



Columbia University

*Department of Economics
Discussion Paper Series*

**Protecting Minorities in Binary Elections:
A Test of Storable Votes Using Field Data**

Alessandra Casella, Shuky Ehrenberg, Andrew Gelman and Jie Shen

Discussion Paper No.: 0708-14

*Department of Economics
Columbia University
New York, NY 10027*

May 2008

Protecting Minorities in Binary Elections: A Test of Storable Votes Using Field Data*

Alessandra Casella[†] Shuky Ehrenberg[‡] Andrew Gelman[§]
Jie Shen[¶]

May 2008.

Abstract

Democratic systems are built, with good reason, on majoritarian principles, but their legitimacy requires the protection of strongly held minority preferences. The challenge is to do so while treating every voter equally and preserving aggregate welfare. One possible solution is *storable votes*: granting each voter a budget of votes to cast as desired over multiple decisions. During the 2006 student elections at Columbia University, we tested a simple version of this idea: voters were asked to rank the importance of the different contests and to choose where to cast a single extra "bonus vote," had one been available. We used these responses to construct distributions of intensities and electoral outcomes, both without and with the bonus vote. Bootstrapping techniques provided estimates of the probable impact of the bonus vote. The bonus vote performs well: when minority preferences are particularly intense, the minority wins at least one of the contests with 15–30 percent probability; and, when the minority wins, aggregate welfare increases with 85–95 percent probability. When majority and minority preferences are equally intense, the effect of the bonus vote is smaller and more variable but on balance still positive.

*We thank A. Sacarny for help with programming, and R. Davidson for discussions on the bootstrap. The research was supported by NSF grants SES-00214013 and SES-0084368 and by the Guggenheim Foundation. A. Casella thanks NBER for its hospitality, and seminar audiences in Marseille and at the 2007 European Summer Symposium in Economic Theory for comments.

[†]Department of Economics, Columbia University, 420 West 118th Street, New York, N.Y. 10027; Greqam, NBER and CEPR; ac186@columbia.edu

[‡]Yale Law School, PO Box 208215, New Haven, CT 06520-8215; shuky.ehrenberg@yale.edu.

[§]Department of Statistics and Department of Political Science, Columbia University, 420 West 118th Street, New York, N.Y. 10027; gelman@stat.columbia.edu.

[¶]Department of Statistics, University of California, Irvine, Bren Hall 2019, Irvine, CA 92697-1250; jshen3@uci.edu.

1 Introduction

When voters must choose between two alternatives, majority voting works well, with a single important drawback: the winning alternative commands the wider support, but not always the more intense. At least since Madison, Mill, and Tocqueville, political thinkers have argued that a necessary condition for the legitimacy of a democratic system is to set limits to the *tyranny of the majority* and provide some expression to intense minority preferences.¹ With increased recourse to direct democracy, the goal becomes particularly important because direct democracy deprives minorities of the protections afforded by the election of a diverse legislature. Devising systems to protect minorities when binary choices are at stake is thus considered central to the establishment of referendums as legitimate tools of policy-making.² The challenge is to do so while treating every voter equally, and avoiding the inertia and inefficiencies of super-majorities or veto powers.

A possible answer comes from *storable votes*, a simple voting mechanism in which individuals voting over multiple binary proposals are granted, in addition to their regular votes, one or more bonus votes to cast as desired over the different proposals. Decisions are then taken according to the majority of votes cast. The bonus votes allow voters to single out the issues that they each consider most important, and make it possible for the minority to win occasionally. But because every voter is treated equally, the majority loses only if, on average, its members consider a given issue a low priority, while members of the minority do not: the majority loses when it should, from an efficiency point of view. Although counterexamples can be found, typically the increase in minority representation comes together with an increase in aggregate expected welfare, relative to simple majority voting.³

Storable votes resemble *cumulative voting*, a voting system used in corporations and in some local jurisdictions in the United States and recommended by the courts exactly to redress violations of fair voting rights. As with storable votes, cumulative voting assigns to each voter a budget of votes to spend freely over multiple choices, but cumulative voting applies to a single election with multiple representatives, while storable votes apply to multiple decisions, each between two alternatives. The strategic game induced by the voting scheme is substantially different, and to our knowledge existing analyses of cumulative voting do not discuss the efficiency properties of the scheme.⁴

¹See for example the discussion in Issacharoff, Karlan and Pildes, 2002, and Guinier, 1994.

²Gerber, 1999, and Matsusaka, 2004.

³Casella, 2005. A similar voting scheme system is proposed, independently, in Hortala-Vallve, 2005. As shown by Jackson and Sonnenschein, 2007, the idea of linking different decisions to elicit the intensity of preferences truthfully can be applied generally. The mechanism proposed by Jackson and Sonnenschein achieves the first best asymptotically, as the number of linked decisions becomes large, but is significantly more complex than storable votes.

