Proxy Manipulation for Better Outcomes[1]

Gili Bielous and Reshef Meir

Abstract

This paper offers a framework for the study of strategic behavior in proxy voting, where non-active voters delegate their votes to active voters. It further studies how proxy voting affects the strategic behavior of non-active voters and proxies (active voters) under complete and partial information. We focus on the median voting rule for single-peaked preferences. Our results show strategyproofness with respect to non-active voters. Furthermore, while strategyproofness does not extend to proxies, we show that under mild restrictions strategic behavior can lead to socially optimal outcomes. For partial information settings, our results show that while convergence is guaranteed, it may be sub-optimal.

1 Introduction

In the age of internet, we see an increase of platforms and mechanisms for collective decision-making. However, many of these platforms suffer from low participation rates [24, 15]. Thus, while there is an increase in the ability of individuals to influence collective decision-making in many areas, most decisions are made by a small, non-elected and non-representative groups of active voters. Partial participation may increase vote distortion [11] (the worst-case ratio between the social cost of the candidate elected and the optimal candidate, first defined in [19]); lead to counter-intuitive equilibria [7]; and significantly decrease the likelihood of selecting the Condorcet winner (when it exists) [10]. Above all, when the outcome of an election only considers a fraction of all opinions, it is unreasonable to assume that they accurately reflect the aggregated opinions of the collective.

Proxy voting, a long standing practice in politics and corporates [21], and an up-and-coming practice in e-voting and participatory democracies [18], aims at mitigating the adverse effects of partial participation. Non-active voters (followers) delegate their vote to another active voter (proxy), thereby at least having some influence on the outcome. In some cases, the outcome of proxy elections provide a better estimate of the aggregated social preference of all voters [5].

However, such delegation changes the power dynamic of voters by shifting some of the voting power to proxies. While much consideration is granted in the literature of social choice for the strategic behavior of voters [12, 23] and candidates [9, 22], there is little consideration of the *strategic behavior of proxies or followers* in proxy-mediated settings. Cohensius et al.[2017] consider strategic participation (i.e. selecting to participate or abstain) with mostly positive results; yet they pose the question of strategic behavior of proxies and followers as an open question, which was part of the inspiration to the current study.

Moreover, it is common to study strategic behavior in adversarial settings assuming complete information. However, this assumption may be unreasonable in the context of proxy voting. By delegating their vote, followers may wish to avoid the cognitive strain, time loss and other costs associated with determining and communicating their position. Thus, a setting that requires followers to explicitly define their positions negates these benefits of proxy voting for followers. Reijngoud and Endriss [20] propose a framework for the study of strategic behavior in partial information settings. We apply it to study strategic behavior of proxies.

[1]This work was previously published at EUMAS2022 [2], an extended version is available at arXiv:2305.10969 [3]

Our model considers a political spectrum over the real line [14, 8], using the median voting rule for single-peaked preferences that was shown to be strategyproof [17]. Our initial study considers strategyproofness and manipulability with respect to both followers and proxies positions. Then, we consider sequences where proxies react to other proxies' actions. Finally, we turn to study strategic behavior in partial information settings. Our contribution is as follows:

- Strategyproofness of the median voting rule for single-peaked preferences extends to followers in proxy voting.
- Proxy voting with the median voting rule is manipulable with respect to proxy positions.
- Under mild restrictions, sequences of manipulations converge to an optimal equilibrium.
- Manipulations under partial information may converge to a worse equilibrium than without delegation.

2 Model and Preliminaries

We define the model of *Strategic Proxy Games (SPG)* as follows. There is a set of voters $N = \{1, ..., n\}$, and a set of proxies (active agents) $\Phi = \{\varphi_1, ..., \varphi_m\} \subseteq N$. Non-active voters, i.e. the set $N \setminus \Phi$ are called followers. Each voter $1 \leq i \leq n$ has a position $p_i \in \mathbb{R}$ along the political spectrum. Voters are assumed to have single-peaked preferences with peak at p_i. That is, for every $x, y \in \mathbb{R}$, if $x < y \leq p_i$, then Voter i prefers y to x, and if $p_i \leq x < y$, then Voter i prefers x to y. A *profile* is a vector $s \in \mathbb{R}^n$, such that s_i is the position Voter i declares. We denote by (s_{-i}, s'_i) the profile that is equal to s except for the strategy of Voter i, that is s'_i. We adopt the model of [5], where followers each delegate their vote to the nearest proxy (as in [25]). That is, given a profile s, each Follower $i \in N \setminus \Phi$ delegates their vote to Proxy $\varphi_j \in \Phi$, where $\varphi_j = \text{argmin}_{\varphi_j \in \Phi} |s_j - s_i|$. All proxies delegate their vote to themselves. Voters' preferences are symmetric single-peaked for followers, that is, for every $x, y \in \mathbb{R}$, if $|x - p_i| < |y - p_i|$, then Follower i prefers x to y. Thus, voters' preferences are consistent with the delegation model. We assume that some tie-breaking scheme exists that only depends on positions of proxies.

Example 2.1. Consider the SPG appearing in Figure 1.

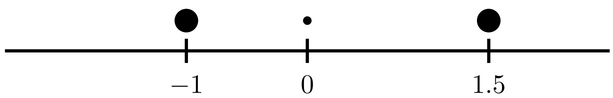

Figure 1: An example SPG. Large dots indicate the positions of proxies, small dots indicate the positions of followers.

There are three voters $N = \{1, 2, 3\}$ with positions $p_1 = -1$, $p_2 = 0$, and $p_3 = 1.5$, where $\varphi_1 = 1, \varphi_2 = 3$ are proxies. In the truthful profile $s = (-1, 0, 1.5)$, the follower (voter 2) delegates their vote to the closer proxy φ_1. Thus, there are two votes to -1 and a single vote to 1.5. ■

Given a finite set $S \subseteq \mathbb{R}$ such that each element $s_i \in S$ has weight $w_{s_i} \in \mathbb{R}^+$, let $W = \sum_{s_i \in S} w_{s_i}$. The weighted median of S is an element $s_i \in S$ such that $\sum_{\{s_j \in S \setminus \{s_i\} : s_j \leq s_i\}} w_{s_j} \leq \frac{W}{2}$ and $\sum_{\{s_j \in S \setminus \{s_i\} : s_j \geq s_i\}} w_{s_j} \leq \frac{W}{2}$. That is, the sum of weights of elements that are smaller than s_i is at most half the total sum of weights, and the same holds for the sum of weights of elements that are larger than s_i.

