

# Efficient Bayesian Learning in Social Networks with Gaussian Estimators

Elchanan Mossel\* and Omer Tamuz†

April 20, 2010

## Abstract

We propose a Bayesian model of iterative learning on social networks that is computationally tractable; the agents of this model are fully rational, and their calculations can be performed with modest computational resources for large networks. Furthermore, learning is efficient, in the sense that the process results in an information-theoretically optimal belief. This result extends Condorcet’s Jury Theorem to general social networks, preserving rationality, computational feasibility and efficient learning.

The model consists of a group of agents who belong to a social network, so that a pair of agents can observe each other’s actions only if they are neighbors. We assume that the network is connected and that the agents have full knowledge of the structure of the network, so that they know the members of the network and their social connections.

The agents try to estimate some state of the world  $S$  (say, the price of oil a year from today). Each agent has a private measurement: an independently acquired piece of information regarding  $S$ . This is modeled, for agent  $v$ , by a number  $S_v$  picked from a Gaussian distribution with mean  $S$  and standard deviation one. Accordingly, agent  $v$ ’s prior belief regarding  $S$  is a normal distribution with mean  $S_v$  and standard deviation one.

The agents start acting iteratively. At each iteration, each agent takes the optimal action given its current belief. This action reveals its mean estimate of  $S$  to its neighbors. Then, observing its neighbors’ actions, each agent updates its belief, using Bayes’ Law.

We show that this process is efficient: all the agents converge to the belief that they would have, had they access to all the private measurements. Additionally, and in contrast to other iterative Bayesian models on networks, it is computationally efficient, so that each agent’s calculation can be easily carried out.

## 1 Introduction

### 1.1 Background

The basic premise of mathematical behavioral economics is that human beings are rational. A natural model of rationality, when faced with uncertainty, is Bayesian inference together with a utility function that depends on the state of the world and the actions of the individuals. This model, from its first days, ran into the difficulties of computational intractability [11].

An obvious source of difficulty is the complexity of the world; it is infeasible to even represent a probability distribution over the possible states of our world. But even when the state of the world

---

\*Weizmann Institute and U.C. Berkeley. E-mail: mossel@stat.berkeley.edu. Supported by a Sloan fellowship in Mathematics, by BSF grant 2004105, by NSF Career Award (DMS 054829) by ONR award N00014-07-1-0506 and by ISF grant 1300/08

†Weizmann Institute. Supported by ISF grant 1300/08

is taken to be very simple - binary for example - the computational challenge an individual faces may still prove insurmountable, when *social networks* are taken into account [8].

Research in this field has been characterized by a tension between rationality and tractability: Rational models are in general intractable, and therefore unrealistic, whereas other (usually boundedly-rational) models are somewhat arbitrary.

For some network geometries the problem is easy. For example, when the network graph is a clique, Condorcet's Jury Theorem [5], a founding work in this field, states that when each individual receives a weak, independent, binary signal on the state of the world, the group can aggregate their information using Majority Rule to recover the true state of the world with high probability, if the group is large enough. More specifically, the probability of correct recovery goes to one as the number of agents goes to infinity, a property of the model known as *asymptotic learning*.

The history of *Social Learning on Networks* can be viewed as an attempt to find extensions of Condorcet's Jury Theorem, to setups where direct interaction is allowed only between some of the agents, so that the structure of social relationships is given by a connected network. An elusive goal has been finding a model that is rational, tractable and results in asymptotic learning, as Condorcet's does.

First models, such as the De-Groot model [6], consider iterative network processes, where each node performs the computationally simple task of averaging its current distribution for the state of the world with that of its neighbors. This leads to convergence of all agents to the same value, which is the average of the original beliefs, as weighted by the degree. This model is elegant, but leaves more to be desired; the nature of the utility maximized by the actions of the different agents is unclear, and so it is hard to see how this model can be classified as rational. Additionally, convergence to a weighted average may be considered as suboptimal, as the true average of the signals is a better estimate of the original signal under standard statistical assumptions. This is most apparent in networks where some nodes have degree which is proportional to the total number of connections in the network. In such networks, no matter how large they are, with constant probability the estimates will converge to values which are bounded away from the true ones.

