Large-Scale Decision-Making via Small Group Interactions: the Importance of Triads[1]

Ashish Goel and David T. Lee

Abstract

We study a framework for large-scale decision-making through small group interactions. In this framework, a crowd of participants interact with each other through a sequence of small group interactions, the composition of which are chosen by an algorithm designer, or in some settings, by nature. We consider the problem of finding the wisdom of the crowd, which we take to be the generalized median, in two settings: opinion formation and strategic bargaining. In both cases, we find a significant difference between groups of two and three.

When the small groups are of size two, we find that there is no sequence of pairwise interactions which can always converge to a non-trivial approximation of the generalized median so long as the small group interactions satisfy a natural property we call local consistency. This holds even in the simple case when participants come from a line.

In contrast, when the small groups are of size three, we find that the generalized median can be tightly and efficiently approximated when participants come from R^d under the l_1 norm or when they come from any median graph, a class of graphs including squaregraphs, trees, and grids. Specifically, suppose that participants of each small group either update their opinions to, or come to consensus on, the generalized median of the group as a result of their interaction. Then by simply choosing the sequence of triads uniformly at random, the process is able to find a $(1+O(\sqrt{\frac{\ln n}{n}}))$-approximation of the global generalized median with high probability. Moreover, this occurs after each participant has only participated in an average of $O(\log^2 n)$ small group interactions. In the strategic setting, we also design a mechanism for the entire extensive form game which implements this behavior under a Nash equilibrium. When participants treat each small group interaction as a separate game, we show that this can be improved to a strong Nash equilibrium.

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 [13]. Based on this observation, he hypothesized that an appropriate aggregation of a crowd's preferences can produce an extremely accurate estimate, an idea that later became known as "the wisdom of the crowd"[30]. This idea has continued to develop and flourish under the umbrella of crowdsourcing, in which simple units of information such as labels, rankings, or predictions are elicited from individuals, and aggregated together (e.g. [20][24][33]).

In the political arena, this idea has also taken root under the auspices of participatory and deliberative democracy, in which decision-making is crowdsourced to the constituents of a governing body through various democratic innovations (e.g. [29][2][22]). One key difference, due to the democratic nature of the domain, is that deliberation, and not just

[1]For an updated version, see http://stanford.edu/ dtlee88/papers/large-scale-decision-making-via-small-groups-comsoc14.pdf

mere information aggregation, has been sought after as an important component to such a process[9][11].

Motivated by this, we consider large-scale decision-making through deliberative processes. Due to the many known limitations of humans in interacting with large groups of individuals or dealing with large amounts of information (e.g. [26][18][28]), we specifically study large scale decision-making through *small group interactions*. In this framework, a global result is achieved through a sequence of small group interactions, the compositions of which are decided by an algorithm designer, or in some cases, by nature. Besides its applicability to deliberative democracies, such a framework may also produce benefits for crowdsourcing in general, in that it allows one to design algorithms which leverage the unique human ability to communicate, bargain, learn, and make creative decisions.

This idea of using small group interactions, while new for information aggregation and the design of algorithms, is a fundamental one in the opinion formation literature. In this context, the sequence of groups can be random, such as in flocking[31], or due to an underlying social network[8].While these models have been studied for their ability to find the wisdom of the crowd, they have either focused on deGroot-like dynamics of repeated averaging for which the wisdom of the crowd is defined as the mean (average) opinion[16], or on Bayesian dynamics, in which the wisdom of the crowd is measured in terms of some ground truth[1]. Also, both of these settings have typically considered the one-dimensional case, where opinions come from the real line. Notably absent from these studies is the case of the median, the quantity which Galton preferred over the mean[13][12]. As he states:

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

and

> "...it may not be sufficiently realized that the suppression of any one value in a series can only make the difference of one half-place to the median, whereas if the series be small it may make a great difference to the mean."

In the first quote, Galton was expressing a concept known in the voting literature as the Condorcet winner, a candidate who receives at least half the votes in any pairwise election. When such a candidate exists, this naturally describes the voice of the people, and is also the optimal choice under a maximum likelihood interpretation[34]. In the second quote, Galton notes that the median is robust to outliers and noise, whereas the mean is not.

In this paper, we consider the ability of small group interactions to find the generalized median when opinions come from either R^d under the l_1 norm or when they come from any median graph, a class of multi-dimensional metric spaces including squaregraphs, trees, and grids. The generalized median is defined as the point which minimizes the sum of distances to the original points, and coincides with the median in the special case when opinions come from a line. It is also robust to outliers[23] and strongly related to the Condorcet winner in that, for any median graph, if a Condorcet winner exists, then it must coincide with the generalized median[27][6]. Our main conclusion is the following: if one would like to find the generalized median through small group interactions, then it is necessary and sufficient to use small groups of size at least three.

1.1 Results summary

We consider two settings, opinion formation and strategic bargaining. In opinion formation, each small group interaction leads participants to update their opinion according to an opinion formation function. In strategic bargaining, all participants start with k tokens, and each small group interaction takes the form of a bargaining dynamic in which the group is required to come to consensus on one participant to support. The chosen participant, who does not have to be one of the small group members, then receives one token from each member of the small group.

For both of these settings, we find a surprising difference when using small group interactions of size two versus those of size three. In a sentence, the generalized median cannot be found using natural pairwise interactions, but can be found in a scalable manner under simple triadic interactions. Specifically, suppose that small groups are of size two and that the small group interactions satisfy a natural property we call *local consistency*. Then there is no sequence of small group interactions which is able to find a non-trivial approximation of the generalized median even for the simple case when participants come from a line.

In contrast, suppose that small groups are of size three and that participants of a small group either update their opinions to, or come to consensus on, the generalized median of the group (in the case of strategic bargaining, this requires the additional assumption, discussed further in Section 4, that the generalized median of every three participants is also represented by a participant). Then by simply choosing the sequence of triads uniformly at random, the process is able to find a $(1+O(\sqrt{\frac{\ln n}{n}}))$-approximation of the global generalized median with high probability. Moreover, this process is scalable to large crowds in that each participant only needs to participate in an average of $O(\log^2 n)$ small group interactions to achieve this approximation. This is significant in that the participant only ever talks to a rapidly vanishing fraction of the total participants, preventing the cognitive overload that would have been experienced if required to talk to all other participants. In the strategic setting, we also design a mechanism for the entire extensive form game which implements this behavior under a Nash equilibrium. When participants treat each small group interaction separately, we show another mechanism which improves this to a strong Nash equilibrium.

Beyond the results themselves, one of the most interesting contributions in our own eyes is the idea of bringing small group interactions into the design of algorithms for crowdsourcing. With the communication capabilities now available through the Internet, it is possible to easily initiate small group interactions between arbitrary individuals in a crowd. The mode of interaction, as well as the composition of groups can all be freely designed. Hopefully, such an approach will produce interesting new algorithms as theoretical tools are brought to bear on complex crowdsourcing tasks, and experimental insights and modeling (of small group interactions) are simultaneously used to inform and motivate algorithm design.

1.2 Other related work

We are not aware of other theoretical work on the design of crowdsourcing algorithms through small group interactions. In practice, however, there exist crowdsourcing systems such as the ESP game[32] which consist of a sequence of small group interactions. Even though their small group interactions do not involve opinion formation or strategic bargaining, one can also imagine modeling the interactions as a function of some participant state. Currently, this game has been analyzed from a game-theoretic perspective[19].

