Heuristic Strategies in Uncertain Approval Voting Environments

Extended Abstract

Jaelle Scheuerman
Tulane University
New Orleans, LA
jscheuer@tulane.edu

Jason L. Harman
Louisiana State University
Baton Rouge, LA
jharman@lsu.edu

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

K. Brent Venable
Institute for Human & Machine Cognition
University of West Florida
Pensacola, FL
bvenabl@ihmc.us

ABSTRACT

In many collective decision making situations, agents use voting to choose an alternative that best represents the preferences of the group. It is often assumed that voters will vote truthfully rather than expending the effort needed to manipulate the outcome in cognitively and computationally complex situations. However, being truthful is just one possible heuristic that agents may employ. We examine how real voters employ heuristics in a variety of approval voting scenarios. In particular, we consider heuristics where a voter ignores information about other voting profiles and makes their decisions based solely on how much they like each candidate. In a behavioral experiment, we show that people vote truthfully in some situations, but prioritize high utility candidates in others. We show how the structure of the voting environment affects how well each heuristic performs as well as how and when humans employ these different heuristics. [1]

KEYWORDS

voting; heuristics; computational social choice

ACM Reference Format:

Jaelle Scheuerman, Jason L. Harman, Nicholas Mattei, and K. Brent Venable. 2020. Heuristic Strategies in Uncertain Approval Voting Environments. In *Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), Auckland, New Zealand, May 9–13, 2020,* IFAAMAS, 3 pages.

1 INTRODUCTION

Computational Social Choice (COMSOC) investigates computational issues surrounding the aggregation of individual preferences and collective decision making [2]. Much of this work has focused on the computational complexity of manipulating elections under different voting rules, but less work has considered how heuristic decision making may affect the outcome of various voting rules [5]. It is often assumed that if the voting rule is sufficiently complicated, people will default to using a truthful strategy, voting for the candidates they like the best [6]. However, voting truthfully is just one possible heuristic that agents may use when faced with complex voting scenarios involving tradeoffs between multiple options and limited amounts of information. In our current line of work we turn to behavioral experiments to explore how and when agents employ different heuristic strategies in approval voting settings [12, 13].

Approval voting is an interesting rule as it offers voters the option of several sincere voting strategies, allowing them to vote for one or more of their preferred candidates [10]. A ballot is regarded as sincere if the voter prefers approved candidates over the disapproved candidates [4]. A recent study of voting behavior in multi-winner approval elections showed that the majority of voters did not vote truthfully or optimally [12]. Instead, the predominant strategy was to use a *take the X best* heuristic which prioritized the highest utility candidates. In decision making, heuristics strategies that ignore some information can sometimes achieve better performance than more complicated optimization strategies in situations that are computationally complex or uncertain [8] (e.g. stock market predictions [3]). Through a behavioral experiment, we investigate the use and effectiveness of three heuristics that are sincere, including *truthful*, *take the X best*, and *regret minimization.*

2 METHODS

Following Aziz et al. [1] and Kilgour [9] we consider a social choice setting (N, C) where we are given a set $N = \{1, \ldots, n\}$ of voting agents and a disjoint set $C = \{c_1, \ldots, c_m\}$ of candidates. Each agent $i \in N$ expresses an approval ballot $A_i \subseteq C$ which gives rise to a set of approval ballots $A = \{A_1, \ldots, A_n\}$, called a profile. We study the multi-winner approval voting rule that take as input an instance (A, C, k) and return a subset of candidates $W \subseteq C$ where $|W| = k$ is called the winning set. Approval Voting (AV) finds the set $W \subseteq C$ where $|W| = k$ that maximizes the total weight of approvals across all ballots, $AV(W) = \sum_{i \in N} |W \cap A_i|$. Informally, the winning set under AV is the set of candidates that are approved by the largest number of voters. In order to align our work with the literature on decision heuristics [7] we assume that each agent $i \in N$ also has a real valued utility function $u_i : C \rightarrow \mathcal{R}$. We also assume that the utility of agent i for a particular set of winning candidates $W \subset C$ is $u_i(W) = \sum_{c \in W} u_i(c)$ (slightly abusing notation). If W

[1]This extended abstract is based on the paper found at: https://arxiv.org/abs/1905.12104

Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), B. An, N. Yorke-Smith, A. El Fallah Seghrouchni, G. Sukthankar (eds.), May 9–13, 2020, Auckland, New Zealand.

is the subset elected by the voting rule we will refer to $u_i(W)$ as agent's i's utility for the election's *outcome*.