A modern approach to the problem is in the Bayesian setup. Founding work in the Bayesian realm (e.g., [3], [4] and [12]) focused on chains of individuals, each of which, in turn and according to a set order, chooses an action based on private information and past actions. Similar models, in which agents have limited sets of social ties, have also been devised, e.g. [1]. Of course these models are somewhat limited in their modeling power, as most realistic network interactions are iterative, and so agents learn and act over and over again.

This problem is partially addressed in the work of Bala and Goyal [2], where iterative interactions are allowed. They have shown that for some priors and some network structures, the agents converge to an optimal action. However, their model assumes that in each round each agent may receive an independent signal of the state of the world, and thus a potentially unbounded amount of information may be available to an individual. Furthermore, they bound the rationality of the agents for various reasons, one of which is that fully rational agents would have to carry out intractable computations.

Gale and Kariv [8] propose a model in which each agent receives an independent signal at the beginning of the process, and thenceforth all agents act and observe simultaneously and repeatedly, improving their knowledge of the state of the world with every iteration. The model is completely Bayesian in the sense that the agents are Bayesian and their actions are aimed at maximizing the expected utility at each round. It requires the agents to know the structure of the network graph.

This is a natural model for studying the paradigm of the Condorcet Jury Theorem, in a network setup. It is shown in [8] that all the beliefs of the agents converge, and that it could not be

that different agents will converge to different actions. However, their results do not rule out the possibility that the actions will not converge at all. Additionally, their model, in its full generality, is computationally intractable.

## 1.2 Our Contribution

We consider a model which is a variant of this model. We insist that the action space is infinite and indeed that the action of each agent reveals the mean of its belief to its neighbors. We further assume that the original signals are Gaussian.

In this model the state of the world is  $S$ , a real number. The agents each have a private Gaussian measurement of  $S$ . Iteratively, they pick an action, and learn their neighbors' estimate for the expected value of  $S$  from *their* actions. Then they perform a fully Bayesian calculation which results in a new estimate of  $S$ .

We prove that this model satisfies the strongest possible Condorcet Jury Theorem:

- The agents are fully rational *and* the agents' calculations are tractable.
- For every connected network, all agents will converge to the same beliefs, and that this belief is optimal, i.e. the belief given by the average of the original signals.
- Finally, convergence takes place in a finite number of rounds - in fact the number of rounds is at most twice the number of agents in the network times the diameter of the network.

After the first draft of our paper was circulated [10] Marcus Möbius (whom we would like to thank) brought to our attention that essentially the same model is briefly mentioned in a paper by DeMarzo, Vayanos and Zweibel [7], who prove that convergence takes place after a number of rounds equal to the square of the number of agents. It appears that the main reason that [7] did not devote more attention to the model was computational, as they write: "We should emphasize that the calculations that agents must perform even in this simple case where the network is common knowledge can be very complicated.". In contrast, we show that the computations are efficient - both theoretically (each update involves linear algebra manipulations in dimension  $n$ ) and in practice (we have efficient code that performs all the agents calculations for large networks).

A second property of this model that DeMarzo et al. found unrealistic is the requirement that all agents know the structure of the social network. While indeed this may be difficult to justify for some large networks, *it is perhaps not strictly necessary*; in order to perform their calculations, the agents need to know the covariance between the estimators of *their neighbors only*. In our model, they derive this knowledge from the structure of the graph, but in principle it may be derived by other means. This observation (which was clarified in discussions with Rafael M. Frongillo, Grant Schoenebeck and Adam Kalai, whom we would like to thank) presents an opportunity for follow-up work involving some variant of this model, which does not require knowledge of the network, but still preserves rationality, tractability and efficient learning. More on this and other future research is presented at the conclusion section.

## 2 The Model

The model can be divided into three parts: the agents and their social network, the state of the world and its measurements, and the agents' behavior.

## 2.1 The Agents and their Social Network

The agents in our model are the nodes of a fixed network of social ties. Formally:

- The agents are a finite set  $V$ .
- The set of social ties  $E$  is a set of pairs of agents, so that  $\{u, v\} = \{v, u\}$  is in  $E$  if agents  $u$  and  $v$  are neighbors.
- Every pair of agents is connected by a chain of neighbors: even if  $u$  and  $v$  are not neighbors, then there exists a chain  $w_1, \dots, w_k$  such  $u = w_1$ ,  $v = w_k$ , and  $w_1$  is a neighbor of  $w_2$ ,  $w_2$  is a neighbor of  $w_3$  etc.