There are also many experiments on decision-making in small groups. In [10], Fiorina and Plott studied majority rule dynamics among groups of five in a euclidean space. They found that the only consistently predictive theory was that when a Condorcet winner exists, it is chosen. But more surprisingly, they show that even when the Condorcet winner does not exist, participants are able to converge to a central point despite the chaos predicted by McKelvey's chaos theorem[25]. This experimental insight regarding a theoretical impossibility result is illustrative of the benefits that experimental work can provide for the design of algorithms via small group interactions.

Though small group interactions do not fall directly into social choice frameworks such as rank or judgment aggregation, our goal is also to aggregate user opinions. As such, many of our results and approaches are inspired from those of the social choice literature. These include the axiomatic approach originating from Condorcet[34], impossibility results such as Arrow's impossibility theorem[3], the goal of reducing cognitive burden on participants[7], and mechanisms to guard against strategic manipulation[14]. Most directly, the ideas of this paper were initially inspired out of an attempt to understand preference elicitation for rank aggregation, as described in Triadic Consensus[15].

1.3 Outline

Following a description of the model (Section 2), our results are divided into two major sections. Section 3 defines a natural property called local consistency and demonstrates the impossibility of finding the generalized median through any such pairwise interactions. In contrast, Section 4 demonstrates that a simple triadic interaction can efficiently generate a close approximation of the generalized median by simply using random interactions. It also shows how one can implement the required triadic interactions under strategic equilibria.

2 Model

2.1 Median graphs and the generalized median

In our paper, we suppose that participants have opinions $x_1, x_2, \dots, x_n \in \mathcal{X}$, where $\mathcal{X}$ can be R^d under the l_1 norm or any median graph. 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. $I_{xy} = \{w \mid d(x,y) = d(x,w) + d(w,y)\}$. A graph is 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. See [5], [21], and [17] for surveys on the median graph literature. Median graphs are interesting because they include many common graphs such as trees, grids, and squaregraphs. We list well-known properties of median graphs which we will use in Appendix A.

Our goal of finding the wisdom of the crowd is technically defined as finding the generalized median of the initial opinions, the point x^* which minimizes $D(x) = \sum_{i=1}^{n} d(x, x_i)$. A $(1+\epsilon)$-approximation of the generalized median is any point x for which $D(x) \leq (1+\epsilon)D(x^*)$. This definition for the wisdom of the crowd generalizes the median on a line (the concept used by Galton), and also is strongly related to the concept of a Condorcet winner in that any Condorcet winner must coincide with the generalized median (in median graphs).

2.2 Small group interactions

We consider a framework in which a sequence of small group interactions take place, during which group "computations" occur. The computation that occurs in each small group interaction is modeled as a function of the current state of the participant and those of the group he interacts with. In this paper, the state of a participant is the current opinion, $x_i \in \mathcal{X}$, that is held. We consider two cases for the computation that occurs: opinion formation and strategic bargaining. In the case of opinion formation, the computation represents influence by the members of the group, and results in an updated opinion. In the case of strategic bargaining, the computation represents a bargaining process in which participants must jointly decide on a participant to support.

Definition 1 (Small group opinion formation)**.** *A small group opinion formation process is represented by the function $f : \mathcal{X} \times \mathcal{X}^k \to P_{\mathcal{X}}$, where $\mathcal{X}$ is the opinion space, k is the size of the small group, and $P_{\mathcal{X}}$ denotes a probability distribution over $\mathcal{X}$.*

$f(x, S)$ is a function which takes an individual's opinion x and the set of individuals S in the small group, and outputs a probability distribution over the new opinion of the individual. At each discrete time step, some small group of participants interact with each other, during which each participant x updates their opinion by drawing from the probability distribution $f(x, S)$. The sequence of interacting participants could be the result of an algorithm, or could arise from natural dynamics such as those existing in a social network. One would like to identify sequences for which opinions are able to converge to the generalized median.

Definition 2 (Small group bargaining)**.** *A small group bargaining process is represented by the function $g : \mathcal{X}^k \rightarrow P_{\mathcal{X}}$, where $\mathcal{X}$ is the opinion space, k is the size of the small group, and $P_{\mathcal{X}}$ denotes a probability distribution over $\mathcal{X}$.*

$g(S)$ is a function which takes the set of individuals in the small group interaction S and outputs a probability distribution over the joint opinion which the group decides on. We consider such small group interactions in the setting where all participants start with k tokens. At every point in time, some small group of participants is chosen (each of which must have one token), and asked to come to consensus on a participant to support. The chosen participant, who does not have to be one of the small group members, then receives one token from each member of the small group, and the process repeats. One would like to design a mechanism that results in a winner which is the generalized median.

We note that, for strategic participants who consider joint strategies that span multiple small groups (an extensive form game), the function g could depend on factors outside of S such as prior small group interactions. Our impossibility result will not apply to such cases. Our possibility results will first consider a function $g(S)$ which achieves the generalized median, and then how that we can implement the same function under a Nash equilibrium.

3 An impossibility result for pairwise interactions

In this section, we consider small group interactions of size two, and ask whether such pairwise interactions can find the generalized median of the initial participant opinions. We find a surprising result: for "natural" small group interactions, as formalized by a mathematical property we describe as *local consistency*, no sequence of pairwise interactions is able to always do better than a trivial (constant) approximation of the generalized median. Our proof for this is constructive: we give a choice of initial participants opinions in $\mathcal{R}$ for which any sequence fails to find a $(1 + o(1))$-approximation.

3.1 Locally consistent small group interactions

As currently defined, the functions f and g are too powerful in their generality. Since f and g are the outputs of a small group interaction, one expects them to conform to certain natural constraints. Before formally defining our notion of local consistency, we first give two examples of *unnatural* behaviors which motivate our definition.

Consider participants on the real line. It would be unnatural for a participant whose opinion lies in $[50, 100]$ to change his opinion to 0 after only interacting with participants whose opinions also lie in $[50, 100]$. Instead, one expects that any opinion update or consensus point should be within the convex hull[2] of the opinions making up the small group.

Another unnatural function would be for a participant to update his or her opinion to the point in the convex hull closest to some external point p (such as the global median).

[2]Define a set S to be convex if for any two points x, y in S, every shortest path between x and y also lies in S. Define the convex hull of a set S to be the smallest convex set that contains S.

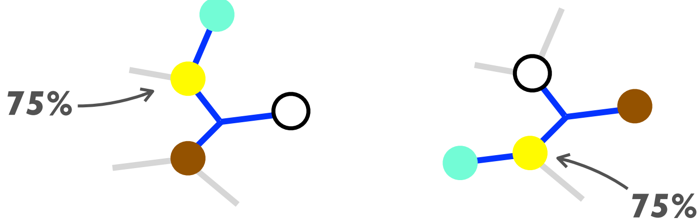

Figure 1: Since the green, brown, and white participants on the left have the same relationship structure to the green, brown, and white participants on the right, they should be expected to return the same participant under that structure, e.g. if yellow is returned with 75% probability on the left, then the corresponding yellow participant on the right should also be returned with 75% probability. Note: the blue region is the convex hull of the green, brown, and white participants.

This is unnatural because p is not part of the decision-making process. One would like to encode the intuition that an opinion update should only depend on the structure of how the opinions of a small group relate to an individual's current opinion.

