Naïve Learning in Social Networks and the Wisdom of Crowds

Benjamin Golub   Matthew O. Jackson*

Abstract

We study learning in a setting where agents receive independent noisy signals about the true value of a variable and then communicate in a network. They naïvely update beliefs by repeatedly taking weighted averages of neighbors' opinions. We show that all opinions in a large society converge to the truth if and only if the influence of the most influential agent vanishes as the society grows. We also identify obstructions to this, including prominent groups, and provide structural conditions on the network ensuring efficient learning. Whether agents converge to the truth is unrelated to how quickly consensus is approached.

1. Introduction

Social networks are primary conduits of information and opinions. In view of this, it is important to understand how the evolution of beliefs and behaviors depends on network structure and whether or not the resulting outcomes are efficient. In this paper we examine one aspect of this broad theme: for which social network structures will a society of agents who communicate and update naïvely come to aggregate decentralized information completely and correctly?

Given the complex forms that social networks often take, it can be difficult for the agents involved (or even for a modeler with full knowledge of the network) to update beliefs properly. For example, Choi, Gale and Kariv [2] find that although subjects in simple three-person networks update fairly well in some circumstances, they do not do so well in evaluating repeated observations and judging indirect information whose origin is uncertain. Given that social communication often involves repeated transfers of information among large numbers of individuals in complex networks, fully rational learning becomes infeasible.

Nonetheless, it is possible that agents using fairly simple updating rules will arrive at outcomes like those achieved through fully rational learning. We discuss general obstacles to this, including the existence of prominent groups that receive a disproportionate share of attention. We also identify social networks for which naïve individuals converge to correct beliefs. These networks exhibit properties of balance and large-scale cohesion. Thus, whether naïve individuals can jointly learn the truth depends crucially on subtle properties of network structure, and we study the precise nature of this dependence.

2. The Model

2.1 Basic Communication

We begin by describing the DeGroot [3] model of belief updating in a fixed, finite network before discussing large societies. A finite set $N = \{1, 2, \ldots, n\}$ of *agents* or *nodes* interact according to a social network. The interaction patterns are captured through an $n \times n$ nonnegative matrix $\mathbf{T}$, where $T_{ij} > 0$ indicates that i pays attention to j. The matrix $\mathbf{T}$ may be asymmetric, and the interactions can be one-sided, so that $T_{ij} > 0$ while $T_{ji} = 0$. We refer to $\mathbf{T}$ as the *interaction matrix*. This matrix is stochastic, so that its entries across each row are normalized to sum to 1.

Agents update beliefs by repeatedly taking weighted averages of their neighbors' beliefs, with T_{ij} being the weight or trust that agent i places on the current belief of agent j in forming his or her belief for the next period. In particular, each agent has a belief $p_i^{(t)} \in \mathbb{R}$ at time $t \in \{0, 1, 2, \ldots\}$. The vector of beliefs at time t is written $\mathbf{p}^{(t)}$. The updating rule is:

$$\mathbf{p}^{(t)} = \mathbf{T}\mathbf{p}^{(t-1)}$$

and so

$$\mathbf{p}^{(t)} = \mathbf{T}^t \mathbf{p}^{(0)}.$$

A standard result from the theory of Markov chains states that if $\mathbf{T}$ is irreducible and aperiodic (the g.c.d. of

*Golub: Graduate School of Business, Stanford University, Stanford, CA, 94305-5015, bgolub@stanford.edu. Jackson: Department of Economics, Stanford University, Stanford, CA, 94305-6072, and the Santa Fe Institute, Santa Fe, NM 87501. We thank Francis Bloch, Antoni Calvó-Armengol, Drew Fudenberg, Tim Roughgarden, Jonathan Weinstein, and two anonymous referees for helpful comments and suggestions. Financial support under NSF grant SES-0647867 and the Thomas C. Hays Fellowship at Caltech is gratefully acknowledged. A full version of this paper may be found at http://www.stanford.edu/∼jacksonm/naivelearning.pdf.

all simple cycles in the directed graph induced by $\mathbf{T}$ is 1) then this process of belief updating will converge to a steady state. For simplicity in this extended abstract, we will assume that all matrices we discuss are strongly connected and aperiodic, even though our results hold in more general settings that do not assume strong connectedness.

2.2 Large Societies and Wisdom

In discussing convergence to true beliefs in large societies, we will be taking a double limit. First, for any fixed network, we ask what its beliefs converge to in the long run. Next, we study limits of these long-run beliefs as the networks grow; the second limit is taken across a sequence of networks.

