

Modeling Voters in Multi-Winner Approval Voting

Jaelle Scheuerman, Jason Harman, Nicholas Mattei, K. Brent Venable

Paper previously published in the
Proceedings of the Thirty-Fifth AAAI Conference on Artificial Intelligence (AAAI-21).

Abstract

In many real world situations, collective decisions are made using voting and, in scenarios such as committee or board elections, employing voting rules that return multiple winners. In multi-winner approval voting (AV), an agent submits a ballot consisting of approvals for as many candidates as they wish, and winners are chosen by tallying up the votes and choosing the top- k candidates receiving the most approvals. In many scenarios, an agent may manipulate the ballot they submit in order to achieve a better outcome by voting in a way that does not reflect their true preferences. In complex and uncertain situations, agents may use heuristics instead of incurring the additional effort required to compute the manipulation which most favors them. In this paper, we examine voting behavior in single-winner and multi-winner approval voting scenarios with varying degrees of uncertainty using behavioral data obtained from Mechanical Turk. We find that people generally manipulate their vote to obtain a better outcome, but often do not identify the optimal manipulation. There are a number of predictive models of agent behavior in the social choice and psychology literature that are based on cognitively plausible heuristic strategies. We show that the existing approaches do not adequately model our real-world data. We propose a novel model that takes into account the size of the winning set and human cognitive constraints; and demonstrate that this model is more effective at capturing real-world behaviors in multi-winner approval voting scenarios.

1 Introduction

Computational Social Choice (COMSOC) investigates computational issues surrounding the aggregation of individual preferences for the purposes of collective decision-making [1]. There is a rich body of work that focuses on the computational complexity of manipulating elections under various voting rules [2]. In some cases, it is easy for a voter to compute the optimal manipulation for a given scenario, e.g., an agent casting a ballot for their second ranked candidate, who would win over their third ranked candidate, when their most preferred candidate has no chance in a plurality election.

However, in cases when it is computationally hard for the agent to manipulate an election optimally, it is assumed that voters will report their true preferences rather than trying to strategize [3, 2]. Voting truthfully is just one possible heuristic that voters may use when faced with complex voting scenarios in which the optimal strategy is not easy to compute. For example, a recent study showed that in a plurality election where there was uncertainty over how many more votes were to be cast and a preferred candidate had no chance to win, voters would compromise and vote for the current leader [4]. In other work on allocations, agents have been observed using local strategies to manipulate when faced with a complex computation to find the optimal strategy [5].

In this paper, we examine the accuracy with which these heuristics model decision making in single-winner and multi-winner approval voting elections. We consider scenarios both with and without uncertainty that is represented, in our case, as missing ballots in the election.

Multi-winner approval voting is particularly interesting as a voter may have multiple sincere ballots they can cast, with some being better than others [6, 7]. Understanding the behavioral component of strategic actions plays an important role in agent models as the complexity of the voting scenarios increases. A computational model can capture contextual information about an election to predict voter’s choices in such settings and provide a better understanding of how voters behave at both the individual and group level. Prediction accuracy is an important metric for evaluating models of how humans vote. Accurate models play a fundamental role in providing more reliable forecasts, plausible explanations of voter behavior, and new models for complexity analysis of voting settings [8].

Contributions. We perform a novel experiment to investigate voter behavior in multi-winner approval voting under uncertainty. We identify cognitively plausible heuristics from the psychology literature that may serve as models of voters’ responses, and we provide a quantitative analysis of when and how often these models are deployed in the real world. We evaluate these heuristics against others from the literature on voting in terms of accuracy in predicting voters’ behavior, showing that none of these existing models are very accurate. Finally, we propose a new model which exhibits a significant performance improvement by taking into account contextual information, such as the size of the winning set, as well as human cognitive constraints.

2 Approval Voting

Following [9] and [10] we consider a social choice setting (N, C) where we are given a set $N = \{1, \dots, n\}$ of voting agents and a disjoint set $C = \{c_1, \dots, c_m\}$ of candidates. Each agent $i \in N$ expresses an approval ballot $A_i \subseteq C$ which gives rise to a set of approval ballots $A = \{A_1, \dots, A_n\}$, called a profile. We study the multi-winner approval voting rule that takes as input an instance (A, C, k) and returns a subset of candidates $W \subseteq C$ where $|W| = k$ called the winning set. Approval Voting (AV) finds the set $W \subseteq C$ where $|W| = k$ that maximizes the approval score, i.e., the total weight of approvals, $AV(W) = \sum_{i \in N} |W \cap A_i|$. Informally, the winning set under AV is the set of candidates approved by most voters.

In some cases, it is necessary to use a tie-breaking rule in addition to a voting rule to enforce that the size of W is indeed k . Tie-breaking is an important topic in COMSOC and can have significant effects on the complexity of manipulation of various rules even under idealized models [11, 12, 13]. Typically in the literature, a lexicographic tie-breaking rule is given as a fixed ordering over C , and the winners are selected in this order. However, in this paper, as discussed in [11], we break ties by selecting a winner uniformly at random from the tied set.

Candidates (C):	A	B	C	D	E
Utility (u_i):	0.40	0	0.20	0.01	0

Table 1: Example of the utility of a voter i for 5 candidates A,B,C, D and E.

