

Social Rankings in Human-Computer Committees

Moshe Bitan¹, Ya'akov (Kobi) Gal³ and Elad Dokow⁴, and Sarit Kraus^{1,2}

¹ Computer Science Department, Bar Ilan University, Israel

² Institute for Advanced Computer Studies University of Maryland, College Park
MD 20742

³ Department of Information Systems Engineering, Ben-Gurion University of the
Negev, Israel

⁴ Department of Economics, Bar Ilan University, Israel

Abstract. This paper provides a study of human and computational strategies in voting systems. Despite committees and elections being widespread in our daily lives, the design of agents that can operate in such settings has received far less attention than the theoretical analysis of voting strategies in such settings. We address this gap by comparing people's behavior in voting systems with that of computer agents playing various strategies. In our setting participants vote by simultaneously submitting a ranking over the set of candidates and the election system uses a plurality rule to select a ranking that minimizes disagreements with participants' votes. We ran an extensive study in which hundreds of people participated in repeated voting rounds with other people as well as computer agents that differed in the extent to which they employ strategic reasoning in their behavior. Our results show that over time, people learned to deviate from truthful voting strategies, and use more sophisticated voting strategies. However, these strategies do not improve people's performance because of the erratic nature of their behavior. In particular, a computer agent using a best response voting strategy to people's actions in the previous round was able to outperform people in the game as well as an agent using truthful voting strategies. This result has implication for agent designers, highlighting the types of strategies that facilitate voting behavior in committees comprising both human and computer participants.

1 Introduction

Voting systems have been used by people for centuries as tools for group decision making, in settings as diverse as politics [1–3], and entertainment [4]. As computers become ubiquitous in people's lives, heterogeneous group activities of computer systems and people are becoming more prevalent. As a result, opportunities arise for computer agents to participate in voting systems, whether as autonomous agents or proxies for individual people.

Past work of human-computer decision-making has focused on the design of computer agents for interacting with people in negotiation and coordination

settings [5]. However, there is no work on modeling people’s voting behavior in heterogeneous systems comprising human and computer participants. This paper addresses this gap by comparing people’s voting behavior to that of computer agents using classic voting strategies from the literature. In our voting system all participants declare their true preferred ranking over a set of candidates prior to commencing a series of voting rounds. At each round all participants vote by simultaneously submitting a ranking over the set of candidates. The election system uses a plurality rule to choose a ranking that minimizes the sum of conflicts with the votes that were submitted by the participants. The utility of participants is proportional to the extent to which the chosen ranking agrees with their preferences. Such settings are analogous to real-world voting scenarios such as rating grant proposals and ranking applicants for positions in academia, industry or competitions.

We designed a three-player game that implemented the voting system described above using a budget allocation analogy. The preferences of participants over the various sectors was chosen such that players could potentially improve their score in the game if they deviated from their truthful vote. We formalized several voting strategies for this game that differ in the extent to which they reason strategically about other’s voting behavior. We conducted an extensive empirical study in which hundreds of human subjects played this game repeatedly with other people as well as computer agents that varied in the extent to which they voted strategically. We hypothesized that over time, people would vote less truthfully, and that computer agents using various levels of strategic voting would be able to outperform people.

Our results show that people deviate more from their truthful voting strategies in later rounds than in earlier rounds, but that this deviation does not necessarily result in an improvement in performance. In some cases, voting truthfully can outperform people. In addition, we identified several heuristics people use to guide their play. A computer agent using a best-response strategy to people’s voting actions in the previous round was able to outperform people. This work has significance for agent designers in demonstrating that best-response strategies are sufficient for agents to outperform other people in voting systems that output a complete ranking. This is the first work to compare the performance of computational voting strategies with people’s voting behavior.

