Voting Systems in Participatory Budgeting: A Comparative Study

Bachelor's Project Thesis

Imme Huitema, i.r.huitema@student.rug.nl,
Supervisor: prof. dr. D. Grossi

Abstract: Originating from Porto Alegre Brazil, participatory budgeting (PB) is "a democratic process in which community members decide on how to spend part of a public budget". It has been used as a way to involve citizens in local decision-making and to improve their understanding of democratic processes. Rutger Bregman posed it as a possible answer to the seven plagues of modern-day democracy, such as a lack of confidence in politicians. Based on recent uses of PB in Groningen, the current research compares multiple voting rules by simulating a PB process. Data was generated to represent voters' opinions and project costs, allowing for a comparison of the utilitarian, egalitarian and maximal social welfares under various circumstances. Results show that voting systems using ranked list ballots offer significant improvements over the more common one using approval ballots. Furthermore, using truncated lists does not significantly impact performance of ranked list algorithms.

1 Introduction

Participatory budgeting is a form of direct democracy, in which members of a community get to directly influence the spending of a public budget, typically by voting on various projects to improve their municipality. These projects could be anything suggested by the citizens, ranging from new playgrounds for children to access to CPR courses*.

Participatory budgeting has been shown to be beneficial for it's participants, improving living conditions. Originating from Porto Alegre in 1989 (de Sousa Santos, 1998), better access to sanitation as a direct result of participatory budgeting caused a significant decrease in child mortality (Gonçalves, 2014). It has since then spread across the Americas and Europe, nowadays being used in over 1500 cities worldwide (Baiocchi and Ganuza, 2014).

1.1 Gap in current literature

As participatory budgeting is used to select projects based on voters preference, it tries to maximize the social welfare (voter satisfaction) of the population within the constraint of the budget.

Considering computational social choice, an important aspect of participatory budgeting is the way in which voters can express their preference in the form of ballots, as well as the way these ballots produce a selection of projects. Most participatory budgeting cases so far use a form of *approval* ballots (Benade, 2018), which allow voters to either approve or disapprove of projects. This simple ballot type has benefits and drawbacks.

The main benefit of using approval ballots is that voters have an easy time filling out their ballot, as they merely need to select a preferred subset of projects and approve them. Voters have a harder time filling in more complex ballots (Camerer, 2011). Some variations of voting systems using approval ballots limit the number of projects that any voter can approve. Others make every voter solve their own knapsack problem, where voters can approve of as many projects they want as long as it fits within the budget (Goel, Krishnaswamy, Sakshuwong, and Aitamurto, 2019). This knapsack voting method is currently the most widely used voting system in participatory budgeting cases around the world.

There is however a drawback to using approval

*Retrieved from
https://stemvan.groningen.nl/budgets/4/results
visited: 2021, March 10

ballots, namely that voters can not be very specific about their preference. By simply approving a subset of projects, no information is given about distinctions between the approved and disapproved projects. There are many different voting systems that can be used which allow voters to be more specific about their preference. This increase in precision may lead to an increase in voter satisfaction, typically measured in utilitarian social welfare.

1.2 Aim of the study

This research compares performance of different voting systems which could be used in participatory budgeting. It specifically looks at those voting systems which use ranked lists as ballot. Ranked list ballots, while still relatively simple for voters to use, provide much more information about a voters preferences than approval ballots. This study measures performance of different voting systems by measuring the social welfare of voters. Results are combined with results from two other studies (Kopmels, 2021) (Pulles, 2021) to allow for comparison between many different voting systems and ballot forms.

The goal of this study is to find out how voter satisfaction changes depending on which voting system is used given certain circumstances. Because participatory budgeting is used in municipalities of various sizes, tests were done with various population sizes and variations in voting behaviour of the population.

2 Preliminaries

This preliminary section provides mathematical definitions of participatory budgeting. It explains how voter preference is modeled, and it covers the implemented voting systems, as well as the different ways in which voter satisfaction was measured.

2.1 Participatory Budgeting

Any participatory budgeting case consists of three important parts: N voters, M projects and a budget B. Every project has a cost, such that the sum of project costs is larger than the budget: if the budget is large enough, there is no reason to do any budgeting. Every voter has some order of preference among projects. This preference order of course has to be transitive: If there are three projects A, B, C a voter can not prefer A over B and B over C, while at the same time prefer C over A.

The goal of every participatory budgeting algorithm is to maximize social welfare (voter satisfaction) within the constraint given by the budget.