Similarly to work in the literature on decision heuristics, e.g., Gigerenzer and Goldstein [14] and COMSOC, e.g., Meir et al. [6], we assume that each agent $i \in N$ also has a real valued utility function $u_i : C \rightarrow \mathcal{R}$; Table 1 shows an example with 5 candidates. We also assume that the utility of agent i for a particular set of winning candidates $W \subseteq C$ is $u_i(W) = \sum_{c \in W} u_i(c)$ (with slight abuse of notation). If W is the subset elected by the voting rule we will refer to $u_i(W)$ as agent i ’s utility for the *outcome* of the election.

2.1 Truthfulness in Approval Voting

Studies of approval voting for multi-winner elections span over 40 years [15]. For nearly that entire period, there has been an intense discussion of the strategic aspects of approval balloting [16]. Researchers over the years have made a variety of assumptions and (re)definitions of what makes a particular vote either *truthful* or *strategic* [17, 18, 16]. According to Brams [16]: “A voter votes sincerely if and only if whenever he votes for some candidate, he votes for all candidates preferred to that candidate.” However, even this definition is debated as there can be multiple sincere strategies [18]. This definition is used in recent COMSOC literature to define *Sincere Ballots* [19], of which there may be many for a given scenario. This is a subtle issue as when it is assumed that agents have dichotomous preferences, then multi-winner approval voting is incentive compatible as a complete ballot, i.e., one for all candidates with positive utility, and a sincere ballot are the same. However, if agents have linear preferences over the candidates, as they do in our settings, then there may be multiple sincere ballots that are not complete [20, 6].

We make the following distinctions which are supported by both psychology research on heuristic strategies discussed in Section 4 and discussions in the COMSOC community about voting in approval voting scenarios.

Complete Ballot: a voter submits a ballot approving all candidates for which they have positive utility.

Sincere Ballot: a voter submits a ballot in which if a voter approves a particular candidate, then she also approves all candidates that are preferred to that candidate [19, 6, 16].

As an example of a *Complete Ballot* versus a *Sincere Ballot*, consider a voter having the set of preferences described in 1. Given these preferences, a *Complete Ballot* is $\{A, C, D\}$, while a *Sincere Ballot* could be either $\{A\}$, $\{A, C\}$, $\{A, C, D\}$, or $\{A, B, C, D\}$, or $\{A, B, C, D, E\}$.

3 Related Work

The complexity of manipulation for various types of AV has received considerable attention in the COMSOC literature [1]. If agents act rationally and have full information about the votes of other agents, when agents have *Boolean utilities* AV is strategy-proof. When agents have general utilities, finding a ballot that maximizes the agent’s utility can be computed in polynomial time [20, 6]. For variants of AV, including Proportional AV, Satisfaction AV, and RAV, the complexity of finding utility-maximizing votes ranges in complexity from easy to coNP-complete [9].

Theoretical work in COMSOC often makes worst-case assumptions, e.g., that manipulators have complete information [1]. There are efforts to expand these worst-case assumptions to include the uncertain information and agents that are not perfectly rational, to more closely model the real-world [8]. In Reijngoud and Endriss [21], agent behaviors are measured when agents are given access to poll information. In Meir et al. [22], agents are modeled as behaving in *locally dominant*, i.e., myopic ways. A survey of recent work on issues surrounding strategic voting is given by [23].

There is a growing effort to use simulations and real-world data to test various decision-making models, e.g., [24, 25, 8]. Within the economics and psychology literature, there have been several studies of approval voting and the behavior of voters. Perhaps the most interesting and relevant to our work are the studies of [26], which focus on the elections of various professional societies where approval balloting is used. Regenwetter et al. [26] find that many voters use a *plurality heuristic* when voting in AV elections, i.e., they vote

as if they are in a plurality election, selecting only their most preferred candidate. Zou et al. [27] use approval voting data from Doodle and conclude that the type of poll has a significant effect on voter behavior. In both of these works, only AV with a single winner was investigated, and both works relied on real-world elections where it was not possible to tease out the relationship between the environment and decisions.

Three recent papers address strategic voting under the plurality rule, where agents are making decisions in uncertain environments. First, Tyszler and Schram [28] studied the voting behavior of agents under the plurality rule with three options. They find that the amount of information available to the voters affects the decision on whether or not to vote strategically and that in many cases, the strategic decisions do not affect the outcome of the plurality vote. Second, in Tal et al. [4] an online system is presented where participants vote for cash payments in a number of settings using the plurality rule under uncertainty. They find that most participants do not engage in strategic voting unless there is a clear way to benefit. Indeed, most voters were lazy, and if they did vote strategically, they would do a one-step look-ahead or perform the best response myopically. Finally, in Fairstein et al. [29], a comprehensive study using both past datasets and newly collected ones examines the actual behavior of agents in multiple settings with uncertainty versus behavior that is predicted by a number of behavioral and heuristic models. The paper proposes a novel model of user voting behavior in these uncertain settings called *attainability utility*, where agents consider how much utility they would gain versus the likelihood of particular candidates winning given an uncertain poll. They conclude that the attainable utility model is able to explain the behavior seen in the experimental studies better than existing models and even perform near the level of state-of-the-art machine learning algorithms in modeling users' actual behavior.

We expand upon these works on plurality to consider heuristics in the more complex setting of multi-winner approval voting with uncertainty. We show that, in this context, simple heuristics are inadequate to capture voting behavior and applying the attainability-utility model directly results in limited predictive capabilities. We then extend the attainability-utility model to incorporate more information of the voting scenarios while maintaining cognitive plausibility, proving that this is key to greatly enhance predictive accuracy.