2 Related Work

Voting systems have been studied extensively in computer science and economics (see for example [6] and [7]). The most widely used voting rule is the plurality rule, in which each voter has one vote and the winner is determined as the candidate that receives the highest number of votes. Other popular voting rules, such as the Borda rule, allow voters to order the candidates, and the winner is determined by the candidate that receives the most points (relative to its positions in all of the voters’ rankings). Both of these voting rules are susceptible to manipulation, that is, self-interested players have an incentive to vote strate-

gically against their true preferences in certain situations. Thus, many studies in behavioral economics which have examined the effect of these voting rules on people’s voting strategies. Specifically, Forsythe et al. [8] studied the effect of different voting rules on people’s voting strategies in three-candidate elections in which a single candidate is elected and there was full information about voters’ preferences. They showed that people generally diverge from truthful voting, and that over time, they learn to cast votes that are consistent with a single equilibrium. In a follow-up study, Bassi [9] showed that people invoked different voting strategies depending on the voting rule implemented by the system. In particular, incorporating a simple plurality voting rule led people to adopt more strategic voting than when incorporating the Borda rule which was based on ranking the candidates.

Our research extends these studies in two ways. First, we consider more complex settings in which the voting system outputs a ranking over the candidates, rather than a single winning candidate. Such settings occur frequently in the real world, but people’s behavior in these voting systems has not been studied. We hypothesized that people’s behavior will significantly diverge from equilibrium, and in order to succeed, computer agents will need to adopt other types of voting strategies. Second, we provide a first study that compares the performance of computational strategies with people’s voting behavior.

3 The Setting

In this section we provide a formal description of our voting system. Let C denote a set of candidates. For example, consider a committee that needs to prioritize categories for the next budget. The candidate categories in C are: Defense (d), education (e), health (h) and transportation (t). For any candidate pair $a, b \in C$, we use notation $a \succ b$ to mean that a is preferred over b . A *ranking* is a relation that defines a total order over C . One such ranking in our committee example is $d \succ e \succ h \succ t$. Given a set of agents N , we denote the preferred ranking of agent i over C as F_i . For example, suppose that there are three agents in the committee. The preference F_1 of agent 1 is $e \succ d \succ h \succ t$; the preference F_2 of agent 2 is $e \succ t \succ d \succ h$; and the preference F_3 of agent 3 is $h \succ t \succ d \succ e$.

Agents vote by submitting a rank over the sets of candidates C . Thus a vote can be any relation in $C \times C$. For example, the vote of agent i may be $e \succ d \succ h \succ t$. Note that the vote of an agent does not need to match its preferred ranking. Let V denote a set of votes $\{V_1, \dots, V_n\}$ for all agents. Now how do we combine the different votes of the agents in the committee to form a “social result”? A natural method is to choose the ranking that agrees with the majority of agents’ votes. Formally, we say that candidates a, b in vote V_i are *consistent* with rank R if $a \succ b$ in V_i and $a \succ b$ in R . Let $I(V_i, R \mid a, b)$ be an indicator function that equals 1 if a, b in vote V_i are consistent with R . We make the following definition.

Definition 1 Given agents' votes $V = \{V_1, \dots, V_n\}$, we say that relation R is a Pairwise Plurality relation over C if for any $a, b \in C$, it is the case that $a \succ b$ in V_i if and only if $\sum_{V_i \in V} I(V_i, R | a, b) \geq \frac{n}{2}$.

For example, suppose that V_1 is $e \succ d \succ h \succ t$; V_2 is $e \succ t \succ d \succ h$; V_3 is $h \succ t \succ d \succ e$. The chosen relation R will satisfy the following condition:

$$R = \{e \succ d, d \succ h, e \succ h, e \succ t, t \succ d, h \succ t\} \quad (1)$$

To see this, consider that for both votes V_1 and V_2 , it holds that

$$e \succ d, d \succ h, e \succ h, e \succ t$$

This is because both agents 1 and 2 prefer education to defense, defense to health, education to health, and education to transportation. Similarly, for both V_2 and V_3 , it holds that $t \succ d$, and for both V_1 and V_3 it holds that $h \succ t$.

We have shown how to combine the votes of the committee members into a single relation R . For R to form a valid ranking, it has to define a total ordering over C . However, using pairwise plurality to produce a chosen ranking may not result in a total ordering. In fact, Arrow [10] has shown that there does not exist a reasonable voting rule that generates a ranking over candidates and is guaranteed to produce a unique order. To see this, consider that a relation defines a total order if it is transitive, anti-symmetric and total. In our example, it is easy to see that the R is anti-symmetric (it is not the case that $a \succ b$ and $b \succ a$) and total (all candidate pairs are either accepted or not accepted by R). However, R is not transitive. For example, we have that $d \succ h$ and $h \succ t$, but it is not the case that $d \succ t$.

