

Does Delegating Votes Protect Against Pandering Candidates?

Extended Abstract

Xiaolin Sun
Tulane University
New Orleans, USA
xsun12@tulane.edu

Jacob Masur
Tulane University
New Orleans, USA
jmasur@tulane.edu

Ben Abramowitz
Tulane University
New Orleans, USA
babramow@tulane.edu

Nicholas Mattei
Tulane University
New Orleans, USA
nsmattei@tulane.edu

Zizhan Zheng
Tulane University
New Orleans, USA
zzheng3@tulane.edu

ABSTRACT

The election of representatives in regular election cycles ostensibly prevents misbehavior by elected officials and keeps them accountable in service of the “will of the people.” This democratic ideal can be undermined if candidates campaign dishonestly when seeking office over one or more election cycles or ‘rounds’. We introduce a novel formal model of *pandering*, or strategic preference reporting by electoral candidates, and examine the resilience of two democratic voting systems to such pandering. The two voting systems we compare are Representative Democracy (RD) and Flexible Representative Democracy (FRD). For each voting system, our analysis centers on the types of strategies candidates employ and how voters update their views of candidates across rounds based on how the candidates have pandered in the past. We provide theoretical results on the complexity of pandering for a single round, formulate our problem for multiple rounds as a Markov Decision Process, and use reinforcement learning to study the effects of pandering by sets of candidates across a number of rounds.

KEYWORDS

Social Choice; Delegative Voting; Manipulation; Reinforcement Learning; Flexible Representative Democracy; Pandering

ACM Reference Format:

Xiaolin Sun, Jacob Masur, Ben Abramowitz, Nicholas Mattei, and Zizhan Zheng. 2023. Does Delegating Votes Protect Against Pandering Candidates?: Extended Abstract. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023)*, London, United Kingdom, May 29 – June 2, 2023, IFAAMAS, 3 pages.

The will of the people shall be the basis of the authority of government; this will shall be expressed in periodic and genuine elections...

Article 21, Universal Declaration of Human Rights, 1948 [4]

1 INTRODUCTION

Modern representative democracies use regular elections to ensure that officials uphold the “will of the people.” Periodic elections are meant to prevent corrupt or ineffective officials from maintaining

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom. © 2023 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

power and to keep them honest. However, elections are arguably not enough as voters only have a say during the election (aside from the potential recalls), which typically occur at regular intervals.

In recent years, variants of *delegative* voting have been advanced in the computer science and social choice literature, including liquid democracy [6], flexible representative democracy [1], and weighted representative democracy [17]. In these delegative voting systems the voters collectively weight their representatives, possibly updating their weights between elections. These voting systems fall under the umbrella of *interactive democracy* [8], which encompasses the idea that modern internet and communication tools allow us to create decision-making procedures more interactive and responsive for large populations in ways not possible in the past. However, much of the work on COMSOC to date has investigated strategic actions in only a single election cycle.

Some proposed delegative voting schemes can interpolate between direct and representative democracy and may be better at keeping representatives accountable [1, 9]. However, the idea that delegative voting will be better at keeping representatives accountable, or that it will be better at expressing the “will of the people”, is largely untested aside from some nascent applications of Liquid Democracy (with transitive delegations) [11, 16]. Little is known about how such systems perform in the presence of agents who are strategic, selfish, and even malicious [5, 20]. Understanding responsiveness to voter preferences and robustness to bad actors is critical for selecting and comparing voting systems.

One of the primary features of representative electoral systems is that candidates campaign while seeking office, making promises and stating positions on future decisions. Unfortunately, politicians lie, especially when trying to get elected or maintain power. This *pandering* is a form of attack on representative democratic systems, and we introduce the first formal model of pandering [18] to the literature on Computational Social Choice (COMSOC), which has previously considered other forms of election attack including manipulation, bribery, and control [7].

Political pandering is a global phenomenon. US citizens consistently rank Congresspeople as occupying the least trustworthy profession [13], the USA has arguably reached its highest level of corruption in a decade [10], and over half of Americans are unsatisfied with representative democracy as it stands [19]. In Spain, a study involving Spanish mayors [12] demonstrated that lying may improve a politician’s chances of being reelected. However,

Australian voters demonstrably decrease support for politicians upon the revelation of their lies [2]. Broadly, voters are often aware of pandering and are suspicious of perceived panders [14, 15].

1.1 Electoral Pandering Model

Whether a democratic system is robust to dishonest candidates depends on how much voters know about the dishonesty, how voters react to such dishonesty, and what tools the voters have at their disposal to respond to it. Our core concern is whether delegative voting systems are more or less vulnerable to dishonest candidates. To study these questions we model two types of democratic voting system; classic Representative Democracy (RD) and Flexible Representative Democracy (FRD) [1]. The difference between RD and FRD is that in FRD the representatives vote using a weighted majority rule where the weights are determined by the voters, while in RD the representatives use unweighted majority rule.

