

Modeling People’s Voting Behavior with Poll Information

Roy Fairstein and Adam Lauz
Ben-Gurion Univ. of the Negev, Israel

Reshef Meir
Technion, Israel Institute of
Technology

Kobi Gal
Ben-Gurion Univ. of the Negev, Israel,
University of Edinburgh, U.K.

ABSTRACT

Despite the prevalence of voting systems in the real world there is no consensus among researchers of how people vote strategically, even in simple voting settings. This paper addresses this gap by comparing different approaches that have been used to model strategic voting, including expected utility maximization, heuristic decision-making, and bounded rationality models. The models are applied to data collected from hundreds of people in controlled voting experiments, where people vote after observing non-binding poll information. We introduce a new voting model, the Attainability-Utility (AU) heuristic, which weighs the popularity of a candidate according to the poll, with the utility of the candidate to the voter. We argue that the AU model is cognitively plausible, and show that it is able to predict people’s voting behavior significantly better than other models from the literature. It was almost at par with (and sometimes better than) a machine learning algorithm that uses substantially more information. Our results provide new insights into the strategic considerations of voters, that undermine the prevalent assumptions of much theoretical work in social choice.

KEYWORDS

[Economic Paradigms] Social choice theory; [Economic Paradigms] Behavioral game theory; [Agent Societies and Societal Issues] Coordination and control models for multiagent systems

ACM Reference Format:

Roy Fairstein and Adam Lauz, Reshef Meir, and Kobi Gal. 2019. Modeling People’s Voting Behavior with Poll Information. In *Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), Montreal, Canada, May 13–17, 2019*, IFAAMAS, 9 pages.

1 INTRODUCTION

Voting is a commonplace tool for group decision making, used in political elections, in professional committees, in local assemblies, and also in online platforms such as Doodle.com and robovote.org. While there is general consensus that people vote strategically, understanding individual voting behavior is a challenging open question. Due to inherent uncertainty about other people’s votes, the strategies that people apply are far from obvious.

Researchers in economics, political science, and more recently in AI and computational social choice, have suggested various models to represent and reason about voters’ decision making under uncertainty. These include models of *utility maximization*, *heuristic*, and *bounded rational* (see below). In a recent paper, Meir, Lev and Rosenschein [19] suggested different criteria for evaluating models

of strategic voting, that included “theoretic criteria” (such as generality and discriminative power among actions), “behavioral criteria” (such as cognitive plausibility), and “scientific criteria” (such as alignment with empirical data). Theoretical analysis of voting models are abundant (see Sec. 2.1), and in this work we focus on the latter two kinds of criteria.

Research Goal. The goal of the paper is to study strategic choices of human voters, and in particular to test how their individual behavior fits different types of models. We use real world data from controlled experiments in which human voters either faced a non-binding poll or played a strategic game versus other people.

We follow Wright and Leyton-Brown [33], who separate collected data from strategic games into training and test data, and compared the predictive power of strategic decision-making models based on their predictive performance on the test data. If a certain model predicts well the behavior of many voters, this is an important indication for the plausibility of this model. Prediction is a standard evaluation metric in behavioral economics [8, 12]. Thus it should be considered in addition to its theoretical properties, cognitive limitations of the voters and so on. By further analyzing which models succeed and when they fail, we hope to better understand the considerations that guide people’s strategic choices.

1.1 Theoretical Models

We briefly describe common approaches for modeling strategic voting behavior in the theoretical literature.

Expected utility maximization. A rational voter maximizes her expected utility with respect to a probability distribution over the actions of the other voters. The distribution itself may be given exogenously (e.g., by a poll), or derived via equilibrium analysis from the uncertain preferences of the other voters. Such models were developed mainly in the economics literature and are sometimes known as the “calculus of voting” [20, 21, 27]. A somewhat different model was suggested by Bowman et al. [6] for voting on multiple binary issues. The model explicitly estimates the “attainability” of each issue (the probability it gets a majority of the votes), and uses this estimate to calculate the expected utility of every possible combinatorial vote.

Heuristic decision-making. A voter uses some function that maps any given situation to an action. The voter is not assumed to be rational, and may not even have a cardinal utility measure or an explicit probabilistic representation of the different outcomes. For example, a voter following the k -pragmatist heuristics behaves as if only the k leading candidates are participating [26].

Bounded rationality. A voter makes a rational strategic decision based on a subjective, rather than accurate, belief. These models present a mid-point between utility maximization and heuristics. One example of such a model is *local dominance* [19], which assumes

Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), N. Agmon, M. E. Taylor, E. Elkind, M. Veloso (eds.), May 13–17, 2019, Montreal, Canada. © 2019 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

that each voter derives a set of possible outcomes based on a poll, and then selects a non-dominated action within these outcomes. Similar probability-free approaches were followed in [11, 22].

All the models we work with assume that individual voters may behave differently, but each one follows a deterministic, consistent voting strategy. They are unable to perfectly explain or predict votes that have a random component or where voter behavior changes over time. Therefore, if the models can still explain the data, it would mean that noise and learning only play a secondary role in people’s voting decisions. We go back to this point when analyzing the results.

Other models, such as *quantal-response equilibrium* (QRE) [16, 17], *sampling equilibrium* [24] or *trembling hand equilibrium* [23] assume voters act stochastically. Evaluation of such models is much easier on the aggregate rather than individual level (see below), and therefore they are outside the scope of the current paper.

1.2 Previous Empirical Work

While the literature is abundant with voting experiments, the vast majority analyze voter behavior in political or organizational elections [1, 5, 13, 25, 31]. These studies test how well historic election results fit various game-theoretic models, without any consideration of individual votes. Further, each voter makes a single strategic decision, and her true preferences are typically unknown.