We encode this through the following mathematical notion which we call local consistency (see Figure 1 for an illustration). Roughly speaking, it says that if two small groups S and S' are equivalent under a rigid "translation", then the result of the bargaining function g should also be equivalent under the same "translation". Similarly, if the relationship of x to S is equivalent to the relationship of x' to S' under a rigid "translation", then the opinion formation function should also be equivalent under the "translation".

Definition 3. *Let S and S' be two sets of participants with convex hulls $C(S)$ and $C(S')$ respectively. S and S' are said to have the same relationship structure under ψ if there is a distance preserving isomorphism $\psi : C(S) \to C(S')$ from the convex set $C(S)$ to $C(S')$ that maps the participants of S onto the participants of S'.*

Definition 4. *The opinion formation function f and the bargaining function g are said to be locally consistent if, for any two sets of participants S and S' with the same relationship structure under ψ, and for any participant $x \in S$ and $x' = \psi(x) \in S'$,*

1. *$f(x, S)$ and $f(x', S')$ are contained in $C(S)$ and $C(S')$ respectively, and are identical under the isomorphism ψ.*
2. *$g(S)$ and $g(S')$ are contained in $C(S)$ and $C(S')$ respectively, and are identical under the isomorphism ψ.*

Example 1. *Consider the opinion formation function in which x updates his opinion based on a weighted average of his current opinion and the average group opinion. Specifically, for any $w \in [0, 1]$, let $f(x, S) = wx + (1 - w)\frac{1}{|S|}\sum_{i \in S} x_i$. Then f is locally consistent. This is essentially deGroot's process, except that all other members of S are treated equally.*

Example 2. *Consider the opinion formation function in which x updates his opinion to a random and uniformly chosen proposal in S. Specifically, for any $y \in S$, let $f(x, S) = y$ with probability $\frac{1}{|S|}$. Then f is locally consistent.*

3.2 The impossibility result

If only two participants are drawn at each step, it turns out that no sequence of locally consistent small group interactions is able to find a $(1 + o(1))$-approximation of the median for all possible initial participant opinions. This is even true of opinions on the real line, for which the median coincides with the concept of a Condorcet winner, the point preferred by at least half of the participants to any other point. This means that one is unable to use small group interactions of size two to find the wisdom of the crowd as measured by either the median or the Condorcet winner. We state the theorem in terms of opinion formation functions, but the results also hold for bargaining functions by noting that any bargaining function $g(S)$ can be interpreted as opinion formation functions $f(x, S)$ where participants move to the same point, i.e. $f(x, S) = g(S)$.

Theorem 1. *Consider a set of points $x_1, x_2, \ldots, x_n \in \mathcal{R}$, and a sequence of sets $S_1, S_2, \ldots$, where $S_i \subset \{1, 2, \ldots, n\}$ and $|S_i| = 2$. Let $x_i^0 = x_i$, and*

$$x_i^{t+1} = \begin{cases} f(x_i^t, S_{t+1}) & \text{if } i \in S_{t+1} \\ x_i^t & \text{otherwise} \end{cases},$$

where f is a locally consistent opinion formation function. Then there exists a choice of x_i such that $\frac{1}{n}\sum_i \mathbb{E}[D(x_i^t)] \geq (\frac{9}{8} + o(1))D(x^*)$ for all t, where x^* is the median of the x_i.

Proof. The first observation to make is that for any two points in $\mathcal{R}$, their convex hull is simply the interval between them. It is easy to then show that any pair of points x, y has the same relationship structure as another pair of points found by translation or reflection. In particular, the interval $[x, y]$ is isomorphic to itself reversed, i.e. $\psi(t) = x + y - t$. The implication of invoking local consistency on this is that $f(x, \{x, y\})$ must be distributed as the reflection of $f(y, \{x, y\})$ about the point $\frac{1}{2}(x + y)$.

Suppose that $n = 2k + 1$ and let $x_i = i$ for $i = 1, 2, \ldots, k$ and $x_i = 0$ otherwise. Then the median $x^* = 0$. Define $X_t = \frac{1}{n}\sum_i x_i^t$ and consider time $t + 1$ for which $S_{t+1} = \{j_1, j_2\}$. Then, $n\mathbb{E}[X_{t+1} \mid x^t] = \sum_i \mathbb{E}[x_i^{t+1} \mid x^t] = \left(\sum_{i \notin S_{t+1}} x_i^t\right) + \mathbb{E}[x_{j_1}^{t+1} \mid x^t] + \mathbb{E}[x_{j_2}^{t+1} \mid x^t] = \left(\sum_{i \notin S_{t+1}} x_i^t\right) + \mathbb{E}[x_{j_1}^{t+1} \mid x^t] + \mathbb{E}[x_{j_1}^t + x_{j_2}^t - x_{j_1}^{t+1} \mid x^t] = nX_t$. Therefore, X_t is a martingale, and for any time t, $\mathbb{E}[X_t] = X_0$. Noting that $D(x)$ is convex, we can apply Jensen's inequality to get $\mathbb{E}\left[\frac{1}{n}\sum_i D(x_i^t)\right] \geq D\left(\mathbb{E}\left[\frac{1}{n}\sum_i x_i^t\right]\right) = D(X_0)$. It is not hard to verify that $D(X_0) = (\frac{9}{8} + o(1))D(x^*)$ which concludes our proof. □

For the x_i chosen in the above proof, we point out that the median opinion is initially represented by a strict majority (over half the participants). In other words, (locally-consistent) pairwise interactions are actually *destroying* the remarkable consensus that existed initially.

4 Finding the median quickly through triads

In this section, we show a stark contrast to the impossibility results of the previous section when using small group interactions of size three. First, for opinions in R^d under the l_1 norm or in any median graph, we show that a simple policy for choosing triadic interactions (uniform random sampling), and a simple locally-consistent opinion update or bargaining function (the generalized median of S), is able to find a $(1+O(\sqrt{\frac{\ln n}{n}}))$-approximation of the generalized median in an average of $O(\log^2 n)$ interactions per participant. For the strategic bargaining setting, we remind the reader that small groups must choose a participant to support. The results, therefore, assume that the generalized median of S is represented by a

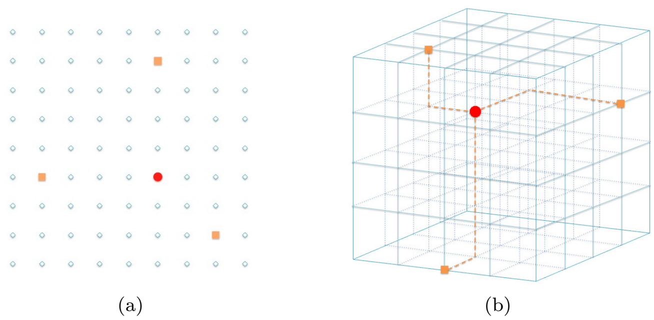

Figure 2: If the three orange squares are sampled, then the red circle is their generalized median.

participant, which is equivalent to assuming that the participant opinions themselves form a median graph (as opposed to simply belonging to one).

Second, we show that these results are obtainable under strategic equilibria. Specifically, we design a mechanism which has a Nash equilibrium over the entire extensive form game and where each triad chooses to support the generalized median of the small group. When participants are myopic in that they treat each small group as a separate game, then we show another mechanism which achieves this under a strong Nash equilibrium and show that any strong Nash equilibrium must satisfy this property.

The strength of the approximation is illustrated in the following example, in which the participants are uniformly distributed in a three-dimensional grid. In this case, the small group interaction is able to select the exact global generalized median with high probability.

