

Towards Large-Scale Deliberative Decision-Making: Small Groups and the Importance of Triads

ASHISH GOEL and DAVID T. LEE, Stanford University

Though deliberation is a critical component of democratic decision-making, existing deliberative processes do not scale to large groups of people. Motivated by this, we propose a model in which large-scale decision-making takes place through a sequence of small group interactions. Our model considers a group of participants, each having an opinion which together form a graph. We show that for median graphs, a class of graphs including grids and trees, it is possible to use a small number of three-person interactions to tightly approximate the wisdom of the crowd, defined here to be the generalized median of participant opinions, even when agents are strategic. Interestingly, we also show that this sharply contrasts with small groups of size two, for which we prove an impossibility result. Specifically, we show that it is impossible to use sequences of two-person interactions satisfying natural axioms to find a tight approximation of the generalized median, even when agents are non-strategic. Our results demonstrate the potential of small group interactions for reaching global decision-making properties.

Additional Key Words and Phrases: participatory and deliberative democracy; decision-making at scale; wisdom of the crowd; triadic consensus; small groups

1. INTRODUCTION

Recent years have seen an increase in the number of democratic innovations [Smith 2009] aimed at increasing the participation of the public in policy-making. While the core argument of participatory democracy states that democracy is made more meaningful when citizens are directly involved [Fung and Wright 2003; Macpherson 1977; Pateman 1979], it has also been argued that participatory practices foster more engaged citizens [Fung and Wright 2003] and can better harness the wisdom of the crowds [Landemore and Elster 2012]. One class of these processes emphasizes deliberation as an essential component to democratic decision-making. Proponents argue that deliberation grants legitimacy to the democratic process [Gutmann and Thompson 2004], can produce outcomes superior to those resulting from mere voting [Fishkin 2011], and may be a means through which one can mediate conversations among polarized participants [Menkel-Meadow 2011]. The latter is particularly desirable in light of the increasing polarization that has been observed in politics and society at large [McCarty et al. 2006; Poole and Rosenthal 1984].

A major challenge for increasing participation is the problem of effectively scaling up to thousands or millions of people. Yet, while there have been many efforts at scaling up voting to large numbers of proposals [Conitzer and Sandholm 2005; Lee et al. 2014; Lu and Boutilier 2011; Service and Adams 2012], the problem of scaling up delibera-

This work was supported by the Army Research office Grant No. 116388, the Office of Naval Research Grant No. 11904718, and by the Stanford Cyber Initiative. David T. Lee was also supported by an NSF graduate research fellowship and a Brown Institute for Media Innovation Magic Grant. The full version (with appendix) can be found at: <https://arxiv.org/abs/1605.08143>

Author emails: {ashishg,davidtlee}@stanford.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

EC'16, July 24–28, 2016, Maastricht, The Netherlands. Copyright © 2016 ACM 978-1-4503-3936-0/16/07 ...\$15.00.

<http://dx.doi.org/10.1145/2940716.2940722>

tion has been largely untouched. One part of the problem is that humans face fundamental cognitive and psychological barriers that prevent effective communication in large groups. For instance, it is known that humans can only hold a small number of concepts in their working memory [Miller 1956], that there are fundamental limits to effective group sizes [Dunbar 1992], and that humans can become more disengaged when in a large group [Latané et al. 1979].

Existing initiatives such as The Deliberative Poll [Fishkin et al. 2000] and the 21st Century Town Hall Meeting [Lukensmeyer and Brigham 2002], while successful on many counts, have typically operated at the scale of hundreds, or at most, thousands of participants. Even for these events, a major component of these processes involve further breaking down the gathered participants into small group discussions.

Motivated by this observation, this paper considers whether large-scale decision-making can take place through a sequence of small group interactions. Essentially, we take the view that the deliberative decision-making process can be, at its core, the interaction of many separate but intertwined deliberations. The use of small groups circumvents communication barriers, changing the problem of scalability from that of increasing group size to that of decreasing the number of small group interactions necessary. Thus, our use of small groups for scaling up deliberation can be thought of as analogous to the use of pairwise comparisons for scaling up voting when the number of proposals is very large [Lee et al. 2014]. From a practical perspective, scaling deliberation via small groups would not have been feasible even 10 years ago. Today, however, advances in communication and Internet-based technologies such as video conferencing have made it trivial for arbitrary small groups of individuals to deliberate with one another, and moreover, to sequence these deliberations based on outcomes of prior small groups.

In our model, participants each have an opinion which together form a graph. Participants prefer opinions that are closer to their own. Our goal is to find the generalized median using a sequence of small group interactions, where the generalized median is defined as the point minimizing sum of distances to the initial opinions. Each small group is required to come to a small group decision, which is modeled by us in two distinct ways. For our impossibility results, we assume that agents do not behave strategically, and that the small group decision is therefore a (possibly stochastic) function of the opinions of the small group members. We assume that this function satisfies natural axioms. For our positive results, we allow strategic agents. This means that the small group decision is the result of the negotiation strategies chosen by the individuals in that group. Under this framework, our main findings uncover a sharp dichotomy among small group processes:

- First, we obtain a surprising impossibility result for *dyadic decision-making* even for non-strategic agents. We show that under natural axioms, no sequence of two-person interactions can accurately estimate the generalized median. This holds even for the simple case when opinions lie on a line.
- In contrast, for *triadic decision-making*, we obtain a strong positive result even for strategic agents. We prove that when the small group decision is decided through majority rule dynamics, a small number of three-person interactions (an average of only $O(\log^2 n)$ triads per person) finds a tight approximation of the generalized median. This holds not only when agents have opinions from a line, but for *every* graph for which every three individuals in the graph have a majority rule equilibrium (we prove this is identical to a class of graphs known as median graphs). The mechanism we give is shown to achieve our results under a subgame perfect Nash equilibrium.

Figure 1 illustrates our mechanism for the simple case where the opinions of the voters lie on a line. Given a triad, our mechanism for this case reduces to each member



Fig. 1: For the special case of triads on a line, our mechanism can be reduced to having each individual vote between the other two participants. This results in the median participant receiving two votes.

voting among the other two. This results in the median participant winning two votes, and getting increased weight in later rounds. Details are described in later sections. To show our positive results for triadic decision-making, we first show that our mechanism induces a behavior in which every triad reports its true generalized median; hence our game-theoretic result is also valid if we postulate an small group decision dynamic where triads agree on their generalized median.

