

# Modeling Voters in Multi-Winner Approval Voting

Jaelle Scheuerman,<sup>1</sup> Jason Harman,<sup>1,2</sup> Nicholas Mattei,<sup>3</sup> K. Brent Venable<sup>4</sup>

<sup>1</sup> U.S. Naval Research Laboratory, Stennis Space Center, MS, USA

<sup>2</sup> Louisiana State University, Baton Rouge, LA, USA

<sup>3</sup> Tulane University, New Orleans, LA, USA

<sup>4</sup> University of West Florida and Institute for Human and Machine Cognition, Pensacola, FL, USA  
jaelle.scheuerman@nrlssc.navy.mil, jharman@lsu.edu, nsmattei@tulane.edu, bvenable@ihmc.org

## 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 (Brandt et al. 2016). There is a rich body of work that focuses on the computational complexity of manipulating elections under various voting rules (Faliszewski and Procaccia 2010). 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 (Faliszewski, Hemaspaandra, and Hemaspaandra 2010; Faliszewski and Procaccia 2010). 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 (Tal, Meir, and Gal 2015). 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 (Mennle et al. 2015).

In this paper, we examine both 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 (Meir et al. 2008; Walsh 2007). 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 (Mattei 2020).

**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.

Candidates ( $C$ ):	A	B	C	D	E
Utility ( $u_i$ ):	0.05	0.10	0.01	0.25	0

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

## 2 Approval Voting

Following (Aziz et al. 2015) and (Kilgour 2010) 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 (Aziz et al. 2013; Mattei, Narodytska, and Walsh. 2014; Obraztsova, Elkind, and Hazon 2011). 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 (Aziz et al. 2013), we break ties by selecting a winner uniformly at random from the tied set.

Similarly to work in the literature on decision heuristics, e.g., Gigerenzer and Goldstein (1996) and COMSOC, e.g., Meir et al. (2008), 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.

### Truthfulness in Approval Voting

Studies of approval voting for multi-winner elections span nearly 40 years (Brams 1980). For nearly that entire period, there has been an intense discussion of the strategic aspects of approval balloting (Brams 1982). Researchers over the years have made a variety of assumptions and (re)definitions of what makes a particular vote either *truthful* or *strategic* (Laslier and Sanver 2010; Niemi 1984; Brams 1982). According to Brams (1982): “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 (Niemi 1984). This definition is used in recent COMSOC literature to define *Sincere Ballots* (Endriss 2007), 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 (Meir, Procaccia, and Rosenschein 2008; Meir et al. 2008).

We make the following distinctions which are supported by both psychology research on heuristic strategies discussed in Section ?? 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 (Endriss 2007; Meir et al. 2008; Brams 1982). Intuitively, this is an assumption of monotonicity over the preferences and captures many of the votes one would cast in the *take the X best* heuristic we will see later.

As an example of a *Complete Ballot* versus a *Sincere Ballot*, consider a voter having the set of preferences  $[A = 0.4, B = 0, C = 0.2, D = 0.01]$ . Given these preferences, a *Complete Ballot* is  $[A, C, D]$ , while a *Sincere Ballot* could be either  $[A]$ ,  $[A, C]$ ,  $[A, C, D]$ , or  $[A, B, C, D]$ .

## 3 Related Work

The complexity of manipulation for various types of AV has received considerable attention in the COMSOC literature (Brandt et al. 2016). 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 vote that maximizes the agent’s utility can be computed in polynomial time (Meir, Procaccia, and Rosenschein 2008; Meir et al. 2008). 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 (Aziz et al. 2015).

Theoretical work in COMSOC often makes worst-case assumptions, e.g., that manipulators have complete information (Brandt et al. 2016). 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 (Mattei 2020). In Reijngoud and Endriss (2012), agent behaviors are measured when agents are given access to poll information. In Meir, Lev, and Rosenschein (2014), 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 (Meir 2018).