How the values for N, M and B, as well as the project costs and voter preferences were chosen or generated for this research is covered in section 3.

2.2 Modeling voter preference

For this research, a Mallows model was used to more realistically represent voter preference. A Mallows model is a family of r-noise models $(P_{d_{swap},p})_{1/2<p<1}$, defined by equation 2.1.

$$P_{d_{swap},p}(v|u) = \frac{1}{\mu_p}\varphi^{-d_{swap}(v,u)} \tag{2.1}$$

Where $\varphi = \frac{p}{1-p}$ and $\mu_p = \sum_{v \in L(A)} \varphi^{-d_{swap}(v,u)}$. Here, μ_p is the normalization constant, which does not depend on u. u itself is a ranking of (in this case) projects, which will be referred to as the "true ranking". This ranking is the ranking as it is supposed to be. The other voters will choose how to rank their projects, and the likelihood of choosing any variation v, is based on the distance $d_{swap}(v,u)$ between v and u. p is a constant, which determines how "true" the true ranking is. p closer to 1 will result in more voters choosing v's that are more similar to u, and a p closer to 0.5 means voters are more likely to choose a v which is not so similar to u.

The use of a Mallows model also allows for multiple true rankings. Using two true rankings that are opposite, one can model a divide in the public opinion on the projects. This allows for testing performance differences between Participatory Budgeting algorithms based on whether the population is split on how the budget should be spent, or somewhat in agreement.

After a Mallows model is made with one or multiple true rankings, every voter is assigned a ranking v based on the probability of selecting v from the set of all permutations (all possible rankings). Based on this v, values from 1 to 100 are assigned to every project for every voter, such that if project A is ranked above project B according to v for a voter, the utility of project A is higher than the utility of project B for that voter.

2.3 Voting rules

For participatory budgeting, there are many ways in which voters can express their preference (using different ballot-types), and how these votes can produce a selection of projects. A voting rule is an algorithm which takes ballots as input, and assigns points to projects based on these ballots. The amount of points a project gets indicates how much the voters favor that project.

This subsection describes the four different voting rules that are compared in this research. One of these makes use of approval ballots, which allow voters to either approve or disapprove of projects. The other three use ranked-list ballots, where voters order projects based on their preference.

2.3.1 Approval voting

Approval voting is perhaps the simplest way to count votes. It takes as input the approval ballots of voters. For every ballot it assigns 1 point to every approved project.

$$P_m = \sum_{n=1}^{N} n_m \tag{2.2}$$

Formula 2.2 shows how the number of points awarded to a project is calculated: P_m is the number of points awarded to project m. N is the total number of voters and n is a voter. Lastly, n_m is 1 if voter n approves of project m, and 0 if voter n does not approve of project m. This formula is repeated for every project m.

2.3.2 Borda voting

Borda voting (Emerson, 2013) is the first and most straightforward voting rule that uses ranked-list ballots. For every ballot it assigns points to every project based on their ranking, where higher ranked projects get more points than lower ranked projects.

$$P_r = M + 1 - r \tag{2.3}$$

Formula 2.3 shows how points are awarded to a project based on their rank using Borda voting. P_r is the amount of points awarded to a project at rank r on a voters ballot, given M total projects. It shows that there is a linear relation between rank and points given to a project (based on this rank).

2.3.3 Dowdall System voting

Dowdall System voting (Fraenkel and Grofman, 2014) works very similar to Borda voting, also giving points to projects based on their rank.

$$P_r = \frac{M}{r} \tag{2.4}$$

Formula 2.4 shows how points are awarded to a project based on their rank using Dowdall System voting. Again, P_r is the amount of points awarded to a project at rank r on a voters ballot, given M total projects. Compared to Borda voting, it heavily favors higher ranked projects. This is because the relation between rank and points is no longer linear, but instead follows an inverse relation.

2.3.4 Eurovision Song Contest voting

The final algorithm which uses a ranked list type ballot, is also a variation of borda voting, used by the Eurovision Song Contest †.

$$P_r = \begin{cases} 14 - 2 * r & \text{if } r < 3, \\ 10 - r & \text{if } 2 < r < 10, \\ 0 & \text{otherwise.} \end{cases} \tag{2.5}$$

Formula 2.5, where again P_r is the amount of points awarded to a project at rank r on a voters ballot, shows how points are awarded to projects based on their rank. It works similarly to Borda voting, with mostly a linear relation between rank and points. It differs however, in that the highest 2 projects are slightly favored, gaining an extra point over projects ranked below them. It also makes no distinction between projects ranked lower than 10^{th} place, giving 0 points to all.

