Consensus via small group interactions: the importance of triads

ASHISH GOEL and DAVID T. LEE, Stanford University

We consider a framework for crowdsourcing in which "the wisdom of the crowd" is elicited through a consensus process. In our setting, a set of participants submit proposals for a given question of interest. An algorithm designer would like to find the (generalized) median of this set, but does not have access to information about the metric space that the proposals lie in. The goal is then to design a consensus mechanism which induces the participants to "find" the median through interactions with each other.

We propose an urn-based algorithm which solves this problem through interactions among groups of three. Each participant is represented as a ball in an urn. At each step, three balls are drawn from the urn, and the corresponding participants are made to take part in a mechanism in which they must decide on one of the submitted proposals. The balls that were drawn from the urn are then relabeled with the chosen winner, and the process is repeated until there is a single participant remaining. When participants form a median graph (e.g. trees, grids, squaregraphs), we show that any triadic mechanism which finds the median of the three participants will lead to a $(1+O(\sqrt{\frac{\log n}{n}}))$-approximation of the true global median with high probability, while only requiring an average of $O(\log n)$ triadic interactions per participant. We give an example of such a three person mechanism, and show that it finds the median of the given participants under a Nash equilibrium.

Finally, we prove an impossibility result which highlights the importance of triads. Namely, if our urn-based reduction is modified to draw only two balls at each step, then there is no two person mechanism for the participants which finds a good approximation of the median, unless the mechanism fails a natural local consistency property. This result has implications not only for the design of consensus mechanisms, but also for opinion formation dynamics on social networks.

Categories and Subject Descriptors: F.2 [**Theory of Computation**]: Analysis of Algorithms and Problem Complexity; J.4 [**Computer Applications**]: Social and Behavioral Sciences—*Economics*

General Terms: Algorithms, Economics, Theory

Additional Key Words and Phrases: Crowdsourcing, wisdom of the crowd, computational social choice, randomized algorithms, 1-median, consensus reduction, Triadic Consensus

ACM Reference Format:
ACM Trans. Algor. X, X, Article X (February 2013), 18 pages.
DOI = 10.1145/0000000.0000000 http://doi.acm.org/10.1145/0000000.0000000

1. INTRODUCTION

In 1907, Sir Francis Galton went to a carnival and observed a competition occuring in which participants could guess the weight of an ox. As people made their estimates, Galton recorded them and observed that the median, which he called the voice of the people (*Vox Populi*), was remarkably close to the correct answer [Galton 1907]. Based on this observation, he hypothesized that a crowd's choice solution may be among the best of any individual, an idea that later became known as "the wisdom of the crowd" [Surowiecki 2005].

In this paper, we consider how to elicit the wisdom of the crowd for abstract goals such as "designing an undergraduate curriculum" or "deciding on a governmental budget". We will suppose that proposals $x_1, x_2, \ldots, x_n$ belong to a metric space with distance metric $d(\cdot,\cdot)$, and our goal will be to find the 1-median, the proposal which minimizes the quantity $\sum_i d(\cdot, x_i)$ (see Section 1.1.2 for its relation to the wisdom of the crowd). We assume that, as in many practical settings, we do not know $d(\cdot,\cdot)$, but that participants have some notion of the proposal space, as reflected in their preferences and interactions between each other. The resulting problem statement is to design a mechanism (through voting, bargaining, etc...) that induces the participants to find the 1-median of their proposals.

ALGORITHM 1: Triadic Consensus

Input: An urn with k labeled balls for each participant
Output: A winning participant
while *there is more than one label* **do**
    Sample balls x, y, z uniformly at random with replacement;
    if *two or more of* x, y, z *have the same label* **then**
        $w =$ the majority candidate;
    else
        $w = \text{TriadicMechanism}(x, y, z)$;
    Relabel all the sampled balls with the winning label w;
return the label of the remaining balls;

Our approach will be to solve this through a *consensus reduction*, in which we reduce our original problem of finding the median of a crowd into the problem of finding the median for three participants. Once this reduction has been made, it is not hard to then design the required triadic median-finding mechanism.

We present an urn-based reduction which finds a $(1 + O(\sqrt{\frac{\log n}{n}}))$-approximation of the 1-median when the proposals can be represented as a median graph, a class of graphs including trees and high-dimensional grids. We then give an example of a three person median-finding mechanism, prove that each participant only needs to participate in an average of $O(\log n)$ of these triadic interactions, and show that all these results hold under a Nash equilibrium, i.e. that this process is resistant to manipulation. Finally, we prove a surprising impossibility result on two person consensus reductions. Namely, there are no two person mechanisms satisfying *local-consistency* that can find a good approximation of the global median under any urn-based reduction. Intuitively, this property states that the consensus outcome of a set of participants should only depend on the relationships between the participants.

Our reduction (Algorithm 1) is a generalization of Triadic Consensus as introduced in Goel and Lee [2012]. It is best described by imagining an urn with balls, each of which is labeled by a participant id. The urn starts with k balls for each of the n participants. At each step of the algorithm, three balls are uniformly sampled at random with replacement, and the participants represented by these balls engage in a three person consensus mechanism in which they decide on one winning proposal. The selected balls are then replaced by three copies of this winner. This is repeated until there is only one participant id left, which is declared the winner. The following examples illustrate the strong convergence properties of this reduction despite its inherently random nature:

Example 1.1 (*A* $n \times n$ *grid*). Consider n^2 participants which make up a $n \times n$ grid. The global 1-median of all participants is the point $(l + 1, l + 1)$, where we assume for simplicity that n is odd and that $n = 2l + 1$. Suppose that the TriadicMechanism returns the 1-median of any three sampled balls. Then Triadic Consensus (with $k = 1$) will produce a winner that is within $\frac{1}{2}\sqrt{\ln \frac{1}{\delta}}$ of the global 1-median with probability at least $1 - \delta$. Similarly, Triadic Consensus with $k = \ln \frac{1}{\delta}$ will produce the exact global 1-median with probability at least $1 - \delta$. Note that this is independent of n.

Example 1.2 (*A* $n \times n \times n$ *grid*). Consider n^3 participants which make up a $n \times n \times n$ grid. The global 1-median of all participants is the point $(l + 1, l + 1, l + 1)$, where we assume for simplicity that n is odd and that $n = 2l + 1$. Then supposing that the TriadicMechanism will return the 1-median of any three sampled balls, Triadic Con-

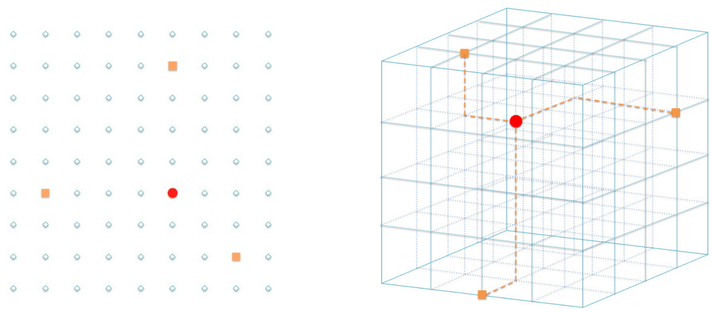

(a) (b)

Fig. 1: If the three orange squares are randomly sampled, then we will assume that TriadicMechanism returns the red circle, which is their 1-median. For Triadic Consensus with $k = 4$ on a $n \times n$ grid, this seemingly random process is able to identify the exact global median with ninety percent probability regardless of the grid width. For a $n \times n \times n$ grid, the results are even stronger. Triadic Consensus with $k = 1$ identifies the exact global median with probability $1 - e^n$.

sensus (with $k = 1$) will produce the exact global 1-median with probability at least $1 - e^{-n}$.

