Debiasing Social Wisdom

Abhimanyu Das
Microsoft Research
Silicon Valley
abhidas@microsoft.com

Sreenivas Gollapudi
Microsoft Research
Silicon Valley
sreenig@microsoft.com

Rina Panigrahy
Microsoft Research
Silicon Valley
rina@microsoft.com

Mahyar Salek
Microsoft Research
Cambridge
mahyar@microsoft.com

ABSTRACT

With the explosive growth of social networks, many applications are increasingly harnessing the pulse of online crowds for a variety of tasks such as marketing, advertising, and opinion mining. An important example is the wisdom of crowd effect that has been well studied for such tasks when the crowd is non-interacting. However, these studies don't explicitly address the network effects in social networks. A key difference in this setting is the presence of social influences that arise from these interactions and can undermine the wisdom of the crowd [17].

Using a natural model of opinion formation, we analyze the effect of these interactions on an individual's opinion and estimate her propensity to conform. We then propose efficient sampling algorithms incorporating these conformity values to arrive at a debiased estimate of the wisdom of a crowd. We analyze the trade-off between the sample size and estimation error and validate our algorithms using both real data obtained from online user experiments and synthetic data.

Categories and Subject Descriptors

H.2.8 [**Database Management**]: Database Applications - Data Mining

General Terms

Algorithms, Experimentation

Keywords

Wisdom of crowd, Social Networks, Opinion Formation

1. INTRODUCTION

The "wisdom of crowd" effect refers to the phenomenon in which the average opinion of a diverse group of individuals is often closer to the truth than the opinion of any single member of the group [10]. The wisdom of the crowd is increasingly finding use in a plethora of contexts ranging from the traditional online surveys to query predicates involving human computation in database applications [20, 19]. In many of these applications, the underlying assumption is that the crowd does not interact and that individuals in the crowd form their opinions independently. With the explosive growth of social networks and online question-answer websites [1], these platforms are becoming good sources for harnessing the collective opinions of online users. However, one significant difference in this new setting is the interacting crowd wherein the user can interact with her neighborhood to arrive at an opinion which might not necessarily reflect her original opinion. Crowdsourcing applications that rely on getting an unbiased opinion from the user will not work effectively in this setting. Lorenz et. al. [17] study how social influence can undermine the wisdom of a crowd for a variety of estimation tasks.

The presence of interactions between a group of online users brings up two important problems that we address in this paper.

The first is that of characterizing the effect of these interactions on the dynamics of opinion formation in online social networks or ad hoc settings such as social crowds. In a social setting, a user can be associated with both an innate opinion and an expressed opinion [5, 11, 3, 22] for any given topic or question. Her innate opinion is typically formed independent of online social interactions, while her expressed opinion could be shaped by the opinion of her online neighbors. This shaping of her expressed opinion is attributed to her *propensity to conform*. There is a line of impressive work that studied the consensus and fragmentation of the expressed opinions at steady state using different models of opinion formation [9, 13, 3, 2]. In this study, we adopt the model due to Friedkin and Johnsen [9], wherein the effect of social influence is captured by a user's propensity to conform to her neighborhood's opinion. However, our focus is fundamentally different from [9] in that we rely the underlying opinion formation dynamics to extract the latent innate opinions from the network.

The second problem is that of factoring out the effect of social influence when estimating the collective wisdom of a crowd. In the presence of social interactions, this wisdom now corresponds to the average innate opinion of the crowd

---

[1] http://www.quora.com

(as opposed to the average expressed opinion after interactions). One important point to note is that the innate opinions of people are not known in general, and the knowledge of the users' propensity to conform is either incomplete or noisy. This brings us to the two questions that we need to address for estimating the average innate opinions in the social network. First, we need efficient sampling algorithms that can obtain good estimates of the true average innate opinion using a small number of samples. Second, and more importantly, since we can only work with the expressed opinion of the nodes and cannot directly observe their innate opinion, our algorithms need to take care of debiasing the expressed opinions of the nodes that they sample.

In this paper, we study the above two problems using both an analytical opinion formation model that accounts for social influence, and a real-world experimental setup based on online surveys among participants that can interact with each other. While these social interactions in our experiments are limited solely to a user being able to share her opinions with a random set of other online users, we show that even these relatively simple interactions can cause users to move away from their original innate opinions for a given question.

Contributions of this study

We make the following contributions in this work

- We analyze a natural model of opinion formation in the context of extracting the wisdom of a social crowd. We show analytically that this model converges to a unique equilibrium under weaker assumptions (Section 3).

- We show how to debias the effect of social influence on the overall wisdom of the crowd which we treat to be the average innate opinion in the crowd. In particular, we design a near-optimal sampling algorithm (Section 4) that estimates the true average innate opinion of the crowd, and study the estimation error performance of the algorithm, both analytically and using experiments on both real and synthetic data.

- We study the effect of social influence in real-world experiments, and empirically validate the opinion formation model. In particular, we ask users from Amazon Mechanical Turk to complete a set of online surveys with two different question topics, where we allow limited interactions of users with each other. We observe the opinion formation process starting from users' initial answers to their final expressed answers after interactions, and compute their propensity to conform. We show that this value is indeed largely specific to a given user and does not change significantly for the different questions that the user answers in her survey (Section 5).

2. RELATED WORK

The findings suggesting that the collective opinion of a group is as good as, and often better than, the answer of an individual to a question are well established. Over a century ago, Galton [10] observed that the median of eight hundred participants' responses to a weight guessing contest was accurate within 1% of the true answer. Surowiecki [23] surveys numerous case studies and experiments from stock markets, political elections, and quiz shows supporting the

above statement, highlighting that the independence of individuals' opinions is a key requirement to form a wise crowd – an assumption undermined by the very nature of social networks [12, 21].

Lazer and Friedman [16] take an agent-based computer simulation approach to argue for a tradeoff between diversity and information flow, showing that the connectedness (in moderation) may improve the social wisdom and performance in several contexts. Lorenz et al. [17] demonstrate through human-subject lab experiments that the connectedness may harm more than benefit by diminishing the diversity and also by false-boosting the confidence of the crowd.