These findings, put together, illustrate the importance of triads for small group decision-making. The fact that only an extremely small number of triads is required is a positive indicator of the feasibility of achieving deliberative decision-making at scale. The fact that it is not achievable with dyadic interactions is surprisingly initially, but makes intuitive sense after the fact: when participating in a two-person interaction, it is impossible to protect against an extreme outlier. If two individuals are to be treated fairly, an arbitrarily unreasonable individual could greatly affect the decision made. On the other hand, three-person interactions can essentially ignore single outliers through simple processes such as majority rule.

Before moving on, we also want to put our results in the broader context of deliberative democracy. Our model is focused on large-scale decision-making through small group interactions, which we believe will play an important part of deliberative systems due to the benefits of back-and-forth dialogue. However, our results do *not* account for an information-based evolution of beliefs. Our main theorem (Theorem 6.3) essentially states that, if groups of three can find their generalized median, then we can find the global generalized median of *initial opinions* by composing a small number of three-person decisions. If beliefs change, then one would really want to find the generalized median of *new opinions* after sufficient information exchange has occurred. From a political theorist’s perspective, deliberative democracy must involve a reasoned exchange of information, and not just negotiation [Gutmann and Thompson 2004]. Therefore, a critical step for the future is to consider models of decision-making where information is taken into account. It is also important to note that, while deliberation was traditionally viewed as an exchange of arguments among individuals, theorists have recently expanded that definition to encompass richer dynamics. In this view, political debates, news media, or dialogue over a social network (that has the potential to influence decision-making outcomes) can also be viewed as a part of a deliberative system. We believe that studying models for decision-making that incorporate richer, deliberative dynamics is an important question for the future.

2. MODEL

2.1. Participant opinions and median graphs

In our model, n participants have opinions $V = \{x_1, x_2, \dots, x_n\}$ which are nodes of a graph $G = (V, E)$. This graph is a representation of the opinion space. Each participant prefers opinions that are closer to his own, where the distance $d(x, y)$ is defined as the shortest path distance between opinions x and y in the graph. The participants participate in an algorithmically defined sequence of small group interactions, each of which results in a *small group decision*. When a stopping criterion is reached, a function maps the sequence of small group decisions to a *final consensus decision*.




Fig. 2: The red circle is the generalized median of the three orange squares.

A common example of an opinion space is to consider opinions on a line representing the liberal-conservative axis. This line could, for example, lie in a high-dimensional space (e.g. specific budget or policy proposals), so long as participants preferences are determined by distances on the line (how conservative or liberal the proposals are).

For our positive results on scaling decision-making using a sequence of triads (Theorem 6.3), our results are only true for graphs where every three individuals in the graph have a majority rule equilibrium (also known as Condorcet winner), i.e. there exists an opinion that would receive two out of the three votes against any other opinion. This turns out to coincide to a class of graphs called median graphs (we prove this in Theorem 6.2). Median graphs include several common classes of graphs such as trees, grids, and squaregraphs. As a toy example, if one were deciding on ice-cream flavors, one might represent the opinions as a two-dimensional grid graph where the two axes correspond to discrete levels of sweetness and fattiness. If people’s preferences for ice-cream flavors are roughly determined independently by sweetness and fattiness (or any finite number of dimensions), then this could be represented by a median graph. However, there are many graphs which are not median graphs. For instance, any graph with a triangle cannot be a median graph.

2.2. The generalized median as the wisdom of the crowd

Our objective is to find the wisdom of the crowd, here defined to be the generalized median x^* , the opinion which minimizes the sum of distances to all of the participant opinions (see Figure 2). Formally, $x^* = \arg \min_{x \in V} D(x)$, where $D(x) = \sum_{i=1}^n d(x, x_i)$. The decision-making process achieves a $(1 + \epsilon)$ -approximation of the generalized median if the final consensus decision is a point \hat{x} for which $D(\hat{x}) \leq (1 + \epsilon)D(x^*)$, i.e. \hat{x} almost minimizes the sum of distances. We seek sequences of small group interactions for which $\epsilon \rightarrow 0$ as the number of participants increases ($n \rightarrow \infty$).

The concept of the wisdom of the crowd [Surowiecki 2005] originated in 1907 from Sir Francis Galton’s paper [Galton 1907b] describing a competition at a carnival 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 *Vox Populi* or voice of the people, was remarkably close to the correct answer. Based on this, he hypothesized that an appropriate aggregation of a crowd’s preferences can produce an extremely accurate estimate.

Much of the following theoretical work on the wisdom of the crowd has chosen to study the mean rather than the median [Golub and Jackson 2010]. We follow Galton, who stated that he preferred the median over the mean due to its robustness to outliers and the fact that the median is the Condorcet winner, the opinion which is preferred by at least half of the participants to any other opinion [Galton 1907a; Hooker 1907]. Since we consider opinions that lie in high-dimensional metric spaces, we take the wisdom of the crowd to be the *generalized* median, which is also robust to outliers [Lopuhaä and Rousseeuw 1991] and closely related to the Condorcet winner [Bandelt and Barthelemy 1984; Saban and Stier-Moses 2012]. We also note that

recent experimental studies also support the position that the median is more accurate than the mean [Lorenz et al. 2011].

2.3. Large-scale decision-making via small groups

We will use an *LDSG process* to refer to a large-scale decision-making process conducted through small group interactions. In our process, each participant starts with k tokens, which can be thought of as k votes. Unless otherwise specified, we will assume $k = 1$. A small group of participants is chosen by a (possibly stochastic) *small group selection function* f to deliberate with one another, with the requirement that the members must come to a small group decision on a single participant to be their representative. This participant, who does not need to be one of the group members, is then given one token from each of the group members. We use S_t to represent the t -th small group and g_t to represent the small group decision that they come to. Another small group is chosen (again, by the function f), which can be dependent on the past history of small group decisions. This process is repeated until one participant has all the tokens. This participant's opinion is returned as the final consensus decision \hat{x} . Formally,