Next, we define the Weighted Median voting rule. The weight if each proxy is defined as the number of delegations to them. Then, the *weighted median voting rule* (WM) selects the position that is the weighted median of proxy positions. Note that in this case, $W = n$. For example, the WM in Example 2.1 is the position -1, as there are 2 votes for -1 (the proxy at -1 and the single follower who delegates to them), and 1 vote for 1.5 (the proxy at 1.5 with no followers).

For a profile s, we denote the unweighted median of s by med_s, or med when clear from context. For a truthful profile p the *median voter* is a voter i such that $p_i = med_p$.

We say that a voter is *truthful* if they declare their true position, i.e. $p_i = s_i$. Voters may lie about their positions, i.e. $p_i \neq s_i$. We assume that voters are rational, that is, voters lie only if the outcome changes in their favor. We say that Voter i *has a manipulation* in p if there is $s_i \neq p_i$ such that Voter i strictly prefers the outcome by reporting s_i to the outcome by reporting p_i. A voting rule is *strategyproof* if for every p, no voter has a manipulation, otherwise, it is *manipulable*. The Median voting rule is known to be (group) strategyproof for single-peaked preferences [4, 17].

3 Strategyproofness for Median Proxy Voting

We begin our analysis by showing that strategyproofness extends to Median Proxy Voting with respect to followers' positions. In [5] the authors show that for an infinite population of non-atomic voters given by some distribution where proxies are randomly selected, the winner of median proxy voting is the closest proxy to the true median. The following Lemma shows that this result extends to our setting. A similar variant appears in [22].

Lemma 3.1. *Let s be a profile, and let med_s be the median of s. Then the reported position s_j of $\varphi_j = \arg\min_{\varphi_i \in \Phi} |s_i - med|$ is the winner by WM.*

Proof. W.l.o.g, assume $s_j \leq med$. As there are at most $\frac{n}{2}$ voters with positions smaller than med_s, the sum of votes to proxies left of s_j is at most $\frac{n}{2}$. As all median followers (followers that report position s_j) delegate their vote to φ_j, and there are at most $\frac{n}{2}$ voters with positions greater than med_s, the sum of votes to proxies with positions greater than s_j is at most $\frac{n}{2}$. Thus, s_j is the weighted median of s. ∎

Next, we prove that for WM, followers do not have manipulations. Note that for followers manipulation implies delegation to another proxy.

Theorem 3.2. *WM is strategyproof w.r.t followers' positions.*

Proof. Assume towards contradiction that for some truthful profile s, there is a follower $i \in N \setminus \Phi$ that has a manipulation. W.l.o.g, assume $p_i \geq med_s$. Let φ be the proxy that is the winner for s, and let s'_i be the manipulation for Follower i. As s'_i is a manipulation, it must be that the winner for $s' = (s_{-i}, s'_i)$ changed. Let $\varphi' \neq \varphi$ be the winner for s'. By Lemma 3.1, $med_s \neq med_{s'}$. Therefore, it must be that $s'_i < med_s$, and therefore $med_{s'} < med_s$, hence $s_{\varphi'} < s_\varphi$. One of the following must hold:

- $p_i < s_{\varphi'} < s_\varphi$: Since $med_s \leq p_i$, it follows that $s_{\varphi'}$ is closer to med_s than s_φ, in contradiction to Lemma 3.1.
- $s_{\varphi'} < s_\varphi < p_i$ by single-peakedness, Follower i prefers s_φ to $s_{\varphi'}$, in contradiction to s'_i being a manipulation.
- $s_{\varphi'} < p_i < s_\varphi$: as $p_i \geq med_s$, and since by Lemma 3.1 $|med_s - s_\varphi| < |med_s - s_{\varphi'}|$, we get $|p_i - s_\varphi| < |p_i - s_{\varphi'}|$. Thus by symmetric single-peakedness Follower i prefers s_φ to $s_{\varphi'}$, in contradiction to s'_i being a manipulation.

∎

As Theorem 3.2 shows that WM is strategyproof with respect to followers' positions, we can henceforth consider them as non-strategic agents. In what follows, followers are considered to always be truthful, and we abuse the notation of a profile restricted to the strategic agents, i.e., the proxies.

We continue to analyze the strategic behavior of proxies. While we obtain a positive result of strategyproofness when only followers are considered strategic, the same does not hold for proxies, as demonstrated by the following example.

Example 3.3. Recall the SPG appearing in Example 2.1. When proxies are truthful, $s_1 = -1$ is the winner by the Weighted Median voting rule.

Next, consider the profile $s' = (-1, 1-\varepsilon)$ for some $0 < \varepsilon < 2$.

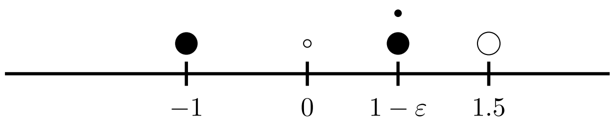

Figure 2: The SPG with φ_3 manipulation. Large empty dot is φ_3's true position, the manipulation is re-positioning strategically at $1-\varepsilon$. Follower delegates their vote to φ_3.

Follower 2 delegates their vote to φ_3. There are two votes to $s_2 = 1-\varepsilon$ and a single vote to $s_1 = -1$, thus, $1-\varepsilon$ is the winner by WM. As preferences are single-peaked and φ_3's peak is at $p_3 = 1.5$, we get that φ_3 prefer $1-\varepsilon$ to -1. Hence, s' is a manipulation for φ_3. ■

To highlight the pervasiveness of this result, we point out that the (unweighted) median voting rule is strategyproof for single-peaked preferences. Example 3.3 shows that even when we relax preferences of followers to be symmetric single-peaked, strategyproofness does not extend to WM with respect to positions of proxies.

Moreover, this example can be easily expanded to any number of followers and proxies. However, rather than formally constructing such example, The following theorem provides a complete characterization of manipulable scenarios. As a consequence, it shows that manipulations exist under very simple and reasonable conditions.