Our problem somewhat relates to consensus formation [14] in the sense that each node aggregates the opinion of her neighborhood. One notable example in this setting is the work by DeGroot [5] which studies how consensus is formed and reached when individual opinions are updated using the average of the neighborhood of a fixed network. Work of Friedkin and Johnsen (FJ) [9], is perhaps the first study to extend the DeGroot model [5] to include both disagreement and consensus by associating with each node an innate opinion in addition to her expressed opinion. In their model, they propose a certain degree g_i with which a user adheres to her initial opinion and by a *susceptibility* of $1 - g_i$ is socially influenced by others in her network. French [8] used a similar model to empirically estimate the susceptibility values.

Budgeted actions and inquiries over social networks have been studied before in the context of influence maximization [7, 15], vaccination [6] and expectation polling [4]. This theme is increasingly motivated by the explosive growth of social networks and the inhibitive cost of covering their members. Our work differs from [4] in the sense that we do not assume that the individuals may explicitly provide an answer to an expectation poll, yet their expressed opinion reflects what they observe in their neighborhood as well as their innate opinion. Moreover, every individual - to express an opinion - take all her neighbors' expressed opinions who take all their own neighbors expressed opinion and so on.

3. OPINION MODEL

We consider an online social network graph $G = (V, E)$ with nodes $\{v_1, v_2, \ldots, v_n\} \in V$. For ease of notation, we will frequently interchange node v_i and index i. The nodes v_i correspond to individuals, and edges $E = [e_{ij}]$, denote social interactions between the individuals. For a node v_i, its set of neighbors is denoted by $N(i) = \{j : e_{ij} = 1\}$, and its degree is denoted by $d_i = |N_i|$. We denote the total number of edges in the graph as $|E|$. We define $D \in \mathbb{R}^{n \times n}$ to be $\text{Diag}(d_1, d_2, \ldots, d_n)$.

We wish to estimate the average wisdom or opinion held in the social network about a particular topic or question of interest. As mentioned earlier, due to social interactions between the individuals in the social network, an opinion stated by an individual might not be the same as the original or true opinion held by the individual, and is often influenced by the expressed opinions in the individual's social neighborhood. All that we are likely to observe in the social network at any instant of time is therefore the stated or expressed opinions held by individuals at the time, and not their original innate opinions. We refer to node v_i's expressed opinion at time t as $Y_i(t)$, and its innate opinion as $Y_i(0)$ (which is assumed to be the same as the initial expressed opinion of the node at time 0). We assume that an opinion is expressed

(or encoded) by a single real quantity, hence $Y_i(t) \in \mathbb{R}$ for all i and t. For each individual v_i, we define a conformity parameter α_i, $0 \leq \alpha_i \leq 1$, which is a measure of how strong her innate opinions are, and how likely will she be influenced by her neighborhood opinions. An α_i value close to 1 implies that the individual is highly opinionated, and her expressed opinion is similar to her innate opinion, while a value close to 0 implies that the individual has a very weak innate opinion and consequently her expressed opinion is largely governed by the opinions of neighbors around her. Thus, the value $1 - \alpha_i$ represents agent i's propensity to conform. We define $\alpha_D \in \mathbb{R}^{n \times n}$ to be $\mathrm{Diag}(\alpha_1, \alpha_2, \ldots, \alpha_n)$.

We analyze a natural model for opinion formation due to Friedkin and Johnsen [9] (under weaker assumptions than in [9]) and show existence and convergence to a unique equilibrium.

In this model, individuals update their expressed opinion in discrete time steps by taking a convex combination of their innate opinion and the expressed opinions of their neighbors. As mentioned earlier, the weights in the convex combination depends on a user's α value. For simplicity, we assume that the individuals don't distinguish between their neighbors and take their opinions equally important. Equation 1 captures the above model for all individuals i.

$$Y_i(t+1) = \alpha_i Y_i(0) + (1 - \alpha_i) \frac{\sum_{j \in N_i} Y_j(t)}{d_i} \qquad (1)$$

We show that the above opinion formation model defines a unique equilibrium as long as all α's are non zero: that is, individuals hold an innate opinion that has some impact on what they express.

Lemma 1 *The above model has a unique equilibrium if $\alpha_i > 0$ for all i.*

Proof. Consider any equilibrium $\mathbf{Y}^* = \{Y_1^*, \ldots, Y_n^*\}^T$ for Equation 1. Then, for $i = 1, 2, \ldots, n$, $\mathbf{Y}^*$ satisfies

$$Y_i^* = \alpha_i Y_i(0) + (1 - \alpha_i) \frac{\sum_{j \in N_i} Y_j^*}{d_i}.$$

Or,

$$\mathbf{Y}^* = \alpha_D \mathbf{Y}^0 + M \mathbf{Y}^*,$$

where $\mathbf{Y}^0 = \{Y_1(0), Y_2(0), \ldots, Y_n(0)\}^T$, and

$$M = [m_{ij}] = \begin{cases} 0 & \text{if } i = j \\ \frac{1 - \alpha_i}{d_i} e_{ij} & \text{otherwise} \end{cases}$$

Thus, we have

$$\mathbf{Y}^* = (I - M)^{-1} \alpha_D \mathbf{Y}^0, \qquad (2)$$

where I is the $n \times n$ identity matrix.

We show that the matrix $I - M \in \mathbb{R}^{n \times n}$ is non-singular. For any row i of $I - M$, the sum of absolute values of its non diagonal elements is $\sum_{j \neq i} \frac{1 - \alpha_i}{d_i} e_{ij} = 1 - \alpha_i$, and its diagonal element is 1. Thus, using the Gersgorin Disc Theorem, every eigenvalue λ of $I - M$ lies within one of the discs $\{z : |z - 1| \leq 1 - \alpha_i\}$ for $i = 1, 2, \ldots n$. Since $\alpha_i > 0$, the eigenvalues of $I - M$ cannot include 0. Thus, $I - M$ is invertible, and the equilibrium is unique. ∎

Hence $\mathbf{Y}^* = (I - M)^{-1} \alpha_D \mathbf{Y}^0$ is the unique equilibrium of the opinion formation model defined in Equation 1. We now

prove that after sufficient iterations, the model will indeed converge to this equilibrium.

