Robustness to Manipulation in Voting Theory

Charlie Hechler, Vraj Patel, Arya Thirodira, Firmiana Wang
Graduate Student Leader: Robert Krueger

July 30, 2022

Abstract

A voting system is a way of aggregating individuals' preferences into a collective decision. The first voting system you may think of is majority rule: if there are two options, we simply choose the option that most voters prefer. When there are more than two options, the voting systems can be of varying complexity, with no clear obvious 'best' system. We study one particular metric of voting systems: on average, what is the fewest number of votes needed for an adversary to change in order to change the result of the election? When there are two candidates, we prove which voting systems are the best under this metric. When there are more than two candidates, we offer some simulation techniques and compare how manipulable some common voting systems are.

1 Introduction

Voting theory is a mathematical investigation into how a group of people make collective decisions. We call the choices *candidates*, and we call the people *voters*. Perhaps the simplest example we are all familiar with is when there are exactly two candidates, and we decide according to 'majority rule': the candidate which the most voters prefer wins. However, this is not the only voting system, even for two candidates — for example, depending on the situation, 'unanimous consent' may be preferred, where we collectively pick candidate A if and only if everyone votes for A.

When there are more than three candidates, voting systems become even more varied. We first need to discuss how the voters vote; in other words, what are the ballots voters can cast? A common model is that every voter ranks the candidates, and we wish to pick one winning candidate. Unfortunately, it is possible that a majority of voters prefer A to B, a majority prefers B to C, and a majority prefers C to A. How do we pick a winner in this case? There are many different voting systems with different nice properties that try to answer this question. Overall, the answer is not very optimistic, because of impossibility theorems. The most famous one is Arrow's Theorem [1], stated in Heilman [3], stating that there are no voting systems with some nice natural properties. In view of this, we should fix ahead of time which criteria we use to judge our voting system by.

We look at the 'manipulability' of various voting systems. When there are two candidates, we are able to prove that certain voting systems are the 'least manipulable.' For three or more candidates, we run some simulations comparing some common voting systems. Since our models are different for these situations, we give the precise definitions below.

1.1 Two Candidates

A *ballot* is a voter's choice of candidate. An *election* is a selection of ballot for each voter. A *voting system* is a function which takes in an election and assigns a winning candidate. We call a voting system *balanced* if there is an equal chance that either candidate wins, if an election is chosen uniformly at random. This means that the number of possible elections where one candidate wins is equal to the number of possible elections where the other candidate wins. The *majority* system is the voting system where the winner is the one candidate which receives the most votes; this system exists only when the number of voters is odd, since we do not allow the case of a 'tie.' The *leave-one-out majority* system is a voting system with an even number of voters, where we remove one voter and then use the majority system. This is done to avoid ties. These systems are balanced voting systems since there is an equal chance of either candidate winning. Another example of a balanced voting system is a *dictatorship*, where one voter completely determines the winner.

We call an election *t-manipulable* if changing at most t ballots can result in a different winner. A balanced voting systems that maximizes the number of 1-manipulable elections is the dictatorship, since dictatorships only require a change in 1 vote to change the result.

Heilman [3] provides proof for determining the most robust voting method for an odd number of voters.

Theorem 1. *[Heilman [3]] Majority minimizes the number of elections for which you can change t votes to change the result among all balanced voting systems, when the number of voters is odd.*

We extend his theorem to an even number of voters.

Theorem 2. *Leave-one-out majority minimizes the number of elections for which you can change t votes to change the result among all balanced voting systems, when the number of voters is even.*

Our proof of Theorem 2 follows Heilman's proof, except that we use a more detailed version of Harper's vertex-isoperimetry theorem as proven by Räty [4]. See Section 2.3 for remarks on some open problems in this direction.

1.2 Three Candidates

When there are more than two candidates, it is much harder to analyze the manipulability of all voting systems abstractly. Instead, we focus on a few classes of voting systems and compare the manipulability of those voting systems. We also only consider the case of three candidates to simplify our analysis.

The λ-*Borda count* voting system takes in a ranked ballot and assigns points to each candidate based on its position. The candidate with the most points is then elected as the winner. For three candidates, we assign the points to be 1, λ, and 0, where λ varies between 0 and 1. By changing the value of λ, this λ-Borda count system can be extended to other voting methods. When $\lambda = 0$, the winner should be consistent with that of a plurality system where the winner has the most votes. The $\lambda = 1/2$ case is also special, as this is the voting system originally devised by Borda. The $\lambda = 1$ case is equivalent to everyone voting for their least-favorite candidate, and then the candidate with the least least-favorite votes wins.

The λ-*Borda elimination* is similar to the Borda count systems in that it takes in a ranked ballot and assigns points to each candidate based on its position. The candidate with the least points is eliminated, and the scores of the remaining candidates are recalculated and eliminated until one candidate remains, which is then named the winner. For three candidates, we assign the points to be 1, λ, and 0, where λ varies between 0 and 1. The $\lambda = 0$ case is known as instant run-off voting or ranked choice voting. The $\lambda = 1$ case is known as survivor or Coombs voting, where effectively voters vote for their least-favorite to be eliminated.

An *election* is defined by every voter selecting a ballot. In the voting systems we consider, all voters are equal and indistinguishable. We also follow the *impartial culture assumption*, where we assume every voter chooses a ballot independently and uniformly at random from all possible ballots.

The ε-*manipulability* of a voting system is defined as the proportion of elections in which we can change at most ϵ-proportion of each type of ballot to change the outcome of the election. For small ε, the manipulability of the λ-Borda count system increases as λ increases from 0 to 1. In contrast, the manipulability of the λ-Borda elimination system approaches its minimum at about $\lambda = 0.6$. Of all these voting systems, the 0-Borda count system — plurality — is the least manipulable. See Section 3.3 for comments on future work and open problems in this direction.

2 Two Candidates

To prove Theorem 2, we make an analogy between manipulability in the voting theory of two candidate elections and the isoperimetric inequalities in the hypercube.