For both voting systems we study, each election cycle consists of (1) voters electing a subset of candidates as a committee of representatives, and (2) the representatives voting sequentially on a fixed set of issues. This stylized setting maps to most representative democracies across the world where representatives decide on a slate of issues between elections [3]. To study the effect of strategic electoral candidates on voting systems we first define the basics of voting systems, introduce our novel model of pandering, and investigate the complexity of this type of attack in a single election. We then move on to studying sequences of elections in which candidates learn pandering strategies by reinforcement learning while the voters' responses to observed pandering are fixed. See Sun et al. [18] for our working paper with a full set of results.

Let V be a set of n voters and C be a disjoint set of m candidates. The voters elect a subset of candidates $D \subset C$ where $|D| = k$ to serve as representatives. The set of representatives will then vote on a sequence of r binary issues. We assume that every voter and candidate has a binary preference over every issue. For voter $v \in V$, we denote their preference vector by $v \in \{0, 1\}^r$. Similarly, for candidate $c \in C$, their preference vector is denoted $c \in \{0, 1\}^r$. The collective preference profiles of the agents are denoted by V and C .

With m candidates, there are $\binom{m}{k}$ possible ways to elect k representatives. However, it is infeasible for voters to express preferences over all possible committees of size k . Therefore, representatives are elected via k -Approval with random tie-breaking. Each voter reports the subset of candidates of which they approve and the k candidates who receive the greatest number of approvals get elected. Following Abramowitz and Mattei [1], voters submit approval preferences over candidates based on the fraction of issues on which they agree. That is, v approves of c if $g(v, c) > 1/2$ where $g(v, c)$ is based on the Hamming distance between preference vectors. Let $d_H(\mathbf{x}, \mathbf{y}) = \sum_{i \leq r} |\mathbf{x}(i) - \mathbf{y}(i)|$ be the Hamming distance between two vectors of length r . For any two vectors \mathbf{x} and \mathbf{y} of length r , we refer to $g(\mathbf{x}, \mathbf{y}) = 1 - \frac{1}{r}d_H(\mathbf{x}, \mathbf{y})$ as their *agreement* and $\frac{1}{r}d_H(\mathbf{x}, \mathbf{y})$ as their *disagreement*. Intuitively, $g(v, c)$ is the fraction of issues the voter and candidate agree upon. We measure the quality of a voting system as the agreement (or disagreement) between the outcomes it produces and the outcomes preferred by the voter majority.

Following Abramowitz and Mattei [1] we model our democratic systems as follows. In classic (RD) the candidates are elected by

k -Approval with random tie-breaking, and each set of elected representatives votes on a sequence of r binary issues using simple majority voting before the next election. By contrast, in FRD the representatives use weighted majority voting on every issue and these weights are determined on every issue by the voters. Each voter has 1 unit of weight to assign to the representatives and may distribute it among the representatives however they wish. The weight of a representative on an issue is then the sum of weights assigned to them. That is, each voter v assigns each representative c a weight $0 \leq w^t(v, c) \leq 1$ on each issue t such that $\sum_{c \in D} w^t(v, c) = 1$ for all t and the weight of a representative is $w_c^t = \sum_{v \in V} w^t(v, c)$. If $c(t) \in \{0, 1\}$ is the preference of c on issue t , then weighted majority voting leads the outcome to be 1 if $\sum_{c \in D} w_c^t c(t) > n/2$, 0 if $\sum_{c \in D} w_c^t c(t) < n/2$, and breaks ties randomly otherwise.

We model two types of agents: selfish agents who want the decision outcomes to agree with their preferences as much as possible and malicious agents whose goal is to maximize disagreement with the voter majority. Voters decide who to vote for in the elections based on their agreement with candidates and the candidates' observed history of pandering (or honesty). In FRD, voters weight the representatives based on the same.

2 CONCLUSIONS

In the full version of this paper [18] we formalize and study a novel model of election attack, *pandering*, where candidates report their preferences strategically in order to get elected. Using tools from COMSOC and reinforcement learning, we analyze two democratic voting systems, representative democracy (RD) and flexible representative democracy (FRD), in terms of their resilience to pandering. We first show that the pandering problem itself is computationally hard (NP-Hard) for a single round and provide an optimization program to solve this problem. We then model the problem of pandering over multiple rounds as a sequential decision making problem, formally a Markov Decision Problem (MDP). We then use techniques from reinforcement learning to solve this problem for pandering candidates and investigate how robust RD and FRD are to these attacks.

We show that FRD resists the attacks of both malicious and selfish candidates better than RD, showing significantly higher agreement with the voter majority across all tested scenarios. The scenarios varied in the number of strategic agents, reactivity of voters to pandering, and distributions of preferences. Thus, we can draw the conclusion that FRD is more resilient than RD facing one or more strategic candidates pandering. Furthermore, we find that the damage from strategic candidates is usually almost linear with the number of strategic candidates, but a high tolerance for pandering leads to more coordination opportunities by sets of malicious candidates.