Lemma 2 *If $\alpha_i > 0$ for all i, then the above opinion formation model converges to its unique equilibrium $\mathbf{Y}^*$.*

Proof. Note that $\mathbf{Y}_{k+1}$, the social opinion state at time $k+1$ can be formulated by the following system:

$$\mathbf{Y}_{k+1} = \alpha_D \mathbf{Y}^0 + M \mathbf{Y}_k$$

where M is the iteration Matrix. Define ϵ_k so that $\mathbf{Y}_k = \mathbf{Y}^* + \epsilon_k$. By definition $\epsilon_{k+1} = M \epsilon_k$. Again we use the property that the sum of all the entries of each row in M is a non-negative quantity smaller than one to show that the error term goes to zero as k grows. Let $\epsilon_k^{\max}$ be the largest coordinate in ϵ_k. Hence, it is sufficient to show that $\epsilon_{k+1}^{\max} < \epsilon_k^{\max}$. Note that

$$\begin{aligned} \epsilon_{k+1}^{\max} &= \frac{1 - \alpha_i}{d_i} \sum_j \epsilon_k^{(j)} \leq \frac{1 - \alpha_i}{d_i} \sum_j \epsilon_k^{\max} \\ &\leq (1 - \alpha_i) \epsilon_k^{\max} < \epsilon_k^{\max} \end{aligned}$$

for some node i as long as $\alpha_j > 0$ for all j. ∎

4. ESTIMATING THE AVERAGE INNATE OPINION

Our goal is to estimate the average innate opinions in the social network (denoted by $\bar{Y^0} = \sum_{i=0}^n Y_i(0)/n$), by factoring out the social influences in the expressed opinions of the social network users. There are several reasons why this estimation of the average innate opinion of the social network is more important than the average expressed opinions. First, for many opinion polls and surveys, pollsters care about the true opinions held by an individual which might be quite different from their expressed opinions. Second, as shown in [17], the wisdom-of-crowd effect can break down when using individual opinions that are not independent and are influenced by social interactions. Hence it is important to use the original innate opinions of individuals when estimating answers using the wisdom-of-crowds phenomenon.

As mentioned earlier, there are two problems that we need to tackle when estimating the average innate opinions in the social network. Firstly, we need efficient sampling algorithms that can obtain good estimates of the true average innate opinion using a small number of samples. More importantly, since we can only work with the expressed opinions and cannot directly observe the innate opinions of individuals, we need to take care of debiasing the expressed opinions of the sampled nodes.

4.1 Sampling Algorithms

We now describe three sampling algorithms that might be used to estimate the average innate opinion in the network.

The simplest and perhaps the most prevalent sampling method is *uniform sampling*, in which a sampling budget is decided and each node is sampled with a uniform probability $\frac{1}{n}$ (with replacement) until the budget is exhausted. This naive algorithm is oblivious of any differences between innate and expressed opinions in the social network.

With the goal of eliciting the innate opinion and assuming that we have access to the α values, one may prefer to sample the nodes with large values of α as they are expected to retain their innate opinion in their expressed opinion to a

---
Algorithm 1 UNIFORMSAMPLING

1: Choose a random sample $S \subseteq V$ of size r with replacement by sampling each node i with probability p_i
2: Output $\hat{Y}^0 = \frac{1}{r} \sum_{i \in S} Y_i(t)$.
---

large extent. We call this *conformity sampling.* and formally define it as follows:

Algorithm 2 CONFORMITYSAMPLING

1: Choose a random sample $S \subseteq V$ of size r with replacement by sampling each node i with probability $p_i = \frac{\alpha_i}{\sum_{j=1}^n \alpha_j}$
2: Output $\hat{Y}^0 = \frac{1}{nr} \sum_{i \in S} \frac{1}{p_i} Y_i(t)$.

We use both sampling criteria as baselines and propose a new algorithm that outperforms the baselines both theoretically and empirically.

4.1.1   Influence Sampling

Our approach is reminiscent of the social sampling ([4]) algorithms where, instead of only using the opinions of the sampled nodes to estimate the true average opinion, the algorithm uses the opinions of the neighborhood of each sampled node via expectation polling. This might however arguably be a strong assumption, since individuals are expected to accurately report their neighborhood's average opinion. Furthermore, the authors do not distinguish between innate and expressed opinions in the social network. As described in [4], the authors analyze sampling algorithms that sample each node v_i with probability proportional to its degree d_i to obtain a sample set S of size r. Assuming that the opinion of each node v_i is Y_i and the goal is to estimate the average $\bar{Y} = \sum_{i=0}^n Y_i/n$, the authors then propose using an estimator $\hat{Y} = \frac{1}{nr} \sum_{i \in S} \frac{2|E|}{d_i} \sum_{j \in N(i)} \frac{Y_j}{d_j}$ and obtain a probabilistic bound for the estimation error $|\hat{Y} - \bar{Y}|$, given a sample size r.

In our case, the goal is to estimate the average innate opinion $\bar{Y}^0 = \sum_{i=0}^n Y_i(0)/n$ by sampling a set S; however we do not have access to the $Y_i(0)$ of the sampled set S or its neighbors, and can only observe the $Y_i(t)$ values instead, at time t. If we knew the $Y_i(0)$ values, the problem would degenerate to simple uniform sampling of the nodes, for which the following result holds.

Fact 3 *If we choose* $r = \frac{1}{\epsilon^2} \log(1/\delta)$ *samples, then with probability* $1 - \delta$, *a uniform sampling strategy will give an estimate* $\hat{Y}^0$, *such that* $|\hat{Y}^0 - \bar{Y}^0| \leq \epsilon$

In our case however, since we have estimates for the α_i values at all nodes in the graph, we would like to incorporate these to guide our sampling strategy and our choice of estimator function $\hat{Y}^0$. Note that since solving Equation 2 to obtain the Y_i^0 values requires explicit knowledge of all the Y_i^* in the network, we cannot directly use the equation to compute the $Y_i(0)$ values from the observed $Y_i(t)$.

Our approach therefore is to use the $Y_i(t)$ values directly in our estimator function $\hat{Y}^0$, and to construct $\hat{Y}^0$ such that the $Y_i(t)$ of the sampled nodes and its neighbors lead to good approximations for the corresponding $Y_i(0)$ values. We will

then carefully choose our sampling strategy in a way that tends to minimize these approximation errors.

