1

Voting by Adaptive Agents in Multi-Candidate Elections

Scott Moser

March 5, 2005

Abstract

Here we investigate how strategic adaptation changes our understanding and predictions of voting systems. We model populations of agents voting in multiple elections over time that revise their voting strategies based on electoral outcomes. Using this framework we examine the dynamic formation of electoral outcomes. Allowing agents to adapt their strategies refines the predictions of rational choice models and allows for a more realistic basis for understanding elections. We find that Approval and Plurality voting rules generate similar, relatively intuitive behavior, while Borda Count results in less intuitive, and often miscoordinated, outcomes.

1 Introduction

Understanding how social choices occur in a society has been a central focus of political science. Here we investigate how such choices emerge from agents who adapt to their environment. We are interested in the strategic effects various voting rules engender, and we focus on three candidate elections[1]. Predominantly, our understanding of voting rules has been grounded on rational choice assumptions of perfectly maximizing agents. However, the usual rationality assumptions may be to strict to describe realistic settings, as experimental evidence shows humans are not perfect utility maximizing agents. Scholars have noted the cognitive and informational limitations of humans [**?**]. We relax the usual maximizing assumption and investigate the role of adaptation in voting. That

[*]The author wishes to thank John Patty, Bill Keech, Maggie Penn and John Miller for their helpful comments and suggestions.

[1]If one shows up to vote, voting sincerely is dominant in 2 candidate elections, so we start with elections with three candidates.

is, when agents are able to adapt their voting strategy in response to past history, what types of strategies, and what types of outcomes are observed?

Formal theory provides a rigorous framework from which to study various voting rules. While rational choice models are wide-spread in decision theory, economics and political science, their usefulness is limited when studying voting games. As an equilibrium concept, rational choice models typically employ Nash equilibria. However, when applying the concept of Nash equilibria to voting games, one finds there are (generically) a plethora of equilibria. This is due to the fact that if an agent is not pivotal in determining the outcome of an election *any* action is a best response since the agent's action has no utility-relevant consequences. Further, in reasonably sized elections the probability of an agent being pivotal is extremely small. Even some Nash refinements do not reduce the set of equilibria to be useful [15]. In summery, while the Nash program is extremely useful in general, when it comes to analyzing voting games the Nash concept gives us little traction due to the large set of predictions, and no suitable *ex ante* justification for limiting the set of predictions.

We relax the usual rational choice assumptions by allowing agents to adapt their strategies. Specifically, we allow agents to make responsive adjustments to their environment and investigate its effect on voting systems. Our model allows us to examine voting outcomes in various environments. A novel feature of our model is its dynamic aspect: agent strategies adapt to a history of votes and outcomes. The adaptation process employed here is a simple dynamic that has been used in other studies to examine various strategic situations with great success [7]. We do not wish to claim that the proposed dynamic is an actual model of *how* agents learn, rather we use stylized elements of learning and adaptation to provide an alternate solution concept that yields tighter predictions. We then use this model to investigate differences between voting rules. One goal is to understand how strategic behavior by voters changes under different voting systems. Approaching this issue from an analytical standpoint has proven difficult, and so we use agent-based models.

Understanding the properties of various electoral systems is vital for democratic societies. While Arrow's famous theorem does dispel the idea of a 'perfect' voting rule [1], it remains open which available voting rule a society should use to aggregate the preference of its citizens. Clearly, the choice of voting rule is crucial in determining the outcome of an election. For example, it has been suggested by Tabarrok and Spector that the Civil War could have been avoided had an alternate voting rule been used in the 1860 presidential election [16]. Along these lines, some scholars have advocated certain voting rules. For example, Brams [2] has enumerated the benefits of using Approval voting, while Donald Saari [12] clearly endorses using Borda Count to aggregate preferences. Since the particular voting rule used can have a great effect on outcomes, we need to understand how various voting rules elicit what types of strategies, and what kinds of outcomes occur.

Prior analysis on electoral systems broadly fits into two categories. First, some scholars have examined how proportional various voting rules are – the "systematic" effect [5] of electoral systems. Such research typically compares vote shares at the district level with representation. Alternatively, we may evaluate voting rules based on their micro-level effects – what Farrell calls their "strategic effect" [5], [4]. Different voting rules may elicit different behavior from voters and candidates and work studying the strategic effect of voting rules examines these differences in agent behavior. For example, Merrill [13] examines the strategic effect of voting rules by using computer simulation and NES data to examine multicandidate election under various voting rules. He studies plurality, Borda, Approval, Coombs, and the Condorcet completion method of Black. He evaluates various voting rules on the basis of Condorcet efficiency and social utility. In his model, the position of candidates are fixed and exogenously given. He finds Codorcet and Hare runoff systems difficult to manipulate and Coombs and Borda count systems to invite strategic voting. When agents are assumed to vote sincerely, he finds plurality voting is categorically inferior to other voting rules.

Cox (1987) also analysis a variety of voting rules in multi-candidate elections, including scoring rules (such as Borda), completion methods, and negative voting. He employs formal game theoretic techniques to compare various voting rules by examining the equilibrium configuration of candidate positions. He finds plurality is alone in having only non-centrist equilibrium, while candidates position centrally in equilibrium for some scoring rules and Condorcet completion methods. Further, when more than three candidates compete in one dimension, he finds a marked difference between plurality, for which only non-centrist equilibria exist and other voting rules, for which candidates position at the median voters ideal point. In addition, prior work has indicated there are differences in the amount of truthful revelation of voters' preferences. Single member plurality systems induces a lot of strategic voting, while proportional representation does not [5].

Our work contributes to existing work comparing various voting rules in multi-candidate elections. We evaluate the strategic effects of voting rules on agents' strategies. The rest of the paper is organized as follows. The model of adaptive voting is explained in section 2, followed by results from example preference distributions in section 3. Experiments are conducted altering some parameters of the model and the results are discussed.

2 The Model

We focus on the outcomes agents' strategies produce in simple voting games. For simplicity, we assume there are three candidates or parties A, B and C, whose positions are fixed. Our model does not consider the motivations of candidates: their positions are fixed and exogenously determined. Following usual practice, we define an agent's type to be the utility an agent would receive from each of the candidates winning an election. Agents are randomly selected to vote but not every agent votes in every election. In order for adaptive voting to have an effect, the outcome must be relevant to one's actions. That is, the probability of being pivotal cannot be trivial. By allowing agents to vote in small groups, useful feedback from the environment may convey information about he success of a voting strategy. A winner is then determined based on the votes of agents and utilities (payoffs) are assigned to voting agents. After several such elections, agents undergo a strategy revision period in which better performing strategies are replicated, but with some perturbation. A novel aspect of this framework is the co-evolution of strategies. A system coevolves if the composition of one group of agents changes in response to a change in another group. Coevolution in this setting allows the frequency of strategies employed by one type of agents to change in response to a change in the strategies employed by agents with different preferences. When comparing strategy fitness, only agents of a similar type are compared. That is, a strategy for an agent who prefers candidate A over candidate B over candidate C (we represent agents of this type by $A \succ B \succ C$) is only compared to the strategies of other $A \succ B \succ C$-type agents.

We begin with a population for each type of agent. Each population consists of 50 agents, and within each population, agents have the same fixed preference ordering over candidates. For simplicity, we fix the utility an agent may possibly get: the utility an agent receives if her first choice candidate is elected is 1, the utility an agent receives if her least preferred candidate wins is 0 and the utility an agent receives if her intermediate candidate wins is .5. Hence, with three candidates, there are $3! = 6$ possible voter types. Agents' strategies adapt based on the cumulative utility an agent earns – strategies that earn higher utility are used with greater frequency.

