Strategic Behavior is Bliss: Iterative Voting Improves Social Welfare

Anonymous Author(s)
Affiliation
Address
email

Abstract

Recent work in iterative voting has defined the difference in social welfare between the truthful winner and worst-case equilibrium winner, due to repeated strategic manipulations, known as the additive dynamic price of anarchy (ADPoA) [Brânzei et al., 2013]. While all iterative plurality winners have been shown to differ from truth by at most one initial vote, it is less understood how agents' welfare changes in equilibrium. To this end, we differentiate agents' utility from their iteration mechanism and determine iterative plurality's ADPoA in the worst- and average-case. We first negatively demonstrate that the worst-case ADPoA is linear in the number of agents. In expectation, rather, equilibrium winners have a constant order welfare advantage over the truthful winner. Our positive results illustrate the prospect for social welfare to increase due to strategic manipulation.

1 Introduction

Voting is one of the most popular methods for a group of agents to make a collective decision based on their preferences. Whether a decision is for a high-stake presidential election or a routine luncheon, agents submit their preferences and a voting rule is applied to select a winning alternative.

One critical flaw of voting is its susceptibility to strategic manipulations. That is, agents may have an incentive to misreport their preferences to obtain a more favorable outcome. Unfortunately, manipulation is inevitable under any non-dictatorial single-round voting systems when there are three or more alternatives, as recognized by the celebrated Gibbard-Satterthwaite theorem [Gibbard, 1973, Satterthwaite, 1975]. Consequently, much research has sought to prevent manipulation in voting, especially by high computational barriers [Bartholdi et al., 1989, Faliszewski et al., 2010, Faliszewski and Procaccia, 2010]; see [Conitzer and Walsh, 2016] for a recent survey of the field.

While there is a large body of literature on manipulation of single-round voting systems, sequential and iterative voting procedures are less understood. Indeed, sequential and iterative voting happens in a variety of applications, such as Doodle polls or presidential elections, where agents decide their votes after previewing others' responses [Meir et al., 2010, Desmedt and Elkind, 2010, Xia and Conitzer, 2010, Reijngoud and Endriss, 2012, Zou et al., 2015]. Our key question is:

What is the effect of strategic behavior in sequential and iterative voting?

A series of work initiated by Meir et al. [2010] characterizes dynamics and equilibria of *iterative voting*, where agents sequentially and myopically update their reported preferences (votes) based on other agents' reports [Reyhani and Wilson, 2012, Lev and Rosenschein, 2012, Brânzei et al., 2013, Grandi et al., 2013, Obraztsova et al., 2013, Meir et al., 2014, Rabinovich et al., 2015, Obraztsova et al., 2015, Endriss et al., 2016, Meir, 2016, Tsang and Larson, 2016, Koolyk et al., 2017]. While the convergence of iterative voting has been investigated for many commonly studied voting rules,

Submitted to 35th Conference on Neural Information Processing Systems (NeurIPS 2021). Do not distribute.

the effect of strategic behavior, in terms of aggregate social welfare, remains largely unclear. A notable exception is the work by Brânzei et al. [2013], who introduced and characterized the *additive dynamic price of anarchy* (ADPoA) of iterative voting with respect to the plurality, veto, and Borda social choice functions. The (additive) DPoA measures the social welfare (difference) ratio between the truthful winner and the iterative plurality (equilibrium) winner, when an adversary controls both (1) the order in which agents make their strategic manipulations and (2) agents' truthful preferences altogether, in order to minimize social welfare. In particular, Brânzei et al. [2013] proved that under iterative plurality voting, the number of agents whose top preference is an equilibrium winner is zero or one less than that of the truthful plurality winner. Therefore, strategic behavior does not have a significant negative impact on the social welfare measured by the sum plurality score of the winner. Nevertheless, it is unclear whether this observation holds for other notions of social welfare.

1.1 Our Contributions

We address the key question discussed above in the iterative voting framework, proposed by Meir et al. [2010], by characterizing Brânzei et al. [2013]'s ADPoA under plurality dynamics and arbitrary *rank-based utility functions*. Given $m \geq 3$ alternatives, a ranked-based utility function is characterized by a utility vector $\vec{u} = (u_1, \ldots, u_m)$ with $u_1 \geq u_2 \geq \cdots \geq u_m$ and $u_1 > u_m \geq 0$, such that every agent receives the same utility u_i if their i-th ranked alternative wins, although different agents may rank different alternatives i-th.

Our first main result (Theorem 1) states that, unfortunately, for any fixed $m \geq 3$ and utility vector $\vec{u}$, the ADPoA is $\Theta(n)$, where n is the number of agents. Therefore, the positive result by Brânzei et al. [2013] does not extend to utility functions beyond plurality utility.

To overcome this negative worst-case result, we introduce the notion of *expected additive dynamic price of anarchy* (EADPoA), which assumes that agents' true preferences are generated from a probability distribution. Our second main result (Theorem 2) is positive and surprises us: for any fixed $m \geq 3$ and utility vector $\vec{u}$, the EADPoA is $-\Omega(1)$ under the i.i.d. uniform distribution, known as *Impartial Culture (IC)* in social choice. In particular, our result suggests that *strategic behavior is bliss* because iterative voting helps agents choose an alternative with higher expected social welfare, regardless of the order of agents' strategic manipulations.

We provide an experimental justification of our second main result in Appendix B.

1.2 Related Work and Discussions

Implications and limitations of the main results. Our main result (Theorem 2) compares the social welfare of the worst-case equilibrium winner to that of the truthful plurality winner, instead of the maximum social welfare. It is possible that social welfare can be improved over the truthful profile when a different iterative voting rule is used – measurable via the Price of Stability [Tsang and Larson, 2016]. Nevertheless, we believe that iterative plurality voting is worth studying because of its simplicity and popularity in practice, and it preserves privacy so that agents do not need to reveal their utility functions to the mechanism's center.

Our main result also assumes rank-based utility functions, which naturally extends [Brânzei et al., 2013]. Even though our main result only works for IC, we are not aware of previous theoretical work on the expected performance of iterative voting under any distribution. We do recognize the empirical results presented by Grandi et al. [2013] that seem to show an increase in Borda score due to plurality dynamics, when agent preferences are independent (IC) or correlated via a Polya-Eggenberger urn distribution [Berg, 1985]. However, our results provide completeness and explanatory power missing in [Grandi et al., 2013]. Generalizing this study to other dynamics, utility functions, and families of distributions is an interesting and important direction for future work.

A second future direction is determining whether the cost of strategic behavior is equitably divided among agents or concentrated among subsets of agents with certain rankings. If the latter, then minority populations with similar preferences might bear the brunt of the iterative voting process, even while the electorate as a whole improves. This disadvantage for unweighted additive social welfare functions could be fixed by incorporating quasi-linear utilities for agents, for example.