Thus, consider a sequence of networks captured by a sequence of n-by-n interaction matrices: we say that a *society* is a sequence $(\mathbf{T}(n))_{n=1}^{\infty}$ indexed by n, the number of agents in each network. We will denote the (i,j) entry of interaction matrix n by $\mathbf{T}_{ij}(n)$, and, more generally, all scalars, vectors, and matrices associated to network n will be indicated by an argument n in parentheses.

Next, we specify initial beliefs. There is a true state of nature $\mu \in [0,1]$. At time $t = 0$, agent i in network n acquires an initial belief $p_i^{(0)}(n) \in [0,1]$. The belief is distributed with mean μ and a variance of at least $\sigma^2 > 0$, and the beliefs $p_1^{(0)}(n), \ldots, p_n^{(0)}(n)$ are independent for each n. No further assumptions are made about the joint distribution of the variables $p_i^{(0)}(n)$ as n and i range over their possible values.

Finally, we define the key concept of our study: *wisdom*. For any given n and realization of $\mathbf{p}^{(0)}(n)$, the belief of each agent i in network n approaches a limit which we denote by $p_i^{(\infty)}(n)$. Each of these limiting beliefs is a random variable that depends on the initial beliefs. We say the sequence of networks is *wise* when the limiting beliefs converge jointly in probability to the true state μ as $n \to \infty$.

Definition 1 *The sequence* $(\mathbf{T}(n))_{n=1}^{\infty}$ *is* wise *if,*

$$\operatorname*{plim}_{n\to\infty} \max_{i \le n} |p_i^{(\infty)}(n) - \mu| = 0.$$

3 Characterizations of Wisdom

We begin with a straightforward characterization of wisdom that is the workhorse for much of our analysis. To each interaction matrix $\mathbf{T}(n)$ corresponds a unique left eigenvector $\mathbf{s}(n)$ with eigenvalue 1, normalized so that its entries sum to 1. We call this the *influence* vector, and $s_i(n)$ is called the *influence* of node i. This eigenvector is related to measures of centrality and prestige that have been developed in sociology [4, 1] and also to the way that Google computes the importance of websites [5]. The eigenvector property just asserts that $s_j(n) = \sum_{i=1}^{n} T_{ij}(n) s_i(n)$ for all j, so that the influence of i is a weighted sum of the influences of various agents j who pay attention to i, with the weight of $s_j(n)$ being the trust of j for i. This is a very natural property for a measure of influence to have and entails that influential people are those who are trusted by other influential people. Without loss of generality, assume agents in each network are numbered so that $s_1(n)$ is the largest entry in $\mathbf{s}(n)$.

In our setting, it follows from standard Markov chain results that, for every i,

$$p_i^{(\infty)}(n) = \sum_{j=1}^{n} s_j(n) p_j^{(0)}(n).$$

This immediately implies the following characterization of wisdom. The proof follows the standard technique for proving weak laws of large numbers using the Chebyshev inequality.

Proposition 1 *The sequence* $(\mathbf{T}(n))_{n=1}^{\infty}$ *is wise if and only if* $s_1(n) \to 0$ *as* $n \to \infty$*.*

The result has a very simple intuition: for the idiosyncratic errors to wash out and for the limiting beliefs – which are weighted averages of initial beliefs – to converge to the truth, nobody's idiosyncratic error should be getting positive weight in the large-society limit.

This characterization is, however, rather abstract in that it applies to influence vectors and not directly to the structure of the social network. It is interesting to see how wisdom is determined by the geometry of the network. Which structures prevent wisdom, and which ones ensure it?

In one simple special case we can give a complete structural characterization in terms of quite familiar quantities – the degrees of various agents.

Proposition 2 *Assume all links are bilateral and that each agent weights all his neighbors equally. In this case, the sequence* $(\mathbf{T}(n))_{n=1}^{\infty}$ *is wise if and only if the maximum degree in network* n *divided by the sum of degrees in network* n *converges to* 0 *as* $n \to \infty$*.*

To state a more general result, we consider an obstacle to wisdom in arbitrary networks: namely, the existence of prominent groups which receive a disproportionate share of attention.