Definition 2.1. An LDSG process L is a triple $L = (V, f, g)$, where V is the set of participants $V = \{x_1, x_2, \dots, x_n\}$, f is the *small group selection function* which maps the sequence of past small groups and small group decisions to the next small group, and $\{g_t, t \geq 1\}$ are the small group decisions chosen by the t -th small group. Let S_t denote the small group chosen to interact in the t -th round and let y_i^t denote the participant holding the i -th token after the t -th round. Let T be the round at which the process ends, and \hat{x} denote the final consensus decision made. Then,

$$\begin{aligned} y_i^0 &= x_i, \\ y_i^{t+1} &= \begin{cases} g_{t+1} & \text{if } y_i^t \in S_{t+1} \\ y_i^t & \text{otherwise} \end{cases}, \\ S_{t+1} &= f(\langle S_1, g_1, S_2, g_2, \dots, S_t, g_t \rangle), \\ T &= \min\{t \mid y_1^t = y_2^t = \dots = y_n^t\}, \\ \hat{x} &= y_1^T \end{aligned}$$

2.4. Small group decision-making

In our model's full generality, the small group decision g_t can be modeled in two distinct ways. For non-strategic agents, we represent it by a *small group decision function* $g : S \rightarrow \mathcal{P}_V$, which maps any small group $S \subseteq V$ to a probability distribution \mathcal{P}_V over the set of opinions V , i.e. $g_t = g(S_t)$. For strategic agents, the decision of a small group is taken to be the output of a mechanism that the small group participates in. This output is a function of the strategic decisions that the small group participants choose to make. Since participants can base their strategies on the entire history of the LDSG game up to that time, g_t is not necessarily a function solely of S_t .

For our dyadic impossibility results, we will consider non-strategic agents and show that no function $g(S_t)$ exists which simultaneously satisfies natural axioms and also results in tight approximation of the generalized median. In our result on triadic interactions, we will consider small groups which make their decision through a simple majority rule process. In the non-strategic case, $g(S_t)$ is assumed to be the majority rule equilibrium for that small group (described further in Section 6). In the strategic case, we will give a formal representation of the entire sequence of small group majority rule processes as an infinite-horizon extensive form game (Section 7).

3. OUTLINE

After describing further related work (Section 4), the remainder of this paper will be structured as follows. Section 5 will present our impossibility result on dyadic interactions. Section 6 contains our results on scaling deliberative decision-making with three-person *non-strategic* interactions. This section will formally define the majority rule process and equilibrium that motivates our model of the small group decision function $g(S_t)$. It will then use this to show that $O(n \log^2 n)$ triads suffice to find a tight approximation of the generalized median for a large class of graphs. Section 7 expands on the prior section by considering the case when agents are strategic and show that we can achieve the same results under a subgame perfect Nash equilibrium. Finally, we conclude in Section 8 with a brief discussion of future directions.

4. RELATED WORK

The problem of scaling up decision-making is a well-studied problem in the context of voting and preference elicitation. In this context, one is typically trying to approximate or calculate the output of a social choice function while only eliciting small amounts of information from voters (e.g. through pairwise comparisons) [Conitzer and Sandholm 2005; Lee et al. 2014; Lu and Boutilier 2011; Service and Adams 2012]. Our paper can be viewed as an attempt to create a thread of research mirroring preference elicitation, but for deliberative processes. Thus far, computational social choice has primarily viewed decision-making and preference elicitation from the perspective of efficiency, accuracy, and strategic issues. We propose that deliberation is a valuable new dimension to consider in social choice and provide a small step in this direction. We discuss this direction more in the concluding section.

Median graphs have been studied before in the context of voting. For instance, it is known that for median graphs, the Condorcet winner is strongly related to the generalized median [Bandelt and Barthelmy 1984; Saban and Stier-Moses 2012; Wendell and McKelvey 1981]. Nehring and Puppe [Nehring and Puppe 2007] show that any single-peaked domain which admits a non-dictatorial and neutral strategy-proof social choice function is a median space. Clearwater et. al. also showed that any set of voters and alternatives on a median graph will have a Condorcet winner (their full result is stronger than this) [Clearwater et al. 2015].

A similar high-level triadic decision-making process was proposed in a prior paper by the authors [Goel and Lee 2012] and analyzed for the restricted case (described in Figure 1) of opinions on a line. In this paper, we consider richer small group decision dynamics such as majority rule. We also consider spaces more complex than the line and provide a general framework for analyzing deliberative decision-making that enables us to prove an impossibility result for dyads.

To our knowledge, we are not aware of other literature on algorithmic or mathematical models for scaling deliberation. However, the problem of scaling deliberation has been discussed in the political science community. Deliberation was initially conceptualized as an exchange of arguments among a small group of rational individuals. Recent developments acknowledge that deliberation should be thought of in richer ways, including as an activity that takes place at the system level, among groups. [Parkinson and Mansbridge 2012]. In other words, it has recently become possible to think of deliberation as a task that can be broken down into various components and performed at different social levels and by large numbers of individuals. There have also been several practical initiatives for scaling deliberation such as Deliberative Polling [Fishkin et al. 2000], in which a single representative sample of participants is brought together to deliberate, and the 21st Century Town Hall Meeting [Lukensmeyer and Brigham




Fig. 3: The green, brown, and white participants in the two separate graphs have the same opinion structure since their positions relative to their convex hulls (the dotted blue lines) are identical. In this example, a locally consistent small group decision function would have to assign the same probability to the event in which the participants choose the yellow participant in either graph.

2002], in which the entire set of participants are divided into tables of size 10 for a single round of deliberation followed by voting.

The study of majority rule dynamics, which is an important part of deliberative decision-making, has also been a long-studied problem. A particular relevant experimental result analyzed the majority rule dynamic in groups of five and showed that the solution concept that performed best (out of 16 considered) was the majority rule equilibrium [Fiorina and Plott 1978]. Our work is a natural next step of this observation applied to the goal of scaling deliberative decision-making. Namely, given experimental evidence for the ability of small groups to come to consensus on the majority rule equilibrium, how can we use these small groups to make good decisions for larger groups?

Finally, we also note that our work has strong connections to opinion formation dynamics. Deliberative decision-making processes are essentially opinion formation dynamics for which one can algorithmically choose the sequence of interactions, rather than having them decided by a given social network [DeGroot 1974] or a nature-induced flocking process [Tsitsiklis 1984]. Most models of opinion formation use a weighted averaging dynamic and consider the mean rather than the median. We note that our main result on the dichotomy between triadic and dyadic decision-making can be adapted for opinion formation dynamics. Namely, for any opinion formation dynamic to have a chance at approximating the median, individuals should talk to at least two people before they update their opinions. Otherwise, the group consensus will be too easily influenced by extreme opinions.