In each election, not all agents vote. Instead, agents vote in small groups and winners are determined based on these votes. In each election, a fraction of the total agents are chosen at random (without replacement) to participate in the current election. The composition of the group is controlled so that a particular distribution of types is realized in each election. This is done according to a preference distribution that is exogenously given. For example, agents may be drawn from all populations identically and independently (i.i.d.). Alternatively, there may be only one type of agent present in each election (of course, these extremes, with interesting preference distributions in between). For a richer example, if a given preference distribution dictates that 5 out of 10 voting agents are $A \succ B \succ C$-type agents, then five agents are randomly selected from the $A \succ B \succ C$-type population. The flexibility of altering the composition of the group of agents participating in each election is explored later. Each agent possesses a pure strategy such that each type is mapped to one particular vote. Based on the agent's strategy, a vote is cast. It is the agents' strategies that are subjected to change and adaptation.

Depending on the aggregation rule used an agent chosen to vote in the current election submits either (i) a single candidate (as in the Plurality rule), (ii) a ranked list of the candidates, from top ranked to least preferred (i.e. the Borda Count), (iii) a ballot consisting of the candidates for which the agent "approves" (Approval Voting) or (iv) a ranked list of candidates as in Instant Runoff (also known as *alternative voting*).

After all chosen agents have voted, the winner of the election is calculated by the appropriate aggregation rule[2]. Agents then receive utility based on the winner of the election and according to their type. All agents are then "returned" to their respective populations and the process is repeated. After several elections (here 600), the populations are modified[3] in the following way: two agents are selected from a population at random (and with replacement) and the agent with the higher average utility (from the past 600 elections) is placed in a new population. This process is repeated until the new population is of the same size as the old. The updating process used here captures a very simple notion of adaptation, which has been shown to be a robust method of locating optima in a variety of settings from control theory to economics (see [6] and [7]). By selecting better performing strategies the system moves in the direction of improvement (from a social welfare perspective). The introduction of mutations is beneficial to the search process by avoiding lock-in on suboptimal strategies. Hence, the simple dynamic used here captures the notions of adaptation and learning,

[2]Ties are broken by randomly choosing a candidate in the win set.

[3]If no agents from a population participated in any of the 600 elections, the population does not undergo modification.

while relaxing the usual strict rationality (or perfect maximizing) assumption typically employed in rational choice models. Also, akin to many evolutionary game theory models [14], we require very little rationality on the part of the agents: they simply respond, at prescribed intervals, to their comparative payoff ranking. We use this process more as an equilibrium-seeking device than as an actual model of how voters learn.

Agents evolve their strategies in separate populations and for a strategy to be successful it need only outperform other strategies in its population. That is, successful agents possess strategies that outperform agents with similar preferences.

3 Results

The framework employed here allows us to control the distribution of preferences of voting agents. For example, half of a group of agents voting in an election may prefer candidate A over candidate B over candidate C ($A \succ B \succ C$-type agents) while the other half may have exactly the opposite preferences ($C \succ B \succ A$-type agents). It is the distribution of agent preferences that we refer to as the "preference distribution". To begin, we consider all equal 2-type preference distributions. That is, all (up to symmetry) preference distributions that consist of two different types of agents, each represented equally in each election.

For each experiment, initial conditions are randomly selected and 600 elections are held before strategies are first updated. This process continues for 300 generations. The system settles down quickly and to ameliorate initial condition-effects the first 30 generations are omitted from analysis. In addition this entire experiment is repeated 50 times, each time initial strategies are randomly assigned. For each run, average data is recorded. We present the aggregate outcomes of the elections including which candidates won and how often, in addition to what 'types' of strategies agents employ. We introduce a taxonomy of voting strategies and show the frequency of strategies used.

Within a given run, each generation yields an average payoff, averaged over all elections and all agents in a population. This generation-average is then averaged over all generations (excluding the first 30 generations as discussed above) to yield a run-average. The average and variance of the run-averages are then calculated and shown below.

3.1 Preference Distributions Involving Two Types of Voters

In this section, we examine situations in which only two different preferences are present. Since there are 3 candidates, there are 5 distinct preference distributions consisting of two types of voters.[4]. They are:

1. $A \succ B \succ C$-type agents and $B \succ A \succ C$-type agents
2. $A \succ B \succ C$-type agents and $C \succ B \succ A$-type agents
3. $A \succ B \succ C$-type agents and $A \succ C \succ B$-type agents
4. $A \succ B \succ C$-type agents and $B \succ C \succ A$-type agents
5. $A \succ B \succ C$-type agents and $A \succ B \succ C$-type agents

Of these five different preference distributions with two types of preferences, only 3 are interesting. The $A \succ B \succ C : A \succ B \succ C$ preference distribution is trivial, as all agents in this situation have the exact same preferences. The $A \succ B \succ C : A \succ C \succ B$ preference distribution is uninteresting as all agents agree on the top-ranked candidate (candidate A).

Using classical non-cooperative game theoretic concepts such as Nash equilibrium [**?**] does not provide much help in predicting outcomes in the various preference distributions. There are often a plethora of (Nash) equilibria in voting games. In fact, we can construct vote vectors supporting any candidate winning such that no agent is pivotal. Hence, when agents have only one of two types of preferences, any candidate winning is a Nash equilibrium of the voting game[5]. As such, this concept is not useful in making predictions. So we consider the theoretical predictions of Nash equilibria when agents use weakly undominated strategies (hereafter WUD strategies). A strategy s_i is weakly undominated for an agent, i, if for any action taken by the other agents in the game, it does at least as well as any other available strategy s'_i. [11] discuss the Nash equilibria when only WUD are used by agents and [**?**] discusses what strategies are WUD in a variety of voting games.

To find Nash equilibria in WUD strategies, we examine vote vectors that are Nash equilibria, when agents only use strategies that are WUD. Recall that if an agent is not pivotal, he is best

[4]Clearly, there are $\binom{6}{2} = 15$ ways of picking 2 types from 6 possible types. But by eliminating symmetries, we are left with the above 5.

[5]Provided there are sufficient number of agents.

responding and so the problem of finding outcomes supported by WUD strategies reduces to finding vote vectors arising from combinations of WUD strategies such that no one is pivotal[6]. As we will see below, this restriction on permissible strategies only marginally helps make more precise predictions. Other scholars have developed equilibria concepts that are more suited to the analysis of voting games, such as Meyerson and Weber's voting equilibria [10] and McKelvey and Palfrey's quantal response equilibria [9]. Our work contributes to the literature on voting equilibria by examining the outcomes resulting from adaptive agents (not necessarily perfectly rational) myopically seeking to improve their welfare.

The $A \succ B \succ C$: $B \succ A \succ C$ preference distribution

We begin by examining an environment in which there are two types of agents: $A \succ B \succ C$-type agents and $B \succ A \succ C$-type agents. That is, all agents agree that candidate C is their least preferred candidate, but the agents are divided on who is the most preferred candidate. With each type of agent equally represented in the population, what outcomes should we expect? Since candidate C is not in the core, we would certainly not expect to see candidate C winning. Further, because of the equal number of each type of agent, we would expect candidate A and candidate B to split equally the elections won.

When plurality voting is used, candidate C winning cannot be a Nash equilibrium when agents use WUD strategies: voting for one's least preferred candidate is dominated in plurality voting (by voting for *any* other candidate) and all agents agree that candidate C is the least preferred candidate. Under Borda Count, candidate A and B winning are both possible Nash equilibrium in WUD strategies, as all types voting (A,B,C) produces an outcome of candidate A winning (respectively B). When approval voting is used, candidate C winning cannot be a Nash in WUD strategies as approving for one's least preferred candidate is dominated[7]. Likewise, under Instant Runoff, candidate C cannot win in a Nash in WUD strategies.