Techniques. To compute the EADPoA, we partition the (randomly generated) profiles according to their *potential winners* – the alternatives that can be made to win by incrementing their plurality

scores by one. We then show that conditioned on the number of potential winners being 2, iterative voting will choose the alternative who beats the other alternative in their head-to-head competition (Lemma 1). This type of "self selection" improves the social welfare over truthful plurality winner by $\Theta(1)$ in expectation (Lemma 2). We further show that when there are three or more potential winners, the expected welfare loss is no more than $o(1)$ (Lemmas 3 and 4). Overall, the social welfare is improved in expectation.

Sequential and iterative voting. Since its inception in 2010, many researchers have studied the convergence properties of iterative voting under differing assumptions and iteration protocols. Meir et al. [2010] first established the convergence of iterative plurality with deterministic tie-breaking from any initial preference profile or with randomized tie-breaking from the truthful profile. However, this result appears quite sensitive to its assumptions, since the authors found counter-examples when allowing agents to manipulate simultaneously, use better- instead of best-replies, or weigh their votes unequally. Lev and Rosenschein [2012] and Reyhani and Wilson [2012] independently showed that no other scoring rule besides veto necessarily converges, while Koolyk et al. [2017] demonstrated the same for common non-scoring rules, such as Maximin, Copeland, and Bucklin.

Similarly, Obraztsova et al. [2013, 2015] and Rabinovich et al. [2015] each analyze the conditions for Nash equilibrium for iterative voting rules and their truth-biased or lazy voting counterparts [Thompson et al., 2013]. Rabinovich et al. [2015] conclude that determining whether a given profile is a reachable Nash equilibrium is NP-complete.

Most iterative voting rules require agents have full information about each others' votes in order to compute their best-responses. Reijngoud and Endriss [2012] and Endriss et al. [2016], inspired by Chopra et al. [2004], therefore introduce communication graphs and poll information functions to restrict the amount information each agent receives and better respect voter privacy. They subsequently provide susceptibility, immunity, and convergence results according to different voting rules. Tsang and Larson [2016] use these concepts to simulate iterative plurality on a social network and permit agents to infer their best-responses based on their neighbors' reports. The authors demonstrate how correlating agents' preferences affects the PoA and DPoA of strategic outcomes.

Sequential and non-iterative voting games were also investigated in the literature. Desmedt and Elkind [2010] characterized the subgame perfect Nash equilibrium of a voting game where agents vote sequentially and are allowed to abstain from voting. Xia and Conitzer [2010] characterized the subgame perfect Nash equilibrium of a similar voting game where agents are not allowed to abstain, and proved that the equilibrium winner is highly unfavorable in the worst case, which can be viewed as an ordinal PoA. Our paper focuses on iterative voting setting proposed by Meir et al. [2010] and is therefore different from these works.

Best-response mechanisms. A separate line of research from iterative voting has studied the convergence and acyclicity properties of best-response mechanisms. Monderer and Shapley [1996] first introduced the finite improvement property applied to games with agents that sequentially change their actions. Apt and Simon [2012] and Fabrikant et al. [2010] subsequently characterized better-response dynamics in weakly acyclic games, which encapsulate potential and dominance-solvable games, and demonstrate bounds on finding their Nash equilibrium. The relationship between iterative voting and best-response mechanisms was explored by Meir et al. [2014] and Meir [2016], who fully characterized the acyclicity and local dominance properties of iterative voting rules.

Brânzei et al. [2013] introduced the DPoA, which we expand upon in this work, as a measure of the quality of strategic Nash equilibrium. However, not all iterative social choice functions necessarily converge, but all games with best-response dynamics have cycles or steady-state equilibrium over agents' joint action space. Such games may instead be characterized by the worst-case ratio (or difference) between the game's truthful outcome and the average welfare of a stationary distribution over each cycle – known as sink equilibrium and the price of sinking [Goemans et al., 2005]. This serves as a plausible extension to the DPoA, left for future work.

2 Preliminaries

Basic setting. Let $\mathcal{A} = [m] = \{1, \ldots, m\}$ denote the set of $m \geq 3$ *alternatives* and $n \in \mathbb{N}$ denote the number of agents. We denote by $\mathcal{L}(\mathcal{A})$ the set of all strict linear orders over $\mathcal{A}$ and use

$R_j \in \mathcal{L}(\mathcal{A})$, $j \leq n$ to represent agents' preference *rankings*. Preferences are aggregated into *profile* vectors $P = (R_1, \ldots, R_n)$, and $top(R_j) \in \mathcal{A}$ denotes agent j's top preferred alternative. For any pair of alternatives $a, b \in \mathcal{A}$, we use $P[a \succ b]$ to denote the number of agents that prefer a to b in P.

Integer positional scoring rules. An *(integer) positional scoring rule* $r_{\vec{s}}$ is characterized by an integer scoring vector $\vec{s} = (s_1, \ldots, s_m) \in \mathbb{Z}^m_{\geq 0}$ with $s_1 \geq s_2 \geq \cdots \geq s_m \geq 0$ and $s_1 > s_m$. For example, *plurality* uses the scoring vector $\vec{s}_{plu} = (1, 0, \ldots, 0)$, *veto* uses $(1, \ldots, 1, 0)$, and *Borda* uses $(m-1, m-2, \ldots, 0)$. In this work we focus on the plurality rule $r_{plu} = r_{\vec{s}_{plu}}$ and denote by $s_P(a) = \sum_{R \in P} \mathbb{1}\{top(R) = a\}$ the *plurality score* of a according to P. We use the resolute function $r_{plu}(P) = \arg\max_{a \in \mathcal{A}} s_P(a)$, breaking ties lexicographically to select a single winning alternative, favoring that with the smallest index. Assume $r = r_{plu}$ unless stated otherwise.

