Social Deliberation: Democratic Deliberation Through Unbiased Social Interactions

Anonymous Author(s)
Submission Id: 3609

Abstract

Deliberative democracy asserts that democratic legitimacy requires authentic deliberation rather than simply aggregating individual preferences through voting. While deliberation plays a crucial role in modern democracy, its implementation often relies on deliberative microcosms formed through random sampling or sortition, since in many settings, direct deliberation among all voters is impractical. Although these approaches have been successful in fostering informed and public-spirited decision making, they suffer from significant limitations. First, they exclude most citizens from direct participation, and citizens in the broader population who have not participated in deliberation may hold different views from their elected representatives. Second, organizing and facilitating large-scale deliberation requires substantial organizational resources.

In this work, we model social interactions of the voters as a form of deliberation and name this process *social deliberation*. In this process, voters start by only having the knowledge of their own utilities, and keep learning about the social welfare of the community through iterative social interactions. We capture the effect of the underlying graph that represents the social network among the voters, as well as voters' behavior in aggregating the information that they receive. Using distortion as a quantitative measure, we assess how social deliberation affects voting outcomes. Our results indicate that under mild assumptions about the network structure, when public-spirited voters engage in unbiased social communication, a constant distortion can be achieved, highlighting the potential efficiency of social deliberation in democratic decision making.

1 Introduction

Deliberative democracy is a cornerstone of modern democratic theory, emphasizing the role of reasoned discussion, mutual respect, and the exchange of diverse perspectives in reaching collective decisions. Unlike traditional democratic models that derive legitimacy primarily from vote aggregation, deliberative democracy shifts the foundation of legitimacy toward reasoned deliberation and informed consensus. The roots of deliberative democracy can be traced back to Aristotle and his notion of politics as "politics as the pursuit of the common good". A variety of mechanisms have been proposed to implement deliberative democracy, such as sortition-based citizen assemblies and deliberative microcosms, where small, representative groups engage in structured discussions. Although these approaches have been successful in fostering informed and public-spirited decision-making, they suffer from significant limitations. First, they exclude most citizens from direct participation, and citizens in the broader population who have not participated in deliberation may hold different views from their elected representatives. Second, organizing and facilitating meaningful large-scale deliberations requires substantial organizational resources.

In contrast to centralized deliberations, Mansbridge [32] emphasizes the importance of *everyday talk* in the deliberative process. She argues that although everyday talk may not meet all formal criteria typically associated with the term deliberation, it still plays a crucial role within the deliberative system. Through these social interactions, voters gain a clearer understanding of both their individual preferences and the collective benefits for society. Today, with the rapid growth of the internet, communication is easier than ever, and people interact constantly across distances through the internet and social media. This implies that when people discuss public decisions in everyday conversations, these interactions can potentially serve as large-scale deliberation. While centralized deliberation has been extensively studied before, there is little to no work in the social choice literature that studies these everyday interactions. We use the term "*Social Deliberation*" to describe the discussions among voters about a public decision, and our objective is to examine the impact of social deliberation on the quality of the final winner or outcome. We show that, under mild assumptions, if voters interact within their local networks, the process can yield a near-optimal decision. Our results highlight the substantial effect of day-to-day talk from this aspect and motivate future research on this phenomenon.

To measure the efficiency of deliberative processes, we employ the distortion framework introduced by Procaccia and Rosenschein [37]. In this framework, we assume that the preferences that voters submit come from the underlying utilities that they associate to each alternative. Distortion of a voting rule captures the worst-case ratio between the social welfare of the output selected by the rule and the social welfare of the optimal choice. This notion has gained a lot of attention in the literature over the past decade. We refer to the recent survey of Anshelevich et al. [5] for an extensive overview of the related work. Mansbridge [32] argues that "If a deliberative system works well, it filters out and discards the worst ideas available on public matters, while it picks up, adopts, and applies the best ideas." With this approach, distortion serves as an effective measure of the efficiency of a deliberative system, given its sensitivity to the selection of suboptimal ideas, even in specific scenarios. If a voting system has low distortion, it is evidence that it never selects an alternative that could be considered "bad" or, as we define it, one with significantly lower social welfare compared to the optimal outcome.

A substantial body of literature on distortion focuses on voting systems with ranked preferences, where each voter submits a ranking of the alternatives. In such settings, and without any assumptions on the utilities, the distortion of any rule will be unbounded. To prevent this strong negative outcome, much of the literature on distortion adopts the assumption that voters' utilities over alternatives sum up to 1. However, this assumption is not realistic in all contexts, as the outcome of the election could affect voters to varying degrees (see Flanigan et al. [26] for a specific example). In this work, rather than this strong assumption, we follow the recent line of work by Flanigan et al. [26] and Bedaywi et al. [9] that instead of making assumptions on the utilities, assume that voters are public-spirited, meaning that they care about the social good as well as their own good. Public-spiritedness, a well-established concept in political science, is considered one of the key prerequisites for effective democratic deliberation. In a deliberative process, voters learn about the preferences of others, but for this new information to influence their own preferences, they must care about the common good and act in a "public-spirited" manner. Flanigan et al. [26] demonstrate that public-spirited voter behavior significantly reduces the distortion of various voting rules. Specifically, they show that under public-spiritedness, the distortion of the Plurality rule remains linear in relation to the number of alternatives, and the Copeland rule can achieve constant distortion if all voters exhibit a nonzero level of public-spiritedness. While they prove these bounds without any assumption on the underlying utilities, they make use of a strong assumption that the voters have complete knowledge of the social welfare associated with each alternative and shape their public-spirited preferences accordingly.

We believe that this assumption is not realistic in large-scale elections, since knowing the social welfare requires knowing the actual utility of each voter for each alternative. In addition, if the social welfare was a common knowledge, why not choose the alternative with the maximum social welfare? To address these concerns, we assume that each voter keeps her own estimates of the social welfare, and updates these estimates through social deliberation. Our key contribution is to demonstrate that even when voters submit their preferences based on these estimates, the results of Flanigan et al. [26] still holds asymptotically.

To model voter interactions and the process by which they update their estimates, we employ a standard discrete-time dynamics model. These interactions are modeled using a graph, where each voter exchanges her estimates with neighboring voters over several time steps. Initially, voters have knowledge only of their own utilities and they iteratively update their estimates based on the information received from their neighbors. Following the seminal work of DeGroot [17], we prove that for voters to reach consensus on their estimates, it is sufficient for the graph that represents the social interactions to be connected and for the voters to assign a positive weight to their own estimates.

Regarding the structure of the underlying graph, we analyze both deterministic and stochastic graphs. It is straightforward to observe that the model of Flanigan et al. [26] is a special case of our model where the underlying network is a complete graph. However, graphs that model social networks have diverse structural properties, often characterized by community formation, small-world effects, and heterogeneous degree distributions. Our main focus in this work is on the Erdős–Rényi model [20, 21], one of the most well-studied models for generating random graphs. This model provides a useful approximation for social networks in settings where interactions are relatively unstructured, connections form randomly, and the focus is on aggregate statistical properties rather than specific community structures. Additionally, we investigate whether our results extend to other stochastic models and deterministic network structures.

1.1 Related Work

Our work is situated at the intersection of several domains which have been extensively explored in prior research. In this section, we provide a brief overview of related studies in these topics.

Democratic Deliberation. Deliberation is widely recognized as a foundational element of modern democracy. Scholars argue that engaging citizens in reasoned discussion enhances democratic legitimacy and decision-making [15, 23, 24]. One of the main practical attempts to facilitate deliberation is to form citizen assemblies via sortition. There is a rich body of literature on sortition from various points of views [13, 18, 25, 35]. Other efforts aim to make deliberation practical, such as the idea of deliberation day by Ackerman and Fishkin [2], where people meet in public spaces for debates. Deliberation has also recently attracted attention in social choice [22, 30, 33].