2.1 Harper's Theorem

A *hypercube* is an n-dimensional version of a square/cube: the vertices are $\{0, 1\}^n$, 0-1 vectors of length n, where two vertices are adjacent if and only if they differ in exactly one coordinate. The *lexicographic order* is an ordering of vertices of a hypercube alphabetically. *Simplicial order* is an ordering of vertices of a hypercube that first orders by the number of zeros/ones, then within those sets of vertices, orders lexicographically. The *boundary* of a subset S of the hypercube are the vertices not in S of the hypercube that neighbor S. The set $\bar{S}$ is the complement of S in the hypercube, meaning the subset of vertices that lead to the other candidate's win. The *neighborhood of a subset, $N(S)$,* is all of the vertices which are adjacent to a vertex in S. The boundary of S is thus $N(S) \setminus S$. The t-neighborhood $N^t(S)$ is all vertices at distance at most t from S. As such, the t-boundary is the set of vertices outside of S that are distance at most t away from S.

Harper's theorem [2] and a strengthening due to Räty [4] are useful for us.

Theorem 3 (Harper [2])**.** *For every ℓ, a subset S of size ℓ of the hypercube of minimum boundary is given by an initial segment of simplicial order.*

One point to notice is that the final segment of initial order is also an intital segment of simplicial order, if we swap the names of the candidates (change 0 to 1 and vice versa). Therefore, if S has minimum boundary, then $\bar{S}$ also has minimum boundary. Räty [4] extends Harper's Theorem to say that the sets with minimum boundary whose complements also have minimum boundary have minimum t-boundary.

Theorem 4 (Räty 4)**.** *Let $A \subseteq Q_n$ be a subset for which every $B \subseteq Q_n$ with $|B| = |A|$ satisfies $|N(B)| \geq |N(A)|$ and $|N(\bar{B})| \geq |N(\bar{A})|$. Then for all $t > 0$ and $B \subseteq Q_n$ with $|B| = |A|$ we have $|N^t(B)| \geq |N^t(A)|$ and $|N^t(\bar{B})| \geq |N^t(\bar{A})|$.*

For an example of simplicial order, let us look at $n = 3$ voters with 2 candidates, named 0 and 1. There will be $2^3 = 8$ possible elections, which correspond to 8 vertices of the hypercube. We list them below in simplicial order:

- 111
- $011, 101, 110$
- $001, 010, 100$
- 000

We begin with the vertices that have zero 0's: there is only 1 vertex. We move to vertices that have one 0: there are 3 vertices as such, and we order these alphabetically (0 comes before 1 when ordering alphabetically). We then move to vertices with two 0's and three 0's. This is our simplicial order of the vertices in the hypercube for 3 voters.

We can use this order in our voting theory context. For 2-candidate voting systems with n voters, the elections are in one-to-one correspondence with the vertices of the n-dimensional hypercube. A voting system is equivalent to a subset of the hypercube given by the elections where candidate 1 wins. There are 2^n possible elections. For the system to be balanced, it must be that $\frac{1}{2}$ of the elections must result in one candidate's victory, while the other $\frac{1}{2}$ must result in the other candidate's victory. The ℓ from Harper's Theorem represents this balance, where we set $\ell = \frac{1}{2} \cdot (2^n)$. Let S be a subset of the vertices of the hypercube of size ℓ, consisting of all elections where candidate 1 wins. The vertices of S which are adjacent to a vertex outside of S are precisely the elections which the result can be changed from 1 to 0 by altering a single vote. Likewise, the vertices of $\bar{S}$ which are adjacent to a vertex outside of $\bar{S}$ are precisely the elections which the result can be changed from 0 to 1 by altering a single vote. The total 1-manipulability of a voting system is the sum of the sizes of these two groups of vertices, divided by 2^n, the total number of elections.

In the example above, Harper's Theorem gives that voting system of minimum 1-manipulability is where the first 4 elections above will result in candidate 1 winning. The other 4 elections result in candidate 0 winning by majority. This is the majority system.

Let us now look at an even number of voters, for example $n = 4$. There will be $2^4 = 16$ possible elections, corresponding to 16 vertices of the hypercube. We list them below in simplicial order:

- 1111
- $0111, 1011, 1101, 1110$
- $0011, 0101, 0110, 1001, 1010, 1100$
- $0001, 0010, 0100, 1000$
- 0000

We find that $\ell = \frac{1}{2} \cdot 2^4 = 8$. Thus, the first 8 vertices in the simplicial list result in one candidates' win. Taking these 8, we see that the first 5 vertices result in candidate 1 winning by majority. Yet, in the latter 3 — vertices $0011, 0101$, and 0110 — there exists a tie between the two candidates. However, we know that candidate 1 must win for these elections in order to keep the voting method balanced by Harper's Theorem. Upon further analysis, it appears that the first voter in these three elections votes the same. If we disregard this voter's vote, we find that candidate 1 wins the three elections by majority. We can make a similar analysis for candidate 0's first three elections. In fact, removing the first voter's vote in all elections and then applying majority rule gives this voting method. This is an example of the leave-one-out majority method to determine the winner of an election with an even number of voters.

2.2 Proof of Theorem 2

We prove the Theorem for an even number n of voters, since the odd case was already shown in [3].

Let $g : \{0,1\}^n \to \{0,1\}$ be a leave-one-out majority function with an even number n of voters, where we remove voter 1's vote. Let f be a balanced voting method $f : \{0,1\}^n \to \{0,1\}$. Let $S = \{v \in \{0,1\}^n : f(v) = 0\}$ and $B = \{v \in \{0,1\}^n : g(v) = 0\}$. Here S refers to the set of vertices such that candidate 0 wins with the voting system f, and B refers to the set of vertices such that candidate 0 wins with the leave-one-out majority voting system g.