4 Heuristics for Approval Voting

We consider three heuristics inspired by the literature in cognitive science and COMSOC, including *Complete*, *Take the X best*, and *Attainability-Utility (AU)*. We study these heuristics experimentally with behavioral data and evaluate them in terms of their predictive accuracy as a model of voter behavior. We show, in Section 5 and Section 6, that though these strategies are frequently employed by real-world decision makers, none of them are able to forecast human votes with reasonable accuracy. To address this, we present a novel heuristic model for multi-winner approval voting, *Attainability-Utility with Threshold (AUT)*.

4.1 Cognitive Heuristics

We start by considering two simple heuristics which ignore the current voting scenario information and use only the voter's utility for each candidate. Experimental data originally collected in [30, 31] shows that these heuristics are indeed adopted by voters in the context of single winner and multi-winner approval voting. Examples of ballots corresponding to these heuristics are shown in Table 3 and 4.

Complete. An agent approves all candidates for which they have positive utility; corresponding to a Complete Ballot.

Take the X best. When an agent i votes with the *Take the X best* heuristic, they vote for a subset of the complete ballot. First, they order the list of candidates by the utility value, $c_j > c_h$ when $u_i(c_j) > u_i(c_h)$. The agent will then vote for the top- X candidates in this ordering with $1 \leq X \leq m$. X could be calculated using a magnitude cut off or a proportional difference between preferences [32], or setting X to be the number of winners in the election ($X = k$). In this paper, we examine situations where the candidates’ utility values are not tied, so no tie-breaking rule is required.

In the example shown in Table 1, the *Complete* ballot is $\{A, B, C, D\}$ and the *Take X best* with $X = 2$ is $\{D, B\}$. Observe that, for every setting of X in *Take the X best* we have an instance of a *Sincere Ballot* as described by Brams [16].

4.2 Attainability-Utility (AU) Heuristic

Fairstein et al. [29] present a heuristic for plurality that explains the voting behavior of individual agents as a trade-off between attainability and utility. They define attainability as the likelihood that a candidate will win an election, given access to uncertain poll information about the current state of the election. The formula for computing attainability was first introduced by Bowman et al. [33] for voting over a set of binary issues and was then extended by Fairstein et al. [29] to the case of m candidates.

Let us first review attainability in the case of plurality. Given a vector containing the current count of ballots in favor of each candidate, s , as well as the number of currently known ballots r , the attainability of candidate c_j is defined as:

$$A_\beta(c_j, s, r) = \frac{1}{\pi} \tan^{-1} \left(\beta \left(\frac{s(c_j)}{r} - \frac{1}{m} \right) \right) + \frac{1}{2} \quad (1)$$

where, $s(c_j)$ is the number of ballots for candidate c_j and β is a parameter modeling how a particular voter perceives candidate attainability. In Figure 1 we depict three different ways in which the attainability values for a candidate vary with settings to β in an election with 5 candidates as a function of the uncertainty remaining in the election. Lower values of β mean the candidate seems more attainable when there is more uncertainty, with attainability increasing as more ballots are known. Higher values of β scale the cotangent curve such that candidates seem very unattainable when there is high uncertainty, but more attainable as the uncertainty decreases.

We expand upon the work of Fairstein et al. [29] by (1) modifying AU to work in a multi-winner approval setting and (2) add a threshold parameter τ to model the minimum AU score that a candidate must have to be approved by the voter.

Translating the AU model to the multi-winner approval voting setting is non-trivial. AU in Equation 1 is only defined for single winners, and so we must modify the model to account for winning sets.

In the original model by Bowman et al. [33] where $1/2$ of the ballots are required to attain the passage of binary proposals, they used the quantity $(s(c_j)/r - 1/2)$ to capture the difference between the proportion of the current ballots for a proposal and the score necessary to win. In the extension to plurality elections by Fairstein et al. [29], the $1/2$ was changed to $1/m$, to capture the idea that a candidate is more attainable if they are closer to having a plurality of the total current ballots, i.e., $(s(c_j)/r - 1/m)$.

In both Bowman et al. [33] and Fairstein et al. [29], the voter is intuitively comparing the proportion of current ballots for candidate c_j to a uniform probability assumption over all possible outcomes.

To apply this notion to winning sets of k winners, we modify the $1/m$ to be $1/(mk)$. Hence, when we have only one winner, $k = 1$, we recover the model of Fairstein et al. [29], and as

the number of winners grows, the denominator mk becomes larger, requiring c_j to have a smaller percentage of ballots to appear attainable. Intuitively, this represents the idea that a desired candidate is more attainable as the size of the winning set increases as they need fewer total ballots cast.




Figure 1: $A_\beta(c_j, s, r)$ for different values of β in any single-winner approval election with five candidates.

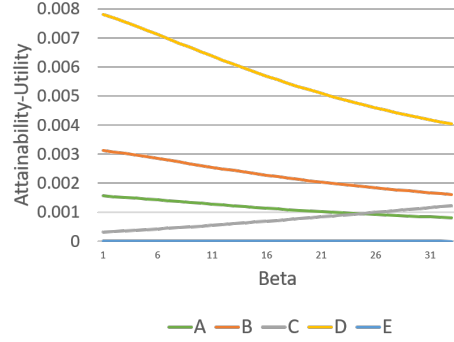


Figure 2: AU score of the five candidates from Scenario B for different values of β when $\alpha = 1$.