Explaining aggregate voting behavior. Some controlled experiments track voters’ decisions in different situations (e.g., [14, 30, 32]). Most of these experiments included groups of 12-70 subjects who played a repeated strategic voting game, knowing the preferences of others but not how they are going to vote. Yet these papers focused on how well *aggregate behavior* fits the equilibrium models. For example, in the QRE model used in Tyszler and Schram [30], voters are assumed to vote for low utility candidates with some probability, which is determined by a parameter of the model. A model is considered to be an adequate explanation for a dataset if there are some parameters that result in a similar distribution of votes (e.g., a similar rate of strategic compromise) to the one observed in the data. However, such models are not designed to track individual behavior and whether it is consistent. For example, a 20% rate of strategic compromise could result either from a small group that is consistently strategic, from all voters being occasionally strategic, or even from some random component in the behavior.

Explaining individual voting behavior. Blaise et al. compare individual behavior to rational models, inferring voters’ parameters from verbal surveys [5] or from carefully designing the conditions of a controlled experiment [4]. However, they focused on voting with two candidates, where the only strategic decision is whether to vote or abstain.

We emphasize that all of the above work tested how well empirical data fits the theory *in retrospect*, without dividing the data into separate training and test sets. This approach may cause overfitting, especially in complex models with many parameters.

Tal et al. [28] study voter behavior under poll information, but did not compare to any existing decision model, neither suggested a new one. They demonstrated empirically that there are voters exhibiting different behaviors: in particular, truthful voters, voters

who compromise strategically, and voters that tend to vote for the poll leader (“leader biased”). These findings were part of our inspiration to focus on understanding individual votes.

Our contributions and results are as follows:

- (1) We provide a new voting model called the Attainability-Utility (AU) heuristic. The model is based on the model suggested by Bowman et al. [6] which considers the attainability of a binary issue (the probability it is accepted in referendum) when computing the expected benefit of a vote. The AU heuristic extends this model in two ways, first by considering multi-candidate voting settings, second by including a parameter that measures the tradeoff between how much the voter values candidates’ attainability given the poll information versus their *utility* (if selected).
- (2) We collect the strategic decisions of 520 people in voting experiments with three candidates, where participants each play up to 36 rounds, each round with different poll information and preferences (more than 14,000 decisions in total). All of the data and code will become available for the public using repositories such as votelib.org.
- (3) Using behavioral data from our experiments as well as from Tal et al. [28] and Tyszler and Schram [30], we compare the performance of the AU approach to that of theoretical decision models from the literature and to benchmarks set by off-the-shelf machine learning algorithms.
- (4) Our results show that the AU model outperforms all other voting models, some of them by a large margin, and gets close to the benchmark set by machine learning algorithms. In particular, AU is able to capture much of the behaviors described by the models in Sec. 2.1. Most errors in the prediction of the AU model can be attributed to participants who played few number of rounds, demonstrated random behavior, and/or changed their strategy during the experiment.
- (5) Our main insight from the success of the AU model is that people *independently* evaluate each candidate and use simple substitutes to probabilistic calculations. These findings are in line with the more general research on decision making under uncertainty, and at odds with the underlying assumptions of most models from the social choice literature.

This is the first paper to provide an empirical evaluation of theoretical decision making models on individual voter behavior under poll information, and the first to test the predictive performance of any voting model in general. Understanding the strategic decisions made by voters of different types, is crucial to the development, analysis and application of voting rules in strategic environments, and can inform the design of agents for making voting decisions with other people [3, 34].

2 PRELIMINARIES

We consider a single voter who faces a decision, to vote for one of several candidates C . We use the Plurality rule which collects the total number of votes for each candidate, and returns the candidate(s) with the largest number of votes.

The voter has a cardinal utility function $u : C \rightarrow \mathbb{R}$, where $u(c)$ is the utility of the voter if candidate c wins (different utility for each candidate). In case of a tie with multiple winners $W \subseteq C$, the

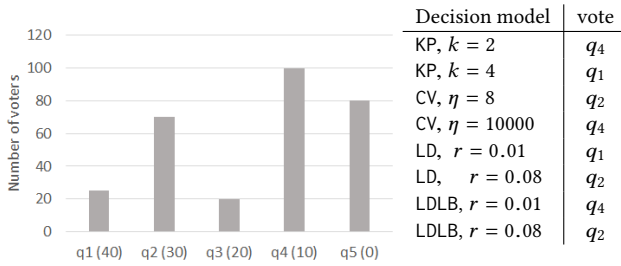


Figure 1: Left: Poll s from Example 2.1. The utility of each candidate to the voter appears in brackets, and the height of the column is the number of votes. Right: The candidate selected by each decision model.

utility to the voter is $u(W) = \frac{1}{|W|} \sum_{c \in W} u(c)$. Denote by $\mathcal{U}(C)$ the set of all utility functions over the set C .

Prior to her vote, the voter is faced with non-binding poll information that reflects the popularity of each candidate. Formally, the poll is a vector $s \in \mathbb{N}^m$, where $s(c)$ is the number of voters expected to vote for c . Denote $n = \sum_{c \in C} s(c)$.

We index the candidates q_1, q_2, \dots from the perspective of the voter, where q_1 is the most preferred, then q_2 , and so on.

A *decision model* (for Plurality with m candidates and a poll) is a function $M : \mathcal{U}(C) \times \mathbb{N}^m \rightarrow C$. Here, $M(u, s) \in C$ is the vote of a voter with utility function u , using decision model M given a poll s . We use a superscript for the name of the decision model, and subscripts to denote voter-specific parameters, if relevant. For example, a voter who is always truthful regardless of the poll follows the decision model $M^{\text{Truth}}(u, s) := \arg \max_{c \in C} u(c)$, which is q_1 .

To illustrate we introduce a running example with 5 candidates, and specify which candidate the voter will choose under every decision model.

Example 2.1. The set of candidates is $C = \{q_1, \dots, q_5\}$, and the voter’s utility is described by the vector $u = (40, 30, 20, 10, 0)$ (preferences are lexicographic). Poll scores are given by $s = (25, 70, 20, 100, 80)$, where $n = 295$ voters. Figure 1 (left) shows the poll scores of all candidates graphically.