Voting and Social Networks. The foundational studies on political voting behavior highlight the critical role of social networks in shaping political choices [10, 38]. For instance, Abrams et al. [1] propose a model exploring how social interactions influence voter turnout. A pioneering study on how social interactions shape voting is Conitzer [14], which models voting as reconstructing a ground-truth outcome from noisy votes within social networks. The nature of these networks supports systems like liquid democracy, where voters can delegate their votes to trusted peers. Building on this, Alouf-Heffetz et al. [3] provide a detailed experimental comparison of different voting systems including direct democracy, sortition, and liquid democracy within social networks. In a slightly different direction, Liu et al. [31] investigates how to identify influential vertices using a ranked voting approach.

Distortion. Distortion was first introduced by Procaccia and Rosenschein [37] in a utilitarian framework and later extended to the metric setting by Anshelevich et al. [4]. Over the past two decades, it has been studied across various settings. In the single-winner case, Caragiannis and Procaccia [12] showed that if we assume the voters' utilities are unit-sum, the optimal distortion is $O(n^2)$, which is achievable by the Plurality rule. In the metric setting Anshelevich et al. [4] showed that the Copeland rule distortion is 5. This bound was later improved to the tight bound of 3 via the Plurality Matching rule [27]. Most recently, Gkatzelis et al. [28] designed a rule that achieves constant metric and $O(n^2)$ utilitarian distortion at the same time. We refer to [5] for a comprehensive overview of the results on distortion.

The closest to ours is Flanigan et al. [26], who study the distortion of several voting rules under public-spirited voter behavior. They show that Plurality has distortion $O(m/\gamma_{\min})$, while Copeland achieves a constant $O(1/\gamma_{\min}^2)$, where m is the number of alternatives and $\gamma_{\min}$ is the minimum public-spirit level. Building on this, Bedaywi et al. [9] extend the analysis to randomized rules and participatory budgeting, and Bagherideloue et al. [6] study distortion in the metric setting. There are also works on sortition distortion under metric assumptions [11, 29].

1.2 Our Contribution and Structure of the Paper

In Section 2, we introduce the model of social deliberation. In this model, voters are connected in a network, each maintaining an estimate of the social welfare of each alternative. They update their estimates iteratively by communicating with their neighbors. Formally, each voter has a weight vector over her neighbors and herself, and in each timestep, updates her estimate to be the weighted average of their estimates from the previous round. Later in Section 3, we prove that these estimates converge to a steady state if the underlying graph is connected and each voter assigns a non-zero weight to her own estimate from the previous round. Additionally, we show that in this state, each voter's contribution to the final estimate is proportional to her degree in the graph.

In Sections 4 and 5, we prove bounds on the distortion of the Plurality and Copeland rules for different classes of deterministic and stochastic graphs. The main observation here is that if voters are unbiased among their neighbors, the distortion of the Copeland rule remains constant, whereas with biased aggregation, the distortion can grow as a polynomial function of n. While our model generalizes that of Flanigan et al. [26] from complete graphs to general graphs, our bounds on distortion match theirs up to a constant factor that depends on the features of the graph.

While our results hold in the steady state of the estimates, in Section 6 we show how close the system must be to this state to effectively behave as if it has converged, and how long it takes for the estimates to reach this vicinity. We also characterize the assumptions under which the dynamics converge in sublinear time with respect to n (number of voters).

2 Preliminaries

For integer t let $[t] = \{1, 2, \dots, t\}$, and for set S let $\Delta(S)$ denote the set of all distributions over S. Let $N = [n]$ represent a set of n voters and $A = [m]$ be a set of m alternatives. Each voter $i \in N$ has a utility over alternative $a \in A$, and together these utilities form the utility matrix $U \in \mathbb{R}_{\geq 0}^{n\times m}$, where $U[i, a]$ denotes the utility of voter i for alternative a. We define the social welfare of alternative a, denoted by $\mathrm{SW}(U, a)$, as the total utility of the voters for a, i.e., $\mathrm{SW}(U, a) = \sum_{i\in N} U[i, a]$. If U is clear from the context, with slight abuse of notation, we use $\mathrm{SW}(a)$ instead of $\mathrm{SW}(U, a)$.

Communication Model. Each voter maintains estimates of social welfares and updates them over rounds by averaging her neighbors' estimates. Formally, $\mathrm{EW}_0[i, a] = nU[i, a]$ and for $t \geq 1$, $\mathrm{EW}_t[i, a] = \sum_{j\in N} T[i, j]$
, $\mathrm{EW}_{t-1}[j, a]$ with row-stochastic T (i.e., $\sum_j T[i, j] = 1$). Communications are bidirectional: if $T[i, j] > 0$ then $T[j, i] \geq 0$ (weights may differ). Writing $\mathrm{EW}_t = T\,\mathrm{EW}_{t-1}$, the estimation dynamics is $D = (U, T)$ with state space $\mathcal{D} = \mathbb{R}_{\geq 0}^{n\times m} \times [0, 1]^{n\times n}$. Under mild conditions (Section 3), EW_t converges to the steady state $\mathrm{EW}^* = \lim_{t\to\infty} \mathrm{EW}_t = (\lim_{t\to\infty} T^t)U = T^*U$.

Network of Voters. Given T, let G_T be the undirected graph on n voters with edge $\{i, j\}$ iff $T[i, j] > 0$. When clear, write G. Let d_i be degrees, and denote $d_{\min}, d_{\max}$, and $d_{\text{avg}} = \frac{1}{n}\sum_{i=1}^n d_i$.

Public-Spirited Behavior. We assume that voters exhibit public-spirited behavior. Following the model introduced by Flanigan et al. [26], we assume that each voter $i \in N$ has a public spirit (PS) level $\gamma_i \in [0, 1]$, which shows the level to which she cares about the common good. These public spirit levels together form *ps-profile* $\vec{\gamma}$. At the time of election, each voter submits her preference based on her own utility over the alternatives and her estimated social welfare for different alternatives. Formally, the PS-value of voter i for alternative $a \in A$ in the steady state is

$$V^*[i, a] = (1 - \gamma_i)U[i, a] + \gamma_i \frac{\mathrm{EW}^*[i, a]}{n}.$$

Note that in this work we mostly focus on analyzing the election after reaching the steady state. However, in Section 6 we define PS-value in each timestep similarly.

Elicitation. We consider *ranked preferences*: each voter i submits a ranking σ_i; together these form $\vec{\sigma}$. Preferences are induced by PS-values: $a \succ_i b$ implies $V^*[i, a] \geq V^*[i, b]$. We say $\vec{\sigma} \triangleright V$ if the profile is consistent with value matrix V; since V^* depends on $\vec{\gamma}$ and D, we write $\vec{\sigma} \triangleright (\vec{\gamma}, D)$ when it is consistent with the resulting V^*.

Aggregation Rule. A deterministic single-winner voting rule is a function f that takes a preference profile as input and selects a single alternative as the winner. In this paper, we consider two well-known voting rules: *Plurality* and *Copeland*. In both rules we give a score to each alternative and the winner is one of the alternatives with maximum score. In Plurality the score of each alternative is the number of voters that prefer him to all other alternatives, and in Copeland the score of each alternative is the number of alternatives that he dominates in pairwise comparison.

Distortion. For subclass $\mathcal{D}' \subseteq \mathcal{D}$, voting rule f, and ps-profile $\vec{\gamma}$,

$$\mathrm{dist}_{\mathcal{D}'}(f, \vec{\gamma}) = \sup_{\substack{\vec{\sigma}\in\Sigma,\ D\in\mathcal{D}':\\ \vec{\sigma}\triangleright(\vec{\gamma}, D)}} \frac{\max_{a\in A} \mathrm{SW}(U, a)}{\mathrm{SW}(U, f(\sigma))},$$

where Σ is the set of all preference profiles. We study how distortion scales with $\gamma_{\min} = \min_i \gamma_i$, and later consider stochastic graphs: for a distribution X over transition matrices, $\mathrm{dist}_X(f, \vec{\gamma}) = \mathbb{E}_{T\sim X}[\mathrm{dist}_{\mathcal{D}(T)}(f, \vec{\gamma})]$, and a decoupled variant in Section 5.