2.4 Measuring voter satisfaction

In order to determine which of the voting rules worked best, voter satisfaction has to be measured

†Retrieved from https://eurovision.tv/about/voting last visited: 2021, Januari 12

in some way. To do this, three different social welfare functions (Sen, 1970) were used. This subsection explains how each of these work and why they are relevant.

2.4.1 Utilitarian social welfare

Utilitarian social welfare is a measurement of the satisfaction of the entire voter population. A higher utilitarian social welfare indicates an over-all happier population with regards to the outcome of the vote.

$$W_u = \sum_{n=1}^{N} S_n \tag{2.6}$$

Formula 2.6 shows that to get the utilitarian social welfare W_u, the satisfaction S of every voter n is summed.

2.4.2 Egalitarian social welfare

Egalitarian social welfare measures the satisfaction of the least satisfied voter:

$$W_e = \min(S_1, S_2, ..., S_N) \tag{2.7}$$

Formula 2.7, where W_e is the social welfare, and S_n is the satisfaction of the n^{th} voter. Measuring the satisfaction of the least satisfied voter is important if users want to ensure every voter has at least some amount of satisfaction. Also, in combination with utilitarian social welfare, egalitarian social welfare gives a good indication of how the satisfaction is divided between voters: having higher egalitarian social welfare while having similar utilitarian social welfare means that while the total satisfaction is the same, the satisfaction is more equally divided among voters.

2.4.3 Maximal social welfare

The third measurement of social welfare is called maximal social welfare W_m. Formula 2.8 shows that W_m measures how satisfied the most satisfied voter is:

$$W_m = \max(S_1, S_2, ..., S_N) \tag{2.8}$$

Similar to egalitarian social welfare, this measurement in combination with utilitarian social welfare says something about the division of satisfaction among voters. In this case, having the same utilitarian social welfare but higher maximal social welfare means satisfaction is less equally divided among voters.

3 Method

To learn about the differences in performance of the different Participatory Budgeting algorithms, a program was created. Figure 3.1 shows a flowchart which breaks this program down into 5 different parts: generating the project costs, generating the votes, counting the votes, determining how the budget is divided based on these votes and lastly analyzing the outcome.

For this research, four different experiments were ran every experiment consisting of 100 repetitions to allow for generalizing of the results. Every repetition would repeat the entire process from figure 3.1. To test whether or not the number of voters has any effect on voter satisfaction, two tests were run with 50 voters and two tests were run with 250. In order to test whether or not having a split opinion among the voters versus voters mostly agreeing has any effect on voter satisfaction, a Mallows model (Moulin, 2016) was used to replicate this effect. This model is explained more in depth in subsection 3.2.

The next 5 subsections describe the theory behind the 5 parts from figure 3.1, as well as the reasoning behind the specific values that were chosen for certain variables.

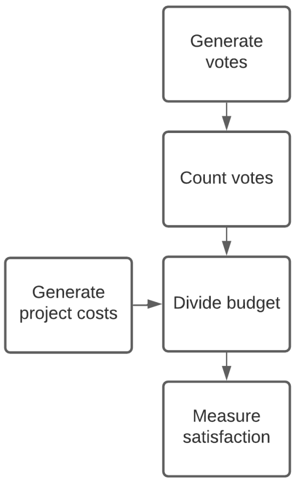

Figure 3.1: Flow chart overview of the Participatory Budgeting program used for this research.

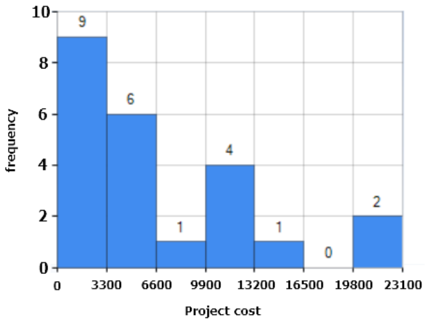

Figure 3.2: Frequency diagram of project costs from participatory budgeting in district Oosterparkwijk in Groningen in 2019.

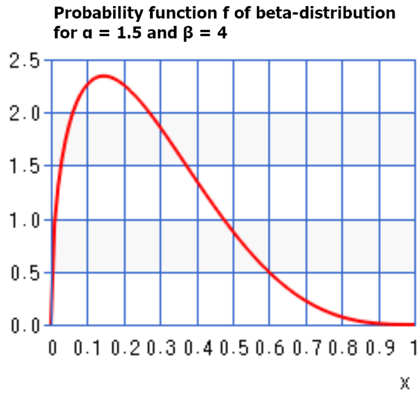

Figure 3.3: Probability function f of beta-distribution with $\alpha = 1.5$ and $\beta = 4.0$.