2.1 Decision models from the literature

We describe decision-making models of voting behavior from the literature. Figure 1 (right) shows the result of using each of the decision models, applied to the voting decision in Example 2.1.

k-pragmatist (KP): Let $B_k(s)$ contain the k candidates with highest score in s , Reijngoud and Endriss [26] formalized the k -pragmatist heuristic (following early work such as [7] which selects the most preferred candidate among k candidates with highest score in $B_k(s)$):

$$M_k^{\text{KP}}(u, s) := \arg \max_{c \in B_k(s)} u(c).$$

We allow k to be an individual parameter that differs from voter to voter. When $k = 1$ the rule always selects the leader of the poll, and for $k = m$, $M_m^{\text{KP}} \equiv M^{\text{Truth}}$. In Figure 1 for $k = 2$ the voter will vote for the candidate that is most preferred among the two leading candidates (q_4 and q_5). For $k = 4$, the voter considers all candidates

except q_3 as possible winners, and will vote for her most preferred candidate q_1 .

Calculus of Voting (CV): The calculus of voting suggests that a rational voter always votes in a way that maximizes her expected utility [21, 27]. The complications of the model usually arise from the fact that the voter is assumed to know the other voters’ preferences, and uses an equilibrium model to predict their votes. We consider a simpler version where the distribution of votes is given exogenously [20], as is the case with poll information.

We denote by $\mathcal{D}(s)$ the distribution on the actual candidate scores, conditional on poll scores s . We say a voter is *pivotal for candidate y over x* , if voting for y makes y a joint or unique winner, whereas any other vote results in the victory of candidate x . Denote by $P_{s, \mathcal{D}}(x, y)$ the probability that the voter is pivotal for y over x given the distribution over candidate scores \mathcal{D} induced by poll s . A voter following the calculus of voting (CV) model maximizes her expected utility:

$$M_{\mathcal{D}}^{\text{CV}}(u, s) := \arg \max_{c \in C} \sum_{c' \neq c} P_{s, \mathcal{D}}(c', c) (u(c) - u(c')).$$

To make the model concrete, we determine a specific distribution \mathcal{D} in a way that depends on the score of the candidates in the poll s . We use $P_{s, \eta}$ as a shorthand for $P_{s, \mathcal{D}}$ when $\mathcal{D}(s)$ is a multinomial distribution with η voters, and the probability for sampling a vote for each candidate c is $s(c)/n$. When $\eta = n$ (i.e., the true number of voters), this means that M_{η}^{CV} selects the candidate that exactly maximizes the voter’s expected utility given the true distribution over candidate scores. However, the M_{η}^{CV} decision model allows for a more flexible, bounded-rational decision: when $\eta < n$ the voter overestimates her true pivot probability, and thus her influence on the outcome, whereas $\eta > n$ means that she underestimates her influence. In Figure 1 when $\eta = 10000$ the resulting vote is q_4 and when $\eta = 8$ the resulting vote is q_2 .

Local Dominance (LD): Under the Local dominance model [18, 19], a bounded-rational voter has an ‘uncertainty parameter’ r . Meir et al. [19] characterize the set of undominated candidates $U(s, u, r)$ in poll s for a voter with utility u and parameter r :

- The set of Possible Winners PW includes all candidates whose score in s is at least $\max_{c \in C} s(c) - 2r \cdot n$.
- If $|PW| \geq 2$, then the undominated candidates are all candidates in PW except the least preferred.
- If $|PW| = 1$, then all candidates are undominated.

The decision model of such a voter selects the most preferred undominated candidate, if more than one exists:

$$M_r^{\text{LD}}(u, s) := \arg \max_{c \in U(s, u, r)} u(c).$$

In Figure 1 we see that for $r = 0.01$ the voter believes that the poll is very accurate (the score of each candidate may change by at most $r \cdot n < 3$ votes), and there is only one possible winner ($PW = \{q_4\}$). In this case, all candidates are undominated and the voter remains truthful ($M_{0.01}^{\text{LD}}(u, s) = q_1$). When $r = 0.08$, the voter believes that the poll is not very accurate and $PW = \{q_2, q_4, q_5\}$. In such a case both q_2, q_4 are undominated and $M_{0.08}^{\text{LD}}(u, s) = q_2$.

Local-Dominance with Leader bias (LDLB): Inspired by the findings of Tal et al. [28] on ‘leader bias’, we modify the local




Figure 2: The attainability A_β for different values of β in a poll with 3 candidates.

dominance model to allow such behavior: when the voter is certain that there is only one possible winner ($|PW| = 1$), she simply votes for the leader (instead of truthfully), i.e., $M_r^{\text{DLB}}(u, s) := M_r^{\text{D}}(u, s)$, and otherwise $M_r^{\text{DLB}}(u, s) := PW$. In Figure 1 we see that this model acts similarly to the LD model. However when there is only one possible winner, a voter following the LDLB model will vote for the leader (d in this case).

Attainability (AT): Bowman et al. [6] provide a model for voting over multiple binary issues. The attainability of issue j is a measure of certainty that the eventual number of votes cast for j will reach the majority threshold required for approval. It is defined as

$$\hat{A}_\beta(j, s) := \frac{1}{\pi} \arctan(\beta \cdot (s_j - \frac{1}{2})) + \frac{1}{2},$$

where s_j is the expected number of votes in favor of issue j , and β is a voter-specific parameter.

The “candidates” considered in Bowman et al.[6] are all possible subsets of issues, i.e., $C = 2^{\{1, \dots, k\}}$, where w.l.o.g. the voter gains some nonnegative utility u_j from each issue j being approved. Then the utility of a candidate $c \subseteq \{1, \dots, k\}$ is the sum of utilities of all issues in c , and its attainability $\hat{A}_\beta(c, s)$ is defined as the product of $\hat{A}_\beta(j, s)$ for all $j \in c$, and $1 - \hat{A}_\beta(j, s)$ for all $j \notin c$. The voter selects the candidate c that maximizes the product of its attainability and utility ($\hat{A}_\beta(c, s) \cdot u(c)$).

