Empirical Analysis of Plurality Election Equilibria

David R. M. Thompson
University of British Columbia
daveth@cs.ubc.ca

Omer Lev
The Hebrew University of Jerusalem
omer.lev@mail.huji.ac.il

Kevin Leyton-Brown
University of British Columbia
kevinlb@cs.ubc.ca

Jeffrey Rosenschein
The Hebrew University of Jerusalem
jeff@cs.huji.ac.il

ABSTRACT

Voting is widely used to aggregate the different preferences of agents, even though these agents are often able to manipulate the outcome through strategic voting. Most research on manipulation of voting methods studies (1) limited solution concepts, (2) limited preferences, or (3) scenarios with a few manipulators that have a common goal. In contrast, we study voting in plurality elections through the lens of Nash equilibrium, which allows for the possibility that any number of agents, with arbitrary different goals, could all be manipulators. This is possible thanks to recent advances in (Bayes-)Nash equilibrium computation for large games. Although plurality has numerous pure-strategy Nash equilibria, we demonstrate how a simple equilibrium refinement—assuming that agents only deviate from truthfulness when it will change the outcome—dramatically reduces this set. We also use symmetric Bayes-Nash equilibria to investigate the case where voters are uncertain of each others' preferences. This refinement does not completely eliminate the problem of multiple equilibria. However, it does show that even when agents manipulate, plurality still tends to lead to good outcomes (e.g., Condorcet winners, candidates that would win if voters were truthful, outcomes with high social welfare).

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Multiagent systems*

General Terms

Economics, Experimentation, Theory

Keywords

Social choice theory, voting protocols, game theory

1. INTRODUCTION

When multiple agents have differing preferences, voting mechanisms are often used to decide among the alternatives. One desirable property for a voting mechanism would be strategy-proofness, i.e., that it would be optimal for agents to truthfully report their preferences. However, the Gibbard-Satterthwaite theorem [12, 27] shows that no non-dictatorial strategy-proof mechanism can exist. Whatever other desirable properties a voting mechanism may have, there will always be the possibility that some participant can gain by voting strategically.

Since voters may vote strategically (i.e., manipulate or counter-manipulate) to influence an election's results, according to their knowledge or perceptions of others' preferences, much research has considered ways of limiting manipulation. This can be done by exploiting the computability limits of manipulations (e.g., finding voting mechanisms for which computing a beneficial manipulation is NP-hard [2, 1, 30]), by limiting the range of preferences (e.g., if preferences are single-peaked, there exist non-manipulable mechanisms [10]), randomization [13, 25], etc.

When studying the problem of vote manipulation, nearly all research falls into two categories: coalitional manipulation and equilibrium analysis. Much research into coalitional manipulation considers models in which a group of truthful voters faces a group of manipulators who share a common goal. Less attention has been given to Nash equilibrium analysis which models the (arguably more realistic) situation where all voters are potential manipulators. One reason is that it is difficult to make crisp statements about this problem: strategic voting scenarios give rise to a multitude of Nash equilibria, many of which involve implausible outcomes. For example, even a candidate who is ranked last by all voters can be unanimously elected in a Nash equilibrium—observe that when facing this strategy profile, no voter gains from changing his vote. Another problem is that finding even a single Nash equilibrium of a game can be computationally expensive, and plurality votes can have exponentially many equilibria (in the number of voters).

Despite these difficulties, this paper considers the Nash (and subsequently, Bayes-Nash) equilibria of voting games. We focus on plurality, as it is by far the most common voting mechanism used in practice. We refine the set of equilibria by adding a small additional assumption: that agents realize a very small gain in utility from voting truthfully; we call this restriction a *truthfulness incentive*. We ensure that this incentive is small enough that it is always overwhelmed by the opportunity to be pivotal between any two candidates: that is, a voter always has a greater preference for swinging an election in the direction of his preference than for voting truthfully. All the same, this restriction is powerful enough

Appears in: *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2013),* Ito, Jonker, Gini, and Shehory (eds.), May 6–10, 2013, Saint Paul, Minnesota, USA.

to rule out the bad equilibrium described above, as well as being, in our view, a good model of reality, as voters might reasonably have a preference for voting truthfully.

We take a computational approach to the problem of characterizing the Nash equilibria of voting games. This has not previously been done in the literature, because the resulting normal-form games are enormous. For example, representing our games (e.g., 10 players and 5 candidates) in the normal form would require about a hundred million payoffs ($10 \times 5^{10} \simeq 9.77 \times 10^7$). Unsurprisingly, these games are intractable for current equilibrium-finding algorithms, which have worst-case runtimes exponential in the size of their inputs. We overcame this obstacle by leveraging recent advances in compact game representations and efficient algorithms for computing equilibria of such games, specifically action-graph games [15, 14] and the support-enumeration method [28].

Our first contribution is an equilibrium analysis of full-information models of plurality elections. We analyze the number of Nash equilibria that exist when truthfulness incentives are present. We also examine the winners, asking questions like how often they also win the election in which all voters vote truthfully, or how often they are also Condorcet winners. We also investigate the social welfare of equilibria; for example, we find that it is very uncommon for the worst-case result to occur in equilibrium. Our approach can be generalized to richer mechanisms where agents vote for multiple candidates (i.e., approval, k-approval, and veto).