Additive rank-based utility and social welfare. We assume that agents have additive utilities characterized by a *rank-based utility vector* $\vec{u} = (u_1, \ldots, u_m) \in \mathbb{R}^m_{\geq 0}$ with $u_1 \geq \ldots \geq u_m \geq 0$ and $u_1 > u_m$. Like the scoring rule, each agent j gets $\vec{u}(R_j, a) = u_i$ utility for the alternative $a \in \mathcal{A}$ ranked i^{th} in R_j. Differing from prior work, however, we do not presume that $\vec{u}$ is the same as the scoring vector $\vec{s}$. We define the additive *social welfare* of a according to P as follows.

$$\text{SW}_{\vec{u}}(P, a) = \sum\nolimits_{j=1}^{n} \vec{u}(R_j, a)$$

Iterative plurality voting. Given agents' truthful preferences P, we consider an iterative process of profiles $(P^t)_{t \geq 0}$ that describe agents' *reported preferences* $(R^t_1, \ldots, R^t_n)$ over time. For each round t, one agent j is chosen by a *scheduler* ϕ to make a myopic *improvement step* $R^t_j \xrightarrow{j} R'_j$ to their report. This step is called a *better-response* if j prefers the new outcome $r(R'_j, R^t_{-j})$ to the prior one $r(R^t_j, R^t_{-j})$, whereas it is a *best-response* (BR) if, additionally, j could not have achieved a more preferred outcome than manipulating to R'_j from P^t. [1]

Following Brânzei et al. [2013], we limit our discussion to strategic manipulations beginning from the truthful profile $P^0 = P$. This guarantees that all improvement steps $R'_j \xrightarrow{j} R''_j$ from profile P' to P'' are best-responses that change the iterative winner: $r(P') \neq top(R'_j) \ \wedge \ r(P'') = top(R''_j)$ [Brânzei et al., 2013]. [2] As a result, any sequence of BR steps converges in $\mathcal{O}(nm)$ rounds [Reyhani and Wilson, 2012]. The profiles $\{P^*\}$ with no further improvement steps are therefore *Nash equilibrium* (NE) with respect to P. Let $\text{EW}(P)$ denote the set of *equilibrium winning* alternatives corresponding with all NE reachable from P via some BR sequence. That is,

$$\text{EW}(P) = \{r(P^*) : \exists \text{ a BR sequence from } P \text{ leading to the NE profile } P^*\}$$

Lastly, we'll define the set of *potential winning* alternatives of any profile P as those who could become a winner if their plurality score were to increment by one, including the current winner. That is, some agent could make these alternatives win by taking a BR step that increases their plurality score, if the agent's ranking permits. Following Rabinovich et al. [2015], we have:

$$\text{PW}(P) = \left\{ a \in \mathcal{A} \ : \begin{cases} s_P(a) = s_P(r(P)) - 1, & a \text{ is ordered before } r(P) \\ s_P(a) = s_P(r(P)), & a \text{ is ordered after } r(P) \end{cases} \right\} \cup \{r(P)\}$$

where the ordering is lexicographical for tie-breaking.

The number of alternatives tied in $|\text{PW}(P)|$ is significant to determining $\text{EW}(P)$. Reyhani and Wilson [2012] demonstrated that $\text{EW}(P) \subseteq \text{PW}(P^0)$ and $\forall t \geq 0$, $\text{PW}(P^{t+1}) \subseteq \text{PW}(P^t)$. As a result, iterative plurality voting acts like a sequential tie-breaking mechanism whose outcome follows from the scheduler ϕ. Once a two-way tie appears in a BR sequence at the profile P^t, with $|\text{PW}(P^t)| = 2$, Lemma 1 can be applied break it and determine which alternative is an equilibrium winner. Furthermore we'll see that the key idea to Theorem 2 is to partition $\mathcal{L}(\mathcal{A})^n$ based on the size of each profile's potential winning set. The following example demonstrates this section's concepts.

[1] Note that ϕ must select a BR step is one exists [Apt and Simon, 2012].

[2] These are characterized as *Type 1* or *direct best replies* in the literature [Meir, 2016]. Conversely, *Type 3* best-responses ($r(P') = top(R'_j) \ \wedge \ r(P'') = top(R''_j)$) do not occur in improvement sequences from the truthful profile. Note that no BR step is of *Type 2* ($r(P') = top(R'_j) \ \wedge \ r(P'') \neq top(R''_j)$).

Example 1. *Let $n = 9$, $m = 3$, and consider the truthful profile P where $R_1 = R_2 = R_3 = [1 \succ 2 \succ 3]$, $R_4 = R_5 = [2 \succ 3 \succ 1]$, $R_6 = [2 \succ 1 \succ 3]$, and $R_7 = R_8 = R_9 = [3 \succ 2 \succ 1]$. We can immediately observe from the score histogram $(s_P(1), s_P(2), s_P(3)) = (3, 3, 3)$ that $r(P) = 1$ and $PW(P) = \{1, 2, 3\}$, representing a three-way tie. Next, there are five BR sequences from P (the following vectors denote agents' reports $(top(R_j^t))_{j \leq n}$; the winner appears in curly brackets):*

1. $(111222333)\{1\} \xrightarrow{4} (111322333)\{3\} \xrightarrow{6} (111321333)\{1\} \xrightarrow{5} (111331333)\{3\}$
2. $(111222333)\{1\} \xrightarrow{5} (111322333)\{3\} \xrightarrow{6} (111321333)\{1\} \xrightarrow{4} (111331333)\{3\}$
3. $(111222333)\{1\} \xrightarrow{7} (111222233)\{2\}$
4. $(111222333)\{1\} \xrightarrow{8} (111222323)\{2\}$
5. $(111222333)\{1\} \xrightarrow{9} (111222332)\{2\}$

We therefore conclude EW$(P) = \{2, 3\}$. Moreover, consider the utility vector $\vec{u} = (u_1, u_2, u_3)$. Then the social welfare for each alternative is $\begin{pmatrix} SW_{\vec{u}}(P,1) \\ SW_{\vec{u}}(P,2) \\ SW_{\vec{u}}(P,3) \end{pmatrix} = \begin{pmatrix} 3u_1 + 1u_2 + 5u_3 \\ 3u_1 + 3u_2 + 3u_3 \\ 3u_1 + 5u_2 + 1u_3 \end{pmatrix}$. □

3 Additive Dynamic PoA under General Utility Vectors

How bad are equilibrium outcomes, given that strategic manipulations inevitably occur by the Gibbard-Satterthwaite theorem [Gibbard, 1973, Satterthwaite, 1975]? Brânzei et al. [2013] sought to answer this question by defining the *additive dynamic price of anarchy (ADPoA)* as the *adversarial loss* – the difference in welfare between the truthful profile P and its worst-case equilibrium profile in EW(P) – according to the worst-case P. For example, if users of a website can regularly log in and update their preferences for some cause, then the ADPoA bounds the welfare loss if a virtual assistant recommends when users should make their changes. Brânzei et al. originally defined the ADPoA for an arbitrary positional scoring rule $r_{\vec{s}}$ and the additive social welfare function respecting $\vec{u} = \vec{s}$. In this case, the ADPoA of plurality was found to be 1, while the (multiplicative) DPoA of veto is $\Omega(m)$ and Borda is $\Omega(n)$ for $m \geq 4$ [Brânzei et al., 2013].

Although these results answer Brânzei et al.'s question and appear optimistic for plurality, they suggest more about the iteration mechanism than agents' collective welfare. For example, an ADPoA for plurality of 1 means that for any truthful profile, the difference in initial plurality scores of any equilibrium winner is zero or one fewer than that of the truthful winner. However, when we extend the utility vector $\vec{u}$ to be arbitrary (subject to the conditions in Section 2), we find in Theorem 1 that the ADPoA is quite poor at $\mathcal{O}(n)$. First we recall Brânzei et al.'s definition of ADPoA using our notation, explicitly stating the adversarial welfare loss D^+ for a particular truthful profile P.