For each voting rule, Figure 1 shows the average fraction of elections each candidate won. For example, under Plurality, Approval Voting and Instant Runoff candidates A and B each win about half of the elections. As can be seen in the Figure 1, the predictions of Nash in WUD strategies

[6]If no one is pivotal, everyone is best responding.

[7]A ballot XX1 that approves of an agent's least preferred candidate is dominated by XX0, a ballot whose only difference is not approving of the least preferred candidate.

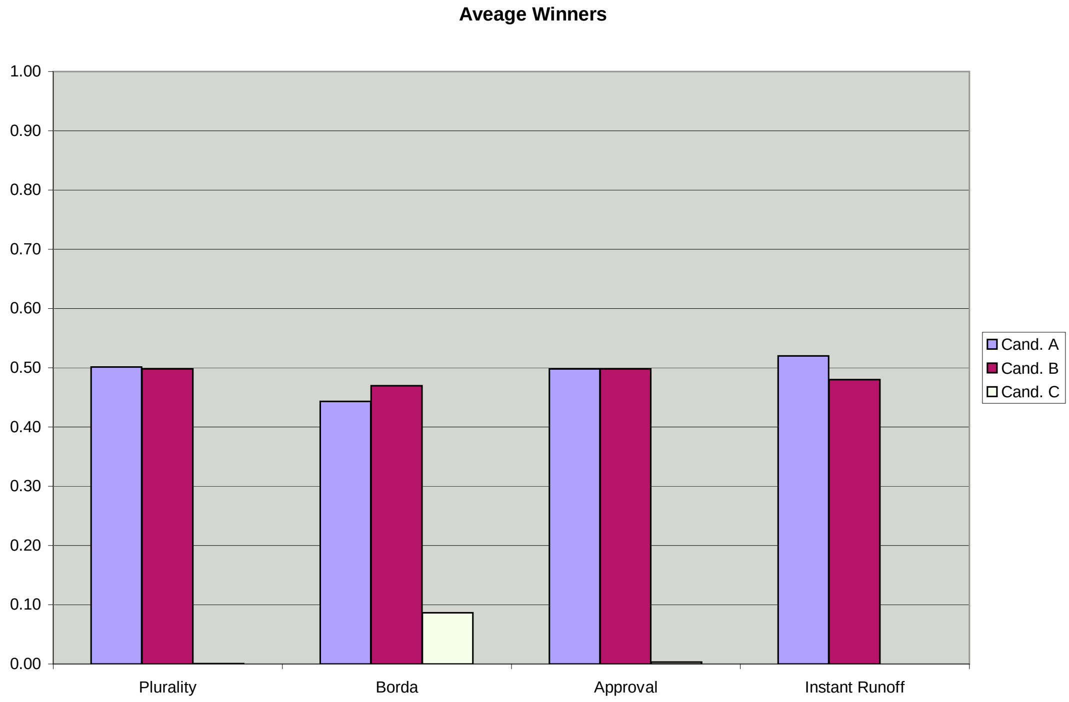

Figure 1: Average Winners for $A \succ B \succ C : B \succ A \succ C$ preference distribution

are generally supported. For each of the aggregation rules, candidate A and candidate B each win about half of the elections. Notice, however, the substantial amount of elections won by candidate C under Borda Count (about 8.7% of the elections).

FIGURE 1 HERE

Hence, when agents agree on their least preferred candidate, that candidate does not win. The Borda voting rule offers a slight exception to this, as a non-trivial amount of elections are won by candidate C. To further explore these results, we examine the types of strategies agents employ. To aid in this investigation, we introduce a taxonomy of strategies that allows us to compare strategies used by agents from different populations. The taxonomy of strategies will allows us to track the kinds of strategies being employed, and to measure the amount of sincere preference revelation on the part of the agents (that is, the agreement between an agent's vote and his preferences)

Recall that under Plurality, an agent may vote for either his most preferred candidate, his second most preferred candidate, or his least preferred candidate (the latter of which is always a weakly dominated strategy). Hence, there are K possible votes under the Plurality rule, where K is the number of candidates. In each election, we record the kind of strategy each voting agent employs. For example, if a $A \succ B \succ C$-type agent votes for candidate A under a Plurality rule, we say that agent is using a strategy of type 0. If he votes for her second most preferred candidate (candidate 1), he is using a strategy of type 1. Likewise, if he votes for his least preferred candidate (2), we say he is using a strategy of type 2. Hence, since there are three different votes an agent may cast under Plurality rule, there are 3 different kinds of strategies an agent may employ: vote for his most preferred candidate; vote for his second most preferred candidate or; vote for his least preferred candidate.

In a similar fashion, with K candidates there are $K!$ votes under Borda Count and Instant Runoff. However, under Instant Runoff, the lower ranked candidates might or night not be influential, depending on if and agent's top ranked candidate is eliminated of not. If an agent's vote agrees with his preference ordering, we say he employs a strategy of type 0 – he truthfully reveals his preferences. If an agent gives 3 votes to his least preferred candidate, 2 to his middle ranked candidate and 1 to his most preferred candidate, we say he uses a strategy of type 5 (and we question the rationality of such an agent). For three candidates, a complete list of the possible votes along with their strategic type is given in Table 1. The key to this taxonomy is to express an agent's vote as a function of his preferences and the strategy employed.

Finally, if we assume voters do not abstain, there are $2^K - 1$ possible votes in an Approval Voting. For three candidates, we order the 7 possible votes as strategies of type 0 to 6, with lower numbers representing "closer" to sincere. Table 1 shows, for three candidates, the possible votes for each aggregation rule and the associated strategic type of each kind of vote.

TABLE 1 HERE

Hence, knowing an agent's preferences and knowing the ballot he casts, we may classify his strategy as described above. This classification does not depend on the the actions of other, but of course the expected utility of a strategy does. Roughly speaking the ordering of strategies goes from "sincere" to "strategic"(or "sophisticated") to "stupid".

Strategic Type	Plurality Vote	Borda and Instant Runoff Vote	Approval Vote
0	first	(first, second, third)	first
1	second	(second, first, third)	first and second
2	third	(first, third, second)	second
3		(second, third, first)	first and third
4		(third, first, second)	all candidates
5		(third, second, first)	second and third
6			third only

Table 1: The taxonomy of votes an agent may cast. Here "first" corresponds to the agent's most preferred candidate, etc. For example, if Approval voting is used and an agent has preferences $A \succ B \succ C$ and vote to approve of candidates A and B, then we say he is using 'strategy 1' (str1). Note that when there are 3 candidates, plurality yields 3 different types of strategies, Borda and Instant Runoff 6 and Approval 7 different strategies.

With this classification of strategies, we can investigate the types of strategies that emerge in our model. Turning our attention back to the $A \succ B \succ C : B \succ A \succ C$ preference distribution, we can examine the strategies that support the outcomes we observe. Figure 2 shows the average proportion of each type of vote, for each population. Within each population, we record the average percentage of each strategy being used, in each generation. For example, when plurality is used, over 90% of the agents in both populations are voting sincerely (str0 in the figure).