Our second contribution involves the possibly more realistic scenario in which the information available to voters is incomplete. We assume that voters know only a probability distribution over the preference orders of others, and hence identify Bayes-Nash equilibria. We found that although the truthfulness incentive eliminates the most implausible equilibria (i.e., where the vote is unanimous and completely independent of the voters preferences), many other equilibria remain. Similarly to Duverger's law (which claims that plurality election systems favor a two-party result [9], but does not directly apply to our setting), we found that a close race between almost any pair of candidates was possible in equilibrium. Equilibria supporting three or more candidates were possible, but less common.

1.1 Related Work

Analyzing equilibria in voting scenarios has been the subject of much work, with many researchers proposing various frameworks with limits and presumptions to deal with both the sheer number of equilibria, and to deal with more real-life situations, where there is limited information. Early work in this area, by McKelvey and Wendell [20], allowed for abstention, and defined an equilibrium as one with a Condorcet winner. As this is a very strong requirement, such an equilibrium does not always exist, but they established some criteria for this equilibrium that depends on voters' utilities.

Myerson and Weber [23] wrote an influential article dealing with the Nash equilibria of voting games. Their model assumes that players only know the probability of a tie occurring between each pair of players, and that players may abstain (for which they have a slight preference). They show that multiple equilibria exist, and note problems with Nash equilibrium as a solution concept in this setting. The model was further studied and expanded in subsequent research [4, 16]. Assuming a slightly different model, Messner and Polborn [22], dealing with perturbations (i.e., the possibility that the recorded vote will be different than intended), showed that equilibria only includes two candidates (Duverger's law). Our results, using a different model of partial information (Bayes-Nash), show that with the truthfulness incentive, there is a certain predilection towards such equilibria, but it is far from universal.

Looking at iterative processes makes handling the complexity of considering all players as manipulators simpler. Dhillon and Lockwood [6] dealt with the large number of equilibria by using an iterative process that eliminates weakly dominated strategies (a requirement also in Feddersen and Pesendorfer's definition of equilibrium [11]), and showed criteria for an election to result in a single winner via this process. Using a different process, Meir et al. [21] and Lev and Rosenschein [19] used an iterative process to reach a Nash equilibrium, allowing players to change their strategies after an initial vote with the aim of myopically maximizing utility at each stage.

Dealing more specifically with the case of abstentions, Desmedt and Elkind [5] examined both a Nash equilibrium (with complete information of others' preferences) and an iterative voting protocol, in which every voter is aware of the behavior of previous voters (a model somewhat similar to that considered by Xia and Contizer [29]). Their model assumes that voting has a positive cost, which encourages voters to abstain; this is similar in spirit to our model's incentive for voting truthfully, although in this case voters are driven to withdraw from the mechanism rather than to participate. However, their results in the simultaneous vote are sensitive to their specific model's properties.

Rewarding truthfulness with a small utility has been used in some research, though not in our settings. Laslier and Weibull [18] encouraged truthfulness by inserting a small amount of randomness to jury-type games, resulting in a unique truthful equilibrium. A more general result has been shown in Dutta and Sen [8], where they included a subset of participants which, as in our model, would vote truthfully if it would not change the result. They show that in such cases, many social choice functions (those that satisfy the No Veto Power) are Nash-implementable, i.e., there exists a mechanism in which Nash equilibria correspond to the voting rule. However, as they acknowledge, the mechanism is highly synthetic, and, in general, implementability does not help us understand voting and elections, as we have a predetermined mechanism. The work of Dutta and Laslier [7] is more similar to our approach. They use a model where voters have a lexicographic preference for truthfulness, and study more realistic mechanisms. They demonstrated that in plurality elections with odd numbers of voters, this preference for truthfulness can eliminate all pure-strategy Nash equilibria. They also studied a mechanism strategically equivalent to approval voting (though they used an unusual naming convention), and found that when a Condorcet winner exists, there is always a pure-strategy Nash equilibrium where the Condorcet winner is elected.

2. DEFINITIONS

Elections are made up of candidates, voters, and a mechanism to decide upon a winner:

Definition 1. *Let C be a set of m candidates, and let A*

be the set of all possible preference orders over C. Let V be a set of n voters, and every voter $v_i \in V$ has some element in A which is his true, "real" value (which we shall mark as a_i), and some element of A that he announces as his value, which we shall denote as $\tilde{a}_i$.
The voting mechanism itself is a function $f : A^n \to C$.

Note that our definition of a voter incorporates the possibility of him announcing a value different than his true value (strategic voting).

In this paper, we restrict our attention to scoring rules, in which each voter assigns a certain number of points to each candidate, and specifically, to scoring rules in which each candidate can get at most 1 point from each voter. Mainly, we will deal with plurality, but we will touch on some more:

- **Plurality**: A point is only given to a single candidate.
- **Veto**: A point is given to everyone except one candidate.
- **k-approval**: A point is given to exactly k candidates.
- **Approval**: A point is given to as many candidates as each voter chooses.

Another important concept is that of a Condorcet winner.

DEFINITION 2. *A Condorcet winner is a candidate $c \in C$ such that for every other candidate $d \in C$ ($d \neq c$) the number of voters that rank c over d is at least $\lfloor \frac{n}{2} \rfloor + 1$.*

Condorcet winners do not exist in every voting scenario, and many voting rules—including plurality—are not Condorcet-consistent (i.e., even when there is a Condorcet winner, that candidate may lose).