Figure 3: Visualization of the ballots cast by 54 subjects who voted in the single-winner and 3-winner elections with 0 missing votes.

In the plurality setting, the number of current ballots r coincides with the current total number of approvals. This is, of course, not the case in approval voting, where people are allowed to approve up to m candidates. Thus, in our setting, r represents the total number of approvals contained in all ballots that have been currently counted. For example, if we consider an example with five candidates, $m = 5$, A, B, C, D , and E , and $n = 5$ voters where two voters have submitted their ballots consisting of $\{A\}$ and $\{A, B, C\}$, we would have $r = 4$. We calculate the attainability in an approval election with missing ballots to be the mean attainability over all possible ways in which the remaining ballots can be cast. We model this by considering the set of all possible total approval counts in the election: $T = \{t \mid r \leq t \leq r + \hat{n}m\}$, where r is the number of approvals made so far, \hat{n} is the number of missing ballots and $\hat{n}m$ represents the number of total possible approvals. In the running example, $\hat{n} = 3$, $\hat{n}m = 15$, and $T = \{4, 5, 6, 7, \dots, 19\}$. We can now define the attainability of candidate c_j in a multi-winner approval voting setting as:

$$A_\beta(c_j, s, T) = \frac{\sum_{t \in T} \frac{1}{\pi} \tan^{-1} \left(\beta \left(\frac{s(c_j)}{t} - \frac{1}{mk} \right) \right) + \frac{1}{2}}{|T|} \quad (2)$$

The utility of a subset of candidates $u(b)$, $b \in \mathcal{P}(C)$ is the sum of the utilities of each candidate c_j in b . Formally, with a slight abuse of notation, $u(b) = \sum_{c_j \in b} u(c_j)$. Slightly abusing notation, we define the attainability of a set of candidates b , as the product of $A_\beta(c_j, s, T)$ for all $c_j \in b$ and $1 - A_\beta(c_j, s, T)$ for all $c_j \notin b$, that is:

$$A_\beta(b, s, T) = \prod_{c_j \in b} A_\beta(c_j, s, T) \cdot \prod_{c_j \notin b} (1 - A_\beta(c_j, s, T)) \quad (3)$$

We then find a candidate subset b that maximize the trade-off between attainability and utility by incorporating an additional parameter $\alpha \in [0, 2]$:

$$M_{\alpha, \beta}^{AU}(s, u, T) = \arg \max_{b \in \mathcal{P}(C)} ((\epsilon + u(b))^\alpha \cdot (A_\beta(b, s, T))^{2-\alpha}) \quad (4)$$

Note that, when α is 0 then the voter only considers attainability, but when α is 2 she only considers utility. A small constant value, ϵ , is also added to account for situations where the left hand component of the equation would result in 0^0 , i.e. when candidates have 0 utility and $\alpha = 0$.

Note that, when α is 0 then the voter only considers attainability, but when α is 2 she only considers utility. A small constant value, ϵ , is also added to account for situations where the left hand component of the equation would result in 0^0 , i.e. when candidates have 0 utility and $\alpha = 0$.

In Fairstein et al. [29], voting behavior was predicted for a plurality election, where voters may approve of only one candidate at a time. In multi-winner AV, voters approve as many candidates as they wish, and the top k candidates win the election. When the AU model is applied in the approval setting, voters cast the ballot described by $M_{\alpha, \beta}^{AU}$ as defined in Equation 4. Table 2 shows which ballot is selected by $M_{\alpha, \beta}^{AU}$ for various settings of $\alpha \in \{0, 1, 2\}$ and $\beta \in \{1, 2, \dots, 32\}$ on the single winner approval voting election scenario (s, u, T) when s and u are described as in Table 3 and $T = \{r\}$ (that is, there are no additional missing ballots). These β values were chosen to represent the broad range of shapes that the attainability function can represent as shown in Figure 1.

$M_{\alpha, \beta}^{AU}$ ballot	α		
	0	1	2
D	[1..32]		
A,B,C,D,E		[1]	
A,B,D,E		[2..8]	
D,E		[9..21]	
B,D,E		[22..32]	
A,B,C,E			[1..32]

Table 2: Ballots predicted for Scenario A using the attainability-utility model (first column) and corresponding β ranges (all other columns) for $\alpha \in \{0, 1, 2\}$.

Observing Table 2 we note how, in the AU model, $\alpha = 0$ predicts a ballot containing only the leading candidate, while $\alpha = 2$ predicts the complete ballot of all candidates with positive utility. When $\alpha = 1$, different ballots are chosen, depending on the value of β .

Behavior that prioritizes higher utility candidates without voting for leading candidates is not predicted by this model.

It is worth noting that the prediction space for Scenario A, shown in Table 2, predicts only a single Take the X Best strategy (A,B,C,E) and does not predict the optimal ballot (E for 1 winner, 0 missing votes as seen in Table 5). This is counter to the observed behavior where most people used different Take the X Best strategies (see Figure 3). Hence, it is clear that when moving into an approval voting setting, the AU heuristic is inadequate as a model of voter behavior.