Recall that $|(N^t(S) \setminus S) \cup (N^t(\bar{S}) \setminus \bar{S})|$ is the number of t-manipulable elections for voting system f. We want to show that

$$|(N^t(S) \setminus S) \cup (N^t(\bar{S}) \setminus \bar{S})| \geq |(N^t(B) \setminus B) \cup (N^t(\bar{B}) \setminus \bar{B})|$$

Theorem 4 tells us that $|N^t(S)| \geq |N^t(B)|$ and $|N^t(\bar{S})| \geq |N^t(\bar{B})|$. Thus:

$$
\begin{aligned}
|(N^t(S) \setminus S) \cup (N^t(\bar{S}) \setminus \bar{S})| &= |N^t(S) \setminus S| + |N^t(\bar{S}) \setminus \bar{S}| \\
&\geq |N^t(S)| - |S| + |N^t(\bar{S}) - |\bar{S}| \\
&\geq |N^t(B)| - |B| + |N^t(\bar{B})| - |\bar{B}| \\
&= |N^t(B) \setminus B| + |N^t(\bar{B}) \setminus \bar{B}| \\
&= |(N^t(B) \setminus B) \cup (N^t(\bar{B}) \setminus \bar{B})|
\end{aligned}
$$

Thus the leave-one-out majority voting system minimizes the number of t-manipulable elections in a condition with an even number of voters.

2.3 Remarks on Open Problems

The main result of Räty's paper [4] is to characterize all subsets A of the hypercube of a given size which minimize $N^t(A)$ and $N^t(\bar{A})$. In our context, this gives all balanced voting systems which minimizes the t-manipulability of an election. For an odd number of voters, it turns out that the majority system is the unique minimizer. But for an even number of voters, there are systems other than leave-one-out majority which are minimizers.

Consider the following example with $n = 4$ voters. Given the vertices where candidate 0 is said to win, for the vertices that result in a tie, change either voter 2, 3, or 4's vote to candidate 1, if not already 1. Then, with these new vertices, after determining the manipulable boundary vertices, we will find a minimized number of them, the same number as if we used the leave-one-out method. For example, if we change the tied elections where voter 2 votes 0 into a vote for 1, our tied vertices (where 0 wins) go from (1001, 1010, and 1100) to (1101, 1110, and 1100). Counting the number of vertices not in the new subset but adjacent to the subset, we find 6. We then count the number of vertices in the subset that are adjacent to vertices not in the subset, which adds to another 6. Therefore, we find 12 manipulable elections, the same number of elections as if we used the leave-one-out majority method. The general minimizers are similar.

We can also look at situations with 3 candidates. There will be 3^n possible elections with candidates named 0, 1, and 2. However, this situation is more complex, mainly because the rate of having ties is much higher than with 2 candidate elections. There are possibilities of having 2-way ties as well as 3-way ties. For example, with $n = 5$, there could be candidates 0 and 1 receiving two votes each and candidate 2 receiving one vote — an example of a 2-way tie. With $n = 6$, we could have all three candidates receiving two votes each — a 3-way tie. Once again, we want to minimize the total boundary of the subsets, of which there will be 3 this time. We believe plurality is the voting system that will minimize this boundary of manipulable elections.

3 Three Candidates: Simulations

For the voting systems considered here, a *ballot* is a total ranking of all the candidates. A voting system takes in all the ballots and returns a single winner. For more than two candidates, we make the simplifying assumption that all voters are equal and indistinguishable for our voting systems. This means that we only have to keep track of the proportion of voters that vote for each ballot. Thus an election is defined by a tuple $(x_1, \ldots, x_k)$ where there are k ballots and $x_1 + \cdots + x_k = 1$. For three candidates, there are $3!$ possible orderings of the candidates, so $k = 6$.

We say an election $(x_1, \ldots, x_k)$ is ε-*manipulable* if there exists another election $(y_1, \ldots, y_k)$ with a different winner than $(x_1, \ldots, x_k)$ such that $|x_i - y_i| \leq \varepsilon$ for all $i \in \{1, \ldots, k\}$. The *manipulability* of an election is the smallest ε such that the election is ε-manipulable. The ε-*manipulability* of a voting system is the proportion of elections which are ε-manipulable. We wish to measure the ε-manipulability of the λ-Borda count system and the λ-Borda elimination system.

These voting systems can be visualized in a equilateral triangle, where each vertex of the triangle represents one candidate receiving all the top-place votes, as shown in Figure 1. The perpendicular distance from a point to that opposite side in the triangle is the proportion of top-place votes that candidate received in a particular election. The winner of the election, as decided by the particular voting method, determines the color of the point.

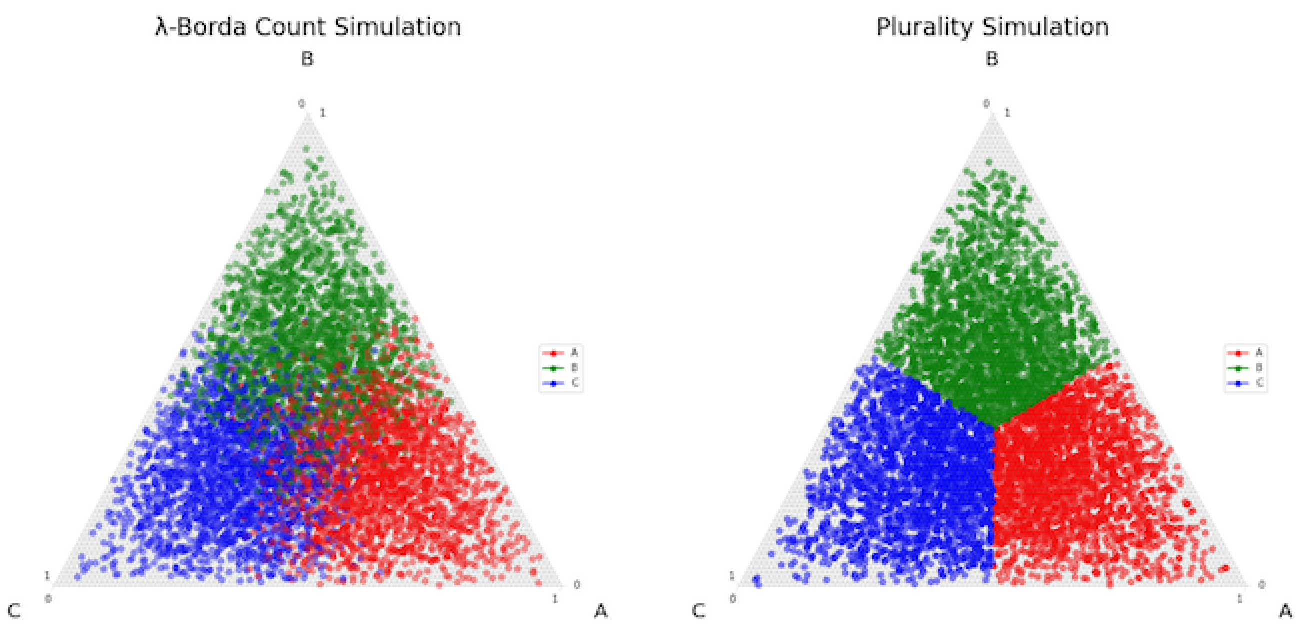

Figure 1: Triangle diagrams for λ-Borda count with $\lambda = 1/2$ (left) and $\lambda = 0$ (right, equivalent to plurality).