To reason about the equilibria of voting systems, we need to formally describe them as games, and hence to map agents' preference relations to utility functions. More formally, each agent i must have a utility function $u_i : A^n \mapsto \mathbb{R}$, where $u_i(a_V) > u_i(a'_V)$ indicates that i prefers the outcome where all the agents have voted a_V over the outcome where the agents vote a'_V. Representing preferences as utilities rather than explicit rankings allows for the case where i is uncertain what outcome will occur. This can arise either because he is uncertain about the outcome given the agents' actions (because of random tie-breaking rules), or because he is uncertain about the actions the other agents will take (e.g., agents behaving randomly; agents play strategies that condition on information i does not observe). Here we assume that an agent's utility only depends on the candidate that gets elected and on his own actions (e.g., an agent can get some utility for voting truthfully). Thus, we obtain simpler utility functions $u_i : C \times A \mapsto \mathbb{R}$, with an agent i's preference for outcome a_V denoted $u_i(f(a_V), \tilde{a}_i)$.

In this paper, we consider two models of games, full-information games and symmetric Bayesian games. In both models, each agent must choose an action $\tilde{a}_i$ without conditioning on any information revealed by the voting method or by the other agents. In a full-information game, each agent has a fixed utility function which is common knowledge to all the others. In a symmetric Bayesian game, each agent's utility function (or "type") is an independent, identically distributed draw from a commonly known distribution of the space of possible utility functions, and each agent must choose an action without knowing the types of the other agents, while seeking to maximize his expected utility.

We consider a plurality voting setting with voters' preferences chosen randomly. We show detailed results for the case of 10 voters and 5 candidates (numbers chosen to give a setting both computable and with a range of candidates), but we also show that changing these numbers results in similar characteristics of equilibria.

Suppose voter i has a preference order of $a^5 \succ a^4 \succ \ldots \succ a^1$, and the winner when voters voted a_V is a^j. We then define i's utility function as

$$u_i(f(a_V), \tilde{a}_i) = u_i(a^j, \tilde{a}_i) = \begin{cases} j & a_i \neq \tilde{a}_i \\ j + \epsilon & a_i = \tilde{a}_i, \end{cases}$$

with $\epsilon = 10^{-6}$.

Note that we use utilities because we need, when computing an agent's best response, to be able to compare nearly arbitrary distributions over outcomes (e.g., for mixed strategies or Bayesian games). This is not meant to imply that utilities are transferable in this setting. Most of our equilibria would be unchanged if we moved to a different utility model, provided that the preferences were still strict, and the utility differences between outcomes were large relative to ϵ. The one key distinction is that agents are more likely to be indifferent to lotteries (e.g., an agent that prefers $A \succ B \succ C$ is indifferent between $\{A, B, C\}$ and $\{B\}$; under a different utility model, the agent might strictly prefer one or the other).

As with perfect information games, we consider Bayesian games with a fixed number of candidates (m) and voters (n). The key difference is that the agents' preferences are not *ex ante* common knowledge. Instead, each agent's preferences are drawn from a distribution $p_i : A \mapsto \mathbb{R}$. Here we consider the case of symmetric Bayesian games, where every agent's preferences are drawn independently from the same distribution, p. Due to computational limits, we cannot study games where p has full support; each agent would have $5^{5!} \simeq 7.5 \times 10^{83}$ pure strategies. Instead, we consider distributions where only a small subset of preference orders are in the support of p. We generate distributions by choosing six preference orderings, uniformly at random (this gives a more reasonable $5^6 = 15625$ pure strategies). For each of these orderings a, we draw $p(a)$ from a uniform $[0, 1]$ distribution. These probabilities are then normalized to sum to one. This restricted support only affects what preference orders the agents can have; agents' action sets are not restricted in any way.

Note that formally the ϵ truthfulness incentive represents a change to a game, rather than a change in solution concept. However, there is an equivalence between the two approaches: for any sufficiently small ϵ, the set of pure-strategy Nash equilibria in the game with ϵ truthfulness incentives is identical to the set of pure-strategy Nash equilibria (of the game without truthfulness incentives) that also satisfy that only the pivotal agents (i.e., agents who, were their vote to change, the outcome would change) deviate from truthfulness. The meaning of sufficiently small depends on the agents' utility functions, and on the tie-breaking rule. If $\underline{u}$ is the difference in utility between two outcomes, and $\underline{t}$ is the minimum probability of any type profile (in a Bayesian game), then ϵ must be less than $\underline{ct}/|C|$ (the $1/|C|$ factor comes from the fact that uniform tie-breaking can select some candidate with that probability).

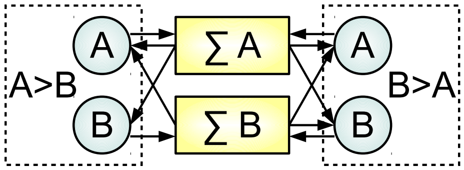

Figure 1: An action graph game encoding of a simple two-candidate plurality vote. Each round node represents an action a voter can choose. Dashed-line boxes define which actions are open to a voter given his preferences; in a Bayesian AGG, an agent's type determines the box from which he is allowed to choose his actions. Each square node is an adder, tallying the number of votes a candidate received.

