Chapter 5

Heuristics and Biases in Approval Voting

The previous chapters considered bias that occurs in the context of single agents. In this chapter, we consider heuristics in a multi-agent voting scenario and show when they are effective and when they may lead to bias. Many real-world situations involve multiple agents participating in collective decision-making tasks. This usually involves aggregating preferences through a voting rule or procedure to choose the alternative that best reflects the preferences of the group. Agents may vote with their true preference, use heuristics (such as voting for the current leader in a poll), or vote strategically to attain a better outcome. In real-world voting scenarios, people often do not have complete information about other voter preferences, and it can be computationally complex to identify a strategy that will maximize their expected utility. In such scenarios, it is often assumed that voters will vote sincerely rather than expending the effort to strategize. Here we examine heuristics and bias in approval voting elections. In an approval election, voters can try to maximize their utility or use a heuristic. Several sincere heuristics are possible, including voting completely truthfully (for all candidates for which the voter has some positive utility) or voting for their top x candidates with the highest utility. We present a behavioral experiment to examine the use and effectiveness of sincere heuristics in multi-winner approval voting scenarios with missing votes. The results show that people generally vote sincerely but used different underlying heuristics that depended on features of the voting scenario, including the number of winners and whether or not there was a strong preference for or against a particular candidate [26]. This work provides key insights on human behavior in voting environments and can inform the development of more realistic simulation tools and more accurate predictions of election outcomes where approval voting is used. This chapter presents work that will be published in the *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2020* [26] and that was presented at the *Behavioral EC Workshop at the 20th ACM Conference on Economics and Computation 2019* [25] and the *Society for Judgement and Decision Making Annual Conference 2019* [88].

5.1 Introduction

Computational Social Choice (COMSOC) investigates computational issues surrounding the aggregation of individual preferences and collective decision making [22]. Much of the work in this area has focused on the computational complexity of manipulating elections under different voting rules. When it is computationally prohibitive to manipulate an election, it is assumed that voters will vote with their true preferences rather than trying to strategize [23].

Voting truthfully is just one possible heuristic that voters may use when faced with complex voting scenarios, where the optimal strategy is not easy to compute. A recent study of voting behavior in multi-winner approval elections showed that the majority of voters did not vote truthfully or optimally [25]. Instead, the predominant strategy was to use a *take the X best* heuristic, which prioritized the highest utility candidates. Another study showed that in a plurality election where a preferred candidate was currently dominated, remaining voters would compromise and vote for the leader [89].

The effectiveness of a particular heuristic depends on the environment in which that heuristic is being used. Decision science research has examined heuristic decision making in complex and uncertain situations. Sometimes, heuristics are viewed as second best shortcuts, when the environment is too complex to use rational strategies [90]. However, researchers have also shown that heuristics are adaptive strategies that work in the real world, leveraging natural cognitive abilities that exploit the structure of the environment, often leading to better outcomes with the use of less information. Key to this view of heuristics is that in uncertain environments, decision strategies that ignore some information can sometimes achieve better performance than more complicated optimization strategies in situations that are computationally complex or uncertain [91]. In many situations, heuristics have been shown to outperform solutions that use more complex algorithms (i.e., stock market predictions [36]).

In this chapter, we examine the effectiveness of heuristics in a multi-agent setting, namely in single winner and multi-winner approval voting elections with uncertainty. In the scenarios we consider, voters are presented with situations where there are missing votes. In approval voting, an agent may vote for as many candidates as they wish. Winners are chosen by tallying up the votes and choosing the top-k candidates receiving the most votes. This setting is interesting to consider because a voter has potentially multiple sincere votes they can cast, with some that are more beneficial than others. Under the basic approval voting rule, optimal manipulation can be computed in polynomial time when an agent has complete information about the preferences of all the voters [92]. However, it has been shown that when information about voting preferences is missing, computing the possible winners or manipulating the vote is computationally complex [93]. Although manipulation may be computationally hard, voting truthfully or using other heuristics may still maximize a voter's expected utility.

5.2 Preliminaries

We give a brief overview of the mathematical formalism used to study approval voting and formally define the heuristics that we will consider in this chapter.

5.2.1 Approval Voting

Following [94] and [95] 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 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 total weight of approvals (approval score),

$$AV(W) = \sum_{i \in N} |W \cap A_i|. \tag{5.1}$$

Informally, the winning set under AV is the set of candidates that are approved by the largest number of voters.

In some cases, it is necessary to use a tie-breaking rule in addition to a voting rule in order 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 [96, 97, 98]. Typical 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 chapter, as discussed in [96], we break ties by selecting the winner randomly to more closely simulate a real-world election.

In order to align our work with the literature on decision heuristics [99] we assume that each agent $i \in N$ also has a real valued utility function $u_i : C \rightarrow \mathcal{R}$ (see example profile in Table 5.1. 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 is the subset elected by the voting rule we will refer to $u_i(W)$ as agent's i's utility for the *outcome* of the election.

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

Table 5.1: A sample approval voting profile

5.2.2 Truthfulness and Sincerity in Approval Ballots

The literature on approval voting for multi-winner elections goes over nearly 40 years to the work of Brams [100]. For nearly that entire period, there has been intense discussion of the strategic aspects of approval balloting [101]. Researchers over the years have made a variety of assumptions and (re)definitions of what makes a particular vote either *truthful* or *strategic*. Much of this commentary is captured in the introductory chapter to the Handbook of Approval Voting [24]. As detailed by Laslier and Sanver [24], Niemi [102] quotes the following definition of Sincere Approval Voting from Brams [101]: ”A voter votes sincerely if and only if whenever he votes for some candidate, he votes for all candidates preferred to that candidate” and writes ”Note that this definition includes nothing about approval as such; it does not require voting only for 'approved' alternatives.” Niemi even writes, ”the existence of multiple sincere strategies almost begs the voter to behave strategically.” [102].

Within the COMSOC community, this issue has arisen a number of times: what does it mean to be sincere and/or truthful in a given situation? This is nicely expressed by [103], "In approval voting, a ballot consists of the names of any subset of the set of candidates standing; these are the candidates the voter approves of. The candidate receiving the most approvals wins. A ballot is considered sincere if the voter prefers any of the approved candidates over any of the disapproved candidates. Hence, there will be multiple sincere ballots for any given preference ordering."