Example 3 (A $n \times n \times n$ grid)**.** *Consider n^3 participants which make up a $n \times n \times n$ grid, so that each point (i, j, k) for $1 \leq i, j, k \leq n$ has one participant with that opinion. The global generalized median of all participants is the point $(l+1, l+1, l+1)$, where we assume for simplicity that n is odd and that $n = 2l + 1$. Suppose that groups of three are uniformly selected at random to interact, and that they update their opinions according to $f(x, S) =$ generalized-median(S). Then every participant will converge to the exact global generalized median with probability at least $1 - e^{-n}$.*

4.1 Why would triads choose their generalized median?

Before stating our results, we briefly discuss the question of why one might expect individuals to update their opinion to, or choose to support, the generalized median of the triad.

Theorem 2. *For every x, y, z which come from either R^d under the l_1 norm or from any median graph, let m denote the generalized median of these three points. Then m is also the Condorcet winner for the preference profiles of x, y, and z. That is, for any other point p in the space, m is closer to at least two of x, y, z as compared to p.*

Proof. See Appendix B for the proof. □

Since the generalized median of three participants is a Condorcet winner, it is the clear theoretical choice for bargaining situations. Moreover, as mentioned in the related work, experiments have shown that majority rule dynamics do tend to converge to the Condorcet winner when it exists[10]. In the opinion formation setting, updating to the generalized median can be interpreted as a participant who is swayed by any majority opinions.

4.2 Convergence and approximation when triads return their local generalized median

Definition 5 (Triadic process). *Define a triadic process on points* $x_1, x_2, \ldots, x_n \in \mathcal{X}$ *as follows. Let* $S_1, S_2, \ldots$ *denote a sequence of (random) multi-sets where each* S_i *is independent and consists of three points in* $\{1, 2, \ldots, n\}$ *drawn uniformly at random with replacement. Let* $x_i^0 = x_i$*, and*

$$x_i^{t+1} = \begin{cases} f(x_i^t, S_{t+1}) & \text{if } i \in S_{t+1} \\ x_i^t & \text{otherwise} \end{cases},$$

where $f(x, S) = \text{generalized-median}(S)$*.*

Theorem 3. *Consider the triadic process described in Definition 5, where* $\mathcal{X}$ *is a median graph with* m *nodes. For* $t = O(cn \log^2 n)$*, and with probability at least* $1 - \frac{m \log m}{n^c}$*,* x_i^t *will have converged to a single point* w *which satisfies*

$$D(x^*) \leq D(w) \leq \left(1 + O\left(\sqrt{\frac{c \log n}{n}}\right)\right) D(x^*)$$

where x^* *is the median of the* x_i*. When* $\mathcal{X}$ *is* R^d *with the* l_1 *norm, then the same results hold with probability at least* $1 - \frac{d}{n^c}$*.*

The proof of this theorem essentially uses a property that can be derived for median graphs (Lemma 1) to reduce the triadic process (Definition 5) to at most $m \log m$ urns, for which we can derive convergence properties using the theorems of [15] and probabilistic recurrences (Lemma 2). These lemmas are then used along with more properties of median graphs to prove our final result. We sketch the proof of Lemma 2 here. Detailed proofs of Lemma 1, Theorem 3 and Lemma 2 can be found in the appendix.

Lemma 1. *Consider any edge* $e = (u, v)$ *of a given median graph. Let* U *denote the set of nodes closer to* u*, and* V *denote the set of nodes closer to* v*. Then for any three nodes* x, y, z*, their generalized median* m *must belong to the set (*U *or* V*) which contains at least two of* x, y, z*.*

Lemma 2. *Consider the triadic process described in Definition 5, where* $\mathcal{X}$ *is a median graph with* m *nodes. Consider any edge* $e = (u, v)$ *of the median graph and let* U *and* V *denote the set of nodes closer to* u *and* v *respectively. Let* T *denote the number of triads selected until all balls left either belong to* U *or* V*, and* w *denote the opinion that all participants eventually converge to. Then,*

$$\mathbf{Pr}[w \in U] = \left(\frac{1}{2}\right)^{n-1} \sum_{j=1}^{|U|} \binom{n-1}{j-1} \tag{4.1}$$

$$\mathbb{E}[T] \leq n \ln n + O(n) \tag{4.2}$$

$$\mathbf{Pr}[T > cn \ln^2 n] < O(n^{-c}) \tag{4.3}$$

Proof. Equations (4.1) and (4.2) are derived in [15]. We prove (4.3) using probabilistic recurrences in the appendix. The essential observation needed (a generalization of the proof in [15]) is that this process can be represented as a random walk on the integers 0, 1, ..., n. Let X_t denote the number of balls in U after t samples. By Lemma 1, whenever two or more of the sampled balls belong to U, the generalized median will belong to U. Similarly,

ALGORITHM 1: Triadic Consensus

Input: k tokens per participant

Output: A winning participant

while *the tokens belong to more than one participant* **do**

  Sample tokens x, y, z uniformly at random with replacement;

  if *two or more of* x, y, z *have the same label* **then**

    $w =$ the majority participant;

  else

    $w =$ TriadicMechanism(x, y, z);

  Give all the sampled tokens to the winning participant w;

return the participant who has all the tokens;

whenever two or more of the sampled balls belong to V, the generalized median will belong to V. This means that,

$$\mathbf{Pr}[X_{t+1} = X_t + \Delta] = \begin{cases} \frac{3X_t^2(n-X_t)}{n^3} & \text{if } \Delta = 1 \\ \frac{3X_t(n-X_t)^2}{n^3} & \text{if } \Delta = -1 \\ \frac{X_t^3+(n-X_t)^3}{n^3} & \text{if } \Delta = 0 \end{cases}$$

□

4.3 Implementation under strategic equilibria

In our analysis of strategic bargaining thus far, we abstracted the bargaining dynamics of the triad into a function $g(S)$ and showed that choosing $g(S) =$ generalized-median(S) results in quick convergence to a tight approximation of the global generalized median. In this section, we analyze this process from a game-theoretic perspective and consider how one can design a mechanism which induces participants to support the generalized median under a strategic equilibrium. For concreteness, the bargaining process is codified in Algorithm 1, with the exception of a TriadicMechanism, which will be designed in the following sections.

4.3.1 Myopic participants

We will first consider participants who are strategic, but short-sighted in that they treat each triadic round as a separate game in which they try to make the supported participant as close to themselves as possible (as opposed to making the overall winner close to themselves).

Consider using Majority Rule as the TriadicMechanism (Algorithm 2). In Majority Rule, an arbitrary participant is chosen by the mechanism to be the initial winner. Participants of the small group then take turns suggesting to either replace the current winner with an alternate participant or to conclude the process. Each time a participant is suggested, a vote is then taken between the suggested participant and the current winner. If at least two of the three small group participants prefer the suggested participant, then it becomes the current winner. Similarly, if a suggestion is made to conclude the process, a vote is taken for or against that suggestion. If two of the three participants agree, then the process concludes and the current winner is returned as the supported participant. If the number of rounds taken exceeds some given parameter T (think of this as large), the process is halted, and the current winner is returned as the supported participant.

Formally, each triadic mechanism is modeled as an independent extensive form game in which each node of the gametree corresponds to either a participant's turn to make a proposal or to vote on another participant's proposal (see Appendix C.1).

