

Empirical Aspects of Plurality Election Equilibria

David R. M. Thompson, Omer Lev, Kevin Leyton-Brown
and Jeffrey S. Rosenschein

Abstract

Social choice functions aggregate the different preferences of agents, choosing from a set of alternatives. 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. We do this through a computational analysis, leveraging recent advances in (Bayes-)Nash equilibrium computation for large games. Although plurality has exponentially many pure-strategy Nash equilibria, we demonstrate how a simple equilibrium refinement—assuming that agents very weakly prefer to vote truthfully—dramatically reduces this set. We also use symmetric Bayes-Nash equilibria to investigate the case where voters are uncertain of each others’ preferences. Although our refinement does not completely eliminate the problem of multiple equilibria, it tends to predict an increased probability that a good candidate will be selected (e.g., the candidate that would win if voters were truthful, or a Condorcet winner).

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 is strategy-proofness, i.e., that it is 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.

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

Dutta and Laslier [7] studied a somewhat similar model, where voters have a lexicographic preference for truthfulness. They demonstrated that for some voting mechanism, a small preference for truthfulness can eliminate all pure-strategy Nash equilibria. We observed a similar occurrence in our results with plurality (which is problematic voting methods designed to reach an equilibrium by an iterative process, e.g., [21; 19]).

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 (10 players and 5 candidates) in the normal form would require about a hundred million payoffs. 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 how many Nash equilibria 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 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. Dutta and Laslier [7] attempted to inject truthfulness directly into a voting rule combined of approval voting and veto, but only found a few existence results that show truthful equilibria exist in that case. 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.

2 Definitions

Before detailing our specific scenario, we first define elections, and how winners are determined.

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 which he announces as his value, which we shall denote as \tilde{a}_i .*

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

Definition 2. *A voting rule is a function $f : A^n \rightarrow 2^C \setminus \emptyset$.*

In this paper, we restrict our attention to plurality, where a point is given to each voter’s most-preferred candidate, and the candidates with the highest score win.

Our definition of voting rules allows for multiple winners. However, in many cases what is desired is a single winner; in these cases, a tie-breaking rule is required.

Definition 3. A tie-breaking rule is a function $t : 2^C \rightarrow C$ that, given a set of elements in C , chooses one of them as a (unique) winner.

There can be many types of tie-breaking rules, such as random or deterministic, lexical or arbitrary. In this work, we use a lexical tie-breaking rule.

Another important concept is that of a Condorcet winner.

Definition 4. 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 $\lceil \frac{n}{2} \rceil$.

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). Note that our definition allows for the possibility of multiple Condorcet winners in a single election, in cases where n is even. Conversely, a Condorcet loser is ranked below any other candidate by a majority of voters.

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 (either because they are behaving randomly, or because they have committed to a strategy that agent i does not observe). In this paper, 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 strictly prefer to abstain when his vote is not pivotal, as in [5], or to vote 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 $u_i(t(f(a_V)), \tilde{a}_i)$.

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 i.i.d. 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 10 voters and 5 candidates (numbers chosen to give a setting both computable and with a range of candidates), and with the voters’ preferences chosen randomly. Suppose voter i has a preference order of $a^5 > a^4 > \dots > a^1$, and the winner when voters voted a_V is a^j . We then define i ’s utility function as

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

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

In the incomplete-information case, we model agents as having one of six possible types (to make the problem more easily computable), each corresponding to a different (randomly selected) preference ordering. The agent’s type draws are i.i.d. but the probability of each type is not necessarily uniform. Instead, the probability of each type is drawn from a uniform distribution, and then normalized; thus, the probabilities ranged from 0.0002 to 0.55.




Figure 1: An action graph game encoding of a simple two-candidate plurality vote. Each round node represents an action that 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 our voting 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).

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 of a game quickly.


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 1,000 voting experiments using plurality with 10 voters and 5 candidates. Such a game might ordinarily have hundreds of thousands of Nash equilibria. However, adding a small truthfulness incentive ($\epsilon = 10^{-6}$) lowers these numbers significantly. Not counting permutations of voters with the same preferences, every game had 25 or fewer equilibria; counting permutations, the maximum number of equilibria was still only 146. Indeed, an overwhelming number of these games (96.2%) had fewer than 10 equilibria (27 with permutations). More surprisingly, a few (1.1%) had no pure Nash equilibria at all.¹ To gauge the impact of the truthfulness incentive, we also ran 50 experiments without it; every one of these games had over a hundred thousand equilibria, without even considering permutations.