3.1 Generating project costs

In order to generate the project costs, data from a recent application of Participatory Budgeting in a district called Oosterparkwijk in Groningen was used as an example [‡]. It had 23 projects, and the project costs ranged from €150 to €21600. Figure 3.2 shows that most of these projects cost somewhere between €0 and €6600, with a few outliers costing more.

In order to best fit the cost of the projects to this example, a beta-distribution as shown in 3.3 was used. The values $\alpha = 1.5$ and $\beta = 4.0$ were chosen to best mach the data from Groningen. Then, based on this distribution, every project was assigned a cost. With this distribution, the minimal cost of the projects was €0, and the maximal cost €20000, with most project costs being on the lower end.

Finally, a simple check was done to ensure that the total costs of the projects exceeded the budget. If this condition explained in section 2.1 is not met a new set of project costs is generated.

The total number of projects was 25 for all 4 runs. This number is based on the 23 total projects in the example from Groningen mentioned earlier in this subsection. The total budget was €30000, based on the budgets of €25000 and €35000 from the two districts in Groningen.

[‡]Retrieved from
https://stemvan.groningen.nl/budgets/4/results
visited: 2021, Januari 12

3.2 Generating the votes

Votes (ballots) are the input of the different Participatory Budgeting algorithms. A Mallow's model as explained in section 2.2 is used to generate utilities from which these ballots are created. This subsection provides the values that were used for different variables in the Mallow's model, and it then explains how the output of this Mallow's model is used to generate ballots.

3.2.1 Mallow's model parameters

For this research, $p = 0.6$ was used for all 4 runs of the experiment. This value is closer to 0.5 than to 1 to make sure that voters do not all follow the exact same trend.

One experiment of 50 voters and one of 250 voters was run with a single true ranking u, and the other two experiments were run with two true rankings. This was done to test whether or not this would have an impact on performance of the algorithms. The true ranking u is randomly chosen. For the experiments with two true rankings u_1 and u_2, u_2 would be the inverse of a randomly chosen u_1.

In the experiments with multiple true rankings, the population was equally divided between these true rankings, meaning half the population would choose their ranking based on u_1, and the other half would choose according to u_2.

Every permutation of the rankings, v in equation 2.1, has their own probability of being selected by a voter based on how similar v is to u. Originally, this probability was calculated for every possible v. However, after deciding the number of projects should be 25, it became impossible to calculate these probabilities, as there were too many, and almost all of them would be extremely close to 0. Therefore the decision was made to only calculate the probabilities for a random subset of all the permutations of v (always including u). The size of this subset was set to be as high as possible without running into runtime errors, which turned out to be 6250.

3.2.2 Generating ballots

After the utilities have been generated, ballots have to be made to be used as input for the algorithms. As mentioned in section 2.3, this research compares 4 different algorithms: three which use ranked-lists as input, and one which uses a list of approved projects as input.

To generate the approval ballot of a voter, the program checks for every project if the utility is above a certain threshold. If it is, the project is approved and added to the ballot. For this research a threshold of 50 was chosen, to have voters on average approve of 50% of the projects.

In order to generate the ranked list ballot of a voter, all their projects are ranked based on the utility. This converts it back to their respective v as discussed in the previous subsection.

Lastly, the three ranked list algorithms were also implemented to take truncated ballots. This simply means that only the top few projects of the ranked list are represented on the ballot, and the rest has to be left out by the voter (and awarded 0 points each). This was done to see if this loss of information has a notable effect on performance. For this research, truncated lists were limited to 8 projects, so that the 17 least preferred projects are left out (and thus receive 0 points each from that voter).

3.3 Counting the votes

After ballots have been created, these are passed to algorithms based on the different voting rules discussed in subsection 2.3. These algorithms take their respective form of ballot as input, and return a ranked list of projects. The project ranked first is the project that, according to the algorithm, the population favors most.

Processing the approval votes is done by simply repeating formula 2.2 for every project m.

In order to process all ranked-list ballots, the calculations shown in formulas 2.3, 2.4 and 2.5 are summed per project as they are repeated for every project on every ballot.

This yields a list of projects with their points. These projects are then simply ordered based on how many points they received from the algorithm used.