Theorem 3.4. *There is a proxy that has a manipulation in the profile s iff it holds that $s_{\varphi_i} \neq med$ for all $1 \leq i \leq m$, and there are proxies $\varphi_i, \varphi_j \in \Phi$ such that $p_{\varphi_i} < med < p_{\varphi_j}$.*

Proof. First, assume the winner is φ^* with position p_{φ^*}, and assume $p_{\varphi^*} < med$. For proxy φ_j we have that $p_{\varphi^*} < med < p_{\varphi_j}$. As preferences are single-peaked, φ_j prefers med to p_{φ^*}. We proceed by showing that $s_{\varphi_j} = med$ is a manipulation for Proxy φ_j, as in Fig. 3.

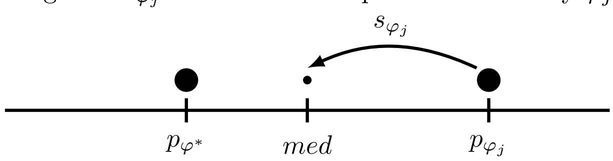

Figure 3: med is a manipulation for φ_j.

The median of (p_{-j}, s_{φ_j}) remains med, as there are at most $\frac{n}{2}$ voters with position that are smaller than med, hence the sum of votes to all proxies with positions smaller than med is at most $\frac{n}{2}$. The same holds for the sum of votes to proxies with position larger than med. Since φ_j reports position $s_{\varphi_j} = med$, by Lemma 3.1 their position med is the winner by WM. Since φ_j prefer med to p^*, that is a manipulation for φ_j. When $p^*_\varphi > med$, by the same reasoning $s_{\varphi_i} = med$ is a manipulation for φ_i. Hence, no proxy has position at med, and there are φ_i, φ_j such that $p_{\varphi_i} < med < p_{\varphi_j}$, then there is a proxy that has a manipulation.

Next, if there is some proxy φ_k such that $s_{\varphi_k} = med$, then by Lemma 3.1 med is the winner and no proxy with position that is not med can change the outcome by reporting a position that is closer to *med*. Furthermore, every proxy with position at *med* have their peak outcome, so they cannot improve the outcome for them. Therefore, the only possible manipulations are by proxies with positions other than *med*, and they can only manipulate by reporting a position that changes the location of the median. Assume towards contradiction that there is such a proxy φ, and w.l.o.g assume that $p_\varphi > med$. Then, to change the location of the median φ must report a position $s_\varphi < med$. Denote the median of $(p{-}\varphi, s_\varphi)$ by med'. We get that $med' < med$. Since s_φ is a manipulation, the outcome is at position $p' \leq med' < med < p_\varphi$. By single-peakedness φ prefers *med* to p', in contradiction to s_φ being a manipulation.

Finally, assume that for all proxies $p_{\varphi_i} \leq p_{\varphi^*} < med$. By the same reasoning as for the case where there is a proxy with position at *med*, no proxy can change the outcome in their favor by reporting a position larger than *med*. Then, by Lemma 3.1 the only way to change the outcome is to get closer to the true median, i.e. set a strategy s_φ such that $p_{\varphi^*} < s_\varphi \leq med$. As proxy preferences are single-peaked, every proxy prefer p_{φ^*} to any position $p_{\varphi^*} < s_\varphi$, thus, they do not have a manipulation. ■

4 Manipulations for Better Outcomes

So far, we showed not only that proxy voting using WM is manipulable, but also that manipulations exist in common voting scenarios. Though strategyproofness is usually considered as a desirable property for voting rules, in what follows we argue that in the case of proxy voting, manipulations can actually be proven useful.

Recall that one of the motivations for proxy voting is to mitigate the caveats of partial participation. In [5], the authors show that proxy voting can better aggregate voter preferences as it can only reduce the distance from the true median over partial participation. The true median is the outcome of the median voting rule. It is both Condorcet consistent and the minimal sum of distances from voters' true preferences. It is common in Hotelling-Downs-like settings [8, 14] to measure the social cost of outcomes using the sum of distances of all agents positions to the outcome. Thus, the median of all voters reflects the social optimum. Moreover, reducing the distance of the outcome from the median may improve the social welfare, and generally can better reflect the collective preference.

4.1 Dynamics and Convergence

While manipulations are actions that agents may take from their truthful profile to get a better outcome, proxies may continue to take actions and reposition themselves to get better outcomes. In this section, we discuss on-going dynamics for proxies.

Our model offers an infinite action set. Hence, the terminology used in Iterative Voting [16], which is the standard framework for studying the ongoing dynamics in voting, cannot be directly applied. We address it when relevant.

In what follows, we assume scenarios where manipulations exist, i.e., that meet the conditions of Theorem 3.4.

For every $\varphi_i \in \Phi$ and every profile $s_{-\varphi_i}$, we say that the position s'_{φ_i} is a *better-response* to s if φ_i prefers the outcome of $\left(s_{-\varphi_i}, s'_{\varphi_i}\right)$ to the outcome of s. We denote the set of better-responses for a profile s and proxy φ_i by $\mathcal{B}_s^{\varphi_i}$. We abuse the terminology of manipulations, such that a better-response from any profile s is a manipulation. A profile s is a *pure Nash equilibrium* (PNE) if for every $\varphi_i \in \Phi$ it holds that $\mathcal{B}_s^{\varphi_i} = \varnothing$, that is, no proxy have a manipulation for s.

A *policy* for proxy $\varphi_i \in \Phi$ is a function that maps a profile to a manipulation in the better-response set. Formally, let $\mathcal{S}$ be the set of all possible profiles for the proxies, and let $\mathcal{S}^* = \bigcup_{k=1}^{\infty} \mathcal{S}^k$. Then, a policy for φ_i is a function $\pi_{\varphi_i} : \mathcal{S} \to \mathbb{R}$ such that $\pi_{\varphi_i}(s) \in \mathcal{B}_s^{\varphi_i}$. One particular policy is the *best-response policy*, which selects a position with an outcome that the proxy prefers to all other positions in the better-response set, when one exists.

A *better-response dynamics* is a series of profiles, such that for every two consecutive profiles s^i, s^{i+1} in the series s^*, there is a proxy φ_j such that $s^{i+1} = \left(s^i_{-\varphi_j}, \pi_{\varphi_j}\left(s^i\right)\right)$. That is, every profile in the series is created from a single manipulation by some proxy, according to that proxy's policy. We say that a dynamics s^* *converges* if the series s^* has a limit.