There is a growing effort to use simulations and real-world data to test various decision-making models, e.g., (Mattei and Walsh 2017, 2013; Mattei 2020). 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 (Regenwetter, Ho, and Tsetlin 2007), which focus on the elections of various professional societies where approval balloting. Regenwetter, Ho, and Tsetlin (2007) 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, Meir, and Parkes (2015) 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 (2016) 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, Meir, and Gal (2015) 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. (2019), 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 *attainable 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 recent works on plurality to consider heuristics in the significantly 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 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 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

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

Table 2: 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 3: 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]

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.; Scheuerman et al. (2019a; 2019b).

In this analysis, we consider two scenarios in particular, Scenario A and Scenario B, shown in detail in Tables 2 and 3. 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  $\chi^2$  analysis, we examined the responses from both scenarios across all conditions and found no significant difference in the distribution of responses between each scenario. This means that even though the values of the candidates and the current leader changed, the voters behaved largely the same in both Scenarios A and B. Within each scenario, there was also no significant difference in how people voted as the number of missing ballots increased. However, significant differences ( $P < 0.005$ ) were found when comparing the strategies used by voters in those conditions electing one or two winners compared to those electing three winners. 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 ?? 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 changed depending on the individual and the size

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

Table 4: The maximum expected utility for Scenarios A and B, and the heuristic that achieves this optimal outcome.  $n$  represents the number of missing ballots.

of the winning set, as seen in Figure ??.

## 5 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 ??, namely, *Complete*, *Take X Best* with  $X$  equal to the size  $k$  of the winning set, AU and AUT.

### 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:

$$M^{Opt}(k, \hat{n}) = \arg \max_{b \in \mathcal{P}(C)} E(b, k, \hat{n}). \quad (1)$$

The optimal ballots and the corresponding maximum expected values for Scenarios A and B can be found in Table 4. 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, the *Take the 1 Best* heuristic is optimal when there is a single winner and no missing ballots. However, it is optimal to *Take the 2 Best* with 3 winners and no missing ballots.

### 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*

	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:	<b>58.1% (11.6%)</b>	<b>68.0% (1.8%)</b>	<b>67.9% (3.4%)</b>

Table 5: Mean prediction accuracy (standard deviation) for each model across conditions.

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  $\beta$  and  $\tau$ . 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 5 and shown in Figure 1.

## 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.

Next, we tested whether or not voters make a trade-off between the utility and attainability of the candidates. We first




Figure 1: Mean and standard deviation of prediction accuracy for each model across conditions.

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 the way that it calculates the attainability-utility for every possible subset of candidates, rather than each individual candidate. The attainability-utility of ballots containing all utility generating candidates is higher than ballots that contain only a subset of them, leading to *Complete Ballots* being predicted much more often than sincere *Take the X Best* behavior (See Table ??).

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.

## 6 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 re-

quired 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.

## Acknowledgements

Nicholas Mattei was supported by NSF Award IIS-2007955 and an IBM Faculty Research Award. K. Brent Venable are supported by NSF Award IIS-2008011.

## References

- Aziz, H.; Gaspers, S.; Gudmundsson, J.; Mackenzie, S.; Mattei, N.; and Walsh, T. 2015. Computational Aspects of Multi-Winner Approval Voting. In *Proc. of the 14th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*.
- Aziz, H.; Gaspers, S.; Mattei, N.; Narodytska, N.; and Walsh, T. 2013. Ties matter: Complexity of manipulation when tie-breaking with a random vote. In *Proc. of the 27th AAAI Conference on Artificial Intelligence (AAAI)*, 74–80.
- Brams, S. J. 1980. Approval voting in multicandidate elections. *Policy Studies Journal* 9(1): 102.
- Brams, S. J. 1982. Strategic information and voting behavior. *Society* 19(6): 4–11.
- Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A. D., eds. 2016. *Handbook of Computational Social Choice*. Cambridge University Press.
- Endriss, U. 2007. Vote manipulation in the presence of multiple sincere ballots. In *Proc. 11th Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*, 125–134.
- Fairstein, R.; Lauz, A.; Meir, R.; and Gal, K. 2019. Modeling People’s Voting Behavior with Poll Information. In *Proc. of the 18th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*.
- Faliszewski, P.; Hemaspaandra, E.; and Hemaspaandra, L. A. 2010. Using Complexity to Protect Elections. *Commun. ACM* 53(11): 74–82. ISSN 0001-0782. doi:10.1145/1839676.1839696. URL <http://doi.acm.org/10.1145/1839676.1839696>.
- Faliszewski, P.; and Procaccia, A. D. 2010. AI’s War on Manipulation: Are We Winning? *AI Magazine* 31(4): 53–64.
- Gigerenzer, G.; and Goldstein, D. G. 1996. Reasoning the fast and frugal way: models of bounded rationality. *Psychological Review* 103(4): 650.
- Kilgour, D. M. 2010. Approval Balloting for Multi-winner Elections. In *Handbook on Approval Voting*, chapter 6. Springer.
- Laslier, J.-F.; and Sanver, M. R., eds. 2010. *Handbook on Approval Voting*. Studies in Choice and Welfare. Springer-Verlag.
- Mattei, N. 2020. 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)*, 5169–5173.

