. 6, pp. 1420–1443, May 1978.
[30] J. Goldenberg, B. Libai, and E. Muller, "Talk of the network: A complex systems look at the underlying process of word-of-mouth," *Market. Lett.*, vol. 12, no. 3, pp. 211–223, 2001.
[31] E. Mossel and S. Roch, "On the submodularity of influence in social networks," in *Proc. 39th Annu. ACM Symp. Theory Comput.*, ACM, 2007, pp. 128–134.
[32] N. Chen, "On the approximability of influence in social networks," *SIAM J. Discrete Math.*, vol. 23, no. 3, pp. 1400–1415, 2009.
[33] S. Bharathi, D. Kempe, and M. Salek, "Competitive influence maximization in social networks," in *Internet and Network Economics.* Springer, 2007, pp. 306–311.
[34] K. Apt and E. Markakis, "Diffusion in social networks with competing products," in *Algorithmic Game Theory.* Springer, 2011, pp. 212–223.
[35] R. J. Aumann, "Correlated equilibrium as an expression of Bayesian rationality," *Econometrica*, vol. 55, no. 1, pp. 1–18, 1987.
[36] F. Vega-Redondo, *Complex Social Networks*, vol. 44, Cambridge, U.K.: Cambridge Univ. Press, 2007.
[37] M. Jackson, *Social and Economic Networks.* Princeton, NJ, USA: Princeton Univ. Press, 2010.
[38] V. Krishnamurthy, M. Maskery, and G. Yin, "Decentralized activation in a ZigBee-enabled unattended ground sensor network: A correlated equilibrium game theoretic analysis," *IEEE Trans. Signal Process.*, vol. 56, no. 12, pp. 6086–6101, Dec. 2008.
[39] M. Maskery, V. Krishnamurthy, and Q. Zhao, "Decentralized dynamic spectrum access for cognitive radios: Cooperative design of a non-cooperative game," *IEEE Trans. Commun.*, vol. 57, no. 2, pp. 459–469, Feb. 2008.
[40] V. Krishnamurthy, "Decentralalized activation in dense sensor networks via global games," *IEEE Trans. Signal Process.*, vol. 56, no. 10, pp. 4936–4950, Oct. 2008.
[41] V. Krishnamurthy, "Bayesian sequential detection with phase-distributed change time and nonlinear penalty—A lattice programming POMDP approach," *IEEE Trans. Inf. Theory*, vol. 57, no. 3, pp. 7096–7124, Oct. 2011.
[42] 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.
[43] S. Hart, "Adaptive heuristics," *Econometrica*, vol. 73, no. 5, pp. 1401–1430, 2005.
[44] J. Predd, S. R. Kulkarni, and H. V. Poor, "A collaborative training algorithm for distributed learning," *IEEE Trans. Inf. Theory*, vol. 55, no. 4, pp. 1856–1871, Apr. 2009.
[45] G. Wang, S. R. Kulkarni, H. V. Poor, and D. Osherson, "Aggregating large sets of probabilistic forecasts by weighted coherent adjustment," *Decis. Anal.*, vol. 8, no. 2, pp. 128–144, Jun. 2011.
[46] G. Ellison and D. Fudenberg, "Rules of thumb for social learning," *J. Polit. Econ.*, vol. 101, no. 4, pp. 612–643, 1993.
[47] G. Ellison and D. Fudenberg, "Word-of-mouth communication and social learning," *Quarterly J. Econ.*, vol. 110, no. 1, pp. 93–125, 1995.
[48] K. C. Chen, M. Chuang, and H. V. Poor, "From technological networks to social networks," *IEEE J. Sel. Areas Commun.*, vol. 31, no. 9, pp. 548–572, Sep. 2013.
[49] L. Smith and P. Sorensen, "Informational herding and optimal experimentation," Economics Group, Nuffield College, Univ. Oxford, 1997, Economics Papers 139 [Online]. Available: http://ideas.repec.org/p/nuf/econwp/139.html
[50] Y. Ephraim and N. Merhav, "Hidden Markov processes," *IEEE Trans. Inf. Theory*, vol. 48, pp. 1518–1569, Jun. 2002.

[51] L. Mui, "Computational models of trust and reputation: Agents, evolutionary games, and social networks," Ph.D. dissertation, Dept. Elect. Eng. Comput. Sci., Massachusetts Institute of Technology, 2002.