3 Convergence of Estimation Dynamics

In this section, we study the convergence of the estimation dynamics. Our goal is to determine a set of conditions under which the system converges to a steady state and further, to characterize the properties of this steady state.

Definition 3.1. We say that an estimation dynamics D and its transition matrix T are:

- **Unbiased**: If voters give the same weight to their neighbors, i.e., for any pair of voters $i, j \in N$, such that $T[i, j] \neq 0$, any other voter $k \in N$ we have $T[i, k] \in \{0, T[i, j]\}$.
- **Fully unbiased**: If the dynamics are unbiased and in addition the voters give the same weight to themselves as to others. That means for any pair of voters $i, j \in N$, $T[i, j] = \frac{1}{d_i+1}$, and $T[i, j] \in \{0, \frac{1}{d_i+1}\}$.
- **Diagonally uniform**: If for any two voters i and j, $T[i, i] = T[j, j]$.
- **Strongly Influenced**: If G_T is connected and for every voter i, $T[i, i] > 0$.

Our model for the evolution of voters' estimations is in fact a Markov process. Formally, a *Markov chain* C is a stochastic process that evolves over discrete time steps on a finite state space X. At each step t, the chain transitions from its current state to the next according to a probability distribution defined by the transition matrix P, where $P[i, j]$ represents the probability of moving from state i to state j. This process satisfies the property that the transition from a state depends only on the current state, not on the chain's history.

The convergence of the voters' estimations in our model relies on two key properties of Markov chains stated in Lemma 3.2. We refer to the work of Norris [36] for more details.

Lemma 3.2 (Convergence of Markov Chains). *Let C be a Markov chain with a finite state space X and a transition matrix P. We say C is* ***irreducible****, if any state of X can be reached from any other state in a finite number of steps. Formally, for all $i, j \in X$, there exists $n > 0$ such that $P^n[i, j] > 0$. In addition, we say that C is* ***aperiodic*** *if, for every state, the possible number of steps in sequences that start and end in that state have greatest common divisor one. Formally, C is aperiodic if for every state $i \in X$, $\gcd\{n \mid P^n[i, i] > 0\} = 1$, where* gcd *denotes the greatest common divisor of a set. Given that C is irreducible and aperiodic, then C has a unique stationary distribution $\vec{\pi}$, such that $\vec{\pi} = \vec{\pi}P$. Furthermore, $\lim_{t\to\infty} P^t = P^*$ exists and $P^*[i, j] = \pi_j$. Note that $\vec{\pi}$ only depends on P and is independent of the initial distribution.*

Building on Lemma 3.2, we establish that being strongly influenced is sufficient for an estimation dynamics to converge to a steady state in Theorem 3.3. Furthermore, under specific assumptions on the estimation dynamics, we derive a closed-form expression for this steady state in Theorem 3.4.

Theorem 3.3. *Let $T_{n\times n}$ be a strongly influenced transition matrix. Then, for any utility matrix U, estimation dynamics $D = (U, T)$ converges to a final welfare estimation matrix EW^*. Furthermore, the voters reach a consensus that is a linear combination of their utilities, i.e., there exists a vector $\vec{w} \in [0, 1]^n$ with $\sum_{i=1}^n w_i = 1$ such that, for any voter i and alternative a, $EW^*[i, a] = \sum_{j=1}^n w_j U[j, a]$. We refer to $\vec{w}$ as the Influence Weight Vector.*

Proof. Treat T as the transition matrix of a Markov chain C on N. Let $k \in \mathbb{N}$ denote the diameter of G_T, that is, the longest shortest-path between any pair i, j of voters. k is well-defined since G_T is connected. Note that, by definition $T^k[i, j] > 0$ and hence C is irreducible. Furthermore, since $T[i, i] > 0$, C is aperiodic as well. By Lemma 3.2, we can conclude that

$$\mathrm{EW}^* = \lim_{t\to\infty} \mathrm{EW}_t = (\lim_{t\to\infty} T^t)U = T^*U$$

exists and there also exists a unique stationary distribution $\vec{w} \in [0, 1]^n$ with $\sum_{i=1}^n w_i = 1$ (since it's a distribution). Since $T^*[i, j] = w_j$, then

$$\mathrm{EW}^*[i, a] = \sum_{j=1}^{n} T^*[i, j]U[j, a] = \sum_{j=1}^{n} w_j U[j, a]. \qquad \square$$

By Theorem 3.3 we know that if the estimation dynamics is strongly influenced, voters reach a consensus which can be identified by U and $\vec{w}$. In this case since all the voters have the same estimated welfare for the same alternative, we can drop i from $\mathrm{EW}^*[i, a]$ and with slight abuse of notation use $\mathrm{EW}^*[a]$ when applicable.

We now shift our focus to computing the influence weight vector. Theorem 3.4 which is borrowed from David and Fill [16] provides a convenient tool called **Markov Tree Formula** that allows us to derive a closed-form solution for $\vec{w}$.

Theorem 3.4 (Markov Chain Tree Formula). *Let C be an irreducible and aperiodic Markov chain with transition matrix $P_{n\times n}$. Consider a* directed spanning tree *of the Markov chain, which is a spanning tree that satisfies the following conditions:*

- *It is directed toward a designated root.*
- *Every directed edge corresponds to a valid transition in the Markov chain.*

The weight *of a directed spanning tree Tr, denoted by $\rho(Tr)$, is defined as the product of the transition probabilities of its edges:*

$$\rho(Tr) = \prod_{\overrightarrow{(i,j)} \in E(Tr)} P[i, j].$$

Let $\mathcal{T}_i$ be the set of all directed spanning trees rooted at i. Then for any i, j we have that

$$\frac{\pi_i}{\pi_j} = \frac{\sum_{Tr\in\mathcal{T}_i} \rho(Tr)}{\sum_{Tr\in\mathcal{T}_j} \rho(Tr)}.$$

In Theorems 3.5 and 3.6 we use Theorem 3.4 to provide a closed form solution for influence weight vectors under different conditions on T.

Theorem 3.5. *Let $T_{n\times n}$ be a strongly influenced, unbiased, and diagonally uniform transition matrix. Then for every $i \in N$ we have, $w_i = \frac{d_i}{n d_{avg}}$.*

Proof. Denote by $\mathcal{T}_i$ the set of all directed spanning trees rooted at i. Let Tr be an undirected spanning tree of G_T. Note that you can take any vertex i as the root and direct the edges of Tr to obtain a directed spanning tree Tr_i of $\mathcal{T}_i$. Therefore, for two vertices i, j we have

$$\frac{\rho(Tr_i)}{\rho(Tr_j)} = \frac{\prod_{k\in N:k\neq i} \frac{1-T[k,k]}{d_k}}{\prod_{k\in N:k\neq j} \frac{1-T[k,k]}{d_k}} = \frac{1 - T[j, j]}{1 - T[i, i]} \frac{d_i}{d_j} = \frac{d_i}{d_j},$$

where we used the fact that T is unbiased and diagonally uniform. Let $\mathcal{T}$ be the set of all spanning trees of G_T. By a similar argument we get

$$\sum_{Tr_i\in\mathcal{T}_i} \rho(Tr_i) = \sum_{Tr\in\mathcal{T}} \rho(Tr_i) = \frac{d_i}{d_j} \sum_{Tr\in\mathcal{T}} \rho(Tr_j) = \sum_{Tr_j\in\mathcal{T}_j} \rho(Tr_j),$$

and hence $\frac{w_i}{w_j} = \frac{\sum_{Tr_i\in\mathcal{T}_i} \rho(Tr)}{\sum_{Tr_j\in\mathcal{T}_j} \rho(Tr)} = \frac{d_i}{d_j}$.
The result follows by noting that $\sum_{i\in N} w_i = 1$. $\square$