1.1. Related work and our contributions

Before continuing, we will describe some of the related literature[1], which will be followed with a summary of our contributions (Section 1.1.3).

1.1.1. Wisdom of the Crowd. The literature on crowdsourcing and the wisdom of the crowd is vast (see Brabham [2008]). The overarching premise is that humans are able to perform tasks that computers are not; therefore, one would like to understand how to harness such abilities in human powered algorithms.

Generally speaking, there are two classes of these systems. One class contains systems which are highly tailored for a given problem domain. These systems often utilize highly specialized domain knowledge to solve tasks as difficult as protein-folding [Khatib et al. 2011]. On the other end of the spectrum are general crowdsourcing frameworks such as Amazon Mechanical Turk, which ask participants to perform simple tasks such as labeling, for which general algorithms can be designed for their aggregation. There are very few general frameworks for the crowdsourcing of complex tasks. One example may be that of prediction markets [Wolfers and Zitzewitz 2004], in which an information market is set up that enables the crowdsourcing of arbitrary prediction tasks, so long as they can be posed in the form of binary questions.

1.1.2. Social choice. Another general framework for the aggregation of potentially complex proposals can be found in the social choice literature. In the most common setting known as rank aggregation (see the review by Brandt et al. [2012]), voters provide rankings for a set of candidates, and some social choice function is then applied to determine a winner or aggregate ranking. The social choice functions are then evaluated according to their ability to meet different social choice criteria. In the spatial voting context, proposals are points in space, and voter preferences are induced by the

[1]The authors would love any pointers to literature on the wisdom of the crowd, crowdsourcing, or consensus that may be related!

distances between proposals, i.e. the closer a proposal is to you, the more you prefer it to another.

One approach social choice has taken towards the wisdom of the crowd is to study social choice functions that are maximum likelihood estimators for certain error functions [Conitzer and Sandholm 2013] [Mao et al. 2013]. Such a voting rule can be interpreted as finding the wisdom of the crowd given some assumption on distribution on voter errors.

Another approach is to note that the Condorcet criterion, named after the Marquis de Condorcet (1743 - 1794), coincides with the notion that Galton had in his concept of the wisdom of the crowd in which he states that:

> "the middlemost estimate expresses the *vox populi*, every other estimate being condemned as too low or too high by a majority of the voters."

He is essentially stating the Condorcet criterion, which states that if there is one candidate who would beat every other candidate in a pairwise election, then this candidate should be chosen as the winner. The Condorcet criterion does not always exist. However, in spatial voting on a line or a tree, the Condorcet winner coincides exactly with the 1-median [Wendell and McKelvey 1981]. More generally, for a large class of graphs known as median graphs, if a Condorcet winner exists, then it is also a 1-median [Bandelt and Barthelemy 1984] [Saban and Stier-Moses 2012].

One of the challenges in aggregation in the crowd context is the cognitive burden that is imposed on the user. If there are one thousand proposals, then even reading each proposal, much less ranking them, becames impractical. In the social choice setting, researchers have studied this problem by asking whether it is possible to find the aggregate winner or ranking with only a small number of comparisons or other partial queries. In the general case, Conitzer and Sandholm essentially showed that it is not possible for most common voting rules, and that determining how to elicit preferences efficiently is also NP-complete, even with perfect knowledge about voter preferences [Conitzer and Sandholm 2005] [Conitzer and Sandholm 2002]. One exception is when preferences are restricted to single-peaked preferences, a setting in which candidates can be represented on an axis, and voter preferences can be described by their ideal point on that axis. In this case, voters can simply state their ideal proposal. If we have access to the candidate axis, then a simple algorithm for finding the Condorcet winner is just to calculate the median proposal stated, and return this candidate. If we do not have access to the candidate axis, [Conitzer 2009] and [Goel and Lee 2012] show how to use comparison queries to elicit the Condorcet winner. There have also been approaches based on machine learning [Lu and Boutilier 2011] that show experimental results.

1.1.3. Our Contributions. The premise of our work is that humans are able to do more than just make comparisons, e.g. find "the proposal closest to my ideal proposal x that lies between y and z". Under this premise, we show that it becomes possible to find a tight approximation of the 1-median in settings far more general than the single-peaked setting. The algorithm we propose requires only an average of $O(\log n)$ interactions per participant and is resistant to manipulation, making it practical from a communication and game-theoretic perspective. This work provides a novel framework for the aggregation of potentially complex proposals via consensus processes.

We also find an important insight on the use of small group interactions for finding the wisdom of the crowd. Specifically, we find that there are no locally-consistent two-person mechanisms that can find a good approximation of the 1-median under any urn-based consensus reduction, even in simple settings. The implications are quite surprising for not only the design of consensus reductions, but also for opinion formation dynamics on social networks. Consider any opinion formation dynamic on a social network where opinion updates happen by picking an edge and having the two individuals selected update their opinions in a manner that is locally consistent. Then there are proposal spaces such that the resulting consensus proposal is far from the 1-median.

We note that the idea of consensus reductions may be of independent interest for the future study and design of consensus systems for large crowds. Potentially, one may be able to leverage such reductions to solve crowd problems through experimental analysis of small groups.

1.1.4. Outline of the paper. The remainder of the paper will be structured as follows. We will discuss our reduction (Algorithm 1) and its convergence properties in Section 2. In Section 3, we show that our reduction inherits Nash equilibrium properties from the triadic mechanisms used and give a concrete example of a three person mechanism. Finally, in Section 4, we formally define *locally-consistent* mechanisms, and show that there are no such two-person mechanisms that can find the 1-median under urn-based consensus reductions. We conclude with discussions and future work in Sections 5.