Definition 1 (Additive Dynamic Price of Anarchy (ADPoA) [Brânzei et al., 2013]). *Given a positional scoring rule $r_{\vec{s}}$, a utility vector $\vec{u} = (u_1, \ldots, u_m)$, and an initial truthful profile P, the* adversarial loss starting from P *is defined as*

$$D^+_{r_{\vec{s}},\vec{u}}(P) = SW_{\vec{u}}(P, r_{\vec{s}}(P)) - \min_{a \in EW(P)} SW_{\vec{u}}(P, a)$$

The additive dynamic price of anarchy (ADPoA) *of $r_{\vec{s}}$ and $\vec{u}$ under n agents is defined as*

$$ADPoA(r_{\vec{s}}, \vec{u}, n) = \max_{P \in \mathcal{L}(\mathcal{A})^n} D^+_{r_{\vec{s}},\vec{u}}(P)$$

We will use ADPoA and D^+ to denote ADPoA$(r_{plu}, \vec{u}, n)$ and $\mathrm{D}^+_{r_{plu},\vec{u}}$ when the context is clear. [3] For example, we saw in Example 1 that $r(P) = 1$ and $EW(P) = \{2, 3\}$. Therefore

$$\begin{aligned} D^+(P) &= \max\{\ SW_{\vec{u}}(P,1) - SW_{\vec{u}}(P,2),\ SW_{\vec{u}}(P,1) - SW_{\vec{u}}(P,3)\ \} \\ &= \max\{\ (3u_1 + 1u_2 + 5u_3) - (3u_1 + 3u_2 + 3u_3)\,, \\ &\qquad\qquad (3u_1 + 1u_2 + 5u_3) - (3u_1 + 5u_2 + 1u_3)\ \} \\ &= -2(u_2 - u_3) \\ &< 0 \end{aligned}$$

[3] The superscript '+' denotes an additive measure instead of multiplicative in the classical definition of PoA.

if $u_2 > u_3$. We conclude that the social welfare of both equilibrium winners is at least that of the truthful winner. As we will see in Theorem 2, this example's conclusion generalizes well.

Theorem 1. *Fix $m \geq 3$ and utility vector $\vec{u} = (u_1, \ldots, u_m)$. Then ADPoA$(r_{plu}, \vec{u}, n)$ is $\Theta(n)$. Specifically,*

$$(u_2 - u_m)\left(\frac{n}{m} - 2\right) \leq \textit{ADPoA}(r_{plu}, \vec{u}, n) \leq nu_1$$

Proof. The ADPoA is trivially upper bounded by the maximum social welfare attainable by any truthful profile P. For example, if P is defined with $R_j = (1, 2, \ldots, m)\ \forall j \leq n$, then $\forall \tilde{P} \in \mathcal{L}(\mathcal{A})^n$,

$$D^+(\tilde{P}) = \text{SW}_{\vec{u}}(\tilde{P}, r(\tilde{P})) - \min_{a \in \text{EW}(\tilde{P})} \text{SW}_{\vec{u}}(\tilde{P}, a) \leq \text{SW}_{\vec{u}}(\tilde{P}, r(\tilde{P})) \leq \text{SW}_{\vec{u}}(P, r(P)) = nu_1$$

To lower bound ADPoA, we will construct a profile P with a two-way tie between alternatives $1, 2 \in \mathcal{A}$ such that $D^+(P) = (u_2 - u_m)\left(\frac{n}{m} - 2\right)$. This implies the desired lower bound of

$$\text{ADPoA} = \max_{\tilde{P} \in \mathcal{L}(\mathcal{A})^n} D^+(\tilde{P}) \geq D^+(P) = (u_2 - u_m)\left(\frac{n}{m} - 2\right)$$

Fix $m \geq 3$ and let $n \in \mathbb{N}$ be even. We denote by $k = \arg\min_{\tilde{k} \in [2, m-1]}(u_{\tilde{k}} - u_{\tilde{k}+1})$ the position in $\vec{u}$ with the minimal difference in adjacent coordinates. Let $\alpha = \frac{1}{m}(n+m-2)$ and $\beta = (\alpha-1)(m-2)$, such that $n = 2\alpha + \beta$. We will then construct P as follows, with α agents that prefer 1 first and 2 last, α agents that prefer 2 first and 1 second, $(\frac{\beta}{2} - 1)$ agents that prefer 1 second and 2 last, and $(\frac{\beta}{2} + 1)$ agents that prefer 2 in their ranking's k-th position and 1 in their ranking's $(k+1)$-th position. We can see here that $s_P(1) = s_P(2) = \alpha$, and $\forall c > 2$, $s_P(c) = \alpha - 1$, thus guaranteeing the two-way tie. Therefore $r(P) = 1$ and $P[2 \succ 1] = \alpha + \frac{\beta}{2} + 1 > \alpha + \frac{\beta}{2} - 1 = P[1 \succ 2]$. This implies $\text{EW}(P) = \{2\}$ by the following lemma.

Lemma 1. *Let $m \geq 2$ and $a, b \in \mathcal{A}$ such that a is ordered before b in tie-breaking. Suppose PW$(P) = \{a, b\}$ for some truthful profile P. Then EW$(P) = \{a\}$ if $P[a \succ b] \geq P[b \succ a]$; otherwise EW$(P) = \{b\}$.*

The proof of the lemma can be found in Appendix A.1. As a result,

$$\begin{aligned}
D^+(P) &= \text{SW}_{\vec{u}}(P, a) - \text{SW}_{\vec{u}}(P, b) \\
&= \alpha(u_2 - u_m) + \left(\frac{\beta}{2} - 1\right)(u_2 - u_m) - \left(\frac{\beta}{2} + 1\right)(u_k - u_{k+1}) \\
&\geq (u_2 - u_m)(\alpha - 2) \\
&= (u_2 - u_m)\left(\frac{n - m - 2}{m}\right) \\
&\geq (u_2 - u_m)\left(\frac{n}{m} - 2\right)
\end{aligned}$$

where the first inequality holds because $(u_k - u_{k+1}) \leq (u_2 - u_m)$. □

4 Expected Additive DPoA

In this section we extend Brânzei et al.'s DPoA to account for the average-case adversarial loss for a positional scoring rule $r_{\vec{s}}$, rather than only the worst-case. This *expected additive dynamic price of anarchy* (EADPoA) bounds the welfare loss of strategic manipulation according to more typical distributions of agents' rankings. Here we distribute profiles i.i.d. uniformly over $\mathcal{L}(\mathcal{A})^n$, known as the Impartial Culture distribution $\pi_n = IC^n$.