4.2 Monotone Policies

For proxies that are on the other side of the median than the outcome, it is reasonable to assume that their policies select a position that is on the side of the median as their position. This is due to the fact that every position on their side of the median is a better-response to every position on the opposite side of the median. In the following discussion, we restrict policies to ones that preserves the integrity of proxies positions with respect to the median.

Formally, we say that a better-response dynamics s^* is *monotone* if for every $\varphi_i \in \Phi$, we have that $p_\varphi \leq med$ iff for every s^t it holds that $\pi_{\varphi_i}(s^t) \leq med$. Note that for every monotone better-response s^* starting from the truthful profile, for every s^t of s^* it holds that the median of s^t is *med*. We discuss non-monotone dynamics in Subsection 4.4.

For a better-response dynamics s^*, and for a profile s^t of s^*, we say that $\varphi(t)$ is the *moving* proxy at t if $s^{t+1} = \left(s^t_{-\varphi(t)}, \pi_{\varphi(t)}\left(s^t\right)\right)$. We denote the manipulation of $\varphi(t)$ at s^t by $s'(t) = \pi_{\varphi(t)}(s^t)$. We say that $\varphi^*(t)$ is a *winning* proxy at s^t if the outcome of s^t is $s_{\varphi^*(t)}$, and denote $s^*(t) = s^t_{\varphi^*(t)}$. Finally, denote $\Delta_t = |med - s^*(t)|$, i.e. the distance between the median and the outcome of s^t.

We next show that any manipulation in a monotone better-response dynamics where the winning proxy is not the moving proxy decreases the distance to the median.

Lemma 4.1. *For every monotone better-response dynamics s^* starting from the truthful profile $s^1 = p$, for every $t \geq 1$ if $\varphi(t) \neq \varphi^*(t)$, then $\Delta_{t+1} < \Delta_t$.*

Proof. By Lemma 3.1, for s^{t+1} it holds that $\varphi^*(t+1) = \arg\min_{\varphi_k \in \Phi} |s^{t+1}_{\varphi_k} - med|$, therefore, for every $\varphi_k \in \Phi$ it holds that $|s^*(t+1) - med| \leq |s^{t+1}_{\varphi_k} - med|$. In particular, this holds for $\varphi^*(t)$. We get:

$$|s^*(t+1) - med| \leq |s^{t+1}_{\varphi^*(t)} - med|$$

Since $\varphi(t) \neq \varphi^*(t)$, it holds that $s'(t)$ is a manipulation for $\varphi(t)$, so $|s^{t+1}_{\varphi^*(t+1)} - med| \neq |s^{t+1}_{\varphi^*(t)} - med|$. Hence:

$$\Delta_{t+1} = |s^*(t+1) - med| < |s^*(t) - med| = \Delta_t. \qquad \blacksquare$$

Lemma 4.1 suggests that manipulations made by proxies that do not have strategic positions at the current outcome reduce the distance to the true median. However, it is possible for winning proxies to manipulate in a way that increase the distance to the median. Figure 4 describe a proxy with 2 consecutive steps. The first makes them the winning proxy, the next is a better-response as they remain the winning proxy with a position that is closer to their true position.

We call sequences of consecutive manipulations by the same winning proxy *meta-move*. The following shows that while local manipulations within a meta-move can increase the current distance to the true median (as Figure 4 demonstrates), meta-moves globally decrease the distance to the true median.

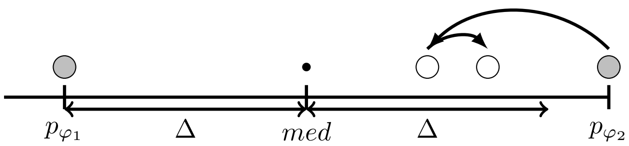

Figure 4: Consecutive steps that increase the distance to the median. Gray dots indicate truthful positions of proxies, empty dots indicate positions of manipulation. Arrows indicate repositions. A small full dot is the position of a (single) follower.

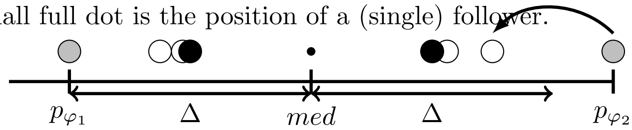

Figure 5: A dynamics that diverges. The two large black dots indicate oscillation positions. The arrow indicates the first manipulation.

Lemma 4.2. *Let s^* be a monotone better-response dynamics starting from the truthful profile $s^1 = p$. Then, every meta-move strictly decreases the distance between the winning position and med.*

Proof. We start by giving a formal description of meta-moves. Let s^k such that $\varphi(k) \neq \varphi^*(k)$ and $\varphi^*(k+1) = \varphi(k)$. That is, in profile s^k, a proxy $\varphi(k)$ manipulates such that the manipulation makes them the winner. Next, let $t \geq 1$ such that for every $1 \leq i \leq t$ it holds that $\varphi(k+i) = \varphi(k) = \varphi^*(k+1)$. That is, once $\varphi(k)$ becomes the winning proxy, they keep making consecutive manipulations for t steps. We show that $\Delta_{k+t} < \Delta_k$.

By Lemma 3.1, monotonicity and since for every $1 \leq i \leq t$ it holds that $\varphi(k+i) = \varphi^*(k+1) \neq \varphi^*(k)$, we get that $\Delta_{k+i} \leq \Delta_k$. Furthermore, since $s'(k)$ is a manipulation for $\varphi(k)$, it must be that the outcome of s^{k+1} is not equal to the outcome of s^k. We get that for every i, s^{k+i} is a manipulation and therefore $s^{k+i} \neq s^k$. Thus $\Delta_{k+i} \neq \Delta_k$. ■

Lemma 4.1 and Lemma 4.2 together provide a complete analysis of the better-response sets for proxies, and show that the better-response set strictly decreases after each (meta) move. However, this alone is not sufficient for convergence.