2. A CONSENSUS REDUCTION FOR THE 1-MEDIAN PROBLEM

The main proof technique relies on a reduction of the Triadic Consensus urn to a class of urns studied in Lee and Bruck [2012], which they describe as fixed size urns with an urn function. Consider an urn with some number of balls, each of which is colored red or blue. Let R_t and B_t denote the number of red and blue balls respectively at time t, where $R_t + B_t = n$. Also, let $p_t = \frac{R_t}{n}$ denote the fraction of red balls. At every discrete time t, either a red ball is sampled with probability $f(p_t)$, a blue ball is sampled with probability $f(1 - p_t)$, or nothing happens with the remaining probability. The function $f : [0, 1] \to [0, 1]$ is called an urn function and satisfies $0 \leq f(x) + f(1 - x) \leq 1$ for $0 \leq x \leq 1$. If a ball was sampled, it is then recolored to the opposite color and placed back into the urn. This process repeats until some time T when all the balls are the same color, i.e. $R_T = n$ or $R_T = 0$.

It turns out that, for median graphs, Triadic Consensus can be reduced to fixed size urns with urn function $f(p) = 3p(1 - p)^2$. The following theorems will be used to prove our main results:

THEOREM 2.1. *Let a fixed size urn start with R_0 red balls out of n total balls and have an urn function $f(p) = 3p(1 - p)^2$. Let T denote the first time when either $R_T = n$ or $R_T = 0$. Then,*

$$\mathbf{Pr}[R_T = n] = \left(\frac{1}{2}\right)^{n-1} \sum_{j=1}^{R_0} \binom{n-1}{j-1} \tag{1}$$

$$\mathbb{E}[T] \leq n \ln n + O(n) \tag{2}$$

$$\mathbf{Pr}[T > 11cn \ln n] < n^{-c} \tag{3}$$

PROOF. Equations (1) and (2) are derived in Goel and Lee [2012] through fairly straightforward applications of general theorems in Lee and Bruck [2012]. Proving (3) requires a bit more work.

Let $T_n(i)$ denote $\mathbb{E}[T]$ given $R_0 = i$. From Lee and Bruck [2012], we have that $X_t = T_n(R_t) + t$ is a martingale, and that

$$\tau_n(i) = T_n(i) - T_n(i-1) = \sum_{j=i}^{\lfloor n/2 \rfloor} \frac{n^3}{3j(n-j)^2} \frac{\binom{n-1}{i-1}}{\binom{n-1}{j-1}}$$

for $i < \frac{n}{2}$ and n odd. We can then derive that $\tau_n(\lfloor \frac{n}{2} \rfloor) = \frac{8}{3}$, and

$$\tau_n(i) = \tau_n(i+1)\frac{i}{n-i} + \frac{n^3}{3i(n-i)^2}$$

For $\frac{n}{2} - c\sqrt{n} < i < \frac{n}{2}$, $n > 4c^2$,

$$\begin{aligned}\tau_n(i) &< \sum_{j=i}^{n/2} \frac{n^3}{3j(n-j)^2} \\ &< c\sqrt{n}\frac{n^3}{3(n^3/8 - c^2n^2/2 + cn^{5/2}/4 - c^3n^{3/2})} \\ &< c\sqrt{n}\frac{8n}{3n-4c^2} \end{aligned} \tag{4}$$

Let $k = c\sqrt{n}\frac{8n}{3n-4c^2}$. Then we will prove by induction that $\tau_n(i) < k + \frac{1}{q_i}$, for all $i < \frac{n}{2}$. By (4), this is trivially true for $\frac{n}{2} - c\sqrt{n} < i < \frac{n}{2}$. For our inductive step, we first note that $k\frac{p_i}{q_i} < k - \frac{8}{3}$ and $\frac{1}{q_{i+1}}\frac{p_i}{q_i} < \frac{8}{3}$ for $i < \frac{n}{2} - c\sqrt{n}$, $c > 1$. Then assuming $\tau_n(i+1) < k + \frac{1}{q_{i+1}}$, we have

$$\begin{aligned}\tau_n(i) &= \tau_n(i+1)\frac{p_i}{q_i} + \frac{1}{q_i} \\ &< k\frac{p_i}{q_i} + \frac{1}{q_{i+1}}\frac{p_i}{q_i} + \frac{1}{q_i} \\ &< k - \frac{8}{3} + \frac{8}{3} + \frac{1}{q_i}\end{aligned}$$

Now, we note that $\frac{1}{q_i} < \frac{4n}{3i} < \frac{\sqrt{n}}{3}$ for $4\sqrt{n} \le i \le \frac{n}{2}$ and that $\tau_n(i) = -\tau_n(n-i+1)$ for all i. Therefore, $\tau_n(i) < 3\sqrt{n}$ for $4\sqrt{n} < i < \frac{n}{2} - 4\sqrt{n}$.

Let $\hat{\tau}$ be the stopping time when the random walk reaches either $4\sqrt{n}$ or $n - 4\sqrt{n}$. By Azuma's inequality,

$$\begin{aligned}\mathbf{Pr}[|T_n(4\sqrt{n}) + \hat{\tau} - T_n(R_0)| \ge \lambda] &< e^{-\lambda^2/\sum_{k=1}^{\hat{\tau}} c_k^2} \\ &= e^{-\lambda^2/9\tau n}\end{aligned}$$

Then, for $c' \ge 1$,

$$\begin{aligned}\mathbf{Pr}[\hat{\tau} \ge 11c'n\ln n] &\le \mathbf{Pr}[\hat{\tau} \ge n\ln n + \sqrt{9n\hat{\tau}c'\ln n}] \\ &\le \mathbf{Pr}[\hat{\tau} \ge T_n(R_0) + \sqrt{9n\hat{\tau}c'\ln n}] \\ &\le \mathbf{Pr}[|T_n(4\sqrt{n}) + \hat{\tau} - T_n(R_0)| \ge \sqrt{9n\hat{\tau}c'\ln n}] \\ &\le e^{-(\sqrt{9n\hat{\tau}c'\ln n})^2/9\hat{\tau}n} \\ &= n^{-c'}\end{aligned}$$

To finish, we just need to bound the time it takes for the random walk to go from $i = 4\sqrt{n}$ to 0. Note that conditioned on the random walk moving either to the left or right, there is a strong drift to the boundary. For states $i < 8\sqrt{n}$, there is probability of at most $\frac{8}{\sqrt{n}}$ of moving right and $1 - \frac{9}{\sqrt{n}}$ of moving left. It is not hard to bound this by $11n\ln n$ with high probability. □

2.1. Triadic Consensus on a Line

The case of a line was proved in Goel and Lee [2012], but with only an upper bound of $O(n \log^2 n)$ for the number of triadic interactions required. Here, we repeat their proof, but also prove their conjecture that $O(n \log n)$ triadic interactions is enough.

A line is a graph whose vertices are indexed from 1 to n, and for which a vertex x is a neighbor of y if $|x - y| = 1$. Surprisingly, the winning distribution of Triadic Consensus is a clean expression.

THEOREM 2.2. *Suppose that TriadicMechanism returns the* 1*-median of any three participants, and suppose that the* n *participants form a line. Then Triadic Consensus (for* $k = 1$*) will produce a winner* w *such that*

$$\mathbf{Pr}[w = i] = \left(\frac{1}{2}\right)^{n-1} \binom{n-1}{i-1}$$

PROOF. Color balls with ids $1, 2, \ldots, i$ red and balls with ids $i+1, i+2, \ldots, n$ blue. Whenever a ball is relabeled due to Algorithm 1, we will also recolor it according to the rule above.