3.1 Generalized Borda Count

We compute the manipulability of the λ-Borda count system in two different ways, each qualitatively verifying the other.

3.1.1 Polytope Calculation

The manipulability of a voting system is related to the boundary between candidates in the triangle simulations since easily manipulable elections are naturally closer to the boundaries between candidates. Thus, by finding the "volume" of this boundary, we can determine the manipulability of the voting system. We are able to find this volume by finding the volume of a polytope represented by seven linear inequalities.

A *polytope* is defined as the set of all points in R^n satisfying some linear inequalities. A n-dimensional polytope is enclosed by a finite number of hyperplanes that are represented by these linear inequalities $Ax \leq b$, where A is a $s \times n$ matrix and b is a real s vector, where s is the number of linear inequalities. For example, a hyperplane in two dimensions is a line where A is an $s \times 2$ matrix.

The matrix A can be described using linear inequalities derived from the ballots of the generalized borda count method, where $x_1 + \cdots + x_6 = 1$ and $x_1, \ldots, x_6 \geq 0$. The seventh inequality represents a tie between two candidates where we set score(A) = score(B) > score(C). When we write these seven linear inequalities in terms of four of the variables, they form a polytope with a 7×4 matrix and a vector with 7 elements that correspond to the seven inequalities. For each of the rows in the matrix, the four elements in the row are the coefficients to a set of four variables (e.g., x_1, x_2, x_3, x_4) in a linear inequality that correspond to an element in the vector, which is the constant in the same linear inequality. The four variables inputted into this matrix can be anything in $(x_1, x_2, x_3, x_4, x_5, x_6)$ as the polytope should give the same plot for each set of four variables. The resulting polytope represents a tie, the edge boundary between candidates A and B (which is precisely one third of the total boundary volume representing ties). An example set of these types of inequalities in a matrix with the corresponding vector is shown below, where the columns correspond to x_2, x_4, x_5, and x_6.

$$\begin{bmatrix} 3\lambda & 3\lambda & 3 & 3 \\ 2-\lambda & -\lambda & 1 & 1-2\lambda \\ -1 & 0 & 0 & 0 \\ -\lambda & 2-\lambda & 1-2\lambda & 1 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix} x \leq \begin{bmatrix} 1+\lambda \\ 1-\lambda \\ 0 \\ 1-\lambda \\ 0 \\ 0 \\ 0 \end{bmatrix}$$

While we could compute the volume of the polytope analytically, it is much easier to use a computer simulation. Since the polytope is four-dimensional, we randomly pick points in a four-dimensional cube bounding the polytope, count how many points are in the polytope, and compute the appropriate ratio. This approach does not yield the correct volume, even approximately. The true polytope lives in six dimensions, and we have calculated the volume of a four-dimensional projection of this polytope, so we also calculate a correction factor based on this projection. Multiplying by this correction factor amplifies error in this volume computation (since we sample points uniformly at random). Depending on the projection, that is, which variables we solve for, this error is amplified for λ near 0 or 1. See the blue plot in Figures 2 and 3.

This noise makes the plot difficult to interpret and thus, a more exact version of the plot is needed to extrapolate accurate data. The noise can be reduced by sampling the polytope volume for every value of λ on the plot multiple times, then averaging the result.

As stated before, the polytope matrix and vector can be derived using any four of the seven variables. In the plots shown in Figures 2 and 3, Figure 2 uses x_2, x_4, x_5, and x_6 while Figure 3 uses x_1, x_2, x_3, and x_4. The difference in these variables creates targeted noise at different tails for each of the plots. As shown in Figures 2 and 3, the plots with noise (shown in blue) differ in the two graphs while the noise reduced representations (shown in orange) consistently have the same values in both Figure 2 and Figure 3 (keep in mind that the axes are not the exact same for both Figures 2 and 3).

3.1.2 Random Perturbation Simulation

A simulation is run for a certain number of elections, perturbations, λ, and ε. Each simulation estimates the manipulability of a certain voting system given the previously mentioned parameters. The estimated manipulability is the proportion of manipulable elections to the total number of elections that are simulated.

Each perturbation represents a manipulation to a simulated election by changing at most ε proportion of ballots. We select the perturbation uniformly at random, given ε. We then check if the perturbation has manipulated the election. The greater the number of perturbations the greater the likelihood of detecting that an election is manipulable.

We needed to choose a number of perturbations to use consistently for our simulations that was large enough that it would not take too long to run the simulations, but also not too small that our data would be inaccurate. To choose the number of perturbations, we ran thousands of simulations each with a different number of perturbations between 10 and 200. Then the resulting data was compiled, as shown in Figure 4, with the number of perturbations for each election on the horizontal axis and the resulting proportion of manipulable elections on the vertical axis. The graph shows a steep increasing concave curve that plateaus around 125 perturbations. This plateau represents the decreasing effect that raising the number of perturbations has on the the outcome of the simulation. Therefore we decided to use 125 perturbations for all the simulations that used the random perturbation method to determine a certain election's manipulability.

Another parameter that we need to determine for our simulations is the epsilon value. As we can see in Figure 5, since manipulability is a proportion, of course it cannot be greater than one. Also an ε of 1 would mean that all the ballots would be changed, almost guaranteeing a manipulation of outcome. One thing that we learned from the graph is that the relationship between ε and manipulability is not perfectly linear as we might suspect. We already knew that we wanted a small ε, since realistically a certain candidate in the real world would only be able to manipulate a small number of ballots. We decided to choose an ε in the range where an increase in manipulability is proportional to the increase in ε. So for the following simulations these are conducted with an ε of 0.1.