To adapt the decision model to Plurality voting with m candidates, we re-define the attainability function as $A_\beta(c, s) := \frac{1}{\pi} \arctan(\beta \cdot (s(c) - \frac{1}{m})) + \frac{1}{2}$, and define the attainability choice function (AT) as

$$M_\beta^{\text{AT}}(u, s) := \arg \max_{c \in C} A_\beta(c, s) \cdot u(c).$$

Figure 2 shows how β affects the attainability score. Candidates that are tied have the same attainability. As shown by the figure, high β means that a small advantage in score translates to a large gap in attainability.

3 THE ATTAINABILITY-UTILITY (AU) MODEL

Bowman et al.’s AT model allows voters some flexibility in how they estimate attainability using the β parameter. However it assumes the same model for each voter. We extend the attainability model by an additional parameter that lets each voter choose a different tradeoff between the attainability and utility of candidates. To this end we define the *Attainability-Utility* (AU) decision rule as

$$M_{\alpha, \beta}^{\text{AU}} := \arg \max_{c \in C} \left((\varepsilon + u(c))^\alpha \cdot (A_\beta(c, s))^{2-\alpha} \right),$$




Figure 3: The AU scores of all five candidates from Example 2.1 for $\beta = 5$ and different values of α .

where ε is a small constant added to handle 0 utility (ε can also be used as a parameter to control the utility range).

Intuitively, the α parameter trades-off the relative importance of attainability and utility, where $\alpha = 0$ means the voter always selects the candidate with maximal score, and $\alpha = 2$ means the voter is always truthful. Figure 3 shows how the relative score (and the selected candidate) changes as we increase α . When α is small, AU will prefer q_4 as it has more votes (higher attainability) and when it is large the AU will prefer q_1 as it got higher utility. Note that we get the AT model as a special case when setting $\alpha = 1, \varepsilon = 0$. We further discuss the meaning of these parameters in Sec. 6.1.

4 METHODOLOGY

Datasets. We evaluated the different models described above on several datasets as follows.

	D32	D36	TMG15	TS16
# participants	187	335	437	144
# voters in poll	1000	1000	8 to 10000	12
# rounds	up to 32	up to 36	up to 20	40
# instances	4886	9478	8011	5760

Three of the datasets (D32, D36 and TMG15) were collected using the framework of Tal et al. [28], in which voters played multiple one-shot voting rounds. A snapshot of the GUI used for this setting is shown in Figure 4.

Each round included a single human participant, that is automatically assigned preferences over candidates, observes a noisy “poll” with the expected votes of the entire population (e.g., 1000 voters), and then votes once.

The outcome of the round was generated by sampling each of the other votes i.i.d using the poll scores as the distribution (e.g. in Fig. 4 we sample 102 “voters”, each of which votes Blue w.p. $\frac{34}{102}$). Participants were only informed that the poll was inaccurate, but not on the exact distribution. The final score of each candidate and the outcome were shown in the end of each round. In all datasets the reward for participants was determined by the position of the winning candidate in their preferences, using the average reward if there was more than one winning candidate.




Figure 4: A snapshot from the experiment (taken from Tal et al. [28]). Blue is the most preferred candidate q_1 , thus if Blue wins, the participant gets 20 coins (10 cents). The bars show the poll scores s_1, s_2, s_3 .

We used this framework to generate datasets D32 and D36, presenting the participants with a different poll each time. All participants were recruited via the Amazon Mechanical Turk platform. The reward was R_j for each round where q_j was elected, where $R_1 > R_2 > R_3$. For most of the participants, we set $R_1 = 10\text{¢}, R_2 = 5\text{¢}, R_3 = 0\text{¢}$. For some participants we varied the rewards.

The dataset TS16 was generated by Tyzslar and Schram [30]. Here, every voting round was a 12-player complete information game with dictated preferences over 3 candidates, and the outcome was the result of all actual votes rather than artificial samples. We used the other 11 voters’ true top preferences (which are visible) as a true “poll” input to the different decision models.

In all datasets, only when q_1 is ranked last at the poll, the voter may have a monetary incentive to vote for q_2 . There is never a monetary incentive to vote for q_3 .

Random Forest (RF) Benchmark. We applied off-the-shelf machine learning algorithm to build predictive models of voting behavior. We used two types of features: those relating to the particular voting round (examples: the gap $s_1 - s_2$ between the two leaders of the poll, the number of votes s_2 in the poll, the winning candidate in the poll); and those aggregating the behavior of the voter (examples: the frequency that the voter chose q_1, q_2 and q_3 in the training set, the frequency of a strategic compromise, and the number of dominated actions). Using these features, we compared the performance of black-box prediction models on the D32 dataset.

We compared Random Forest, Neural Network, AdaBoost algorithms, CART (Decision Tree), Support Vector Machines and Logistics Regression.¹ The best performance was exhibited by a random forest ensemble model using 100 weak trees as subclassifiers, and a Gini splitting criterion. We thus used this algorithm (henceforth, *RF*) as our benchmark.

Evaluation. We used a ten-fold cross validation method. We divided the data of each voter into 10 folds (when possible). KP has only three parameter values. For the other models, we discretized the parameter space.

For each of the decision models KP, CV, AT, LD, LDLB and AU, we used a *basic fitting procedure* to train each model separately for each voter: 9 folds of data for this voter were used to fit the

¹We used the sklearn ensemble python package for this purpose [9]. The full list of features is available at https://github.com/AdamLauz/OneShotVoting/blob/master/Documentation/One_Shot_ML_features_description.pdf.