3. METHOD

Before we can use any Nash-equilibrium-finding algorithm, we need to represent our games in a form that the algorithm can use. Because normal form games require space exponential in the number of players, they are not practical for games with more than a few players. The literature contains many "compact" game representations that require exponentially less space to store games of interest, such as congestion [26], graphical [17], and action-graph games [15]. Action-graph games (AGGs) are the most useful for our purposes, because they are very compactly expressive (i.e., if the other representations can encode a game in polynomial-space then AGGs can as well), and fast tools have been implemented for working with them.

Action-graph games achieve compactness by exploiting two kinds of structure in a game's payoffs: anonymity and context-specific independence. Anonymity means that an agent's payoff depends only on his own action and the number of agents who played each action. Context-specific independence means that an agent's payoff depends only on a simple sufficient statistic that summarizes the joint actions of the other players. Both properties apply to our games: plurality treats voters anonymously, and selects candidates based on simple ballot counts.

Encoding plurality games as action-graph games is relatively straightforward. For each set of voters with identical preferences, we create one action node for each possible way of voting. For each candidate, we create an adder node that counts how many votes the candidate receives. Directed edges encode which vote actions contribute to a candidate's score, and that every action's payoff can depend on the scores of all the candidates (see Figure 1). The same approach generalizes to approval-based mechanisms; if an action involves approving more than one candidate, then there must be an edge from that action node to each approved candidate's adder node. Similarly, positional scoring rules (e.g., Borda) can be encoded by using weighted adders.

A variety of Nash-equilibrium-finding algorithms exist for action-graph games [15, 3]. In this work, we used the support enumeration method [24, 28] exclusively because it allows Nash equilibrium enumeration. This algorithm works by iterating over possible supports, testing each for the existence of a Nash equilibrium. In the worst case, this requires exponential time, but in practice SEM's heuristics (exploiting symmetry and conditional dominance) enable it to find all the pure-strategy Nash equilibria (PSNEs) of a game quickly.

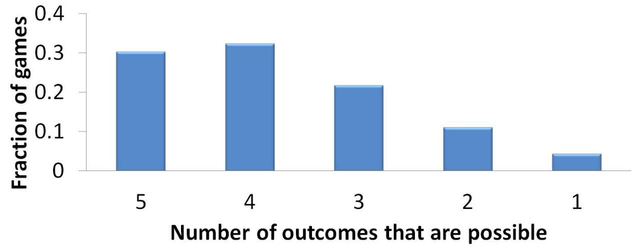

Figure 2: Even with the truthfulness incentive, many different outcomes are still possible in equilibrium.

We represented our symmetric Bayesian games using a Bayesian game extension to action-graph games [14]. Because we were concerned only with *symmetric* pure Bayes-Nash equilibria, it remained feasible to search for every equilibrium with SEM.

4. PURE-STRATEGY NASH EQUILIBRIUM RESULTS

To examine pure strategies, we ran 1000 voting experiments using plurality with 10 voters and 5 candidates.

4.1 Selectiveness of the truthfulness incentive

The logical first question to ask about our truthfulness incentive is whether or not it is effective as a way of reducing the set of Nash equilibria to manageable sizes. As a baseline, each plurality game had over a million PSNEs, when we did not use any equilibrium selection method. A stronger baseline is the number of PSNEs that survive removal of weakly dominated strategies (RDS). RDS reduces the set by an order of magnitude, but still allows over 100,000 PSNEs per game to survive. In contrast, the truthfulness incentive reduced the number of PSNEs down to 30 or less, with the median game having only 3. Interestingly, a handful of games (1.1%) had no PSNEs. Laslier and Dutta [7] had shown that PSNEs were not guaranteed to exist, but only when the number of voters is odd (and at least 5). Our results show that the same phenomenon can occur, albeit infrequently, when the number of voters is even.

One of the problems with unrestricted Nash equilibria is that there are so many of them; the other problem is that they are compatible with any outcome. Given that the TI is so effective at reducing the set of PSNEs, one could wonder whether or not TI is helpful for the second problem. Unfortunately, TI has only limited effectiveness in ruling out some outcomes as impossible (only 4.4% of games support exactly one outcome in equilibrium). However, nearly always ($> 99\%$ of the time) some outcomes occur more frequently than others. See Figures 2 and 3.

4.2 Equilibrium outcomes

With a workable equilibrium selection method, we can now consider the question of what kinds of outcomes occur in plurality. We shall examine two aspects of the results: the preponderance of equilibria with victors being the voting method's winners, and Condorcet winners. Then, moving to the wider concept of social welfare of the equilibria (which we can consider since we work with utility functions), we examine both the social welfare of the truthful voting rule vs. best and worse possible Nash equilibria and the average rank of the winners in the various equilibria.

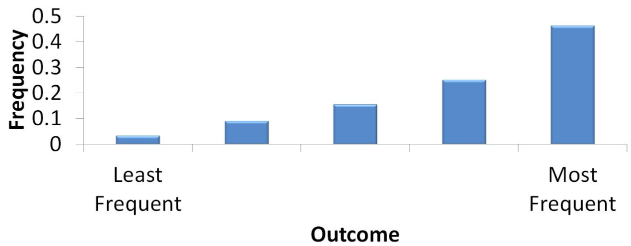

Figure 3: With the truthfulness incentive, some outcomes occur much more frequently than others.