However, this does rest on an assumption about the underlying preference model, as discussed by [92], when agents are only endowed with binary utilities, a truthful vote is always a strategic vote [104, 92], i.e., approval voting is incentive compatible. A *strategic* vote is one in which an agent maximizes their total (expected) utility given a particular decision setting. However, as [92] continues, "manipulation in Approval is a subtle issue, since the issue may be ill-defined when the voters are assumed to have linear preferences over the candidates. In this case, there are multiple sincere ballots (where all approved candidates are preferred to all disapproved candidates)."

Given this discussion, we make the following distinctions:

1. In the presence of additive utilities and multiple winners, we assume that a *completely truthful* vote is one where the voter approves all candidates for which they have positive utility.

2. A *sincere vote*, which includes the definitions of Endriss [103], Meir et al. [92], and Brams [101], is one in which if a voter prefers a particular candidate, then he approves all candidates that are preferred to that particular candidates. 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 discussed in the following.

As an example of a truthful vote versus sincere votes, consider a voter has the set of preferences $[A = 0.4, B = 0, C = 0.2, D = 0.01]$. Given these preferences, a completely truthful voter vote for $[A, C, D]$, whereas a sincere voter could vote for either $[A]$, $[A, C]$ or $[A, C, D]$.

We argue and will use the terminology that any vote that is not completely truthful by our definition is considered strategic. While it is the case that these votes may be sincere, we argue they are not completely truthful as, given the definitions above, it is strategically leaving some information out. In what follows, we consider, as does much of the literature, the question of which strategic vote to use, and what internal heuristics one may be using to decide it.

5.2.3 Heuristics in Approval Voting

We present three heuristics inspired by the literature that we believed a priori could be used in single winner and multi-winner approval voting. These include *truthful*, *take the X best*, and *regret minimization*. Each of these strategies ignores information (i.e., the total votes so far) and use only the utility of each candidate to decide whom to vote for.

Truthful.

We define a *truthful* vote as one where an agent approves of all candidates for which they have positive utility. This corresponds to the notion of *completely truthful* above.

Take the X Best.

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. Formally, top x candidates $T = t_1 > ... > t_x$ where $u_1(t_1) > ... > u_X(t_X)$. The agent will then vote for the top-X candidates in the list. X could be calculated using a magnitude cut off or a proportional difference between preferences [105]. We do not use any specific rule to choose X, opting instead to test all sizes of X. In this chapter, we examine situations where there candidates' utility values are not tied, so no tie-breaking rule is assumed. In the future, it would be interesting to explore if and how voters choose between candidates with equal utility.

Regret Minimization.

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 may choose to minimize the chance that the disliked candidate(s) will win by voting for all other candidates, whether they generate positive utility or not [106, 107].

5.3 Related Work

Approval voting is a set of methods for aggregating group preferences that is particularly popular among economists, computer scientists, psychologists, and beyond [24, 108]. There are even multiple political action committees (PACs) in the United States, e.g., The Center for Election Science[1], that are committed to seeing the United States change voting procedures to approval voting. One reason for this popularity is the idea that participants are allowed to express a preference over a set of candidates and not just a single one. In France, a large study was run parallel to the 2002 election, showing that many voters would have preferred approval ballots to traditional plurality ballots [109].

The complexity of manipulation for various types of approval voting (AV) has received considerable attention in the COMSOC literature [22]. Assuming that agents act rationally and have full information about the votes of other agents, when agents have *Boolean utilities*, i.e., when all agents either have utility 1 or 0 for candidates they approve or disapprove of, respectively, AV is strategy-proof. When agents have general utilities, finding a vote that maximizes the agent's utilities can be computed in

---

[1] `https://www.electionscience.org/`

polynomial time [104, 92]. For variants of AV, including Proportional Approval Voting, Satisfaction Approval Voting, and the Repeated Approval Voting, the complexity of finding utility-maximizing votes ranges in complexity from easy to coNP-complete [94].

Many theoretical works in COMSOC make worst-case computational assumptions: manipulators have complete information, all votes are known, etc. However, there also are several efforts to expand these worst-case assumptions and strategic issues to include the presence of uncertain information or when agents are not perfectly rational. In [110], agents are given access to poll information, and agent behaviors are modeled as being k-pragmatist, i.e., they only look at the top k candidates when deciding whether or not to make a strategic decision. In [111], agents are modeled as behaving in *locally dominant* ways, i.e., they take into account only a small number of possible outcomes when deciding whether or not to act strategically in a particular voting setting. A survey of other recent work on issues surrounding strategic voting is given by [112].

There is a growing effort to use simulations and real-world data to test various decision-making models, e.g., [113, 114]. 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 is the studies of [115], which focus on elections of various professional societies where approval balloting was used and the work of [116] where many approval voting settings were obtained from Doodle, an online polling platform. In [115] election data is used along with proposed heuristics for individual choice behavior, the conclusion is 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. 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 relationships between environment and decision. To our knowledge, the work presented in this chapter is the first that examines human voting behavior in multi-winner approval settings. The behavioral experiment presented here examines how voters select different strategies, depending on the underlying environmental factors (i.e., number of winners and number of missing votes), and this parameterization is also novel.

Three recent papers address strategic voting under the plurality rule, where agents are making decisions in uncertain environments. First, [117] study 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 [89] an online system is presented where participants vote for cash payments in a number of settings using the plurality rule under uncertainty. Two specific scenarios are studied: one where a user votes after being given access to a large pre-election poll and the second where agents vote simultaneously and can update their votes. They find that most participants do not engage in strategic voting unless there is a clear way to benefit. In the iterative setting, 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 [118], 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 agents 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 this work on plurality to consider heuristics and bias in the significantly more complex setting of approval voting with uncertainty, showing that it may be ecologically rational for voters to use heuristics over more complex optimization strategies.

5.4 Behavioral Experiment Design

We design specific scenarios that a voter may encounter in approval voting and in which we predict users may use the heuristics described in Section 5.2.3. For each scenario, we explore which strategies people use, if they maximize expected utility, and whether people vote truthfully for all candidates with positive utility or use some other approach.

Each scenario consists of a set of candidates $C = \{c_1, ..., c_5\}$, the utility of agent i, and the current number of votes for each candidate, with i's utility for each candidate in $[-1.0, 0.25]$. We manipulate two environmental features, including the number of winners in the election ($k = 1, 2, 3$) and the number of missing votes ($n = 0, 1, 3$), in addition to i's vote. When the final ballots result in a tie, the winner(s) are chosen randomly.

For each scenario, the maximum expected utility can change for different numbers of winners and missing votes. We calculate the expected utility by generating the set of all possible votes that i could cast over C, i.e. the power set $V = \mathcal{P}(C\ )$.