Example 4.3. Recall the setting appearing in Example 2.1. Define $\alpha_1 = \frac{1}{4}$, and for every $t \in \mathbb{N}$, define $\alpha_{t+1} = \frac{1}{2}\alpha_t$. We define the following policy for $\varphi_i \in \Phi$:

$$\pi_{\varphi_i}\left(s^t\right) = med - sign\left(med - p_{\varphi_i}\right)\left(\Delta_t - \alpha_t\right)$$

For every $t \in \mathbb{N}$ we get that

$$\begin{aligned}\Delta_{t+1} = |s^{t+1}_{\varphi_{t+1}} - med| &= |med - sign\left(med - p_{\varphi_{t+1}}\right)\left(\Delta_t - \alpha_t\right) - med| \\ &= |-sign\left(med - p_{\varphi_{t+1}}\right)\left(\Delta_t - \alpha_t\right)| = \Delta_t - \alpha_t\end{aligned}$$

As $\alpha_t = \frac{1}{2}\alpha_{t-1}$, we get $\Delta_{t+1} = \Delta_1 - \sum_{i=0}^{t-2} \frac{1}{2^i}\alpha_1 = \Delta_1 - \alpha_1 \sum_{i=0}^{t-2} \frac{1}{2^i}$. As $t \to \infty$, we get that the distance to the median converges to $\Delta_1 - 2\alpha_1 = \Delta_1 - 2\frac{1}{4}\Delta_1 = \frac{1}{2}\Delta_1$, and the outcome oscillates between $-\frac{1}{2}$ and $\frac{1}{2}$, thus the best-response dynamics diverges. Figure 5 shows a schematic of this dynamic. ■

Note that Example 4.3 not only shows that monotone better-response dynamics need not converge, it also shows a key difference between our setting and Iterative Voting. We say that a dynamic is *acyclic* if there are no recurring states. For finite action sets, i.e., when the space of available manipulations for each agent is finite, acyclicity implies convergence. Example 4.3 shows that for infinite action spaces this does not hold.

In effect, α_t is the amount by which the outcome gets closer to the true median between steps. As Δ_t decreases, so does the leeway that proxies have to improve the outcome for themselves. While it is reasonable that α_t decreases as Δ_t decreases, Example 4.3 captures the behavior in which α_t decreases at a higher rate than Δ_t.

By restricting policies such that α_t and Δ_t decrease at the same rate, we can obtain convergence. Moreover, this guarantees that Δ_t itself converges to 0, meaning that the outcome converges to the true median.

Let s^* be a monotone better-response dynamics from the truthful profile $s^1 = p$ with policies π_{φ_i}, and let $0 < \alpha < 1$. We restrict policies such that manipulations must create a noticeable difference in the outcome. In particular, as Δ_t defines the interval in which manipulations are possible, such a restriction bounds the outcome away from Δ_t. Formally, for every $t \geq 1$, if $\varphi^*(t) \neq \varphi(t)$ then we require $|med - \pi_{\varphi_i}(s^t)| \leq \alpha \cdot \Delta_t$. For every t, l such that for every $0 \leq i \leq l$ it holds that $\varphi(t+i) = \varphi(t) = \varphi^*(t+1) \neq \varphi^*(t)$, we require $|med - \pi_{\varphi_i}(s^{t+i})| \leq \alpha \cdot \Delta_t$.

Theorem 4.4. *Under the above restrictions, every monotone better-response dynamics from the truthful profile converges, and the limit is a PNE where the outcome is the true median.*

Proof. First, if the series $\{\Delta_t\}_{t=1}^{\infty}$ converges to 0, then by definition of Δ_t, the distance between the outcome and the true median in s^* converges to 0. By Lemma 3.1, this implies that no proxy can change the outcome, thus, the better response set of every proxy is empty, and this is a PNE. Moreover, the outcome is the true median *med*.

Next, we argue that under the policy restrictions, $\Delta_t \to 0$ as $t \to \infty$. We construct a series $\{\Gamma_t\}_{t=1}^{\infty}$ as follows. $\Gamma_1 = \Delta_1$. If $\varphi^*(t) \neq \varphi(t)$, then set $\Gamma_{t+1} = \alpha \cdot \Gamma_t$. Otherwise, set $\Gamma_{t+1} = \Gamma_t$. We get that for every $t \geq 1$, it holds that $\Delta_t \leq \Gamma_t$. For the case where $\Gamma_{t+1} = \alpha \cdot \Gamma_t$ due to assumption and Lemma 4.1, and for the case where $\Gamma_{t+1} = \Gamma_t$ by Lemma 4.2.

Note that as long as $\Delta_t > 0$, there is a proxy with position not in $s^*(t)$, therefore, every manipulation for them strictly reduces the distance to the median. We therefore get that the amount of cases where $\Gamma_{t+1} = \Gamma_t$ is finite, and therefore for convergence it is sufficient to consider only the case where $\Gamma_{t+1} = \alpha \cdot \Gamma_t$. We get that $\Gamma_{t+1} = \alpha^t \cdot \Gamma_1 = \alpha^t \Delta_1$. As $\alpha < 1$, we get that $t \to \infty$ implies $\Gamma_t \to 0$.

Finally, since $\{\Gamma_t\}_{t=1}^{\infty}$ bounds $\{\Delta_t\}_{t=1}^{\infty}$ from above, and $\{\Gamma_t\}_{t=1}^{\infty}$ converges to 0, then $\{\Delta_t\}_{t=1}^{\infty}$ also converges to 0. ■

As the true median of voters is the socially optimal outcome, Theorem 4.4 implies that the strategic behavior of proxies can in fact produce a socially optimal outcome.

4.3 Discretization

In many real-world applications, the assumption that voters can express any position on the political spectrum $\mathbb{R}$ is unreasonable. Voters are unlikely to distinguish between positions that are too similar, and this is the case both for selecting their truthful position, and distinguishing between different proxy positions for delegation. In computerized settings, there is some limited resolution to the expression of preferences (e.g. a temperature or a monetary amount). As it turns out, any such limit eliminates the possibility of oscillation we encountered in the previous section. In this section, we assume w.l.o.g that the political spectrum is restricted to the set of all integers $\mathbb{Z}$.

For discrete spaces, every policy meets the conditions of Theorem 4.4. This is due to the fact that every manipulation made by a proxy with position that is not the current weighted median must decrease the distance to the true median by at least 1 (as the minimal distance between every distinct possible positions). Thus, the conditions are met for $\alpha = 1 - \frac{1}{\Delta_1}$. Therefore, for discrete spaces, every better-response dynamics converges, and the outcome is the true median, which is the socially optimal outcome.