In order to derive an "optimal" sampling algorithm we rewind the opinion formation process as shown in the following lemma. The following lemma shows how to get the average innate opinion from the expressed opinions

Lemma 4 $\bar{Y}^0 = \sum_{i=0}^n c_i Y_i^*$ *where*
$$c_i = \frac{1 - \sum_{j \in N(i)}(1 - \alpha_j)/d_j}{n\alpha_i}$$

Proof.   Using Equation 2, we have $\mathbf{Y}^0 = \alpha_D^{-1}(I - M)\mathbf{Y}^*$. Thus
$$\bar{Y}^0 = \frac{\sum_{i=0}^n Y_i^0}{n} = \frac{1}{n}\mathbf{1}\mathbf{Y}^0 = \frac{1}{n}\alpha_D^{-1}(I - M)\mathbf{Y}^*$$
Expanding the right hand side of the last step into a sum form completes the proof ∎

Note that the c_i above might be negative. Guided by the above lemma, we now define the sampling probabilities p_i of our INFLUENCESAMPLING algorithm as

$$p_i = \frac{|c_i|}{\sum_{j=1}^n |c_j|} = \frac{|\frac{1-\sum_{j \in N(i)}(1-\alpha_j)/d_j}{n\alpha_i}|}{\sum_{j=1}^n |\frac{1-\sum_{k \in N(j)}(1-\alpha_k)/d_k}{n\alpha_j}|}$$

Algorithm 3 INFLUENCESAMPLING

1: Choose a random sample $S \subseteq V$ of size r with replacement by sampling each node i with probability p_i proportional to $|c_i| = |\frac{1-\sum_{j \in N(i)}(1-\alpha_j)/d_j}{n\alpha_i}|$
2: Output $\hat{Y}^0 = (\sum_{j=1}^n |c_j|) \cdot \frac{1}{r} \sum_{k \in S} Y_k(t) \cdot \text{sgn}(c_k)$, where $\text{sgn}(c_k) = 1$ if $c_i \geq 0$ and -1 otherwise.

We now show that estimator $\hat{Y}^0$ of the INFLUENCESAMPLING algorithm is an unbiased estimator of $\bar{Y}^0$

Lemma 5 *The* INFLUENCESAMPLING *estimator* $\hat{Y}^0$ *is an unbiased estimator of* $\bar{Y}^0$

Proof.   Since we sample each node with probability proportional to $|c_i|$, $E[\hat{Y}^0] = (\sum_{j=1}^n |c_j|) \cdot \sum_{i=1}^n p_i Y_i(t) \text{sgn}(c_i) = \sum_{i=1}^n |c_i| Y_i(t) \text{sgn}(c_i) = \sum_{i=1}^n c_i Y_i(t) = \bar{Y}^0$ ∎

Theorem 6 *Using* $\Theta(\frac{1}{H_\alpha \epsilon^2} \log(1/\delta))$ *samples (where* $H_\alpha = \frac{n}{\sum_{i=1}^n 1/\alpha_i}$ *denotes the harmonic mean of the* α_i), *the estimator* $\hat{Y}^0$ *output by the* INFLUENCESAMPLING *algorithm satisfies* $|\hat{Y}^0 - \bar{Y}^0| < \epsilon$ *with probability* $1 - \delta$, *when the* c_i *are non-negative[2]. Furthermore, this number of samples is optimal up to constant factors.*

Proof.   Note that we sample each node with probability proportional to $|c_i|$. Let D_c denote this sampling distribution. The number of samples required follows from the Hoeffding inequality [18]. This is because $\hat{Y}^0$ is the mean of i.i.d. random variables of the form $(\sum_{j=1}^n |c_j|) \cdot Y_i(t) \text{sgn}(c_i)$ where $i \sim D_c$. Each of these is a bounded random variable in $[-l, l]$ where $l = \sum_{i=1}^n |c_i|$. Hence, using the Hoeffding inequality and Lemma 5, $\text{Prob}(|\hat{Y}^0 - E[\hat{Y}^0]| > \epsilon) =$

---
[2]we address the case when the c_i are negative in Remark 7.

Prob($|\hat{Y}^0 - \bar{Y}^0| > \epsilon$) < $e^{-k\epsilon^2/2l^2}$, where k is the number of samples. Thus, $2\frac{l^2}{\epsilon^2}\log(1/\delta)$ is sufficient (and necessary [18]) to get an additive ϵ approximation estimate with confidence $1 - \delta$. For non-negative c_i, $l \leq \sum_{i=1}^{n} \frac{1}{n\alpha_i} = \frac{1}{H_\alpha}$. The bound follows. ∎

Note that the number of samples in the above theorem matches (up to a constant factor of $\frac{1}{H_\alpha}$) the bound that we get from Fact 3. Thus, INFLUENCESAMPLING performs almost as well as this optimal ideal uniform sampling algorithm, even though the latter has a huge advantage in terms of assuming accessing to the innate opinion. This is due to the fact that INFLUENCESAMPLING can take advantage of implicitly accessing neighbors' innate opinions encoded in the node's expressed opinion, as long as it is debiased properly. This shows the optimality of our algorithm.

Remark 7 *If the c_i values are negative, then the number of samples needed in Theorem 6 could be larger (since l is then bounded by $\sum_{i=1}^{n} \frac{d_i}{n\alpha_i}$ instead of $\sum_{i=1}^{n} \frac{1}{n\alpha_i}$). Note that nodes with negative c_i correspond to subgraphs having large star topologies with node i in the center. We observe that in our experiments (and social networks in general) the number of such nodes is never too large. In fact, in our experiments, the c_i values that we empirically calculate are always positive.*

A insightful special case is when all degrees are equal to d. Note that in this case $c_i = \frac{\sum_{j\in N(i)} \alpha_j}{nd\alpha_i}$. This means that we give advantage to weakly opinionated nodes that are surrounded by strongly opinionated ones as they aggregate their innate opinions efficiently.
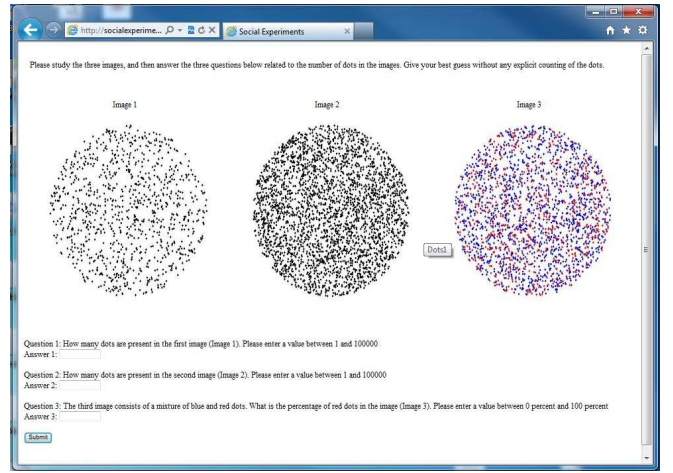
Remark 8 *Note that in practice we may not have the exact values of α_i, but only estimates that are appromations of the true value. If these estimates are within an ϵ additive factor of the true values then our estimator for the mean Y_i^0 will also have a bias within at most ϵ additive factor of the true value.*