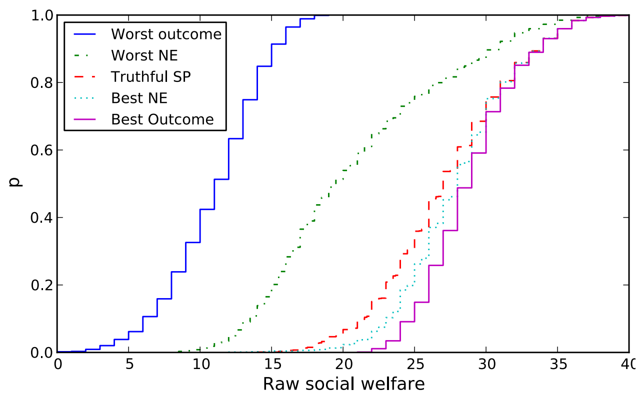

Figure 4: CDF of social welfare.

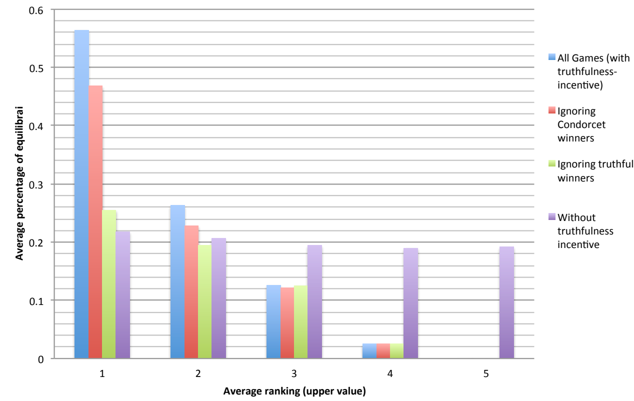

Figure 5: The average proportion of equilibria won by candidates with average rank of 0–1, 1–2, etc.

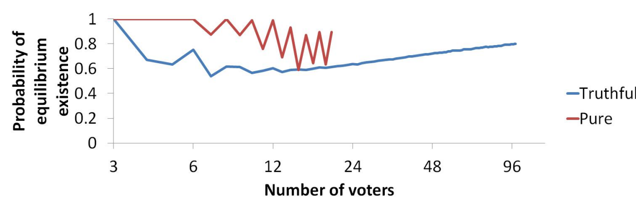

Figure 6: Percentage of games with a Nash equilibrium of a given type with 3 candidates and varying number of voters.

The first issues to consider are to what extent truthful voting is an equilibrium, and to what extent the agents cancel out each other's manipulations (i.e., when there are non-truthful Nash equilibria that lead to the same outcome as truthful voting). We call a candidate the "truthful winner" iff that candidate wins when voters vote truthfully. For 63% of the games, the truthful preferences were a Nash equilibrium, but more interestingly, many of the Nash equilibria reached the same result as the truthful preferences: 80% of the games had at least one equilibrium where the truthful winner wins, and looking at the multitudes of equilibria, 42% elected the truthful winner (out of games with a truthful result as an equilibrium, the share was 52%). Without the truthfulness incentive, the truthful winner only wins in 22% of equilibria.

Next, we turn to the question of whether or not strategic voting leads to the election of good candidates, starting with Condorcet winners. 55% of games had Condorcet winners, which would be elected by truthful voting in 49% of the games (not a surprising result; plurality is known not to be Condorcet consistent). However, the combination of truthful voting both being an equilibrium and electing a Condorcet winner only occurred in 42% of games. In contrast, a Condorcet winner could win in some Nash equilibrium in 51% of games (though only 36% of games would elect a Condorcet winner in every equilibrium).

Turning to look at the social welfare of equilibria, once again, the existence of the truthfulness incentive enables us to reach "better" equilibria. In 93% of the cases, the worst-case outcome was not possible at all (recall that without the truthfulness incentive, every result is possible in some Nash equilibrium), while only in 30% of cases, the best outcome was not possible. While truthful voting led to the best possible outcome in 59% of cases, it is still stochastically dominated by best-case Nash equilibrium (see Figure 4).

When looking at the distribution of welfare throughout the multitudes of equilibria, one can see that the concentration of the equilibria is around high-ranking candidates, as the average share of equilibria by candidates with an average ranking (across all voters in the election) of less than 1 was 56%. (See Figure refaverageRankWinner.) Fully 72%, on average, of the winners in every experiment had above (or equal) the median rank, and in more than half the experiments (52%) all equilibria winners had a larger score than the median. As a comparison, the numbers from experiments without the truthfulness incentive, are quite different: candidates—whatever their average rank—won, with minor fluctuations, about the same number of equilibria (57% of winners, were, on average, above or equal to the median rank).

4.3 Scaling behavior and stability

We next varied the number of voters and candidates. Our main finding was that when voter number was odd, the probability of having no equilibria at all increased dramatically, as the truthfulness incentive causes many such situations to be unstable (see Figure 6). Less surprisingly, as the number of candidates increased PSNEs were less likely to exist.

Nevertheless, these equilibria equilibria retained the properties we have seen—a concentration of equilibria around "quality" candidates, such as truthful and Condorcet winners, as can be seen in Figure 7a, for truthful winners (it

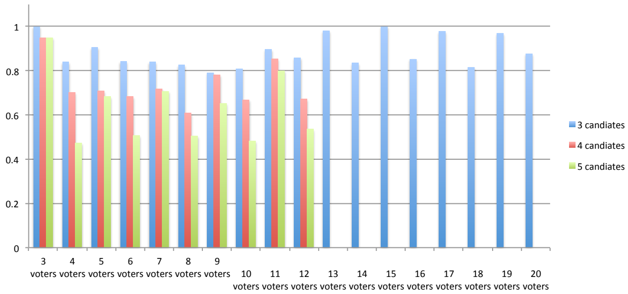

(a) Percentage of equilibria that elect the truthful winner.