Hence  $G = (V, E)$  is a finite, simple, connected, undirected graph. Note that this graph does not change with time.

## 2.2 The State of the World and its Measurements

The agents reside in a world characterized by some number  $S$ . They have some information on  $S$  which is not certain.

- Let  $S \in \mathbb{R}$  be some state of the world.
- For each agent  $v$ , let  $S_v$  be  $v$ 's private measurement, so that  $S_v$  is picked from the normal distribution with mean  $S$  and standard deviation  $\sigma$ .
- The different  $S_v$ 's are independent.

## 2.3 The Agents' Behavior

The agents are Bayesian, so they initially have some prior belief regarding  $S$ , and update it to a posterior belief, according to Bayes' Law, with each additional piece of information they encounter. Both prior and posterior beliefs are distributions on the possible values of  $S$ .

In our model each agent  $v$  has a *different* prior, which is the Gaussian distribution with mean  $S_v$  and variance one. An equivalent model, in which all agents initially have *identical* priors, is the one in which the common ("improper") prior is the uniform measure on  $\mathbb{R}$ . After learning the private measurement, each agent would update its belief and would at that point have the prior belief of the agents in our model. Roughly speaking, this improper prior is well approximated by the normal distribution with some extremely large variance.

At each iteration, each agent picks an action  $A_x$ , where  $x$  is some real number. We assume that when  $x$  is equal to the expectation of an agent's current belief, then  $A_x$  is its optimal action. This can be achieved by, for example, setting the agents' utility to be  $U(A_x) = -(S - x)^2$  and assuming that the agents want to maximize their expected utility. We also assume that one can learn  $x$  by observing that an agent has picked  $A_x$ . That is  $A_x$  is different from  $A_y$  when  $x$  is different than  $y$ .

Having carried out some action, each agent observes its neighbors' actions, and calculates a new posterior distribution, based on all the information it has come across so far.

Formally:

- Agent  $v$ 's prior belief regarding  $S$  is the Gaussian distribution with mean  $S_v$  and standard deviation one. Denote this belief  $B_{v,0}$ .
- At time  $t \in \mathbb{N}$ , agent  $v$  takes action  $A_{x_v(t)}$ , where  $x_v(t) = \mathbf{E}[S|B_{v,t}]$ .

- Next, it observes its neighbors actions, and learns  $x_u(t)$  for each of its neighbors  $u$ . It then calculates, using Bayes' Law, a new posterior distribution  $B_{v,t+1}$ , based on all it has observed so far - its neighbors' actions in the previous iterations.

Note that the fact that each agent knows the structure of the graph allows its to know what calculation each of the other agents performed; not the actual numbers involved in the calculation, but rather “the formula” that was used.

This model is similar to the one presented in [9]. The agents in that model, however, were not Bayesian and had no memory of their observations in past iterations.

### 3 Results

We prove two results:

1. *Efficient Computation*: Each agent's calculation is computationally efficient: it can be achieved using simple linear algebra operations, involving matrices whose size is the size of the network.
2. *Efficient Learning*: The agents' posterior beliefs all converge to the same value. This is the value that they would have converged to, had they all access to each other's private measurements. Furthermore, the process converges in at most  $2n \cdot d$  iterations, where  $d$  is the diameter of the graph.

#### 3.1 Computational Efficiency

We begin with an informal discussion of the computational aspects of the learning process. The agents' private signals are distributed normally, with means  $S_v$  and standard deviations one. When an agent observes its neighbor's action, it learns the mean of its neighbor's current belief. This is an estimator of  $S$ , which can inductively be shown to also be normally distributed. It can also be shown to be a weighted average, and hence a linear combination, of the estimators seen so far, and hence a linear combination of the different  $S_v$ 's. Knowing the structure of the graph, an agent can know the coefficients of these linear combinations, coefficients that are independent of the actual values of the  $S_v$ 's.