Since the 1-median of three balls on a line is simply the standard median, the winning proposal returned from every triad must have the majority color of the three sampled balls. Then the probability that a red ball is recolored blue is $3p_r p_b^2$, where p_r and p_b are the fraction of red and blue balls respectively. Similarly, the probability that a blue ball is recolored red is $3p_r^2 p_b$. Then the dynamics of recoloring is exactly the same as that of a fixed size urn with urn function $f(p) = 3p(1-p)^2$.

Note that the event $w \le i$ is exactly the probability of convergence to red urn. Applying Theorem 2.1 gives us expressions for $\mathbf{Pr}[w \le i]$ and $\mathbf{Pr}[w \le i-1]$. We now get our desired expression by subtracting the second expression from the first. □

2.2. Triadic Consensus on a Grid

A grid is a graph whose vertices are indexed by d dimensions, each of which takes an integer from 1 to n_i. A vertex x is a neighbor of y if there is some dimension i for which $|x_i - y_i| = 1$. A square grid is a grid for which $n_i = n^{1/d}$ for all dimensions i and has a global 1-median of $(b{+}1, b{+}1, \ldots, b{+}1)$, where $n^{1/d} = 2b{+}1$[2]. The convergence properties on a grid are quite counter-intuitive. One would not expect that such a random process like Triadic Consensus could elicit the exact 1-median of arbitrarily sized square grids. However, this is exactly what we show.

THEOREM 2.3. *Suppose that TriadicMechanism returns the* 1*-median of any three participants, and suppose that the* n *participants form a* d*-dimensional square grid. Then Triadic Consensus (with* $k = 1$*) produces a winner* w *such that for* $d = 2$*,*

$$\mathbf{Pr}[\forall i, |w_i - (b+1)| \le c] \ge 1 - 4 \cdot \textit{exp}(-(.5 + c)^2)$$

and for $d \ge 3$*,*

$$\mathbf{Pr}[w = (b+1, b+1, \ldots, b+1)] \ge 1 - 2d \cdot \textit{exp}\left(-\frac{1}{4} n^{\frac{d-2}{d}}\right)$$

PROOF. For a given dimension i, we can use the same proof method as Theorem 2.2. Color all balls with $x_i < (b+1) - c$ red and all balls with $x_i \ge (b+1) - c$ blue. It is not hard to verify that the winning color of any triad will always be the majority color.

[2]We will assume b is an integer for simplicity.

Then we can use Theorem 2.1 to get

$$\mathbf{Pr}[w_i < (b+1) - c] = \left(\frac{1}{2}\right)^{n-1} \sum_{j=0}^{(b-c)\frac{n}{n^{1/d}}-1} \binom{n-1}{j}$$

Since binomials can be interpreted as resulting from sequences of coin flips, we can apply Chernoff bounds to get

$$\mathbf{Pr}[w_i < (b+1) - c] \leq e^{-m^2/2\mu}$$

where $m = (c+1/2)\frac{n}{n^{1/d}}$ and $\mu = n/2$. Plugging in the appropriate values and applying union bounds twice for each dimension gives us our desired expression. □

2.3. Triadic Consensus on a Tree

A tree is a connected graph without any cycles. A d-ary tree of height h is a tree where each node has d children, and the maximum distance from the three is h. For n nodes, the height of a d-ary tree is $\sim \log_d n$. It is not hard to see that the root of the tree (the only node with height zero) is the global 1-median. Again, our intuition is surprised by the revelation that Triadic Consensus gets the exact 1-median of arbitrarily sized d-ary trees (or is one off in the case of $d = 2$).

THEOREM 2.4. *Suppose that TriadicMechanism returns the* 1*-median of any three participants, and suppose that the* n *participants form a* d*-ary tree. Then Triadic Consensus (with* $k = 1$*) produces a winner* w *such that for* $d = 2$*,*

$$\mathbf{Pr}[\mathit{height}(w) \leq 1] \geq 1 - 4 \cdot \mathit{exp}\left(-\frac{n}{16}\right)$$

and for $d \geq 3$*,*

$$\mathbf{Pr}[w = \mathit{root}] \geq 1 - d \cdot \mathit{exp}\left(-\frac{n}{4}\left(\frac{d-2}{d}\right)^2\right)$$

PROOF. For a given edge $e = (x, y)$, we can use the same proof method as Theorem 2.2. Color all balls that are closer to x red and all the balls closer to y blue. It is not hard to verify that the winning color of any triad will always be the majority color. Now let $\rho(i, j)$ denote the number of nodes closer to j than to i and let $N(i)$ denote the neighbors of i. Then we can use Theorem 2.1 to get

$$\mathbf{Pr}[w = i] = 1 - \left(\frac{1}{2}\right)^{n-1} \sum_{j \in N(i)} \sum_{k=0}^{\rho(i,j)-1} \binom{n-1}{k}$$

For the root node r and any of his neighbors s, we know that $\rho(r, s) = (n-1)/d$. Then we can simply plug this in to the above expression and use Chernoff bounds to get our desired result for $d \geq 3$. The case of $d = 2$ is similar, except that we calculate the probability for a node at a height of 1. □

2.4. Triadic Consensus on a Median Graph

Before going into the main results, we will give some background on median graphs, define some notation, and list some lemmas that we will use.

2.4.1. Median Graphs. Define the interval I_{xy} between points x and y to be the set containing all points lying on a shortest path between x and y, i.e. $I_{xy} = \{w \mid d(x, y) = d(x, w) + d(w, y)\}$. Define a graph to be a median graph if, for every x, y, z, $|I_{xy} \cap I_{xz} \cap$

$I_{yz}| = 1$, i.e. there is exactly one point which lies on some shortest path between x and y, x and z, and y and z. Median graphs are a fairly broad class of graphs including trees, multi-dimensional grids, squaregraphs, and distributive lattices. See [Bandelt and Chepoi 2008] and [Knuth 2011] for surveys on the median graph literature.

It is interesting to note that, for median graphs that do not contain cubes, the Condorcet winner always exists and coincides with the 1-median. Indeed, for all median graphs, it is known that whenever the Condorcet winner exists, it coincides with the 1-median [Saban and Stier-Moses 2012]. So in some sense, our results on approximating the 1-median can also be interpreted as results on an approximation of the Condorcet winner in median graphs.

2.4.2. Notation and Lemmas

Definition 2.5. For $u, v \in V$, define the win sets of u and v, $W_{uv} = \{w \in V \mid d(w,u) < d(w,v)\}$, $W_{vu} = \{w \in V \mid d(w,v) < d(w,u)\}$, $W_{u=v} = \{w \in V \mid d(w,u) = d(w,v)\}$, to be the set of nodes that are closer to u, closer to v, or equidistant to them respectively.