5. DYADS CANNOT TIGHTLY APPROXIMATE THE GENERALIZED MEDIAN

We now prove our impossibility result for dyadic decision-making. This result holds for all small group decision functions g satisfying an axiom we call local consistency. Informally, local consistency captures the fact that the decision made among a set of non-strategic participants should only depend on the relationship structure of the participant opinions to each other. In other words, if one set of opinions is a “translation” of a second set of opinions, then the small group decision of the second set should be the same “translation” of the small group decision of the first set. This is depicted visually in Figure 3.

More formally, we use the notion of convex sets and convex hulls to capture the *opinion structure* of a set of participants. A convex set is any set such that for any x, y in the set, all shortest paths between x and y also lie in the set. The convex hull of S , denoted as C_S , is the smallest convex set containing S .

Definition 5.1. S and S' are said to have the same opinion structure under ψ if $\psi : C_S \rightarrow C_{S'}$ is an isometry (distance-preserving map) between C_S and $C_{S'}$ such that $S' = \psi(S)$.

Definition 5.2. The small group decision function g is said to be locally consistent if, (1) for any S , $g(S) \in C_S$, and (2) for any S and S' with the same opinion structure under ψ , $g(S') = \psi(g(S))$.

For equally influential participants, this definition encompasses well-known mathematical models in opinion formation such as DeGroot and flocking [DeGroot 1974; Tsitsiklis 1984]. We also note that, while we only focus on equally influential participants in this paper due to space constraints, our impossibility results also generalize to participants with varying degrees of influence (with an appropriate generalization of the model to account for an influence parameter).

THEOREM 5.3. *Consider the LDSG process $L = (V, f, g)$ where f is defined in any way such that $|S_t| = 2$ for all t , and $g_t = g(S_t)$ where $g(\cdot)$ is any locally consistent function. Then there exists a configuration of the participant opinions on a line graph such that the final consensus decision \hat{x} satisfies*

$$\mathbb{E}[D(\hat{x})] \geq \left(\frac{9}{8} - o(1)\right) D(x^*)$$

where x^* is the generalized median.

PROOF. The proof is by constructing a counterexample (see Appendix D). \square

In other words, simple examples exist for which the error ϵ is always lower bounded by $1/8$, regardless of how many participants are involved or how many small group interactions take place. We point out that a constant factor approximation is not satisfying since, in many scenarios, even the *worst* candidate is a constant factor approximation.

6. $O(N \log^2 N)$ TRIADS TIGHTLY APPROXIMATE THE GENERALIZED MEDIAN

In contrast to the impossibility result concerning dyadic decision-making, we find that triadic decision-making under a simple majority rule process is able to find an extremely tight approximation of the generalized median while only requiring an average of $O(\log^2 n)$ triads per participant.

The main step in the proof is to first assume that every triad decides on their (local) generalized median. We show that if this is true, then we can find a tight approximation of the (global) generalized median. This assumption is justified by showing that the generalized median is the equilibrium of the majority rule process for any small group coming from the preference domain we are considering (median graphs). This section will only consider non-strategic agents. Section 7 will extend this to the strategic case.

6.1. Triadic decision-making

The small group selection function f we use for selecting small groups is a simple randomized process. At each step, three tokens will be selected uniformly at random with replacement. The participants holding these tokens form the group S_t at time t .

Our result (Theorem 6.3) in Section 6.2 does not depend on any specific small group process, only on the assumption that every three-person interaction decides on the generalized median. This section, however, will motivate the generalized median by considering small groups which make their final decision through a specific process: majority rule (described below, and depicted in Figure 4). We model the small group decision function as the majority rule equilibrium and consider only those graphs for which a majority rule equilibrium exists for every group of three. For such graphs, it




Fig. 4: A block diagram of triadic decision-making with majority rule. The process starts (S) with selecting a small group (“Small group selection”). After selecting a small group (a), the small group participates in the majority rule process which starts with a “Propose” action. After a proposal (c), all three participants “Vote” on the proposal. If there was no motion to end which was accepted by two of the three participants, then another proposal is taken (d). This repeats until a participant motions to end, and two of the three participants accept. If all tokens now belong to a single participant, the overall process ends (E). Otherwise, another small group is selected (b) and the new group of participants engage in the majority rule process. Extensive form representations of each block are sketched in Figure 6.

turns out that the majority rule equilibrium coincides with the generalized median. Thus, we have the following formalization.

Definition 6.1. Define triadic decision-making as an LDSG process $L = (V, f, g)$ where $g_t = g(S_t) = \text{generalized-median}(S_t)$. For round $t + 1$, let $I_{t+1} = \{u, u', u''\}$ denote three tokens selected uniformly at random with replacement, i.e. $u, u', u'' \stackrel{i.i.d.}{\sim} U_n$, where U_n denotes the uniform distribution over integers $1, 2, \dots, n$. Then define f such that, for all t ,

$$S_{t+1} = \{y_u^t, y_{u'}^t, y_{u''}^t\}$$

Majority rule is a classical process for making a decision in a group. Majority rule occurs in rounds in which participants repeatedly suggest an alternate decision to some status quo. If at least two of the three participants vote for the alternate, then the status quo is replaced with the alternative. In practice, there are many variations in the way majority rule is implemented. For concreteness, we consider the following majority rule process for groups of three participants.

- (*Initial*) At the beginning, two participants are chosen uniformly at random. One is chosen to be the *current winner* and the other is chosen to be the *proposer*,
- (*Propose*) At the start of each round, the proposer proposes an alternative winner,
- (*Vote*) All three participants then simultaneously vote for either the current winner or the suggested alternative,
- (*Update*) The proposal which received the majority of the votes becomes the new current winner. If the vote was unanimous, then the proposer remains the same. Otherwise, the participant who lost the vote becomes the new proposer.
- (*Repeat*) Another round takes place starting from the Propose step,
- (*Termination*) During any round, the proposer can propose to end the process instead of proposing an alternative. If he does, and if a majority votes to accept, then the process ends. Otherwise, the proposer is updated according to the same rule in the Update step and another round takes place starting from the Propose step.