Our simulation has some degree of uncertainty in calculating the manipulability because there is always a chance of not detecting an ε-manipulable election. We calculated our theoretical uncertainty using the formula $4\sqrt{p(1-p)/n}$ from the width of a binomial distribution, where p is the simulated manipulability. Then we created a graph of 30 of the same simulations to see how much the outcomes will differ because of the inherent randomness of each simulation. See Figure 6

From Figure 7, we can conclude that the lower the λ, the less manipulable the voting system is. This conclusion suggests that plurality is the least manipulable among all λ-Borda count systems.

3.1.3 Non-random Perturbation Simulation

We ran simulations using a non-random perturbation method for more accurate results. The method used in Figure 8 used specific perturbations, where ε is subtracted and added to each possible combination of pairs of ballots in an election to recreate the most extreme possible manipulations to the election to check if any cause a change in outcome. This not only is more accurate for detecting manipulability, but also makes our code much faster as the number of perturbations can only be up to 30 per election, while the random perturbation method uses over 100. The resemblance of Figures 7 and 8 to Figures 2 and 3 confirms our findings.

3.2 Generalized Borda Elimination

The Borda elimination system required us to run simulations since it is a multiple round voting system which is nearly impossible to calculate by hand, and is much quicker and easier using the techniques we developed for the Borda count systems. Figure 9 displays some unexpected properties: a λ of 0.6 appears to minimize the manipulability.

3.3 Open Problems and Future Work

We have seen good evidence that among all Borda count and Borda elimination methods, plurality is the least manipulable for three candidate voting systems. We think this trend should be true (in our simulations) for more than three candidates, but these voting methods are characterized by several parameters describing the points received by the second, third, etc. ranked candidates. Of course, the main open problem is to prove that plurality is the least manipulable (under the impartial culture assumption), even among all three candidate voting systems which treat all voters and candidates equally.

It would also be interesting to run simulations with assumptions on the ballot distribution other than the impartial culture assumption. For example, how manipulable are elections where a particular candidate is strongly preferred? What if there are two very similar candidates, which always get ranked together? What if the candidates lie on a spectrum, and voters rank the candidates in terms of (ideological, in this case, geometric) distance? There is a lot of room to adapt our simulation techniques to these questions.

4 Acknowledgements

We thank David Frankel (Uni High class of 1976) whose gift made this experience possible for University Laboratory High School students. We also thank the Illinois Geometry Lab for organizing this program.

References

[1] K. J. Arrow, A Difficulty in the Concept of Social Welfare. *Journal of Political Economy*, 58 (4): 328-346, 1950.

[2] L. Harper, Optimal numberings and isoperimetric problems on graphs. *J. Comb. Theory*, 1 (3): 385-393, 1966.

[3] S. Heilman, Designing Stable Elections: A Survey. *arXiv:2006.05460*.

[4] E. Räty, Uniqueness in Harper's vertex-isoperimetric theorem. *Discrete Mathematics* 343.4: 111696, 2020.

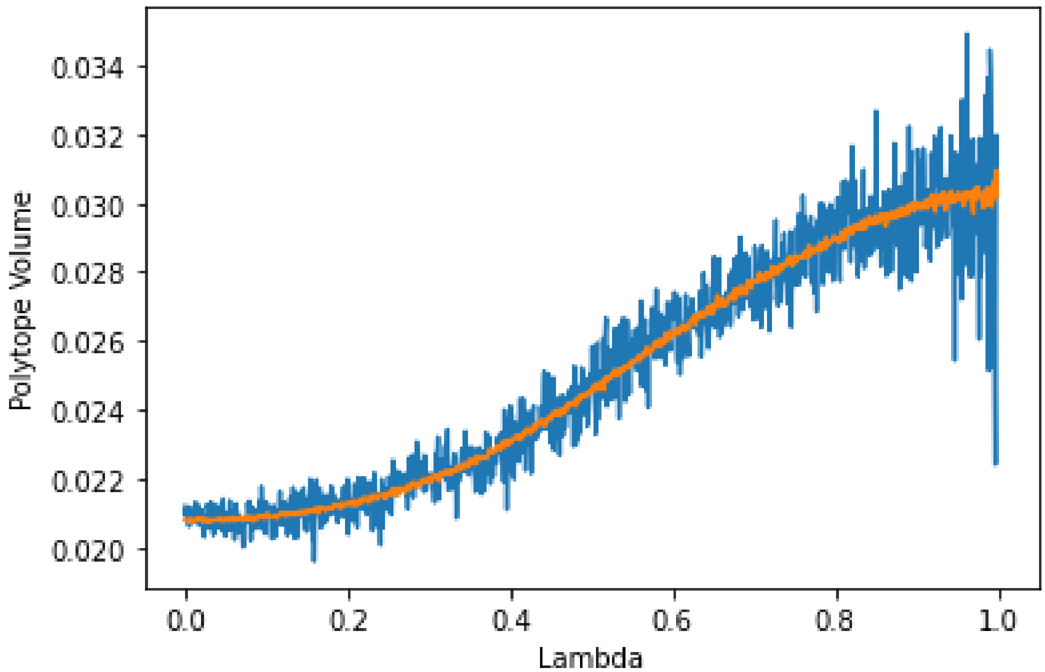

Figure 2: Noise Reduction of Polytope Volume (for variables x_2, x_4, x_5, and x_6). Orange line represents noise reduced plot; blue line is plot with noise.