To be able to generate a total order over C while using the pairwise plurality rule, we need to transform R to a transitive relation over the candidates C . To this end, we make the following definition: A *pair reversal* of $a \succ b$ in R modifies R such that $b \succ a$ holds, but not $a \succ b$. We make the following Lemma.

Lemma 1 Any pairwise plurality consistent relation R can be transformed to a total ordering over C by a process of pair reversals.

We can now make the following definition:

Definition 2 Let R be the pairwise plurality ranking over C . Let the chosen ranking $R^*(V)$ be the ranking that transforms R to a total ordering using the minimal number of pair reversals.

For example, by performing a single pair reversal in the relation from $d \succ h$ to $h \succ d$ to achieve the following total ranking of the candidates given participants' votes:

$$R^*(V) = \{e \succ h \succ t \succ d\} \quad (2)$$

This process aligns with the Kemeny Young method [11] commonly used in economics.

We now specify a scoring function which will be used to measure the distance between the chosen ranking and each of the individual preferences. An agent i receives one point for each candidate pair in its preferences that is consistent with the chosen ranking. Formally, the score of i , denoted sc_i is defined as follows:

$$sc(F_i, V) = \sum_{a,b \in C} I(F_i, R^*(V) | a, b) \quad (3)$$

Suppose that the preferred ranking F_i of agent i is $e \succ d \succ h \succ t$ and the chosen ranking R^* is $e \succ h \succ t \succ d$. There are four candidate pairs in V_i that are consistent with R^* , namely $e \succ d; e \succ h; e \succ t; h \succ t$. Thus, the score of agent i is 4 points.

3.1 Voting Strategies

Given that the chosen ranking combines agents' votes, and their scores depend on the extent to which the chosen ranking reflects their preference, how should agents vote in order to maximize their scores? The most obvious way to vote is according to each agent's preferences.

A vote V_i is said to be *truthful* if $V_i = F_i$. In this case the agent votes according to its preferences. In our example, if all members vote truthfully, then $V = (V_1, V_2, V_3)$ and the chosen ranking $R^*(V)$ will be $e \succ h \succ t \succ d$. In this case, agent 1 will incur a score of 4, agent 2 will incur a score of 4, and agent 3 will incur a score of 3. However, agents may be able to do better when they vote strategically. For example, if V_1 is $d \succ e \succ h \succ t$, and both agents 2 and 3 vote truthfully ($V_2 = F_2, V_3 = F_3$) then the chosen ranking $R^*(V)$ will be $d \succ e \succ h \succ t$ which will give agent 1 a score of 5. This score is higher than its score for voting truthfully given that agents 2 and 3 vote truthfully. The following definitions formalize this intuition by describing strategic voting patterns.

Let $p(V_i)$ denote the probability that agent i submits vote V_i . Let V_{-i} denote a voting profile of all agents other than i . Assuming agents' votes are independent, let $p(V) = p(V_1) \cdot p(V_2) \cdots p(V_n)$ denote the joint probability that agents submits the voting profile V . The *best response* vote of agent i , denoted BR_i , is defined as the vote that maximizes the score of i given a probability distribution $p(V_{-i})$ over the votes of the other agents. Formally, we write

$$BR_i \in \operatorname{argmax}_{V_i \subset C \times C} \sum_{V_{-i}} p(V_{-i}) \cdot sc(F_i, V) \quad (4)$$

where $V = (V_{-i}, V_i)$.

The *Level-0* vote of i , denoted $L0_i$, is defined to be the best response of i given that all other agents are believed to vote truthfully, that is, for any agent $j \neq i$, we have that $p(V_j) = 1$ if $V_j = F_j$; otherwise, $p(V_j) = 0$. For example, the level-0 vote of agent 1, given that it believes that agent 2 and 3 vote truthfully, is $d \succ e \succ h \succ t$, as stated above. The level-0 vote of agent 2 is $t \succ e \succ d \succ h$. The *Level-1* vote of i , denoted $L1_i$, is defined to be the best response of i given that

all other agents are believed to be level 0 voters, that is, $p(V_j) = 1$ if $V_j = L0_j$. For example, the level-1 vote for agent 3, given that it believes that agents 1 and 2 submit level-0 votes, is $h \succ d \succ t \succ e$.