3.4 Dividing the budget

The budget has to be divided over the projects in such a way that projects preferred most by the voters get prioritized. In order to do this knapsack budgeting is used on the list of ranked projects produced as a result of counting the votes. Knapsack budgeting is an algorithm easiest explained with an example of filling a knapsack. Imagine going on a road trip, and you have with you a knapsack that can only hold a certain weight of items. There are multiple different items you want to take with you, weighing different amounts, some of which you value more than others. Knapsack budgeting solves this problem by filling the knapsack with the most valued items first, until the weight limit is reached.

This is also what was used to ultimately decide which projects would get approved and which would not. Every project has a certain cost, as described in section 3.1. Each of the algorithms covered in section 3.3 returns a list of projects ranked from most to least favored by the population. The budgeting algorithm will approve these projects one by one, starting with the most favored project, as long as they fit within the total budget. Any projects that do not fit within the budget are not approved.

3.5 Collecting the results

In order to compare performance of the different Participatory Budgeting algorithms, the three social welfare functions covered in subsection 2.4 were used.

In order to measure how satisfied a voter was with the outcome of a vote, the utilities of a voter for all projects that ended up being approved are summed up to portray this voter's welfare. If this is done for every voter, formulas 2.6, 2.7 and 2.8 can be used to calculate the utilitarian, egalitarian and maximal welfare of the population given a vote.

3.6 Data analysis

For this research, R-studio version 1.3.1093 was used for all statistical analysis as well as visualization of the results. In order to determine whether or not the data was normally distributed, a Shapiro-Wilk test was performed. The statistical test used to compare differences in mean social welfare given different circumstances was a multivariate analysis of variance.

3.7 General remarks

All the programming was done in Python 3, working on an Ubuntu (20.04) operating system. The main program was made in collaboration with colleagues Lonneke Pulles and Marieke Kopmels.

4 Results

This section discusses the results gathered during the 4 runs. Note that the plots show results for many more voting systems than discussed in this paper: These results were gathered with the same program, and the different algorithms are discussed in more detail in papers by L. Pulles and M. Kopmels (Kopmels, 2021) (Pulles, 2021).

4.1 Average difference in social welfare

First, the differences in average social welfares between each algorithm are presented in the form of heatmaps, in figures 4.1, 4.2 and 4.3. In these heatmaps, the colored squares represent the difference in welfare between the algorithms shown on the left-hand side and the algorithms shown above. A blue square indicates the algorithm on the left has higher social welfare compared to the algorithm above. In addition, a red square means the algorithm on the left has lower social welfare than the algorithm above. The darker this color, the bigger the difference. The differences were calculated using one way analysis of variance. Significant differences ($p - value < 0.05$) in social welfare are highlighted with stars.

4.1.1 Utilitarian social welfare

First off, figure 4.1 shows the difference in utilitarian social welfare between the different algorithms. It shows that there are no significant differences in utilitarian social welfare between approval, borda, dowdall system and eurovision voting. The truncated variants of borda, dowdall system and eurovision voting also do not perform significantly better or worse compared to each other or their non-truncated variants.

The fact that no significant differences were found in utilitarian social welfare indicates that, on average, the welfare of the entire population combined is the same for any of the voting systems discussed in this paper.

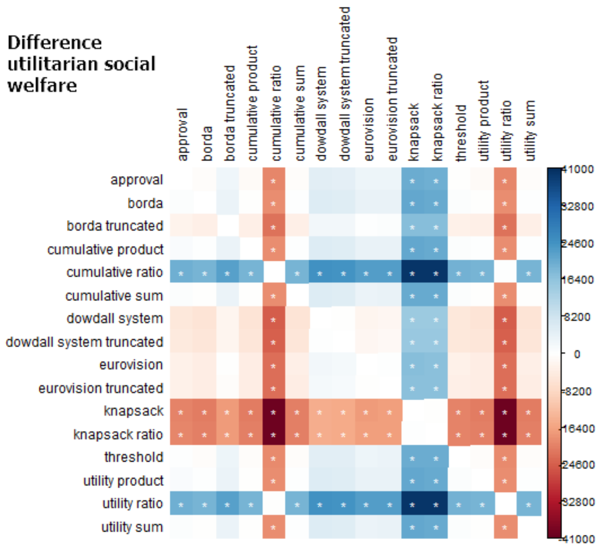

Figure 4.1: Heatmap of differences in utilitarian social welfare between the voting rules, averaged over 4 different runs.