The expected utility is then calculated for each of i's votes in V for every combination of k winners, and n remaining voters.

$$E[u(v, k)] = p_1 u_i(c_1) + ... + p_m u_i(c_m)$$

$$\forall v \in V,\ 0 < k \leq 3,\ 0 < m \leq 5.$$

Here, p_j refers to the probability that candidate c_j is in the current profile. For all combinations of numbers of winners and missing votes, we calculate 1) the expected utility for each heuristic, 2) the maximum expected utility, and 3) any votes in V not represented by the heuristics that maximized i's expected utility.

This calculation shows how the computation of an expected utility maximizing strategy, requires many calculations, which is cognitively demanding and for which heuristics can be low effort alternatives.

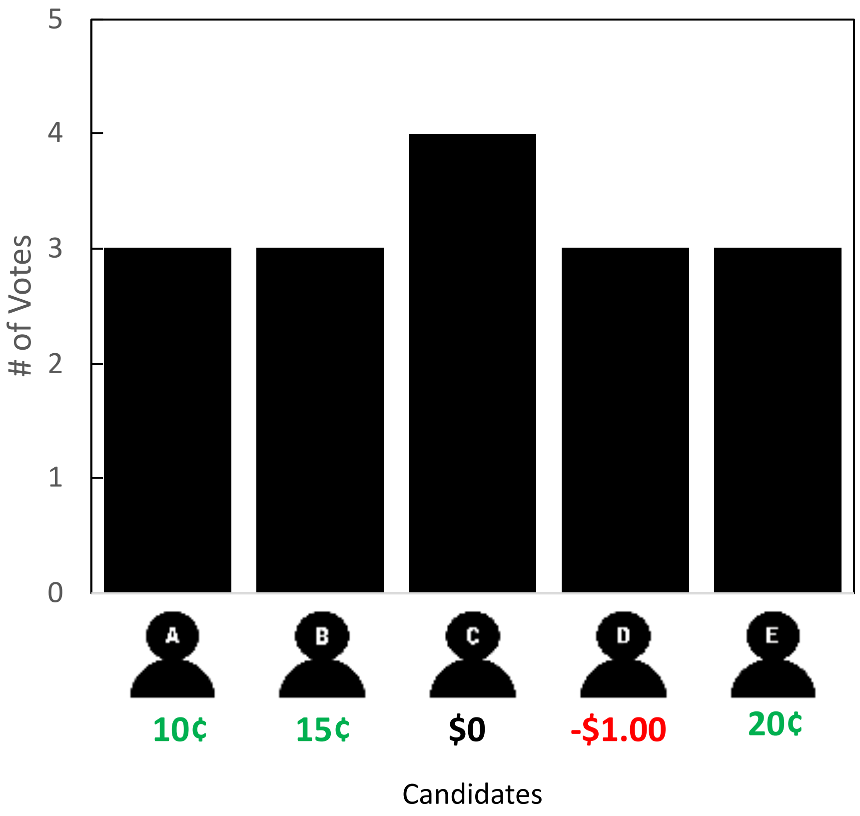

Figure 5.1: Example of a subjects' view of a scenario's details, including the candidates, utility and votes.

5.4.1 Scenarios

Below we detail the candidates, utilities, the number of current votes, and the number of missing votes (if applicable) for each candidate for several partial profiles, which we designed to study specific behaviors as the number of winners (k) and the number of missing votes (n) change. In the behavioral study, the scenarios are presented as depicted in Figure 5.1 along with text describing how many voters remain to vote.

For each scenario described, we show the maximum possible utility for each combination of missing votes (n) and numbers of winners (k) (for example, see Table 5.3). The tables are annotated with the heuristic strategy that would lead to the maximum expected utility, that is, the one we would expect each participant to employ when presented with a particular voting scenario.

Scenario 1a: Candidate with Trivial Utility

This scenario (Tables 5.2, 5.3) represents a situation where a non-leading candidate generates a trivial amount of utility if elected.

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

Table 5.2: Scenario 1a details, including candidates, utilities and votes. *Heuristic votes*: Truthful: [A,B,C,E], Take 1 Best: [E], Take 2 Best: [E,B], Take 3 Best: [E,B,A]

	# winners (k)		
n	1	2	3
0	0.12 **Take 1**	0.22 **Take 1**	0.31 **Take 2**
1	0.11 **Take 1**	0.21 **Take 2**	0.30 **Take 2**
3	0.11 **Take 1**	0.20 **Take 2**	0.29 **Take 2**

Table 5.3: Scenario 1a: Maximum expected utility and the voting strategies that achieve it. n represents the number of missing votes.

Scenario 1b: Leader with Trivial Utility

This scenario (Tables 5.4, 5.5) examines a situation where a leading candidate will generate a trivial amount of utility of elected.

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

Table 5.4: Scenario 1b details, including candidates, utilities and votes. *Heuristic votes*: Truthful: [A,B,C,D], Take 1 Best: [D], Take 2 Best: [D,B], Take 3 Best: [D,B,A]

	# winners (k)		
n	1	2	3
0	0.13 **Take 1**	0.26 **Take 1**	0.36 **Take 2**
1	0.12 **Take 1**	0.22 **Take 2**	0.31 **Take 2**
3	0.11 **Take 1**	0.21 **Take 2**	0.29 **Take 2**

Table 5.5: Scenario 1b: Maximum expected utility and voting strategies that achieve it. n represents the number of missing votes.

Scenario 2a: Dominated for One and Two Winners

This scenario (Tables 5.6, 5.7) examines a situation where neutral candidates dominate the preferred candidates. When only 1 or 2 candidates can win, there is no possibility of electing a preferred candidate, except when there are 3 missing votes. See Table 5.6 and Table 5.7 for scenario details.

Candidate:	**A**	**B**	**C**	**D**	**E**
Utility:	0.05	0.10	0	0	0.25
# Votes:	1	1	4	4	1

Table 5.6: Scenario 2a details, including candidates, utilities and votes. *Heuristic votes:* Truthful: [A,B,E], Take 2 Best: [E,B], Take 1 Best: [E]

	# winners (k)	
n	1	2
0	–	–
1	–	–
3	0.01 **Truth**	0.04 **Truth**

Table 5.7: Scenario 2a: Maximum expected utility and voting strategy that achieve it. n represents the number of missing votes. When $n = 0$ and $n = 1$, it is impossible to elect a preferred candidate, and all voting strategies lead to an expected utility of 0.