Our results are consistent with the intuition that holding regular elections is, in fact, important in upholding the "will of the people." In future work it would be interesting to look at how voters can learn strategies for their approval votes in the election, and how representatives can learn strategies for when to vote according to their true preferences and when to vote according to their reported preferences to appear more honest.

ACKNOWLEDGMENTS

Nicholas Mattei was supported by NSF Awards IIS-RI-2007955, IIS-III-2107505, and IIS-RI-2134857, as well as an IBM Faculty Award and a Google Research Scholar Award. Ben Abramowitz was supported by the NSF under Grant #2127309 to the Computing Research Association for the CIFellows Project. Xiaolin Sun and Zizhan Zheng were supported by NSF awards CNS-1816495 and CNS-2146548, and Tulane University Jurist Center for Artificial Intelligence.

REFERENCES

- [1] Ben Abramowitz and Nicholas Mattei. 2019. Flexible Representative Democracy: An Introduction with Binary Issues. arXiv:1811.02921 [cs.MA]
- [2] Michael J. Aird, Ullrich K. H. Ecker, Briony Swire, Adam J. Berinsky, and Stephan Lewandowsky. 2018. Does truth matter to voters? The effects of correcting political misinformation in an Australian sample. *Royal Society Open Science* 5, 12 (2018), 180593.
- [3] Christopher J Anderson, Damien Bol, and Aurelia Ananda. 2021. Humanity’s Attitudes about Democracy and Political Leaders: Patterns and Trends. *Public Opinion Quarterly* 85, 4 (11 2021), 957–986.
- [4] UN General Assembly et al. 1948. Universal declaration of human rights. *UN General Assembly* 302, 2 (1948), 14–25.
- [5] Daan Bloembergen, Davide Grossi, and Martin Lackner. 2019. On rational delegations in liquid democracy. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 1796–1803.
- [6] Christian Blum and Christina Isabel Zuber. 2016. Liquid democracy: Potentials, problems, and perspectives. *Journal of Political Philosophy* 24, 2 (2016), 162–182.
- [7] Felix Brandt, Vincent Conitzer, Ullrich Endriss, Jérôme Lang, and Ariel D Procaccia. 2016. *Handbook of computational social choice*. Cambridge University Press.
- [8] Markus Brill. 2018. Interactive democracy. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*. 1183–1187.
- [9] Bryan Alexander Ford. 2002. *Delegative democracy*. Technical Report.
- [10] Cailey Griffin and Amy Mackinnon. 2021. Report: Corruption in U.S. at worst levels in almost a decade. <https://foreignpolicy.com/2021/01/28/report-transparency-international-corruption-worst-decade-united-states/>
- [11] Reinhard Hainisch and Alois Paulin. 2016. Civocracy: Establishing a Competent and Responsible Council of Representatives Based on Liquid Democracy. In *2016 Conference for E-Democracy and Open Government (CeDEM)*. 10–16. <https://doi.org/10.1109/CeDEM.2016.27>
- [12] Katharina A. Janezic and Aina Gallego. 2020. Eliciting preferences for truth-telling in a survey of politicians. *Proceedings of the National Academy of Sciences* 117, 36 (2020), 22002–22008.
- [13] Niall McCarthy. 2021. America’s most & least trusted professions [infographic]. *Forbes* (6 2021).
- [14] Kathleen M. McGraw, Milton Lodge, and Jeffrey M. Jones. 2002. The Pandering Politicians of Suspicious Minds. *The Journal of Politics* 64, 2 (2002), 362–383. <https://doi.org/10.1111/1468-2508.00130> arXiv:<https://doi.org/10.1111/1468-2508.00130>
- [15] Joseph McMurray. 2017. Polarization and Pandering in Common-value Elections. *Brigham Young University. Manuscript* (2017).
- [16] Alois Paulin. 2020. An overview of ten years of liquid democracy research. In *The 21st Annual International Conference on Digital Government Research*. 116–121.
- [17] Marcus Pivato and Arnold Soh. 2020. Weighted representative democracy. *Journal of Mathematical Economics* 88 (2020), 52–63.
- [18] Xiaolin Sun, Jacob Masur, Ben Abramowitz, Nicholas Mattei, and Zizhan Zheng. 2022. Pandering in a Flexible Representative Democracy. <https://doi.org/10.48550/ARXIV.2211.09986>
- [19] Richard Wike and Shannon Schumacher. 2021. Democratic rights popular globally but commitment to them not always strong. *Pew Research Center’s Global Attitudes Project* (3 2021).
- [20] Yuzhe Zhang and Davide Grossi. 2021. Power in liquid democracy. In *Proceedings of the AAAI conference on Artificial Intelligence*, Vol. 35. 5822–5830.