When an agent observes a neighbor's action, it adds to its memory an additional estimator of  $S$ , and in particular one that is a linear combination of the original  $S_v$ 's. If this estimator is already in the space spanned by the estimators in the agent's memory, then the agent gains no new information. Otherwise, the agent increases the dimension of the space spanned by its memory.

In this multivariate Gaussian case, an agent's belief, at each iteration, is the unique linear combination, over a basis of the estimators in its memory, that minimizes its belief's variance while keeping it unbiased. This calculation involves inverting an  $n$  by  $n$  matrix, where  $n$  is the number of linearly independent estimators observed so far (so at most equal to  $n$ , the number of agents). This can be done very efficiently as the more rigorous analysis below establishes.

##### 3.1.1 The agents' calculation

In this subsection we take the alternative, but equivalent, point of view that the agents' prior is the improper uniform measure over the reals, and that their private signals are their first observation. It is easy to see that this is indeed equivalent to having the private signal as the prior.

Let  $\mathcal{W}$  be the vector space spanned by the different  $S_v$ 's:

$$\mathcal{W} = \left\{ \sum_{v \in V} \beta_v S_v \text{ s.t. } \forall v : \beta_v \in \mathbb{R} \right\}. \quad (1)$$

It is easy to convince oneself that this indeed is a vector space of finite dimension. Note that all the random variables in this space are normally distributed, since a linear combination of independent normal random variables is in turn normal too. Denote by  $\mathcal{W}^1$  the subset of unbiased estimators in  $\mathcal{W}$ :

$$\mathcal{W}^1 = \left\{ \sum_{v \in V} \beta_v S_v \in \mathcal{W} \text{ s.t. } \sum_{v \in V} \beta_v = 1 \right\}. \quad (2)$$

**Theorem 3.1.** *For all agents  $w$  and times  $t$ , it holds that  $x_w(t) \in \mathcal{W}^1$ , with  $x_w(t) = \sum_v \beta_{wv}(t) S_v$  for some  $\beta_{wv}(t)$ . Moreover, given the graph structure only, there exists an efficient algorithm to calculate the coefficients  $\beta_{wv}(t)$ .*

*Proof.* We shall prove this by induction on  $t$ . At time  $t = 0$  the claim is true since  $\beta_{wv}(0)$  is one when  $w = v$  and zero otherwise. Assume that the claim is true until time  $t$ .

Consider an agent  $w$ , and denote by  $r_0, \dots, r_k$  the random variables that agent  $w$  has observed up to time  $t$ , with  $r_0 = S_w = x_w(0)$ . Those are  $w$ 's own and its neighbors past estimators, and so knowledge of the graph structure is required here. By our assumption these are all in  $\mathcal{W}^1$ , and we can write  $r_i = \sum_v \alpha_{iv} S_v$ , where the coefficients  $\alpha_{iv}$  are a simple re-indexing of the coefficients  $\beta_{wv}(t)$ , by some relation that maps each  $w$  and  $t$  to some  $i$ . Since by assumption  $r_i \in \mathcal{W}^1$  then  $\sum_v \alpha_{iv} = 1$ .

Denote by  $\mathbf{r}$  the vector  $(r_0, \dots, r_n)$ , denote by  $\mathbf{1}$  the vector  $(1, \dots, 1) \in \mathbb{R}^n$ , and denote by  $C_{ij}$  the covariance between  $r_i$  and  $r_j$ , so that  $C_{ij} = \sum_v \alpha_{iv} \alpha_{jv}$ . Then  $\mathbf{r}$ 's distribution is the normal multivariate distribution with covariance matrix  $\mathbf{C}$  and mean  $\mathbf{1}S$  (since  $r_i \in \mathcal{W}^1$ ), and the likelihood of  $S$  given that agent  $w$  has observed  $\mathbf{r}$  is

$$\mathcal{L}(S|\mathbf{r}) = p(\mathbf{r}|S) = \frac{1}{(2\pi)^{n/2} |\mathbf{C}|^{1/2}} e^{-\frac{1}{2}(\mathbf{r}-\mathbf{1}S)'\mathbf{C}^{-1}(\mathbf{r}-\mathbf{1}S)}. \quad (3)$$

Note that in the case that  $\mathbf{C}$  is not invertible (equivalently,  $\mathbf{r}$  is not linearly independent) we remove from it (and correspondingly from  $\mathbf{r}$ ) a minimal set of columns and rows such that it becomes invertible. By corollary,  $\mathbf{C}$  is never larger than  $n \times n$ .