Scenario 2b: Dominated for Three Winners

Like Scenario 2a, this scenario (Tables 5.8,5.9) examines a situation where neutral candidates dominate the preferred candidates. In this particular scenario there is no possibility of electing a preferred candidate in the 3-winner case. See Table 5.8 and Table 5.9 for scenario details.

Candidate:	**A**	**B**	**C**	**D**	**E**
Utility:	0.10	0	0	0	0.25
# Votes:	1	4	4	4	1

Table 5.8: Scenario 2b details, including candidates, utilities and votes. *Heuristic votes:* Truthful: [A,E], Take 1 Best: [E]

	# winners (k)
n	3
0	–
1	–
3	0.05 **Truth**

Table 5.9: Scenario 2b: Maximum expected utility and voting strategy that achieve it. n represents the number of missing votes. It is impossible to elect a preferred candidate for $n = 0$ and $n = 1$.

Scenario 3: Disliked Candidate

This scenario (Tables 5.10,5.11) examines a situation where a candidate will generate negative utility if elected, representing a situation where the voter i dislikes the candidate.

Candidate:	**A**	**B**	**C**	**D**	**E**
Utility:	0.05	0.10	0	-1.00	0.25
# Votes:	3	3	4	4	4

Table 5.10: Scenario 3 details, including candidates, utilities and votes. *Heuristic votes:* Truthful: [A,B,E], Take 1 Best: [E], Take 2 Best: [E,B], Regret Minimization: [A,B,C,E]

	# winners (k)		
n	1	2	3
0	0.25 **Truth** **Take 1** **Take 2**	0.25 **Regret** **[C,E]**	-0.03 **Regret**
1	0.10 **Regret**	0.06 **Regret**	-0.10 **Regret**
3	0.03 **Regret**	-0.03 **Regret**	-0.17 **Regret**

Table 5.11: Scenario 3: Maximum expected utility and the voting strategies that achieve it. n represents the number of missing votes. [C,E] represents a vote that maximizes expected utility, but does not fall into one of our defined heuristics.

Scenario 4: Neutral Leader

This scenario (Tables 5.12,5.13) examines a situation where a neutral candidate is leading the election.

Candidate:	**A**	**B**	**C**	**D**	**E**
Utility:	0.10	0	0.15	0.20	0
# Votes:	3	4	3	3	3

Table 5.12: Scenario 4 details, including candidates, utilities and votes. *Heuristic votes:* Truthful: [A,C,D], Take 1 Best: [D], Take 2 Best: [C,D]

	# winners (k)		
n	1	2	3
0	0.11 **Truth**	0.23 **Truth**	0.32 **Take 2**
1	0.11 **Truth**	0.22 **Take 2**	0.31 **Truth**
3	0.11 **Take 2**	0.21 **Truth**	0.31 **Truth**

Table 5.13: Scenario 4: Maximum expected utility and voting strategies that achieve it. n represents the number of missing votes.

5.4.2 Implementation

Participants. 104 participants were recruited through Mechanical Turk to participate in the voting heuristics study. Participants were paid $1.00 to complete the survey. They also received a bonus of no more than $8.00 that was determined by the outcome of the hypothetical elections.

Procedure. In the study, participants were asked to vote in a series of unrelated hypothetical elections, using instances of the scenarios described in Section 5.4.1. 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.

Participants were asked to give informed consent and then proceeded to the study. Instructions explained approval voting and the tie-breaking mechanism with examples. After reading the instructions, participants proceeded through single-winner scenarios, first encountering scenarios with 0 missing votes, then 1 and finally, 3 missing votes. From there, the survey presented each participant with a series of multi-winner scenarios for their assigned group (2 or 3-winner), in order of increasing uncertainty.

Each election displayed an image showing the candidates, the number of votes cast for each candidate so far, and how much money the participant would earn for each candidate if they were elected. Figure 5.1 is an example of what the participants saw.

When voting, subjects could vote for 0 or more (up to five) of the five candidates. After voting, they would see the election results, including the winners, the amount earned, and the ballots cast by any missing voters (when applicable).

The experiment was designed so that participants could choose to vote truthfully (for all candidates with positive utility) or manipulate their vote to achieve a higher utility. We expected that most people would try to vote strategically, but since the situations involved varying degrees of uncertainty and were cognitively complex, participants would not perform all of the necessary computations to identify the strategy that maximizes their utility. Instead, we expected that people would use heuristics, such as being *truthful* or using *take the X best*, to prioritize the highest priority candidates.

5.5 Results & Discussion

The results of the behavioral experiment described above showed unique patterns of behavior in each scenario, particularly across the different conditions. The next subsections describe the results for each scenario.

5.5.1 Scenarios 1a, 1b: Trivial Utilities

In these scenarios, we wanted to see whether or not people would vote for a candidate with a trivial utility (represented as a candidate that would earn 1¢ if elected). Scenario 1a (see Tables 5.2, 5.3) examined the case when the trivial candidate was not leading the election and Scenario 1b (see Tables 5.4, 5.5) examined when it was. We found that in both scenarios, people generally did not vote truthfully for all candidates with positive utility, including the trivial candidate. In Scenario 1a, only 15.4% voted truthfully in the 1-winner election, 16.0% in the 2-winner election, and 9.9% in the 3-winner election. Scenario 1b was similar, with only 14.7% voting truthfully in the 1-winner election, 11.3% in the 2-winner election, and 8.0% in the 3-winner election.