Proof of Theorem 6.1. Let $t \in \mathbb{N}$, $a, b \in A$ be any two alternatives, and $i \in N$ be any voter. We first show that if

$$\left|\frac{\gamma_i}{n}\left[(\mathrm{EW}^*[a] - \mathrm{EW}^*[b]) - (\mathrm{EW}_t[i, a] - \mathrm{EW}_t[i, b])\right]\right| \quad (1)$$

$$\leq \left|V^*[i, a] - V^*[i, b]\right|, \quad (2)$$

then $a \succ_i^t b \Rightarrow a \succ_i b$; Let $X = z_{\gamma_i}(U[i,a] - U[i,b])$, $Y = \frac{1}{n}(\mathrm{EW}^*[a] - \mathrm{EW}^*[b])$, and $Z = \mathrm{EW}_t[i,a] - \mathrm{EW}_t[i,b]$, then Equation (1) becomes $|Y - Z| \leq |Y + X|$. Assume the contrary: Let Equation (1) hold and, without loss of generality, $X + Z = V_t[i,a] - V_t[i,b] \geq 0$ while $X + Y = V^*[i,a] - V^*[i,b] < 0$. This means that $Z > Y$. By Equation (1) and $A + B < 0$ we get $-X - Y > Z - Y$ and hence $X + Z < 0$ which is a contradiction.

Since $\lim_{t\to\infty} \mathrm{EW}_t = \mathrm{EW}^*$, then for any voter i and alternative a, $\lim_{t\to\infty} \mathrm{EW}_t[i,a] = \mathrm{EW}^*[i,a]$. Hence for any $\epsilon > 0$, there exists a timestep $t_i(\epsilon)$ such that if $t \geq t_i(V^*)$, then $|\mathrm{EW}_t[i,a] - \mathrm{EW}^*[i,a]| < \epsilon$. Set

$$\epsilon_{\mathrm{mix}} = \min_{i\in N, a,b\in A: V^*[i,a]\neq V^*[i,b]} \left| z_{\gamma_i}(U[i,a] - U[i,b]) \right| \quad (3)$$

$$+ \frac{1}{n}(\mathrm{EW}^*[a] - \mathrm{EW}^*[b]). \quad (4)$$

Note that if $V^*[i,a] = V^*[i,b]$, then no matter how i ranks a, b, the resulting profile σ_i^t would be consistent with V^*. Let $t_{\mathrm{mix}}(D) = \max_{i\in N} t_i(\epsilon_{\mathrm{mix}})$. Then for any $t \geq t_{\mathrm{mix}}(D)$, voter $i \in N$ and alternatives $a, b \in A$, Equation (1) holds and hence $\vec{\boldsymbol{\sigma}}_t = \vec{\boldsymbol{\sigma}}$. □

A similar result as to Theorem 3.5 can be derived when the estimation dynamics is strongly influenced and fully unbiased. This is posted in Theorem 3.6.

Theorem 3.6. *Let $T_{n\times n}$ be a strongly influenced and fully unbiased transition matrix. Then for every $i \in N$ we have, $w_i = \frac{d_i+1}{n(d_{avg}+1)}$.*

Theorem 3.5 reveals a key insight: under the conditions given in the statement of Theorem 3.5, each voter's influence weight is directly proportional to their degree in the underlying graph. This leads to a straightforward formula for calculating the estimated welfare of each alternative a:

$$\mathrm{EW}^*(a) = \sum_{i=1}^{n} \frac{d_i}{d_{\mathrm{avg}}} U[i,a].$$

Therefore, when analyzing how the underlying graph G_T shapes the system's overall behavior, it is sufficient to consider the degree sequence rather than the precise edge set E. In other words, the structure of individual connections matters less than the overall distribution of influence across voters. Second, in the special case where G_T is a regular graph, every alternative's estimated welfare $\mathrm{EW}^*(a)$ aligns exactly with its social welfare $\mathrm{SW}(a)$. As a result, the problem simplifies to a classic public-spirited voting scenario discussed by Flanigan et al. [26].

Remark. A similar argument can be used to calculate $\vec{\boldsymbol{w}}$ for a strongly influenced and unbiased (but not necessarily diagonally uniform) transition matrix using Theorem 3.4. In particular, one can show that $w_i \propto \frac{d_i}{1-T[i,i]}$. The detailed derivation has been omitted due to space constraints.

Remark. The unbiasedness of the estimation dynamics is required for achieving a closed form expression for the influence weights. Furthermore, it is necessary for achieving meaningful distortion bounds. We establish this in Section 4 by showing that the absence of this assumption results in drastic degradation of the distortion of the system (see Theorem 4.5).

This concludes our discussion on the conditions for convergence of the perception dynamics. In Section 6, we discuss the convergence rate of the process. Prior to that, however, we analyze the performance of the voting mechanism in the steady state under various criteria.

4 Deterministic Guarantees

In this section, we analyze the impact of the underlying graph's structure on the distortion. Our objective is to establish upper and lower bounds on the distortion for some specific classes of the underlying graph. These results show the dependency of distortion on different features of the network. Before we present our main results in this section, let us start with some definitions and lemmas which we will use throughout the rest of this section. Definition 4.1 is our main assumption about the estimation dynamics.

Definition 4.1. We say a transition matrix $T_{n\times n}$ is *well-structured* if it is strongly influenced, unbiased, and diagonally uniform. Furthermore, we say that estimation dynamics D is *well-structured* if its corresponding transition matrix T is *well-structured.*

Next, in Definition 4.2 we introduce a notion that captures the extent to which an estimation dynamics is biased.

Definition 4.2. We say that a transition matrix T is α*-approximately balanced* or α*-balanced* for short, if for any distinct triplet $i, j, j' \in [n]$, we have $T[i,j]/T[i,j'] \leq \alpha$. In other words, the weight that each voter gives to her neighbors are within an α factor of each others. Furthermore, we say that an estimation dynamics D is α-balanced if its corresponding transition matrix is α-balanced.

In Lemma 4.3, we introduce a tool that reduces the problem of finding a lower bound on the distortion to identifying the most unbalanced estimation dynamics.

Lemma 4.3. *Let $T_{n\times n}$ be a strongly influenced transition matrix and $\mathcal{D}(T)$ be the subclass of all estimation dynamics with transition matrix T. Let $\vec{\boldsymbol{w}}$ denote the influence weight vector associated with T. Then for any voting rule f and ps-profile $\vec{\boldsymbol{\gamma}}$,*

$$\mathrm{dist}_{\mathcal{D}(T)}(f, \vec{\boldsymbol{\gamma}}) \geq \max_{i,j\in N} \frac{w_i}{w_j}.$$

Finally, we establish Lemma 4.4 which provides an upper bound on the ratio between the social welfare of any two alternatives. This result is a key tool that helps us in bounding distortion and is analogous to Lemma 4.2 of Flanigan et al. [26], adopted to our social interaction setting.

Lemma 4.4. *Let $D = (U, T)$ be a well-structured estimation dynamics. Then for any two alternatives a, b and ps-profile $\vec{\boldsymbol{\gamma}}$ with $\gamma_{\min} > 0$,*

$$\frac{SW(b)}{SW(a)} \leq \frac{d_{avg}}{d_{\min}} \frac{EW^*(a)}{SW(a)} + \frac{d_{avg}}{d_{\min}} \frac{n}{|\{i : a \succ_i b\}|} z_{\gamma_{\min}},$$

where $\{i : a \succ_i b\}$ is the set of voters who prefer alternative a over alternative b, $\gamma_{\min} = \min_{i\in N} \gamma_i > 0$, and $z_{\gamma_{\min}} = \frac{1-\gamma_{\min}}{\gamma_{\min}}$.