FIGURE 2 HERE

Notice that, perhaps not surprisingly, the dominant strategies under Plurality and Approval Rules are sincere voting. Also notice the symmetry between the two populations – this is exactly what we would expect given the symmetry in the preference distribution. More surprisingly, however, are the strategies developed by agents voting under Borda Count: they give three votes to their most preferred candidate, two vote for their *least preferred* candidate, and one vote to their second most preferred. Notice that one population of agents' first choice is the other population's second choice. Hence, giving more votes to one's least preferred candidate than to one's opponents most preferred candidate is payoff superior to voting sincerely. Consider the following: if all agents

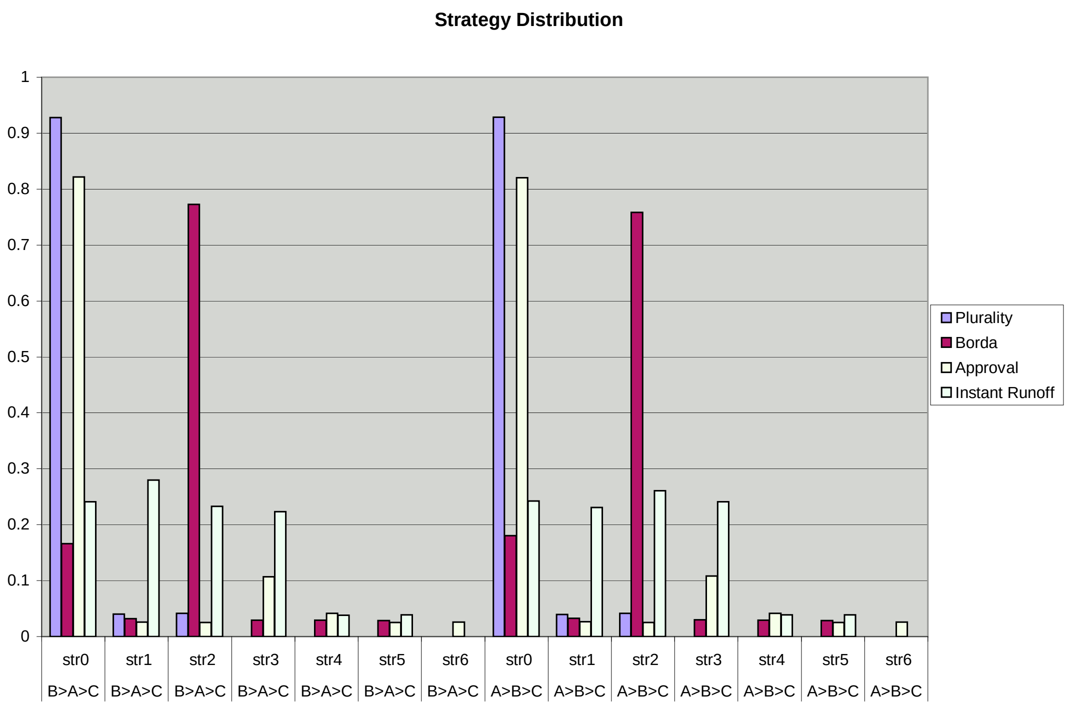

Figure 2: Average Frequency of Voting Strategies for the $A \succ B \succ C$: $B \succ A \succ C$ preference distribution. The abscissa lists for each population the various kinds of strategies. For example, "str 1, $A > B > C$" shows, for each voting rule, the frequency of strategy 1 used by the agents with preferences $A > B > C$.

with different preferences than myself are voting sincerely, then by voting (first, last, second), I [8] can guarantee my most preferred candidate wins. If others are voting (first, last, second), then myself also voting (first, last, second) causes a A-B tie.

In this preference distribution, agents voting via Plurality and Approval Rule settle on truthfully revealing their preferences. However, agents voting repeatedly under Borda Count develop a sort of "race to the bottom" where they give three votes to their most preferred candidate, two votes to their *least preferred* candidate, and one vote to their opponent's most preferred candidate. The action of voting for one's least preferred candidate over one's opponent's most preferred candidate helps explain the substantial amount of elections candidate 2 wins under Borda Count (see Figure 1). It is unclear if this strategic behavior would lead agents to rank as second their least preferred candidate in general, or if this is an artifact of only having 3 candidates. This curious "race to the bottom" phenomena needs to be explored further.

The $A \succ B \succ C$: $C \succ B \succ A$ preference distribution

We next turn out attention to a situations of diametrically opposed preferences: the preference distribution in this section involves only $A \succ B \succ C$-type agents and $C \succ B \succ A$-type agents. Here, there are Nash equilibria in WUD strategies that support any candidate winning under Plurality rule, Borda Count, Approval voting and Instant Runoff. Hence, in this setting, the predictions of Nash in WUD strategies are vacuous. However, when agents are allowed to adapt their strategies we see a large effect of aggregation rule on the average outcome of elections. As Figure 3 shows, the "extreme" candidates A and B win most of the time under Approval, Plurality and Instant Runoff voting, while the middle ranked candidate (candidate B) wins most of the elections under Borda Count. The "smoothing" effect of the Borda Count method has been noted [12]. In this setting we see that Plurality and Approval voting generate similar outcomes, although a substantial amount of elections are won by candidate A under Approval voting, whereas virtually no elections are won by candidate A under plurality. The reason for this is that Approval voting allows for "hedging" of one's second most preferred candidate, and plurality does not. In fact, the strategies that agents use under Approval, Plurality and Instant Runoff rules are qualitatively similar, with sincere voting emerging under both systems. However, Approval voting allows the possibility of

[8] Conditional on being pivotal

voting for more than one candidate, and we see approximately 26% of agents voting for their second most preferred candidate, as well as their top ranked candidate (see Figure 4). This structural feature allows agents to gain more social utility under Approval voting than the other voting rules examined[9] (see Figure 7). Notice there is "drift" in strategies employed under Instant Runoff. Both strategies of type 0 and type 2 list one's most preferred candidate first, so that if there is no runoff, or if one's most preferred is does not receive the minimum of votes, strategies of type 0 and 2 are functionally equivalent. However, since the agents are voting for their most preferred candidate, we consider such strategies as sincere. Lastly, notice the outcomes under Borda Count are qualitatively different: candidate A, the commonly agreed upon second best candidate wins more than any other candidate (candidate A wins about 38% of the time, with the remaining elections being split evenly by candidates B and C). Again looking to the strategies that agents use under Borda counts, we see that sincerity is by far the most popular strategy being used (almost 90% of all agents truthfully reveal there preference ordering under Borda here – see Figure 4).

When agents have diametrically opposed preferences, we see a large effect of voting rule on outcomes. Under Plurality, the extreme candidate split the elections, under Approval and similar outcome is observed, with some non trivial hedging, and under Borda, the commonly agreed upon second best candidate, B, wins most of the time. Despite the effect of voting rule on outcomes, the strategies that emerge in this preference distribution are by and large all sincere.

FIGURE 3 HERE.

FIGURE 4 HERE.

As can clearly be seen, this preference distribution induces sincere voting regardless of aggregation rule. Hence, when preferences are exact opposites, sincere voting results, with some "hedging the bet" under Approval. Additionally, the average utility enjoyed by agents is comparable across the various voting rules and is shown in figure 7. As can be seen in figure 7, average societal happiness if roughly .5.

The $A \succ B \succ C$: $B \succ C \succ A$ preference distribution

The predictions of Nash are as follows. Under Plurality, Borda, Approval and Instant Runoff, any candidate may win as a WUD Nash. However, under Approval Voting, candidate A can only

[9]The average social utility under Plurality was 0.496, 0.513 under Borda and 0.515 under Approval.