Furthermore, for discrete spaces (in contrast to continuous) there is a well-defined best-response, that is to reposition at a distance that is one step closer to the true median than the current winner on their opposite side of the median. In particular, the best-response is monotone. Following the terminology of [16], a game has the *Finite Best Response Property (FBRP) from truth* if from any truthful profile, when restricted to best-responses, the dynamics converges. Thus, SPGs with WM are FBRP from truth.

4.4 Non-Monotone Policies

In the previous sections we restricted the set of policies to those that maintain the integrity of proxies. That is, proxies always position themselves in the same side of the median as their truthful positions. However, there may be cases where it might be beneficial for a proxy to deviate to a position that is on the other side of the median. Proxies might attempt this in an intention to divert the positions of proxies on the opposite side of the median, or they might be willing to shift the median a little in an attempt to cause convergence to this new position. The following example demonstrate such a scenario.

Example 4.5. Consider the SPG appearing in Figure 6

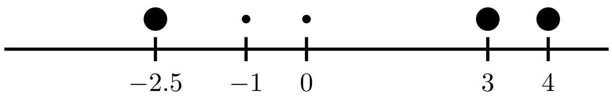

Figure 6: The SPG, large dots indicate proxies, small dots are followers.

The positions are $p = (-2.5, -1, 0, 3, 4)$. There are 3 proxies $\Phi = \{\varphi_1, \varphi_2, \varphi_3\}$ with positions $p_{\varphi_1} = -2.5, p_{\varphi_2} = 3, p_{\varphi_3} = 4$. There are 2 followers with positions $p_2 = -1, p_3 = 0$. The median is 0, and the outcome by the Weighted Median voting rule is the position $p_{\varphi_1} = -2.5$.

Assume that φ_2 manipulates to $s^1_{\varphi_2} = -2$. Now, the median is $p_2 = -1$, and the weighted median is $s^1_{\varphi_2}$. Note that φ_2 prefers this outcome to -2.5 by single-peakedness. Therefore, this is a manipulation for φ_2.

Next, assume that φ_3 manipulates to $s^2_{\varphi_3} = -1.5$. Now, the median in $s^2_{\varphi_3} = -1.5$, and it is also the position of the closest proxy, thus, this is the weighted median. This outcome is preferable to φ_3 than -2.

Finally, φ_2 manipulates to $s^3_{\varphi_2} = -1$. Now both the median and the weighted median is -1. Furthermore, both φ_1 and φ_3 can change the outcome in their favor.

Note that this does not imply convergence, as φ_2 still has a manipulation. ■

This example shows that a similar potential argument as used in the proof of Theorem 4.4, even with the added assumption that proxies make substantial enough steps to decrease the distance to the median, is unlikely to work. Our conjecture is that convergence holds for the unrestricted case as well, and that ultimately proxies would have an incentive to deviate back to their original side of the median, yet this is a matter of future research.

5 Partial Information

In previous sections we assumed that the proxies have complete information about the positions of proxies and followers alike. This assumption is common when analyzing adversarial behavior. However, is it reasonable in a proxy voting setting?

Recall that one of the applications of proxy voting is to mitigate the adverse effects of partial participation, where voters choose not to report their positions, rather only delegate their vote. Moreover, followers may not even know their exact position, rather they only know how to rank proxies by proximity. Thus, followers can still delegate their vote without the added cognitive strain of figuring what is their exact position.

When proxies have no information about positions of followers, then proxy voting is strategyproof. To see this, consider the profiles appearing in Figure 7

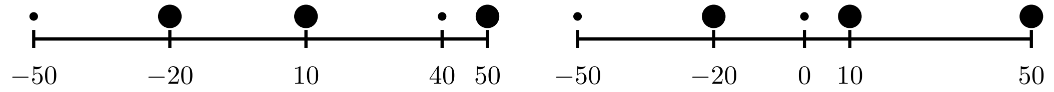

Figure 7: Example of profiles that are indistinguishable if followers positions are not public.

For the bottom profile, the proxy at -20 has a manipulation by deviating to -5. However, for the top profile, the proxy has no manipulation. When proxies have no information except proxies positions, the proxies cannot distinguish between the two profiles. Thus, proxies do not even know if the have a valid manipulation, let alone their better-responses.

This example shows that restricting the information available to proxies to reported proxy positions is too severe of an assumption.

In this section we relax the assumption of complete information. We first describe formally a less restrictive setting for the study of partial information. Next, we show that when only partial information is made available to voters, the strategic behavior of proxies may converge to a worse position than without delegation.

We employ the framework described in [20]. In this setting, a *poll information function (PIF)* σ maps each profile s to an information set $\sigma(s)$. For example, σ returns the outcome by WM, the number of delegations for each proxy, and even s itself. The set $\sigma(s)$ is then communicated to all voters.

In this setting, proxies cannot distinguish between profiles that yield the same information by σ. Recall the two profiles from Figure 7. When only proxy positions are communicated by σ, the profiles are indistinguishable by the proxies. Therefore, they must assume the profiles are equally likely. However, proxies can deduce an equivalent set of profiles that are consistent with the information they have. In particular, both profiles in Figure 7 would be in the same set.

Formally, Each proxy φ_i, based on their knowledge of their own preferences and the additional information $\sigma(s)$, deduce a set $W_{\varphi_i}^{\sigma(s)}$ of the possible positions of other followers (and proxies) that are compatible with the information set. That is, each profile $s_{-i}{}'$ in $W_i^{\sigma(s)}$ is a profile of all voters except i such that $(s_{-i}{}', s_i)$ is mapped by σ to the same information set they received $\sigma(s)$, i.e. $\sigma(s_{-i}{}', s_i) = \sigma(s)$.

Following the terminology of [6], we say that a position $s^*_{\varphi_i}$ is a *dominating manipulation* for Proxy φ_i if by reporting $s^*_{\varphi_i}$, there is some profile in $W_i^{\sigma(s)}$ that will produce a preferable outcome, and for all other profiles in $W_i^{\sigma(s)}$, it holds that their outcome is weakly preferred by φ_i over the current outcome. More formally, let F be a resolute voting rule, and let $\succ_{\varphi_i}$ be a full order over all possible outcomes that define φ_i's true preferences. Then, $s^*_{\varphi_i}$ is a dominating strategy if there is a profile $s_{-i}{}' \in W_i^{\sigma(s)}$ such that $F\left(s_{-i}{}', s^*_{\varphi_i}\right) \succ_i F\left(s_{-i}{}', s_i\right)$ and for all profiles $s_{-i}{}' \in W_i^{\sigma(s)}$ it holds that $F\left(s_{-i}{}', s^*_{\varphi_i}\right) \succcurlyeq_i F\left(s_{-i}{}', s_i\right)$. Note that if σ returns the profile s, then the set of dominating strategies coincides with the set of better-responses. Moreover, dominating manipulations are the only rational actions that a risk-averse agent may take.