To introduce this concept, we need some definitions and notation. Define

$$T_{B,C} = \sum_{\substack{i \in B \\ j \in C}} T_{ij}$$

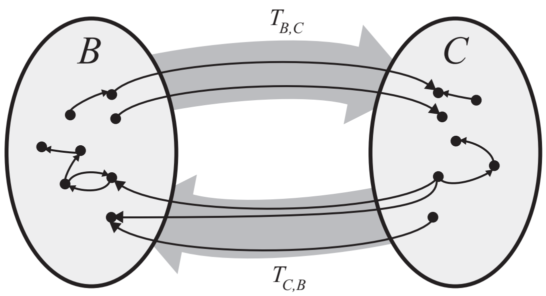

Figure 1: The large arrows illustrate the concept of the weight of one group on another.

which is the weight that the group B places on the group C (see Figure 1). We say that a set of agents B is *prominent in t steps* relative to $\mathbf{T}$ if $(\mathbf{T}^t)_{i,B} > 0$ for each $i \notin B$. Call $\pi_B(\mathbf{T};t) := \min_{i \notin B} (\mathbf{T}^t)_{i,B}$ the *t-step prominence* of B relative to $\mathbf{T}$. Next, we define a *family* to be a sequence of sets (B_n) such that $B_n \subset \{1, \dots, n\}$ for each n. With these notions in hand, we can define a uniformly prominent family.

Definition 2 *The family (B_n) is* uniformly prominent *relative to $(\mathbf{T}(n))_{n=1}^{\infty}$ if there exists a constant $\alpha > 0$ such that for each n there is a t so that the group B_n is prominent in t steps relative to $\mathbf{T}(n)$ with $\pi_{B_n}(\mathbf{T}(n);t) \geq \alpha$.*

We also define a notion of finiteness for families: a family is finite if it stops growing eventually, i.e. if there is a q such that $\sup_n |B_n| \leq q$.

We can now state a general structural characterization of wisdom.

Proposition 3 *The sequence $(\mathbf{T}(n))$ is wise if and only if there does not exist a finite, uniformly prominent family relative to it.*

While this result – especially the direction that says finite, uniformly prominent families destroy wisdom – is useful, it still leaves something to be desired. In particular, it requires considering possibly large powers of the interaction matrices, which takes us far from basic intuitions about how the network looks. We thus formulate sufficient conditions for wisdom that work with the basic network and not its powers. The first one is as follows.

Property 1 (Balance) *There exists a sequence $j(n) \to \infty$ such that if $|B_n| \leq j(n)$ then*

$$\sup_n \frac{T_{B_n^c,B_n}(n)}{T_{B_n,B_n^c}(n)} < \infty.$$

The balance condition says that no family below a certain size limit captured by $j(n)$ can be getting infinitely more weight from the remaining agents than it gives to the remaining agents. The sequence $j(n)$ can grow very slowly, which makes the condition reasonably weak.

In addition to imbalances of trust, one also has to worry about large-scale asymmetries of a different sort, which can be viewed as small groups focusing their attention too narrowly. The next condition deals with this.

Property 2 (Minimal Out-Dispersion) *There is a $q \in \mathbb{N}$ and $r > 0$ such that if B_n is finite, $|B_n| \geq q$, and $|C_n|/n \to 1$, then $T_{B_n,C_n}(n) > r$ for all large enough n.*

The minimal out-dispersion condition requires that any large enough finite family must give at least some minimal weight to any family which makes up almost all of society. Having stated these two conditions, we can give our final result, which states that the conditions are sufficient for wisdom.

Proposition 4 *If $(\mathbf{T}(n))_{n=1}^{\infty}$ is a sequence of convergent stochastic matrices satisfying balance and minimal out-dispersion, then it is wise.*

Finally, we can consider the relationship between the speed of learning and its accuracy. Simple examples show that, in general, these two quantities are not necessarily related. For each possible specification of the speed of learning and its accuracy, there exists a society which manifests that combination.

References

[1] Phillip B. Bonachich. Power and centrality: A family of measures. *American Journal of Sociology*, 92:1170–1182, 1987.

[2] Syngjoo Choi, Douglas Gale, and Shachar Kariv. Social learning in networks: A quantal response equilibrium analysis of experimental data. *Journal of Economic Theory*, 143(1):302–330, December 2008.

[3] Morris H. DeGroot. Reaching a consensus. *Journal of the American Statistical Association*, 69(345):118–121, 1974.

[4] Leo Katz. A new status index derived from sociometric analysis. *Psychometrika*, 18:39–43, 1953.

[5] Amy N. Langville and Carl D. Meyer. *Google's PageRank and Beyond: The Science of Search Engine Rankings.* Princeton University Press, Princeton, N.J., 2006.