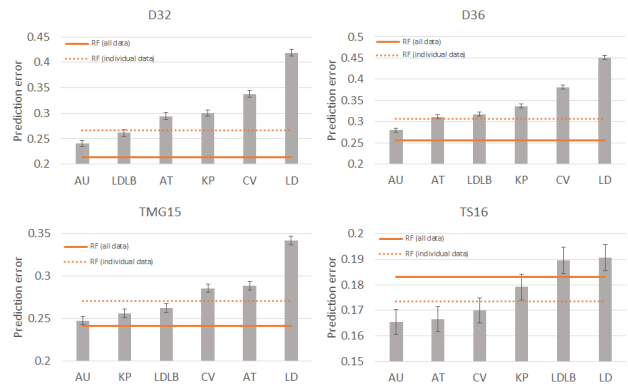


Figure 5: Prediction error for each of the decision models on all four datasets. The horizontal orange lines mark the performance of the RF benchmark when trained on the entire data (solid) or restricted to individual data (dotted).

parameters of the model, and applied the obtained model on the tenth fold to predict the voter’s actions. Since each voter has only few samples and the parameter space of each decision model is small, we used a brute-force search to find the best parameters for each model. For example, for the LD model, we found for each voter the parameter r such that M_r^{LD} agrees with the largest number of rounds in the training set.

The *prediction error* of a model is the number of wrong predictions on the test set, divided by the total number of rounds.

5 RESULTS AND ANALYSIS

Figure 5 shows the performance of all decision models on the datasets. We report the prediction error of the models, adding error bars of two standard deviations.

We can see that the AU model outperforms all other decision models, with the LDLB model second, and the models that ignore leader bias (CV, LD) far behind. These results are statistically significant in all datasets ($p < 0.05$) except in TS16, where there was no significant difference between the performance of AU and AT.

5.1 AU performance vs. the benchmark

Random Forest (RF) uses many features and can create an arbitrarily complicated model, which learns from the entire population of hundreds of voters rather than just from several individual samples. In addition, it uses temporal features and can thus in principle predict even behavior that changes over time. This is why we use RF as a benchmark that is supposed to be hard to beat. Even so, RF does not perform uniformly better than the behavioral models.

Figure 6 breaks down the error of the AU model by individual voters (we refer to the different colors in the next subsection). It demonstrates visually that AU beats RF for many individual voters (about a 100 out of 335), and that the advantage of RF is mainly due to a group of voters for which AU seems to perform substantially worse (those below the dashed line). Indeed it is possible that the AU model is appropriate for most voters but not for all (see discussion in Sec. 6.2).

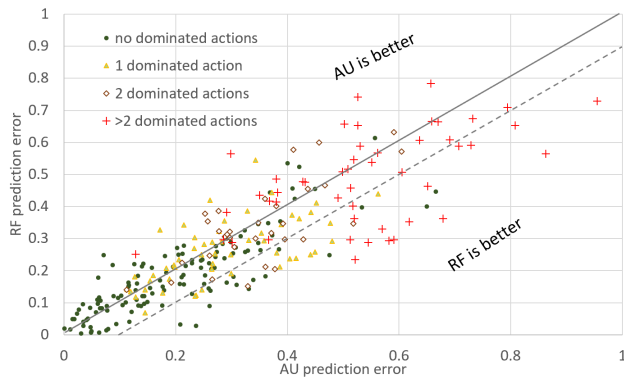


Figure 6: Prediction error of AU versus RF for each voter. Only voters with at least 16 rounds are shown. The prediction of RF for voters below the dashed line was better by at least 10 percentage point.

Another factor is the data used for learning: while AU and all other decision models fit their parameters for a particular voter solely based on her own behavior in other rounds, the black-box algorithms had access to votes of other voters as well. When restricted to learn only from the samples that belong to the same individual, the error of RF leaped dramatically (see dotted lines in Fig. 5). This indicates that information about the entire population could be exploited to further improve the behavioral models. In addition, the performance of RF (as well as the other black-box algorithms we tried) reduces more rapidly when we learn from a small or non-representative sample.

5.2 Where are the errors?

We analyze the factors that contributed most to prediction error, with a focus on the AU model.

Some voters are harder to predict. We say that a candidate is *dominated* (in a particular round) if there is another candidate that is associated with a higher score in the poll and is also more preferred by the voter. E.g. q_3 in Example 2.1 is dominated by q_2 . We count the number of dominated actions each voter performed throughout the experiment. Note that a dominated action is never predicted by any of the decision models we considered. It is hard to think of any rational justification for voting to a dominated candidate. We thus conjecture that dominated actions are indication for some random component in the behavior of the voter.

We classified voters by the number of times they used a dominated action. In Fig. 6 we can see that the number of dominated actions substantially affects prediction accuracy not just for AU but also for the benchmark RF (and in fact for all models). The prediction error of AU for voters who completely avoid dominated actions is less than 18%, and increases to almost 50% for voters with more than 2 dominated actions, indicating that their behavior is almost completely unpredictable. We emphasize that every additional dominated action results in *more than one* (about 1.6-3) predictions errors. This, together with the low performance of RF, corroborates our conjecture that dominated actions are merely an indication for noisy or random voting patterns.

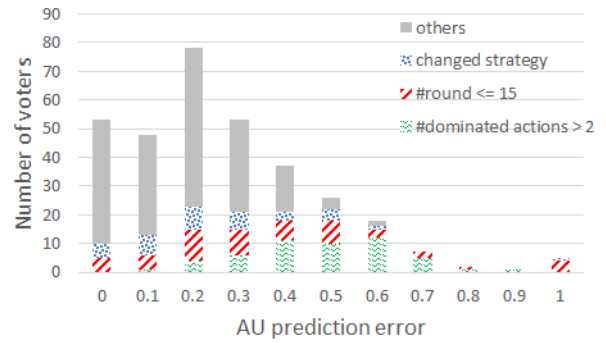


Figure 7: A histogram of all D36 voters by their AU prediction error. We colored groups of voters for which certain conditions apply.

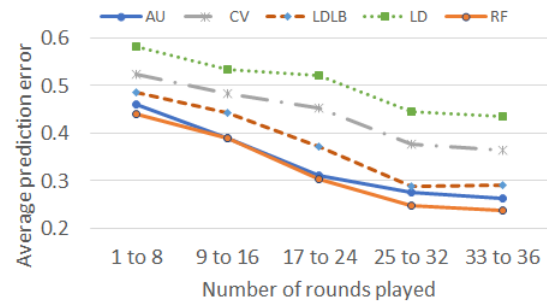


Figure 8: Prediction error as a function of the number of rounds in D36.