A profile of rankings V is *Nash equilibrium* if the following holds for every agent i

$$\forall V'_i \subset C \times C \text{ it holds that } sc(F_i, V) \geq sc(F_i, (V_{-i}, V'_i)) \quad (5)$$

In our example, the profile by which agent 1 submits a truthful vote ($e \succ d \succ h \succ t$), agent 2 submits a level-0 vote ($t \succ e \succ d \succ h$), and agent 3 submits a truthful vote ($h \succ t \succ d \succ e$) is Nash equilibrium in which the chosen ranking $R^*(V)$ is $t \succ e \succ d \succ h$. This profile incurs a score of three points for agent 1, five points for agent 2 and two points for agent 3.

4 The Budget Allocation Game

To study people’s voting behavior we designed a budget allocation game in which n agents vote to allocate a budget among C categories. Each agent is assigned a ranking that represents its preference over the four categories and this information is common knowledge among all agents. The game comprises a finite number of rounds. In each round, all agents simultaneously submit a ranking $V = (V_1, \dots, V_n)$ over the categories. The chosen ranking $R^*(V)$ is computed using Definition 2, and each agent’s score is computed using Equation 3. Agents’ votes, the chosen ranking, and their scores are made visible to agents at the end of each round. Agents’ preferences remain constant across rounds.

We implemented a version of the budget allocation game in which there are three players and four categories: education, transportation, health and defense. A snapshot of the main game board is shown in Figure 1 from the point of view of Player 1. The board shows the preferences of the three players in the game, as well as an editable ranking that player can modify and submit as its vote. The bottom panel of the Figure shows the result of one of the rounds in the game, specifying the votes for all players.

4.1 Rules of the Game

The budget allocation game is played repeatedly for five rounds. Participants simultaneously submit their preferences at each round. The default vote for each participant in the first round is simply their true vote. The default vote for each consecutive rounds is the ranking that the player submitted in the previous round. The default vote is selected for a round if the participant has not submitted a ranking after a three minute deadline. After all participants have submitted their rankings, the chosen ranking and scores are computed as explained above, and displayed to all of the participants. The assigned preferences of the participants are constant and do not change from round to round. The bottom panel of Figure 1 shows the resulting ranking when all participants vote according to their true preferences. As shown by the Figure, the resulting ranking $R^*(V)$ is $e \succ h \succ t \succ d$. Lastly, to help people reason about their decisions in the game,



Fig. 1. Snapshots of the Budget Allocation Game: The main voting panel (top); announcement of participants' vote, the chosen ranking, and obtained score (bottom)

we designed a decision support tool that allows people to query the scores for different voting strategies for themselves and other players in the game.

There are several advantages to using this game to study human and computational voting strategies. First, it includes the minimal number of candidates such that players have incentive to vote strategically. (In fact, it can be shown that voting truthfully is the optimal strategy for voters in the case that there are three candidates.) Second, it provides an analogy to voting scenarios in the real world such as ranking applicants for positions in academia or industry and deciding on the allocation of resources in political committees. Third, the fact that players vote repeatedly allows them to adapt their voting behavior over time, and reflects settings such as annual budget decisions and recurring elections.

4.2 Preference Profiles

As described above, players' scores for each round of voting depend on the extent to which the chosen ranking agrees with their preferred ranking that is assigned to them at the onset of the interactions. In real world voting scenarios, some participants may be in better positions than others to affect the voting outcome. In the budget allocation game, we can define different power conditions between committee participants by varying their assigned preference profile. We used two preference profiles in the study that differed in the extent to which they allowed players to affect the voting result by deviating from their truthful vote.