Remark. A robust version of Lemma 4.4 shows that the bound there does not worsen drastically when (1) a small fraction of voters are not public-spirited ($\gamma_i = 0$), and (2) voter weights are unequal but within a constant factor of each other (i.e., $\beta = \max_{i,j\in N} \frac{1-T[i,i]}{1-T[j,j]}$

is bounded). The detailed derivation has been omitted due to space constraints.

We now prove the main results of this section which revolve around well-structured estimation dynamics defined in Definition 4.1. Theorem 4.5, shows that unbiased estimation dynamics is indeed necessary to achieve meaningful bounds on distortion.

Theorem 4.5. *Let $\mathcal{G}_k$ for some $k \geq 3$, denote the set of all k-regular connected graphs of size n. Let $\mathcal{D}(k,\alpha)$ for some $\alpha \geq 1$ denote the subclass of all strongly influenced, diagonally uniform, α-balanced estimation dynamics $D = (U,T)$ with underlying graph $G_T \in \mathcal{G}_k$. Then for any voting rule f and $\vec{\gamma}$, $\mathrm{dist}_{\mathcal{D}(k,\alpha)}(f,\vec{\gamma}) \geq \alpha^{\Theta(\log n)}$.*

Remark. The strategy to prove Theorem 4.5 consists of two parts; Given $\mathcal{G}_k$, we find a graph with high diameter, and then we prove the desired lower bound on this graph. The second part can be adapted separately to provide lower bounds on the distortion of any graph G with strongly influenced, diagonally uniform and α-balanced estimation dynamics. In particular, when G is a path, one can establish an exponential (in n) lower bound on the distortion.

It is noteworthy that in many real-world scenarios where the number of voters n is large (e.g., $n \approx 10^7$), $\alpha^{\Theta(\log n)}$ can grow substantially. This rapid growth hampers our ability to establish meaningful guarantees on distortion. This is in sharp contrast to the case where voters behave unbiasedly; As discussed below Theorem 3.5, if the estimation dynamics $D = (U,T)$ is unbiased, then by Theorem 3.5, regularity of G_T would necessarily lead to a balanced influence weight vector $\vec{w} = \frac{1}{n}\vec{1}$ and reduce the problem to public-spirited voting. In this case distortion can be bounded from above for many important rules (see Corollary 4.7 and Flanigan et al. [26] for more details). This indicates that for a reasonable voting rule like Plurality, the distortion $\mathrm{dist}_{\mathcal{D}(k,\alpha)}$ does not depend on n. This highlights the importance of unbiased aggregation of opinions in the estimation dynamics and its role in democratic deliberation. This is the main takeaway of this work for the voters.

Now that we established the necessity of unbiasedness, we focus on bounding the distortion of two important voting rules: Plurality and Copeland under this assumption. We begin by applying Lemma 4.4 to derive the following theorem, which establishes an upper bound on the distortion of Plurality and Copeland rules.

Theorem 4.6. *For $d_{\min} \leq d_{avg} \leq d_{\max}$, let $\vec{d} = (d_{\min}, d_{avg}, d_{\max})$ and $\mathcal{D}(\vec{d})$ be the subclass of all well-structured estimation dynamics $D = (U,T)$ with G_T satisfying $d^T_{\min} = d_{\min}, d^T_{avg} = d_{avg}, d^T_{\max} = d_{\max}$. Denote the Plurality and Copeland voting rules as $f_{\mathrm{Plu}}, f_{\mathrm{Cop}}$ respectively. Then for ps-profile $\vec{\gamma}$ with $\gamma_{\min} > 0$,*

(1) $\mathrm{dist}_{\mathcal{D}(\vec{d})}(f_{\mathrm{Plu}}, \vec{\gamma}) \leq \frac{d_{\max}}{d_{\min}} + \frac{d_{avg}}{d_{\min}} m z_{\gamma_{\min}}$.
(2) $\mathrm{dist}_{\mathcal{D}(\vec{d})}(f_{\mathrm{Cop}}, \vec{\gamma}) \leq \left(\frac{d_{\max}}{d_{\min}} + \frac{d_{avg}}{d_{\min}} 2 z_{\gamma_{\min}}\right)^2$.

Proof. Our proofs for both rules are inspired by Flanigan et al. [26] and adopted to our setting. For plurality voting rule, let a denote the alternative chosen by plurality. By definition, a is the top vote of at least $\frac{n}{m}$ voters. Thus for any other alternative b, $|\{i : a \succ_i b\}| \geq \frac{n}{m}$. Let b be the alternative with maximum social welfare (i.e., the optimal one). Note that for any alternative c, $\mathrm{EW}^*[c] \leq \frac{d_{\max}}{d_{\mathrm{avg}}}\mathrm{SW}(c)$, hence Lemma 4.4 the distortion is bounded as follows:

$$\begin{aligned}\mathrm{dist}(f_{\mathrm{Plu}}, G) = \frac{\mathrm{SW}(b)}{\mathrm{SW}(a)} &\leq \frac{d_{\mathrm{avg}}}{d_{\min}}\frac{\mathrm{EW}^*(a)}{\mathrm{SW}(a)} + \frac{d_{\mathrm{avg}}}{d_{\min}}\frac{n}{|\{i : a \succ_i b\}|} z_{\gamma_{\min}} \\ &\leq \frac{d_{\max}}{d_{\min}} + \frac{d_{\mathrm{avg}}}{d_{\min}}\frac{n}{|\{i : a \succ_i b\}|} z_{\gamma_{\min}} \\ &\leq \frac{d_{\max}}{d_{\min}} + \frac{d_{\mathrm{avg}}}{d_{\min}} m z_{\gamma_{\min}}.\end{aligned}$$

This proves the bound for the Plurality rule.

For the Copeland voting rule, let a and b be two alternatives. We say *a weakly pairwise-dominates b* if for $|\{i : a \succ_i b\}| \geq \frac{n}{2}$. It is well known that if a is the Copeland winner, then for any other alternative b, either a weakly pairwise-dominates b or that there exists an alternative c such that a weakly pairwise-dominates c and c weakly pairwise-dominates b. In both cases by applying Lemma 4.4 and letting b to be the optimal alternative, we have

$$\mathrm{dist}(f_{\mathrm{Cop}}, G) = \frac{\mathrm{SW}(b)}{\mathrm{SW}(a)} \leq \left(\frac{d_{\max}}{d_{\min}} + \frac{d_{\mathrm{avg}}}{d_{\min}} 2 z_{\gamma_{\min}}\right)^2. \quad \square$$

Note that when the underlying graph is k-regular, the problem reduces to public-spirited voting and the bounds in Theorem 4.6 match the bounds provided in Flanigan et al. [26]. This is formalized in Corollary 4.7.

Corollary 4.7. *Let $\mathcal{D}(k)$ denote the subclass of all well-structured estimation dynamics $D = (U,T)$ with a k-regular underlying graph G_T. Then for every ps-profile $\vec{\gamma}$ with $\gamma_{\min} > 0$,*

(1) $\mathrm{dist}_{\mathcal{D}(k)}(f_{\mathrm{Plu}}, \vec{\gamma}) \leq 1 + m z_{\gamma_{\min}}$.
(2) $\mathrm{dist}_{\mathcal{D}(k)}(f_{\mathrm{Cop}}, \vec{\gamma}) \leq \left(1 + 2 z_{\gamma_{\min}}\right)^2$.

Remark. Similar to Section 4, we can present a robust version of Theorem 4.6, showing that the distortion bounds remain stable even when voters are not fully public-spirited and have non-uniform self-confidence.