Definition 2.6. Define the median $m(x, y, z)$ of three points x, y, z to be the point (or set of points) that lies on some shortest path between x and y, y and z, and x and z.

Definition 2.7. Define a set S to be convex, if for every $x, y \in S$, all the shortest paths between x and y lie in S.

LEMMA 2.8. *For any x, y, z in a median graph, $m(x, y, z)$ a 1-median for x, y, z.*

LEMMA 2.9. *Let x, y, z be nodes in a median graph. Then $m(x, y, z)$ is a Condorcet winner for x, y, z. That is, for every other node w, $|W_{mw}| \geq |W_{wm}|$.*

PROOF. Consider any node w. If $|W_{wm}| = 0$, then this is trivially true. Now suppose $|W_{wm}| > 0$. Without loss of generality, suppose $x \in W_{wm}$, i.e. $d(x, w) < d(x, m)$.

By the definition of the median m, we know that $y \to m \to x$ must be a shortest path, so $d(y, w)+d(w, x) \geq d(y, m)+d(x, m) \implies d(y, w) \geq d(y, m)+d(x, m)-d(x, w) > d(y, m)$. By the same argument, $d(z, w) \geq d(z, m)$. Therefore, $|W_{mw}| = 2 > 1 = |W_{wm}|$. □

LEMMA 2.10. *For any convex set S in a median graph and any nodes x, y, z, two of which are in S, $m(x, y, z) \in S$.*

LEMMA 2.11. *For any median graph G with n nodes, there is some isometric (i.e. distance preserving) embedding $\phi : V^G \to V^H$ of the nodes in G to the nodes of a hypercube H such that*

(*1*) *The dimension of the hypercube is at most n.*
(*2*) *$\phi(m(x, y, z)) = m(\phi(x), \phi(y), \phi(z))$, i.e. the median of points x, y, z in G is the same as in H.*
(*3*) *For any dimension i of the hypercube, the set of participants $\{p \mid \phi(p)_i = 0\}$ and $\{p \mid \phi(p)_i = 1\}$ is convex.*

Definition 2.12. Given a set of participants $p_1, p_2, \ldots, p_n$,

(1) Let $D(x)$ denote the sum of distances from x to all participants p_i. Specifically, $D(x) = \sum_{i=1}^{n} d(x, p_i)$.
(2) Let w^* denote the 1-median, i.e. $w^* = \min_x D(x)$.
(3) Let p_{ij} denote the j-th bit of the embedding of p_i into the hypercube described in Lemma 2.11.
(4) Let $N(b, j)$ denote the number of participants whose j-th bit is equal to the bit b.

2.4.3. Main Result

THEOREM 2.13. *Suppose that TriadicMechanism returns the 1-median of any three participants, and suppose that the n participants form a median graph. Then Triadic Consensus (with $k = 1$) produces a winner w such that,*

$$D(w) \leq \left(1 + 6\sqrt{\frac{2\ln n}{n}}\right) D(w^*)$$

PROOF. For a given dimension i of the hypercube embedding, we can use the same proof method as Theorem 2.2. Color all balls of participants where $p_{ij} = 0$ red, and all balls of participants where $p_{ij} = 1$ blue. Then because of Lemma 2.10, Lemma 2.11, and Lemma 2.8, we know that the winning color of any triad will always be the majority color.

Now suppose that $N(b, j) \leq \frac{n}{2} - \sqrt{2n \ln n}$ for some bit value b at dimension j. Then by Theorem 2.1, we know that

$$\begin{aligned}
\mathbf{Pr}[w_j = b] &\leq \left(\frac{1}{2}\right)^{n-1} \sum_{j=1}^{\frac{n}{2}-\sqrt{2n\ln n}} \binom{n-1}{j-1} \\
&\leq \exp\left(-\frac{1}{2} \cdot \frac{n}{2} \cdot \left(\frac{\sqrt{2n\ln n}}{n/2}\right)^2\right) \qquad (5) \\
&= \frac{1}{n^2}
\end{aligned}$$

where (5) is by applying Chernoff's bound. Then by applying the union bound to each dimension of w (which is at most n by Lemma 2.11), we have

$$\mathbf{Pr}[N(w_j, j) > \frac{n}{2} - \sqrt{2n\ln n}, \ \ \forall j] \geq 1 - \frac{1}{n}$$

Therefore, we have that with probability at least $1 - \frac{1}{n}$, $N(w_j, j) > \frac{n}{2} - \sqrt{2n \ln n}, \ \ \forall j$. All we need to show is that this condition implies that all these participants are approximate 1-medians.

To do this, we first note that the distance $d(x, y)$ from x to y can be written in a simple form because of the isometric hypercube embedding (Lemma 2.11) of median graphs. Namely,

$$d(x, y) = \sum_{i=1}^{d} \mathbb{1}_{x_i \neq y_i}$$

This means that,

$$D(x) = \sum_{i=1}^{n} d(x, p_i) = \sum_{i=1}^{n}\sum_{j=1}^{d} \mathbb{1}_{x_j \neq p_{ij}} = \sum_{j=1}^{d}(n - N(x_j, j))$$

Now, note that for any p_i satisfying $N(p_{ij}, j) > \frac{n}{2} - \sqrt{2n \ln n}$, it must be true that for any bit b,

$$n - N(p_{ij}, j) \leq (n - N(b, j)) \cdot \frac{n/2 + \sqrt{2n\ln n}}{n/2 - \sqrt{2n\ln n}}$$

But then for these p_i, it must be true that

$$\begin{aligned} D(p_i) &\le \sum_{j=1}^{d}(n - N(w_j^*, j)) \cdot \frac{n/2 + \sqrt{2n\ln n}}{n/2 - \sqrt{2n\ln n}} \\ &< \left(1 + \frac{2\sqrt{2n\ln n} + \sqrt{2n\ln n}}{n/2 - \sqrt{2n\ln n} + \sqrt{2n\ln n}}\right) D(w^*) \\ &= \left(1 + 6\sqrt{\frac{2\ln n}{n}}\right) D(w^*) \end{aligned}$$

and we are done. □

3. STRATEGIC CONSIDERATIONS AND A THREE PERSON MECHANISMS

We will first prove that our urn reduction *preserves strategy-proofness* in a formal sense. We show that, as long as the TriadicMechanism satisfies a specified notion of strategy-proofness, we will be able to find a Nash equilibrium for the overall n person urn reduction which returns the 1-median in each triad. Following this proof, we will give an example of a TriadicMechanism which satisfies the needed property as an illustration.

3.1. Triadic Consensus inherits strategy-proof properties from the TriadicMechanism

The main method in our reduction is to couple two urns together, one in which the agent follows the Nash equilibrium, and one in which the agent deviates. We will then show that, in every coupled history, the agent performs better in the urn in which he follows the Nash equilibrium. We first give the following supporting definition and lemma.