In the first profile, called “symmetric”, the preferences of player 1 were $e \succ d \succ h \succ t$; the preferences of player 2 were $e \succ t \succ d \succ h$; the preferences of player 3 were $h \succ t \succ d \succ e$. These preferences are shown in the the main game board in Figure 1. This profile provides a symmetric outcome for players 1 and 2. If all players vote truthfully (we call this the “naive” voting baseline), player 3 is at a disadvantage, because the chosen ranking will be $e \succ h \succ t \succ d$, incurring a score of 4, 4, and 3 for players 1, 2 and 3, respectively. Moreover, the naive voting baseline is not stable, in the sense that at least one player can improve its score by voting strategically. Specifically, player 1 can improve its score by voting its level-0 strategy of $d \succ e \succ h \succ t$, given that other players vote truthfully. In this case, the scores will be 5, 4 and 3 for players 1, 2 and 3, respectively. In a similar way, player 2 can improve its score over the naive baseline by voting its level-0 strategy of $t \succ e \succ d \succ h$, given that the other players vote truthfully. In this case, the scores will be 4, 5 and 3 for players 1, 2 and 3, respectively. In fact, this voting profile in which player 2 deviates from its truthful vote, while player 1 and player 3 vote truthfully, is one of the Nash Equilibrium for this preference profile.

Player 3 is at a further disadvantage because there is no level-0 strategy that can improve its score over the baseline when other players are truthful. However, player 3 can improve its score when other players vote strategically. Specifically, when players 1 and 2 vote their level-0 strategy of $d \succ e \succ h \succ t$ and $t \succ e \succ d \succ h$ then player 3 can improve its score over the baseline by voting its level-1 strategy of $h \succ d \succ t \succ e$, incurring a score of 5, 4, and 4 points for players 1, 2 and 3, respectively.

In the second profile, called “non-symmetric”, the preferences of player 1 were $e \succ d \succ t \succ h$; the preferences of player 2 were $d \succ h \succ e \succ t$; the preferences of player 3 were $t \succ h \succ e \succ d$. If all participants vote truthfully, the chosen ranking will be $e \succ d \succ t \succ h$. In this case the scores will be 6, 3, and 2 for players 1, 2 and 3 respectively, putting player 1 in a significant advantage relative to the other participants. If player 2 votes its level-0 strategy of $d \succ h \succ t \succ e$, given that the other players vote truthfully, then players 2 and 3 will improve their score and player 1 will lose its advantage. In this case the chosen ranking will be $d \succ t \succ h \succ e$, and the scores will be 3, 4 and 3 for players 1, 2 and 3 respectively. This is also one of the Nash equilibrium for this game.

5 Empirical Methodology

We recruited 335 human subjects from the U.S. using Amazon Mechanical Turk. All participants were provided with an identical tutorial of how to play the budget allocation game, and their participation in the study was contingent on passing a quiz which tested their knowledge of the rules of the game. Participants were paid in a manner that was consistent with their performance, measured by aggregating their scores over 5 voting rounds.

The subjects were randomly divided into three different groups. The first group consisted of people playing the budget allocation game with other people.

Another group consisted of two people playing the game with another computer agent. The third group consisted of one person playing the game with two other computer agents. Each subject played five rounds of the game. In each group, there were between 13-17 games played, making for 65-85 rounds. We ran separate studies for both the symmetric and non-symmetric preference profiles described in Section 4.2.

We now describe the strategies used by two different computer agents that played the budget allocation game with other people. Let V_i^t denote the vote of player i at round t . The *truthful* (TR) agent ranked the candidates according to its preferences, that is, at each round t , agent i submits a vote V_i^t such that $V_i^t = P_i$. The *Previous Round Best Response* (PRBR) agent used the best-response vote of Equation 4 to rank the candidates, under the assumption that all other players repeat their vote in the previous round. Formally, at each round t , agent i submits a vote V_i^t such that $V_i^t = BR_i$ where for any player $j \neq i$, we have that $p(V_j^t) = 1$ if $V_j^t = V_j^{t-1}$ and 0 otherwise.

6 Results

We hypothesized that (1) people’s strategies will become more complex over time (involve less truthful strategies, and more best-response strategies as defined above); (2) that computational strategies using strategic reasoning (such as the PRBR agent) would be more successful when playing against people than computer agents that vote truthfully. All reported results in the upcoming section are significant in the $p < 0.05$ range using Analysis of Variance (ANOVA) tests.

6.1 Analysis of Human Behavior

We first present an analysis of people’s behavior in the game. People’s voting strategies were highly erratic. Out of 80 rounds of the budget allocation game that were played by three people, 64 rounds represented unique votes that appeared only once. In general, people’s strategy significantly deviated from the Nash equilibrium voting strategy. For example, there were only 7 out of 80 rounds played in the 3-person group configuration in which a Nash equilibrium strategy was played, which is not significantly different than random.