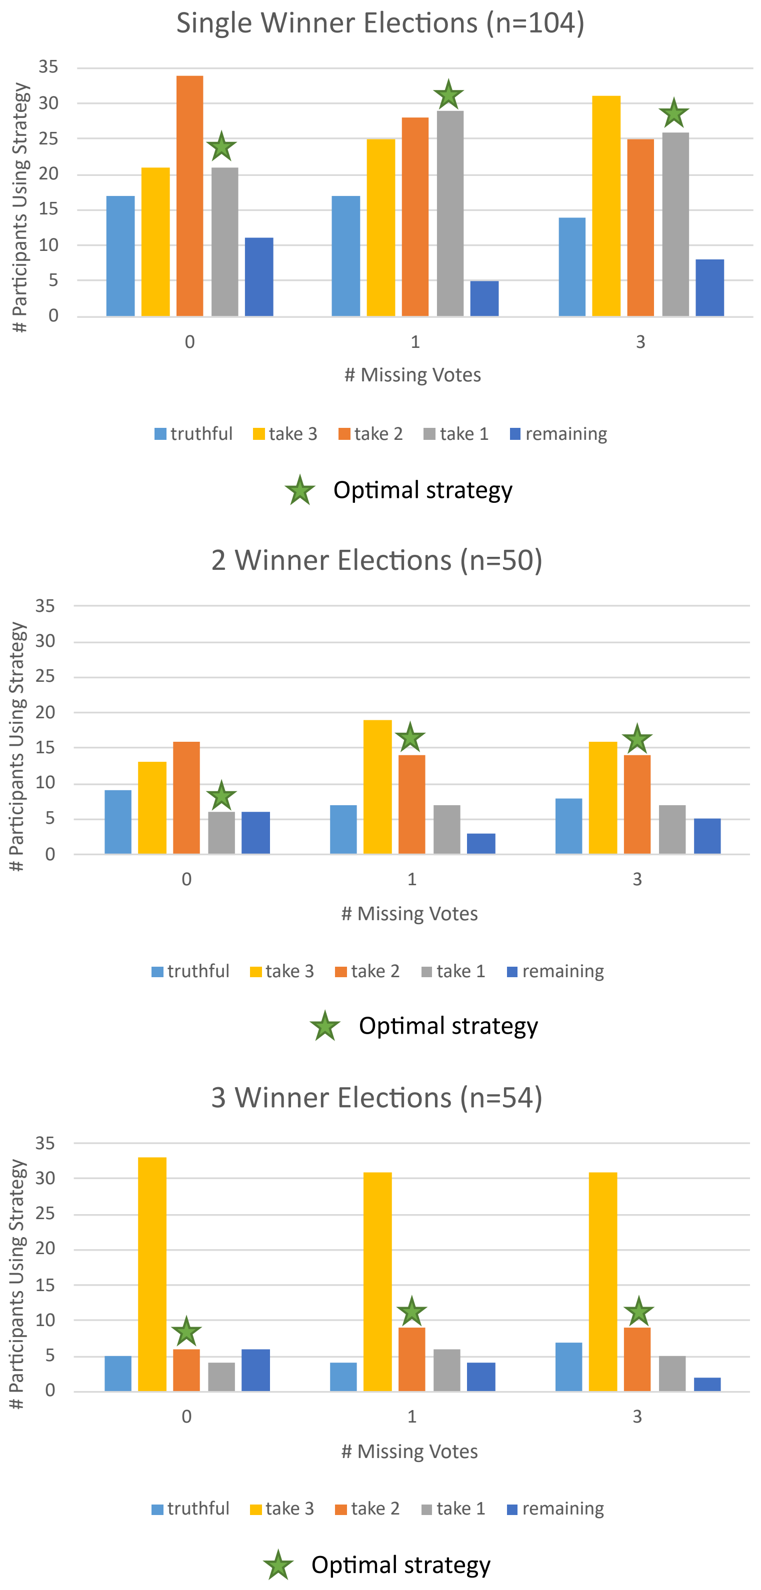

Figure 5.2: Scenario 1a (Trivial Utility) Results. Maximizing strategies are marked with a star. Most voters used a *take the X best* approach. There was a significant difference (P ¡ 0.0005) in voting strategies as the number of winners changed, but not as the number of missing voters changed.