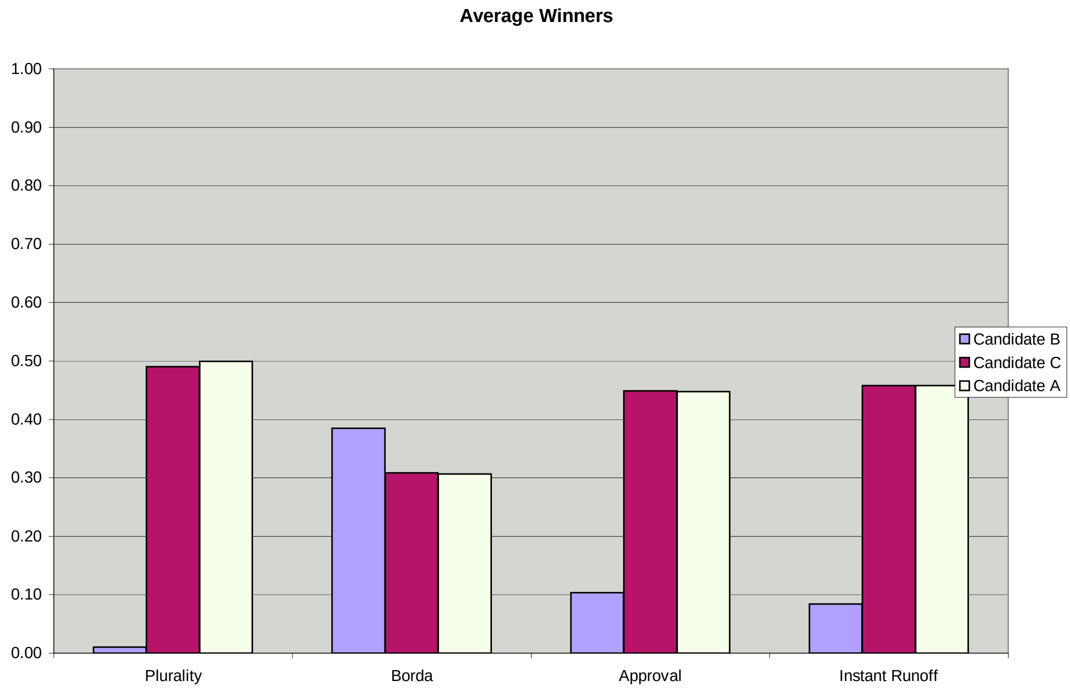

Figure 3: Frequency of Elections Won in the $A \succ B \succ C : C \succ B \succ A$ preference distribution

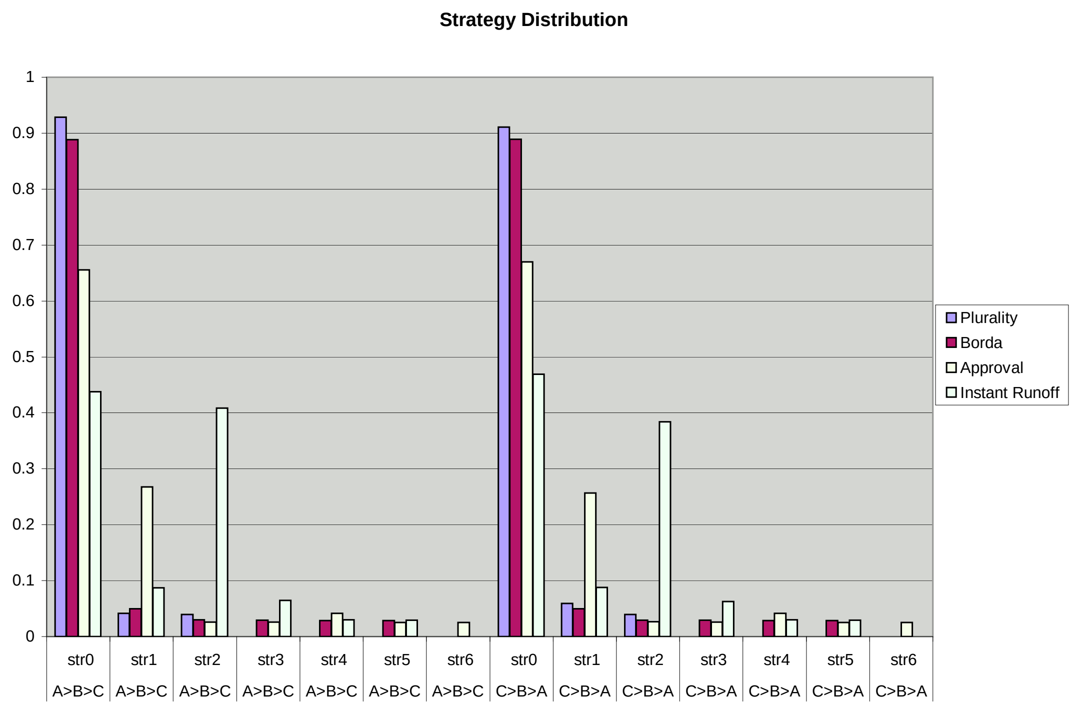

Figure 4: Average Frequency of Voting Strategies for the $A \succ B \succ C$: $C \succ B \succ A$ preference distribution. The abscissa lists the various kinds of strategies for each population. For example, "str 1, A¿B¿C" shows, for each voting rule, the frequency of strategy 1 used by the agents with preferences A¿B¿C.

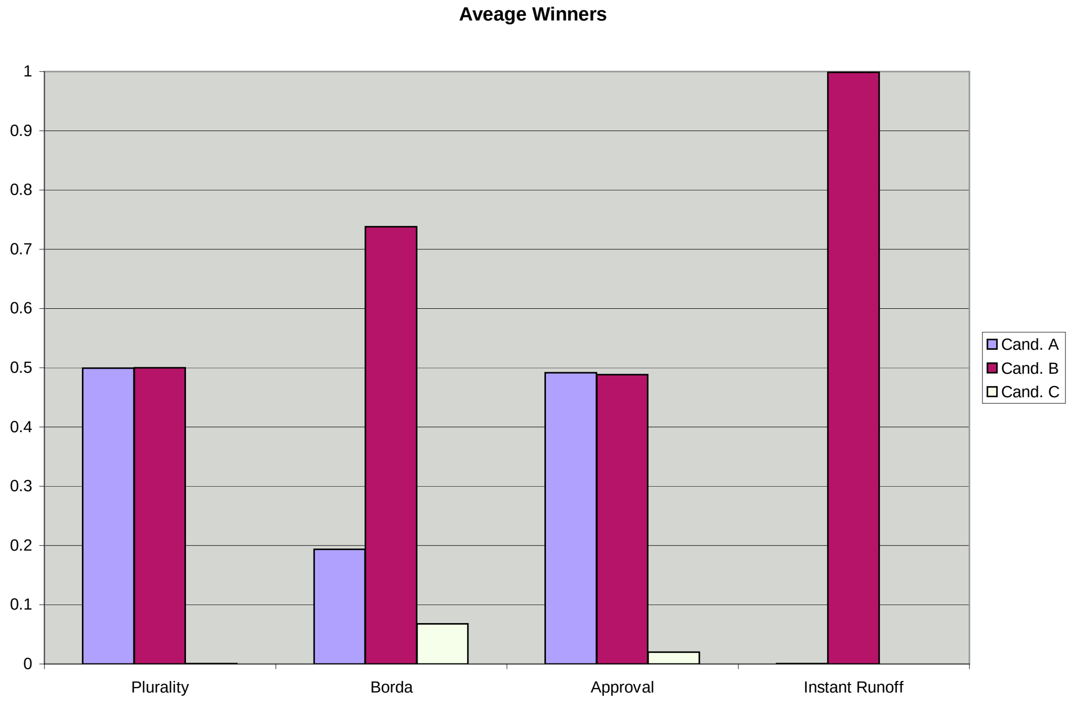

Figure 5: Frequency of Elections Won in the $A \succ B \succ C : B \succ C \succ A$preference distribution

win as the result of a tie being broken.