Definition 2 (Expected Additive DPoA (EADPoA)). *Given a positional scoring rule $r_{\vec{s}}$, a utility vector $\vec{u}$, n agents, and a distribution π_n over $\mathcal{L}(\mathcal{A})^n$ for agents' preferences, the* expected additive dynamic price of anarchy *is defined as follows:*

$$\textit{EADPoA}(r_{\vec{s}}, \vec{u}, n, \pi_n) = \mathbb{E}_{P \sim \pi_n}\left[D^+_{r_{\vec{s}}, \vec{u}}(P)\right]$$

Like before, we will use EADPoA and D^+ to denote $\mathrm{EADPoA}(r_{plu}, \vec{u}, n, \mathrm{IC}^n)$ and $\mathrm{D}^+_{r_{plu},\vec{u}}$ respectively. We will also drop the subscript "$P \sim IC^n$" to simplify notation when the context is clear.

Theorem 2. *For any $m \geq 3$, utility vector $\vec{u} = (u_1, \ldots, u_m)$, and $n \in \mathbb{N}$, we have*

$$\textit{EADPoA}(r_{plu}, \vec{u}, n, \textit{IC}^n) = -\Omega(1)$$

Proof. The key idea is to partition $\mathcal{L}(\mathcal{A})^n$ according to each profile's potential winner set. More precisely, for every $W \subseteq \mathcal{A}$ with $W \neq \emptyset$, we define:

$$\overline{\mathrm{PoA}}(W) = \Pr(\mathrm{PW}(P) = W) \times \mathbb{E}[\mathrm{D}^+(P) \mid \mathrm{PW}(P) = W]$$

By the law of total expectation, then

$$\mathrm{EADPoA} = \mathbb{E}[\mathrm{D}^+(P)] = \sum\nolimits_{\alpha=1}^{m} \sum\nolimits_{W \subseteq \mathcal{A}:|W|=\alpha} \overline{\mathrm{PoA}}(W) \tag{1}$$

where α denotes the number of potential winners in P. It is straightforward to see that when $\alpha = 1$, any profile P with $|\mathrm{PW}(P)| = 1$ is already a NE, which implies $\mathrm{D}^+(P) = 0$. The rest of the proof proceeds as follows. For any $n \in \mathbb{N}$ we will show in Lemma 2 that for $\forall W \subseteq \mathcal{A}$ with $|W| = 2$ $\overline{\mathrm{PoA}}(W) = -\Omega(1)$. We will then demonstrate that $\overline{\mathrm{PoA}}(W) = o(1)$ $\forall W \subseteq \mathcal{A}$ whenever $|W| = 3$ (Lemma 3) and $|W| \geq 4$ (Lemma 4). Finally, these result combine to conclude

$$\mathrm{EADPoA} = \underbrace{0}_{\alpha=1} - \underbrace{\Omega(1)}_{\alpha=2} + \underbrace{o(1)}_{\alpha=3} + \underbrace{o(1)}_{\alpha\geq 4} = -\Omega(1)$$

Lemma 2 ($\alpha = 2$). *Given $m \geq 3$ and a utility vector $\vec{u}$, for any $W \subseteq \mathcal{A}$ with $|W| = 2$ and any $n \in \mathbb{N}$, we have $\overline{\textit{PoA}}(W) = -\Omega(1)$.*

Proof. Without loss of generality let $W = \{1, 2\}$ and suppose $u_2 > u_m$. There are two possible cases of $\mathrm{PW}(P) = \{1, 2\}$: $\mathcal{E}_1 = \mathbb{1}\{s_P(1) = s_P(2)\}$ and $\mathcal{E}_2 = \mathbb{1}\{s_P(1) = s_P(2) - 1\}$. This suggests $\overline{\mathrm{PoA}}(W) = \Pr(\mathcal{E}_1) \times \mathbb{E}[D^+(P) \mid \mathcal{E}_1] + \Pr(\mathcal{E}_2) \times \mathbb{E}[D^+(P) \mid \mathcal{E}_2]$, but we'll focus on the former summand where alternatives 1 and 2 are tied, since the latter's proof is similar.

We believe this proof is challenging due to the dependence in agents' rankings once we condition on profiles that satisfy two-way ties. As a result, standard approximation techniques that assume independence, such as the Berry-Esseen inequality, no longer apply and may also be too coarse to support our claim. Instead, we will further condition agents' rankings ($\vec{Q}$) on their top-ranked alternatives ($\vec{T}$) and whether they satisfy ($1 \succ_j 2$) or ($2 \succ_j 1$) ($\vec{Z}$) using a Bayesian network. This guarantees agents' rankings' conditional independence so we may identify their individual contributions to $D^+(\vec{Q})$.

At a high level, for each pair $\vec{t}$ and $\vec{z}$ that satisfy certain conditions, agents are distinguished by four groups based on whether *top*(R_j) is 1, 2, or neither, and whether ($1 \succ_j 2$) or ($2 \succ_j 1$). It turns out that by $\mathcal{E}_1$, the number of agents in the first two groups are equal and their contributions are inverse, so they cancel out. Given $\mathcal{E}_1$, we will then show by Lemma 1 that $D^+(\vec{Q}) = \vec{u}(\vec{Q}, 1) - \vec{u}(\vec{Q}, 2)$ and with high probability there are $\Theta(n)$ third-party agents (with *top*$(R_j) \neq 1, 2$), more of whom prefer ($2 \succ 1$) than ($1 \succ 2$). This enables our use of approximations such as Stirling's approximation and allows us to conclude that agents' contributions toward alternative 2 outweigh those for 1.

More precisely, for any $j \leq n$, we represent agent j's ranking distribution (i.i.d. uniform over $\mathcal{L}(\mathcal{A})$) by a Bayesian network of three random variables (see Figure 1a):

$T_j \in \mathcal{A}$ represents the top-ranked alternative and follows a uniform distribution

$Z_j \in \{1, 2\}$ indicates ($1 \succ_j 2$) or ($2 \succ_j 1$) conditioned on T_j, and has probability 0.5 if $T_j \neq 1, 2$

Q_j follows the uniform distribution over linear orders that uphold both T_j and Z_j

It is not hard to verify that (unconditional) Q_j follows the uniform distribution over $\mathcal{L}(\mathcal{A})$, which implies that $\vec{Q} = (Q_1, \ldots, Q_n)$ follows the same distribution as P.