The expression  $(\mathbf{r} - \mathbf{1}S)'\mathbf{C}^{-1}(\mathbf{r} - \mathbf{1}S)$  can be rewritten as

$$\frac{1}{\mathbf{1}'\mathbf{C}^{-1}\mathbf{1}} \cdot \left( S - \frac{\mathbf{1}'\mathbf{C}^{-1}\mathbf{r}}{\mathbf{1}'\mathbf{C}^{-1}\mathbf{1}} \right)^2 + B$$

with  $B$  a normalization factor. Denote by  $\gamma$  the vector  $\frac{\mathbf{1}'\mathbf{C}^{-1}}{\mathbf{1}'\mathbf{C}^{-1}\mathbf{1}}$ . Note that  $\sum_i \gamma_i = 1$ .

To calculate its posterior distribution, agent  $w$  uses Bayes' Law. Then

$$\begin{aligned} p(S|\mathbf{r}) &= \frac{p(\mathbf{r}|S) \cdot p(S)}{p(\mathbf{r})} \\ &= \mathcal{L}(S|\mathbf{r}), \end{aligned}$$

because the prior distribution  $p(\cdot)$  is uniform. This can be written as

$$p(S|\mathbf{r}) = \frac{1}{\sqrt{2\pi\tau^2}} e^{-(S-x)^2/2\tau^2} \quad (4)$$

where

$$x = \frac{\mathbf{1}'\mathbf{C}^{-1}}{\mathbf{1}'\mathbf{C}^{-1}\mathbf{1}}\mathbf{r} = \sum_i \gamma_i r_i \quad (5)$$

and

$$\tau^2 = \frac{1}{\mathbf{1}'\mathbf{C}^{-1}\mathbf{1}}.$$

Note that  $x$  is a linear combination of the observations that  $w$  made up to time  $t$ . The expected value of this distribution (4) is  $x$ , and therefore  $x_w(t+1) = x$ . Then

$$x_w(t+1) = \sum_i \gamma_i \sum_v \alpha_{iv} S_v \quad (6)$$

and therefore

$$\beta_{wv}(t+1) = \sum_i \sum_v \gamma_i \alpha_{iv}. \quad (7)$$

Since  $\sum_i \gamma_i = 1$  and  $\sum_v \alpha_{iv} = 1$  then  $\sum_v \beta_{wv}(t+1) = 1$ . We have shown then that  $x_w(t+1) \in \mathcal{W}^1$ . We have also shown that to calculate  $\beta_{wv}(t+1)$ , given the coefficients at time  $t$ , one need only invert  $n$  matrices (one for each agent), of size at most  $n \times n$  - certainly an efficient calculation. Furthermore, no knowledge of the  $S_v$ 's is needed, but only of the graph structure.  $\square$

All the agents, if they know the graph structure, can perform this calculation efficiently. In particular an agent can calculate its coefficients vector  $\gamma$ , and from it calculate its next estimator for time  $t+1$ , given that it has performed the calculation above up to time  $t$ .

## 3.2 Learning Efficiency

### 3.2.1 Convergence in $n^2$

To show that the beliefs of the agents converge, we need only note that being conditional probabilities over increasingly large probability spaces, these beliefs are martingales. Then, because these martingales are bounded in  $L_2$ , they converge. However, the following proof, which does not require the power of martingales, shows that convergence in fact takes places in at most  $n^2$  iterations, and that furthermore all agents converge to the same belief. This proof is similar to the one presented by DeMarzo et al. [7].

When two neighboring agents have different beliefs, then at least one of them will learn from the other and improve its estimator: Assume agents  $u$  and  $v$  are neighbors with different estimators, and agent  $v$ 's belief has variance lower than or equal to that of agent  $u$ . Then agent  $v$ 's estimator is necessarily not in the space spanned by the estimators previously seen by  $u$ . Hence the dimension spanned by  $u$ 's memory will increase at this iteration. We have thus shown that in each iteration, unless all the agents have the same estimator, at least one of them increases the dimension of its space by at least one. Since the maximum dimension possible is  $n$  then convergence will occur after at most  $n^2$  steps, and all agents will converge to the same belief.