Definition 3.1. Given two urns R and S with n balls labeled $r_1, r_2, \ldots, r_n$ and $s_1, s_2, \ldots, s_n$ respectively, R x-dominates S if $r_i \in I_{xs_i}$ for all i.

In other words, R x-dominates S if for every pair of balls r_i, s_i there is some shortest path from x to s_i which contains r_i.

LEMMA 3.2. *Suppose that R x-dominates S. Then the 1-median of r_i, r_j, r_k must be on some shortest path from x to the 1-median of s_i, s_j, s_k.*

PROOF. Consider the isometric hypercube embedding of the median graph, and let x be the zero vector without loss of generality. By Lemma 2.10 and Lemma 2.11, we know that the hypercube embedding of the median of any three points can be found by simply taking the maximum bit for each dimension.

Let m denote the 1-median of r_i, r_j, r_k and m' denote the 1-median of s_i, s_j, s_k. For any dimension d for which $s_{id} = 0$, it must also be true that $r_{id} = 0$. Therefore, if $m'_d = 0$, it must also be true that $m_d = 0$. Because the embeddings are isometric, it follows that m is on a shortest path from x to m'. □

THEOREM 3.3. *Consider a TriadicMechanism which returns the true 1-median of the triad if all agents follow some specified strategy t. For some agent x, let m denote the resulting winner of the triad if x follows strategy t, and let m' denote the resulting winner of the triad given some deviation of x from the strategy. Suppose that $m \in I_{xm'}$ whenever the other two agents follow t. Then playing t in every triad is a Nash equilibrium for the overall urn process.*

PROOF. Since we are proving a Nash equilibrium result, we can assume that all other agents are playing according to the strategy t. Our proof strategy will be to use a coupling argument. We consider two urns R and S. In urn R, x plays according to the given strategy t. In urn S, x plays according to any other strategy. In the initial configuration, we index each ball in R as $r_1, r_2, \ldots, r_n$ and each ball in S as $s_1, s_2, \ldots, s_n$. We index it in a way such that each pair of balls r_i, s_i correspond to the same participant. We now couple the balls drawn in Alg. 1 so that when r_i, r_j, r_k are randomly drawn from urn R, balls s_i, s_j, s_k will be drawn from urn S.

Suppose that at some time, R x-dominates S and then each undergoes a coupled TriadicMechanism where balls r_i, r_j, r_k are selected from R and s_i, s_j, s_k are selected from S. After the TriadicMechanism conclude, we show that the resulting urns R' and S' must still satisfy R' x-dominates S'.

Case 1: x is drawn twice in urn S. Since R x-dominates S, it must be true that x is also drawn twice in urn R. Then all the balls r_i, r_j, r_k and s_i, s_j, s_k will simply be relabeled with x automatically, so R' trivially x-dominates S'.

Case 2: x is not drawn in urn S. Suppose that r_i, r_j, r_k are drawn from urn R and s_i, s_j, s_k are drawn from urn S. Since x is not drawn in urn S, and all other agents follow the strategy t, the returned ball in urn S is the 1-median of s_i, s_j, s_k. Also, since all agents follow t in urn R (including x), it must be true that the returned ball is the 1-median of r_i, r_j, r_k. Then by Lemma 3.2, R' must still x-dominate S'.

Case 3: x is drawn once in urn S. Suppose that r_i, r_j, r_k are drawn from urn R and s_i, s_j, s_k are drawn from urn S. Let m denote the winner in urn S which is returned if x follows strategy t and m' denote the winner in urn S which is returned given some deviation. By our reduction assumption, it must be true that m lies on some shortest path from x to m'. Since all agents follow t in urn R, we know that the winner in urn R must be the median of r_i, r_j, r_k. By Lemma 3.2, the winner in urn R must lie on some shortest path from x to m. But this also means that the winner in urn R lies on some shortest path from x to m', so R' must still x-dominate S'.

Finally, we note that before any TriadicMechanism take place, R and S are identical, which means that R x-dominates S. Then, the winner of R must also x-dominate the winner of S, which means that urn R is better for x in every coupled history. □

3.2. Example TriadicMechanisms and a Nash equilibrium result

We give an example mechanism for each triad which results in a Nash equilibrium for the overall n person process that produces the 1-median in each triad. We highlight the fact that this result follows from Theorem 3.3, which shows that Algorithm 1 is able to inherit certain strategy-proof properties from the given TriadicMechanism. If the given TriadicMechanism is not suited to the particular assumptions of a given situation, one can always design another TriadicMechanism which satisfies the required properties.

The mechanism we give (Algorithm 2) can be thought of in terms of two "proposers", who are each trying to find proposals close to themselves that a "chooser" would also vote for. When a triad is randomly selected, the mechanism arbitrarily assigns one of them to be the chooser, which leaves the remaining two as the proposers. At the start of the mechanism, an arbitrary proposer is chosen to be the starting proposer, and the proposal of the other proposer is set as the "current winning proposal". The starting proposer can then suggest an alternate proposal to the chooser. If the chooser accepts this suggestion, then this suggestion will become the new "current winning proposal", and the other proposer is now given a chance to suggest an alternative. This process repeats until the chooser decides not to accept the given suggestion, or until the proposer whose turn it is decides to pass.

When participants are in a median graph, there is an intuitive strategy (Algorithm 3) for this mechanism which is a Nash equilibrium for the overall urn process, and will also produce the 1-median of the three participants in the triad. The chooser will just respond truthfully, accepting an alternative if he prefers it to the current winning

ALGORITHM 2: Chooser-Proposer TriadicMechanism

Input: Participant x, y, z from a set of n proposals
Output: A winner from the set of all proposals
Let x be the "chooser", and let y and z be the "proposers";
Initialize the current winner w to y and the current proposer to be z;
while *true* **do**
    The current proposer chooses any proposal not equal to w, or $\emptyset$ (call this w');
    if $w' \neq \emptyset$ *and the chooser accepts* w' **then**
        $w = w'$;
        Set the current proposer to be y if it is currently z and vice versa;
    else
        return w;

ALGORITHM 3: A strategy for the Chooser-Proposer TriadicMechanism

Input: p as one of x, y, z in Alg. 2
Output: The appropriate action (depending on whether p is a chooser or proposer)
if p *is the chooser* **then**
    return 'accept' if and only if p prefers candidate w' over w;
else
    if *There is no proposal that is closer to* p*, and lies on a shortest path between* w *and the chooser* **then**
        return $\emptyset$;
    else
        return the closest proposal to p which lies on a shortest path between w and the chooser;

proposal, and rejecting it otherwise. Any proposer will suggest the alternative that is closest to himself out of the set of proposals which lie on some shortest path from the chooser to the current winning proposal. If there is no alternative that satisfies this property, then he will pass.

