A Study of Human Behavior in Online Voting

Maor Tal
Dept. of Information Systems Engineering
Ben-Gurion University, Israel
maorta@post.bgu.ac.il

Reshef Meir
School of Engineering and Applied Sciences
Harvard University, USA
rmeir@seas.harvard.edu

Ya'akov (Kobi) Gal
Dept. of Information Systems Engineering
Ben-Gurion University, Israel
kobig@bgu.ac.il

ABSTRACT

Plurality voting is perhaps the most commonly used way to aggregate the preferences of multiple voters. The purpose of this paper is to provide a comprehensive study of people's voting behaviour in various online settings under the Plurality rule. Our empirical methodology consisted of a voting game in which participants vote for a single candidate out of a given set.

We implemented voting games that replicate two common real-world voting scenarios: In the first, a single voter votes once after seeing a large pre-election poll. In the second game, several voters play simultaneously, and change their vote as the game progresses, as in small committees. The winning candidate in each game (and hence the subject's payment) is determined using the plurality rule. For each of these settings we generated hundreds of game instances, varying conditions such as the number of voters and their preferences.

We show that people can be classified into at least three groups, two of which are not engaged in any strategic behavior. The third and largest group tends to select the natural "default" action when there is no clear strategic alternative. When an active strategic decision can be made that improves their immediate payoff, people usually choose that strategic alternative. Our study has insight for multi-agent system designers in uncovering patterns that provide reasonable predictions of voters' behaviors, which may facilitate the design of agents that support people or act autonomously in voting systems.

Categories and Subject Descriptors

I.2.11 [**Distributed Artificial Intelligence**]: Multiagent systems; J.4 [**Social and Behavioral Sciences**]: Economics

General Terms

Experimentation; Social Choice

Keywords

voting; behavioral studies

Appears in: *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015), Bordini, Elkind, Weiss, Yolum (eds.), May 4–8, 2015, Istanbul, Turkey.*

1. INTRODUCTION

Voting and preference aggregation systems have been used by people for centuries as tools for group decision-making in settings as diverse as politics and entertainment [30, 6, 24]. Computers have assumed an increasingly significant role as platforms and mediators for preference aggregation, such as scheduling applications,[1] aggregating search results from the web [8] and collaborative filtering [25], and more recently as autonomous voters in multi-agent systems [2].

While most theorists, whether in political science, economics, game-theory or computational social choice agree that people do not vote truthfully, it is not clear what voting strategy they actually employ, and what type of environmental factors affect this strategy. Indeed, even under the simple Plurality rule there is no agreement on how voters should vote or would vote given their preferences, and different studies suggest different conclusions [27, 1, 12].

Part of the difficulty lies in the radically different contexts in which people vote, such as political voting, a hiring committee, or in an online survey on which movie to watch, as it is possible that people's voting behavior strongly depends on these contexts. Even if the setting is fixed, there are preciously few publicly available benchmarks that researchers can use to evaluate the assumptions and predictions of various theories from the social choice literature.

Understanding people's voting strategies requires knowledge of their underlying preferences as well as their voting behavior. In recent years we witnessed a huge leap forward in the availability and accessibility of *preference profiles*, largely due the PrefLib project [17]. This database contains over 3,000 datasets from a variety of sources and locations, and is freely available on the web.[2] However the typical dataset contains either reported preferences (e.g. over Sushi orders in PrefLib), or strategic votes (e.g. referee ratings in ice-skating championship in PrefLib), but not both. In contrast, combined datasets used in the social choice literature typically contain very few independent polls or elections (sometimes just one), where the number of responders in each poll ranges from dozens to thousand [3, 26, 33].

Our goal is to fill this gap by collecting and analyzing human strategic voting behavior in a variety of online settings in which voters interact over the internet. There are several benefits for such settings. First, they reflect the growing use of computerized systems in the aggregation of people's preferences and voting behavior. Second, it allows us to create a controlled environment that abstracts away (as much

[1] `www.doodle.com/`

[2] `www.preflib.org/`

as possible) the influence of context and the only factors affecting people's voting behavior are their preferences and the information that is available to them. There is no interference due to dependency relationships between candidates, sense of duty, coalition formation and other factors that are common e.g., in political voting. Third, it allows us to run experiments on a large scale using crowdsourcing and to control people's preferences as well as the information that is presented to them.

We based our empirical study on two interactive voting games that are devoid of context and are easy to explain to subjects. In both settings, voters are automatically assigned a preference order over a fixed set of candidates, which is private information (unknown to the other voters). The first setting consists of a one-shot voting scenario in which a single human voter faces a dictated preference order. We completely control the data available to the voter by providing her with a (non-binding) pre-election poll of others' votes, and record her voting behavior under conditions that vary the information in the poll. The second setting consists of a group of human participants in an iterative voting game. As in the previous scenario, the preference profile is dictated to the voters, but they are free to change their votes at will until they reach an agreed outcome (or a timeout). As in the poll game, we recorded the decision of each voter along with the information available to her at that point in time. Both games put voters under uncertainty, but the source of uncertainty varies: in the first game voters only have access to an inaccurate poll. In the second game a voter directly observes the *current* votes of her peers, but does not know how they will vote *eventually* at the final round (or when will the final round arrive).

Both of our games were constructed so that there is always a "default" action that requires no cognitive effort: voting for the most-preferred candidate in the single-shot game, or keep voting for the same candidate in the iterative setting. In addition, the voter may be presented a situation where there is a clear "strategic" decision that can be made: In the poll setting it is to compromise for a less-preferred but more popular candidate according to the poll. In the iterative setting, it is to change the vote to whatever candidate that would maximize the profit if this was the final round. In game-theoretic terms, this latter strategy is known as *myopic best-response* (MBR).