For the rest of this discussion, We also assume that the positions of the proxies are known as a choice of modeling, as followers need to know their positions for delegation.

In this section, we assume that followers keep their positions as private information, and only delegate their vote to the proxy that is nearest to them. Therefore, for feasibility of delegation, we assume that proxies reveal their reported positions to all voters. Finally, we use σ_{winner}, the PIF that maps a profile s to the outcome of s by WM.

Next, we derive a similar positive result of convergence as in the complete information setting. First, as the PNE is defined with respect to better-responses, we consider convergence to a stable state (equilibrium) where none of the proxies have a dominating manipulation. However, as the position of the median is unknown to proxies, there is no straighforward interpretation of monotonicity. Instead, we consider a setting where proxies do not have a vote themselves. That is, votes are only delegated to them by followers. In this case, the position of the median is not affected by manipulations as is the case for monotone dynamics. This setting is closely related to the models of [5, 22].

We begin our analysis by characterizing the set of dominating manipulations for proxies.

Theorem 5.1. *For any profile s and proxy $\varphi \in \Phi$, the set of dominating manipulations of φ is the interval between the position of the current winner and the closest proxy to the winner on the same side as p_φ (including their truthful position).*

Proof. First, every position s'_φ in the set is a dominating manipulation. Consider the profile $\left(s_{-\varphi}, s'_\varphi\right)$. There are only two possible winners, either $s_{\varphi^*(s)}$ (the position of the current winner) or s'_φ. W.l.o.g assume $p_\varphi < s_{\varphi^*(s)}$. By single-peakedness we get $p_\varphi \leq s'_\varphi < s_{\varphi^*(s)}$, thus φ weakly prefers the outcome. Next, there is a profile where φ wins, thus it is a dominating manipulation.

Next, for every position that is farther than the closest proxy on the same side as φ's truthful position, the outcome of $\left(s_{-\varphi}, s'_\varphi\right)$ is $\varphi^*(s)$ no matter the positions of followers. Thus, it is not a dominating manipulation. For positions that are on the other side of the current winner than φ's truthful position, consider the profile where $s_{\varphi^*(s)}$ is the median, and there are no followers between $s_{\varphi^*(s)}$ and the position of the closest proxy on the other side of $s_{\varphi^*(s)}$ than φ. The outcome must be a position that is further from the truthful position of φ than $s_{\varphi^*(s)}$, thus it is not a dominating manipulation. ■

Next, we show that dominating manipulations decrease the distance to the true median.

Theorem 5.2. *Let s^* be a dynamics, then for every $t \geq 1$ it holds that $\Delta_{t+1} \leq \Delta_t$.*

Proof. Following Theorem 5.1, by repositioning to a dominating manipulations, the outcome either does not change, in which case $\Delta_{t+1} = \Delta_t$, or the moving proxy becomes the winner, in which case by Lemma 3.1 we get $\Delta_{t+1} < \Delta_t$. ■

We get that the distances between the winner and the true median is a decreasing monotone sequence bounded from below, thus, it converges. Therefore, under weak additional assumptions similarly to those made in the previous section (e.g. discretization) the step size is lower bounded so the dynamics converges. However, it is not guaranteed to converge to the true median, and in fact, may converge to a worse position than without delegation.

Consider the SPG appearing in Figure 8.

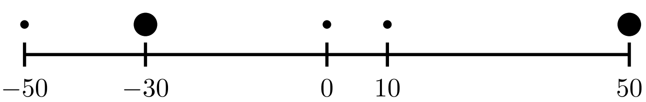

Figure 8: The SPG, large dots are proxies, small dots are followers.

The true positions are $p = (-50, -30, 0, 10, 50)$. There are two proxies $\Phi = \{\varphi_1, \varphi_2\}$ with positions $p_{\varphi_1} = -30$ and $p_{\varphi_2} = 50$, and 3 followers. Note that proxies and followers are unaware to the positions of other followers. The median is 0, and the weighted median is -30. The social cost, or sum of distances from the weighted median to each position is

$$SC = |-50 - (-30)| + |-30 - (-30)| + |-30 - 0| + |-30 - 10| + |-30 - 50| = 170$$

Next, consider the position ${s^1_{\varphi_2}}' = 25$. By Theorem 5.1, it is a dominating manipulation.

For φ_1 in $s^2 = \left(s^1_{-\varphi_2}, {s^1_{\varphi_2}}'\right)$, the only information that φ_1 has is that their position is -30, and that the position of φ_2 in s^2 is 25, and it is the outcome of s^2. Consider the position ${s^2_{\varphi_1}}' = 20$. Again, by Theorem 5.1 this is a dominating manipulation for φ_1. Moreover, φ_1 has no dominating manipulation in $s^3 = \left(s^2_{-\varphi_1}, {s^2_{\varphi_1}}'\right)$. Additionally, this holds for every $s^t = \left(s^2_{-\varphi_2}, {s^t_{\varphi_2}}'\right)$ where ${s^t_{\varphi_2}}' \in \left(s^2_{\varphi_1}, s^1_{\varphi_2}\right] = (20, 25]$.

Finally, for every $s^t = \left(s^2_{-\varphi_2}, {s^t_{\varphi_2}}'\right)$ where ${s^t_{\varphi_2}}' \in \left(s^2_{\varphi_1}, s^1_{\varphi_2}\right] = (20, 25]$, the set of dominating strategies for φ_2 is $\left(20, {s^t_{\varphi_2}}'\right]$. Figure 9 demonstrates the dynamics.

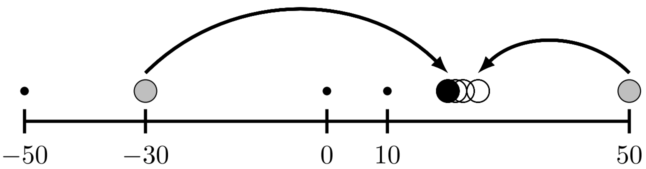

Figure 9: An example that converges to a worse position than without delegation. Gray dots indicate truthful positions of proxies, empty dots indicate positions of dominating manipulations, full dot indicate convergence positions. Small full dots are followers.