(a) CDF showing the fraction of games having the given number of pure-strategy Nash equilibria or fewer (including permutations)



(b) The social welfare distribution



(c) The number of truthful and Condorcet winning equilibria, depending on total number of equilibria per experiment. Note that in the “tail”, the data is based on only a few experiments.

Figure 2: Equilibria and social welfare in Plurality

We shall examine two aspects of the results: the preponderance of equilibria with victors being the voting method’s winners,² and Condorcet winners. Then, moving to the wider concept of social welfare of the equilibria (possible due to the existence of utility functions), we examine both the social welfare of the truthful voting rule vs. best and worse possible Nash equilibria, as well as the average rank of the winners in the various equilibria.

¹This is especially relevant to voting procedures relying on the existence of pure Nash equilibrium, and seeking to “find” one, such as the one proposed in [21].

²This, when expanded to more voting rules, may be an interesting comparative criterion between voting mechanisms.

For 63.3% of the games, the truthful preferences were a Nash equilibrium, but more interestingly, many of the Nash equilibria reached, in fact, the same result as the truthful preferences: 80.4% of the games had at least one equilibrium with the truthful result, and looking at the multitudes of equilibria, the average share of truthful equilibrium (i.e., result was the same as with truthful vote) was 41.56% (out of games with a truthful result as an equilibrium, the share was 51.69%). Without the truthfulness incentive, the average share of truthful equilibrium was 21.77%.

Looking at Condorcet winners, 92.3% of games had Condorcet winners, but they were truthful winners only in 44.7% of the games (not a surprising result, as plurality is far from being Condorcet consistent). However, out of all the equilibria, the average share of equilibria with a victorious Condorcet winner was 40.14% (of games which had a Condorcet winner the average share is 43.49%; when the Condorcet winner was also the truthful winner, its average share of equilibria is 56.96%).

Looking at the wider picture (see Figure 2c), the addition of the truthful incentive made possible games with very few Nash equilibria. They, very often, resulted in the truthful winner. As the number of equilibria grows, the truthful winner part becomes smaller, as the Condorcet winner part increases.

Turning to look at the social welfare of equilibria, once again, the existence of the truthfulness incentive enables us to reach “better” equilibria. In 92.8% 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 29.7% of cases, the best outcome was not possible. We note that 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 2b).

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 56.38%. Even if we exclude Condorcet winners (as they, on many occasions, are highly ranked), the average ranking of less than 1 was 46.56% (excluding truthful winners resulted in 27.48% with average ranking less than 1). Fully 71.65%, on average, of the winners in every experiment had above (or equal) the median rank, and in more than half the experiments (52.3%) 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).

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 i.i.d. (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.

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

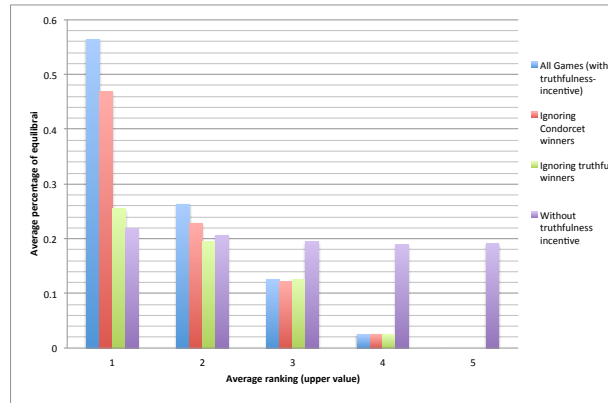


Figure 3: The average proportion of equilibria won by candidates with average rank of 0–1, 1–2, etc.

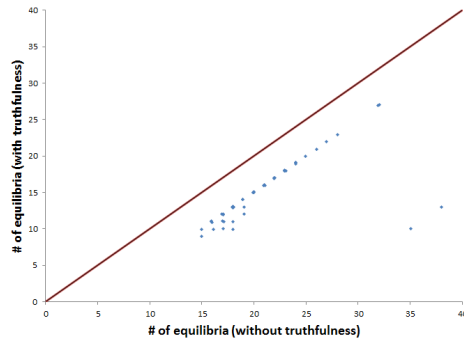


Figure 4: The number of symmetric pure-strategy Bayes-Nash equilibria in plurality votes with and without the ϵ -truthfulness incentive

were exactly five new equilibria: one for each strategy profile where all types vote for a single candidate (see Figure 4).