Main Findings

We conducted an extensive empirical study in which over 300 human subjects played over 2500 game instances in both game settings. We varied the number of voters, subjects' preferences over candidates and (in the one-shot case) the poll information that was made available to them prior to voting. We analyzed under what conditions subjects choose the strategic action (or an unexpected, "irrational," action) over the default action. Our main findings are as follows:

- In both settings we found large interpersonal differences, identifying three distinct groups. A small number of subjects seemed to be voting randomly, and often voted for their least-preferred candidate. The second group consistently voted for their most preferred candidate. The third and largest group demonstrated more complex behavior in choosing when to strategize, and we focused our analysis on this third group.

- Between 65%-70% of voters compromise by strategically voting for their second preferred candidate when their most preferred candidate is ranked last in the poll.

- When the most preferred candidate is ranked second, a significant fraction of the voters voted for the (less preferred) leader of the poll. This phenomena is referred to as "herding" [15].

- In the iterative settings, 90% of subjects in the third group simply keep their current vote when this is consistent with the MBR strategy. When the MBR heuristic suggest to change the vote, only about 35% of the voters keep their vote, and nearly all others follow the MBR.

Taken together, these results provide new and valuable insights about the situations in which people vote strategically. They show that in the vast majority of cases, people follow their default action when there is no clear strategic alternative. When an active strategic decision can be made that improves their immediate payoff, people act strategically in over half of the cases. Moreover, some small amount of people *always* choose the default, whereas most of the others are strategic in 60-70% of the cases.

We are creating a public library called "VoteLib," which will include all of the collected data, and will be made freely available to the research community. This will allow researchers to test their own theories and train their models on our data without incurring the overhead of collecting the data, and will advance research in MAS and computational social choice.

2. RELATED WORK

Our study relates to work in the computational social choice and behavioral economics literature focusing on strategic voting.

2.1 Theoretical work

The literature on strategic voting can be roughly divided into "game-theoretic" models and "decision-theoretic" models. The first class of theories derive voting equilibrium concepts that are based on utility and rationality. Predominant examples include the Myerson and Weber model [22] (a variation of Bayes-Nash equilibrium), trembling hand equilibrium [20], strong equilibrium [31], and subgame-perfect equilibrium [9, 7]. The second class focuses on the strategic decision that a single voter faces, and the heuristics she may apply. These heuristics may range from simple best-response and other myopic heuristics [5, 13] to regret minimization [11], complex decision diagrams [23, 10] and so on. Some of these models specifically consider voters that are faced with poll information rather than with the preferences of their peers [4, 28]. Recent work considers voters that are faced with both poll information and the votes of their neighbors in a social network [32].

The game- and decision-theoretic lines of research are not entirely disjoint: it was shown that under the Plurality rule, in an iterative setting where voters may change their vote one at a time, voters who follow the simple myopic best-response (MBR) heuristic are guaranteed to converge to a Nash equilibrium [19]. Consequently, other heuristics have been shown to converge, giving rise to new notions of equilibrium [29, 13, 18]. Our work complements these papers by studying the conditions under which human voters in an iterative setting tend to follow MBR or other heursitics.

2.2 Experimental work

There is little prior work analyzing people's behavior in iterative voting settings. Most works focus on settings in which the same game is played several times and utilities are assigned at each game after the winning candidate is determined. We mention some formidable examples. Forsythe et al. [12] and Bassi [1] showed that over time people learned to strategize in a way that was consistent with a single equilibrium, and that they manipulated their vote significantly more often for the plurality voting rule than for a Borda voting rule. Van der Straeten et al. [34] performed lab experiments comparing how voters with single-peaked preferences behave under different voting rules and distributions over preferences. They survey previous experiments supporting the view that voters are rational agents whose actions can be predicted by equilibrium concepts, and conclude that this is only true in very simple settings, but that voters rely on simple heuristics when strategizing requires complex computations. Our study represents an interesting middle-ground, where the strategic possibilities are simple but with unkown preferences.

Kearns et al. [14] conducted experiments in which subjects arranged in networks were financially motivated to reach global consensus to one of two opposing choices. They showed that there are some network topologies in which the minority preference consistently wins globally and correlated behavioral characteristics of subjects (e.g., stubbornness) with payoffs. Bitan et al. [2] studied people's behavior in repeated settings in which a vote consisted of a complete ranking over the candidates, people's preferences were known, and their votes were aggregated using the Kemeny-Young rule. They showed that people learned to strategize and deviate from truthful reporting over time, but were outperformed by computer agents using various best-response methods. Lastly, we mention the work by Mattei et al. [16] who evaluated the Plurality, Borda, k-Approval, and Repeated Alternative Vote rules on millions of elections collected from publicly available data of people's reported preferences (the Netflix prize) but did not study strategic behavior.

3. THE SETTING

In this section, we describe how we adapted a popular voting system from the computational social choice literature to be used in committees that include both humans and computer agents. We first make the following definitions:

We denote $[x] = \{1, 2, \ldots, x\}$. Let M be a set of m candidates and let N be a set of n voters. A *social choice correspondence* is a function $f : C^N \longrightarrow \{2^C \setminus \emptyset\}$ that returns the set of winning candidates given a voting profile. A *voting profile* consists of a vector $\mathsf{a} : N \to M$, where $a_i \in M$ is the vote of voter i. The score of a candidate $c \in M$ given the voting profile a is defined as $s_{\mathsf{a}}(c) = |\{i \in N : a_i = c\}|$. A score vector s_{a} given voting profile a contains the scores for all voters summarizes all the relevant information on the outcome. We use the Plurality rule to choose the winning candidates $W(\mathsf{a})$ with maximal score given the voting profile a, that is $W(\mathsf{a}) = \operatorname{argmax}_c s_{\mathsf{a}}(c)$.

Let P be the set of all strict total orders over M. The preference ordering of voter i is a strict total order $Q_i \in P$ over the candidates (which is known only to i). Let $Q_i(a) \in [m]$ be the rank of candidate $a \in M$.