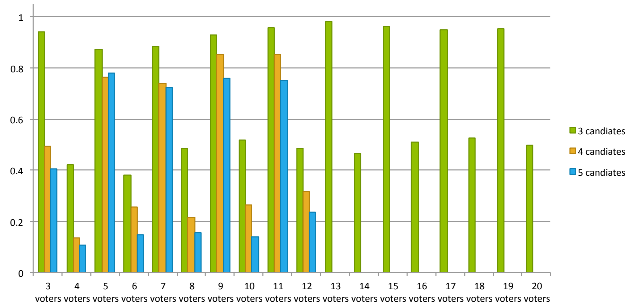

(b) Percentage of equilibria that elect the Condorcet winner

Figure 7: Varying the number of voters and candidates.

is never below 40%) and in Figure 7b for Condorcet winners. These effects were even more pronounced with an odd number of voters, as the number of equilibria was so small.

4.4 Richer Mechanisms

Approval (and variants such as k-approval and veto) are straightforward extensions of plurality, so it is natural to consider whether our approach will work similarly well for these mechanisms. Also, Laslier and Dutta [7] were able to resolve the existence problem for plurality and approval, but noted that the existence of PSNEs for k-approval was an open problem.

Thus, we started by investigating k-approval and veto. As was the case with plurality, the truthfulness incentive kept the number of equilibria manageable. As can be seen from Figure 8, in all cases more than 75% of games had 35 equilibria or fewer, and it seems that the number of equilibria roughly increases with the number of candidates. Our data also allowed us to resolve the open problem of equilibrium existence: for every value of k and m there was at least one instance without any PSNE.

We also considered approval voting. Laslier and Dutta had already shown that a Condorcet consistent equilibria is always guaranteed to exist. However, this raises an interesting question: are there other Nash equilibria? In our experiments, we found that approval voting had an extremely large number of equilibria (over 200,000 per game), so it seems that the addition of another dimension (allowing each voter to decide how many candidates to vote for), results in a large enough flexibility, reducing the effect of the truthfulness incentive.

Looking at the equilibria, we found that it maintains some of the qualities that we discussed: truthful winners won, on average, in over 30% of the equilibria in every setting, and sometimes more (e.g., 2-approval with 4 candidates resulted

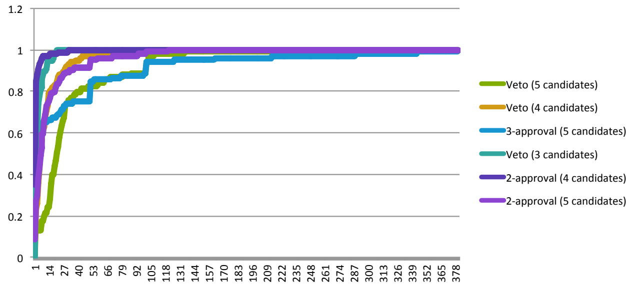

Figure 8: Proportion of games with certain number of equilibria (accumulative).

in almost 50% of equilibria, on average, electing truthful winners). However, Condorcet winners did not show similar strength, and it seems that as the number of candidates grows, and as the number of candidates for which voters allot points to increases, the percent of equilibria with a Condorcet winner drops (so 2-approval has fairly high percentages, 3-approval somewhat less, and 4-approval even less). It will be interesting, in our view, to see how other voting rules compare to plurality in this case, as it seems to trump k-acceptance in this regard.

5. BAYES-NASH EQUILIBRIA RESULTS

Moving beyond the full-information assumption, we considered plurality votes where the agents have incomplete information about each other's preferences. In particular, we assumed that the agents have independent, identically distributed (but not necessarily uniformly distributed) preferences, and that each agent knows only his own preferences and the commonly-known prior distribution. Again, we considered the case of 10 voters and 5 candidates, but now also introduced 6 possible types for each voter. For each of 50 games, we computed the set of all symmetric pure-strategy Bayes-Nash equilibria, both with and without the ϵ-truthfulness incentive.[1]

Our first concern was studying how many equilibria each game had and how the truthfulness incentive affected the number of equilibria. The set of equilibria was small (< 28 in every game) when the truthfulness incentive was present. Surprisingly, only a few equilibria were added when the incentive was relaxed. In fact, in the majority of games (76%), there were exactly five new equilibria: one for each strategy profile where all types vote for a single candidate. Looking into the structure of these equilibria, we found two interesting, and seemingly contradictory, properties. First, most equilibria (95%) only involved two or three candidates (i.e., voters only voted for a limited set of candidates). Second, every candidate was involved in some equilibrium. Thus, we can identify an equilibrium by the number of candidates it involves (see Figure 9). Notably, most equilibria involved only two candidates, with each type voting for their most preferred candidate of the pair. Further, most games had 10 such equilibria, one for every possible pair. There were

[1]We omitted two games from our results. The omitted games each have a type with very low probability. For some profiles, the probability of agents with these types being pivotal was less than machine-ϵ. This led to SEM finding "Bayes-Nash equilibria" that were actually only ϵ-Bayes-Nash.