A majority rule equilibrium (MRE), also known as the Condorcet winner, is a proposal which would beat any other proposal, i.e. always receives a majority of the three votes cast. Whenever this proposal is voted on, it will win. Therefore, it is a commonly accepted solution concept for the decision resulting from a majority rule dynamic.

For this paper, we will consider only graphs for which every group of three participants has a majority rule equilibrium. It turns out that this requirement coincides

exactly with a class of graphs known as *median graphs* [Bandelt and Chepoi 2008; Knuth 2011] (see Theorem 6.2), which include common classes of graphs such as trees, grids, and squaregraphs. Moreover, for median graphs, the majority rule equilibrium of any three participants coincides with the generalized median.

THEOREM 6.2. *Consider any unweighted and undirected graph G . Let P denote the property in which every set of three nodes has a Majority Rule Equilibrium (Condorcet winner). Then,*

G satisfies P if and only if G is a median graph

Moreover, the Condorcet winner of any set of three nodes is the unique generalized median for that set.

PROOF. See Appendix E for the proof. \square

Thus, we let $g_t = g(S_t) = \text{majority-rule-equilibrium}(S_t) = \text{generalized-median}(S_t)$. As is also mentioned in the proof of Theorem 7.3, for any one of the three individuals (say, x), the generalized median is also the closest point to x that lies between y and z . Because of this, it is reasonable to suppose that the small group of individuals can find their equilibrium quickly.

6.2. Scalability theorem for triadic decision-making

We now show that a small number of triadic interactions can tightly approximate the generalized median. We highlight the stark contrast between the approximation achieved by triadic decision-making, for which the error $\sqrt{c \log n/n}$ tends to 0 as n becomes large, and dyadic decision-making, for which the error is lower bounded by $1/8$.

THEOREM 6.3. *Consider a triadic decision-making process $L = (V, f, g)$ where n participants V form a median graph. Then, with probability at least $1 - \frac{n \log n}{n^c}$, $c \geq 2$,*

$$T = O(cn \log^2 n) \quad \text{and} \quad D(x^*) \leq D(\hat{x}) \leq \left(1 + O\left(\sqrt{\frac{c \log n}{n}}\right)\right) D(x^*)$$

PROOF. A high-level understanding of our proof is that, for median graphs, the given urn dynamic can be reduced to several (simpler) coupled urn dynamics, represented by edges in E^* . E^* (defined in Appendix B) can be thought of roughly as “independent dimensions” of a median graph. Most of the proof is then tying the approximation factor or time of the overall urn process to those of the simpler coupled urn dynamics.

Our proof will be split into several lemmas. The lower bound $D(\hat{x}) \geq D(x^*)$ is trivial by the definition of x^* as the generalized median. To prove the upper bound, we first rewrite $D(x)$ as,

$$D(x) = \sum_{i=1}^n d(x, x_i) = \sum_{i=1}^n \sum_{e \in E^*} \mathbb{1}\{e \in I_{xx_i}\} = \sum_{e \in E^*} |\{i \mid e \in I_{xx_i}\}|$$

where I_{xy} is the set of nodes lying on a shortest path between x and y (see Appendix A). The second equality comes from a property of median graphs stating that $d(x, y) = \sum_{e \in E^*} \mathbb{1}\{e \in I_{xy}\}$ (Lemma B.7 in the Appendix). From this, it follows that

$$\frac{D(\hat{x})}{D(x^*)} \leq \max_{e \in E^*} \frac{|\{i \mid e \in I_{\hat{x}x_i}\}|}{|\{i \mid e \in I_{x^*x_i}\}|}$$

In Lemma F.1, we show that

$$\Pr \left[\frac{|\{i \mid e \in I_{\hat{x}_i}\}|}{|\{i \mid e \in I_{x^*}\}|} > \left(1 + O \left(\sqrt{\frac{c \log n}{n}} \right) \right) \right] \leq \frac{1}{n^c},$$

We can now achieve our upper bound via a union bound on all edges of E^* , noting that $|E^*| \leq |E| \leq n \log n$ (Lemma B.6).

To bound T , we first define $T_{e=(u,v)}$ as the first time t at which all tokens y_i^t lie in either W_{uv} or W_{vu} (informally, either closer to node u or node v , see Definition A.3). By Lemma B.5, we have

$$T = \max_{e \in E^*} T_e$$

We show in Lemma F.2 that $T_e = O(cn \log^2 n)$ with probability at least $1 - \frac{1}{n^c}$. Then we can achieve our result $T = O(cn \log^2 n)$ via a union bound on all edges of E^* . \square

When one considers grid graphs and trees, one can derive remarkably strong results using roughly the same proof techniques. Triadic decision-making is essentially able to find the *exact* generalized median with high probability, a very surprising result given the extreme randomness of the process.

THEOREM 6.4. *Consider $n = 9k^2$ participants which make up a $k \times k$ grid, so that each point in the grid has 9 participants. Assume k is odd. Then with probability greater than 99.4%, the winning participant \hat{x} is the exact generalized median x^* (the middle point of the grid) regardless of how large k is.*

THEOREM 6.5. *Consider n participants which make up a binary tree of height h . Then with probability greater than $1 - 4e^{-n/16}$, the winning participant \hat{x} is either the root node (the generalized median) or one of its two children. This is again true regardless of how large h is.*

Even stronger results hold for higher-dimensional grids or higher degree trees.

6.3. A simpler, restricted variant

The downside of the above setting is that participants will need to take time to discover the opinions of other participants. An implementation of this in practice would require more discovery mechanisms and tools. An accurate analysis would require a model of this discovery process, which is out of the scope of this paper.

In this section, we consider a simpler setting in which the small group decision is restricted to one of the members of the triad itself, i.e. $g_t \in S_t$. It turns out that an extremely simple process exists for finding the majority rule equilibrium under this restriction: simply ask each participant to vote on one of the other two participants. The participant who receives the largest number of votes wins. If there is a tie, then there is no exchange in tokens. This dynamic was analyzed in [Goel and Lee 2012] and shown to obtain essentially similar results for the case of opinions on a line. Here, we extend these results to star nodes (Theorem 6.6) and also give simulations indicating that it also works for trees and grids (Figure 5). Extending this to all median graphs (and perhaps beyond) remains an interesting open problem.