FIGURE 5 HERE

FIGURE 6 HERE

In this preference distribution, Instant Runoff affords agents a greater average social utility (.77 under Instant Runoff versus 0.627 under Plurality, 0.69 under Borda and 0.646 under Approval). Hence, *ex ante* agents would prefer to vote in a Instant Runoff system. In addition, Plurality and Approval generate similar outcomes as well as generating similarly sincere strategies.

Now, within a profile of agents' preferences, we can compare the average maximizing utility agents are able to obtain. Below is a graph comparing the average social utilities agents were able

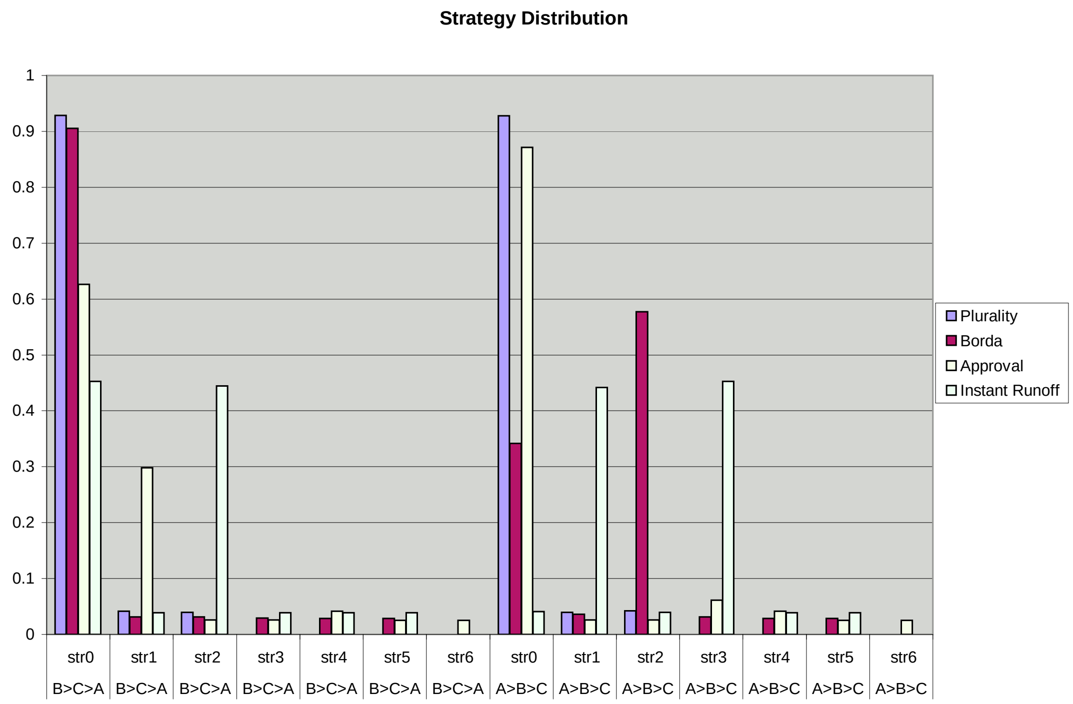

Figure 6: Frequency of Strategies in the $A \succ B \succ C : B \succ C \succ A$ preference distribution

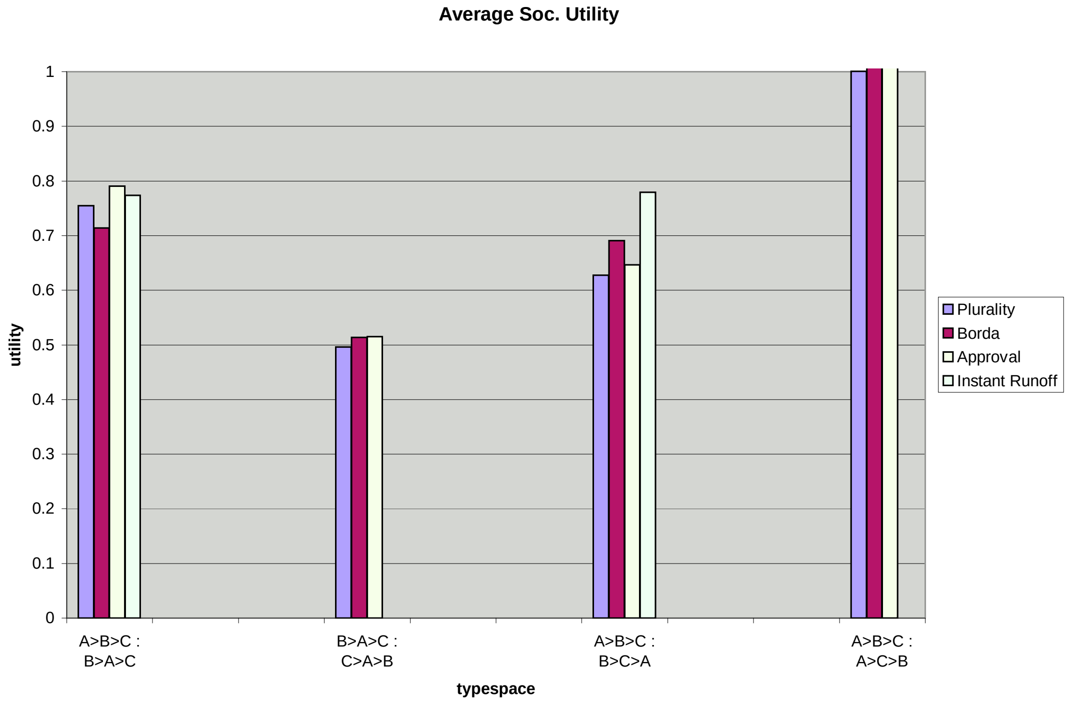

Figure 7: Average social utility.

to achieve, on average, while maximizing utilities in different aggregation rules.

FIGURE 7 HERE

Note the comparison across preference distributions in Figure 7. From this figure, it appears some preference distributions are easier for agents to 'figure out' in the sense that some preference profiles afford higher utility to agents in our model. The extreme case is the $A \succ B \succ C : A \succ C \succ B$ preference distribution.

Hence, it is not clear that one voting rule dominates another (from a social utility standpoint), even in these simple preference profiles. It should be stressed the preference distributions used thus far are extremely simple and exhaust all profiles where only two different preferences are present.

3.2 Preference Distributions With Three Types

While examining 2-type preference distributions provides some initial lessons, situations in which there are only two types of agents are in some ways limited in the scope of political situations they can model. In this section, we present the results from some simple, but politically interesting environments in which there are three types of agents. To begin with, we look at an example in which candidate C is preferred last by the majority of the agents, but sincere voting yields candidate C winning. Specifically, in each election five candidates prefer $C \succ B \succ A$ and three candidates each have preferences $B \succ A \succ C$ and $A \succ B \succ C$ respectively. The interesting aspect of this setting lies in the large gain to coordination *across* preference types (across populations): Do the "minority" agents coordinate their behavior to avoid candidate C being elected? If so, how? And on which candidate, A or B? Once again, the predictions of Nash equilibria in weakly undominated strategies are vacuous: any candidate may be elected under any voting rule. Figure 8 shows the average winners in our model.

FIGURE 8 HERE

As we can see, the agents are able to coordinate to prevent candidate C from winning on most occasions, with type of voting rule playing a large role in the number of elections the majority least-preferred candidate (C) is elected. Specifically, in this case Borda Count is superior in minimizing the amount of candidate C winning and maximizing candidate B winning (which is optimal from a social welfare perspective – see figure **??**).