Step 1: Identify $\mathcal{E}_1$. Let $\mathcal{T}_2 \subseteq [m]^n$ denote the set of vectors $\vec{t} = (t_1, \ldots, t_n)$ such that alternatives 1 and 2 have the maximum plurality score. $\mathcal{E}_1$ holds for $\vec{Q}$ if and only if $\vec{T}$ takes a value in $\mathcal{T}_2$.

$$\mathcal{T}_2 = \left\{\vec{t} \in [m]^n : \forall 3 \leq i \leq m, |\{j : t_j = 1\}| = |\{j : t_j = 2\}| > |\{j : t_j = i\}|\right\}$$

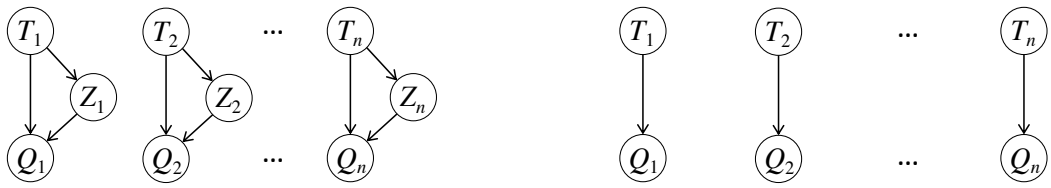

(a) $\alpha = 2$ case (b) $\alpha = 3$ case

Figure 1: Bayesian network representation of P as $\vec{T}$ and $\vec{Q}$

Step 2: Select profiles with non-zero loss. By Lemma 1 given $\vec{t} \in \mathcal{T}_2$, $D^+(\vec{Q})$ is non-zero if and only if $\vec{Q}[2 \succ 1] > \vec{Q}[1 \succ 2]$ – thus $EW(\vec{Q}) = \{2\}$ is unique. Let $\mathrm{Id}_1(\vec{t}) = \{j \leq n : t_j = 1\}$, $\mathrm{Id}_2(\vec{t})\{j \leq n : t_j = 2\}$, and let $\mathrm{Id}_3(\vec{t}) = \{j \leq n : t_j \neq 1, 2\}$ be called *third-party agents*. Since $\mathcal{E}_1$ implies $|\mathrm{Id}_1(\vec{t})| = |\mathrm{Id}_2(\vec{t})|$, there must be more third-party agents with $(2 \succ 1)$ than those with $(1 \succ 2)$. For every $\lceil \frac{|\mathrm{Id}_3(\vec{t})|+1}{2} \rceil \leq k \leq |\mathrm{Id}_3(\vec{t})|$, we then define $\mathcal{Z}_{\vec{t},k} \subseteq \{1,2\}^n$ as follows.

$$\mathcal{Z}_{\vec{t},k} = \{\vec{z} \in \{1,2\}^n : \forall j \in \mathrm{Id}_1(\vec{t}) \cup \mathrm{Id}_2(\vec{t}), z_j = t_j, \text{ and } |\{j \in \mathrm{Id}_3(\vec{t}) : z_j = (2 \succ 1)\}| = k\}$$

By the law of total expectation and noting the independence of $\vec{Q}$'s components,

$$\Pr(\mathcal{E}_1) \times \mathbb{E}[D^+(P) \mid \mathcal{E}_1] = \sum_{\vec{t} \in \mathcal{T}} \sum_{k=\lceil \frac{|\mathrm{Id}_3(\vec{t})|+1}{2} \rceil}^{|\mathrm{Id}_3(\vec{t})|} \sum_{\vec{z} \in \mathcal{Z}_{\vec{t},k}} \Pr(\vec{T} = \vec{t}, \vec{Z} = \vec{z}) \sum_{j=1}^{n} E_{t_j,z_j} \tag{2}$$

where

$$E_{t_j,z_j} = \mathbb{E}_{\vec{Q}_j}[\vec{u}(Q_j, 1) - \vec{u}(Q_j, 2) \mid T_j = t_j, Z_j = z_j]$$

Step 3: Identify contribution per agent. E_{t_j,z_j} only depends on the values of t_j and z_j, but not j. The cases for $(t_j = z_j = 1)$ and $(t_j = z_j = 2)$ negate each other with $E_{1,1} + E_{2,2} = 0$. If $t_j \notin \{1,2\}$ and $z_j = 1$, then $E_{t_j,1} = \eta > 0$ because $u_2 > u_m$. Similarly, it follows that if $t_j \notin \{1,2\}$ and $z_j = 2$, then $E_{t_j,2} = -\eta$. Therefore Equation (2) becomes

$$\sum_{\vec{t} \in \mathcal{T}} \sum_{k=\lceil \frac{|\mathrm{Id}_3(\vec{t})|+1}{2} \rceil}^{|\mathrm{Id}_3(\vec{t})|} \sum_{\vec{z} \in \mathcal{Z}_{\vec{t},k}} \Pr(\vec{T} = \vec{t}, \vec{Z} = \vec{z}) \times (|\mathrm{Id}_3(\vec{t})| - 2k)\eta \tag{3}$$

where we've inserted $\sum_{j=1}^{n} E_{t_j,z_j} = |\mathrm{Id}_1(\vec{t})|E_{1,1} + |\mathrm{Id}_2(\vec{t})|E_{2,2} - k\eta + (|\mathrm{Id}_3(\vec{t})| - k)\eta$.

Step 4: Select on profiles with $\Theta(n)$ third-party agents. Let $\alpha_{\vec{t}} = |\mathrm{Id}_3(\vec{t})|$, although we drop the subscript for ease of notation. α is then equivalent to the sum of n i.i.d. binary random variables, each of which is 1 with probability $\frac{m-2}{m} > \frac{1}{3}$. By Hoeffding's inequality, with exponentially small probability we have $\alpha < \frac{1}{6}n$. Therefore, we can focus on the $\alpha \geq \frac{1}{6}n$ case in (3), which becomes:

$$\begin{aligned}
&\leq e^{-\Omega(n)} + \sum_{\vec{t} \in \mathcal{T}_2 : \alpha \geq \frac{1}{6}n} \sum_{k=\lceil \frac{\alpha+1}{2} \rceil}^{\alpha} \sum_{\vec{z} \in \mathcal{Z}_{\vec{t},k}} \Pr(\vec{T} = \vec{t}, \vec{Z} = \vec{z}) \times (\alpha - 2k)\eta \\
&= e^{-\Omega(n)} - \eta \sum_{\vec{t} \in \mathcal{T}_2 : \alpha \geq \frac{1}{6}n} \left(\frac{1}{2}\right)^{\alpha} \left(\left\lceil \frac{\alpha+1}{2} \right\rceil\right) \binom{\alpha}{\lceil \frac{\alpha+1}{2} \rceil} && (4) \\
&= e^{-\Omega(n)} - \eta \sum_{\vec{t} \in \mathcal{T}_2 : \alpha \geq \frac{1}{6}n} \Pr(\vec{T} = \vec{t}) \times \Theta(1) && (5) \\
&= e^{-\Omega(n)} - \eta \Pr\left(\alpha \geq \frac{1}{6}n\right) \times \Theta(1) = -\Omega(1)
\end{aligned}$$

where Equation (4) follows from Claim 1 (see Appendix A.3) and Equation (5) follows from applying Stirling's approximation (see Appendix A.4). This proves Lemma 2. The full proof can be found in Appendix A.2. □