Voter i prefers candidate a to b, denoted $a \succ_i b$, iff $Q_i(a) < Q_i(b)$. In this paper we focus on $m = 3$ candidates, therefore we refer to the most preferred, second, and least preferred candidates for i as q_i, q_i', and q_i'', respectively. That is, $Q_i = q_i \succ_i q_i' \succ_i q_i''$.

We say that voter i is voting *truthfully* in profile a if $a_i = q_i$; otherwise i is voting *strategically*. We say $W(\mathsf{q})$ is the set of *truthful winner(s)* given the voting profile q in which all voters are truthful.

To define agents' rewards we need to extend preferences over candidates to preferences over subsets. To this end we impose linear rewards that depend on a single constant r. The reward to to voter i when candidate c wins is defined as $r_i(c) = (m - Q_i(c)) \cdot r$. We extend this definition for a subset of candidates $C \subseteq M$ as the average reward obtained over all candidates C: $r_i(C) = \frac{1}{|C|}\sum_{c \in C} r_i(c)$. In game-theoretic terms, the utility for voter i in voting profile a is $u_i(\mathsf{a}) = r_i(W(\mathsf{a}))$.

To illustrate our setting we present the following example in which four voters vote over a set of three candidates: Red (r), Grey (g) and Blue (b). The preference profile of the four voters is as follows:

$$Q_1 = \mathsf{r} \succ \mathsf{g} \succ \mathsf{b};\ Q_2 = \mathsf{r} \succ \mathsf{b} \succ \mathsf{g};\\ Q_3 = \mathsf{g} \succ \mathsf{b} \succ \mathsf{r};\ Q_4 = \mathsf{b} \succ \mathsf{g} \succ \mathsf{r} \quad (1)$$

Suppose that each of the voters votes for its most preferred candidate and that the reward constant is $r = 10¢$. In this the winning candidate is $W(\mathsf{r},\mathsf{r},\mathsf{g},\mathsf{b}) = \{\mathsf{r}\}$, and thus $q_1(\mathsf{r}) = q_2(\mathsf{r}) = 1$, and $q_3(\mathsf{r}) = q_4(\mathsf{r}) = 3$. The rewards for all voters are $r_1 = r_2 = 20¢, r_3 = r_4 = 0¢$. Suppose voter 4 voted for g rather than b. In this case there are multiple winners: $W(\mathsf{r},\mathsf{r},\mathsf{g},\mathsf{g}) = \{\mathsf{r},\mathsf{g}\}$. Consequently, the rewards are $r_1 = \frac{1}{2}(20¢ + 10¢) = 15¢, r_2 = r_3 = \frac{1}{2}(0¢ + 20¢) = 10¢$, and $r_4 = \frac{1}{2}(0¢ + 10¢) = 5¢$.

4. METHODOLOGY

Our methodology consisted of behavioral experiments involving people playing two types of voting games that follow the setting described in the last section. To support both of these settings, we designed an online infrastructure for voting that allows us to configure the number of voters, candidates and voters' preferences over the candidates, and whether the voters are humans or computer agents.

Our study focused on two settings, each comprising a family of games. The first setting consisted of a one-shot voting scenario in which the participant observes a pre-election poll prior to casting her vote. The second study consisted of an iterative voting scenario in which voters can change their vote until convergence. In both settings, voters are automatically assigned a preferred ranking over the candidates, which is private information unknown to the other voters.

4.1 Voting with polls

The first voting game allows each participant to vote (only once) for one of the predefined candidates. A pre-election poll in the form of a voting profile a was presented to each subject prior to casting his or her vote.

Fig. 1a shows a snapshot of the first voting game, that

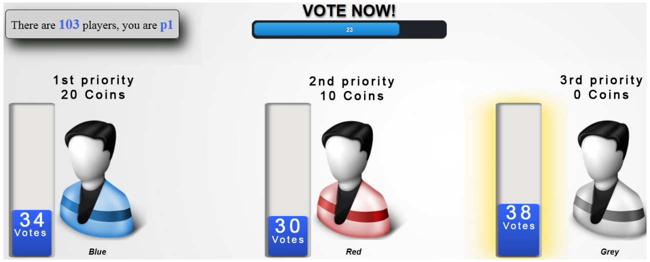

(a) poll voting setting

(b) Iterative voting setting

Figure 1: Voting game interface.

is configured to include three candidates (Red, Grey, and Blue) and 103 voters. The game interface is shown from the perspective of a human subject playing the game. The candidates are displayed in order of the preferences for the voter, from left (the most preferred candidate) to right (the least preferred candidate). The number of votes for each candidate in the voting poll is presented in the voting bar to the left of each candidate. For example, the red candidate has 30 votes. The winning candidate of the poll according to the Plurality rule (the Grey candidate in the figure, with 38 votes) is marked by a glowing voting bar.

Suppose that according to the poll scores s, we have $s(\mathsf{g}) > s(\mathsf{b}) > s(\mathsf{r})$ (as in Fig. 1a). For every instance of the game, a poll was generated using the following three parameters that were manipulated throughout the study:

- The total number of voters n, which ranged over 103, 1,009 and 10,007. This is in order to avoid round numbers in the poll, as such numbers may be treated differently by subjects [21]. Thus $s(\mathsf{g}) + s(\mathsf{b}) + s(\mathsf{r}) = n$.
- The gap between the number of votes for the leader and the runner-up, denoted "gap-leader" ($s(\mathsf{g}) - s(\mathsf{b})$). This gap was set to 3%, 5% or 7% of the value of $s(\mathsf{g})$.
- The gap between the runnerup and the least popular candidate in the poll, denoted "gap-last" ($s(\mathsf{b}) - s(\mathsf{r})$). This number set to one of the four categories: half of gap-leader, same as gap-leader, twice as much as gap-leader, and "large" (meaning $s(\mathsf{r}) = s(\mathsf{g})/2$).