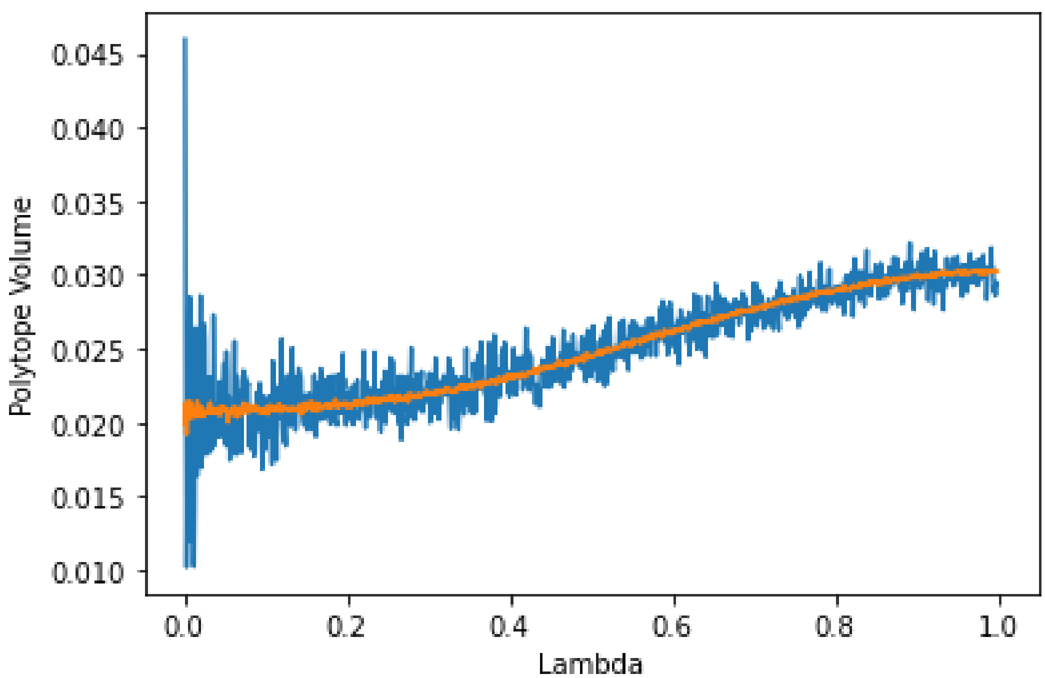

Figure 3: Noise Reduction of Polytope Volume (for variables x_1, x_2, x_3, and x_4). Orange line represents noise reduced plot, matching Figure 2; blue line is plot with noise.

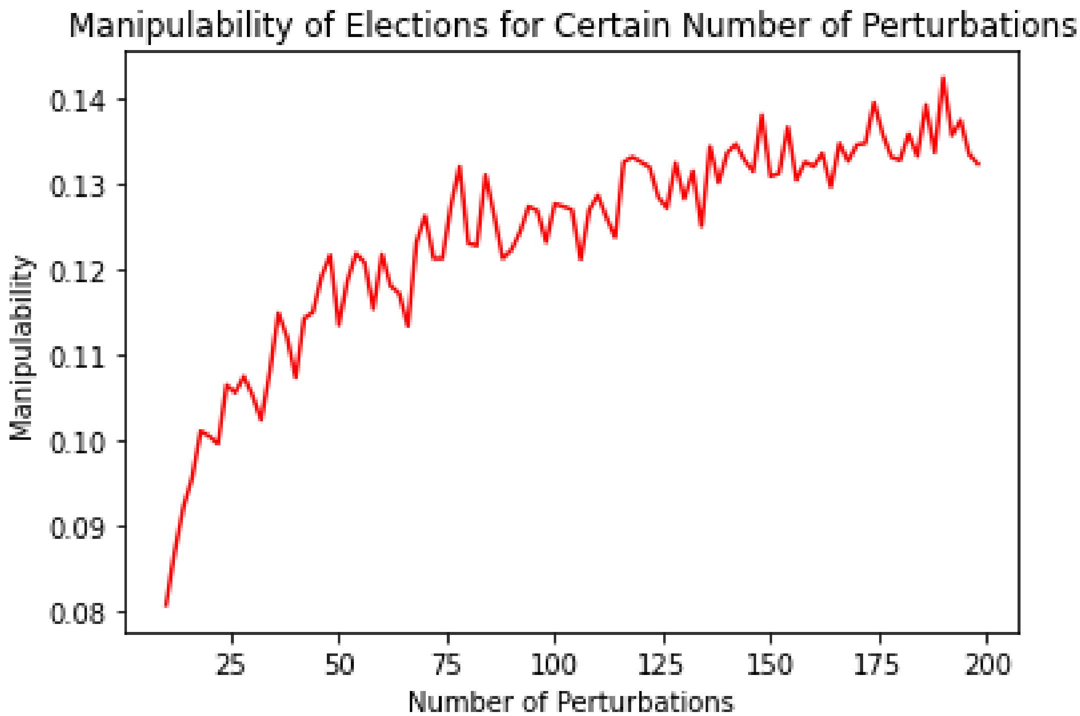

Figure 4: Increasing the number of sampled perturbations increases the accuracy of the manipulability calculation. Diminishing returns occur as the number of perturbations increases. Here we take $\lambda = 0.5$, $\varepsilon = 0.05$, and sample ten thousand elections per data point to estimate the ε-manipulability.