ALGORITHM 2: Majority Rule mechanism

Input: Participants x, y, z from a set of n participants
Output: A participant to support
Initialize w to z and t to 0;
while $t < T$ **do**
    `/* Take turns proposing starting from` x`, then` y`, then` z`, etc...  */`
    The participant whose turn it is chooses w' from any of the $n-1$ participants not equal to w, or suggests to end ($\emptyset$);
    if $w' \neq \emptyset$ *and at least two people vote for* w' *over* w **then**
        $w = w'$;
    else if $w' = \emptyset$ *and at least two people vote to end the process* **then**
        return w;
    $t = t + 1$;
return w;

Theorem 4. *Consider using Majority Rule (Alg. 2) as the mechanism which decides which participant a small group supports. Suppose that the n participants form a median graph, and that they are myopic. In other words, participants consider how to maximize their utility for the participant supported by the current small group. Then the following strategy finds the generalized median and is resistant to deviations from coalitions of any size, i.e. a strong Nash equilibrium:*

1. *If it is your turn, propose the median if it is not the current winner, and propose to end the process otherwise.*
2. *If you need to vote between proposals, vote for the median if it is one of the proposals, and vote truthfully otherwise.*
3. *If you need to vote on ending the process, vote to end if the current winner is the median, and vote to continue otherwise.*

Moreover, all strong Nash equilibria must produce the median as the output.

Proof. The proof rests on the fact that the generalized median of any three participants is also a Condorcet winner (Theorem 2) and the fact that "majority rules". The detailed proof can be found in Appendix C.1. □

In practical crowdsourcing applications, participants may not be aware of all the possible participants from the beginning, implying that the model should be generalized further to consider the set of proposals available to each participant over time. In any modification for which the participants are able to identify the median before time T (that is, the median becomes available as an action before time T), a slight modification of the above strategy is also able to find the median under a strong Nash equilibrium.

4.3.2 Participants with complete knowledge

When participants have complete knowledge, the strategic possibilities become complex. We consider a variant of the Majority Rule, which we call the Chooser-Proposer mechanism (Algorithm 3). In this mechanism, two participants are arbitrarily assigned to be "proposers" and the remaining participant is assigned to be the "chooser". At the start of the small group, one of the proposers is chosen to be the current winner, and the other proposer starts by suggesting an alternate proposal. If the chooser accepts this proposal, then it replaces