We define a *truthful* vote as one where an agent approves all candidates for which they have positive utility. When an agent votes with the *take the X best* heuristic, they vote for a subset of the truthful vote. First, they order the list of candidates by the utility value. The agent will then vote for the top-X candidates in the list. *Regret minimization* takes into account the voter's anticipated regret if a particular disliked candidate were to win the election. Rather than try maximize their utility, the voter chooses to minimize the chance that the disliked candidate(s) will win by voting for all other candidates, whether they generate positive utility or not [14].

		Candidates				
Scenarios:		**A**	**B**	**C**	**D**	**E**
Trivial Candidate	Utility:	0.05	0.10	0.01	0	0.25
	# Votes:	3	3	3	4	3
Trivial Leader	Utility:	0.05	0.10	0.01	0.25	0
	# Votes:	3	3	4	3	3
Dominated Prefs (1-2 winners)	Utility:	0.05	0.10	0	0	0.25
	# Votes:	1	1	4	4	1
Dominated Prefs (3 winners)	Utility:	0.10	0	0	0	0.25
	# Votes:	1	4	4	4	1
Disliked Candidate	Utility:	0.05	0.10	0	-1.00	0.25
	# Votes:	3	3	4	4	4
Neutral Leader	Utility:	0.10	0	0.15	0.20	0
	# Votes:	3	4	3	3	3

Table 1: Six hypothetical approval voting scenarios.

Using this formalization, we developed the six scenarios listed in Table 1. Each scenario consists of a set of candidates $C = \{c_1, ..., c_5\}$ with i's utility for each candidate in $[-1.0, 0.25]$. We manipulate two environmental features, including the number of winners ($k = 1, 2, 3$) and the number of missing ballots ($n = 0, 1, 3$). When the final ballots result in a tie, the winner(s) are chosen randomly.

2.1 Experimental Implementation

104 participants were recruited through Mechanical Turk to participate in the voting heuristics study. Participants were paid between \$1.00 and \$8.00 to complete the survey which required them to submit an approval ballot based on hypothetical elections. Participants were paid according to how much utility they accrued over all presented scenarios [11]. In the study, participants were asked to vote in elections based on the scenarios described in Table 1. All participants voted in the single winner scenarios (n=104) and were randomly assigned to be part of a 2-winner (n=50) or 3-winner (n=54) election for the remainder of the study. All participants completed each scenario with 0, 1 or 3 missing ballots. Each election displayed an image showing the candidates, the number of votes cast for each candidate so far, and the amount of money earned for each candidate if they were elected. When voting, subjects could vote for up to all five candidates or abstain. After voting, they saw the election results, including the winners, the amount earned, and the ballots cast by any missing voters (when applicable).

The experiment was designed to test whether or not participants would vote truthfully, manipulate their vote to achieve a higher utility, or vote with another sincere heuristic. We expected that most people would try to vote strategically, but since the situations were cognitively complex with varying degrees of uncertainty, they would not perform all of the necessary computations to identify the maximizing strategy. Instead, we expected that people would use a heuristic strategy.

3 RESULTS

For each scenario, we calculated the expected utility for using *truthful*, *take the X best* and *regret minimization*. We also calculated the maximum expected utility possible, and any votes in V that maximized i's expected utility. We analyzed the behavioral data gathered from the experiment described in Section 2.1 to determine which strategies people used in each scenario.

Our results showed some distinct patterns of voting across all scenarios. The majority of participants in the 1-winner (25.6%) and 2-winner (38.4%) conditions *did not vote using a strategy that maximized expected utility*. In the 3-winner condition, 49.6% used a maximizing strategy. We also found that *as the number of possible winners increased, participants were more likely to vote truthfully*, i.e., for all candidates with positive utility (1-winner: 33.6%, 2-winner: 33.6%, 3-winner: 46.1%). We also found that when participants were not entirely truthful, they tended to use a *take the X best* heuristic (1-winner: 50.6%, 2-winner: 43.8%, 3-winner: 34.4%).