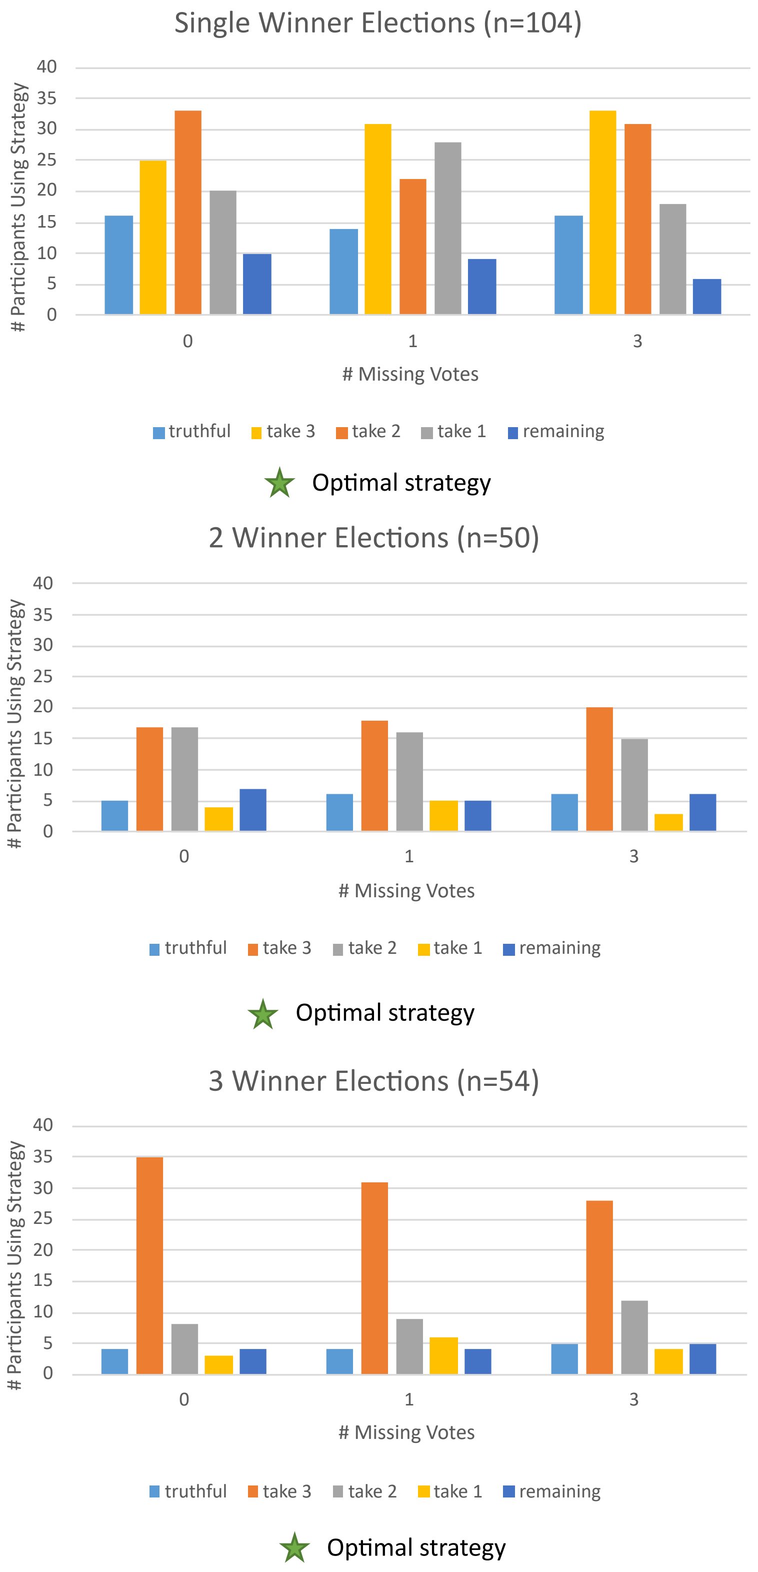

Figure 5.3: Scenario 1b (Trivial Leader) Results. Maximizing strategies are marked with a star. Most voters used a *take the X best* approach. There was a significant difference ($P < 0.0005$) in voting strategies as the number of winners changed, but not as the number of missing voters changed.

In both of these scenarios, it was optimal to use a *take the X best* approach, with X increasing as the number of winners and missing votes increased (see Tables 5.3 and 5.5). Although the majority of people used a *take the X best* strategy (Scenario 1a: 77.8%, Scenario 1b: 78.8%), they rarely prioritized the X candidates that would maximize the expected utility. In Scenario 1a, only 21.5% of participants chose a strategy that would lead to an optimal outcome. In Scenario 1b, only 18.4% chose a maximizing strategy.

Using χ^2 analysis, we found no significant difference in how people voted as the number of missing votes increased, even in the 2-winner elections where increased uncertainty led to a different maximizing strategy (take the 1 best for 0 missing votes vs. take the 2 best for 3 missing votes). However, significant differences ($P < 0.005$) were found when comparing the voting strategies used by those electing one or two winners compared to those electing three winners. In general, when voting in the 1-winner and 2-winner elections, participants voted for 2 or 3 candidates (1-winner: 57.9%, 2-winner: 70.7%) more often than other strategies. When participants voted in the 3-winner election, they usually voted for 3 candidates (61.7%) (see Figures 5.2 and 5.3).

5.5.2 Scenarios 2a, 2b: Dominated Preferences

In these scenarios, we wanted to see if people would vote truthfully when neutral candidates dominated their preferred candidates. Scenario 2a (see Tables 5.6, 5.7) examined this in the context of 1 and 2-winner elections, where Scenario 2b (see Tables 5.8, 5.9) looked at 3-winner contexts. In both Scenario 2a and 2b it was possible to elect a preferred candidate when there were 3 missing votes, where the maximizing strategy was to vote truthfully (see Tables 5.7 and 5.9). Voting truthfully was also the participants' dominant strategy no matter the numbers of winners or missing votes (Scenario 2a: 44.2%, Scenario 2b: 62.3%). The second most common strategy was to abstain (Scenario 2a: 16.5%, Scenario 2b: 20.4%) (see Figures 5.4 and 5.5).

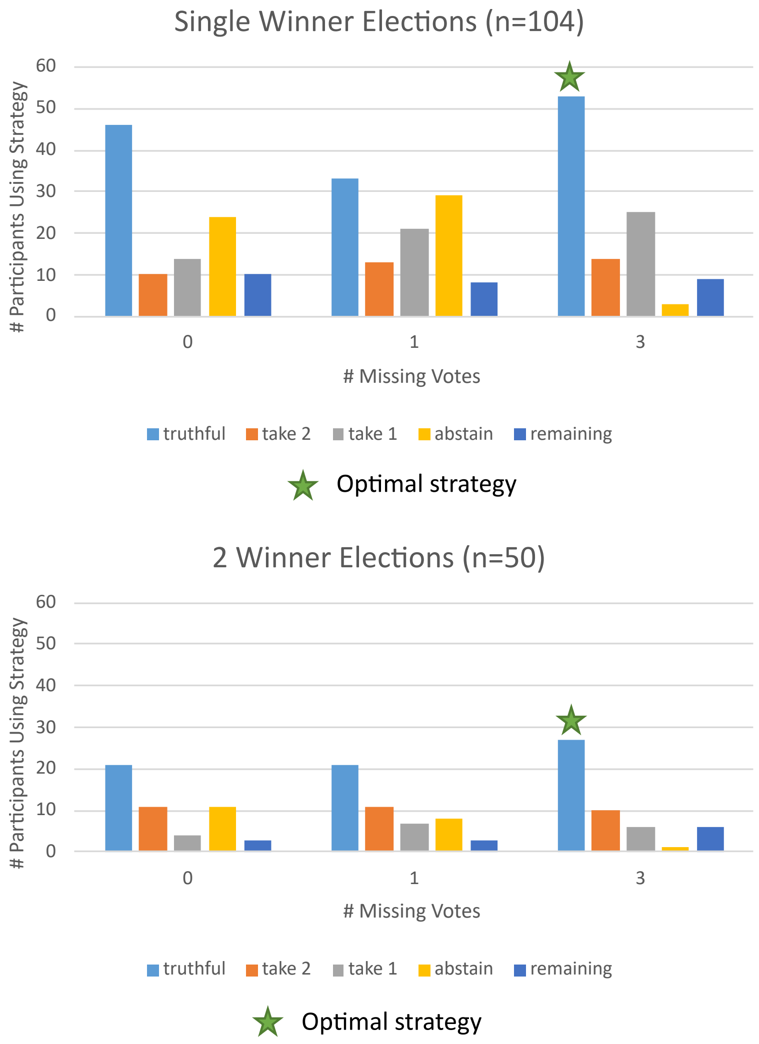

Figure 5.4: Scenario 2a (1 and 2-winner Dominated Preferences) Results. Maximizing strategies are marked with a star. Participants tended to vote truthfully, with abstention coming in as the second-place strategy.