In total, the parameter space defines 36 different possible poll configurations. Fig. 1a shows an example of a poll configuration in which the number of voters was set to 103, gap-leader was 7%, and gap-last was twice as high (i.e., 14%). The outcome of the vote was generated by sampling each voter i.i.d using the poll scores as the distribution. Thus the poll provided a noisy indication of the results. We emphasize several design choices. First, the subjects were not

Step t	1	2	3	4	5	6	7	8	9	10
Voter i	V_1	V_2	V_3	V_4	V_1	V_2	V_3	V_4	V_1	V_2
vote a_i	r	g	g	b	r	r	g	b	r	r
winner $W(\mathsf{a})$	r	r,g	g	g	g	r	r	r	r	r

Table 1: Example of iterative voting process, with convergence at step 10, after two and a half rounds.

informed on the accuracy of the poll (only that the poll was non-binding and that the poll results may not reflect the final score of each candidate), but could see the outcome after each game. Second, the actual probability that the participant would affect the outcome is quite small even for $n = 103$, since the voter is pivotal only in case of a tie.[3] For $n = 1009$ and higher, the probability is completely negligible. Therefore the performance of the subject (their average reward, which depended on the voting results and their preferences) was almost completely independent of their strategy, which is the common situation in wide-scale elections in the real world.

4.2 Iterative voting

In an iterative setting [19], voters start from some initial state a^0, but are then given repeated opportunities to change their vote. In the simplest form of iterative voting, a single voter may change her vote at each step according to some fixed order. The game ends either after a predetermined number of rounds, or if voters *converged* to an agreed outcome (see details below). It is important to note that voters' preferences do not change over the course of the game.

Formally, we denote the voting profile at step t by a^t, and the score vector and winner set derived from it by $\mathsf{s}^t = \mathsf{s}_{\mathsf{a}^t}, W^t = W(\mathsf{a}^t)$. Since only one voter may change her vote at each step, $\mathsf{a}^t, \mathsf{a}^{t+1}$ differ by at most one entry. A *round* is a sequence of n steps (one step for each voter). *Convergence* is defined as the case in which all voters do not change their votes in two consecutive rounds. Formally, if $\mathsf{a}^{t-t'} = \mathsf{a}^t$ for all $t' = 0, 1, \ldots, 2n-1$.

For example, Table 1 shows a history of votes for the above example for steps 1 through 10, in which convergence occurred. In this example, the game converged because the vote for each voter in steps 3 to 6 ($\mathsf{g}, \mathsf{b}, \mathsf{r}, \mathsf{r}$) repeated in steps 7 to 10.

The iterative voting experiments were performed on groups of several human voters, who are using iterative voting to select a winner(s) out of three possible candidates. We chose voter group sizes of 3, 5, and 7 voters, and designed 6 preference profiles according to the interplay between two selection criteria, the Plurality and Condorcet winners.[4] Specifically, The *NoCond* configuration class referred to preference profiles in which there was no Condorcet winner. The "CwP" configuration class referred to preference profiles with Condorcet winner who is also the winner when all voters vote truthfully and the winner is selected according to the Plurality voting rule (we call this the "truthful plurality rule").

[3] For example, if the poll scores are $s(\mathsf{g}) = 46; s(\mathsf{b}) = 31; s(\mathsf{r}) = 26$, then the probability of tie between g and b is $\sim 0.9\%$, and the probability of a tie between g and r is $\sim 0.15\%$.

[4] The Condorcet winner of an election is the candidate who, when compared with every other candidate, is preferred by more voters. It does not always exist.

CwP	V_1	V_2	V_3	V_4	V_5
q_i	r	r	r	g	g
q_i'	b	g	g	b	b
q_i''	g	b	b	r	r

(a) Condorcet winner agrees with Truthful Plurality winner

CnP	V_1	V_2	V_3	V_4	V_5
q_i	b	r	r	g	g
q_i'	r	b	b	b	b
q_i''	g	g	g	r	r

(b) Condorcet winner disagrees with Truthful Plurality winner

Figure 2: An example of preference profiles in iterative voting study

The "CnP" configuration class referred to preference profiles in which the Condorcet winner is not the winner under "truthful plurality." We created two profiles of each class, see a sample in Fig. 2.

Subjects each played up to 6 games, each time assigned to random group and a random preference profile. The game was played according to the protocol of iterative voting described in Section 3, starting from the truthful voting profile. Subjects could not see the actual preferences of the other voters, but could see the voting profile at each step (that is, which voter votes for which candidate). The game GUI is shown in Fig. 1b for an example configuration with 5 voters from the point of view of voter $p1$. The voting bar to the left of each candidate displayed the number of votes for the candidate at each round, as well as the identity of the voters who voted for the candidate. For example, the Red candidate is the current leader, with 3 of the votes, cast by voters $p4, p1$ and $p5$.

4.3 Data collection

The subjects for all experiments were recruited using Amazon mechanical Turk (all from the U.S.). Subjects were given a detailed tutorial of the voting game and their participation in the study was contingent on passing a comprehension quiz about the game.[5]

In the poll game subjects played up to 20 instances in sequence, with the same n parameter. Each game sampled a random poll configuration from the 12 configurations of gap-leader and gap-last. In addition, the subject was assigned a random preference order over the candidates. After each game we showed the subject the true outcome of the election and the winning candidate. The subject could choose to play a new game or to stop and collect her earnings on the games played.

In the iterative voting game subjects could play up to 6 games in a sequence, each time with a different preference profile and with a different group of subjects (matched at random). The games terminated when the voters converged, as described in Section 4.2, or if the number of rounds reached a predetermined threshold unknown to the participants (uniformly distributed between 5 and 10). The average session time per subject (excluding tutorial) for either condition was about 2-3 minutes.