4.1.2 Egalitarian social welfare

Secondly, figure 4.2 shows the difference in egalitarian social welfare between the different voting systems. The pattern is similar to figure 4.1, and there is no significant difference in egalitarian social welfare between any of the algorithms discussed in this paper.

This means that, on average, the least satisfied voter has the same welfare for any of the 7 voting systems discussed in this paper.

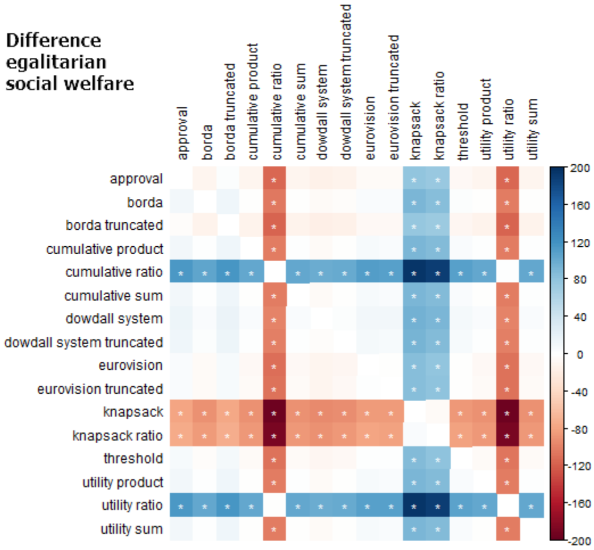

Figure 4.2: Heatmap of differences in egalitarian social welfare between the voting rules, averaged over 4 different runs.

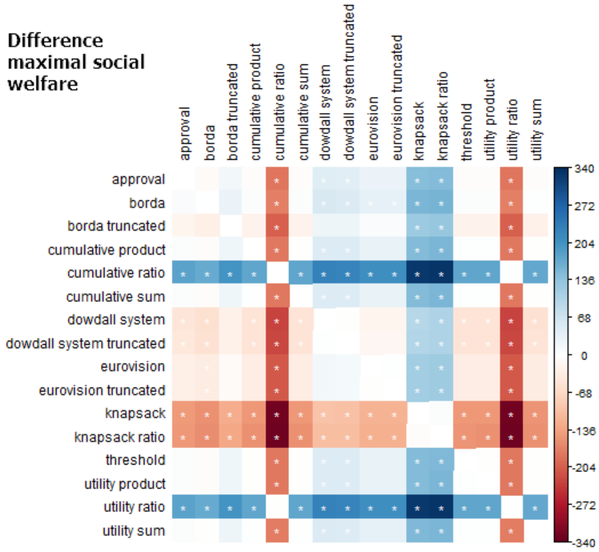

Figure 4.3: Heatmap of differences in maximal social welfare between the voting rules, averaged over 4 different runs.

4.1.3 Maximal social welfare

Lastly, the difference in maximal social welfare between the voting systems is shown in figure 4.3. Again, the voting systems discussed in this paper have mostly insignificant differences in performance. This time there are, however, a few exceptions: Approval voting has significantly higher maximal social welfare than both dowdall system voting and its truncated variant. Borda also scores significantly higher than dowdall system voting and eurovision voting, as well as their truncated variants.

These results, especially in combination with the fact that no significant differences were found in both utilitarian and egalitarian social welfare, say something about the distribution of welfare. In this case, because the difference in utilitarian social welfare is not significant, it means that some voters had to sacrifice some of their welfare to add to that of the most satisfied/welfaring voter. And because the egalitarian social welfare is also not significantly different, these voters that sacrificed their welfare did not give up so much that the least welfaring voter became significantly worse off.

4.2 Effect of population size on social welfare

To see whether or not population size has a significant impact on social welfare, and to see if this differed for the different voting rules, two-way analysis of variance was performed. For this, the data gathered from the two tests that had a population size of 50 voters was compared with the data gathered from the two tests that had a population size of 250 voters. These results are shown in tables 4.1 and 4.2. Note that the difference shown in these tables is the difference going from 50 to 250 voters, or the difference going from one to two true rankings. Also shown in the table are the p-values of the statistical test, this is used to determine whether the difference in mean social welfare is actually significant ($p < 0.0001$).

4.2.1 Utilitarian social welfare

In order to see the effect of the number of voters on utilitarian social welfare, the mean utilitarian social welfare of the population was calculated. The mean is used to account for the fact that a population of 250 voters would of course have 5 times the utilitarian social welfare of a population of 50 voters if number of voters does not affect utilitarian social welfare. The results show that changing from 50 to 250 voters caused no significant change in utilitarian social welfare for any of the voting rules, see the p-values shown in table 4.1.