4.3 The Attainability-Utility with Threshold (AUT) Heuristic

The trade-offs between attainability and utility proposed by Fairstein et al. [29] are common among cognitive models of decision making [34, 35] and are cognitively plausible for settings involving a single winner. When applied to multi-winner contexts, the formalization of the trade-offs entails calculating a value for every possible missing ballot (see Formula 4) which becomes increasingly implausible and cognitively demanding as m increases. Hence, we introduce a new model which instead assumes that participants calculate the AU value for each individual candidate, and then choose those candidates that surpass a threshold. This change makes the required mental computations cognitively plausible and is consistent with the adaptive toolbox theory of human decision making [36, 37].

More formally, rather than considering the attainability-utility score (AU score) of each $b \in \mathcal{P}(C)$, we calculate the AU score for each individual candidate c_j :

$$AU_{\alpha,\beta}(c_j, s, u, T) = (\epsilon + u(c_j))^\alpha \cdot (A_\beta(c_j, s, T))^{2-\alpha} \quad (5)$$

We introduce a threshold τ to set the minimum AU score that will lead to a voter approving c_j . The introduction of the threshold to the model was inspired by concepts in cognitive psychology [38]. Decision thresholds are parameters that vary predictably based on factors in the environment (such as uncertainty). We expected that using thresholds in our AUT model would outperform AU because it more closely models the cognitive process used to vote in the current paradigm. In approval voting scenarios, the model predicts that a voter will submit a ballot approving of all candidates with an AU score greater than τ . We will refer to this heuristic as *attainability-utility with threshold (AUT)*, and define the corresponding ballot for threshold τ as:

$$M_{\alpha,\beta}^{AUT}(s, u, T, \tau) = \{c_j \in C \mid ((\epsilon + u(c_j))^\alpha \cdot (A_\beta(c_j, s, T))^{2-\alpha}) \geq \tau\} \quad (6)$$

Figure 2 shows the AU score of each of the five candidates in Scenario B for different values of β when $\alpha = 1$. We can see that if $\beta = 2$, when $\tau = 0.007$, and $T = \{r\}$ the model predicts voting for only D , while when $\tau = 0.001$, then the model predicts voting for $\{A, B, D\}$. Note that for low values of beta, $\{A, B, D\}$ can be predicted by $\tau = 0.001$, while high values of beta will lead to a prediction of $\{A, B, C\}$. This exemplifies the flexibility AUT has in generating different ballots. Of special note, is that in single-winner plurality elections where voters can only vote for one candidate, the AUT model simplifies to the AU model, accounting for human voting behavior in both plurality and approval voting.

5 Behavioral Data

Our behavioral study aimed at investigating approval voting heuristics and included 104 participants recruited through Mechanical Turk. Participants were asked to cast ballots in a voting game where they were presented with a number of multi-winner approval voting

Candidate:	A	B	C	D	E
Utility:	0.05	0.10	0.01	0	0.25
# Votes:	3	3	3	4	3

Table 3: Scenario A design, including candidates, utilities and votes. *Heuristic votes*: Complete: [A,B,C,E], Take 1 Best: [E], Take 2 Best: [E,B], Take 3 Best: [E,B,A]

Candidate:	A	B	C	D	E
Utility:	0.05	0.10	0.01	0.25	0
# Votes:	3	3	4	3	3

Table 4: Scenario B design, including candidates, utilities and votes. *Heuristic votes*: Complete: [A,B,C,D], Take 1 Best: [D], Take 2 Best: [D,B], Take 3 Best: [D,B,A]

scenarios with a monetary value that was paid out when certain candidates won. Participants were paid \$1.00 to participate in the study and received a bonus of no more than \$8.00, which was determined by the outcome of hypothetical elections. All participants voted in the single winner scenarios (n=104). Participants were then randomly assigned to be part of a 2-winner (n=50) or 3-winner (n=54) election for the remainder of the study. More information about this study can be found in Scheuerman et al. [30, 31].

In this analysis, we consider two scenarios in particular, Scenario A and Scenario B, shown in detail in Tables 3 and 4. In these scenarios, the participant faces a situation where each candidate generates a different amount of utility $u \in \{0, 0.1, 0.05, 0.10, 0.25\}$, paid as a reward if that candidate is in the winning set. In both scenarios, none of the participant’s high utility candidates are leading the election, but all are within 1 approval of being tied for the lead. We observe how people respond, particularly considering if they cast a *Complete Ballot* or a *Sincere Ballot*, and how many sincere candidates they choose to approve. The experimental data consists of responses to Scenarios A and B in 9 different conditions: where there are 1, 2 and 3-winners with 0, 1 or 3 missing ballots. Hence we can examine how varying both the number of winners as well as the amount of uncertainty affects voter behavior.

Using χ^2 analysis, we compared the subject’s responses from both scenarios and found that there was no significant difference in the distribution of responses. Even though the values of the candidates and the current leader changed, there was no evidence that the subjects voted differently in Scenarios A and B. Within each scenario, we compared the responses between each missing ballot condition, and found no significant difference in how people voted as the number of missing ballots increased. However, when comparing the responses between conditions with differing numbers of candidates, significant differences ($P < 0.005$) were found when comparing the strategies used by voters in the single-winner and 3-winner conditions, and the 2 and 3-winner conditions. In general, when voting in single-winner and 2-winner elections, participants cast a ballot for 2 or 3 candidates (single-winner: 57.9%, 2-winner: 70.7%) more often than other strategies. When participants voted in the 3-winner election, they usually cast a ballot for 3 candidates (61.7%). Figure 3 shows how an individual voter’s ballots changed as the number of winners increased from 1 to 3.