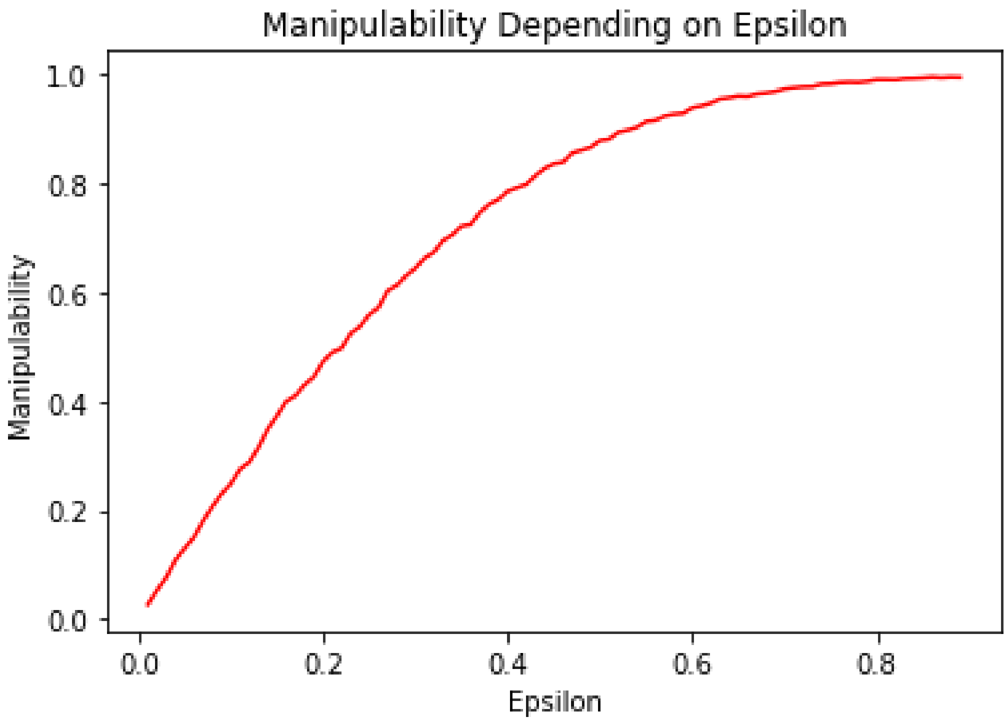

Figure 5: Running 125 random perturbation for each of ten thousand elections for each data point, we plot the ε-manipulability of 0.5-Borda count. The plateau for large ε is due to the fact that manipulability is a proportion which cannot be larger than 1.

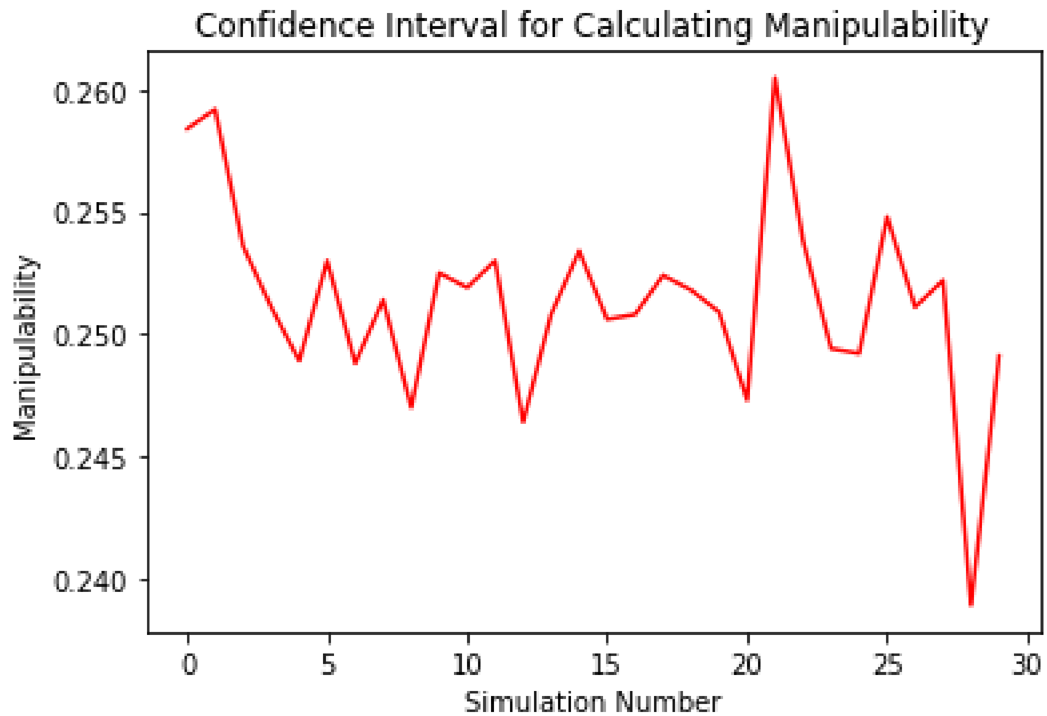

Figure 6: Running thirty identical simulations with ten thousand elections each, 125 perturbations per election, $\lambda = 0.5$, and $\varepsilon = 0.1$. Here we can directly see the kind of errors our simulations produce with these parameters. Theoretically, we expect an interval of width 0.016 to contain 95% of the data points, which is accurately represented in this plot.

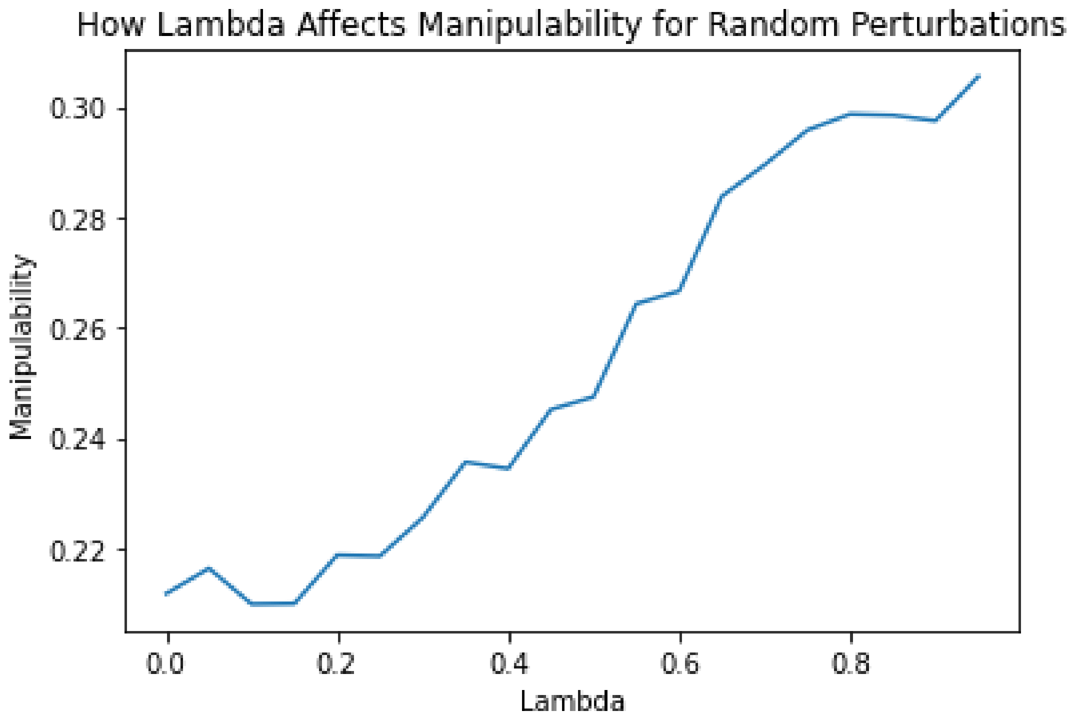

Figure 7: Each data point is ten thousand elections, each with 125 random perturbations, with $\varepsilon = 0.1$. The general shape resembles Figures 2 and 3.