We measured the change in the number of “naive” votes (votes that are truthful and consistent with participants’ preferences) and best-response votes (votes that are a best-response to the votes of the other participants in the previous round). Figure 2 shows the difference in the average number of naive and best-response votes for each role between rounds 4-5 and rounds 1-2 for games that included three people or two people and one computer agent.

As shown in the Figure, there was a drop in the number of naive votes for all players between earlier and later rounds in the game, confirming our hypothesis. In addition, the figure also shows an increase in the number of best-response votes between earlier and later rounds in the game. We conjecture that the reason for

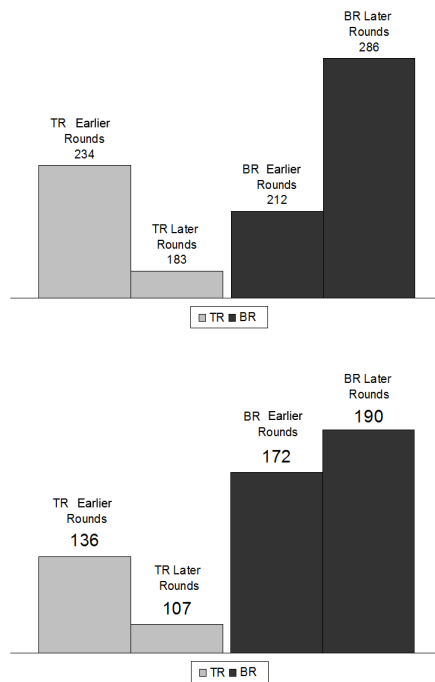


Fig. 2. Difference in Naive and Best-Response votes between earlier and later rounds in the game for symmetric (top) and non-symmetric preference profiles (bottom)

this discrepancy is that participants learned to make more sophisticated voting strategies. However, there was no significant increase in people’s scores as rounds progress. This aligns with past results in behavior economics studying complex aggregation rules [9]. Interestingly, (and not shown by the Figure) there was no increase in the number of best-response votes for people playing the role of Player 3 in the symmetric preference profile. We attribute this to the inherent disadvantage of this role in the game, in that it has a limited number of voting strategies that can improve its score in the game, as we described in Section 4.2. As shown in the bottom panel of the Figure, a similar pattern was also apparent for the non-symmetric preference profile.

6.2 Analysis of Performance

We now compare the performance of computer agents and people in groups comprising two other people (that is, each game included a person or a computer agent voting with two other people). Figure 3 shows the average performance of people and agents across all roles in the game for both preference profiles. As shown in the figure, the PRBR agent was able to outperform the TR agent, and both PRBR and TR agents were able to outperform people. Next, Figure 4




Fig. 3. Performance of computer agents and people in groups that include two other people for symmetric (left) and non-symmetric (right) preferences.

shows the performance of computer agents and people in groups comprising two other agents (that is, each game included a person or computer agent voting with two other agents). As shown by the Figure, the PRBR agent also outperformed people and the TR agent in this additional group configuration, demonstrating that the best-response strategy was independent of the group structure.

The PRBR agent also outperformed people and the TR agent in groups comprising another person and computer agent (that is, each game included a person or a computer agent interacting with another person and a computer agent). Figure 4 shows the average performance of people and agents across all roles in the game for both preference profiles. This result demonstrates that the best-response strategy was independent of the group structure. Interestingly, the TR agent was able to outperform people in the non-symmetric profile but not in the symmetric profile. Because of the structure of the non-symmetric profile, people lost more points from deviating from truthful behavior in this setting, to the benefit of the TR agent.

To compare performance for different roles, we present Table 1 which compares performance for each role in groups comprising a computer or person interacting with two other people for the symmetric preference profile. As shown

Type	Player 1	Player 2	Player 3
People	4.56	3.69	1.28
PRBR	4.87	4.04	2.78
TR	4.33	4.18	2.82

Table 1. Performance for different player roles in the symmetric preference profile

by the Table, in the role of Player 1, the PRBR agent was significantly more successful than the TR agent, while the TR agent was significantly more successful than people in both Player 2 and Player 3 roles. Note that although the TR agent scored higher than the PRBR agent in both Player 1 and Player 2 roles,



Fig. 4. Performance of computer agents and people in groups that include two other computer agents for symmetric (top) and non-symmetric (bottom) preferences.

this difference was not significant. The results for the asymmetric preference profile exhibited a similar pattern.