We found that, in general, the majority of people vote sincerely with a *Take the X Best* strategy: 78.8% in Scenario A and 77.8% in Scenario B. However, the value of X used by the voters depended on the individual and the size of the winning set (see Figure 3).

n	# winners (k)		
	1	2	3
0	Take 1 A: 0.12 B: 0.13	Take 1 A: 0.22 B: 0.26	Take 2 A: 0.31 B: 0.36
1	Take 1 A: 0.11 B: 0.12	Take 2 A: 0.21 B: 0.22	Take 2 A: 0.30 B: 0.31
3	Take 1 A: 0.11 B: 0.11	Take 2 A: 0.20 B: 0.21	Take 2 A: 0.29 B: 0.29

Table 5: The maximum expected utility for Scenarios A and B, and the heuristic that achieves this optimal outcome. n represents the number of missing ballots.

6 Evaluation of Heuristics as Models

Using our experimental data, we evaluated five approaches for predicting voter behavior over single-winner, two-winner, and three-winner conditions. First, we consider the effectiveness of a model that predicts optimal votes, i.e., the one that always maximizes the expected utility. We evaluate four other predictive models, each of which corresponds to the heuristics described in Section 4, namely, *Complete*, *Take X Best* with X equal to the size k of the winning set, AU and AUT.

6.1 Optimal Baseline

As a baseline, we assume that people vote optimally, approving the ballot that maximizes their expected utility, which varies with the number of winners and missing ballots. Given the number of winners $k \in \{1, 2, 3\}$ and the number of missing ballots $\hat{n} \in \{0, 1, 3\}$ the expected utility E of ballot $b \in \mathcal{P}(C)$ is: $E(b, k, \hat{n}) = \sum_{c_j \in b} p(c_j, k, \hat{n})u(c_j)$. Assuming that all potential missing ballots are equally likely, $p(c_j, k, n)$ refers to the probability that candidate c_j is elected, given that there are k possible winners, \hat{n} missing votes and that ties are broken uniformly at random.

For each possible number of winners and the number of missing ballots, the Optimal Baseline, selects the ballot maximizing the expected utility:

$$M^{Opt}(k, \hat{n}) = \arg \max_{b \in \mathcal{P}(C)} E(b, k, \hat{n}) \quad (7)$$

The optimal ballots and the corresponding maximum expected values for Scenarios A and B can be found in Table 5. Note that the Optimal Baseline in each of these scenarios and conditions corresponds to a variant of the *Sincere Ballot* of *Take the X best*. For example, in both scenarios, *Take the 1 Best* is optimal when there a single winner and no missing ballots. When there are 3 winners and no missing ballots, it is optimal to *Take the 2 Best*.

6.2 Fitting Heuristics to Data

In addition to the Optimal baseline, we fit each of the four heuristics to the data and tested their accuracy as predictive models of voter behavior. In addition to testing *Complete* and *Take the X Best*, where X is set to be the number of winners $X = k$, we trained *AU* and *AUT* on the data collected from Scenarios A and B for each individual. The parameters and ranges we considered are:

Attainability-Utility (AU). Using a grid search, we fit $\alpha \in \{0, 1, 2\}$ and $\beta \in \{1, 2, \dots, 32\}$.

Attainability-Utility with Threshold (AUT). Since the behavioral results did not show a tendency to vote based on attainability alone, i.e., people rarely voted for only leading candidates with no utility or low utility, we choose to set $\alpha = 1$ and fit only β and τ . Using a grid search, we tested values for $\beta \in \{1, 2, \dots, 32\}$ and $\tau \in \{0, 0.0005, \dots, 0.10\}$.

We train the parameters of AU and AUT as follows. For each individual voter, when we fix k , we have 6 conditions to use for training and testing our model. These 6 conditions correspond to the number of possible missing ballots, $\hat{n} \in \{0, 1, 3\}$, for Scenarios A and B. We use five of these observations to train the parameters of the AU or AUT model and predict the sixth. Using a leave-one-out methodology, we do this for all possible splits of the data. We compute the accuracy over these six splits for each individual. We then average this accuracy over all individual voters for each of the $k \in \{1, 2, 3\}$ winning set size conditions. Each of the models were evaluated for their accuracy in predicting voter behavior in Scenarios A and B. The *Optimal*, *Complete*, and *Take k Best* models are deterministic, depending on the scenario, the number of winners and the number of missing ballots. The average prediction accuracy and standard deviation over the six responses for each of the winning set size conditions are reported in Table 6 and shown in Figure 4.

6.3 Evaluation Results and Discussion

The model evaluation results show that a model of optimal behavior using expected utilities is not a good representation of voter behavior and supports the idea that people do not take the time to perform the calculations necessary to identify optimal solutions. We also found that people tend to not vote a *Complete Ballot*. The experimental data indicated that voters tend to approve more candidates as the size of the winning set increases. Thus, we conjectured that the *Take k Best* model, predicting that people would vote for the number of top candidates that was equal to the number of winners (k) would perform well. However, this was refuted by our experimental results which showed only a modest improvement in performance in the 3-winner condition.