All subjects received a show-up fee of $0.4 and a bonus that depended on their total rewards in the game. The reward (utility) of each candidate for a voter in a given game was set based on her preferences, as explained in Section 3. The reward constant was $r = 10¢$ for the iterative games (the maximum bonus was $1.2), and $r = 5¢$ for the one-shot games (the maximum bonus was $2).

[5]The tutorial can be found at `http://goo.gl/6rJJ4i`

Num. of voters	Num. games played	Num. of players
103	722	41
1,009	966	51
10,007	691	42
Total	2379	134

Table 2: Poll-voting Game Statistics.

5. RESULTS AND ANALYSIS

We will describe our empirical findings at several levels. For the poll-voting game we will analyze individual behavior, while for the iterative voting game we will analyze individual behavior as well as group behavior. Unless noted otherwise, all results are statistically significant within the $p < 0.05$ range.

5.1 Individual behavior under polls

We begin by analyzing the class of poll-voting games with a pre-election poll. Each game in this class included a single human subject. The other votes were determined by perturbing the poll according to a stochastic model. Specifically, votes where sampled from a multinomial distribution whose parameters are the poll scores s. This means that each of the voters (except the subject) chooses each candidate c with a probability $s(c)/n$.

The final result of the election was a single winning candidate that was determined according to the Plurality rule. The reward to the subject was determined solely by the winning candidates.

We collected at least 10 game instances for any combination of poll parameters and voter's preferences. Table 2 summarizes the number of games and participants for each value of n. We hypothesized the following:

1. People never vote for the least-preferred option q_i''.
2. People vote truthfully when their most-preferred candidate q_i is ranked 1st or 2nd in the poll.
3. When q_i candidate is ranked last, some people will remain truthful and some will compromise q_i for q_i'.
4. People compromise more when q_i is trailing further behind in the poll.
5. People compromise more when their two leading candidates q_i and q_i' are close in the poll.
6. The number of voters n has no effect on behavior.

These hypotheses are consistent with most theoretical models of strategic voting [22; 29; 18, cf.]. We report our findings for the $n = 1009$ condition, highlighting the differences from the other conditions when they exist.

Voters' types

Interestingly, in contrast to our first hypothesis, about 5% of votes were cast for the least-preferred candidate q_i'', even though this action can never be monetary beneficial to the

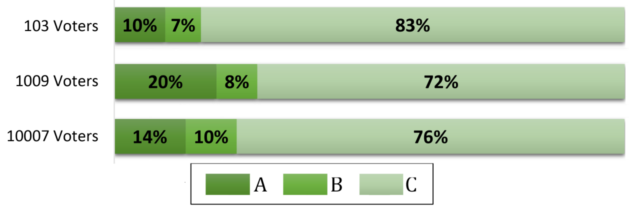

(a) poll voting

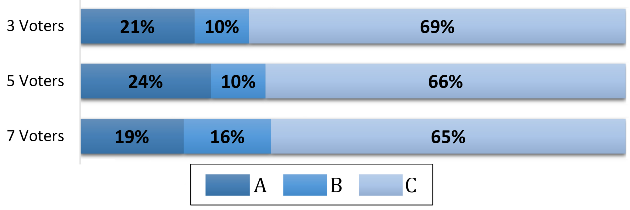

(b) iterative voting

Figure 3: Distribution over subject types.

subject (it is a *dominated action* in game-theoretic terms). This prompted us to check whether this irrational behavior was only expressed by a small number of participants. It turned out that nearly all of these irrational moves were due to a small group of voters, whose voting pattern was close to random. In particular, in the $n = 1009$ condition, 6 of these "random" voters account for over half of the total q_i'' votes. We classified those subjects who voted for q_i'' at least 20% of the time as "group A."

We also identified a second type of voters ("group B"), which consistently voted for their most-preferred candidate. Fig. 4 shows the distribution over the most-preferred candidate votes in the poll-voting games. As can be seen, there was a distinct group of subjects who consistently voted for their most-preferred candidate q_i. The rest of the voters, which formed the vast majority, were identified as Group C. Fig. 3a shows the distribution over group types among subjects, showing that the distinction between groups A, B and C was consistent for all n. We further checked if the subjects changed their behavior by comparing the statistics of the first 10 games to those of the last 10 games. We did not observe any statistically significant difference, indicating that at least at the aggregate level behavior was the same. In an exit survey, we found that the verbal descriptions of those individuals who described their strategy generally matched this group distinction.

Since group B behaved in a completely predictable way, and group A in a completely unpredictable way (they seemed to be clicking almost at random), we focused our analysis on group C. Thus the remaining results in this section refer to group C only.

Most-Preferred candidate leads

Fig. 5 shows the voting ratio for q_i, q_i', and q_i'' for the cases in which the q_i coincided with the poll leader, or coincided with the runner-up of the poll. As shown by the table, subjects overwhelmingly voted for their most-preferred candidate when it coincided with the leader of the poll (92%) and most subjects also voted for the most-preferred candidate when it was the runner-up in the poll (70%). This confirms the second hypothesis, but the asymmetry indicates that subjects are also inclined to vote for the leader of the poll.

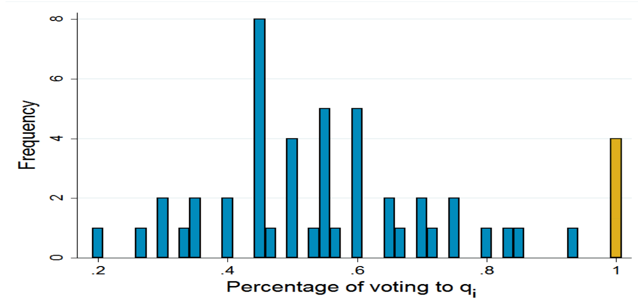

Figure 4: Histogram of votes for q_i by subject (poll voting). We can see that most voters are distributed around 0.5, whereas the voters of group B (the rightmost bar) are a clear outlier from the distribution.