Results similar to Theorem 4.8 can be derived when the estimation dynamics is strongly influenced and fully unbiased.

Theorem 4.8. *For $d_{\min} \leq d_{avg} \leq d_{\max}$, let $\vec{d} = (d_{\min}, d_{avg}, d_{\max})$ and $\mathcal{D}'(\vec{d})$ be the subclass of all strongly influenced and fully unbiased estimation dynamics $D = (U,T)$ with G_T satisfying $d^T_{\min} = d_{\min}, d^T_{avg} = d_{avg}, d^T_{\max} = d_{\max}$. Then for every ps-profile $\vec{\gamma}$ with $\gamma_{\min} > 0$,*

(1) $\mathrm{dist}_{\mathcal{D}'(\vec{d})}(f_{\mathrm{Plu}}, \vec{\gamma}) \leq \frac{d_{\max}+1}{d_{\min}+1} + \frac{d_{avg}+1}{d_{\min}+1} m z_{\gamma_{\min}}$.
(2) $\mathrm{dist}_{\mathcal{D}'(\vec{d})}(f_{\mathrm{Cop}}, \vec{\gamma}) \leq \left(\frac{d_{\max}+1}{d_{\min}+1} + \frac{d_{avg}+1}{d_{\min}+1} 2 z_{\gamma_{\min}}\right)^2$.

We now focus on proving a lower bound on the distortion. Note that if estimation dynamics is well-structured, an immediate consequence of Lemma 4.3 is a lower bound on every voting rule, established in Theorem 4.9.

Theorem 4.9. *For $d_{\min} \leq d_{avg} \leq d_{\max}$, let $\vec{d} = (d_{\min}, d_{avg}, d_{\max})$ and $\mathcal{D}(\vec{d})$ be the subclass of all well-structured estimation dynamics $D = (U,T)$ with G_T satisfying $d^T_{\min} = d_{\min}, d^T_{avg} = d_{avg}, d^T_{\max} = d_{\max}$. Then for every rule f and ps-profile $\vec{\gamma}$, $\mathrm{dist}_{\mathcal{D}(\vec{d})}(f, \vec{\gamma}) \geq \frac{d_{\max}}{d_{\min}}$.*

Proof. Since the estimation dynamics are well-structured, then by Theorem 3.5, $w_i = \frac{d_i}{n d_{\text{avg}}}$ and hence Lemma 4.3 we get that for every rule f and ps-profile $\vec{\boldsymbol{\gamma}}$, $\text{dist}_{\mathcal{D}(\vec{d})}(f, \vec{\boldsymbol{\gamma}}) \geq \frac{w_{\max}}{w_{\min}} = \frac{d_{\max}}{d_{\min}}$. □

5 Stochastic Guarantees

In Section 4, we characterized distortion as a function of the degree sequence of the underlying graph of the estimation dynamics. We analyzed how distortion depends on G_T and proved that the degree sequence alone suffices to capture this relationship. In this section, we extend our analysis by relaxing the assumption of having deterministic knowledge of G_T. Instead, we treat G_T as the outcome of a known random process. This approach allows us to study scenarios where the exact structure of G_T is unknown, but its random properties—such as degree distribution—are known. This framework is particularly relevant for modeling social networks. We start with the Erdős-Rényi random graph process.

5.1 Erdős–Rényi Model

In the Erdős–Rényi model, introduced by Erdős and Rényi [20], a graph consists of n vertices, and each edge (i, j) is included in the graph independently with probability p. Parameter p is known as the *edge probability*. Formally, let $\text{ER}(n, p)$ denote the discrete probability space consisting of all $2^{\binom{n}{2}}$ possible undirected simple graphs with n fixed and labeled vertices. The probability of a particular graph with k edges is given by $p^k(1-p)^{\binom{n}{2}-k}$.

We start with Definition 5.1, which introduces the key concepts for this section. For a distribution $X \in \Delta([0, 1]^{n \times n})$ over transition matrices, let d_i^T denote the degree of voter i in G_T with $T \sim X$, and let $\text{d}_{\min}^T$, $\text{d}_{\max}^T$, and $\text{d}_{\text{avg}}^T = \frac{1}{n}\sum_{i=1}^n \text{d}_i^T$ denote the minimum, maximum, and average degrees in G_T.

Definition 5.1. A distribution $X \in \Delta([0, 1]^{n \times n})$ is well-structured if $T \sim X$ is almost surely well-structured. Additionally, given a probability distribution Y on labeled graphs of size n, we say X is consistent with Y and write $X \rhd Y$ if $G_T \sim Y$ when $T \sim X$.

Next, we prove Theorem 5.2 which provides an upper bound on the distortion of Plurality and Copeland rules for Erdős–Rényi random graphs.

Theorem 5.2. *Let $ER(n, p)$ denote the Erdős–Rényi probability distribution with parameters n and p, and $X \rhd ER(n, p)$ be a well-structured distribution. Then for every ps-profile $\vec{\boldsymbol{\gamma}}$ with $\gamma_{\min} > 0$*

(1) $\lim_{n\to\infty} \text{dist}_X(f_{\text{Plu}}, \vec{\boldsymbol{\gamma}}) \leq 1 + m z_{\gamma_{\min}}$.
(2) $\lim_{n\to\infty} \text{dist}_X(f_{\text{Cop}}, \vec{\boldsymbol{\gamma}}) \leq \left(1 + 2z_{\gamma_{\min}}\right)^2$.

Theorem 5.2 shows that for large n, the expected distortion of any distribution X consistent with the Erdős–Rényi model approaches that of a deterministic, regular G_T (i.e., $\text{EW}[a] = \text{SW}(a)$). This implies that even with a stochastic deliberation graph, the system's expected performance is comparable to simpler cases, such as a regular graph or public-spirited voting [26].

5.2 Other Stochastic Models

It is important to note that the distortion depends heavily on the distribution of the underlying graph G_T. To illustrate this, let Y be any probability distribution on labeled graphs of size n and let $X \rhd Y$ be any well-structured distribution consistent with Y. By Theorem 4.9 we know that for any voting rule f and ps-profile $\vec{\gamma}$,

$$\text{dist}_X(f, \vec{\gamma}) = \mathbb{E}_{T\sim X}\left[\text{dist}_{\mathcal{D}(T)}(f, \vec{\gamma})\right] \geq \mathbb{E}_{T\sim X}\left[\frac{\text{d}_{\max}^T}{\text{d}_{\min}^T}\right] = \mathbb{E}_{G\sim Y}\left[\frac{\text{d}_{\max}}{\text{d}_{\min}}\right].$$

Let $E_Y = \mathbb{E}_{G\sim Y}\left[\frac{\text{d}_{\max}}{\text{d}_{\min}}\right]$. The key to Theorem 5.2 is that $E_Y \to 1$ as $n \to \infty$, so any constant upper bound on distortion requires E_Y to be independent of n. This does not hold for several important real-world network models. For instance, in the Preferential Attachment model [8], Móri [34] shows that E_Y grows as n^c for some $c \in (0, 1)$ as $n \to \infty$. More generally, a well-known property of complex networks is that they exhibit *Power-law* degree distribution. A distribution P on natural numbers is power-law if $P(k) \propto k^{-\eta}$ for some $\eta > 0$. Because of the heavy-tail property of power-law distributions, one cannot establish a similar result as Theorem 5.2 for them. One notable property of $\text{dist}_X(f, \vec{\boldsymbol{\gamma}})$ is that it does not assume any relationship between the sampled transition matrix T and the utility matrix U; the worst-case utility is chosen *after* observing the graph, rendering the measure rather pessimistic. Alternatively, assuming little or no dependency between the graph and utility matrix, we can define a more realistic measure called *Decoupled Distortion*, as follows.