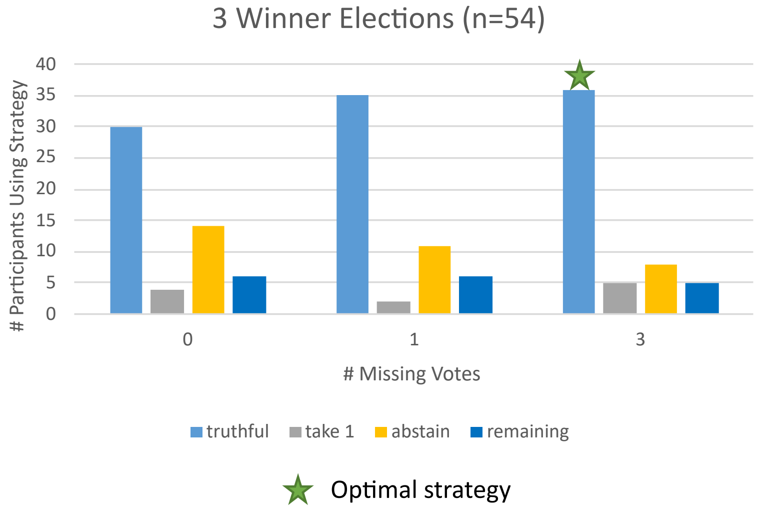

Figure 5.5: Scenario 2b (3-winner Dominated Preferences) Results. Maximizing strategies are marked with a star. Participants mostly voted truthfully, with abstention coming in as the second-place strategy.

Using χ^2 analysis, we found a significant difference ($P < 0.0005$) when comparing the voting strategies that people used in Scenario 2a when voting for one winner versus two winners. There was also a significant difference in voting strategies when there was 0 or 1 missing vote, compared to 3. When there were 0 missing votes, participants chose to abstain 22.7%, but when there were 3 missing votes, only 2.6% of participants abstained. This seems to indicate that voters in Scenario 2a recognized that they had a small chance to elect a preferred candidate in the 3-winner condition and voted accordingly. In Scenario 2b, the number of abstentions decreased as the level of uncertainty increased (0 missing votes: 25.9% abstain, 3 missing votes: 14.8% abstain), but it was not enough to result in a significant difference in each groups' voting strategies.

5.5.3 Scenario 3: Disliked Candidate

In this scenario, we explored how people would vote in the presence of a disliked candidate that would generate negative utility if elected (see Tables 5.10, 5.11). Here, *regret minimization* was a maximizing strategy for most combinations of numbers of winners and missing votes. However, in the single-winner election with 0 missing votes, it was possible to achieve the optimal strategy both by being truthful or using *take the X best*. When one vote was missing in the single winner election, it was best to be truthful. Voting [C,E], a strategy that did not align with any of the heuristics defined in this chapter, was also a maximizing strategy in the 2-winner scenario with 0 missing votes.

This scenario was interesting as people's voting strategy changed significantly ($P < 0.005$) when comparing the strategies used by those voting in 1-winner elections with 0 or 1 missing votes, to those voting in elections missing 3 votes. In all three single winner groups, more people responded with a truthful (0 missing votes: 29.8%, 1 missing vote: 36.5%, 3 missing votes: 30.8%) or take the 1 best strategy (0 missing votes: 46.1%, 1 missing votes: 35.6%, 3 missing votes: 23.1%), than any other strategy. However, the number of voters using regret minimization (0 missing votes: 5.8%, 1 missing vote: 4.8% and 3 missing votes: 17.3%) increased so that it was the 3rd most popular strategy when 3 voters were missing (see Figure 5.6).

The responses to the 2-winner election were more variable, with maximizing strategies being more popular than other strategies (0 missing votes: 38.0%, 1 missing vote: 48.0%, 3 missing votes: 50.0%), but still not used by a majority of the candidates. In the 3-winner election, being truthful was the most popular response, whereas the optimal strategy (regret minimization) was used only 20.4% of the time.

In general, it was common for participants in this scenario to vote for as many candidates as there were winners in the election. When voting in the 1-winner election, participants voted for one candidate 37.9% of the time. In the 2-winner election, voting for two candidates was also the most common (40.0%), and participants in the 3-winner election mostly voted for three candidates (53.7%).