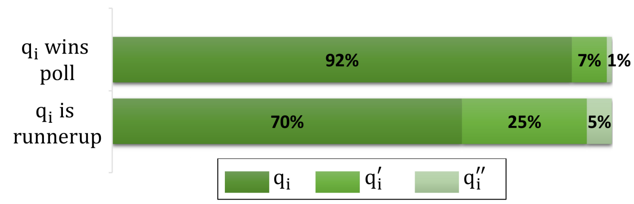

Figure 5: Voting behavior for 1,009 voters.

Most-Preferred candidate loses

To address the third hypothesis, we present Fig. 6 which provides an analysis of voting behavior when q_i is ranked last in the poll. As we can see, most of the voters decided to vote for q_i'. This decision can be rationalized by either of the two following arguments: it reduces the chance of the worst outcome (selection of q_i''); and it has the highest expected utility (payment). However recall that the difference from voting for other candidates is completely negligible due to the high number of voters. We interpret this decision to prefer q_i' over q_i as a "compromise." As shown in the figure, compromises are slightly more common (73%) when q_i' is leading the poll (bottom bar), than when it is the runner-up (64% top bar). This is a further, weaker, confirmation that voters have a bias towards voting for the leader.

To test hypotheses 4 and 5 we controlled the conditions of gap-leader and gap-last. We observed some correlation between gap-last and compromise ratio (supporting hypothesis 3), and some correlation between gap-leader and compromise ratio (contradicting hypothesis 4). However neither of these correlations was statistically significant, and thus the effect of both parameters remains inconclusive. It is clear however that this effect is *much weaker* than the overwhelming effect of the ordinal ranking of candidates in the poll.

As for the sixth hypothesis, we observed very similar qualitative and quantitative patterns in the other two conditions (n=103, n=10007). The only difference was regarding hypotheses 3 and 4: in both conditions there was negligible effect of gap-leader. The effect of gap-last was stronger (statistically significant) for n=103, but practically disappeared for n=10007. We conjecture on possible interpretations of this preliminary finding in the discussion.

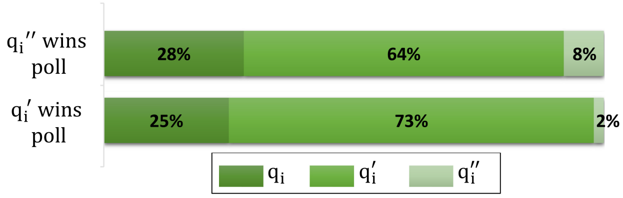

Figure 6: Voting behavior for 1,009 voters when most-preferred candidate q_i is ranked last in the poll.

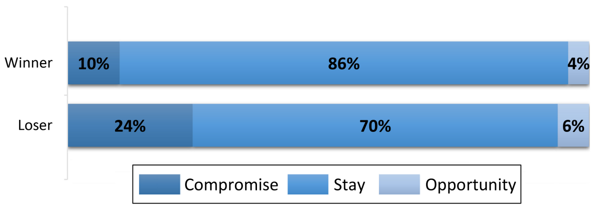

Figure 7: Moves by initial position.

5.2 Individual behavior in a group

In this section we analyze the iterative voting games across all conditions. We report our findings for groups of 5 voters. Strikingly, our findings for groups of 7 and 3 voters were similar and exhibited the same patterns. Due to brevity concerns we do not show them in the paper.

Following [18], we denote a *compromise move* as a change in vote to a less-preferred candidate, and an *opportunity move* as a change in vote to a more-preferred candidate. We denote a *stay move* as no change in voting compared to the previous round. Out of 940 moves performed by all of the subjects in all games, 117 (12%) were compromise moves, 53 (6%) were opportunity moves, and the rest were stay moves. Intuitively, we ask under what conditions a voter performs such moves.

Interpersonal differences and voters' types

As with the poll experiments, we identified a small number of participants who voted randomly, by singling out voters who selected q_i'' at least 20% of the times (group A). There was also a significant portion of the voters who consistently voted for their most-preferred candidate (group B) in all of their votes. As before, all other subjects where grouped together (group C). Fig. 3b shows the distribution over types for the iterative voting games.

Within group C there were still large interpersonal differences in terms of the behavior, as we later discuss. Our remaining results in this subsection are stated for group C.

Initial position vs. current position

Our first attempt was to look at the initial position of the voter in the game. Roughly speaking, a voter can be either a "winner," meaning that her most preferred candidate q_i is the truthful Plurality winner (and hence the winner of round 0), or else, the voter is a "loser."

Fig. 7 lists the ratio of compromise and opportunity moves for winners and losers. As can be seen by the figure, losers engage in significantly more compromise moves than do winners. At a first glance, this seem to indicate that the behavior in the game is largely affected by the initial position.

In order to get a more accurate picture, we considered the position of the voter just before taking the step. Rather than looking at winners and losers, we tested whether the voter is *pivotal*. That is, whether there is a move that can improve the voter's reward *myopically* (assuming that the game ends after the next step). Such a step is called *myopic best response* (MBR) [19]. Fig. 8 classifies all moves according to whether an MBR exists, and whether the MBR itself is a compromise or an opportunity move.

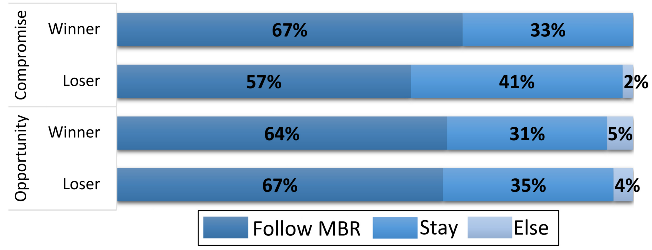

Figure 8: MBR Analysis. We analyze a pivotal voter's action in four conditions: when the MBR heuristic was a compromise/opportunity move (current position); and when the voter is a winner/loser (initial position).