ALGORITHM 3: 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 z;
while true do
    /* y and z take turns proposing starting from y                      */
    The participant whose turn it is chooses w' from any proposal not equal to w, or passes (∅);
    if w' ≠ ∅ and the chooser accepts w' then
        w = w';
    else
        return w;
```

the past proposal as the current winner, 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. We define the entire bargaining process formally in the Appendix as an infinite-horizon extensive form game.

To prove the Nash equilibrium result, we start by showing that we can reduce this problem to a simpler one. Our method is to show that, if some strategy t defined for a single small group has certain properties (defined below), then the extensive form strategy consisting of playing t in *every* small group is a Nash equilibrium (proof in Appendix C).

Lemma 3. *Consider a TriadicMechanism and a strategy t for the extensive form game. Suppose that the (local) generalized median is produced in each small group when all players follow t. Now, consider the supported participant m' produced in some given small group if a single agent x deviates from t. If the generalized median of the small group always lies on some shortest path from x to m', no matter what deviation x makes, then the strategy t is a Nash equilibrium.*

Proof. The proof involves coupling two urns together: one in which every agent follows the strategy t, and the other in which some agent deviates. We then show that the outcome produced in the non-deviating urn is at least as good as the outcome produced in the deviating urn for every coupled history, which gives us our desired result. □

The above Lemma essentially says that if we can find a mechanism and corresponding strategy that 1) produces the generalized median in each small group, and 2) ensures that any deviation by one player would result in a "strictly further" outcome, then this strategy is a Nash equilibrium for the overall urn process.

Theorem 5. *Consider using the Chooser-Proposer mechanism (Alg. 3) as the TriadicMechanism in Triadic Consensus (Alg. 1). Suppose that the n participants form a median graph. Then the following strategy finds the generalized median in each round and is also a Nash equilibrium for the overall urn process:*

1. *If you are the chooser, accept if you prefer the suggested alternative, and reject otherwise.*
2. *If you are the proposer, suggest the proposal closest to yourself that is between the current winner and the chooser. If this is identical to the current winner, then pass.*

Proof. Our proof involves showing (Lemmas 8, 9) that Algorithm 3 satisfies the properties specified in Lemma 3, which gives us our result. Details are in Appendix C. □

5 Conclusion and future directions

We introduced the framework of small group dynamics and applied it to the problem of finding the generalized median of a set of points. In doing so, we found a stark contrast between the possibility of obtaining our goal when using pairwise or triadic interactions. It would be interesting to perform experiments to verify whether such a stark difference really occurs, and whether triadic interactions can be helpful for aggregating opinions.

Many exciting directions also remain with regards to the use of the small group framework for the design of algorithms for complex crowdsourcing tasks.

A Median Graphs

Define a set S to be convex if for any two points x, y in S, every shortest path between x and y also lies in S. Define the convex hull of a set S to be the smallest convex set that contains S. The following properties on median graphs are used in various proofs in this paper.

Lemma 4. *Any convex set S in a median graph is gated. That is, for any x in the graph, there exists some gate $g(x) \in S$ such that for any node $y \in S$, some shortest path from x to y goes through $g(x)$.*

Lemma 5. *The interval I_{xy} between any two nodes x and y in a median graph is convex.*

Lemma 6. *For any edge $e = (u, v)$, define the win sets $W_{uv} = \{w \in V \mid d(w, u) < d(w, v)\}$ and $W_{vu} = \{w \in V \mid d(w, v) < d(w, u)\}$ to be the set of nodes that are closer to u or v respectively. Then,*

1. *W_{uv} and W_{vu} are convex sets that partition the nodes.*
2. *For any two unique nodes, there exists at least one edge that partitions them.*
3. *For a median graph with n nodes, there are at most n edges that uniquely partition the nodes.*
4. *Let E' denote a maximal set of edges that uniquely partition the graph. Then the distance between two nodes x and y is equal to the number of edges in E' that separate x and y.*
5. *$x \in I_{yz} \iff x \in W_{uv}$ if $y, z \in W_{uv}$ and $x \in W_{vu}$ if $y, z \in W_{vu}$.*

The above properties follow from a stronger property stating that the win sets of E' provide an isometric embedding of the median graph into a hypercube of dimension at most n. Specifically, every edge $e = (u, v)$ corresponds with one dimension of the hypercube. If a node falls in the partition W_{uv}, then label it with 0 in that dimension. Otherwise, if it falls in the partition W_{vu}, then label it with 1 in that dimension. It will often be more convenient to use this representation. We will use the following notation in proofs:

Definition 6. *Given a set of participants $p_1, p_2, \ldots, p_n$,*

1. *Let p_{ij} denote the j-th bit of the embedding of p_i into the hypercube, i.e. the label corresponding to the partition that p_i belongs to for the j-th edge.*
2. *Let $N(b, j)$ denote the number of participants whose j-th bit is equal to the bit b, i.e. for partitions induced by the win sets of edge j, $N(b, j)$ is the number of participants belonging to the partition with label b.*

B Approximation and convergence

(Proof of Lemma 1). Without loss of generality, suppose that $x, y \in U$. By known properties of median graphs (Lemma 4), U is a *gated set.* That is, for any point p and gated set U, there exists a gate $g(p) \in U$ such that for every $q \in U$, there exists a shortest path from p to q that goes through $g(p)$. Let $g(m)$ denote the gate of m into U. Then,

$$\begin{aligned} & d(m,x) + d(m,y) + d(m,z) \\ & \quad = [d(m,g(m)) + d(g(m),x)] + [d(m,g(m)) + d(g(m),y)] + d(m,z) \\ & \quad \geq d(m,g(m)) + d(g(m),x) + d(g(m),y) + d(g(m),z) \end{aligned}$$

where the inequality comes from the triangle inequality. If $m \in V$, then $d(m, g(m)) > 0$ since $g(m) \in U$. But then the sum of distances from $g(m)$ to x, y, z would be less than that of the generalized median, which is impossible by definition. Therefore, m belongs to U. $\square$

(Proof of Theorem 2). Recall that by the definition of median graphs, for every x, y, z, there is a unique point that lies on a shortest path between x and y, x and z, and y and z. Call this point m. We will show that m is the only node satisfying all the remaining properties of Theorem 2, which shows their equivalence.

We first show that m is the unique generalized median of x, y, and z. Consider any other node m'. Since m is the unique node lying on the shortest paths of each pair, there must be one pair (x and y WLOG) such that m' does not lie on a shortest path between them. Then,

$$\begin{aligned} 2D(m') &= 2(d(m',x) + d(m',y) + d(m',z)) \\ &= (d(x,m') + d(m',y)) + (d(x,m') + d(m',z)) + (d(y,m') + d(m',z)) \\ &> (d(x,m) + d(m,y)) + (d(x,m) + d(m,z)) + (d(y,m) + d(m,z)) \\ &= 2D(m) \end{aligned}$$

where the inequality comes from the fact that m lies on a shortest path between each pair. Therefore, m is the unique generalized median.

Now we show that m is the unique point closest to x among the points lying between y and z. By Lemma 5, the set of points lying on a shortest path between y and z is convex. By Lemma 4, this means that it is a gated set, and the point closest to x must be the (unique) gate of x into I_{yz}. But by the definition of the gate, it must lie on a shortest path between x and y, and x and z. Also, since it belongs to I_{yz}, it must also lie on a shortest path between y and z. By the definition of median graphs, m is the only node for which this can be true.

Finally, we show that m is the Condorcet winner. Consider any other node m'. As argued before, there must be one pair (x and y without loss of generality) such that m' does not lie on a shortest path between them. Therefore, $d(x, m) + d(m, y) < d(x, m') + d(m', y)$ and at least one of x and y will prefer m over m' by the pigeonhole principle. Let this node be x without loss of generality. If neither of y or z is closer to m' than m, then m clearly beats m' in a pairwise election. Now suppose one of y or z is closer to m' than m (y without loss of generality). Since m lies on a shortest path between y and z, we know that $d(y, m) + d(m, z) \leq d(y, m') + d(m', z)$. But then if y is closer to m' than m, it must follow that z is closer to m than m' and m still beats m' in a pairwise election. Therefore m beats every other node m' in a pairwise election, which means he is the Condorcet winner. $\square$

(Proof of Theorem 3). We will first prove that all opinions will have converged in $O(n \log^2 n)$ triads. For any given edge, Lemma 2, Eqn. 4.3 tells us that $O(cn \log^2 n)$ triads is sufficient for ensuring that all opinions are closer to one side of the edge with probability at least

$1 - n^{-c}$. Since median graphs have at most $m \log m$ unique partitions (Lemma 6, Eqn. 3), a union bound tells us that $O(cn \log^2 n)$ triads is sufficient for ensuring that, for every edge-based partition, all opinions will belong to the same partition with probability at least $1 - \frac{m \log m}{n^c}$. Since any two unique nodes must have one partition separating them (Lemma 6, Eqn. 2), this means that all opinions are on the same node, and the process has converged. For points in R^d with the l_1 norm, there are at most nd unique partitions, which gives us the analogous result for this case.