4.2.2 Egalitarian social welfare

For egalitarian social welfare, table 4.2 shows that an increase in voters lowers egalitarian social welfare. However, this negative change was insignificant for any of the voting systems discussed in this paper, indicated by the p-values being higher than 0.0001.

4.3 Effect of voter preference on social welfare

Two-way analysis of variance was also used to learn more about the effect of a divided opinion among voters on social welfare. For this, the mean social welfares are compared between the two tests in which voters followed one true ranking, and the two tests in which voters had a split opinion: voting according to two opposite true rankings, as explained in section 3.2.1.

4.3.1 Utilitarian social welfare

Table 4.3 shows that having a split opinion in the population leads to significantly lower mean utilitarian social welfare compared to no split opinion for some algorithms. The algorithms for which the difference was significant are truncated borda, and both Eurovision and truncated Eurovision.

4.3.2 Egalitarian social welfare

For all voting systems discussed in this paper, changing from one true ranking to two opposite true rankings caused a significantly lower egalitarian social welfare, as indicated by the low p-values in table 4.4.

Table 4.1: Decrease in mean utilitarian social welfare changing from 50 voters to 250 voters, per voting rule.

Voting rule	Change	p-value
Approval	-2.3639	$8.179 * 10^{-1}$
Borda	6.00	$5.600 * 10^{-1}$
Borda truncated	5.72	$5.727 * 10^{-1}$
Dowdall	19.02	$6.761 * 10^{-2}$
Dowdall truncated	14.23	$1.594 * 10^{-1}$
Eurovision	6.32	$5.304 * 10^{-1}$
Eurovision truncated	8.69	$3.771 * 10^{-1}$

Table 4.2: Decrease in egalitarian social welfare changing from 50 voters to 250 voters, per voting rule.

Voting rule	Difference	p-value
Approval	31.15	$5.51 * 10^{-3}$
Borda	30.59	$5.615 * 10^{-3}$
Borda truncated	29.73	$2.747 * 10^{-3}$
Dowdall	33.43	$8.614 * 10^{-4}$
Dowdall truncated	30.44	$1.831 * 10^{-3}$
Eurovision	27.43	$8.823 * 10^{-3}$
Eurovision truncated	28.35	$5.114 * 10^{-3}$

Table 4.3: Decrease in mean utilitarian social welfare changing from 1 to 2 true rankings, per voting rule.

Voting rule	Difference	p-value
Approval	50.91	$4.735 * 10^{-7}$
Borda	48.67	$1.547 * 10^{-6}$
Borda truncated	34.51	$6.082 * 10^{-4}$
Dowdall	48.29	$2.663 * 10^{-6}$
Dowdall truncated	47.41	$2.056 * 10^{-6}$
Eurovision	33.91	$7.040 * 10^{-4}$
Eurovision truncated	30.94	$1.563 * 10^{-3}$

Table 4.4: Decrease in egalitarian social welfare changing from 1 to 2 true rankings, per voting rule.

Voting rule	Difference	p-value
Approval	153.25	$2.2 * 10^{-16}$
Borda	146.21	$2.2 * 10^{-16}$
Borda truncated	110.85	$2.2 * 10^{-16}$
Dowdall	105.89	$2.2 * 10^{-16}$
Dowdall truncated	105.72	$2.2 * 10^{-16}$
Eurovision	121.51	$2.2 * 10^{-16}$
Eurovision truncated	114.19	$2.2 * 10^{-16}$

4.4 Other voting systems

4.4.1 Ratio voting

Both cumulative ratio voting and utility ratio voting (Pulles, 2021) outperform all other voting systems, having significantly higher social welfare than all other voting systems (see figures 4.1, 4.2 and 4.3). This is as expected as these voting systems allow users to very precisely specify their preferences.

4.4.2 Knapsack voting

Interestingly, figures 4.1, 4.2 and 4.3 show that both knapsack voting and knapsack ratio (Kopmels, 2021) score significantly lower social welfare compared to all other voting systems. This is unexpected because, as stated in section 1.1, knapsack voting is currently the most widely used voting system for participatory budgeting.

5 Conclusion

To conclude, voter satisfaction does not differ much between any of the voting systems discussed in this paper. Approval voting, borda, dowdall system and the voting rule based on the eurovision song contest, as well as the truncated list variations of these rules, do not differ significantly with regards to utilitarian and egalitarian social welfare.