THEOREM 3.4. *When Triadic Consensus (Algorithm 1) is run with the Chooser-Proposer TriadicMechanism (Algorithm 3), then Algorithm 3 is a Nash equilibrium which produces the* 1*-median of every triad.*

PROOF. We will break down our proof into two parts. We first prove that Algorithm 2 will find the 1-median of the three participants, assuming that all agents follow Algorithm 3 as their strategy (Lemma 3.5). Then, we show that, if there is one participant who deviates from this strategy, a winner w will be produced, such that the 1-median of the three participants will lie on a shortest path between the deviating participant and w (Lemma 3.6). We can then directly apply Theorem 3.3 to get our result. □

LEMMA 3.5. *Suppose that three participants in a round of the Chooser-Proposer TriadicMechanism (Algorithm 2) all follow Algorithm 3 as their strategy. Then the winner returned will be the* 1*-median of these participants.*

PROOF. Let c denote the chooser, and p_1, p_2 denote the proposers. Let p_1 denote the first proposer, which means that the current winning proposal will be set to p_2 at the beginning of the mechanism. Then p_1 will propose the closest point to him that lies in the interval from c to p_2. Because the participants are in a median graph, this is simply the 1-median of p_1, c, p_2, which will be accepted by c. Now, in p_2's turn, there cannot be any other alternatives that are closer to both c and p_2 since the 1-median of p_1, c, p_2 is a Condorcet winner of the three points. Therefore, p_2 will pass, and the process ends. □

LEMMA 3.6. *Suppose that three participants in a round of the Chooser-Proposer TriadicMechanism (Algorithm 2) all follow Algorithm 3 as their strategy, except for one participant who deviates. Then the* 1*-median of these participants will lie on a shortest path between the deviating participant and the resulting winner.*

PROOF. Let c denote the chooser, and p_1, p_2 denote the proposers. Let p_1 denote the first proposer, which means that the current winning proposal will be set to p_2 at the beginning of the mechanism. We will show that, at the beginning of p_1's turn to propose, the current winning proposal must be in the interval of c and p_2 in addition to being in the interval of p_1 and p_2. Similarly, at the beginning of p_2's turn to propose, the current winning proposal must be in the interval of c and p_1 in addition to being in the interval of p_1 and p_2. We prove this by induction.

Case One: The chooser (c) deviates from Algorithm 3. If c deviates in the first round when p_1 proposes the 1-median, then the resulting winner will be p_2, which has a shortest path containing the 1-median. If c does not deviate in the first round, then the TriadicMechanism will end in the second round since p_2 will pass, so c does not have any other opportunities to deviate.

Case Two: One of the proposers deviates from Algorithm 3. Let p_s denote the deviator, and p_t denote the other proposer. We will show that, regardless of what p_s does, p_t will always propose an alternative in the interval of p_t and c, which is accepted. Suppose that p_s proposes some alternative a, which is accepted. Then, since p_t is following Algorithm 3, he will propose the closest alternative to himself which is also in the interval of c and a. This alternative is the median of c, a, and p_t, which is in the interval of c and p_t. Note that whenever a is not in the interval of c and p_t, this alternative will always be unique to a, and preferred by both c and p_t, which means that it will get accepted. Therefore, if p_s ever passes, the resulting winner must be in the interval of c and p_t. But this process must eventually converge since every accepted step goes up c's ranking of alternatives. Therefore, either p_t or p_s will pass, or someone will propose an alternative that is rejected by c. Both of these cases will result in a winner that lies on the interval of p_t and c. Since the 1-median is the gate of p_s into the interval of p_t and c, it must be true that the 1-median lies on some shortest path from p_s to the winning alternative, and we are done. □

4. AN IMPOSSIBILITY RESULT FOR CONSENSUS REDUCTIONS USING PAIRWISE INTERACTIONS

In this section, we consider the possibility of designing consensus reductions for the 1-median problem using interactions between pairs of participants. Surprisingly, we find that for several consensus reductions, this is not possible for any two person mechanism satisfying a consistency property. We introduce this with a simple example.

Example 4.1. Consider a natural variant of Triadic Consensus, in which everything proceeds identically, except that two balls (as opposed to three) are selected at random at each step (Alg. 4). The two chosen participants take part in a two person mechanism in which they are made to decide on some winning proposal. The sampled balls are then relabeled with the proposal chosen, and the process is repeated until the urn contains balls of a single color.

Suppose that the $n = 2k+1$ participants form a tree G, where G is the graph resulting from taking the center of a star S with k leaves and connecting it to one of the endpoints of a path P of length k. It is easy to see that the 1-median is the center node of the star.

ALGORITHM 4: Dyadic Consensus

Input: An urn with k labeled balls for each participant
Output: A winning participant
while *there is more than one label* **do**
  Sample balls x, y uniformly at random with replacement;
  $w = \text{DyadicMechanism}(x, y)$;
  Relabel all the sampled balls with the winning label w;
return the id of the remaining balls;

Each time two people are chosen to participate in a DyadicMechanism, we assume that the mechanism is designed such that they come to consensus on a median proposal on a shortest path btween the two participants[3].

It turns out that the winning outcome will be highly concentrated around a point on path P, which is a distance of $\frac{n}{8}$ from the 1-median, and is a $\frac{9}{8}$ approximation. Note that a constant factor approximation is very bad since even the worst proposal is a 3 approximation. The proof of this is simple. Let

$$f(x) = \begin{cases} -1 & \text{if } x \text{ is a leaf of the star} \\ d(x,c) & \text{otherwise} \end{cases}$$

where c is the center of the star, and $d(x, c)$ is the distance of x from c. Let $F_t = \sum_x f(x)$ where the summation is over all proposals at time t. Then it is easy to verify that $\mathbb{E}[F_{t+1}] = \mathbb{E}[F_t]$ for any t. Therefore, at convergence time $\hat{\tau}$, we have $\mathbb{E}[F_{\hat{\tau}}] = \mathbb{E}[F_0] \approx \frac{n^2}{8}$. Since there are n points at a single position, the expected position must be at $\frac{n}{8}$ distance from the center of the star. It is not hard to show that the result is actually concentrated at that point.

In order to study general two person mechanisms, we first define a natural property on mechanisms.

A set of proposals S is convex if every shortest path between $x, y \in S$ lies in S. The convex hull of a set of proposals $C(S)$ is defined as the smallest convex set containing S. Two sets of participants S and S' are locally consistent if the there is a distance preserving isomorphism between their convex hulls which map the proposals lying in $C(S)$ to those in $C(S')$. Intuitively, if two sets S and S' are locally consistent, it means that the structure of their relationships to each other are the identical, regardless of the global space that they lie in.

We call a mechanism locally-consistent if, for two locally consistent sets of participants S and S', the mechanism returns the same probability distribution over outcomes under the isomorphism.