We get that the distance between the position of φ_1 and φ_2 converges to 0, and therefore the dynamics ultimately converges to a PNE where the positions of both proxies is $s^2_{\varphi_1} = 20$. The social cost of this outcome is $\sum_{i=1}^{5} |20 - p_i| = 180$, that is greater than the social cost of the outcome of p.

Note that in the case of complete information, this counter-example would not converge in the same way. This is due to the fact that once φ_1 repositions and becomes the winner again, as they know the position of the median, their better-response set is not empty.

6 Conclusions and Future Work

We introduced *Strategic Proxy Games*, a framework to study strategic behavior of proxies in voting mechanisms.

First, we demonstrated that in this model, the extension of the median voting rule to the weighted median voting rule via proxy voting maintains strategyproofness with respect to followers' positions. In particular, this suggests that with respect to follower positions, the delegation scheme is optimal for followers preferences. Our study uses the Tullock delegation scheme, however other delegation models have been studied in the literature. In the one-step delegation domain, Green-Armytage [13] consider delegation that accounts for small errors in assessment of positions, and Alon et al. [1] consider social connections that influence the weight of proxies. It would be interesting to see how the delegation model affects the outcome of proxy voting and the strategic behavior of followers and proxies.

We continued to study the strategic behavior of proxies, and showed that while strategyproofness does not extend to proxy voting, when proxies maintain the integrity of their positions with respect to the median, the outcome converges to the true median of all voters. This result implies that by relaxing truthfulness to integrity, strategic behavior can improve the outcome with respect to the truthful profile. In future work we would like to study the outcome without this restriction, and it is our conjecture that the outcome converges to the true median as well.

Finally, we study the implications of partial information to the strategic behavior of proxies. While we get a positive result of convergence, our results also show that in this case the outcome may increase the social cost.

In this research we focused on the median voting rule. We plan to study the implication of strategic proxy behavior in higher dimensions, as well as with other voting rules.

References

[1] Noga Alon, Michal Feldman, Omer Lev, and Moshe Tennenholtz. How robust is the wisdom of the crowds? In *Twenty-Fourth International Joint Conference on Artificial Intelligence*. Citeseer, 2015.

[2] Gili Bielous and Reshef Meir. Proxy manipulation for better outcomes. In *Multi-Agent Systems: 19th European Conference, EUMAS 2022, Düsseldorf, Germany, September 14–16, 2022, Proceedings*, pages 79–95. Springer, 2022.

[3] Gili Bielous and Reshef Meir. Strategic proxy voting on the line. *arXiv preprint arXiv:2305.10969*, 2023.

[4] Duncan Black. On the rationale of group decision-making. *Journal of political economy*, 56(1):23–34, 1948.

[5] Gal Cohensius, Shie Manor, Reshef Meir, Eli Meirom, and Ariel Orda. Proxy voting for better outcomes. In *AAMAS'17*, 2017.

[6] Vincent Conitzer, Toby Walsh, and Lirong Xia. Dominating manipulations in voting with partial information. In *Twenty-Fifth AAAI Conference on Artificial Intelligence*, 2011.

[7] Yvo Desmedt and Edith Elkind. Equilibria of plurality voting with abstentions. In *Proceedings of the 11th ACM conference on Electronic commerce*, pages 347–356, 2010.

[8] Anthony Downs. *An economic theory of democracy*. Harper & Row New York, 1957.

[9] Bhaskar Dutta, Matthew O Jackson, and Michel Le Breton. Strategic candidacy and voting procedures. *Econometrica*, 69(4):1013–1037, 2001.

[10] William V Gehrlein and Dominique Lepelley. *Voting paradoxes and group coherence: the Condorcet efficiency of voting rules*. Springer Science & Business Media, 2010.

[11] Mohammad Ghodsi, Mohamad Latifian, and Masoud Seddighin. On the distortion value of the elections with abstention. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 1981–1988, 2019.

[12] Allan Gibbard. Manipulation of voting schemes: a general result. *Econometrica: journal of the Econometric Society*, pages 587–601, 1973.

[13] James Green-Armytage. Direct voting and proxy voting. *Constitutional Political Economy*, 26(2):190–220, 2015.

[14] Harold Hotelling. Stability in competition. *The Economic Journal*, 39(153):41–57, 1929.

[15] Anna Maria Jönsson and Henrik Örnebring. User-generated content and the news: Empowerment of citizens or interactive illusion? *Journalism Practice*, 5(2):127–144, 2011.

[16] Reshef Meir. Iterative voting. *Trends in computational social choice*, pages 69–86, 2017.

[17] Hervé Moulin. On strategy-proofness and single peakedness. *Public Choice*, 35(4): 437–455, 1980.

[18] Klaus Petrik. Participation and e-democracy how to utilize web 2.0 for policy decision-making. In *10th Annual International Conference on Digital Government Research: Social Networks: Making Connections between Citizens, Data and Government*, pages 254–263, 2009.

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

[20] Annemieke Reijngoud and Ulle Endriss. Voter response to iterated poll information. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*, pages 635–644, 2012.

[21] Floyd Millard Riddick and Miriam H Butcher. *Riddick's Rules of procedure: a modern guide to faster and more efficient meetings*. Madison Books, 1991.

[22] Itay Sabato, Svetlana Obraztsova, Zinovi Rabinovich, and Jeffrey S Rosenschein. Real candidacy games: a new model for strategic candidacy. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, pages 867–875, 2017.

[23] Mark Allen Satterthwaite. Strategy-proofness and arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of economic theory*, 10(2):187–217, 1975.

[24] L Christian Schaupp and Lemuria Carter. E-voting: from apathy to adoption. *Journal of Enterprise Information Management*, 2005.

[25] Gordon Tullock. Proportional representation. *Toward a mathematics of politics*, pages 144–157, 1967.

Gili Bielous
Technion- Israel Institute of Technology
Haifa, Israel
Email: `gili.bielous@campus.technion.ac.il`

Reshef Meir
Technion- Israel Institute of Technology
Haifa, Israel
Email: `reshefm@iem.technion.ac.il`