The lower bound of the approximation result, $D(w^*) \leq D(w)$, follows directly from the definition of the generalized median. We now prove the upper bound. We will use notation from Section A on median graphs. Suppose that for an edge $e = (u, v)$ there are less than $\frac{n}{2} - \sqrt{cn \ln n}$ nodes in one of its partitions (W_{uv} without loss of generality). Then by Lemma 2, we know that

$$\begin{aligned}
\mathbf{Pr}[w \in W_{uv}] &\leq \left(\frac{1}{2}\right)^{n-1} \sum_{j=1}^{\frac{n}{2} - \sqrt{cn \ln n}} \binom{n-1}{j-1} \\
&\leq \exp\left(-\frac{1}{2} \cdot \frac{n}{2} \cdot \left(\frac{\sqrt{cn \ln n}}{n/2}\right)^2\right) \qquad \text{(B.1)} \\
&= \frac{1}{n^c}
\end{aligned}$$

where (B.1) is found by interpreting the expression in terms of coin flips and applying Chernoff's bound. Then by applying the union bound to each edge (of which there are at most $m \log m$ by Lemma 6(3)), we know that every win set w belongs to will have at least $\frac{n}{2} - \sqrt{cn \ln n}$ members in it with probability at least $1 - \dfrac{m \log m}{n^c}$. In the notation corresponding to the hypercube embedding of the median graph (Definition 6),

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

All we need to show is that this condition implies that w is an approximate generalized median. Let d denote the number of dimensions in the hypercube embedding, i.e. the number of edges that uniquely partition the graph. From Lemma 6(4), we know that the distance $d(x, y)$ from x to y can be written as $d(x, y) = \sum_{i=1}^{d} \mathbb{1}_{x_i \neq y_i}$. This means that,

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

Since w satisfies $N(w_j, j) > \frac{n}{2} - \sqrt{cn \ln n}$, then it must be true that for any bit b,

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

and

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

and we are done. The proof for points in R^d with the l_1 norm is similar. □

(Proof of Lemma 2). Equations (4.1) and (4.2) are derived in [15]. (4.3) follows from (4.2) through the use of common techniques in probabilistic recurrences (see, for example, [4]). Let $T_n(i)$ denote the time it takes for all balls to belong to either U or V given i initial balls in U ($X_0 = i$). Then $T_n(i)$ satisfies the probabilistic recurrence

$$
T_n(i) = p_i T_n(i+1) + q_i T_n(i-1) + r_i T_n(i) + 1
$$

where p_i, q_i, r_i are the probabilities of having one more ball in U, one less ball in U, or no change. From Equation 4.2, we have that $\mathbb{E}[T_n(i)] \leq T^*$ for all i, and for $T^* = n\ln n + O(n)$.

By Markov's inequality, $\mathbf{Pr}[T_n(i) \geq \alpha T^*] \leq \frac{1}{\alpha}$, for any $\alpha > 0$. By conditional probability, $\mathbf{Pr}[T_n(i) \geq 2\alpha T^*] = \mathbf{Pr}[T_n(i) \geq 2\alpha T^* \mid T_n(i) \geq \alpha T^*]\mathbf{Pr}[T_n(i) \geq \alpha T^*]$. But note that $\mathbf{Pr}[T_n(i) \geq 2\alpha T^* \mid T_n(i) \geq \alpha T_n^*] \leq \frac{1}{\alpha}$ since the remaining random walk conditioned at the point when $T_n(i) = \alpha T^*$ is just distributed as $T_n(i')$ for some i'. Repeating this logic gives $\mathbf{Pr}[T_n(i) \geq k\alpha T^*] \leq \alpha^k$. Choosing $\alpha = e$ and $k = c\ln n$ gives us our desired result. Note that the point at which $T_n(i) = \alpha T^*$ exists for any T^* chosen such that αT^* is an integer since $T_n(i)$ increases in integer increments. □

C Strategic behavior and Nash equilibria

C.1 Myopic participants

The majority rule extensive form game (Algorithm 2) can be defined formally as follows:

Let y_0, y_1, y_2 denote the three participants in the small group. Then,

1. The set of players in the game is $N = \{y_0, y_1, y_2\}$.
2. The entire set of actions is $A = \mathcal{X} \cup \{\emptyset, Y, N\}$. Actions $a \in \mathcal{X}$ and $\emptyset$ are played at proposing nodes and represent a suggestion for a as an alternate participant to support or suggestion to conclude the process respectively. Actions Y and N are played at voting nodes and represent yes or no votes respectively.
3. Before defining the game tree, we need to define a voting tree $V(G_1, G_2, \ldots, G_8)$ which will be used as a subtree in our definition. This voting tree takes eight gametrees $G_1, \ldots, G_8$ as inputs which represent the remainder of the game after the conclusion of the eight possible voting outcomes.
   (a) The root node is a nonterminal choice node, played by y_0, and with actions $\{Y, N\}$.
   (b) The two nodes at height one are nonterminal choice nodes, played by y_1, and with actions $\{Y, N\}$. Both of these nodes are a single information set.

(c) The four nodes at height two are nonterminal choice nodes, played by y_2 and with actions $\{Y, N\}$. All four of these nodes are a single information set.

(d) The eight nodes at height three are the root nodes for the gametrees $G_1, G_2, \ldots, G_8$. Specifically, $G_1, G_2, \ldots, G_8$ are located at the nodes resulting from actions NNN, NNY, NYN, NYY, YNN, YNY, YYN, and YYY respectively. Note that the G_i can potentially include terminal nodes.

4. We can now define the game tree $G(T, y_r)$ as a function of the parameter T (number of rounds) and the root node player y_r (the starting participant). Define $r' = (r + 1) \bmod 3$ (think of this as the next proposer) and $G' = G(T-1, y_{r'})$ (think of this as the remainder of the game). Let $-$ denote a gametree which has only a single terminal node as its root.

   (a) The root node is a nonterminal choice node, played by y_r, and with actions $\mathcal{X} \cup \{\emptyset\}$. This represents a decision to propose an alternate candidate or to propose to end the process.

   (b) Each node at height one resulting from an action $a \in \mathcal{X}$ is the root node for the voting tree $V(G', G', \ldots, G')$.

   (c) Each node at height one resulting from the action $a = \emptyset$ is the root node for the voting tree $V(G', G', G', -, G', -, -, -)$.

   (d) For any node at height $h \equiv 0 \pmod 4$, the winner w is found by calculating the votes leading up to the node. If this node is a terminal node, then assign utility to player p equal to any $u_p(w)$ which decreases as the distance $d(p, w)$ increases.

(Proof of Theorem 4). If all participants follow this strategy, it is easy to see that the winner will be the median. If only one person deviates, it is also easy to see that he cannot change the outcome (since majority rules so that once it is another participant's turn, the median will be proposed and voted for). If two or more people deviate, it is possible to get a winner that isn't the median; however, since the median is a Condorcet winner, at least one of the deviating participants must prefer the median to the resulting winner, which means that the deviation does not benefit every member of the coalition. Therefore, the strategy is a strong Nash equilibrium.

Similarly, for any set of strategies that doesn't produce the median, there must be two participants who prefer the median over the resulting winner since the median is a Condorcet winner. Then these two participants could deviate to the above strategy to achieve the median, which means that the original set of strategies cannot be a strong Nash equilibrium. □

C.2 Participants with complete knowledge

The strategic bargaining game (Algorithms 1 and 3) for participants with complete knowledge can be defined formally as follows:

1. The set of players in the game is $N = \{1, 2, \ldots, n, R\}$, where R represents the random choice of a small group played by nature (in our setting, the moves by nature are known by each player, i.e. they are separate information sets).

2. The actions available are $\mathcal{X}^3 \cup \mathcal{X} \cup \{\emptyset, Y, N\}$. Actions $S \in \mathcal{X}^3$ are played at nature nodes and represent the triad that is selected to participant in the small group. Actions $a \in \mathcal{X}$ and $\emptyset$ are played by the proposers and represent a suggestion for a as an alternate participant to support or a decision to pass respectively. Actions Y and N are played by the chooser and represent accepting or rejecting the proposal.

3. The game G is defined recursively as follows.
   (a) The root node is a nonterminal choice node, played by nature (R), and with actions $S \in \mathcal{X}^3$. This represents the next small group that is chosen to interact. Note that the fact that nature chooses this action uniformly at random is known.
   (b) Each node at height one resulting from an action $S = (x, y, z) \in \mathcal{X}^3$ is the root node for the chooser-proposer tree $CP(x, y, z, G)$, which is a subtree representing the small group interaction.
4. The chooser-proposer tree $CP(c, p_0, p_1, G)$ is an infinite horizon game tree which takes as input the chooser c, proposers p_0, p_1, and a game tree G which represents the remainder of the game after the end of this small group.
   (a) The root node is a nonterminal choice node, played by p_0, and with actions $\mathcal{X} \cup \{\emptyset\}$. This represents a decision to propose an alternate candidate or to pass.
   (b) Each node at height one resulting from an action $a \in \mathcal{X}$ is played by c, and with actions $\{Y, N\}$. This represents a decision to accept or reject the proposal.
   (c) The node at height one resulting from the action $\emptyset$ represents a conclusion of this small group from a pass action. The token distribution at this node is calculated from the actions since the start of the game. If it results in all tokens given to one player, then this node is a terminal node. Otherwise, this node is the root node of the subtree G.