Looking into the structure of these equilibria, we found two interesting, and seemingly contradictory, properties: most equilibria (95.2%) only involved two or three candidates (i.e., voters only voted for a limited set of candidates), but every candidate was involved in some equilibrium. Thus, we can identify an equilibrium by the number of candidates it involves (see Figure 5). 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 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 were extremely low probability (relative to ϵ); in such cases, agents preferred to be deterministically truthful than pivotal with very small probability.³ This observation allowed us to derive a theoretical

³There were two outlier games where one of the types had a very low probability (< 0.001). Because of this, the probability of a realization where half the agents had this type approached machine- ϵ . Thus, any pure strategy profile where this type votes one way and all the other types vote another way will result in a 2-candidate equilibrium (20 such 2-candidate combinations exist, so these games had 20 additional two-candidate equilibria.)

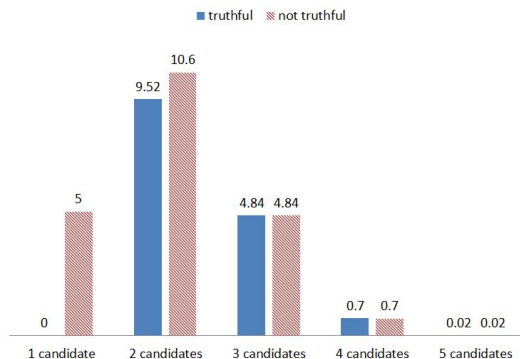


Figure 5: Every instance had many equilibria, most of which only involved a few candidates.

result about when a 2-candidate equilibrium will exist.

Let ℓ be the minimal difference between the utility of 2 different candidates, across all voters (in our scenarios, this minimal difference is 1).

Proposition 5. *In a plurality election with a truthfulness incentive of ϵ , as long as $(\frac{1}{n})^{\lfloor \frac{n}{2} \rfloor} \ell \geq \epsilon$, for every $c_1, c_2 \in C$, either c_1 Pareto dominates c_2 (i.e., all voters rank c_1 higher than c_2), or there exists a pure Bayes-Nash equilibrium in which each voter votes for his most preferred among these two candidates.*

Due to space constraints, we provide only proof sketch.

Proof sketch. Let us define a strategy as follows: every voter that prefers c_1 over c_2 votes for c_1 ; otherwise, he votes for c_2 . Obviously, if c_2 is Pareto dominated, every individual voter believes that he will be better off voting truthfully, and this may not be an equilibrium. However, if c_2 is not Pareto dominated, then there is a probability larger than (or equal to) $\frac{1}{n}$ that there is a voter who prefers c_2 to c_1 . Hence, the probability that a voter who prefers c_1 to c_2 will be pivotal is at least $(\frac{1}{n})^{\lfloor \frac{n}{2} \rfloor}$. If the benefit to all voters from being pivotal in this way is larger than ϵ , the value of the truthfulness incentive, the voter will not deviate from that strategy. Thus, when $(\frac{1}{n})^{\lfloor \frac{n}{2} \rfloor} \ell \geq \epsilon$ they do not deviate. \square

These two-candidate equilibria have some interesting properties. Because they can include any two candidates that do not Pareto-dominate each other, it is possible for them to exclude a third candidate that Pareto-dominates both. In this way, 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 the issues of voting manipulation by combining two less-common approaches: assuming all voters are manipulators, rather than just a subset with a shared

goal, and looking at Nash equilibria as a whole, rather than searching for other solution concepts or a specific equilibrium. We utilized only 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 that 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).

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 high (on average, more than 50% of winners in every experiment had a rank less than 1). This, in a sense, raises the issue of the rationale of seeking 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 generally 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 have introduced in this paper. This includes examining the effects of varying the number of voters and candidates, 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 other 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 useful to assess the usefulness of various strategies available to candidates (e.g., it might be more productive for a candidate to attack a weak candidate to alter the distribution).

7 Acknowledgments

Authors would like to thank the iAGT semester in the Hebrew University, where this work was initiated, and would also like to thank Reshef Meir for helpful discussions on the topic of this work. This work was partially supported by Israel Ministry of Science and Technology grant #3-6797, and the Google Inter-University Center for Electronic Markets and Auctions.

References

- [1] John J. Bartholdi III and James B. Orlin. Single transferable vote resists strategic voting. *Social Choice and Welfare*, 8:341–354, 1991.
- [2] John J. Bartholdi III, Craig A. Tovey, and Michael A. Trick. The computational difficulty of manipulating an election. *Social Choice and Welfare*, 6(3):227–241, 1989.