THEOREM 4.2. *For Alg. 4, there is no locally-consistent two person mechanism which can do better than an expected constant approximation of the* 1*-median when the proposal space is as a tree.*

PROOF. The intuition of the proof is the same as the example given. Suppose that the proposal space is the same tree given in the example, where all the weights are uniform. In this setting, the convex hull of two participants is simply the unique path between them. It is easy to argue that the probability distribution of the outcome of any two person locally-consistent mechanism must be symmetric about the median of the path. Because of this, the same argument shows that the expected winner will be the point that is $\frac{n}{8}$ distance on the path from the center of the star. □

[3]If there are two median points (for odd-length paths), we assume that they will return a random one.

In fact, this result can be strengthened. Suppose that instead of choosing the two balls at each step randomly, we allowed an arbitrary policy for deciding what two balls are chosen. Even in this general case, the expected approximation factor of the 1-median cannot be better than a constant.

5. DISCUSSION AND FUTURE DIRECTIONS

5.1. Experiments

One of the interesting directions to take this work is to take an experimental approach to the design of crowd algorithms. If one could build an experimentally accurate model of consensus dynamics in small groups of participants for different mechanisms, then these models could be used as black box units in algorithm design regardless of theoretical predictions.

An illustrative example of this idea can be found in the experiments of Fiorina and Plott [1978] on Majority Rule in the euclidean spatial voting setting. Despite predictions of chaos [McKelvey 1976] in situations when a Condorcet winner does not exist, experiments show that groups of five participants are able to quickly converge to a "central" point [Grofman et al. 1987]. If such results are robust, one could then use Majority Rule in small groups as building blocks for eliciting a central proposal from the crowd.

5.2. When participants may not be representative of proposals

One of the assumptions that may seem somewhat constraining is that we are assuming that participant preferences are representative of distances from their proposal to other proposals. In some practical settings, this may be an invalid assumption.

In the "wisdom of the crowd" context, this assumption may not be that constraining. Recall that our goal is to find the 1-median of the proposals, which are the estimates produced by the crowd. The only reason we needed our assumption is that we needed a way to get a handle on the distances between the proposals. If strategic considerations are put aside, one could imagine getting around this restriction by having independent agents "act out" consensus between three proposals. For example, instead of asking whether they prefer y or z, they could be asked whether y or z is closer to x. Similarly, they could be asked to think of a compromise proposal close to x but between y or z. Even in the strategic setting, it may be possible to override the agents individual preferences by compensating them with money for the closeness of the resulting winner to the proposal they are acting on behalf of.

5.3. Other directions

Another natural way to extend this work includes generalizing the reduction results of Triadic Consensus past median graphs. For instance, the NP-hard problem of finding the Kemeny-Young ranking is equivalent to finding the 1-median of a set of rankings under the Kendall-tau distance metric. It would be interesting to prove something on the approximation properties of Triadic Consensus in this setting.

Finally, we want to incorporate psychological phenomena into our model. Since the wisdom of the crowd is often affected by social influences [Lorenz et al. 2011], it will be important to understand how algorithms can be designed to counter such effects.

ACKNOWLEDGMENTS

The authors would like to thank Ivy Lee for providing the figures, and the anonymous reviewers at EC for their constructive feedback.

REFERENCES

BANDELT, H.-J. AND BARTHELEMY, J.-P. 1984. Medians in median graphs. *Discrete Applied Mathematics 8,* 2, 131–142.

BANDELT, H.-J. AND CHEPOI, V. 2008. Metric graph theory and geometry: a survey. *Contemporary Mathematics 453*, 49–86.

BRABHAM, D. C. 2008. Crowdsourcing as a model for problem solving: an introduction and cases. *Convergence 14,* 1, 75–90.

BRANDT, F., CONITZER, V., AND ENDRISS, U. 2012. Computational social choice. In *Multiagent Systems*. MIT Press.

CONITZER, V. 2009. Eliciting single-peaked preferences using comparison queries. *Journal of Artificial Intelligence Research 35*, 161–191.

CONITZER, V. AND SANDHOLM, T. 2002. Vote elicitation: Complexity and strategy-proofness. In *Proceedings of the 17th AAAI Conference*.

CONITZER, V. AND SANDHOLM, T. 2005. Communication complexity of common voting rules. In *Proceedings of the 6th ACM Conference on Electronic Commerce (ACM-EC)*. 78–87.

CONITZER, V. AND SANDHOLM, T. 2013. Common voting rules as maximum likelihood estimators. In *Proceedings of the 21st Annual Conference on Uncertainty in Artificial Intelligence (UAI)*.

FIORINA, M. P. AND PLOTT, C. R. 1978. Committee decisions under majority rule: An experimental study. *The American Political Science Review 72,* 2, 575–598.

GALTON, F. 1907. Vox populi. *Nature 75*, 450–451.

GOEL, A. AND LEE, D. T. 2012. Triadic consensus: A randomized algorithm for voting in a crowd. In *Proceedings of the 8th Workshop on Internet and Network Economics (WINE)*.

GROFMAN, B., OWEN, G., NOVIELLO, N., AND GLAZER, A. 1987. Stability and centrality of legislative choice in the spatial context. *The American Political Science Review 81,* 2, 539–553.

KHATIB, F., COOPER, S., TYKA, M. D., XU, K., MAKEDON, I., POPOVIČ, Z., BAKER, D., AND PLAYERS, F. 2011. Algorithm discovery by protein folding game players. *Proceedings of the National Academy of Sciences (PNAS) 108,* 47, 18949–18953.

KNUTH, D. E. 2011. *The Art of Computer Programming*. Vol. 4A: Combinatorial Algorithms, Part 1. Pearson Education, Inc.

LEE, D. T. AND BRUCK, J. 2012. Modeling biological circuits with urn functions. In *Proceedings of the International Symposium on Information Theory (ISIT)*.

LORENZ, J., RAUHUT, H., SCHWEITZER, F., AND HELBING, D. 2011. How social influence can undermine the wisdom of crowd effect. *Proceedings of the National Academy of Sciences 108,* 22, 9020–9025.

LU, T. AND BOUTILIER, C. 2011. Robust approximation and incremental elicitation in voting protocols. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence*.

MAO, A., PROCACCIA, A. D., AND CHEN, Y. 2013. Better human computation through principled voting. In *Proceedings of the 27th Conference on Artificial Intelligence (AAAI)*.

MCKELVEY, R. D. 1976. Intransitivities in multidimensional voting models and some implications for agenda control. *Journal of Economic Theory 12,* 3, 472–482.

SABAN, D. AND STIER-MOSES, N. 2012. The competitive facility location problem in a duopoly: Connections to the 1-median problem. In *Proceedings of the 8th Workshop on Internet and Network Economics (WINE)*.

SUROWIECKI, J. 2005. *The Wisdom of Crowds*. Anchor.

WENDELL, R. E. AND MCKELVEY, R. D. 1981. New perspectives in competitive loca-

tion theory. *European Journal of Operational Research 6,* 2, 174–182.
WOLFERS, J. AND ZITZEWITZ, E. 2004. Prediction markets. *Journal of Economic Perspectives, American Economic Association 18,* 2, 107–126.