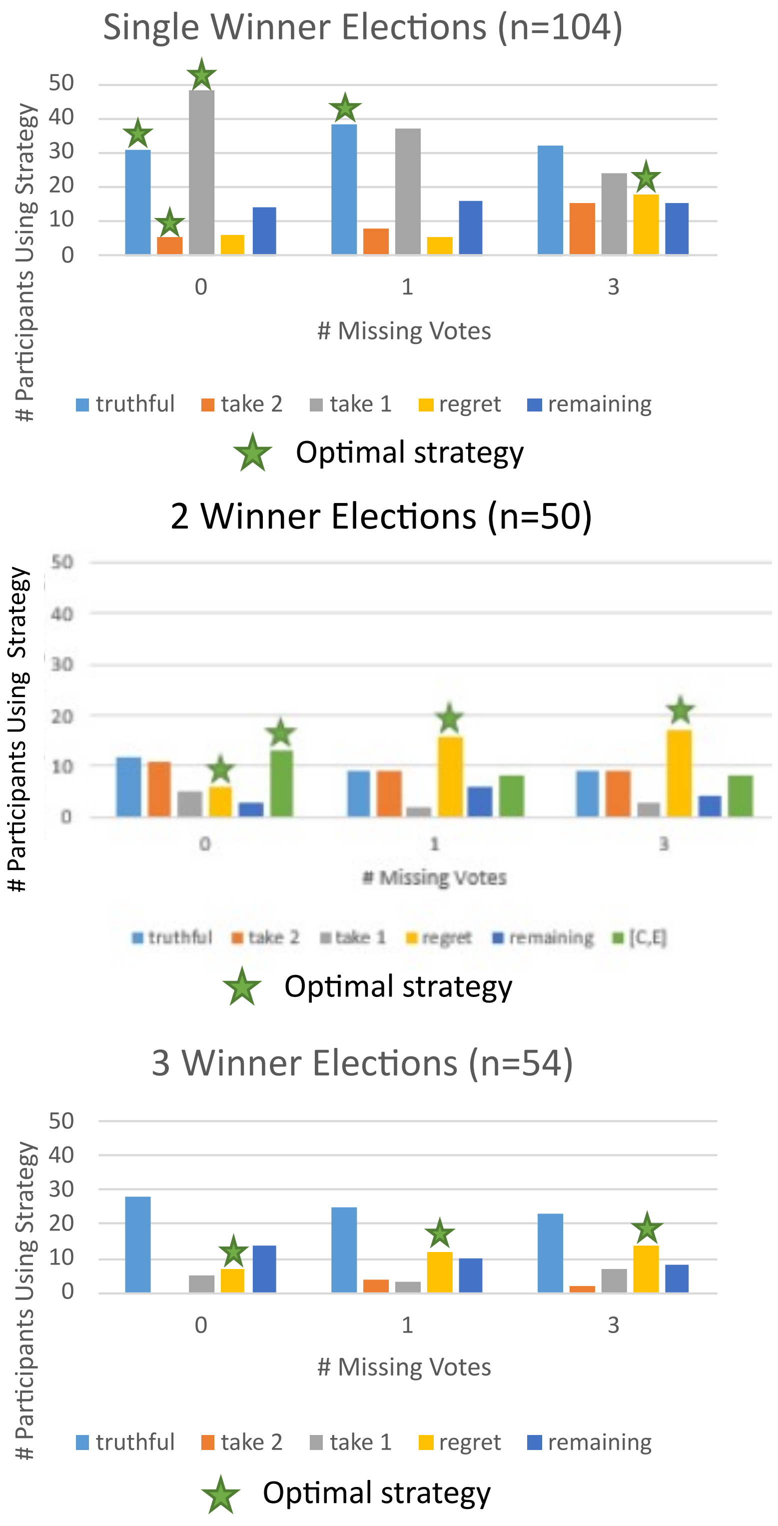

Figure 5.6: Scenario 3 (Disliked Candidate) Results. Maximizing strategies are marked with a star. People tended to vote truthfully in the 1 and 3 winner scenarios, whereas many voters in the 2-winner scenario tended to vote optimally, using regret minimization or just [C,E].

5.5.4 Scenario 4: Neutral Leading Candidate

In this scenario, tested if people would vote truthfully when a neutral candidate is leading, even in situations when *take the 2 best* was the maximizing strategy, e.g., when there is 1 winner with 3 missing votes or 3 winners with 0 missing votes (see Tables 5.12, 5.13). For this scenario, participants voted in single winner elections with 0, 1 or 3 missing votes, and in 1- 2- or 3-winner elections with 0 missing votes.

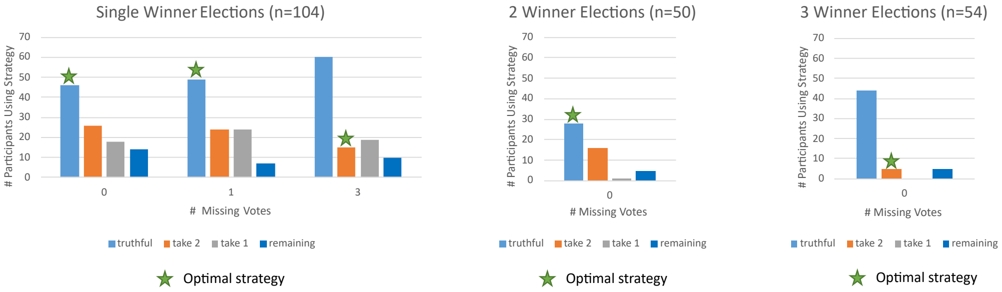

Figure 5.7: Scenario 4 Results. Maximizing strategies are marked with a star.

We found that when there were 0 missing votes, people's strategies changed significantly ($P < 0.005$) depending on the number of winners in the election. Overall, being truthful dominated the other strategies (1-winner: 49.7%, 2-winner: 56%, 3-winner: 81.5%), especially in the 3-winner election, even though it would have been better to use *take the 2 best* in this instance. In fact, *take the 2 best* represented only 9.2% of votes in the 3-winner election. There was no significant difference in people's strategies as the number of missing votes increased. Being truthful was the dominant strategy, even in the 3-winner election, where using take the 2 best had a higher expected utility (see Figure 5.7).

5.5.5 General Discussion

Behavioral results showed some distinct patterns of voting across all scenarios. The majority of participants *did not vote using a strategy that maximized expected utility*, especially in the 1-winner (25.6% maximized) and 2-winner (38.4% maximized) conditions. In the 3-winner condition, 49.6% voted using 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 still tended to use a *take the X best* heuristic, and this captured a significant portion of their responses (1-winner: 50.6%, 2-winner: 43.8%, 3-winner: 34.4%).

We found that people generally used different heuristics in different scenarios, and as the numbers of winners changed. For example, in Scenarios 1 (trivial utilities, see Section 5.5.1) and 3 (disliked candidate, see Section 5.5.3), a significant portion of voters did not vote completely truthfully, and chose to use 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 were choosing a heuristic that aligned with the number of winners. However, in Scenarios 2 (dominated preferences, see Section 5.5.2) and 4 (neutral leader, see Section 5.5.4), 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.

We found that people were not very sensitive to changes in uncertainty. In Scenarios 1 and 4, *participants' behavior did not significantly change as the number of missing votes increased, even when this resulted in using a non-optimal strategy.* In Scenario 2, voters in the 2-winner elections were sensitive to the fact that they had some chance of electing a candidate when there were 3 missing votes, leading to fewer abstentions in that condition. In Scenario 3, some voters were able to identify that the underlying optimal strategy changed, increasing the number of voters using *regret minimization* from 5.8% when there were 0 missing votes to 17.3% when there were 3 missing votes.

5.6 Summary and Future Directions

In this chapter, we showed how specific underlying features of voting environments affect how well different heuristics perform compared to maximizing expected utility. In particular, we looked at heuristic strategies including *truthful*, *take the X best*, and *regret minimization* in scenarios where less preferred, neutral and disliked candidates lead in an election. In a behavioral experiment of 104 subjects on Mechanical Turk, we showed how the number of winners, the current leader(s), and the number of missing votes affects the heuristic strategies that people use in approval voting. In particular, such as when neutral candidates were leading, people tended to be completely truthful. When people did not vote completely truthfully, they tended to vote sincerely, using a *take the X best* heuristic and were generally not very effective at choosing a heuristic that maximized their utility.

Our work provides key insights on human behavior in voting environments that can lead to more realistic simulation tools and more accurate predictions of election outcomes when approval voting is used. Our study can also inform the design of automated decision support systems by providing evidence about which heuristics humans may be inclined to use in different contexts and help in designing suggestions that take these behavioral aspects into account. Heuristics adopted by humans can also inspire the design of fast and frugal algorithms for tackling problems of prohibitive size or complexity.

While the results presented in this chapter provide insights into the use and effectiveness of certain heuristics in approval voting, there are many other voting rules and heuristics. It would be interesting to continue exploring 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). In decision science, taxonomies have been created to show which heuristics may be more or less useful in which environment [119]. We believe that a similar approach could prove beneficial to our understanding of voting heuristics, which is important for factoring them into a more realistic analysis of the voting rules.