Definition 5.3. For distribution $X \in \Delta([0, 1]^{n \times n})$, and given a voting rule f, and ps-profile $\vec{\boldsymbol{\gamma}}$, we define

$$\text{dist}_X^{dc}(f, \vec{\gamma}) = \sup_{U \in \mathbb{R}_{\geq 0}^{n\times m}} \mathbb{E}_{T\sim X}\left[\sup_{\sigma \in \Sigma : \vec{\sigma} \rhd (\vec{\gamma}, (U,T))} \frac{\max_{a\in A} \text{SW}(U, a)}{\text{SW}(U, f(\sigma))}\right].$$

We term this *decoupled distortion* because it separates the utility matrix from the transition matrix. To remove all dependencies, we assume the transition matrix is anonymous, achieved by a random one-to-one matching of voters to graph vertices—equivalently, uniformly shuffling the rows of the utility matrix.

Note that for any distribution X, rule f, and ps-profile $\vec{\gamma}$ we have $\text{dist}_X^{dc}(f, \vec{\gamma}) \leq \text{dist}_X(f, \vec{\gamma})$. Theorem 5.4 analyzes the distortion under this assumption.

Theorem 5.4. *Let Y be a probability distribution on labeled graphs of size n and $X \rhd Y$ be a well-structured distribution on transition matrices that is consistent with Y. Then for any ps-profile $\vec{\boldsymbol{\gamma}}$ with $\gamma_{\min} > 0$ we have*

$$\text{dist}_X^{dc}(f_{\text{Plu}}, \vec{\boldsymbol{\gamma}}) \leq \frac{m}{\gamma_{\min}} \mathbb{E}_{G\sim Y}\left[\frac{\text{d}_{avg}}{\text{d}_{\min}}\right].$$

Theorem 5.4 shows that if the graph structure and utility matrix are decoupled, bounded distortion is achieved when the expected ratio of average to minimum degree is constant. This holds for most real-world complex network models; for instance, the ratio is 2 in the preferential attachment model [8] and a small constant in power-law graphs with $\eta > 2$. Indeed, bounded first moments are a defining feature of such graphs [7].

6 Convergence Rate and Practical Infinity

Up to this point, our analysis has assumed that voters' welfare estimates have reached a steady state. This assumption is problematic for two reasons: (1) EW_t may approach but never equal EW^*, and

(2) even if a close approximation is possible, the convergence may be too slow to be realistic when time is limited. In this section, We address these concerns by showing that 1. For a given estimation dynamics and ps-profile, there exists a timestep after which the preference profile generated based on the estimated welfare matches the preference profile generated by V^*. This indicates that our analysis of distortion based on the final estimation matrix will eventually be true and 2. For many important cases, this timestep has logarithmic dependence on problem parameters.

Formally, let the PS-value of voter $i \in N$ for alternative $a \in A$ at timestep $t \in \mathbb{N}$ be

$$V_t[i,a] = (1-\gamma_i)U[i,a] + \gamma_i \frac{\mathrm{EW}_t[i,a]}{n},$$

and σ_i^t be the preference of voter i at timestep t. Similar to Section 2, we denote the preference profile at timestep t as $\vec{\sigma}_t$ and assume that voters' preferences are consistent with their PS-value at time t—that is, if voter i ranks a higher than b at timestep t, $(a \succ_i^t b)$ then $V_t[i,a] \geq V_t[i,b]$. We denote this consistency by $\vec{\sigma}_t \rhd V_t$.

Our first result, stated in Theorem 6.1, indicates that any strongly influenced estimation dynamics (see Definition 3.1) will eventually reach a point after which the voters rank alternatives as if they knew the final estimation matrix.

Theorem 6.1. *Let $D = (U,T)$ be any strongly influenced estimation dynamics and $\vec{\gamma}$ be any ps-profile. Then there exists a mixing time for D denoted as $t_{mix}(D)$ such that for every $t \geq t_{mix}(D)$, $\vec{\sigma}_t \rhd V^*$.*

Note that after $t_{\text{mix}}(D)$ the input of the voting rule won't change and hence the distortion achieved would not change either.

Before stating the second result of this section, we establish Definitions 6.2 and 6.3 and lemma 6.4 which provide the concepts and background required for Theorem 6.5.

Definition 6.2. Let P be the transition matrix of an irreducible and aperiodic Markov chain C (see below Definition 3.1 and Lemma 3.2) with state space X and steady state π. P is *reversible* if for any two states $x, y \in X$, $\pi_x P[x,y] = \pi_y P[y,x]$.

The following definition is the key condition for establishing fast-mixing results.

Definition 6.3. Let $D = (U,T)$ be an estimation dynamics. We say T is *self-confident* if for any voter i, $T[i,i] \geq 0.5$. Furthermore, we say D is self-confident if T is self-confident. Note that in Markov chain literature, this condition is known as laziness—That is, a transition matrix P of a Markov chain C with state space X is called *lazy* if for any $x \in X$, $P[x,x] \geq 0.5$.

Finally, Lemma 6.4 provides the necessary linear-algebraic tools for 6.5. Note that for the rest of this section, for a vector $\vec{v}$ we use the notation $\vec{v}_i$ to denote its i-th component.

Lemma 6.4. *Let P be a reversible transition matrix of an irreducible and aperiodic Markov chain with steady state π and state space $[n]$. Then*

(1) *P has n eigenvalues λ_i, $1 = \lambda_1 \geq \lambda_2 \geq \ldots \geq \lambda_n \geq -1$. Additionally, define $\lambda_* = \max_{2 \leq i \leq n} |\lambda_i|$, and if P is lazy, then $\lambda_* = \lambda_2$.*
(2) *Let $v^{(i)}$ denote the eigenvector corresponding to λ_i. The eigenvectors can be chosen such that they form an orthonormal basis of $\mathbb{R}$ with respect to $\langle\cdot,\cdot\rangle_\pi$:*
$$\langle v^{(i)}, v^{(j)} \rangle_\pi = \sum_{k \in [n]} v_k^{(i)} v_k^{(j)} \pi_k = \begin{cases} 1 & i = j \\ 0 & i \neq j \end{cases}.$$
(3) *For any $t \in \mathbb{N}$ and $i, j \in [n]$, $\frac{P^t[i,j]}{\pi_j} - 1 = \sum_{2 \leq k \leq n} \lambda_k^t v_i^{(k)} v_j^{(k)}$.*

Our next result, is concerned with the value of $t_{\text{mix}}(D)$ for a given well-structuredestimation dynamics.

Theorem 6.5. *Let $D = (U,T)$ be a well-structuredestimation dynamics and $t_{mix}(D)$ be its mixing time as in Theorem 6.1. then*

$$t_{mix}(D) \leq \frac{\ln\left(\sqrt{n\frac{d^T_{avg}}{d^T_{\min}}}\kappa\left(U,\vec{\gamma}\right)\right)}{\ln\left(\frac{1}{\lambda^*}\right)},$$

where

$$\kappa\left(U,\vec{\gamma}\right) = \frac{2\max_{a\in A}\sqrt{\sum_{i\in N} U[i,a]^2}}{\min_{i\in N, a,b\in A, V^*[i,a]\neq V^*[i,b]} |V^*[i,a] - V^*[i,b]|},$$

and λ^* is the second largest absolute value of the eigenvalues of T.

Theorem 6.5 provides a strong and general result that can used to calculate the mixing time of any deterministic graph by simply finding its λ^*. For the stochastic models we are interested in, Corollary 6.6 posits that they have fast mixing

Corollary 6.6. *Let $D = (U,T)$ be a well-structured and self-confident estimation dynamics. Then there exists a constant $c < 1$ such that $\lambda^* < c$, and hence the dynamics has*

$$t_{mix}(D) \in O\left(\ln\left(\sqrt{n\frac{d^T_{avg}}{d^T_{\min}}}\kappa\left(U,\vec{\gamma}\right)\right)\right)$$

if:

(1) *$G_T \sim ER(n,p)$ for $p > \frac{r\log(n)}{n}$ for some $r > 1$.*
(2) *$G_T \sim PA(d,\delta)$ for some $d \geq 2$ and $\delta > -1$.*

The proof of Corollary 6.6 can be found in Durrett [19].