Also note that if we can compute c_i for a small fraction of the nodes, that can be used to approximately estimate $\sum_i c_i$ that is needed in INFLUENCESAMPLING.
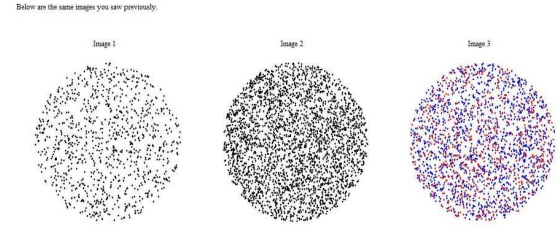
A word about the assumption in our sampling algorithm on knowing the estimates of α values of users: it is not entirely unrealistic to get reasonable estimates for a user's propensity to conform in real social networks. For example, on the Twitter Social Network, one might use a function of a user's tweet and retweet frequency as a proxy for her α values. Similarly, on online sites that support discussion threads [1], a user's participation history could be used to obtain an approximate estimate of her α values. We leave the characterization of algorithms that mine users' social posts to obtain an estimate of her α value for future research. Furthermore, as the above remark shows, our techniques gracefully handle scenarios when these α values are either missing or noisy for some users.

5. EXPERIMENTS

We now validate our opinion formation model and our sampling algorithms using synthetic and real datasets. We first describe our datasets.

(a) Screen 1 of the DOTSREGULAR survey

(b) Screen 2 of the DOTSREGULAR survey

Figure 1: Screenshots of the DOTSREGULAR **survey**

5.1 Datasets

Our first set of experiments consisted of an interacting network of online users built using Amazon Mechanical Turk (mturk) and a personal website used to host the experiments. Subjects were recruited from Amazon Mechanical Turk and asked to take part in three online surveys hosted on the external website. Figure 1 illustrates the one of the online experiments. Figure 1(a) shows the first set of questions asked to users while the subsequent screen in the experiment showing the "interactions" is captured in Figure 1(b). We label these surveys DOTSREGULAR, DOTSRANDOM and TABLETSREGULAR. The aim was to ask subjects' opinions on questions in the surveys both before and after interactions with other subjects, and analyze their innate (initial) and expressed (final) answers. To show that our opinion models are not specific to a particular type of questions or topics, we ran our experiments using two different topics for our questions - the first one dealt with questions about properties of an image, and the second one dealt with the opinions about a class of consumer products. We detail the surveys next.

5.1.1 User Surveys

The DOTSREGULAR survey consisted of three questions about three images containing a set of dots. The first image consisted of 1000 randomly distributed black dots in a circle. The second image consisted of 3000 randomly distributed black dots in a circle, and the third image consisted of a mixture of 900 red dots and 1800 blue dots randomly distributed in a circle. There were three questions in the surveys. The first two questions were about guessing the number of dots in the first and second image respectively. The third question was about guessing the percentage of red dots in the third image. In all the images, the dots were finely spaced enough to discourage explicit counting of dots by the participants to estimate the answers.

The survey participants were asked to log onto the survey site and complete the survey only during a pre-specified 30-minute time window. This ensured roughly simultaneous participation by all the subjects in the online survey. In this time duration, the subjects were first asked to provide their answers to the three questions (which was treated as their innate opinions about the questions). Then, for each participant we randomly assigned a set of 5 other participants as her "neighbors", and showed her their current answers. We then gave her an opportunity to update her current answers. This process was repeated for each survey participant for up to 3 iterations (the answers of the participants changed very little after 3 iterations, hence we chose to report the answers provided by the users after the third iteration.) Thus we created a social graph among the survey participants that is 5-regular, within which each participant interacts with her neighbors during the process of converging to her final answer. Figure 2(a) illustrates the resulting graph. It has 125 nodes corresponding to the number of participants in the experiment and 625 edges.

The DOTSRANDOM survey was identical to the DOTSREGULAR survey, except that instead of fixing 5-neighbors for each participant, we randomly selected a varying number of neighbors (from 1 through 15) for a given participant. This experiment elicited response from 63 participants resulting in graph with 63 nodes and 584 edges. We note that in both the surveys, the original answers provided by the participants (before any interactions with other participants) were treated as their innate opinions and the final answers were treated as their expressed opinions.

The final survey, TABLETSREGULAR, had a setup similar to DOTSREGULAR, except that we asked a different set of questions related to two recently introduced tablet computer products (A and B) in the market. The first question asked about the opinion of the participant regarding the potential success of a tablet computer in the market. The participant was asked to choose between 5 options: Strongly Positive, Positive, Neutral, Negative and Strongly Negative. The question also asked the participant to provide a sentence justifying their choice. Thus, these justifications could include positive, negative, and neutral comments about the product corresponding to the star rating provided by the user. Another unrelated question asked participants about what they thought would be the total sales volume of another tablet in the month of January 2013. These questions are markedly different from the dots surveys in the sense that they are more subjective. In fact, as part of the survey associated with the first question, we showed a user with the justifications written by her randomly selected set of neigh-

bors as a proxy for the influence she might experience as part of her opinion forming process. As before, the original answers for the first and second provided by the participants (before any interactions with other participants) were treated as their innate opinions and the final answers were treated as their expressed opinions. Again, the total number of participants for this survey was 125 users resulting in a graph with 125 nodes and 625 edges.

5.1.2 Synthetic Data

The second set of experiments were performed using synthetic data. We generated random graphs with 200000 nodes and two different types of degree distributions: regular graphs with degree 20, and power-law graphs degree of each node was obtained from a Zipf distribution truncated to lie between 20 and 200. For each graph, we generated the α parameters for each node from a uniform distribution in $[0, 1]$. We also generated an initial (innate) opinion for each node from a Gaussian distribution with mean 50 and varying values of variance ranging between 0 and 1000. We then ran our opinion formation dynamics on the graph until the expressed opinions converged to the equilibrium value and used the starting (innate) values and the converged values as the expressed opinions in our analysis of the model and the sampling algorithms.