Num Voters	Converged	Nash equilibrium	Avg. Rounds
3-Voters	0.84	all	3.49
5-Voters	0.89	0.98	3.40
7-voters	0.75	0.97	4.16

Table 3: Iterative game Statistics.

From the figure we can conclude the following observations: (a) when the MBR indicates the voter is pivotal, voters perform the MBR move in about 60% of the cases, across all 4 conditions (regardless of whether the MBR is a compromise or an opportunity move) (b) there is no significant difference in behavior between "winners" and "losers" once we control for the current position.

As with the rate of strategic response in the poll games, there where large interpersonal differences with some voters having much greater tendency to follow the MBR heuristic.

Therefore it seems that a significant part of voters' behavior is guided by myopic reasoning, that is, by following the best response. To complete this picture, we analyzed the behavior when the voter was not pivotal, and when the MBR was to keep the same vote. Indeed, the probability of keeping the same vote in either condition was over 90% (compared to 35% when the MBR move was to change a candidate, see Fig. 8), where most of the voters who decided not to follow the MBR performed a compromise move.

5.3 Group behavior

For all group sizes, we find that most games converge (and fast), see Table 3. When a game converges it is almost always to a Nash equilibrium: a state where no voter has an MBR move. There was no statistically significant difference between the number of rounds played in the 3- 5- and 7-voter groups until convergence.

What behavior is better for the individual?

We considered the average reward for voters, partitioning once based on initial roles (winner/loser), and once based on subject type (Group A/B/C). The average reward for winners under the CwP configuration was about 5 times higher than that of plurality losers, presumably since the outcome was often the truthful plurality outcome. There was no significant difference in the score for the other configurations.

The "random" group A did most poorly w.r.t. the rewards, but its small size did not allow us to determine if this effect is significant. Interestingly, groups B and C seemed to be doing equally well on average, indicating that there is no clear advantage to follow the MBR heuristic or to remain truthful at all cost.

What behavior is better for the society?

We did find a statistically significant difference between the average reward for all voters and the hypothetical reward they would get for just voting truthfully.

There was a more pronounced effect when we looked at the identity of the winner. In roughly half the games overall the winner did not change (the truthful plurality winner was selected). However under the CnP profiles the Condorcet winner was selected more often than the plurality winner, indicating a compromise at the group level.

6. DISCUSSION AND FUTURE WORK

Our results demonstrate that although there is no single function that maps the information and preferences of a voter to a perfect prediction of her action, there are still simple heuristics or patterns that provide reasonable predictions of voters' behaviors, even in distinct settings such as one-shot and iterative games.

First, different people behave differently when it comes to strategic voting, where some express a strong bias towards their most preferred candidate, which can be thought of as the "default" action. Our partition to types was surprisingly robust over all six experiment conditions (Figure 3).

Second, the people who *are* strategic, seem to base their decision almost entirely on the availability of an obvious strategic move (compromise in polls, MBR in the iterative setting). The specific details of the situation, such as candidate's scores in the poll, are largely ignored. This observation should facilitate the generation of hypotheses regarding voters' behavior in more complex experiments, e.g. with a larger number of candidates.

Lastly, at least in polls, people demonstrate moderate but clear bias towards voting for the leader of the poll (Figures 5,6). This phenomenon, known as *herding* has also been observed in other voting scenarios [15].

This third phenomenon is particularly surprising, since a common explanation for herding is that the voter is unsure about the quality of the candidate, whereas in our setting the rewards were known and fixed. In fact, given the noise model we used to generate the outcome, voting for the second-preferred candidate in the instances used for Fig. 5 invariably decreases the expected utility, and hence the "herding" moves cannot be rationalized by purely economic terms. A possible explanation is that voting for the leader is perceived by some voters as an alternative "default" option, that does not require cognitive effort (just like voting for the most preferred).

We can conclude that our results generally support the "decision-theoretic" models of strategic voting. Indeed, it seems that for the large part, human voters follow relatively simple heuristics, that ignore and sometimes directly contradict economic, or "game-theoretic" reasoning. However, when looking for a theory to explain and predict voters' behavior, it is crucial that the model will allow for a wide range of behaviors, as specified above.

In the future we intend to perform a deeper analysis of interpersonal differences, whether by identifying finer subgroups of voters, or some "personal parameters" that affect voters' behavior (such as different levels of risk-aversion or tendency for herding).

Sophisticated voting

We observed two preliminary patterns that may indicate more sophisticated voting behavior then what described thus far. In the poll games, it seems that a small "gap-last" plays some role in the decision of the voter to remain truthful, especially when n is low. This may be due to the voter's perception that in these cases "everything can happen," and by deserting the most-preferred candidate q_i there is some real risk of failing to support q_i at a critical tie.

In the iterative games, while less than 10% of the "MBR stay" moves were violated by voters, these violations still account for almost a third of the total compromise moves in the game. Thus we cannot just ignore them as noise, and there may be more sophisticated reasoning behind it.

Further research is required to better understand these and possibly other sophisticated patterns of strategic voting. In addition, we plan to analyze how individual behavior changes over time.

Extensions

Since in our experiment there were only three candidates, the range of available strategic decisions was very limited. Running experiments with larger sets of candidates will enable us to study what strategic actions voters prefer when there are several plausible alternatives.

We would also like to check if full or partial knowledge of the other players' *preferences* affects the voter's strategy.

Finally, a better understanding of how people behave strategically in online voting settings can guide the design and implementation of better platforms for preference aggregation. Our experimental infrastructure can be used to test such mechanisms in a context-free environment.

7. ACKNOWLEDGEMENTS

This work was supported in part by EU grant no. FP7-ICT-2011-9 #600854 and by Israeli Science Foundation (ISF) grant no. 1276/12.