[52] E. Gudes, N. Gal-Oz, and A. Grubshtein, "Methods for computing trust and reputation while preserving privacy," in *Data and Applications Security XXIII.* Springer, 2009, pp. 291–298.

[53] A. R. Cassandra, "Exact and approximate algorithms for partially observed Markov decision process," Ph.D. dissertation, Dept. Comput. Sci., Brown Univ., 1998.

[54] W. Lovejoy, "Some monotonicity results for partially observed Markov decision processes," *Oper. Res.*, vol. 35, no. 5, pp. 736–743, Sep./Oct. 1987.

[55] U. Rieder, "Structural results for partially observed control models," *Meth. Model. Oper. Res.*, vol. 35, no. 6, pp. 473–490, 1991.

[56] V. Krishnamurthy and D. Djonin, "Structured threshold policies for dynamic sensor scheduling: A partially observed Markov decision process approach," *IEEE Trans. Signal Process.*, vol. 55, no. 10, pp. 4938–4957, Oct. 2007.

[57] V. Krishnamurthy, "How to schedule measurements of a noisy Markov chain in decision making?" *IEEE Trans. Inf. Theory*, vol. 59, no. 7, pp. 4440–4461, Jul. 2013.

[58] W. Lovejoy, "On the convexity of policy regions in partially observed systems," *Oper. Res.*, vol. 35, no. 4, pp. 619–621, Jul./Aug. 1987.

[59] 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.

[60] B. Tuttle, "Fact-checking the crowds: How to get the most out of hotel-review sites," *Time Magazine,* Jul. 29, 2013.

[61] G. Adomavicius and A. Tuzhilin, "Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions," *IEEE Trans. Knowl. Data Eng.*, vol. 17, no. 6, pp. 734–749, Jun. 2005.

[62] I. Konstas, V. Stathopoulos, and J. M. Jose, "On social networks and collaborative recommendation," in *Proc. 32nd Int. ACM SIGIR Conf. Res. Dev. Inf. Retr.*, ACM, 2009, pp. 195–202.

[63] Y. Kanoria and O. Tamuz, "Tractable Bayesian social learning on trees," in *Proc. IEEE Int. Symp. Inf. Theory* (ISIT), Jul. 2012, pp. 2721–2725.

[64] J. Pearl, "Fusion, propagation, and structuring in belief networks," *Artif. Intell.*, vol. 29, no. 3, pp. 241–288, 1986.

[65] K. Murphy, Y. Weiss, and M. Jordan, "Loopy belief propagation for approximate inference: An empirical study," in *Proc. 15th Conf. Uncertainty Artif. Intell.,* 1999, pp. 467–475.

[66] 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.

[67] R. J. Aumann, "Agreeing to disagree," *Ann. Statist.*, vol. 4, no. 6, pp. 1236–1239, Nov. 1976.

[68] J. Geanakoplos and H. Polemarchakis, "We can't disagree forever," *J. Econ. Theory*, vol. 28, no. 1, pp. 192–200, 1982.

[69] V. Borkar and P. Varaiya, "Asymptotic agreement in distributed estimation," *IEEE Trans. Autom. Control*, vol. 27, no. 3, pp. 650–655, Jun. 1982.

[70] R. Bond, C. Fariss, J. Jones, A. Kramer, C. Marlow, J. Settle, and J. Fowler, "A 61-million-person experiment in social influence and political mobilization," *Nature*, vol. 489, pp. 295–298, Sep. 2012.

[71] M. Hamdi and V. Krishnamurthy, "Removal of Data Incest in Multi-agent Social Learning in Social Networks," *ArXiv e-prints,* Sep. 2013.

[72] M. Benaïm and J. Weibull, "Deterministic approximation of stochastic evolution in games," *Econometrica*, vol. 71, no. 3, pp. 873–903, 2003.

[73] V. Krishnamurthy and A. Aryan, "Quickest detection of market shocks in agent based models of the order book," in *Proc. 51st IEEE Conf. Decis. Control,* Maui, Hawaii, Dec. 2012.

[74] H. V. Poor and O. Hadjiliadis, *Quickest Detection.* Cambridge, U.K.: Cambridge Univ. Press, 2008.

[75] A. Shiryaev, "On optimum methods in quickest detection problems," *Theory Probab. Appl.*, vol. 8, no. 1, pp. 22–46, 1963.