Another factor that substantially affects prediction error is the number of rounds that a voter has played, where prediction error for voters who completed fewer rounds is much higher. See Fig. 8 which lists error as a function of number of rounds per voter. A likely explanation is that these voters are more prone to overfitting.

The histogram in Fig. 7 shows that for almost all voters where AU had high error, the reason was random behavior (indicated by dominated actions), or few rounds.

Behavior in some polls is hard to predict. The behavior in polls that present the voter with an obvious dilemma (e.g., when her favorite candidate is trailing behind) is naturally harder to predict. In Table 1, we classified all polls into 6 poll types, based on the the order of candidates' popularity in the poll.

poll type	D32	D36	TMG15	TS16
$q_1 > q_2 > q_3$	0.085	0.113	0.076	0.047
$q_1 > q_3 > q_2$	0.089	0.108	0.070	0.053
$q_2 > q_1 > q_3$	0.224	0.261	0.268	0.254
$q_3 > q_1 > q_2$	0.202	0.268	0.268	0.338
$q_2 > q_3 > q_1$	0.233	0.250	0.258	0.296
* $q_3 > q_2 > q_1$	0.363	0.403	0.419	0.470

Table 1: AU error for each poll type. The order reflects the popularity of each candidate in the poll.




Figure 9: AU and AU_ϵ accuracy in D32. AU_ϵ is identical to AU, except that ϵ is used as another parameter.

In the scenario most difficult to predict (where $s(q_3) > s(q_2) > s(q_1)$, marked with *), the poll order is reversed to the preference order of the candidates, and all three actions are frequently selected by the voters. For this case the prediction error of the AU model is above 35% (and remains high even if we focus on voters who played all rounds and avoided dominated actions). The results of the other models behaved similarly. Note that while in Fig. 5 prediction accuracy varies considerably between datasets, this is explained by the frequency of different poll types in each dataset (see Table 1).

We emphasize that in TS16, most of the rounds people played as part of the majority group (see first two rows of Table 1), and thus faced a trivial decision where all models predicted the same. This could be the reason we did not obtain statistically significant results on this dataset, and also explains the poor performance of RF (which had few non-trivial rounds to learn from).

Negative reward is harder to predict. In D32, we varied the reward R_3 , to see the effect of positive/zero/negative reward (see Fig. 9). Higher reward R_3 results in higher accuracy (the only statistically significant difference was between negative and zero reward, as we only varied the reward for 25 participants).

A closer look revealed the reason for the excess failures: with a negative reward, $u_3 + \epsilon$ is still negative, and thus the AU model would never select q_3 . Perhaps surprisingly, participants do not care much about R_3 being negative. Adding ϵ as a third optimization parameter that may get higher values (so that $u_3 + \epsilon > 0$) completely negates that effect, as can be seen by the striped columns in Fig. 9. We also varied the reward R_2 to get convex and concave utility scales rather than linear. While higher R_2 does lead to more frequent votes to q_2 , we did not observe a consistent effect on prediction error.

5.3 Subjective reporting by participants

From each participant in our experiment (datasets D32 and D36), we asked to report their subjective answers about how well they understood the instructions; which strategy they used in the study; and whether they changed a strategy during the game.

Descriptions of strategies. Some of the participants described strategies that are similar to the models we tested from the voting literature. Some primary examples appear below.

- “I tried to vote for the person most likely to beat the candidate that would give me no coins.” Describes KP with $k = 2$.
- “I voted for either my first or second priority candidate. I was more likely to vote for the one that appeared to have the highest probability of winning.” - Describe behavior similar to AT.

- “My strategy was to mainly vote for who was leading except when it was a close race and then I voted for who would earn me the most points.” - describes LDLB with low r .

Interestingly, people did not adhere to their reported strategies in all rounds, and often their behavior was predicted more accurately by a different model than the one they verbally describe.

For example, some people who explicitly declared that they would not vote for the least preferred candidate did in fact choose this option. It is not clear whether this results from noisy behavior, from changing the behavior over time, or from poor self-reflection.

Do voters use consistent strategies? Identifying individual changes in strategy from the data is very difficult with only a handful of samples per voter. However from the subjective self-reports, about 63% of those who responded in D36, answered that they did not change their strategy, whereas only 18% did.²

There was a strong correlation between how well people understood the instructions (by their self-report) and their consistency: in D36, more than 80% of those who reported perfect understanding, claimed they kept their strategy. Results in D32 were similar.

Recall the questions from the introduction about consistency and predictability. The level of (reported) consistency strongly affected the empirical error: the average error of AU for consistent voters in D36 was about 24% vs. 30% for inconsistent ones (and 22% vs. 34% in D32). This can also be partly seen in Fig. 7, where voters who reported strategy change (dotted blue) are responsible for slightly more errors.

6 DISCUSSION

Regenwetter et al. [25] observe that “...individual choice research finds actors to behave **worse than normative theory requires**, whereas the sparse empirical research on social choice appears to suggest that electorates may **outperform normative expectations**.” However, most research they refer to considered *aggregated behavior*, as discussed in the early sections.

Our Attainability-Utility (AU) model explains well (and in particular much better than calculus of voting) the behavior of most subjects in the data, except those with inherent inconsistencies in their actions. This partly settles the discrepancy observed by Regenwetter et al.: on the individual level, most voters follow AU or other heuristics that do not maximize expected utility, just like decision makers in other domains, even if on the aggregate level the vote distribution can be explained by more rational theories like calculus of voting [5, 14], or quantal response equilibrium [30]. Interestingly, quantal response can account for the frequency of dominated actions at the aggregate level, even if it cannot predict when a particular action will be dominated. We may therefore get a more complete picture of voters’ behavior by combining individual and aggregate analysis (see also future work below).