From figure 9, we see this is accomplished by $B \succ A \succ C$ and $C \succ B \succ A$ agents voting sincerely, while $A \succ B \succ C$ agents are both sincere and strategic voters. The majority group has an incentive to vote sincerely, while the minority group with diametrically opposed preferences from the majority has a strategic incentive to not vote their true preferences.

FIGURE 9 HERE

To get more traction on the performance of the difference voting rules when three different voter preferences are present, we ran additional experiments exploring a much wider range of preference distributions. Consider a spatial model in two dimensions in which the ideal point of each of three types of voters are fixed. With the voters' ideal points fixed, we investigate how successful various candidate positions are. Picking 50 candidate positions on a grid, we pitted each of these 50 candidate positions against 20 randomly drawn pairs of opponents. Figure 10 shows the results.

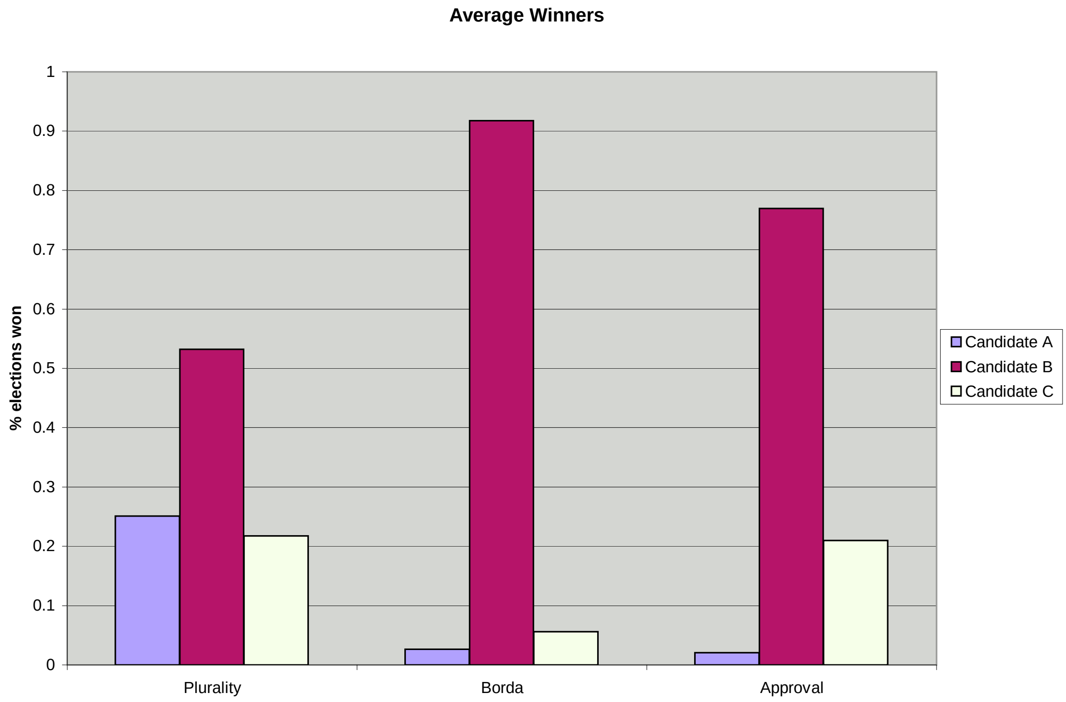

Figure 8: Average winners for the T2 preference distribution.

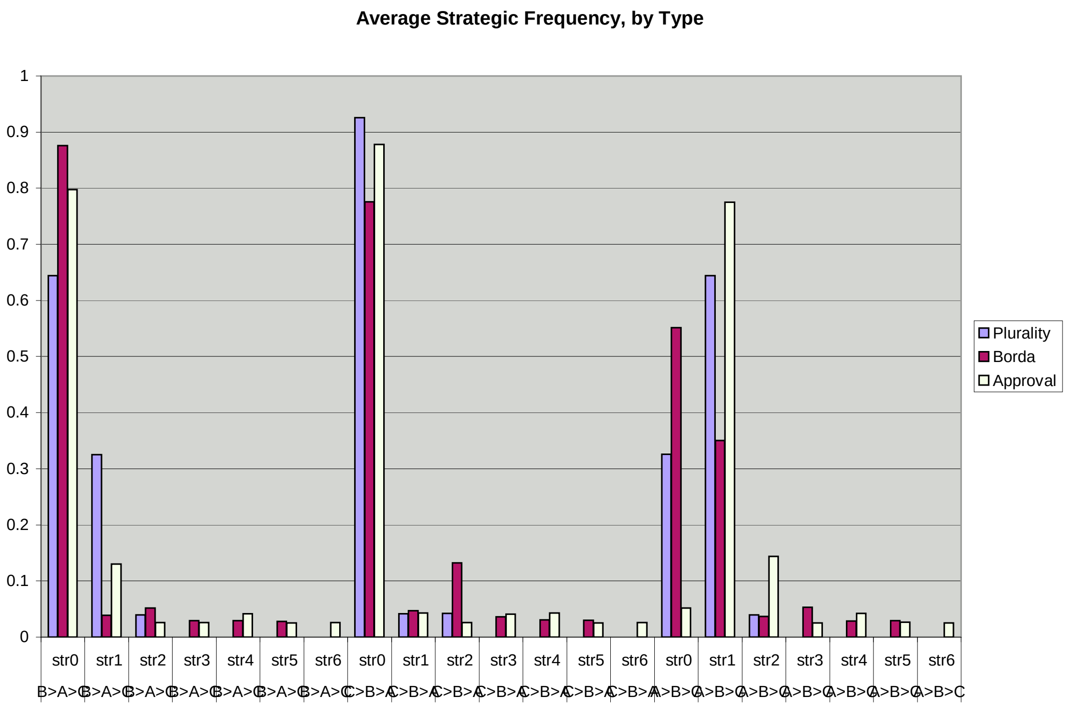

Figure 9: Average strategic frequency.

Each dot represents a candidate position and the intensity of the dot represents the average number of elections each position won (when competing against randomly drawn opponents). The voters ideal points are the vertices of the triangle, located at $(0,0)$, $(0,1)$ and $(1/2, \sqrt{3}/2)$.

FIGURE 10 HERE

Notice that positions in the interior are the most successful candidate positions, across voting rules. In fact, the centroid appears to be the most robust to randomly drawn opponents. So, we see a differentiation of candidate positions by how many elections they win. But in Figure 10 we do not see the difference in voting rules observed earlier. The aggregation of multiple preference distributions appears to have "washed out" differences across voting rules. Of course, differences in optimal candidate positions is just one difference we may be interested in when comparing electoral systems [3], [5]. Future work will examine the social utility and strategic implications of such aggregate voting results.

4 Conclusions

We have devolved a model of adaptive agents voting and adjusting strategies over time. The model has two big advantages: it allows for a more sensible behavioral assumptions than rational choice models and it helps refine predictions. It should be emphasized that when there is a clear prediction (when the core is non-empty), our model coincides with the prediction. Being a computational model, we are able to see the process clearly, but are unable to draw as general conclusions as in analytic theory. That is, we are able to study particular settings (and in the 2-type world, we are able to study all 2-type scenarios with three alternatives), but are unable to say without qualification how the system performs in generic settings.