[76] S. Hart and A. Mas-Colell, "A reinforcement procedure leading to correlated equilibrium," in *Economic Essays: A Festschrift for Werner Hildenbrand*, G. Debreu, W. Neuefeind, and W. Trockel, Eds. Springer, 2001, pp. 181–200.

[77] E. Lieberman, C. Hauert, and M. A. Nowak, "Evolutionary dynamics on graphs," *Nature*, vol. 433, no. 7023, pp. 312–316, Jan. 2005.

[78] J. Hofbauer and K. Sigmund, "Evolutionary game dynamics," *Bull. Amer. Math. Soc.*, vol. 40, no. 4, pp. 479–519, 2003.

[79] S. Hart and A. Mas-Colell, "A simple adaptive procedure leading to correlated equilibrium," *Econometrica*, vol. 68, no. 5, pp. 1127–1150, 2000.

[80] T. N. Cason and T. Sharma, "Recommended play and correlated equilibria: An experimental study," *Econ. Theory*, vol. 33, no. 1, pp. 11–27, Oct. 2007.

[81] J. Duffy and N. Feltovich, "Correlated equilibria, good and bad: An experimental study," *Int. Econ. Rev.*, vol. 51, no. 3, pp. 701–721, Aug. 2010.

[82] C. Griffin and A. Squicciarini, "Toward a game theoretic model of information release in social media with experimental results," in *Proc. IEEE Symp. Security Privacy Workshops,* San Francisco, CA, USA, May 2012, pp. 113–116.

[83] D. Fudenberg and D. Levine, *The Theory of Learning in Games.* Cambridge, MA, USA: MIT Press, 1999.

[84] D. Topkis, *Supermodularity and Complementarity.* Princeton, NJ, USA: Princeton Univ. Press, 1998.

[85] J. Hofbauer and W. Sandholm, "On the global convergence of stochastic fictitious play," *Econometrica*, vol. 70, no. 6, pp. 2265–2294, Nov. 2002.

[86] A. Blum and Y. Mansour, "From external to internal regret," *J. Mach. Learn. Res.*, vol. 8, pp. 1307–1324, 2007.

[87] O. Namvar, V. Krishnamurthy, and G. Yin, "Distributed tracking of correlated equilibria in regime switching noncooperative games," *IEEE Trans. Autom. Control*, vol. 58, no. 10, pp. 2435–2450, Oct. 2013.

[88] O. Namvar, V. Krishnamurthy, and G. Yin, "Distributed energy-aware diffusion least mean squares: Game-theoretic learning," *IEEE J. Sel. Topics Signal Process.*, vol. 7, no. 5, pp. 821–836, Oct. 2013.

[89] A. Drexl, "Scheduling of project networks by job assignment," *Manag. Sci.*, vol. 37, no. 12, pp. 1590–1602, Dec. 1991.

[90] P. Auer, N. Cesa-Bianchi, and P. Fischer, "Finite-time analysis of the multiarmed bandit problem," *Mach. Learn.*, vol. 47, no. 2–3, pp. 235–256, 2002.

[91] O. Namvar, V. Krishnamurthy, and G. Yin, "Diffusion based collaborative decision making in noncooperative social network games," in *Proc. 1st Global Conf. Signal Inf. Process.,* Austin, Dec. 2013.

[92] J. Surowiecki, *The Wisdom of Crowds.* New York. NY, USA: Anchor, 2005.

[93] C. Anderson and G. J. Kilduff, "Why do dominant personalities attain influence in face-to-face groups? The competence-signaling effects of trait dominance," *J. Pers. Soc. Psychol.*, vol. 96, no. 2, pp. 491–503, 2009.

[94] M. Ottaviani and P. Sørensen, "Information aggregation in debate: Who should speak first?" *J. Public Econ.*, vol. 81, no. 3, pp. 393–421, 2001.

[95] H. Carlsson and E. van Damme, "Global games and equilibrium selection," *Econometrica*, vol. 61, no. 5, pp. 989–1018, Sep. 1993.

[96] S. Morris and H. Shin, "Global games: Theory and applications," in *Advances in Economic Theory and Econometrics: Proceedings of Eight World Congress of the Econometric Society,* Cambridge, U.K.: Cambridge Univ. Press, 2000, pp. 56–114.

[97] L. Karp, I. Lee, and R. Mason, "A global game with strategic substitutes and complements," *Games Econ. Behav.*, vol. 60, pp. 155–175, 2007.