6.1 Is AU cognitively plausible?

There are two seemingly “irrational” components in the AU model (both inherited from Bowman [6]), that become apparent when we compare it to the “rational” Calculus of Voting method. The

²The others provided an answer that could not be easily classified, e.g. “It took me a few rounds to get the hang of it.”

first is the fact that the voter assesses the chances of each *candidate to win*, rather than of each possible *tie*. The second is that this chance is estimated using a somewhat arbitrary transformation of the candidate’s score (the logit-shaped “attainability function”), rather than by explicit probabilistic calculations.

Both observations are much less surprising when we recall Kahnemann and Tversky’s account of judgment under uncertainty [29]: they explain that people often use simple substitutes for probabilistic calculation, that require low cognitive effort. For example, rely on how *representative* each event is (in our case, the score of each candidate in the poll).³ Future experiments can further test this hypothesis by making candidates more prominent in other ways than higher score (e.g., using graphic features), and see if voters’ behavior can still be explained when we translate this to greater attainability. The extensive literature on the various heuristics people use to evaluate likely outcomes (e.g. [2, 10]) can also be used to develop better models of voting behavior.

It is worth mentioning that trying to use a simple substitute for probabilistic calculations was the main motivation behind the Local Dominance model [19], but Local Dominance (like the Calculus of Voting) still focuses on ties. The fact that AU better explains the behavior of most voters (and in particular that LD fails to predict leader-biased actions) suggests that perhaps even LD is too cognitively prohibiting. Indeed, except for KP (which is perhaps too simple), AU is the cognitively easiest heuristic to apply, as it *independently* evaluates each candidate. We note that the differences in cognitive burden become even more accentuated in elections with more candidates. We therefore expect the differences in performance to become more significant as well, and are currently collecting more data to test this hypothesis.

6.2 Are there voters of different types?

While AU had the best performance overall, there are many individual voters that are better predicted by one of the other models (not necessarily the leading one). This can be seen in the bottom bar in Figure 10 which shows the number of voters that were optimally predicted by each model. In case of ties we ‘split’ the voter among all leading models. To illustrate, although 129 voters were best predicted by the AU heuristic, almost as many (120 voters) were best predicted by the KP heuristic.

At this point we face a dilemma when trying to explain the reason that so many voters are better predicted by other models: One hypothesis (H1) is that AU can in principle account for the behavior of all voters, but is overfitting its parameters due to the small dataset of each voter. An alternative hypothesis (H2) is that there are indeed voters with different inherent behaviors that are better captured by other models, such as LDLB, KP and so on.

Some evidence for H2 is in the self reports where participants described distinct strategies. However, we believe there is stronger evidence for H1: first, the self reports are often inconsistent with the actual behavior, and AU in fact predicts well many of the voters who described specific strategies. Also, AU can in principle explain (for some parameter values) almost all voters in the data, but we often fail to select the optimal parameters due to the small training

³We tried a variation of the AU model, where attainability was replaced with the actual winning probability. This did not improve the performance of the model.




Figure 10: The bottom bar shows, for each model, the number of participants in D32 and D36 for which this model achieved the best accuracy (possibly tied with other models). The top bar shows the same information, for the subset of voters for which the best prediction error was at most 0.2.

sample so other models have fewer prediction errors. This can be seen in Fig. 8 where AU improves faster as voters have more samples to train on. Lastly, the advantage of AU becomes more clear once we focus on voters with low error (top bar in Fig. 10). These are the voters for which the selected model is more meaningful.

To better answer this question, richer datasets that better distinct between decision models should be generated.

6.3 Discussion and future work

Finding a model that perfectly explains the behavior of all voters is probably impossible. Yet, our AU model does well both on the “behavioral” and on the “scientific” criteria presented in [19]: It is a fairly simple and cognitively plausible model, that captures the behavior of most voters well enough to predict their individual actions in various situations, and even to compete with machine-learning algorithms that use hundreds of features from the entire population. This model trades-off the popularity of a candidate (as a proxy for its winning chances) and its utility to the voter.

Future voting models should be extended to allow behavior that changes over time in some predictable way. More importantly, deterministic decision models should be combined with stochastic ones like quantal response and trembling hand perfection [17, 23] to explain both consistent individual choices and random departures from those choices. New evaluation methods are needed for these combined aggregate and individual choices.

Our findings can inform the development new and better models for strategic voting, much like the PrefLib project [15] is contributing to the study of preference structure, as well as to inform the design of agents for making voting decisions with other people, which is a growing area of research [3, 34].

Most of the decision models we used, including the new AU heuristic, naturally extend to more candidates and other voting rules. We intend to run experiments in more diverse settings (e.g. more than 3 candidates). Those experiments can expose behaviors that do not exist in the current data, can help differentiate between the decision models, and serve as a benchmark for the development of new models.

7 ACKNOWLEDGEMENTS

This work was supported in part thanks to the Israeli Science Foundation grant number 773/16.