THEOREM 6.6. *Consider n participants on a star graph, for which $j = \Omega(\sqrt{cn \log n})$ of the participants are located at the root, and for which the remaining $n - j$ participants are distributed evenly among the k leaves. Then with probability greater than $1 - n^{-c}$, the winning participant \hat{x} is a participant at the root node.*




Fig. 5: These plots show the number of wins for nodes over 1000 runs of restricted triadic decision-making. We consider two graphs: (a) a binary tree of height 10, and (b, c) a grid of size 21x21. Each of these are simulated with two initial conditions: one or five initial tokens per participant. With five initial tokens, the root node of the tree wins 63.9% of the time and the top three nodes win 99.9% of the time. Similarly, with five initial tokens, the center node of the grid wins 44.7% of the time and the middle 9 nodes win 99.6% of the time.

PROOF. We use $\{R_t, t \geq 0\}$, $R_t \in \mathbb{Z}^{k+1}$, to denote a markov chain which is induced by restricted triadic decision-making (RTD) on the star. At time t ,

- (1) R_{t1} is the number of tokens at the center (root node) of the star,
- (2) R_{tj} is the number of tokens at the $j - 1$ -th leaf node of the star for $j = 2, \dots, k + 1$.

We use $\{R'_t, t \geq 0\}$, $R'_t \in \mathbb{Z}^2$, to denote a truncated representation which encodes the number of tokens at the center of the star and the *max* number of tokens at the leaf nodes. Specifically,

$$R'_{t1} = R_{t1}, \quad \text{and} \quad R'_{t2} = \max_{2 \leq j \leq k+1} R_{tj}$$

It is easy to verify that $\{R_t, t \geq 0\}$ is a markov chain because restricted triadic decision-making only depends on the positions of the tokens. However, the truncated representation $\{R'_t, t \geq 0\}$ is *not* a markov chain since transition probabilities depend on the distribution of the tokens among the leaves.

A major step in our proof will be to couple $\{R_t, t \geq 0\}$ to another stochastic process $\{S_t, t \geq 0\}$, $S_t \in \mathbb{Z}^{k+1}$, which we will call the concentrated variant of restricted triadic decision-making on the star (CRTD). The state space of S_t is the same as that of R_t and also describes a partition of n tokens over the $k + 1$ nodes of the star graph. At time t ,

- (1) S_{t1} is the number of tokens at the center (root node) of the star,
- (2) S_{tj} is the number of tokens at the $j - 1$ -th leaf node of the star for $j = 2, \dots, k + 1$.

A step of this markov chain involves first taking a step according to the normal RTD transition resulting from a triad. Once this is done, the maximum number of tokens on the leaf nodes is calculated and the tokens on the leaves are then *concentrated* while preserving the calculated maximum and ensuring that $S_{t2} \geq S_{t3} \geq \dots \geq S_{t(k+1)}$. As a simple example, suppose there are 5, 4, 3, and 2 tokens on leaf nodes $S_{t2}, S_{t3}, S_{t4}, S_{t5}$ respectively after a step of the RTD transition. Then the max number of tokens on the leaf nodes is 5, and concentrating the tokens would result in 5, 5, 4, and 0 tokens on leaf nodes $S_{t2}, S_{t3}, S_{t4}, S_{t5}$ respectively. We define a corresponding truncated representation

to be $\{S'_t, t \geq 0\}$, $S'_t \in \mathbb{Z}^2$, where

$$S'_{t1} = S_{t1}, \quad \text{and} \quad S'_{t2} = \max_{2 \leq j \leq k+1} S_{tj}$$

It is easy to verify that *both* $\{S_t, t \geq 0\}$ and $\{S'_t, t \geq 0\}$ are markov chains. With these definitions, we can state the following three lemmas (proof in Appendix G):

- (1) (Lemma G.1): Consider any two CRTD dynamics X_t, Y_t , and their corresponding truncated representations X'_t, Y'_t . Suppose that $X'_{j1} \leq Y'_{j1}$ and $X'_{j2} \geq Y'_{j2}$ for some j . Then we show that a coupling exists such that $X'_{(j+1)1} \leq Y'_{(j+1)1}$ and $X'_{(j+1)2} \geq Y'_{(j+1)2}$.
- (2) (Lemma G.2): Consider a CRTD dynamic Y_t and a RTD dynamic Z_t and their corresponding truncated representations Y'_t, Z'_t . Suppose that $Y'_{j1} = Z'_{j1}$ and $Y'_{j2} = Z'_{j2}$ for some j . Then we show that a coupling exists such that $Y'_{(j+1)1} \leq Z'_{(j+1)1}$ and $Y'_{(j+1)2} \geq Z'_{(j+1)2}$.
- (3) (Lemma G.3): Consider a CRTD dynamic X_t and a RTD dynamic Z_t and their corresponding truncated representations X'_t, Z'_t . Suppose that $X'_{01} \leq Z'_{01}$ and $X'_{02} \geq Z'_{02}$. Then we show that a coupling exists such that $X'_{t1} \leq Z'_{t1}$ and $X'_{t2} \geq Z'_{t2}$ for all t and for every history of the markov chain.

We can now wrap up our proof. Consider a CRTD dynamic X_t and a RTD dynamic Z_t and their corresponding truncated representations X'_t, Z'_t . Suppose that $X_0 = (a, (n-a)/2, (n-a)/2, 0, \dots, 0)$ and that $X'_{01} \leq Z'_{01}$ and $X'_{02} \geq Z'_{02}$. By Lemma G.3, a coupling exists such that whenever the center of the star wins in X_t , i.e. X'_t converges to $(n, 0)$, the center of the star must also win in Z_t , i.e. Z'_t has converged to $(n, 0)$. Therefore the probability that the center of the star wins in Z_t is at least as high as the probability that the center of the star wins in X_t .

We now calculate the probability that the center of the star wins in X_t . The main observation is that, for our initial condition, the star being considered in X_t only has two leaves. It is easy to verify that in this case, the CRTD dynamic is identical to the RTD dynamic (the leaves are always already “concentrated”). Moreover, since a star with two leaves is simply a line, we already have the probability that the center of the star wins in X_t . For $a = \Omega(\sqrt{cn \log n})$, this probability is greater than $1 - n^{-c}$, which concludes our proof. \square

7. TRIADIC DECISION-MAKING SUCCEEDS WITH STRATEGIC AGENTS

One cannot naively apply the prior results to strategic agents in a full LDSG process. This is because a strategic agent wants to ultimately maximize the utility he receives from the global decision \hat{x} . Thus, an agent might settle for a worse consensus decision in a given small group if it could help in the long run. We show that this does not happen and that, in fact, we can achieve the same results described in Theorem 6.3 for strategic agents under a subgame perfect Nash equilibrium [Shoham and Leyton-Brown 2008].