	1 winner	2 winner	3 winner
Optimal:	24.7% (3.9%)	19.7% (9.1%)	14.2% (5.8%)
AU:	13.8% (2.6%)	16.3% (2.0%)	7.4% (2.0%)
Complete:	15.1% (1.2%)	13.7% (2.7%)	9.0% (2.0%)
Take k best:	22.8% (4.1%)	30.7% (2.2%)	58.3% (4.0%)
AUT:	58.1% (11.6%)	68.0% (1.8%)	67.9% (3.4%)

Table 6: Mean prediction accuracy (standard deviation) for each model across conditions.

Next, we tested whether or not voters make a trade-off between the utility and attainability of the candidates. We first tried this by using the attainability-utility heuristic (AU), which considers the trade-off of attainability and utility for all subsets of candidates. We found that this also does not describe the voting behavior in multi-winner approval voting election. The poor performance of the AU model can likely be attributed to how the attainability-utility is calculated for every possible subset of candidates, rather than individual candidates. The attainability-utility of ballots containing all utility generating candidates is higher than ballots that contain only a subset of them. Thus, *Complete Ballots* are predicted much more often than sincere *Take the X Best* behavior (See Table 2).

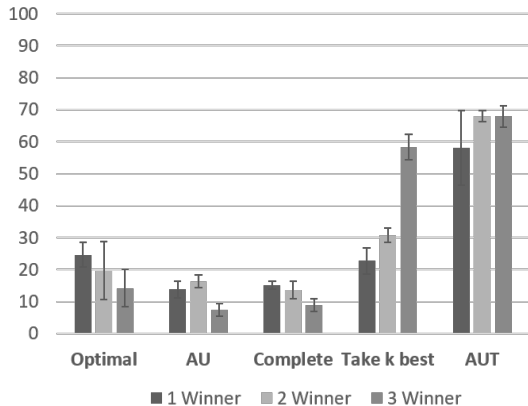


Figure 4: Mean and standard deviation of prediction accuracy for each model.

Finally, we addressed this with a novel model of the attainability-utility heuristic with a threshold for approval voting. This model takes into account human cognitive constraints by generating an attainability-utility score for each candidate, rather than for every possible subset, and choosing to vote for only the candidates above a certain threshold. This model performed the best of all the models evaluated. On average, our AUT model performed 80.7% better than the AU model. For single-winner voting, it performed 57.5% better than the second best model (*Optimal*). In 2-winner elections, it performed 55.0% better than the second best model (*Take k Best*). Finally, in the 3-winner elections, it performed 14.3% better than the second best model (*Take k Best*). Hence, our cognitively inspired heuristic model accurately predicts the voting behavior of agents in multi-winner approval voting environments more often than any existing models in the literature.

7 Conclusion and Future Directions

We have evaluated several heuristics as models of voter behavior in multi-winner approval voting. We found that our novel model of the attainability-utility heuristic with threshold provided the best predictive power of the models tested. This model simulates voters ranking each candidate based on a trade-off between attainability and utility, and then approving candidates ranked above a threshold.

To enhance prediction and understanding, a full taxonomy of (internal) cognitive strategies and capabilities along with (external) voting contexts and elements of uncertainty is required to fully explore the interaction between the two. Our AUT model, which describes the trade-off between attainability and utility at the candidate level while generalizing across different conditions of uncertainty, number of winners, and voting rules (i.e. approval and plurality), is an initial step in this direction. Going forward, we will explore how cognitively plausible models like AUT can be used to develop hybrid machine learning models that leverage models of cognitively plausible heuristics to predict voter behavior with even greater accuracy.

8 Acknowledgements

Nicholas Mattei and K. Brent Venable are supported by NSF Award IIS-2007955.

References

- [1] F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.
- [2] P. Faliszewski and A. D. Procaccia. AI’s war on manipulation: Are we winning? *AI Magazine*, 31(4):53–64, 2010.
- [3] Piotr Faliszewski, Edith Hemaspaandra, and Lane A. Hemaspaandra. Using complexity to protect elections. *Commun. ACM*, 53(11):74–82, November 2010. ISSN 0001-0782. doi: 10.1145/1839676.1839696. URL <http://doi.acm.org/10.1145/1839676.1839696>.
- [4] M. Tal, R. Meir, and Y. Gal. A study of human behavior in online voting. In *Proc. of the 14th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 665–673, 2015.
- [5] T. Mennle, M. Weiss, B. Philipp, and S. Seuken. The power of local manipulation strategies in assignment mechanisms. In *Proc. of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*, 2015.
- [6] R. Meir, A. D. Procaccia, J. S. Rosenschein, and A. Zohar. Complexity of strategic behavior in multi-winner elections. *Journal of Artificial Intelligence Research*, 33:149–178, 2008.
- [7] Toby Walsh. Uncertainty in preference elicitation and aggregation. In *AAAI*, volume 7, pages 3–8, 2007.
- [8] Nicholas Mattei. Closing the loop: Bringing humans into empirical computational social choice and preference reasoning. In *Proc. of the 29th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 5169–5173, 2020.
- [9] H. Aziz, S. Gaspers, J. Gudmundsson, S. Mackenzie, N. Mattei, and T. Walsh. Computational aspects of multi-winner approval voting. In *Proc. of the 14th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 2015.
- [10] D. M. Kilgour. Approval balloting for multi-winner elections. In *Handbook on Approval Voting*, chapter 6. Springer, 2010.
- [11] H. Aziz, S. Gaspers, N. Mattei, N. Narodytska, and T. Walsh. Ties matter: Complexity of manipulation when tie-breaking with a random vote. In *Proc. of the 27th AAAI Conference on Artificial Intelligence (AAAI)*, pages 74–80, 2013.
- [12] N. Mattei, N. Narodytska, and T. Walsh. How hard is it to control an election by breaking ties? In *Proc. of the 21st European Conference on Artificial Intelligence (ECAI)*, pages 1067–1068, 2014.
- [13] S. Obraztsova, E. Elkind, and N. Hazan. Ties matter: Complexity of voting manipulation revisited. In *Proc. of the 10th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 71–78, 2011.
- [14] Gerd Gigerenzer and Daniel G Goldstein. Reasoning the fast and frugal way: models of bounded rationality. *Psychological Review*, 103(4):650, 1996.
- [15] Steven J Brams. Approval voting in multicandidate elections. *Policy Studies Journal*, 9(1):102, 1980.