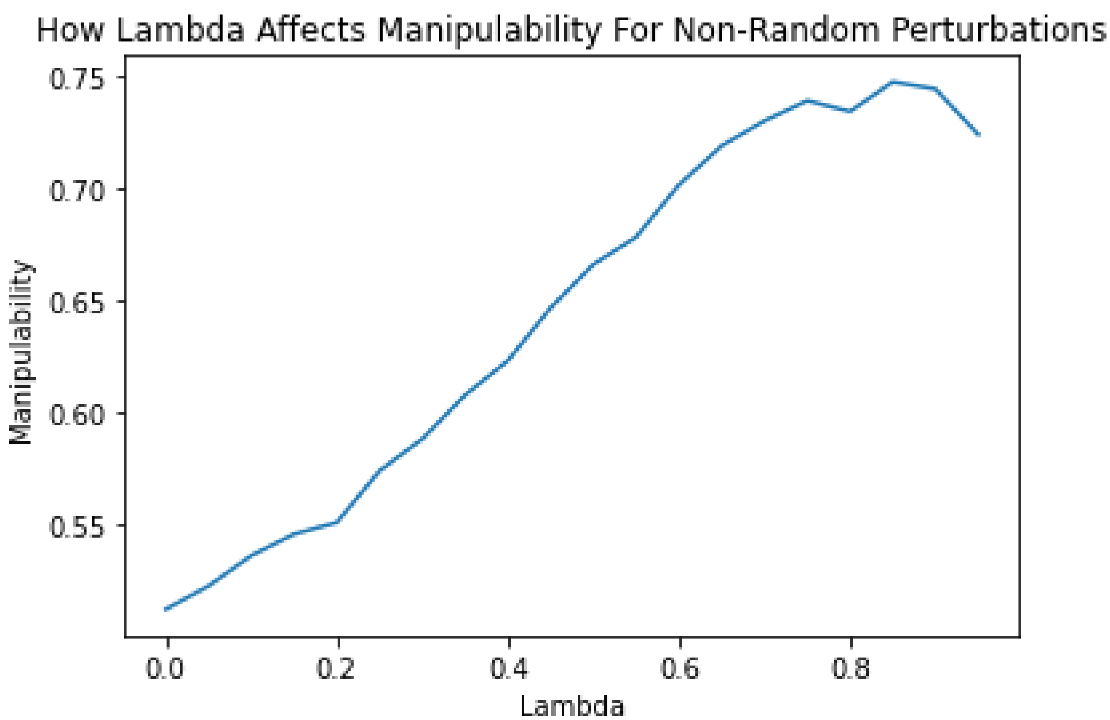

Figure 8: Each data point is ten thousand elections, each with 30 pre-selected perturbations, with $\varepsilon = 0.1$.

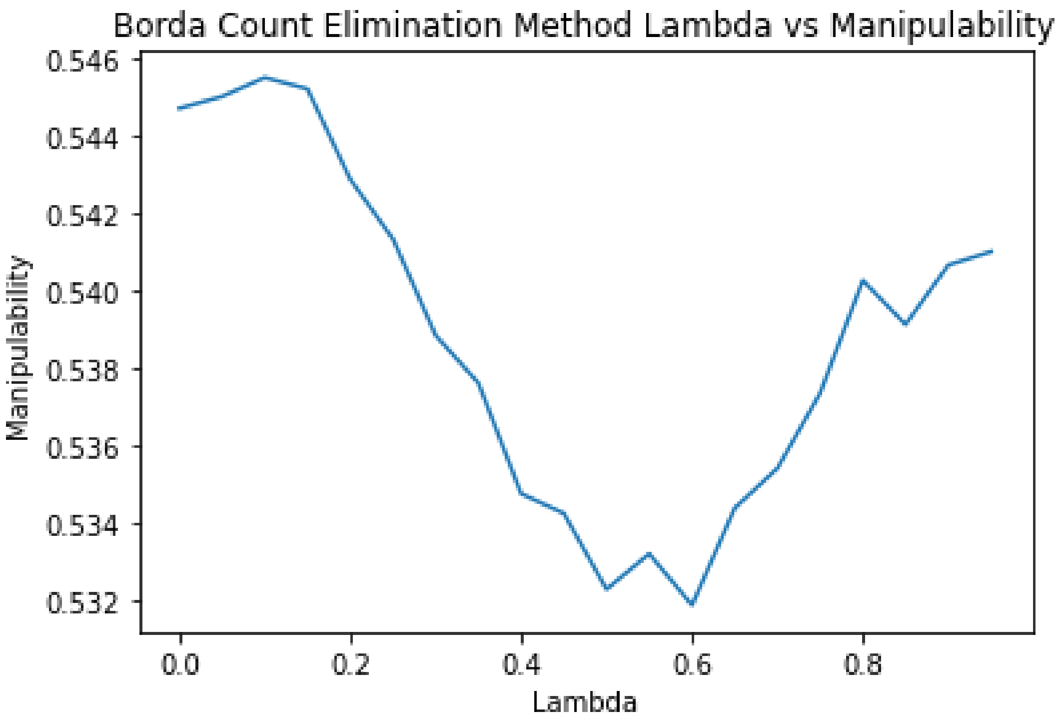

Figure 9: Each data point is five hundred thousand elections, each with 30 pre-selected perturbations, with $\varepsilon = 0.1$.