7.1. Game theoretic formalization

The entire decision-making process induced by Definition 6.1 and the majority rule process (Section 6.1) can be formalized as an infinite horizon extensive form game (see Figures 4 and 6). A full strategy for the game involves choosing, given every possible past history, a strategy for how one votes and proposes motions. For each individual




Fig. 6: The game-theoretic representation of triadic decision-making with majority rule is an infinite horizon extensive form game made out of several blocks as depicted in Figure 4. The “Small group selection” block is played by the small group selection function f who selects one of n^3 possible triads according to the probability distribution described by Definition 6.1. The “Propose” block for a player u consists of $n + 1$ actions corresponding to either proposing to replace the current winner with one of the n participants or motioning to end the process. The “Vote” block for the small group x, y, z consists of a simultaneous vote to either accept or reject the proposal. The fact that these votes are simultaneous is represented formally by making the nodes for y and for z a single information set.

i , the utility assigned from the game is defined as $u_i(\hat{x}) = -d(x_i, \hat{x})$ if a winner \hat{x} is produced, and $-\infty$ if the game never terminates.¹

7.2. Approximation and time results for strategic agents in the unrestricted setting

We show that a strategy which we call truthful bargaining achieves a subgame perfect Nash equilibria for the aforementioned game while also returning the generalized median as the consensus decision of each small group. This means that the results of Theorem 6.3 carry over to this equilibrium.

The truthful bargaining strategy is simple in that it is independent of past small group interactions. Roughly speaking, it requires a participant to vote truthfully between opinions, and to propose his best “bargaining point” when one exists, and to motion to end the process otherwise. When voting on a motion to end, truthful bargaining requires a participant to accept the motion when he has no bargaining points, and reject it otherwise.

Definition 7.1. Suppose that w is the current winner in a triadic majority rule process between participants x, y, z . Define the bargaining points for participant x to be any point which is not the current winner w , but lies on a shortest path between w and x as well as a shortest path between w and at least one of y or z . The best bargaining point for x is the bargaining point which is closest to himself.

Definition 7.2. Define truthful bargaining to be the strategy in which, independent of the past history,

- If a vote is taken for a proposed alternative, the agent votes truthfully according to his preference,
- If it is the agent’s turn to propose a motion, he proposes his best bargaining points if one exists, and motions to end otherwise,
- If a vote is taken on a motion to end, he accepts if he has no bargaining points, and rejects it otherwise.

For all of these decisions, ties are broken arbitrarily.

¹We use this definition of utility due to its simplicity and space constraints; however, our results hold for a much broader class of utilities known as single-peaked preferences.

THEOREM 7.3. *Consider a triadic decision-making process $L = (V, f, g)$ where participants V form a median graph and g_t is determined by a majority rule process between the participants S_t . Then the truthful bargaining strategy is a subgame perfect Nash equilibrium. Moreover, $g_t = \text{generalized-median}(S_t)$ for all t .*

PROOF. We prove this through several lemmas. We first show that, if all participants follow truthful bargaining, then the consensus decision of each triad S will be the generalized median of S (Lemma I.1). Following this, we consider any triadic round of any subgame in which one participant u deviates. We show that regardless of the strategy that u deviates to, either the round never ends or the generalized median of S lies on a shortest path from u to the resulting winner w' (Lemma I.8). Finally, we leverage Lemma I.8 to showing that, for any subgame, if one participant u deviates, then the utility received is stochastically dominated by the utility received if u did not deviate (Lemma I.9), which concludes our proof. \square

Since we have shown that $g_t = \text{generalized-median}(S_t)$, the approximation and time bounds of Theorem 6.3 extend to the strategic case.

8. CONCLUDING DISCUSSION

In this paper, we proposed a model in which sequences of small group interactions are used to scale up decision-making. We showed that small groups can find the wisdom of the crowd efficiently — sometimes in situations when the group at large is unable to come to a stable decision — but we also showed that the feasibility of doing so can depend strongly on the size of the small groups. This phenomenon, which also applies to simple opinion formation dynamics, is succinctly summarized by the following aphorism: while two heads are better than one, three heads are better than two.

There are many questions remaining that we view as promising directions for future research. In our paper, we studied a particular approach and dynamic to scaling deliberative decision-making. We believe that the problem of scaling deliberative decision-making deserves a broader and more fundamental framework. Perhaps one can define a model for *deliberative social choice*, in which a global choice is calculated from a subset of small group choices. This should be done first for participants whose opinions do not change (like in this paper and in majority rule dynamics), but then should be extended to true deliberative models which account for transfer of information and evolution of opinions. Such mathematical work should also be coupled to systems work which consider motivational challenges and interaction interfaces. Ultimately, all of these should work towards a realization of deliberative democracy in practice.

ACKNOWLEDGMENTS

The authors would like to thank Helene Landemore for discussing and pointing us to relevant deliberative democracy literature, and the reviewers (over the course of several submissions) for pushing us to motivate and communicate our paper more clearly. This work was supported by the Army Research office Grant No. 116388, the Office of Naval Research Grant No. 11904718, and by the Stanford Cyber Initiative. David T. Lee was also supported by an NSF graduate research fellowship and a Brown Institute for Media Innovation Magic Grant.

REFERENCES

- James Aspnes. 2013. Notes on Randomized Algorithms. (2013). <http://www.cs.yale.edu/homes/aspnes/classes/469/notes.pdf>.
- Hans-Jürgen Bandelt and Jean-Pierre Barthelemy. 1984. Medians in median graphs. *Discrete Applied Mathematics* 8, 2 (1984), 131–142.