5.2 Results

For the social graphs corresponding to both the survey and the synthetic datasets, we then ran the various sampling algorithms using different sample sizes r, read the expressed opinions of the sampled nodes, and output an estimate of the average innate opinion in the graph as a function of these expressed opinions. We compute the estimation error to be the absolute difference between this estimate and the true average innate opinion. We repeat the sampling process over 100 runs and report both the average mean and variance of the estimation error over all the runs.

5.2.1 Validating the Opinion Formation Model

We began by validating the opinion formation model from Section 3. Note that the α parameter is central to the model as it captures the intrinsic behavior of a user when it comes to interacting with her neighborhood in the process of forming an opinion. Specifically, we set out to measure how innate is this parameter to the user, i.e., how consistent was the user in her behavior for different questions. This consistency would reflect in a rather stable value of the α parameter. To study this, we use the opinion model equation (1 and the initial and final answers for each of the survey questions, to obtain an estimate of the α parameter for each participant: we denote them by α^1, α^2 and α^3. We then measure how similar these three estimated α values are for each participant. For every (α^i, α^j) pair for a participant, we plot the relative difference in values (to handle to case where one of the α^i values might be 0, we use $1 - \frac{\min(\alpha^i, \alpha^j)}{\max(\alpha^i, \alpha^j)}$ to measure the relative difference) Figure 3(a) plots the relative pairwise difference in the computed α values from the three questions in the DOTSREGULAR Survey (the plots corresponding to the DOTSRANDOM and TABLETSREGULAR experiments were similar). As the figure illustrates, the relative difference in the computed α values are indeed small (less than 0.2) for almost 60% of the participants, which shows that the α values indeed are largely independent of the questions
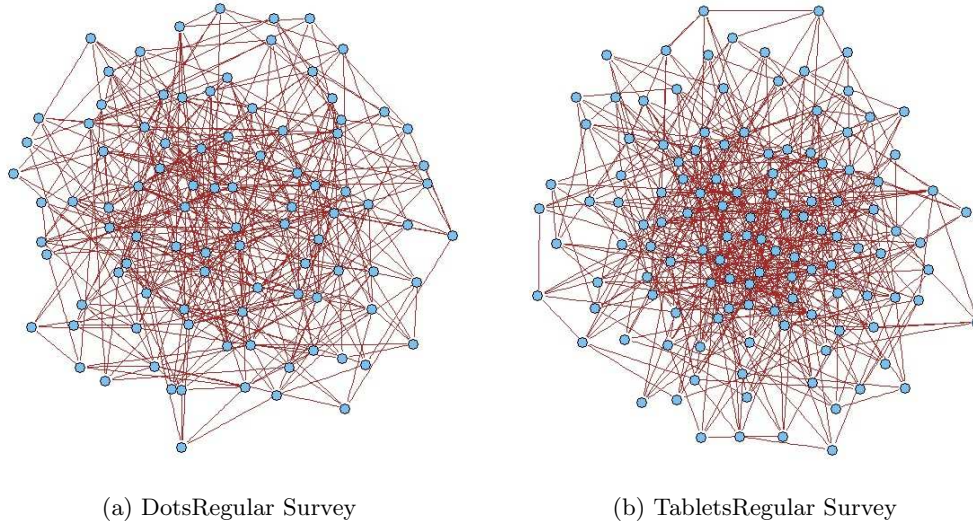
(a) DotsRegular Survey

(b) TabletsRegular Survey

Figure 2: Social graphs resulting from the surveys

(a) Relative difference in α estimates per user

(b) Histogram of α values

Figure 3: α values for DotsRegular

(a) Relative difference in α estimates per user

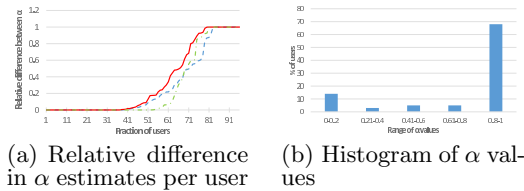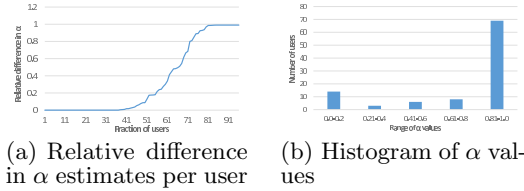(b) Histogram of α values

Figure 4: α values for TabletsRegular

asked. In particular, the average relative difference over all users and all (α^i, α^j) pairs was observed to be 28%.

Figure 3(b) plots the distribution of α^1 values among the survey population of DotsRegular. It is interesting to note that while a majority of the users have a high α value, there is a non-trivial fraction of users that value the opinion of their neighbors or have a higher propensity to conform.

Figures 4(a) and 4(b) plots the corresponding relative pairwise difference in computed α values and the distribution of the α^1 values for the TabletsRegular Survey. Notice that the distribution is very similar to the DotsRegular case, which again seems to indicate that the α values are largely independent of the questions for a given topic.

5.2.2 Sampling Innate Opinions

Next, we move to the question of estimating the average innate opinion of the survey population (i.e., the wisdom of the social crowd) by using only a small sample of the expressed opinions. Using the α parameter value associated with the first question of each of the three surveys (Dot-

sRegular, DotsRandom, and TabletsRegular) we run the three sampling algorithms described in Section 4, *viz.*, InfluenceSampling, UniformSampling and ConformitySampling, and compare the estimates ($\hat{Y^0}$) output by each of the algorithms against the true average innate opinion ($\bar{Y^0}$). For a given sample size, we run each sampling algorithms over 100 runs and plot the mean and variance of the estimation error $|\hat{Y^0} - \bar{Y^0}|$.

Figure 5 plots the average mean and variance of the estimation errors for each sampling algorithm as a function of the samples sizes for the DotsRegular, DotsRandom and TabletsRegular experiments. As seen in the figure, InfluenceSampling outperforms both UniformSampling and ConformitySampling in all three surveys in terms of the mean estimation error, by margins ranging from 10% to almost 30%. This validates our theoretical results showing optimality of the InfluenceSampling algorithm. More intuitively, our algorithm outputs a more accurate estimate by picking nodes that implicitly aggregate their innate neighborhood opinion.