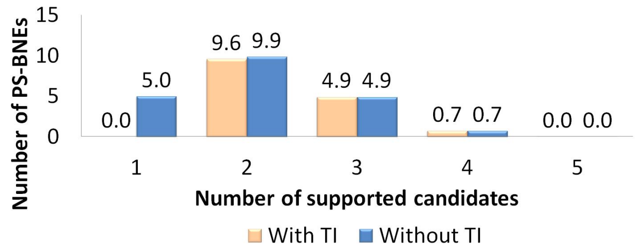

Figure 9: Every instance had many equilibria, most of which only involved a few candidates.

two reasons why some pairs of candidates did not have corresponding equilibria in some games. First, sometimes one candidate Pareto-dominated the other (i.e., was preferred by every type). Second, sometimes the types that liked one candidate were so unlikely to be sampled that close races occurred with extremely low probability (relative to ϵ); in such cases, agents preferred to be deterministically truthful than pivotal with very small probability. This observation allowed us to derive a theoretical result about when a 2-candidate equilibrium will exist.

Theorem 3. *In any symmetric Bayesian plurality election game (with $n \geq 2$), for any pair of candidates c_1, c_2, one of the following conditions is true:*

- *With some positive ϵ truthfulness incentive, there exists a Bayes-Nash equilibrium where each voter votes for his most preferred of c_1, c_2.*
- *One of the candidates Pareto dominates the other ex ante (i.e., the probability that a voter prefers the second candidate to the first is zero).*

We provide only a proof sketch due to space constraints.

Proof Sketch. If c_1 Pareto dominates c_2, then every agent would vote for c_1 and no agent could influence the outcome. For any non-zero ϵ, voters would deviate to honest voting instead of c_1.

So as long as there is non-zero probability of some agent preferring c_2 to c_1, every agent has a non-zero probability of being pivotal between those two outcomes (and a zero probability of being pivotal between any two other outcomes). For a sufficiently small ϵ, the value of influencing the outcome overwhelms the value of truthful voting. $\square$

These two-candidate equilibria have some interesting properties. Because they can include any two candidates where one does not Pareto-dominate the other, they can exist even when a third candidate Pareto-dominates both. Thus, it is possible for two-candidate equilibria to fail to elect a Condorcet winner. However, because every two-candidate equilibrium is effectively a pairwise runoff, it is impossible for a two-candidate equilibrium to elect a Condorcet loser.

Equilibria supporting three or more candidates are less straightforward. Which 3-candidate combinations are possible in equilibrium (even without ϵ-truthful incentives) can depend on the specific type distribution and the agents' particular utilities. Also, in these equilibria, agents do not always vote for their most preferred of the three alternatives (again, depending on relative probabilities and utilities). Finally, 3-candidate equilibria can elect a Condorcet loser with non-zero probability.

6. DISCUSSION AND FUTURE WORK

Our work approaches issues of voting manipulation by combining two less-common approaches: assuming all voters are manipulators, rather than a subset with a shared goal, and looking at the set of Nash equilibria as a whole, rather than searching for other solution concepts or a specific equilibrium. We leveraged a small and realistic assumption—that users attach a small value to voting their truthful preferences. Using the AGG framework to analyze the Nash equilibria and symmetric Bayes-Nash equilibria of plurality, we can extrapolate from the data and reveal properties of such voting games.

We saw several interesting results, beyond a reduction in the number of equilibria due to our truthfulness incentive. One of the most significant was the "clustering" of many equilibria around candidates, which can be viewed as resembling the voters' intention. A very large share of each game's equilibria resulted in winners that were either truthful winners (according to plurality) or Condorcet winners. Truthful winners were selected in a larger fraction of equilibria when the total number of equilibria was fairly small (as was the case in a large majority of our experiments), and their share decreased as the number of equilibria increased (where we saw, in cases where there were Condorcet winners, that those equilibria took a fairly large share of the total). Furthermore, these results held up even when varying the number of candidates and voters, and many of them appear to also hold with other voting systems, such as veto and k-approval.

Looking at social welfare enabled us to compare equilibrium outcomes to all other possible outcomes. We observed that plurality achieved nearly the best social welfare possible (a result that did not rely on our truthfulness incentive). While another metric showed the same "clustering" we noted above, most equilibrium results concentrated around candidates that were ranked, on average, very highly (on average, more than 50% of winners in every experiment had a rank less than 1). This suggests that one should question whether it is even important to minimize the amount of manipulation, as we found that manipulation by all voters very often results in socially beneficial results.

In the Bayes-Nash results, we saw that lack of information often pushed equilibria to be a "battle" between a subset of the candidates—usually two candidates (as Duverger's law would indicate), but occasionally more.

There is much more work to be done in the vein we introduced in this paper. This includes examining the effects of changing utility functions, as well as looking at more voting rules and determining properties of their equilibria. Voting rules can be ranked according to their level of clustering, how good, socially, their truthful results are, and similar criteria. Furthermore, it would be worthwhile to examine other distributions of preferences and preference rules, such as single-peaked preferences. Computational tools can also be used to assess effectiveness of various strategies available to candidates (e.g., it might be more productive for a candidate to attack another weak candidate to alter the distribution).

7. ACKNOWLEDGMENTS

The authors would like to thank the iAGT semester at the Hebrew University, where this work was initiated, and also Reshef Meir for helpful discussions on the topic of this work. This work was partially supported by Israel Science Foundation grant #1227/12, Israel Ministry of Science and Technology grant #3-6797, the Google Inter-University Center for Electronic Markets and Auctions, and Intel Collaborative Research Institute for Computational Intelligence (ICRI-CI).