7 Future Works

We introduce social deliberation with public-spirited voters and analyzed distortion for Plurality and Copeland rules. We prove that social deliberation shows a potential (at least from this POV) to replace centralized deliberations. This opens the door to study the other aspects of this approach. Other directions that can advance our understanding of distortion in this model include: (i) evaluating other voting rules or designing rules tailored to this setting that match our lower bounds; (ii) extending to richer decision problems such as participatory budgeting [9]; (iii) modeling mixed motives and cooperative AI via a cooperativeness parameter γ and partial-information learning dynamics; (iv) studying strategic manipulation of shared estimates and votes and its impact on distortion; and (v) integrating democratic design choices—e.g., sortition-based subgraphs, role distinctions (activists vs. regular voters), and budgeted edge additions—to optimize network structure (independent of utilities) for minimal distortion.

References

[1] Samuel Abrams, Torben Iversen, and David Soskice. 2011. Informal social networks and rational voting. *British Journal of Political Science* 41, 2 (2011), 229–257.

[2] Bruce Ackerman and James S Fishkin. 2008. *Deliberation day*. Yale University Press.

[3] Shiri Alouf-Heffetz, Ben Armstrong, Kate Larson, and Nimrod Talmon. 2022. How Should We Vote? A Comparison of Voting Systems within Social Networks.. In *IJCAI*. 31–38.

[4] Elliot Anshelevich, Onkar Bhardwaj, Edith Elkind, John Postl, and Piotr Skowron. 2018. Approximating optimal social choice under metric preferences. *Artificial Intelligence* 264 (2018), 27–51.

[5] Elliot Anshelevich, Aris Filos-Ratsikas, Nisarg Shah, and Alexandros A Voudouris. 2021. Distortion in Social Choice Problems: The First 15 Years and Beyond. In *30th International Joint Conference on Artificial Intelligence*. International Joint Conferences on Artificial Intelligence Organization, 4294–4301.

[6] Amirreza Bagheridelouee, Marzie Nilipour, Masoud Seddighin, and Maziar Shamsipour. 2024. Metric Distortion Under Public-Spirited Voting. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems*. 2144–2146.

[7] Albert-László Barabási. 2013. Network science. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 371, 1987 (2013), 20120375.

[8] Albert-László Barabási and Réka Albert. 1999. Emergence of scaling in random networks. *science* 286, 5439 (1999), 509–512.

[9] Mark Bedaywi, Bailey Flanigan, Mohamad Latifian, and Nisarg Shah. 2025. Public Spirited Participatory Budgeting. In *Proc. of the 39th AAAI Conference on Artificial Intelligence (AAAI)*.

[10] Paul Burstein. 1976. Social networks and voting: Some Israeli data. *Social Forces* 54, 4 (1976), 833–847.

[11] Ioannis Caragiannis, Evi Micha, and Jannik Peters. 2024. Can a Few Decide for Many? The Metric Distortion of Sortition. *arXiv preprint arXiv:2406.02400* (2024).

[12] Ioannis Caragiannis and Ariel D Procaccia. 2011. Voting almost maximizes social welfare despite limited communication. *Artificial Intelligence* 175, 9-10 (2011), 1655–1671.

[13] Lyn Carson and Brian Martin. 1999. Random selection in politics. (1999).

[14] Vincent Conitzer. 2012. Should social network structure be taken into account in elections? *Mathematical Social Sciences* 64, 1 (2012), 100–102.

[15] Nicole Curato, John S Dryzek, Selen A Ercan, Carolyn M Hendriks, and Simon Niemeyer. 2017. Twelve key findings in deliberative democracy research. *Daedalus* 146, 3 (2017), 28–38.

[16] Aldous David and J Fill. 1995. Reversible Markov——– chains and random walks on graphs. *University of California, Berkeley* (1995).

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

[18] Gil Delannoi and Oliver Dowlen. 2016. *Sortition: Theory and Practice*. Vol. 3. Andrews UK Limited.

[19] Rick Durrett. 2010. *Random graph dynamics*. Vol. 20. Cambridge university press.

[20] Paul Erdős and A. Rényi. 1959. On random graphs. *Publicationes Mathematicae (Debrecen)* 6 (1959), 290.

[21] P. Erdös and A. Rényi. 2006. *On the evolution of random graphs*. Princeton University Press, Princeton, 38–82.

[22] Brandon Fain, Ashish Goel, Kamesh Munagala, and Sukolsak Sakshuwong. 2017. Sequential deliberation for social choice. In *Web and Internet Economics: 13th International Conference, WINE 2017, Bangalore, India, December 17–20, 2017, Proceedings 13*. Springer, 177–190.

[23] James S Fishkin. 1991. *Democracy and deliberation: New directions for democratic reform*. Yale University Press.

[24] James S. Fishkin. 2009. *When the People Speak: Deliberative Democracy and Public Consultation*. Oxford University Press.

[25] Bailey Flanigan, Paul Gölz, Anupam Gupta, Brett Hennig, and Ariel D Procaccia. 2021. Fair algorithms for selecting citizens' assemblies. *Nature* 596, 7873 (2021), 548–552.

[26] Bailey Flanigan, Ariel D Procaccia, and Sven Wang. 2023. Distortion Under Public-Spirited Voting. In *Proceedings of the 24th ACM Conference on Economics and Computation*. 700–700.

[27] Vasilis Gkatzelis, Daniel Halpern, and Nisarg Shah. 2020. Resolving the optimal metric distortion conjecture. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 1427–1438.

[28] Vasilis Gkatzelis, Mohamad Latifian, and Nisarg Shah. 2023. Best of both distortion worlds. In *Proceedings of the 24th ACM Conference on Economics and Computation*. 738–758.

[29] Ashish Goel, Mohak Goyal, and Kamesh Munagala. 2025. Metric Distortion of Small-group Deliberation. *arXiv preprint arXiv:2502.01380* (2025).

[30] Christian List. 2018. Democratic deliberation and social choice. *The Oxford handbook of deliberative democracy* (2018), 463–489.

[31] Panfeng Liu, Longjie Li, Shiyu Fang, and Yukai Yao. 2021. Identifying influential nodes in social networks: A voting approach. *Chaos, Solitons & Fractals* 152 (2021), 111309.

[32] Jane Mansbridge. 1999. Everyday Talk in the Deliberative System. In *Deliberative Politics: Essays on Democracy and Disagreement*. Oxford University Press.

[33] Kanav Mehra, Nanda Kishore Sreenivas, and Kate Larson. 2023. Deliberation and voting in approval-based multi-winner elections. In *International Conference on Autonomous Agents and Multiagent Systems*. Springer, 77–93.

[34] Tamás F Móri. 2005. The maximum degree of the Barabási–Albert random tree. *Combinatorics, Probability and Computing* 14, 3 (2005), 339–348.

[35] Christoph Niessen and Min Reuchamps. 2019. Designing a Permanent Deliberative Citizens' Assembly: The Ostbelgien Modell in Belgium. (2019).

[36] James R Norris. 1998. *Markov chains*. Number 2. Cambridge university press.

[37] Ariel D Procaccia and Jeffrey S Rosenschein. 2006. The distortion of cardinal preferences in voting. In *International Workshop on Cooperative Information Agents*. Springer, 317–331.

[38] Carl A Sheingold. 1973. Social networks and voting: the resurrection of a research agenda. *American Sociological Review* (1973), 712–720.