REFERENCES

- [1] Paul R Abramson, John H Aldrich, Phil Paolino, and David W Rohde. 1992. "Sophisticated" voting in the 1988 presidential primaries. *American Political Science Review* 86, 1 (1992), 55–69.
- [2] Maya Bar-Hillel. 1973. On the subjective probability of compound events. *Organizational behavior and human performance* 9, 3 (1973), 396–406.
- [3] Moshe Bitan, Ya'akov Gal, Sarit Kraus, Elad Dokow, and Amos Azaria. 2013. Social Rankings in Human-Computer Committees. In *AAAI*.
- [4] André Blais, Jean-Benoit Pilet, Karine Van der Straeten, Jean-François Laslier, and Maxime Hérroux-Legault. 2014. To vote or to abstain? An experimental test of rational calculus in first past the post and PR elections. *Electoral studies* 36 (2014), 39–50.
- [5] André Blais, Robert Young, and Miriam Lapp. 2000. The calculus of voting: An empirical test. *European Journal of Political Research* 37, 2 (2000), 181–201.
- [6] Clark Bowman, Jonathan K Hodge, and Ada Yu. 2014. The potential of iterative voting to solve the separability problem in referendum elections. *Theory and decision* 77, 1 (2014), 111–124.
- [7] Steven J Brams and Peter C Fishburn. 1978. Approval voting. *American Political Science Review* 72, 3 (1978), 831–847.
- [8] Eduard Brandstätter, Gerd Gigerenzer, and Ralph Hertwig. 2006. The Priority Heuristic: Making Choices Without Trade-Offs. *Psychological Review* 113, 2 (2006), 409–432.
- [9] Lars Buitinck, Gilles Louppe, Mathieu Blondel, Fabian Pedregosa, Andreas Mueller, Olivier Grisel, Vlad Niculae, Peter Prettenhofer, Alexandre Gramfort, Jaques Grobler, Robert Layton, Jake VanderPlas, Arnaud Joly, Brian Holt, and Gaël Varoquaux. 2013. API design for machine learning software: experiences from the scikit-learn project. In *ECML PKDD Workshop: Languages for Data Mining and Machine Learning*. 108–122.
- [10] Nick Chater, Joshua B Tenenbaum, and Alan Yuille. 2006. Probabilistic models of cognition: Conceptual foundations. *Trends in Cognitive Sciences* 10, 7 (2006), 287–291.
- [11] Ulle Endriss, Svetlana Obraztsova, Maria Polukarov, and Jeffrey S Rosenschein. 2016. Strategic voting with incomplete information. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*. AAAI Press, 236–242.
- [12] Ido Erev, Eyal Ert, Ori Plonsky, Doron Cohen, and Oded Cohen. 2017. From anomalies to forecasts: Toward a descriptive model of decisions under risk, under ambiguity, and from experience. *Psychological review* 124, 4 (2017), 369.
- [13] Dan S Felsenthal, Zeev Maoz, and Amnon Rapoport. 1993. An empirical evaluation of six voting procedures: do they really make any difference? *British Journal of Political Science* 23, 1 (1993), 1–27.
- [14] Robert Forsythe, Thomas Rietz, Roger Myerson, and Robert Weber. 1996. An experimental study of voting rules and polls in three candidate elections. *International Journal of Game Theory* 25, 3 (1996), 355–383.
- [15] Nicholas Mattei and Toby Walsh. 2013. PrefLib: A Library for Preferences <http://www.preflib.org>. In *ADT'13*. 259–270.
- [16] Richard D McKelvey and Thomas R Palfrey. 1995. Quantal response equilibria for normal form games. *Games and economic behavior* 10, 1 (1995), 6–38.
- [17] Richard D McKelvey and John W Patty. 2006. A theory of voting in large elections. *Games and Economic Behavior* 57, 1 (2006), 155–180.
- [18] Reshef Meir. 2015. Plurality Voting under Uncertainty. In *AAAI'15*.
- [19] Reshef Meir, Omer Lev, and Jeffrey S. Rosenschein. 2014. A Local-dominance theory of voting equilibria. In *ACM-EC'14*. 313–330.
- [20] Samuel Merrill. 1981. Strategic decisions under one-stage multi-candidate voting systems. *Public Choice* 36, 1 (1981), 115–134.
- [21] Roger B. Myerson and Robert J. Weber. 1993. A Theory of Voting Equilibria. *The American Political Science Review* 87, 1 (1993), 102–114.
- [22] S. Obraztsova, O. Lev, M. Polukarov, Z. Rabinovich, and J. S. Rosenschein. 2016. Non-myopic voting dynamics: An optimistic approach. In *Proc. of the 10th Multidisciplinary Workshop on Advances in Preference Handling (M-PREF)*.
- [23] Svetlana Obraztsova, Zinovi Rabinovich, Edith Elkind, Maria Polukarov, and Nicholas R Jennings. 2016. Trembling hand equilibria of plurality voting. *AAAI Press/International Joint Conferences on Artificial Intelligence*.
- [24] Martin J Osborne and Ariel Rubinstein. 2003. Sampling equilibrium, with an application to strategic voting. *Games and Economic Behavior* 45, 2 (2003), 434–441.
- [25] Michel Regenwetter, Moon-Ho R Ho, and Ilia Tsetlin. 2007. Sophisticated approval voting, ignorance priors, and plurality heuristics: A behavioral social choice analysis in a Thurstonian framework. *Psychological Review* 114, 4 (2007), 994.
- [26] Annemieke Reijngoud and Ulle Endriss. 2012. Voter response to iterated poll information. In *11th*. 635–644.
- [27] William H Riker and Peter C Ordeshook. 1968. A Theory of the Calculus of Voting. *American political science review* 62, 1 (1968), 25–42.
- [28] Maor Tal, Reshef Meir, and Ya'akov (Kobi) Gal. 2015. A Study of Human Behavior in Online Voting. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2015, Istanbul, Turkey, May 4-8, 2015*. 665–673. Full version available from <https://tinyurl.com/yezxxgoj>.
- [29] Amos Tversky and Daniel Kahneman. 1974. Judgment under uncertainty: Heuristics and biases. *science* 185, 4157 (1974), 1124–1131.
- [30] Marcelo Tyszler and Arthur Schram. 2016. Information and strategic voting. *Experimental economics* 19, 2 (2016), 360–381.
- [31] Karine Van der Straeten, Jean-François Laslier, and André Blais. 2013. Vote au Pluriel: How people vote when offered to vote under different rules. *PS: Political Science & Politics* 46, 2 (2013), 324–328.
- [32] Karine Van der Straeten, Jean-François Laslier, Nicolas Sauger, and André Blais. 2010. Strategic, sincere, and heuristic voting under four election rules: an experimental study. *Social Choice and Welfare* 35, 3 (2010), 435–472.
- [33] James R Wright and Kevin Leyton-Brown. 2010. Beyond equilibrium: predicting human behaviour in normal form games. In *AAAI*.
- [34] David Ben Yosef, Lihi Naamani-Dery, Svetlana Obraztsova, Zinovi Rabinovich, and Marina Bannikova. 2017. Haste makes waste: a case to favour voting bots. In *Proceedings of the International Conference on Web Intelligence*. ACM, 419–425.