REFERENCES

[1] A. Bassi. Voting systems and strategic manipulation: an experimental study. Technical report, mimeo, 2008.

[2] M. Bitan, Y. Gal, S. Kraus, E. Dokow, and A. Azaria. Social rankings in human-computer committees. In *AAAI*, 2013.

[3] J. R. Chamberlin, J. L. Cohen, and C. H. Coombs. Social choice observed: Five presidential elections of the american psychological association. *The Journal of Politics*, 46(02):479–502, 1984.

[4] S. Chopra, E. Pacuit, and R. Parikh. Knowledge-theoretic properties of strategic voting. Presented in JELIA-04, Lisbon, Portugal, 2004.

[5] V. Conitzer, T. Walsh, and L. Xia. Dominating manipulations in voting with partial information. In *AAAI'11*, pages 638–643, 2011.

[6] G. Cox. *Making votes count: strategic coordination in the world's electoral systems*, volume 7. Cambridge Univ Press, 1997.

[7] Y. Desmedt and E. Elkind. Equilibria of plurality voting with abstentions. In *EC'10*, pages 347–356, 2010.

[8] C. Dwork, R. Kumar, M. Naor, and D. Sivakumar. Rank aggregation methods for the web. In *WWW'01*, pages 613–622. ACM, 2001.

[9] R. Farquharson. *Theory of Voting*. Yale Uni. Press, 1969.

[10] D. S. Felsenthal, A. Rapoport, and Z. Maoz. Tacit co-operation in three-alternative non-cooperative voting games: a new model of sophisticated behaviour under the plurality procedure. *Electoral Studies*, 7(2):143–161, 1988.

[11] J. A. Ferejohn and M. P. Fiorina. The paradox of not voting: A decision theoretic analysis. *The American political science review*, pages 525–536, 1974.

[12] R. Forsythe, T. Rietz, R. Myerson, and R. Weber. An experimental study of voting rules and polls in three candidate elections. *International Journal of Game Theory*, 25(3):355–383, 1996.

[13] U. Grandi, A. Loreggia, F. Rossi, K. B. Venable, and T. Walsh. Restricted manipulation in iterative voting: Condorcet efficiency and borda score. In *ADT'13*, pages 181–192. Springer, 2013.

[14] M. Kearns, S. Judd, J. Tan, and J. Wortman. Behavioral experiments on biased voting in networks. *Proceedings of the National Academy of Sciences*, 106(5):1347–1352, 2009.

[15] J. Leskovec, D. P. Huttenlocher, and J. M. Kleinberg. Governance in social media: A case study of the wikipedia promotion process. In *ICWSM*, 2010.

[16] N. Mattei, J. Forshee, and J. Goldsmith. An empirical study of voting ruk and manipulation with large datasets. *Proc. of COMSOC-12*, 2012.

[17] N. Mattei and T. Walsh. Preflib: A library for preferences http://www.preflib.org. In *ADT'13*, pages 259–270. 2013.

[18] R. Meir, O. Lev, and J. S. Rosenschein. A local-dominance theory of voting equilibria. In *EC'14*, pages 313–330. ACM, 2014.

[19] R. Meir, M. Polukarov, J. Rosenschein, and N. Jennings. Convergence to equilibria in plurality voting. In *AAAI'10*, 2010.

[20] M. Messner and M. K. Polborn. Robust political equilibria under plurality and runoff rule, 2002. Mimeo, Bocconi University.

[21] J. Mitchell. Clustering and psychological barriers: The importance of numbers. *Journal of Futures Markets*, 21(5):395–428, 2001.

[22] R. B. Myerson and R. J. Weber. A theory of voting equilibria. *The American Political Science Review*, 87(1):102–114, 1993.

[23] R. G. Niemi and A. Q. Frank. Sophisticated voting under the plurality procedure. In *Political Equilibrium*, pages 151–172. Springer, 1982.

[24] T. Palfrey. Laboratory experiments in political economy. *Annual Review of Political Science*, 12:379–388, 2009.

[25] D. M. Pennock, E. Horvitz, C. L. Giles, et al. Social choice theory and recommender systems: Analysis of the axiomatic foundations of collaborative filtering. In *AAAI/IAAI*, pages 729–734, 2000.

[26] M. Regenwetter and B. Grofman. Approval voting, borda winners, and condorcet winners: Evidence from seven elections. *Management Science*, 44(4):520–533, 1998.

[27] M. Regenwetter and E. Rykhlevskaia. A general concept of scoring rules: general definitions, statistical inference, and empirical illustrations. *Social Choice and Welfare*, 29(2):211–228, 2007.

[28] A. Reijngoud and U. Endriss. Voter response to iterated poll information. In *AAMAS'12*, pages 635–644, 2012.

[29] R. Reyhani, M. C. Wilson, and J. Khazaei. Coordination via polling in plurality voting games under inertia. In *COMSOC'12*, 2012.

[30] W. Riker and P. Ordeshook. A theory of the calculus of voting. *The American Political Science Review*, 62(1):25–42, 1968.

[31] M. R. Sertel and M. R. Sanver. Strong equilibrium outcomes of voting games are the generalized condorcet winners. *Social Choice and Welfare*, 22(2):331–347, 2004.

[32] S. Sina, N. Hazon, A. Hassidim, and S. Kraus. Adapting the social network to affect elections. In *Proc. Of AAMAS15*, 2015.

[33] A. M. van Deemen and N. P. Vergunst. Empirical evidence of paradoxes of voting in Dutch elections. *Public Choice*, pages 475–490, 1998.

[34] K. Van der Straeten, J.-F. Laslier, N. Sauger, and A. Blais. Strategic, sincere, and heuristic voting under four election rules: an experimental study. *Social Choice and Welfare*, 35(3):435–472, 2010.