Again, in terms of the standard deviation of the estimation error, InfluenceSampling achieves significantly lower standard deviations than either UniformSampling and ConformitySampling. Furthermore, the corresponding curves for InfluenceSampling are much better behaved (in terms of the expected monotonic decay as the number of samples increase) compared to the other sampling algorithms. The standard deviation plots for UniformSampling and ConformitySampling algorithms does not decay smoothly with the number of samples, and the reason for this was not intuitively obvious to us.

5.2.3 Moving to Larger Graphs

For the synthetic datasets, our main focus was to evaluate the performance of the sampling algorithms at scale. As mentioned earlier, we use a uniform distribution for the α values, and a Gaussian distribution (with variance ranging from 0 to 1000) for the innate opinions at each node. We conduct our experiments for both regular graphs and power-law graphs. For each type of graph, we report our estimation er-
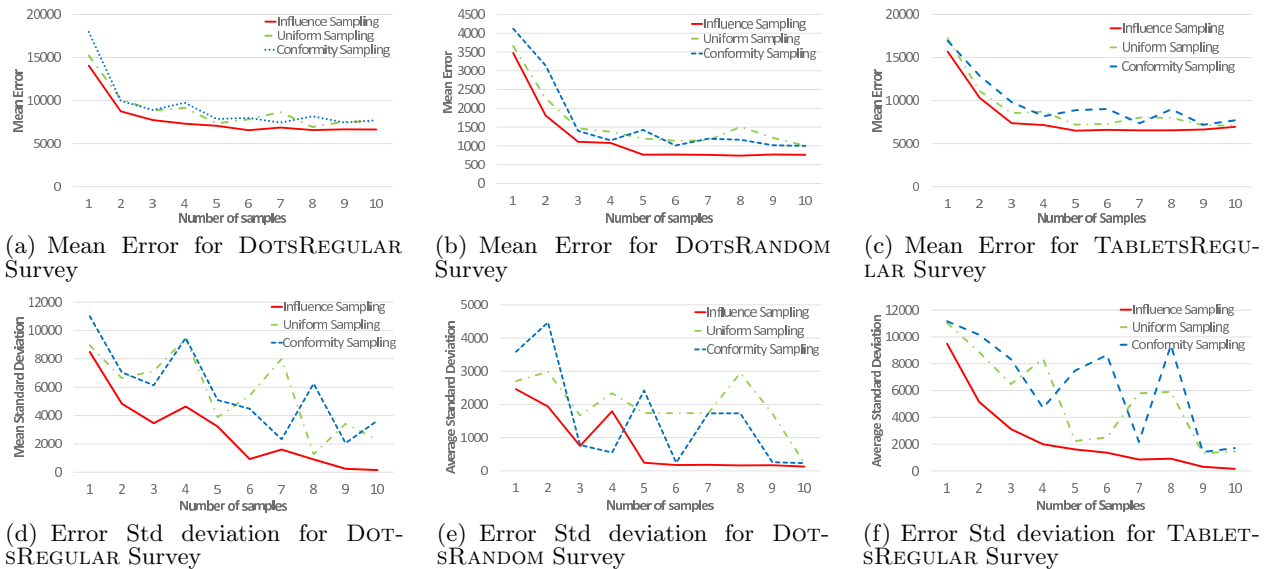
(a) Mean Error for DotsRegular Survey

(b) Mean Error for DotsRandom Survey

(c) Mean Error for TabletsRegular Survey

(d) Error Std deviation for DotsRegular Survey

(e) Error Std deviation for DotsRandom Survey

(f) Error Std deviation for TabletsRegular Survey

Figure 5: Mean Errors and Standard Deviation for the Surveys

ror results both for a high-variance case and a lower-variance case for the innate opinion distribution. We compare the estimates ($\hat{Y}^0$) output by each of the algorithms against the true average innate opinion ($\bar{Y}^0$). For a given sample size, we run each sampling algorithms 100 times and report the mean and variance of the estimation error $|\hat{Y}^0 - \bar{Y}^0|$.

Figure 6 plots the mean and variance of the estimation errors for each sampling algorithm as a function of the samples sizes for regular graphs and for power-law graphs.

Again, as with the survey experiments, InfluenceSampling outperforms both UniformSampling and ConformitySampling in all three surveys in terms of the mean estimation error, by as much as 200% for 10 samples, and around 30% for 100 or more samples. As Figure 6 illustrates, the gap in performance of InfluenceSampling is less pronounced as the number of samples increases.

In terms of the standard deviation of the estimation error too, InfluenceSampling achieves significantly lower standard deviations than either UniformSampling and ConformitySampling for a small number of samples, though this gap goes down as the number of samples increases.

Finally, we studied the convergence of the opinion formation dynamics of Equation 1 in terms of the average change in the expressed opinion of nodes over two successive iterations. As seen from Figure 7, expressed opinions even for regular graphs of size 200000 converges to the equilibrium values in as little as 3 steps something which was also observed in our real-world survey experiments.

As might be expected, the improved performance of Influence Sampling over the other sampling algorithms is even more significant in these large scale synthetic graphs, than for the small-scale survey experiments on Mechanical Turk.

6. CONCLUSIONS

In this paper, we considered the problem of analyzing and debiasing the "wisdom of crowd" phenomenon in the presence of online social interactions. We adopted a natural opinion formation model that depends on users' propensity
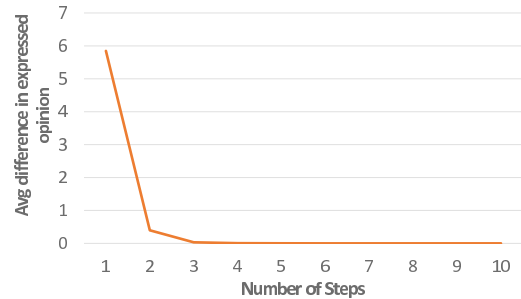

Figure 7: Convergence of the opinion formation process for synthetic regular graphs

to conform (as characterized by their α parameters), and designed a provably-efficient sampling algorithm (InfluenceSampling) that uses these α values to estimate the average innate opinion of the social crowd with a small number of samples. We validated the opinion formation model on User Opinion Surveys, and evaluated our sampling algorithm on both real and synthetic data.

One direction of future work is to validate the model and sampling techniques on large social networks under more richer social interactions.

7. REFERENCES

[1] L. Backstrom, J. Kleinberg, L. Lee, and C. Danescu-Niculescu-Mizil. Characterizing and curating conversation threads: expansion, focus, volume, re-entry. In *Proceedings of the sixth ACM international conference on Web search and data mining*, pages 13–22, 2013.