[98] G. Angeletos, C. Hellwig, and A. Pavan, "Dynamic global games of regime change: Learning, multiplicity, and the timing of attacks," *Econometrica*, vol. 75, no. 3, pp. 711–756, 2007.

[99] V. Krishnamurthy, "Decentralized spectrum access amongst cognitive radios-an interacting multivariate global gametheoretic approach," *IEEE Trans. Signal Process.*, vol. 57, no. 10, pp. 3999–4013, Oct. 2009.

[100] D. Easley and J. Kleinberg, *Networks, Crowds, and Markets: Reasoning About a Highly Connected World.* Cambridge, U.K.: Cambridge Univ. Press, 2010.

[101] E. Sun, I. Rosenn, C. Marlow, and T. M. Lento, "Gesundheit! modeling contagion through facebook news feed.," in *Proc. 3rd Int. AAAI Conf. Weblogs Soc. Media,* San Jose, CA, USA, May 2009.

[102] D. M. Romero, B. Meeder, and J. Kleinberg, "Differences in the mechanics of information diffusion across topics: Idioms, political hashtags, and complex contagion on Twitter," in *Proc. 20th Int. Conf. World Wide Web,* Hyderabad, India, Mar. 2011, pp. 695–704.

[103] R. Pastor-Satorras and A. Vespignani, "Epidemic spreading in scale-free networks," *Phys. Rev. Lett.*, vol. 86, no. 14, p. 3200, 2001.

[104] D. Aldous, "Interacting particle systems as stochastic social dynamics," *Bernoulli*, vol. 19, no. 4, pp. 1122–1149, 2013.

[105] J. Leskovec and C. Faloutsos, "Sampling from large graphs," in *Proc. 12th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining,* Philadelphia: ACM press, 2006, pp. 631–636.

[106] A. Dasgupta, R. Kumar, and D. Sivakumar, "Social sampling," in *Proc. 18th ACM SIGKDD Int. Conf. Knowl. discovery Data Mining,* Beijing, ACM, 2012, pp. 235–243.

[107] D. D. Heckathorn, "Respondent-driven sampling: A new approach to the study of hidden populations," *Soc. Probl.*, vol. 44, pp. 174–199, 1997.

[108] D. D. Heckathorn, "Respondent-driven sampling ii: Deriving valid population estimates from chain-referral samples of hidden populations," *Soc. Probl.*, vol. 49, pp. 11–34, 2002.

[109] S. Lee, "Understanding respondent driven sampling from a total survey error perspective," *Surv. Pract.*, vol. 2, no. 6, 2009.

[110] S. Goel and M. J. Salganik, "Respondent-driven sampling as Markov chain Monte Carlo," *Statist. Med.*, vol. 28, pp. 2209–2229, 2009.

[111] A. Lansky, A. Abdul-Quader, M. Cribbin, T. Hall, T. Finlayson, R. Garffin, L. S. Lin, and P. Sullivan, "Developing an HIV behavioral surveillance system for injecting drug users: The National HIV behavioral surveillance system," *Public Health Rep.*, vol. 122, no. S1, pp. 48–55, 2007.

Vikram Krishnamurthy (S'90–M'91–SM'99–F'05) received the Ph.D. degree from the Australian National University, Acton, ACT, Australian, in 1992. He is a Professor and Canada Research Chair at the Department of Electrical Engineering, University of British Columbia, Vancouver, Canada. His current research interests include statistical signal processing, computational game theory, and stochastic control in social networks.

Dr. Krishnamurthy served as distinguished Lecturer for the IEEE Signal Processing Society and Editor in Chief of *IEEE Journal Selected Topics in Signal Processing*. He received an honorary doctorate from KTH (Royal Institute of Technology), Sweden, in 2013.

H. Vincent Poor (S'72–M'77–SM'82–F'87) received the Ph.D. degree from the Princeton University, Princeton, NJ, USA, in 1977. He is Dean of Engineering and Applied Science at Princeton University, Princeton, NJ, USA, where he is also the Michael Henry Strater University Professor. His interests include statistical signal processing and information theory, with applications in several fields.

Dr. Poor is a member of the National Academy of Engineering, the National Academy of Sciences, and the Royal Academy of Engineering, U.K. Recent recognition includes the 2010IET Fleming Medal, the 2011 IEEE Sumner Award, the 2011 Society Award of IEEE SPS, and honorary doctorates from Aalborg University, the Hong Kong University of Science and Technology, and the University of Edinburgh, respectively.