- Mattei, N.; Narodytska, N.; and Walsh, T. 2014. How Hard Is It To Control an Election by Breaking Ties? In *Proc. of the 21st European Conference on Artificial Intelligence (ECAI)*, 1067–1068.
- Mattei, N.; and Walsh, T. 2013. 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)*.
- Mattei, N.; and Walsh, T. 2017. A PREFLIB.ORG Retrospective: Lessons Learned and New Directions. In Endriss, U., ed., *Trends in Computational Social Choice*, chapter 15, 289–309. AI Access.
- Meir, R. 2018. *Strategic Voting*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool Publishers.
- Meir, R.; Lev, O.; and Rosenschein, J. 2014. A local-dominance theory of voting equilibria. In *Proc. of the 15th ACM conference on Economics and Computation (ACM EC)*, 313–330.
- Meir, R.; Procaccia, A. D.; and Rosenschein, J. S. 2008. 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)*, 991–998.
- Meir, R.; Procaccia, A. D.; Rosenschein, J. S.; and Zohar, A. 2008. Complexity of Strategic Behavior in Multi-Winner Elections. *Journal of Artificial Intelligence Research* 33: 149–178.
- Mennle, T.; Weiss, M.; Philipp, B.; and Seuken, S. 2015. The power of local manipulation strategies in assignment mechanisms. In *Proc. of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*.
- Niemi, R. G. 1984. The Problem of Strategic Behavior under Approval Voting. *The American Political Science Review* 78(4): 952–958.
- Obraztsova, S.; Elkind, E.; and Hazon, N. 2011. Ties matter: Complexity of voting manipulation revisited. In *Proc. of the 10th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 71–78.
- Regenwetter, M.; Ho, M.-H. R.; and Tsetlin, I. 2007. Sophisticated approval voting, ignorance priors, and plurality heuristics: A behavioral social choice analysis in a Thurstonian framework. *Psychological Review* 114(4): 994–1014.
- Reijngoud, A.; and Endriss, U. 2012. Voter response to iterated poll information. In *Proc. of the 11th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 635–644.
- Scheuerman, J.; Harman, J.; Mattei, N.; and Venable, K. B. 2019a. Heuristics in Multi-Winner Approval Voting. Behavioral EC. Phoenix, AZ.
- Scheuerman, J.; Harman, J. L.; Mattei, N.; and Venable, K. B. 2019b. Heuristic Strategies in Uncertain Approval Voting Environments. *CoRR* abs/1912.00011. URL <http://arxiv.org/abs/1912.00011>.
- Tal, M.; Meir, R.; and Gal, Y. 2015. A Study of Human Behavior in Online Voting. In *Proc. of the 14th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 665–673.
- Tyszler, M.; and Schram, A. 2016. Information and strategic voting. *Experimental Economics* 19(2): 360–381.
- Walsh, T. 2007. Uncertainty in preference elicitation and aggregation. In *AAAI*, volume 7, 3–8.
- Zou, J.; Meir, R.; and Parkes, D. 2015. Strategic Voting Behavior in Doodle Polls. In *Proc. of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*, 464–472.