- [16] Steven J Brams. Strategic information and voting behavior. *Society*, 19(6):4–11, 1982.
- [17] J.-F. Laslier and M. R. Sanver, editors. *Handbook on Approval Voting*. Studies in Choice and Welfare. Springer-Verlag, 2010.
- [18] R. G. Niemi. The problem of strategic behavior under approval voting. *The American Political Science Review*, 78(4):952–958, 1984.
- [19] Ulle Endriss. Vote manipulation in the presence of multiple sincere ballots. In *Proc. 11th Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*, pages 125–134, 2007.
- [20] R. Meir, A. D. Procaccia, and J. S. Rosenschein. A broader picture of the complexity of strategic behavior in multi-winner elections. In *Proc. of the 7th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 991–998, 2008.
- [21] A. Reijngoud and U. Endriss. Voter response to iterated poll information. In *Proc. of the 11th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 635–644, 2012.
- [22] R. Meir, O. Lev, and J.S. Rosenschein. A local-dominance theory of voting equilibria. In *Proc. of the 15th ACM conference on Economics and Computation (ACM EC)*, pages 313–330, 2014.
- [23] R. Meir. *Strategic Voting*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool Publishers, 2018.
- [24] N. Mattei and T. Walsh. A PREFLIB.ORG Retrospective: Lessons Learned and New Directions. In U. Endriss, editor, *Trends in Computational Social Choice*, chapter 15, pages 289–309. AI Access, 2017.
- [25] N. Mattei and T. Walsh. PrefLib: A library for preferences, [HTTP://WWW.PREFLIB.ORG](http://www.preflib.org). In *Proc. of the 3rd International Conference on Algorithmic Decision Theory (ADT)*, 2013.
- [26] M. Regenwetter, M.-H. R. Ho, and I. Tsetlin. Sophisticated approval voting, ignorance priors, and plurality heuristics: A behavioral social choice analysis in a Thurstonian framework. *Psychological Review*, 114(4):994–1014, 2007.
- [27] J. Zou, R. Meir, and D. Parkes. Strategic voting behavior in Doodle polls. In *Proc. of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*, pages 464–472, 2015.
- [28] M. Tyszler and A. Schram. Information and strategic voting. *Experimental Economics*, 19(2):360–381, 2016.
- [29] R. Fairstein, A. Lauz, R. Meir, and K. Gal. Modeling people’s voting behavior with poll information. In *Proc. of the 18th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 2019.
- [30] Jaelle Scheuerman, Jason Harman, Nicholas Mattei, and K. Brent Venable. Heuristics in multi-winner approval voting. Behavioral EC, Phoenix, AZ, 2019.
- [31] Jaelle Scheuerman, Jason L. Harman, Nicholas Mattei, and Kristen Brent Venable. Heuristic strategies in uncertain approval voting environments. *CoRR*, abs/1912.00011, 2019. URL <http://arxiv.org/abs/1912.00011>.

- [32] Eduard Brandstätter, Gerd Gigerenzer, and Ralph Hertwig. The priority heuristic: making choices without trade-offs. *Psychological review*, 113(2):409, 2006.
- [33] Clark Bowman, Jonathan K. Hodge, and Ada Yu. The potential of iterative voting to solve the separability problem in referendum elections. *Theory and Decision*, 77(1): 111–124, June 2014. ISSN 1573-7187. doi: 10.1007/s11238-013-9383-2. URL <https://doi.org/10.1007/s11238-013-9383-2>.
- [34] Amos Tversky and Daniel Kahneman. Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157):1124–1131, 1974.
- [35] Claudia Gonzalez-Vallejo. Making trade-offs: A probabilistic and context-sensitive model of choice behavior. *Psychological Review*, 109(1):137, 2002.
- [36] John W Payne, John William Payne, James R Bettman, and Eric J Johnson. *The adaptive decision maker*. Cambridge university press, 1993.
- [37] Gerd Gigerenzer and Reinhard Selten. *Bounded rationality: The adaptive toolbox*. MIT press, 2002.
- [38] Jerome R Busemeyer and James T Townsend. Decision field theory: a dynamic-cognitive approach to decision making in an uncertain environment. *Psychological review*, 100(3):432, 1993.

Jaelle Scheuerman
U.S. Naval Research Laboratory
Stennis Space Center, MS
Email: jaelle.scheuerman@nrlssc.navy.mil

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

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

K. Brent Venable
IHMC & University of West Florida
Pensacola, FL
Email: bvenable@ihmc.org