These results have implications for agent designers, suggesting that a fixed best-response strategy is sufficient towards enabling agents to perform well in voting systems in which participants submit full rankings. To show this, we provide a post-hoc analysis of the benefit of using the PRBR strategy given the observed behavior of people in the game for both preference profiles. Figure 5 shows the number of times where using the PRBR strategy would provide a positive or negative gain to player given people’s actual behavior in the game. As shown by the Figure, using the PRBR strategy was beneficial in the vast majority of cases, despite the fact that people do not actually repeat their vote from the previous round.

7 Conclusion

This paper described a first study comparing people’s voting strategies to that of computer agents in heterogeneous human-computer committees. In our setting participants vote by simultaneously submitting a ranking over the set of

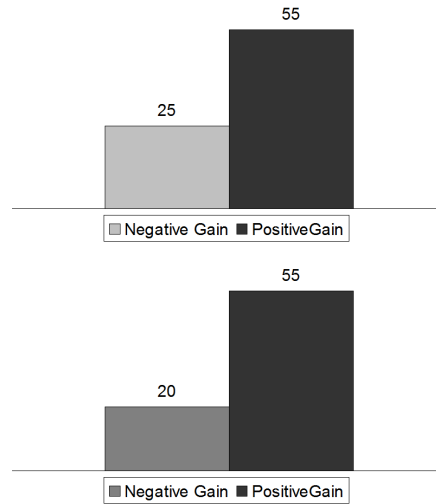


Fig. 5. Measuring the benefit of Previous-Round-Best-Response (PRBR) strategy for symmetric (top) and non-symmetric (bottom) preference

candidates and the election system uses a plurality rule to select a ranking that minimizes disagreements with participants' votes. Our results show that over time, people learned to deviate from truthful voting strategies, and use more sophisticated voting strategies. A computer agent using a best response voting strategy to people's actions in the previous round was able to outperform people in the game. In future work, we intend to design computer agents that adapt to people's play in settings of incomplete information.

8 Acknowledgments

This work is supported in part by the following grants: The Google Inter-university center for Electronic Markets and Auctions, ARO grants W911NF0910206, W911NF1110344 and MOST #3-6797 and Marie Curie grant #268362.

References

1. Riker, W., Ordeshook, P.: A theory of the calculus of voting. *The American Political Science Review* **62**(1) (1968) 25–42
2. Cox, G.: *Making votes count: strategic coordination in the world's electoral systems*. Volume 7. Cambridge Univ Press (1997)
3. Palfrey, T.: Laboratory experiments in political economy. *Annual Review of Political Science* **12** (2009) 379–388
4. Gatherer, D.: Comparison of eurovision song contest simulation with actual results reveals shifting patterns of collusive voting alliances. *Journal of Artificial Societies and Social Simulation* **9**(2) (2006)

5. Lin, R., Kraus, S.: Can automated agents proficiently negotiate with humans? *Communications of the ACM* **53**(1) (2010) 78–88
6. Meir, R., Polukarov, M., Rosenschein, J., Jennings, N.: Convergence to equilibria in plurality voting. In: *Proceedings of AAAI*. Volume 10. (2010) 823–828
7. Dhillon, A., Lockwood, B.: When are plurality rule voting games dominance-solvable? *Games and Economic Behavior* **46**(1) (2004) 55–75
8. Forsythe, R., Rietz, T., Myerson, R., Weber, R.: An experimental study of voting rules and polls in three-candidate elections. *International Journal of Game Theory* **25**(3) (1996) 355–383
9. Bassi, A.: Voting systems and strategic manipulation: an experimental study. Technical report, mimeo (2008)
10. Arrow, K.: Social choice and individual values. Number 12. Yale Univ Pr (1963)
11. Kemeny, J.: Mathematics without numbers. *Daedalus* **88**(4) (1959) 577–591