### 3.2.2 Convergence in $2n \cdot d$ iterations

A slightly more subtle argument proves a better bound for the convergence rate, namely  $2n \cdot d$ , where  $d$  is the diameter of the graph. The idea of the proof is that the current estimator of an agent  $u$  cannot remain unchanged for many steps, unless a growing neighborhood around  $u$  also remains stagnant. The formal proof uses the following lemma.

**Lemma 3.2.** *If some agent’s estimator has not changed for  $2d$  steps then the process has converged.*

*Proof.* Assume agent  $u$ ’s estimator does not change from iteration  $t_0$  to  $t_0 + 2d$ , so that

$$x_u(t_0) = x(t_0 + 1) = \dots = x(t_0 + 2d).$$

Denote  $x := x_u(t_0) = \dots = x(t_0 + 2d)$ , and let  $\mathcal{U}$  be the space spanned by the estimators in  $u$ ’s memory at time  $t_0 + 2d$ . Then by definition of the process  $x$  is the optimal unbiased estimator in  $\mathcal{U}$ .

Let  $w$  be a neighbor of  $u$ . Then  $w$ ’s estimator at time  $t_0 + 1$ ,  $x_w(t_0 + 1)$ , is in  $\mathcal{U}$ , since  $u$  observes  $x_w(t_0 + 1)$  at time  $t_0 + 2$ . Now  $x$  by definition is better than any estimator in  $\mathcal{U}$ , and so, since  $w$  has observed  $x$  at time  $t_0$ , it must be that  $x_w(t_0 + 1) = x$ . By the same argument  $x_w(t) = x$  for  $t_0 + 1 \leq t \leq t_0 + 2d - 1$ .

Applying this argument inductively, it follows that at time  $t_0 + i \leq t \leq t_0 + 2d - i$  all the agents at distance  $i$  from  $u$  have estimator  $x$ , and so at time  $t_0 + d$  all agents have the same estimator.

Recalling that at each iteration an agent’s estimator is a weighted average of those of its neighbors, we conclude that all nodes will have estimator  $x$  for all times  $t \geq t_0 + d$ . The proof follows.  $\square$

**Theorem 3.3.** *The process stops after  $2n \cdot d$  iterations.*

*Proof.* Every time an agent’s estimator changes, the dimension of the space the agent’s memory spans increases by at least one, and so this cannot happen more than  $n$  times. Since This must happen every  $2d$  steps as long as the process hasn’t converged, the process must stop after  $2n \cdot d$  iterations.  $\square$

### 3.2.3 Convergence to the Optimum

At any particular iteration, any node  $v$  contains  $S_v$  in the space of its estimators. At each iteration the estimator at  $v$  is then of the form  $aS_v + bS$  where  $S$  is an unbiased linear estimator based on some signals but  $S_v$ , and  $a + b = 1$ . Note that the variance of this estimator is  $a^2 + b^2Var(S)$  and it is minimized when  $a = Var(S)/(1 + Var(S))$ . Since  $S$  depends on all the signals but  $S_v$  its variance is at least  $1/(n - 1)$  and therefore  $a$  is at least  $1/n$ .

Hence all the agents, at all iterations, give their own estimators weight which is at least  $1/n$ . Since they all converge to the same estimator, and since the sum of the weights in this estimator must be one (since it, too, is unbiased), then the weights must all be  $1/n$ , and the limiting estimator is the simple average of the private measurements, as stated.

## 4 Conclusion and Future Work

In this work we presented a first learning model on social networks which is both rational and computationally efficient. Our work raises a number of future research direction which we briefly discuss.

### 4.1 Network Structure

A shortcoming of the model introduced here is the assumption that all agents have complete access to the network structure. Future work may relax this assumption in a number of ways. In particular:

### 4.1.1 Random Networks

In the Bayesian literature it is common to model the unknown network structure by a common prior distribution on network structures. One may model this by picking the network randomly from some distribution, which is known to the agents. The agents then proceed similarly, calculating conditional expectations Bayesianly, and incorporating into their posterior distributions whatever information they may have gathered about the structure of the network.