Although people generally voted sincerely, they adjusted their underlying strategy as the numbers of winners changed. In scenarios with trivial utilities or disliked candidates, a significant portion of voters did not vote completely truthfully, and used another strategy such as *take the X best* or *regret minimization*. Voters in these scenarios also tended to vote for a number of candidates equal to the number of winners they were electing, indicating that they chose a heuristic that aligned with the number of winners. However, in the scenarios with dominated preferences or neutral leading candidates, being truthful was the dominant strategy by a wide margin, and there was no relation between the number of candidates voted for and the number of winners.

Our results showed that people were not very sensitive to changes in uncertainty. In scenarios with trivial utilities or a neutral leading candidate, *participants' behavior did not significantly change as the number of missing votes increased, even when this resulted in a non-optimal strategy*. In the 2-winner election with dominated preferences, voters were sensitive to the fact that they had some chance of electing a candidate when there were 3 missing votes. This led to fewer abstentions in that condition. When there was a disliked candidate, an increased number of voters used *regret minimization* in the 3-winner condition (from 5.8% at 0 missing votes to 17.3% at 3 missing votes), indicating that some voters could identify that the underlying optimal strategy changed as the uncertainty increased.

4 FUTURE WORK

While these results provide insights into the use and effectiveness of certain heuristics in approval voting, there are many other voting rules, scenarios and heuristics. It would be interesting to explore heuristics under other voting rules, including those that are known to be computationally complex to manipulate with complete information, such as the single transferable vote (STV). We believe that understanding voting heuristics is important for a more realistic analysis of the voting rules.

REFERENCES

[1] H. Aziz, S. Gaspers, J. Gudmundsson, S. Mackenzie, N. Mattei, and T. Walsh. 2015. Computational Aspects of Multi-Winner Approval Voting. In *Proceedings of the 14th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*.

[2] F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia (Eds.). 2016. *Handbook of Computational Social Choice*. Cambridge University Press.

[3] Victor DeMiguel, Lorenzo Garlappi, and Raman Uppal. 2007. Optimal versus naive diversification: How inefficient is the 1/N portfolio strategy? *The review of Financial studies* 22, 5 (2007), 1915–1953.

[4] Ulle Endriss. 2007. Vote manipulation in the presence of multiple sincere ballots. In *Proc. 11th Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*. 125–134.

[5] Roy Fairstein, Adam Lauz, Reshef Meir, and Kobi Gal. [n. d.]. Modeling People's Voting Behavior with Poll Information.

[6] Piotr Faliszewski, Edith Hemaspaandra, and Lane A. Hemaspaandra. 2010. Using Complexity to Protect Elections. *Commun. ACM* 53, 11 (Nov. 2010), 74–82. https://doi.org/10.1145/1839676.1839696

[7] Gerd Gigerenzer and Daniel G Goldstein. 1996. Reasoning the fast and frugal way: models of bounded rationality. *Psychological Review* 103, 4 (1996), 650.

[8] Gerd Gigerenzer and Peter M Todd. 1999. Fast and frugal heuristics: The adaptive toolbox. In *Simple heuristics that make us smart*. Oxford University Press, 3–34.

[9] D. M. Kilgour. 2010. Approval Balloting for Multi-winner Elections. In *Handbook on Approval Voting*. Springer, Chapter 6.

[10] J.-F. Laslier and M. R. Sanver (Eds.). 2010. *Handbook on Approval Voting*. Springer-Verlag.

[11] W. Mason and S. Suri. 2012. Conducting behavioral research on Amazon's Mechanical Turk. *Behavior Research Methods* 44, 1 (2012), 1–23.

[12] Jaelle Scheuerman, Jason L. Harman, Nicholas Mattei, and Kristen Brent Venable. 2019. Heuristics in Multi-Winner Approval Voting. *CoRR* abs/1905.12104 (2019). arXiv:1905.12104 http://arxiv.org/abs/1905.12104

[13] Jaelle Scheuerman, Jason L. Harman, Nicholas Mattei, and K. Brent Venable. 2019. Heuristics in Multi-Winner Approval Voting. (2019). arXiv:cs.GT/1905.12104

[14] Marcel Zeelenberg. 1999. Anticipated regret, expected feedback and behavioral decision making. *Journal of behavioral decision making* 12, 2 (1999), 93–106.