   (d) Each node at height two resulting from an action Y contains the subtree $CP(c, p_1, p_0, G)$. In other words, the same process repeats with p_1 as the first proposer.
   (e) Each node at height two resulting from an action N represents a conclusion of this small group from a reject action. The token distribution at this node is calculated from the actions since the start of the game. If it results in all tokens given to one player, then this node is a terminal node. Otherwise, this node is the root node of the subtree G.
5. For any terminal node, assign utility to player p equal to any $u_p(w)$ which decreases as the distance $d(p, w)$ increases. Assign utility $-\infty$ to the case when the game does not end (Note: any utility can be assigned here without changing the results).

To prove Lemma 3, we first introduce a definition and an intermediate lemma.

Definition 7. *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 7. *Suppose that R x-dominates S. Then the generalized median of r_i, r_j, r_k must be on some shortest path from x to the generalized median of s_i, s_j, s_k.*

Proof. Let m_s denote the generalized median of s_i, s_j, s_k. Let m_r denote the generalized median of r_i, r_j, r_k. Now fix any edge e in the median graph and consider the partitions induced by its win sets. By Lemma 1, we know that m_s is in the same win set as x if and only if at least two of s_i, s_j, s_k are in the same win set as x. By Lemma 6(5), we know that this must imply that at least two of r_i, r_j, r_k are in the same win set as x, which means that m_r must also be in the same win set as x (again, by Lemma 1).

Therefore, for every edge, whenever m_s is in the same win set as x, m_r is also in the same win set. By Lemma 6(5) this implies that m_r lies on a shortest path from x to m_s, which concludes our proof. □

(Proof of Lemma 3). 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 generalized 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 generalized median of r_i, r_j, r_k. Then by Lemma 7, 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 7, 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 initially, 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. □ □

Lemma 8. *Suppose that three participants in a round of the Chooser-Proposer TriadicMechanism (Algorithm 3) all follow the strategy described in Theorem 5 as their strategy. Then the winner returned will be the generalized 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. If p_2 is the generalized median of c, p_1, p_2, then the closest point to p_1 which lies between p_2 and c will just be p_2, and he will pass, which results in p_2 being returned. Otherwise, p_1 will propose the closest point to him that lies in the interval from c to p_2, which is simply the generalized median of p_1, c, and p_2 by Theorem 2.

Now, in p_2's turn, there cannot be any other alternatives that are closer to both c and p_2 since the generalized 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 9. *Suppose that three participants in a round of the Chooser-Proposer TriadicMechanism (Algorithm 3) all follow the strategy described in Theorem 5 as their strategy, except for one participant who deviates. Then the generalized 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.

Case One: The chooser (c) deviates. If c deviates in the first round when p_1 proposes the generalized median, then the resulting winner will be p_2, which has a shortest path containing the generalized median by Theorem 2. 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. Let p_s denote the deviator, and p_t denote the other proposer. We will show that, regardless of what p_s does, the process will converge and the winner will be some proposal that is in the interval of p_t and c. Since the generalized median is the gate of p_s into the interval of p_t and c, it follows that the generalized median lies on a shortest path between p_s and the winner.

Consider a proposal by p_s which is accepted by c and not in the interval of p_t and c. Then p_t will propose the closest point to himself that is also between the current winner and c. This point is just the gate of current winner to the interval p_t and c, which must be unique from the current winner. This means that c will also accept the new proposal. Now, consider a proposal by p_s which is accepted by c and is in the interval of p_t and c. Then p_t will simply pass and that proposal will be returned as the winner. Finally, note that whenever p_s starts his turn, the current winning proposal will always be in the interval of p_t and c. Therefore, if the process ever converges, the winner must be from the interval of p_t and c. Since each accepted proposal must be higher on c's preference list, the process must converge, which finishes our proof. □

References

[1] Daron Acemoglu and Asuman Ozdaglar. Opinion dynamics and learning in social networks. *Dynamic Games and Applications*, 1:3–49, 2011.

[2] Tanja Aitamurto and Helene Landemore. Democratic participation and deliberation in crowdsourced legislative processes: the case of the law on off-road traffic in finland. In *The 6th Conference on Communities and Technologies (C&T). Workshop: Large-Scale Idea Management and Deliberation Systems*, 2013.

[3] Kenneth Arrow. A difficulty in the concept of social welfare. *Journal of Political Economy*, 58(4):328–346, 1950.

[4] James Aspnes. Notes on randomized algorithms. http://www.cs.yale.edu/homes/aspnes/classes/469/notes.pdf, 2013.

[5] Hans-Jürgen Bandelt and Victor Chepoi. Metric graph theory and geometry: a survey. *Contemporary Mathematics*, 453:49–86, 2008.

[6] Berno Buechel. Condorcet winners on median spaces. *Social Choice and Welfare*, 2013.

[7] Vincent Conitzer and Tuomas Sandholm. Communication complexity of common voting rules. In *Proceedings of the 6th ACM Conference on Electronic Commerce (ACM-EC)*, pages 78–87, 2005.

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

[9] Jon Elster. *Deliberative Democracy.* Cambridge University Press, 1998.

[10] Morris P. Fiorina and Charles R. Plott. Committee decisions under majority rule: An experimental study. *The American Political Science Review*, 72(2):575–598, 1978.

[11] James S. Fishkin. *When the People Speak: Deliberative Democracy and Public Consultation*. Oxford University Press, 2011.

[12] Francis Galton. Letters to the editor: The ballot-box. *Nature*, 75(1952):509–510, 1907.

[13] Francis Galton. Vox populi. *Nature*, 75:450–451, 1907.

[14] Allan Gibbard. Manipulation of voting schemes: A general result. *Econometrica*, 41(4):587–601, 1973.

[15] Ashish Goel and David T. Lee. Triadic consensus: A randomized algorithm for voting in a crowd. In *Proceedings of the 8th Workshop on Internet and Network Economics (WINE)*, 2012.

[16] Benjamin Golub and Matthew O. Jackson. Naïve learning in social networks and the wisdom of crowds. *American Economic Journal: Microeconomics*, 2(1):112–149, 2010.

[17] Wilfried Imrich and Sandi Klavžar. *Product graphs, structure and recognition*. Wiley, 2000.

[18] Sheena Iyengar. *The Art of Choosing*. Twelve, 2011.

[19] Shaili Jain and David C. Parkes. A game-theoretic analysis of the esp game. *ACM Transactions on Economics and Computation (TEAC)*, 1(3), 2013.

[20] David R Karger, Sewoong Oh, and Devavrat Shah. Budget-optimal task allocation for reliable crowdsourcing systems. *Operations Research*, 2013.

[21] Donald Ervin Knuth. *The Art of Computer Programming*, volume 4A: Combinatorial Algorithms, Part 1. Pearson Education, Inc., 2011.

[22] Helene Landemore. *Democratic Reason: Politics, Collective Intelligence, and the Rule of Many*. Princeton University Press, 2012.

[23] Hendrik P. Lopuhaä and Peter J. Rousseeuw. Breakdown points of affine equivariant estimators of multivariate location and covariance matrices. *The Annals of Statistics*, 19(1):229–248, 1991.

[24] Andrew Mao, Ariel D. Procaccia, and Yiling Chen. Better human computation through principled voting. In *Proceedings of the 27th Conference on Artificial Intelligence (AAAI)*, 2013.

[25] Richard D. McKelvey. Intransitivities in multidimensional voting models and some implications for agenda control. *Journal of Economic Theory*, 12(3):472–482, 1976.

[26] George A. Miller. The magical number seven, plus or minus two: some limits on our capacity for processing information. *Psychological review*, 63:81–97, 1956.

[27] Daniela Saban and Nicolas Stier-Moses. 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)*, 2012.

[28] Barry Schwartz. *The Paradox of Choice - Why More is Less*. Harper Perennial, 2004.

[29] Graham Smith. *Democratic Innovations: Designing Institutions for Citizen Participation*. Cambridge University Press, 2009.

[30] James Surowiecki. *The Wisdom of Crowds*. Anchor, 2005.

[31] John N. Tsitsiklis. *Problems in Decentralized Decision Making and Computation*. PhD thesis, MIT, 1984.

[32] Luis von Ahn and Laura Dabbish. Labeling images with a computer game. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI)*, pages 319–326, 2004.

[33] Justin Wolfers and Eric Zitzewitz. Prediction markets. *Journal of Economic Perspectives, American Economic Association*, 18(2):107–126, 2004.

[34] H. P. Young. Condorcet's theory of voting. *The American Political Science Review*, 82(4):1231–1244, 1988.

Ashish Goel
Management Science and Engineering and Computer Science
Stanford University
Stanford, USA
Email: `ashishg@stanford.edu`

David T. Lee
Electrical Engineering
Stanford University
Stanford, USA
Email: `davidtlee@stanford.edu`