- [3] Constantinos Daskalakis, Grant Schoenebeck, Gregory Valiant, and Paul Valiant. On the complexity of Nash equilibria of action-graph games. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2009.
- [4] Francesco De Sinopoli. On the generic finiteness of equilibrium outcomes in plurality games. *Games and Economic Behavior*, 34(2):270–286, February 2001.
- [5] Yvo Desmedt and Edith Elkind. Equilibria of plurality voting with abstentions. In *Proceedings of the 11th ACM conference on Electronic commerce (EC '10)*, pages 347–356, Cambridge, Massachusetts, USA, June 2010.
- [6] Amrita Dhillon and Ben Lockwood. When are plurality rule voting games dominance-solvable? *Games and Economic Behavior*, 46(1):55–75, January 2004.
- [7] Bhaskar Dutta and Jean-François Laslier. Costless honesty in voting. in 10th International Meeting of the Society for Social Choice and Welfare, Moscow, 2010.
- [8] Bhaskar Dutta and Arunava Sen. Nash implementation with partially honest individuals. *Games and Economic Behavior*, 74(1):154–169, January 2012.
- [9] Maurice Duverger. *Political Parties: Their Organization and Activity in the Modern State*. Methuen Publishing, 1959.
- [10] Bruno Escoffier, Jérôme Lang, and Meltem 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. IOS Press.
- [11] Timothy Feddersen and Wolfgang Pesendorfer. Voting behavior and information aggregation in elections with private information. *Econometrica*, 65(5):1029–1058, September 1997.
- [12] Allan Gibbard. Manipulation of voting schemes. *Econometrica*, 41(4):587–602, July 1973.
- [13] Allan Gibbard. Manipulation of schemes that mix voting with chance. *Econometrica*, 45(3):665–681, April 1977.
- [14] Albert Xin Jiang and Kevin Leyton-Brown. Bayesian Action-Graph Games. In *Proceedings of the 24th Conference on Neural Information Processing Systems (NIPS)*, 2010.
- [15] Albert Xin Jiang, Kevin Leyton-Brown, and Navin A.R. Bhat. Action-graph games. *Games and Economic Behavior*, 71:141–173, 2011.
- [16] Kei Kawai and Yasutora Watanabe. Inferring strategic voting. Technical report, Northwestern, 2010.
- [17] M.J. Kearns, M.L. Littman, and S.P. Singh. Graphical models for game theory. In *Proceedings of the 17th Conference in Uncertainty in Artificial Intelligence (UAI)*, 2001.
- [18] Jean-François Laslier and Jörgen W. Weibull. A strategy-proof condorcet jury theorem. forthcoming, Scandinavian Journal of Economics, 2012.
- [19] Omer Lev and Jeffrey S. Rosenschein. Convergence of iterative voting. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, Valencia, Spain, June 2012.
- [20] Richard D. McKelvey and Richard E. Wendell. Voting equilibria in multidimensional choice spaces. *Mathematics of Operations Research*, 1(2):144–158, May 1976.
- [21] Reshef Meir, Maria Polukarov, Jeffrey S. Rosenschein, and Nicholas 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] Matthias Messner and Mattias K. Polborn. Miscounts, duverger’s law and duverger’s hypothesis. Technical Report 380, IGIER (Innocenzo Gasparini Institute for Economic Research), Bocconi University, 2011.
- [23] Roger B. Myerson and Robert 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] Ariel D. Procaccia. Can approximation circumvent Gibbard-Satterthwaite? In *Proceedings of the 24th National Conference on Artificial Intelligence (AAAI)*, pages 836–841, Atlanta, Georgia, USA, July 2010.
- [26] R. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *Int. J. Game Theory*, 2:65–67, 1973.
- [27] Mark Allen 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] David R. M. Thompson, Samantha Leung, and Kevin 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] Lirong Xia and Vincent 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] Lirong Xia, Michael Zuckerman, Ariel D. Procaccia, Vincent Conitzer, and Jeffrey 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.

David R. M. Thompson
 Department of Computer Science
 University of British Columbia
 Vancouver, B.C., Canada, V6T 1Z4
 Email: daveth@cs.ubc.edu

Omer Lev
 School of Computer Science and Engineering
 The Hebrew University of Jerusalem
 Jerusalem 91904, Israel
 Email: omer1@cs.huji.ac.il

Kevin Leyton-Brown
 Department of Computer Science
 University of British Columbia
 Vancouver, B.C., Canada, V6T 1Z4
 Email: kevinlb@cs.ubc.ca

Jeffery S. Rosenchein
 School of Computer Science and Engineering
 The Hebrew University of Jerusalem
 Jerusalem 91904, Israel
 Email: jeff@cs.huji.ac.il