- Hans-Jürgen Bandelt and Victor Chepoi. 2008. Metric graph theory and geometry: a survey. *Contemp. Math.* 453 (2008), 49–86.
- Adam Clearwater, Clemens Puppe, and Arkadii Slinko. 2015. Generalizing the Single-Crossing Property on Lines and Trees to Intermediate Preferences on Median Graphs. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*.
- Vincent Conitzer and Tuomas Sandholm. 2005. Communication complexity of common voting rules. In *Proceedings of the 6th ACM Conference on Electronic Commerce (EC)*. 78–87.
- Morris H. DeGroot. 1974. Reaching a Consensus. *J. Amer. Statist. Assoc.* 69 (1974), 118–121. Issue 345.
- Robin Dunbar. 1992. Neocortex size as a constraint on group size in primates. *Journal of Human Evolution* 22 (1992), 469–493. Issue 6.
- Morris P. Fiorina and Charles R. Plott. 1978. Committee Decisions under Majority Rule: An Experimental Study. *The American Political Science Review* 72, 2 (1978), 575–598.
- James S. Fishkin. 2011. *When the People Speak: Deliberative Democracy and Public Consultation*. Oxford University Press.
- James S. Fishkin, Robert C. Luskin, and Roger Jowell. 2000. Deliberative polling and public consultation. *Parliamentary Affairs* 53 (2000), 657–666. Issue 4.
- Archon Fung and Erik Olin Wright. 2003. *Deepening Democracy: Institutional Innovations in Empowered Participatory Governance*. Verso Press.
- Francis Galton. 1907a. Letters to the Editor: The Ballot-Box. *Nature* 75, 1952 (1907), 509–510.
- Francis Galton. 1907b. Vox populi. *Nature* 75 (1907), 450–451.
- Ashish Goel and David T. Lee. 2012. Triadic Consensus: A Randomized Algorithm for Voting in a Crowd. In *Proceedings of the 8th Workshop on Internet and Network Economics (WINE)*.
- Benjamin Golub and Matthew O. Jackson. 2010. Naïve Learning in Social Networks and the Wisdom of Crowds. *American Economic Journal: Microeconomics* 2, 1 (2010), 112–149.
- Amy Gutmann and Dennis Thompson. 2004. *Why Deliberative Democracy?* Princeton University Press.
- Reginald Hawthorne Hooker. 1907. Mean or median. *Nature* 75 (1907), 487–488.
- Wilfried Imrich and Sandi Klavžar. 2000. *Product graphs, structure and recognition*. Wiley.
- Sandi Klavžar and Henry Martyn Mulder. 1999. Median graphs: characterizations, location theory and related structures. *Journal of Combinatorial Mathematics and Combinatorial Computing* 30 (1999), 103–128.
- Donald Ervin Knuth. 2011. *The Art of Computer Programming*. Vol. 4A: Combinatorial Algorithms, Part 1. Pearson Education, Inc.
- Hélène Landemore and Jon Elster (Eds.). 2012. *Collective Wisdom: Principles and Mechanisms*. Cambridge University Press.
- Bibb Latané, Kipling Williams, and Stephen Harkins. 1979. Many hands make light the work: The causes and consequences of social loafing. *Journal of Personality and Social Psychology* 37 (1979), 822–832. Issue 6.
- David T. Lee, Ashish Goel, Tanja Aitamurto, and Hélène Landemore. 2014. Crowdsourcing for Participatory Democracies: Efficient Elicitation of Social Choice Functions. In *Proceedings of the 2nd AAAI Conference on Human Computation and Crowdsourcing (HCOMP)*.
- Hendrik P. Lopuhaä and Peter J. Rousseeuw. 1991. Breakdown points of affine equivariant estimators of multivariate location and covariance matrices. *The Annals of*

- Statistics* 19, 1 (1991), 229–248.
- Jan Lorenz, Heiko Rauhut, Frank Schweitzer, and Dirk Helbing. 2011. How social influence can undermine the wisdom of crowd effect. *Proceedings of the National Academy of Sciences* 108, 22 (2011), 9020–9025.
- Tyler Lu and Craig Boutilier. 2011. Robust Approximation and Incremental Elicitation in Voting Protocols. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*.
- Carolyn J. Lukensmeyer and Steve Brigham. 2002. Taking Democracy to Scale: Creating a Town Hall Meeting for the Twenty-First Century. *National Civic Review* 91 (2002), 351–366. Issue 4.
- Crawford Brough Macpherson. 1977. *Life and Times of Liberal Democracy*. Oxford University Press.
- Nolan McCarty, Keith T. Poole, and Howard Rosenthal. 2006. *Polarized America: The Dance of Ideology and Unequal Riches*. MIT Press.
- Carrie Menkel-Meadow. 2011. Scaling up Deliberative Democracy as Dispute Resolution in Healthcare Reform: A Work in Progress. *Law and Contemporary Problems* 74 (2011), 1–30.
- George A. Miller. 1956. The magical number seven, plus or minus two: some limits on our capacity for processing information. *Psychological review* 63 (1956), 81–97. Issue 2.
- Henry Martyn Mulder. 1978. The structure of median graphs. *Discrete Mathematics* 24, 2 (1978), 197–204.
- Klaus Nehring and Clemens Puppe. 2007. The structure of strategy-proof social choice Part I: General characterization and possibility results on median spaces. *Journal of Economic Theory* 135 (2007), 269–305. Issue 1.
- John Parkinson and Jane Mansbridge. 2012. *Deliberative Systems: Deliberative Democracy at the Large Scale*. Cambridge University Press.
- Carole Pateman. 1979. *The Problem of Political Obligation: A Critique of Liberal Theory*. New York: John Wiley & Sons.
- Keith T. Poole and Howard Rosenthal. 1984. The polarization of american politics. *The Journal of Politics* 46 (1984), 1061–1079. Issue 4.
- Daniela Saban and Nicolas Stier-Moses. 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)*.
- Travis C. Service and Julie A. Adams. 2012. Communication complexity of approximating voting rules. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 593–602.
- Yoav Shoham and Kevin Leyton-Brown. 2008. *Multiagent Systems: Algorithmic, Game-theoretic, and Logical Foundations*. Cambridge University Press.
- Graham Smith. 2009. *Democratic Innovations: Designing Institutions for Citizen Participation*. Cambridge University Press.
- James Surowiecki. 2005. *The Wisdom of Crowds*. Anchor.
- John N. Tsitsiklis. 1984. *Problems in Decentralized Decision Making and Computation*. Ph.D. Dissertation. MIT.
- Richard E. Wendell and Richard D. McKelvey. 1981. New perspectives in competitive location theory. *European Journal of Operational Research* 6, 2 (1981), 174–182.