While there are many options for choosing such a distribution over networks, an interesting challenge is to see for which distributions the agents' calculations remain tractable. We leave this as an open question. We note only that the standard martingale arguments in the area [12] for convergence remains valid for any distribution of connected networks and that therefore, in any case, we have convergence in this scenario, too.

### 4.1.2 Learning Covariance Structure

A different approach is to study natural mechanisms by which neighbors learn the covariance between their neighbors' signals at different iterations. As we noted above, this knowledge is sufficient for calculating optimal estimators, even without any understanding of the graph structure. We have begun to explore such models, and their computational feasibility, together with Rafael M. Frongillo and Grant Schoenebeck.

## 4.2 Convergence Rate

The results established in this paper show convergence in  $O(n \cdot d)$ . A natural question is whether this bound can be improved. Certainly, convergence cannot happen faster than  $O(d)$  - the time it takes information to propagate through the network. For binary trees, where the diameter is  $O(\log n)$ , convergence does happen in  $O(d)$ , as it does for cliques and stars. However, simulations have led us to believe that convergence in general is not that fast, and requires - we conjecture -  $O(n)$  steps.

In our simulations we sampled a population of regular graphs of degree three and diameter  $\log n$ . The result almost always was convergence in time  $n/3$ , with every agent increasing the dimension of the space its memory spanned by three, at every iteration. This may hint that convergence time may, in some sense, be inversely proportional to the degrees of the graph vertices. We thus conclude with the following conjecture and open problem:

**Conjecture 4.1.** *For any graph the learning process converges in  $O(n)$  iterations.*

**Open Problem 4.2.** *Does the process converge in  $O(n/d^*)$  iterations for all graphs, where  $d^*$  is the minimal degree of the graph?*

## 5 Acknowledgments

We thank Shachar Kariv for introducing us to the literature on learning on networks and for fascinating discussions. Thanks go to Marcus Möbius for an interesting discussion and for directing us to the work of DeMarzo et al., after a draft of this paper had been circulated. We also thank Grant Schoenebeck, Rafael M. Frongillo and Adam Kalai for interesting discussions regarding follow-up work. Finally, we are indebted to Yaron Singer for much support and helpful suggestions on writing the results.

## References

- [1] D. Acemoglu, M. A. Dahleh, I. Lobel, and A. Ozdaglar. Bayesian learning in social networks. Preprint, 2008.
- [2] V. Bala and S. Goyal. Learning from neighbours. *Review of Economic Studies*, 65(3):595–621, July 1998.
- [3] A. V. Banerjee. A simple model of herd behavior. *The Quarterly Journal of Economics*, 107(3):797–817, 1992.
- [4] D. H. Bikhchandani, S. and I. Welch. A theory of fads, fashion, custom, and cultural change as informational cascade. *Journal of Political Economy*, 100(5):992–1026, 1992.
- [5] J.-A.-N. Condorcet. *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*. De l'Imprimerie Royale, 1785.
- [6] M. H. DeGroot. Reaching a consensus. *Journal of the American Statistical Association*, 69(345):118–121, 1974.
- [7] P. DeMarzo, D. Vayanos, and J. Zwiebel. Persuasion bias, social influence, and unidimensional opinions. *Quarterly Journal of Economics*, 118:909–968, 2003.
- [8] D. Gale and S. Kariv. Bayesian learning in social networks. *Games and Economic Behavior*, 45(2):329–346, November 2003.
- [9] E. Mossel and O. Tamuz. Iterative maximum likelihood on networks. To appear in the proceedings of Allerton, 2009.
- [10] E. Mossel and O. Tamuz. Efficient bayesian learning in social networks with gaussian estimators. Preprint at <http://arxiv.org/abs/1002.0747>, 2010.
- [11] H. Simon. *Reason in Human Affairs*. Stanford University Press, 1982.
- [12] L. Smith and P. Sorensen. Pathological outcomes of observational learning. *Econometrica*, 68(2):371–398, 2000.