[2] K. Bhawalkar, S. Gollapudi, and K. Munagala. Coevolutionary opinion formation games. In *ACM Symp. on Theory of Computing*, 2013.

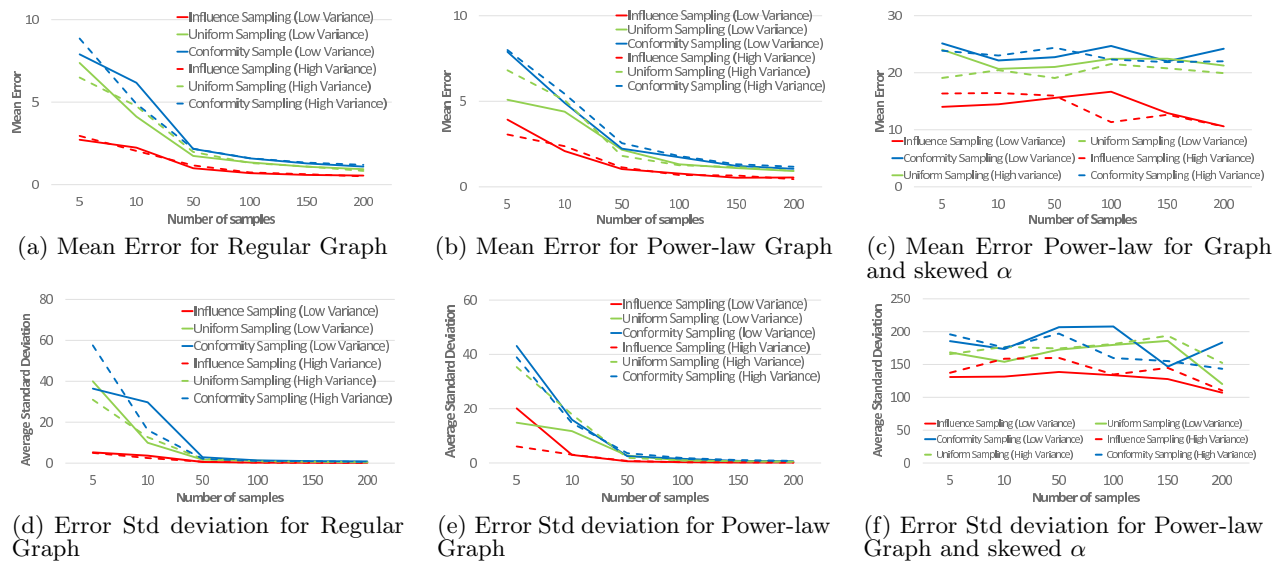[3] D. Bindel, J. Kleinberg, and S. Oren. How bad is forming your own opinion? In *Foundations of*

(a) Mean Error for Regular Graph (b) Mean Error for Power-law Graph (c) Mean Error Power-law for Graph and skewed α

(d) Error Std deviation for Regular Graph (e) Error Std deviation for Power-law Graph (f) Error Std deviation for Power-law Graph and skewed α

Figure 6: Mean Errors and StandardDeviation for Synthetic Graphs

Computer Science (FOCS), 2011 IEEE 52nd Annual Symposium on, pages 57–66, 2011.

[4] A. Dasgupta, R. Kumar, and D. Sivakumar. Social sampling. In *Proc. 18th Intl. Conf. on Knowledge Discovery and Data Mining*, pages 235–243, 2012.

[5] M. H. DeGroot. Reaching a consensus. *J. American Statistical Association*, 69:118–121, 1974.

[6] Z. Dezső and A.-L. Barabási. Halting viruses in scale-free networks. 65, 2002.

[7] P. Domingos and M. Richardson. Mining the network value of customers. In *Proc. 7th Intl. Conf. on Knowledge Discovery and Data Mining*, pages 57–66, 2001.

[8] J. French. A formal theory of social power. *Psychological Review*, 63:181–194, 1956.

[9] N. E. Friedkin and E. C. Johnsen. Social influence and opinions. *J. Mathematical Sociology*, 15(3-4):193–205, 1990.

[10] F. Galton. Vox populi. *Nature*, 75:450, 1907.

[11] S. Goel, W. Mason, and D. Watts. Real and perceived attitude agreement in social networks. In *Journal of personality and social psychology, vol 99, no. 4, pp. 611*, 2010.

[12] M. Granovetter. Threshold models of collective behavior. *American Journal of Sociology*, 83:1420–1443, 1978.

[13] R. Hegselmann and U. Krouse. Opinion dynamics and bounded confidence models, analysis, and simulations. *Journal of Artificial Societies and Social Simulation*, 5(3), 2002.

[14] S. Judd, M. Kearns, and Y. Vorobeychik. Behavioral dynamics and influence in networked coloring and consensus. *Proceedings of the National Academy of Sciences*, 107(34):14978–14982, 2010.

[15] D. Kempe, J. Kleinberg, and E. Tardos. Maximizing the spread of influence in a social network. In *Proc. 9th Intl. Conf. on Knowledge Discovery and Data Mining*, pages 137–146, 2003.

[16] D. Lazer and A. Friedman. The network structure of exploration and exploitation. *Administrative Science Quarterly*, 52(4):667–694.

[17] J. Lorenz, H. Rauhut, F. Schweitzer, and D. Helbing. How social influence can undermine the wisdom of crowd effect. *Proc. Natl. Acad. Sci. USA*, 108(22), 2011.

[18] R. Motwani and P. Raghavan. *Randomized algorithms*. Cambridge university press, 1995.

[19] A. G. Parameswaran, H. Garcia-Molina, H. Park, N. Polyzotis, A. Ramesh, and J. Widom. Crowdscreen: Algorithms for filtering data with humans. In *Proceedings of the 2012 international conference on Management of Data*, pages 361–372, 2012.

[20] H. Park, H. Garcia-Molina, R. Pang, N. Polyzotis, A. Parameswaran, and J. Widom. Deco: a system for declarative crowdsourcing. *Proceedings of the VLDB Endowment*, 5(12):1990–1993, 2012.

[21] T. Schelling. *Micromotives and Macrobehavior*. Norton, 1978.

[22] B. Shiv. The lonely shopper - `ttp://www.gsb.stanford.edu/news/researc`/shiv_lonely_2011.html, 2011.

[23] J. Surowiecki. *The Wisdom of Crowds*. Anchor, 2005.