Lemma 3 ($\alpha = 3$). *Given $m \geq 3$ and a utility vector $\vec{u}$, for any $W \subseteq \mathcal{A}$ with $|W| = 3$ and any $n \in \mathbb{N}$, we have $\overline{PoA}(W) = o(1)$.*

Proof. Without loss of generality, suppose $W = \{1, 2, 3\}$ and $\mathcal{E} = \mathbb{1}\{s_P(1) = s_P(2) = s_P(3)\}$. Like Lemma 2's proof, we split agents' rankings into a Bayesian network by their top preferences ($\vec{T}$) and remaining rankings ($\vec{Q}$), so that $\vec{Q}$ follows the same IC^n distribution as P (see Figure 1b). We want to show that given $\vec{t} \in \mathcal{T}_3$, the set of top preferences representing $\mathcal{E}$, the maximum score difference between any pair of alternatives in W is $o(n)$ with high probability that is close to 1. Let

$$\mathcal{T}_3 = \left\{\vec{t} \in [m]^n : \forall 4 \leq i \leq m, |\{j : t_j = 1\}| = |\{j : t_j = 2\}| = |\{j : t_j = 3\}| > |\{j : t_j = i\}|\right\}$$

Therefore

$$\begin{aligned}
\overline{\text{PoA}}(\{1,2,3\}) &= \Pr\left(\text{PW}(\vec{Q}) = \{1,2,3\}\right) \times \mathbb{E}[\text{D}^+(\vec{Q}) \mid \text{PW}(\vec{Q}) = \{1,2,3\}] \\
&= \sum_{\vec{t} \in \mathcal{T}_3} \Pr(\vec{T} = \vec{t}) \times \mathbb{E}[\text{D}^+(\vec{Q}) \mid \vec{T} = \vec{t}] \\
&= \Pr(\vec{T} \in \mathcal{T}_3) o(n) = o(1) \qquad (6)
\end{aligned}$$

where $\Pr(\vec{T} \in \mathcal{T}_3)$ is equivalent to the probability of three-way ties under plurality w.r.t. IC, which is known to be $\Theta(n^{-1})$ [Gillett, 1977, Xia, 2021].

Given $\vec{t} \in \mathcal{T}_3$ and $i_1, i_2 \in \{1, 2, 3\}$ with $i_1 \neq i_2$, let $D^{t_j}_{i_1,i_2} = \vec{u}(Q_j, i_1) - \vec{u}(Q_j, i_2)$. We assert the $o(n)$ term by arguing that $\forall j \leq n$, $\mathbb{E}[\sum_{j=1}^n D^{t_j}_{i_1,i_2}] = 0$ and $D^{t_j}_{i_1,i_2} \in [u_m - u_1, u_1 - u_m]$. This implies $\Pr(|\sum_{j=1}^n D^{t_j}_{i_1,i_2}| > n^{0.6}) = \exp(-\Theta(n^{0.2}))$ by Hoeffding's inequality, and, by law of total probability, $\Pr\left(\text{D}^+(\vec{Q}) > n^{0.6} \mid \vec{T} = \vec{t}\right) \leq 6\exp(-\Theta(n^{0.2}))$. Equation (6) thus follows.

The full proof for this lemma can be found in Appendix A.5. □

Lemma 4 ($\alpha \geq 4$). *Given $m \geq 4$ and a utility vector $\vec{u}$, for any $W \subseteq \mathcal{A}$ with $|W| \geq 4$ and any $n \in \mathbb{N}$, we have $\overline{PoA}(W) = o(1)$.*

Proof. The proof follows after noticing that:

(1) $\Pr(\text{PW}(P) = W) = \Theta(n^{-1.5})$, following a similar proof using the polyhedron representation as described in the proof of Lemma 3 (see Appendix A.5).

(2) For any profile P, $\text{D}^+(P) = \mathcal{O}(n)$. □

Putting things together. Recall that we assume m to be fixed. Therefore, the total number of subsets of $\mathcal{A}$ is viewed as a constant. Finally, we have:

$$\text{EADPoA} \leq \underbrace{0}_{\alpha=1} - \underbrace{\binom{m}{2}\Omega(1)}_{\alpha=2} + \underbrace{\binom{m}{3}o(1)}_{\alpha=3} + \underbrace{2^m o(1)}_{\alpha\geq 4} = -\Omega(1)$$

This completes the proof of Theorem 2. □

5 Further Work

As discussed in the Introduction, there are many open questions on the (expected) additive dynamic PoA along three dimensions. First, the dynamics could be extended to other voting rules besides plurality [Lev and Rosenschein, 2012, Reyhani and Wilson, 2012, Koolyk et al., 2017], where we would instead determine Goemans et al. [2005]'s price of sinking among voting rules that don't necessarily converge. A second dimension would be to modify the the utility function, for example being normalized [Procaccia and Rosenschein, 2006, Caragiannis and Procaccia, 2011] or in a metric space [Anshelevich et al., 2018]. Thirdly, the choice of agents' ranking distribution could be generalized from IC, for example using smoothed analysis [Xia, 2020].

References

Elliot Anshelevich, Onkar Bhardwaj, Edith Elkind, John Postl, and Piotr Skowron. Approximating Optimal Social Choice under Metric Preferences. *Artificial Intelligence*, 264(27–51), 2018.

Krzysztof R. Apt and Sunil Simon. A classification of weakly acyclic games. In Maria Serna, editor, *Algorithmic Game Theory*, pages 1–12, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg.

John Bartholdi, III, Craig Tovey, and Michael Trick. The computational difficulty of manipulating an election. *Social Choice and Welfare*, 6(3):227–241, 1989.

Sven Berg. Paradox of voting under an urn model: The effect of homogeneity. *Public Choice*, 47 (2):377–387, Jan 1985.

Simina Brânzei, Ioannis Caragiannis, Jamie Morgenstern, and Ariel D. Procaccia. How bad is selfish voting? In *Proceedings of the 27th AAAI Conference on Artificial Intelligence*, pages 138–144, 2013.

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

Samir Chopra, Eric Pacuit, and Rohit Parikh. Knowledge-theoretic properties of strategic voting. In José Júlio Alferes and João Alexandre Leite, editors, *Logics in Artificial Intelligence, 9th European Conference, JELIA 2004, Lisbon, Portugal, September 27-30, 2004*, volume 3229, pages 18–30. Springer, 2004.

Vincent Conitzer and Toby Walsh. Barriers to Manipulation in Voting. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel Procaccia, editors, *Handbook of Computational Social Choice*, chapter 6. Cambridge University Press, 2016.