Perhaps unintuitively, by weakening assumptions of rationality, we get more precise outcome predictions. These outcomes can then be used to compare voting rules, in terms of outcomes and the strategies that emerge. Under such a set up, Plurality and Approval voting look very similar, both in terms of outcomes, and in terms of the strategies that agents employ. The strategies we see emerging under Approval voting are essentially "Plurality plus some small hedging the bet". Both Plurality and Approval generate generally sincere strategies in a variety of preference distributions. Borda behaves differently, sometimes selecting the social utility maximizing candidate (or the "middle ground"), but also allowing for more miscoordination. Despite its similarity to the Borda ballot,

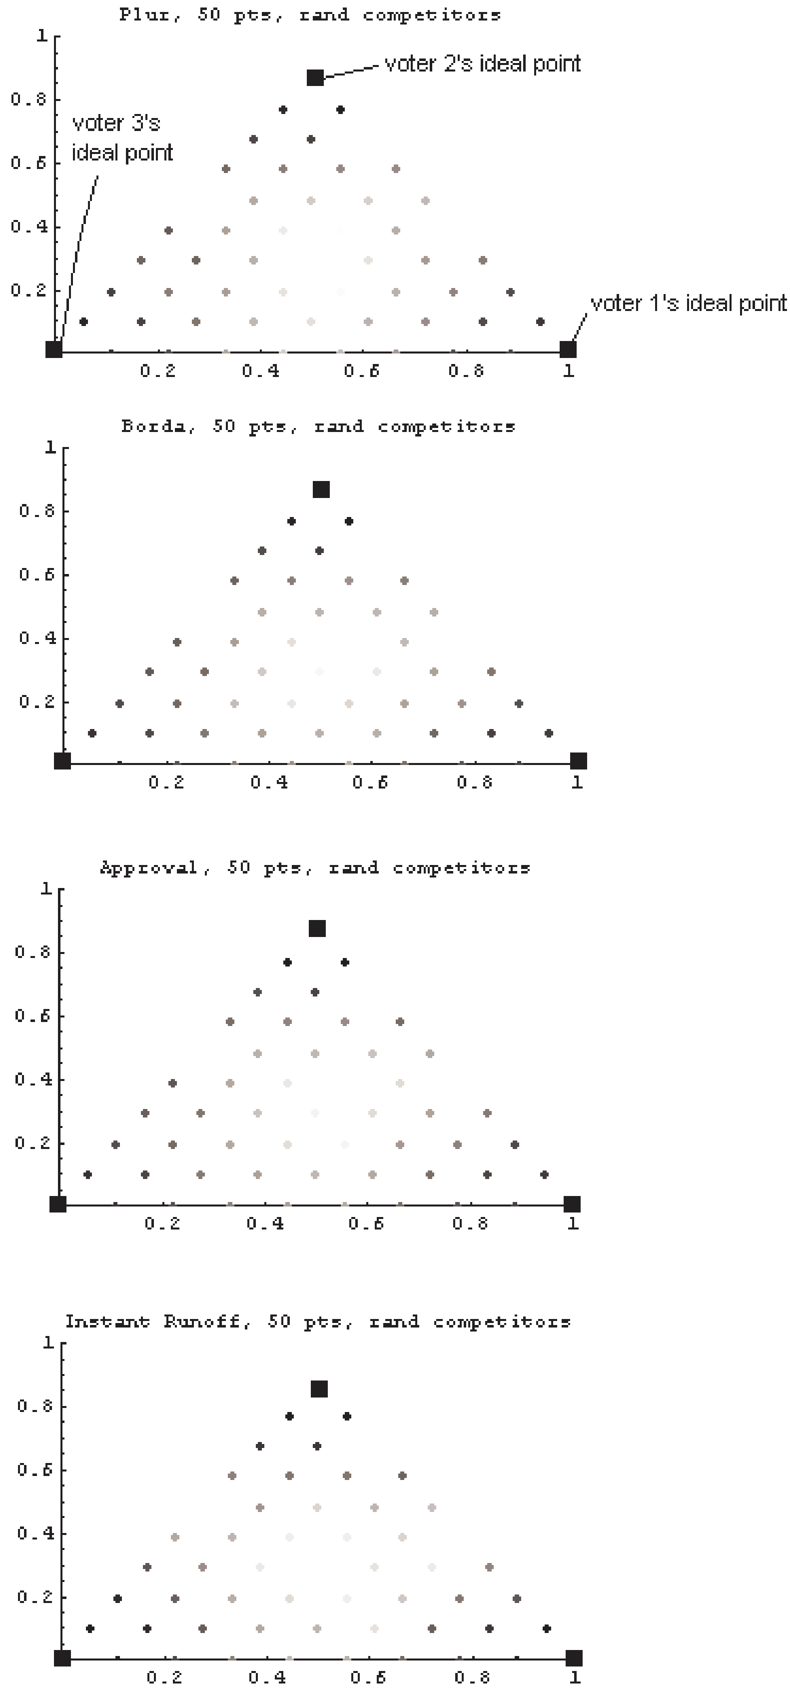

Figure 10: Average winners with 3 types. For each candidate position (each ‘dot’ on the grid), the average number of elections won against random opponents was recorded. Darker dots indicate fewer elections were won.

Instant runoff appears as a Plurality-Approval hybrid when considering outcomes and strategies.

The model is extremely flexible and allows for a panoply of other voting questions to be address. We anticipate two immediate extensions. First, we plan to use this model to study turnout and abstention in voting games. Second, this paper has largely focused on static comparisons of various rules. However, there is a rich dynamic story telling *how* various outcomes are achieved that would be interesting to examine. In addition, future work will explore more generally the differences in various voting rules. For example, we have presented results when there are only two or three different preferences in a society, but this certainly can be extent to other preference distributions. In addition, future work will examine arbitrary (but more than three) candidates. The sometime un-intuitive results observed under Borda voting (the "race to the bottom" phenomena discussed earlier) can be more fully developed with arbitrary candidates, and future work will seek to illuminate the mechanisms by which such phenomena occur.

References

[1] Kenneth Arrow, *Social choice and individual values*, Yale University Press, 1970.

[2] Steven J. Brams and Peter C. Fishburn, *Approval voting*, Birkhauser, 1983.

[3] Gary W. Cox, *Electoral equilibrium under alternative voting institutions*, American Journal of Political Science **31** (1987), no. 1, 82–108.

[4] ———, *Making votes count: Strateig coordination in the world's electoral systems*, Cambridge University Press, 1997.

[5] David M. Farrell, *Electoral systems: a comparative introduction*, Palgrave, 2001.

[6] John H. Holland, *Hidden order: How adaptation builds complexity*, Addison Wesley Publishing Company, 1996.

[7] John H. Holland and John H. Miller, *Artificial adaptive agents in economic theory*, American Economic Review, Papers and Proceedings **81** (1991), 365–370.

[8] Richard McKelvey and Thomas Palfrey, *Quantal response equilibria for normal form games*, Games and Economic Behavior **10** (1995), 6–38.

[9] Roger B. Myerson and Robert J. Weber, *A theory of voting equilibria*, The American Political Science Review **87** (1993), no. 1, 102–114.

[10] Thomas R Palfrey and Sanjay Srivastava, *Nash implementation using undominated strategies*, Econometrica **59** (1991), no. 2, 479–501.

[11] Donald Saari, *Basic geometry of voting*, Springer-Verlag, 1995.

[12] III Samuel Merrill, *Making multicandidate elections more democratic*, Princeton University Press, 1988.

[13] Larry Samuelson, *Evolutionary games and equilibrium selection*, Economic Learning and Social Evolution, The MIT Press, 1998.

[14] Herbert Simon and Allen Newell, *Human problem solving*, Prentice Hall, 1972.

[15] Fancesco De Sinopoli, *Sophisticated voting and equilibrium refinements under plurality rule*, Social Choice and Welfare **17** (2000), 655–672.

[16] Alexander Tabarrod and Lee Spector, *Would the Borda count have avoided the civil war?*, Journal of Theoretical Politics **11** (1999), no. 2, 261–288.