8. REFERENCES

[1] J. J. Bartholdi III and J. B. Orlin. Single transferable vote resists strategic voting. *Social Choice & Welfare*, 8:341–354, 1991.

[2] J. J. Bartholdi III, C. A. Tovey, and M. A. Trick. The computational difficulty of manipulating an election. *Social Choice & Welfare*, 6(3):227–241, 1989.

[3] C. Daskalakis, G. Schoenebeck, G. Valiant, and P. Valiant. On the complexity of Nash equilibria of action-graph games. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2009.

[4] F. De Sinopoli. On the generic finiteness of equilibrium outcomes in plurality games. *Games and Economic Behavior*, 34(2):270–286, February 2001.

[5] Y. Desmedt and E. Elkind. Equilibria of plurality voting with abstentions. In *Proceedings of the 11th ACM conference on Electronic commerce (EC '10)*, pages 347–356, Cambridge, Massachusetts, June 2010.

[6] A. Dhillon and B. Lockwood. When are plurality rule voting games dominance-solvable? *Games and Economic Behavior*, 46(1):55–75, 2004.

[7] B. Dutta and J.-F. Laslier. Costless honesty in voting. in 10th International Meeting of the Society for Social Choice and Welfare, Moscow, 2010.

[8] B. Dutta and A. Sen. Nash implementation with partially honest individuals. *Games and Economic Behavior*, 74(1):154–169, 2012.

[9] M. Duverger. *Political Parties: Their Organization and Activity in the Modern State*. Methuen Publishing, 1959.

[10] B. Escoffier, J. Lang, and M. Öztürk. Single-peaked consistency and its complexity. In *Proceedings of the 2008 conference on ECAI 2008: 18th European Conference on Artificial Intelligence*, pages 366–370, Amsterdam, The Netherlands, 2008.

[11] T. Feddersen and W. Pesendorfer. Voting behavior and information aggregation in elections with private information. *Econometrica*, 65(5):1029–1058, September 1997.

[12] A. Gibbard. Manipulation of voting schemes. *Econometrica*, 41(4):587–602, July 1973.

[13] A. Gibbard. Manipulation of schemes that mix voting with chance. *Econometrica*, 45(3):665–681, April 1977.

[14] A. X. Jiang and K. Leyton-Brown. Bayesian Action-Graph Games. In *Proceedings of the 24th Conference on Neural Information Processing Systems (NIPS)*, 2010.

[15] A. X. Jiang, K. Leyton-Brown, and N. A. Bhat. Action-graph games. *Games and Economic Behavior*, 71:141–173, 2011.

[16] K. Kawai and Y. Watanabe. Inferring strategic voting. Technical report, Northwestern, 2010.

[17] M. Kearns, M. Littman, and S. Singh. Graphical models for game theory. In *Proceedings of the 17th Conference in Uncertainty in Artificial Intelligence (UAI)*, 2001.

[18] J.-F. Laslier and J. W. Weibull. A strategy-proof condorcet jury theorem. forthcoming, Scandinavian Journal of Economics, 2012.

[19] O. Lev and J. S. Rosenschein. Convergence of iterative voting. In *Proceedings of the 11th International Coference on Autonomous Agents and Multiagent Systems (AAMAS)*, volume 2, pages 611–618, Valencia, Spain, June 2012.

[20] R. D. McKelvey and R. E. Wendell. Voting equilibria in multidimensional choice spaces. *Mathematics of Operations Research*, 1(2):144–158, May 1976.

[21] R. Meir, M. Polukarov, J. S. Rosenschein, and N. R. Jennings. Convergence to equilibria of plurality voting. In *Proceedings of the 24th National Conference on Artificial Intelligence (AAAI)*, pages 823–828, Atlanta, July 2010.

[22] M. Messner and M. K. Polborn. Miscounts, Duverger's law and Duverger's hypothesis. Technical Report 380, IGIER (Innocenzo Gasparini Institute for Economic Research), Bocconi University, 2011.

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

[24] R. W. Porter, E. Nudelman, and Y. Shoham. Simple search methods for finding a Nash equilibrium. *Games and Economic Behavior*, 63:642–662, 2009.

[25] A. D. Procaccia. Can approximation circumvent Gibbard-Satterthwaite? In *Proceedings of the 24th National Conference on Artificial Intelligence (AAAI)*, pages 836–841, Atlanta, Georgia, July 2010.

[26] R. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *Int. J. Game Theory*, 2:65–67, 1973.

[27] M. A. Satterthwaite. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10(2):187–217, April 1975.

[28] D. R. M. Thompson, S. Leung, and K. Leyton-Brown. Computing nash equilibria of action-graph games via support enumeration. In *Proceedings of the 7th International Conference on Internet and Network Economics (WINE)*, 2011.

[29] L. Xia and V. Conitzer. Stackelberg voting games: Computational aspects and paradoxes. In *Proceedings of the 24th National Conference on Artificial Intelligence (AAAI)*, pages 805–810, Atlanta, Georgia, USA, 2010.

[30] L. Xia, M. Zuckerman, A. D. Procaccia, V. Conitzer, and J. S. Rosenschein. Complexity of unweighted coalitional manipulation under some common voting rules. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI)*, pages 348–353, Pasadena, California, USA, July 2009.