Some significant differences were found with regards to maximal social welfare. Approval voting had significantly higher maximal social welfare compared to dowdall system voting and truncated dowdall system voting. Borda voting also had significantly higher maximal social welfare compared to dowdall system voting and eurovision voting as well as their truncated variants. This, in combination with the fact that no significant differences in utilitarian social welfare was found, means that dowdall system voting might be preferred over borda and approval voting, in case a more equal distribution of welfare among voters is desired.

As to the effect of population size, both utilitarian and egalitarian social welfare of the voting systems did not change significantly when changing from 50 to 250 voters.

Lastly, the results showed that having a split opinion among voters as opposed to voters being mostly in agreement led to a significant drop in egalitarian social welfare for all voting systems. For utilitarian social welfare, this negative change was only significant for borda voting, dowdall system voting, and truncated dowdall system voting.

6 Discussion

6.1 Limitations

Some of the limitations which make it difficult to generalize the findings of this research to what might occur in a real world participatory budgeting case are the following:

Firstly, for this research the assumption was made that voters would not vote strategically. A strategic voter does not vote true to their preference, but instead changes their vote based on what they might believe would lead to the best outcome, taking into account what they expect others to vote. This is difficult to model and therefore this research assumed no strategic voting behaviour would occur.

Secondly, the way project costs were generated might not be realistic. The beta-distribution used for this research is based on only two real participatory budgeting cases, and therefore a very rough estimate of what it might be like for other participatory budgeting cases.

Thirdly, it is difficult to translate the Mallows model to real world voting preferences. The hyperparameters chosen are at best educated guesses to what would lead to the most realistic model, as the model is quite abstract. For example, perhaps having $p = 0.8$ would lead to more realistic results, but this is difficult to say.

Lastly, due to limitations in computing power, this research was limited with regards to number of projects and number of voters. Many real world applications of participatory budgeting have more projects and more voters.

6.2 Future research

As participatory budgeting is still in its infancy, there are still many aspects that could be researched further.

One very interesting takeaway from the results of this research is that there barely seems to be a drawback for using truncated lists over full ranked list ballots. The fact that no significant differences in social welfare were found between any of the ranked list voting rules and their truncated variants shows that apparently the big loss of precision (no distinction is made between any of the 17 lowest ranked projects) does not impact voter satisfaction. This makes truncated lists preferable over full ranked lists, as it is easier for voters to fill in such a ballot. Future research could try to find the point where the loss of precision starts to lead to lower voter satisfaction.

Another suggestion for future research would be to use these different voting rules in a real participatory budgeting case. Surveys could be used to find out voter satisfaction, and results could be compared to this research to get a better understanding of the performance of the different voting rules.

Future research could also reverse-engineer outcomes of participatory budgeting cases that have already taken place, to feed as input into a simulation such as the one presented in this paper. This would provide insight into how the outcome might have changed if a different voting system had been used instead, and how this might affect voter satisfaction.

References

Gianpaolo Baiocchi and Ernesto Ganuza. Participatory budgeting as if emancipation mattered. *Politics & Society*, 42(1):29–50, 2014.

Gerdus Benade. Efficiency and usability of participatory budgeting methods. 2018.

Colin Camerer. *Behavioral Game Theory: Experiments in Strategic Interaction.* Princeton University Press, 2011.

Boaventura de Sousa Santos. Participatory budgeting in porto alegre: Toward a redistributive democracy. *Politics & Society*, 26(4):461–510, 1998.

Peter Emerson. The original borda count and partial voting. *Social Choice and Welfare*, 40(2): 353–358, 2013.

Jon Fraenkel and Bernard Grofman. The borda count and its real-world alternatives: Comparing scoring rules in nauru and slovenia. *Australian Journal of Political Science*, 49, 04 2014.

Ashish Goel, Anilesh K. Krishnaswamy, Sukolsak Sakshuwong, and Tanja Aitamurto. Knapsack voting for participatory budgeting. 7(2), 2019.

Sónia Gonçalves. The effects of participatory budgeting on municipal expenditures and infant mortality in brazil. *World Development*, 53:94–110, 2014.

Marieke Kopmels. The effect of different voting methods in participatory budgeting. 2021.

Hervé Moulin. *Handbook of Computational Social Choice*, chapter 8.3, pages 184–195. Cambridge University Press, 2016.

Lonneke Pulles. 2021.

Amartya Sen. *Collective Choice and Social Welfare*, chapter 3, pages 33–46. Holden-Day, 1970.