Yvo Desmedt and Edith Elkind. Equilibria of plurality voting with abstentions. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 347–356, Cambridge, MA, USA, 2010.

Ulle Endriss, Svetlana Obraztsova, Maria Polukarov, and Jeffrey S. Rosenschein. Strategic voting with incomplete information. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*, IJCAI'16, page 236–242. AAAI Press, 2016.

Alex Fabrikant, Aaron D. Jaggard, and Michael Schapira. On the structure of weakly acyclic games. In Spyros Kontogiannis, Elias Koutsoupias, and Paul G. Spirakis, editors, *Algorithmic Game Theory*, pages 126–137, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg.

Piotr Faliszewski and Ariel D. Procaccia. AI's war on manipulation: Are we winning? *AI Magazine*, 31(4):53–64, 2010.

Piotr Faliszewski, Edith Hemaspaandra, and Lane A. Hemaspaandra. Using complexity to protect elections. *Communications of the ACM*, 53:74–82, 2010.

Allan Gibbard. Manipulation of voting schemes: A general result. *Econometrica*, 41:587–601, 1973.

Raphael Gillett. Collective Indecision. *Behavioral Science*, 22(6):383–390, 1977.

M. Goemans, Vahab Mirrokni, and A. Vetta. Sink equilibria and convergence. In *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS'05)*, pages 142–151, 11 2005.

Umberto Grandi, Andrea Loreggia, Francesca Rossi, Kristen Brent Venable, and Toby Walsh. Restricted manipulation in iterative voting: Condorcet efficiency and borda score. In Patrice Perny, Marc Pirlot, and Alexis Tsoukiàs, editors, *Algorithmic Decision Theory*, pages 181–192, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.

Aaron Koolyk, Tyrone Strangway, Omer Lev, and Jeffrey S. Rosenschein. Convergence and quality of iterative voting under non-scoring rules. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, pages 273–279, 2017.

Omer Lev and Jeffrey S. Rosenschein. Convergence of iterative voting. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems - Volume 2*, AAMAS '12, page 611–618, Richland, SC, 2012. International Foundation for Autonomous Agents and Multiagent Systems.

Reshef Meir. Strong and weak acyclicity in iterative voting. In Martin Gairing and Rahul Savani, editors, *Algorithmic Game Theory*, pages 182–194, Berlin, Heidelberg, 2016. Springer Berlin Heidelberg.

Reshef Meir, Maria Polukarov, Jeffrey S. Rosenschein, and Nicholas R. Jennings. Convergence to Equilibria of Plurality Voting. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 823–828, Atlanta, GA, USA, 2010.

Reshef Meir, Omer Lev, and Jeffrey S. Rosenschein. A Local-Dominance Theory of Voting Equilibria. In *Proceedings of the 15th ACM Conference on Electronic Commerce*, pages 313–330, Palo Alto, CA, USA, 2014.

Dov Monderer and Lloyd S. Shapley. Potential games. *Games and Economic Behavior*, 14(1): 124–143, 1996.

Svetlana Obraztsova, Evangelos Markakis, and David R. M. Thompson. Plurality voting with truth-biased agents. In Berthold Vöcking, editor, *Algorithmic Game Theory*, pages 26–37, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.

Svetlana Obraztsova, Evangelos Markakis, Maria Polukarov, Zinovi Rabinovich, and Nicholas R. Jennings. On the convergence of iterative voting: How restrictive should restricted dynamics be? In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, AAAI'15, page 993–999. AAAI Press, 2015. ISBN 0262511290.

Ariel D. Procaccia and Jeffrey S. Rosenschein. The distortion of cardinal preferences in voting. In *Proceedings of the 10th International Conference on Cooperative Information Agents*, CIA'06, page 317–331, Berlin, Heidelberg, 2006. Springer-Verlag.

Zinovi Rabinovich, Svetlana Obraztsova, Omer Lev, Evangelos Markakis, and Jeffrey Rosenschein. Analysis of equilibria in iterative voting schemes. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence*. AAAI Press, 2015.

Annemieke Reijngoud and Ulle Endriss. Voter response to iterated poll information. In *Proceedings of the Eleventh International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, volume 1, pages 635–644, 06 2012.

Reyhaneh Reyhani and Mark C. Wilson. Best reply dynamics for scoring rules. In *Proceedings of the 20th European Conference on Artificial Intelligence*, ECAI'12, page 672–677, NLD, 2012. IOS Press.

Mark Satterthwaite. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10: 187–217, 1975.

David R.M. Thompson, Omer Lev, Kevin Leyton-Brown, and Jeffrey Rosenschein. Empirical analysis of plurality election equilibria. In *Proceedings of the Thirteenth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 391–398, Paris, France, 2013.

Alan Tsang and Kate Larson. The echo chamber: Strategic voting and homophily in social networks. In *Proceedings of the 2016 International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '16, page 368–375, Richland, SC, 2016. International Foundation for Autonomous Agents and Multiagent Systems.

Lirong Xia. The Smoothed Possibility of Social Choice. In *Proceedings of NeurIPS*, 2020.

Lirong Xia. How Likely Are Large Elections Tied? In *Proceedings of the 22nd ACM Conference on Economics and Computation*, Budapest, Hungary, 2021.

Lirong Xia and Vincent Conitzer. Stackelberg voting games: Computational aspects and paradoxes. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 921–926, Atlanta, GA, USA, 2010.

James Zou, Reshef Meir, and David Parkes. Strategic voting behavior in doodle polls. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*, CSCW '15, page 464–472, New York, NY, USA, 2015.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes] The abstract and introduction accurately reflects the paper's main technical contributions - Theorems 1 and 2.
   (b) Did you describe the limitations of your work? [Yes] See Subsection 1.2.
   (c) Did you discuss any potential negative societal impacts of your work? [No] We mentioned that future work could determine if social welfare cost is divided equitably among all agents, although we do not approach the topic in this work. See Subsection 1.2.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes] Assumptions for each Theorem and Lemma are listed in their respective statements. The iterative plurality voting model assumptions are described in more detail in the Preliminaries (Section 2).
   (b) Did you include complete proofs of all theoretical results? [Yes] The proofs to Theorems 1 and 2 and Lemma 4 are included in full in the main text of the paper. Lemmas 2 and 3 have proof overviews included in the main text of the paper and full proofs in Appendices A.2 and A.5 respectively. Lemma 1 is stated in the main text and proved in Appendix A.1.
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] Experimental results are included in Appendix B.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] Experimental results are included in Appendix B.
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] Experimental results are included in Appendix B.
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] Experimental results are included in Appendix B.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [N/A]
   (b) Did you mention the license of the assets? [N/A]
   (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]

(b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]

(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]