⁴Experiences with cumulative voting in local elections are discussed in Bowler et al., 2003, and Pildes and Donoghue, 1995. Cox, 1990, studies the scheme theoretically, Gerber et al., 1998, test it experimentally. Cumulative voting was advocated particularly by Guinier, 1994.

Storable votes are simple and, because voters are routinely presented several referendums at the same time, could be implemented as a minor modification of existing voting practices. So far, testing has been limited to small groups in laboratory settings.⁵ But the tight controls of the laboratory should be complemented by field studies where the number of voters is larger, and voters' preferences are observed, as opposed to being induced. This paper reports on a test based on field data collected during an actual election. The goal is to evaluate the impact of storable votes on the probability of minority victories, aggregate welfare, and voters' ex-post inequality.

In the spring of 2006, we attached a short survey to students' election ballots in two different schools at Columbia University, asking students to rank the importance assigned to all binary contests on the ballot, and to indicate where they would have cast an additional bonus vote, had one been available. An identifier connected responses and actual voting choices, allowing us to construct distributions of intensities and electoral outcomes, both without and with the bonus vote. Bootstrapping techniques provide estimates of the bonus vote's probable impact. As a robustness check, in addition to the use of the bonus vote reported in the survey, we study three alternative plausible rules for casting the bonus vote. For each of the four cases, we estimate three measures: (1) the frequency with which the bonus vote allows the minority to win at least one election in each set; (2) the difference in aggregate welfare, comparing the hypothetical outcome using bonus votes to simple majority voting; and (3) the impact of the bonus vote on ex post inequality. We find that the bonus vote works well: when minority preferences are particularly intense, the minority wins at least one of the contests with 15–30 percent probability, ex post inequality falls, and yet when the minority wins aggregate welfare increases with 85–95 percent probability. When majority and minority preferences are equally intense, the effect of the bonus vote is smaller and more variable, but on balance still positive.

Student elections are rarely considered worth of study because they are low-stake contests, measuring students' popularity more than the attractiveness of their electoral platforms, in a population that is hardly representative of a typical electorate. But our work is not about extrapolating political tendencies from campus elections, nor, because it focuses on binary contests only and non-representative samples of students, is it really about the campus elections themselves. Our objective is to study the voting mechanism and its comparison to simple majority voting in realistic large scale elections where the distribution of preferences *is not controlled by the researchers* — this is the central requirement here. For our purposes, the true determinants of the voters' preferences are secondary. Even with the different methodology, the test remains experimental, and the elections we study artificial.

The paper proceeds as follows. In section 2 we discuss briefly the theory behind the idea of granting a bonus vote in large elections; section 3 describes

Chwe, 1990, presents an efficiency argument in favor of minority representation that arises from information revelation in common values environment. His approach is unrelated to ours, where majority and minority interests are opposed.

⁵Casella et al., 2006.

the design of the survey and the data; section 4 studies the bonus vote choice; section 5 describes the bootstrapping exercise and its results, and section 6 concludes. A copy of the questionnaire and some additional data and results are reported in the Appendix.

2 The Theory

We begin by summarizing briefly the theory of storable votes in large elections. The results described in this section are drawn from Casella and Gelman (forthcoming), where the model is phrased in terms of contemporaneous referendums over several proposals. Here, in line with Columbia student elections, the two outcomes of each election are labeled as two different candidates.

A large number n of voters are asked to vote, contemporaneously, on a set of K unrelated elections (with $K > 1$). Each election E_k is between two candidates, a_k and b_k , with $k = 1, \dots, K$. Voters are asked to cast one vote in each election, but in addition are given a single bonus vote that can be spent on any of the elections. Each election is won by the candidate with most votes, including bonus votes.

The valuation that voter i attaches to election k is summarized by v_{ik} . By convention, a negative valuation indicates that i favors candidate a , and a positive valuation that i favors b . The valuation's absolute value, denoted by v_{ik} , is the voter's differential utility from having his or her preferred candidate win the election — the *intensity* of the voter's preferences. Voter i 's utility function is then $U_i = \sum_{k=1}^K u_{ik}(E_k)$ where $u_{ik}(E_k) = v_{ik}$ if i 's favorite candidate wins E_k , and 0 otherwise.

Individual valuations are drawn independently across individuals from a joint distribution $\mathcal{F}(v_1, \dots, v_K)$. The assumption of a common distribution is in practice equivalent to assuming that we have no additional knowledge about individual voters (which, for example, would allow different distributions for men and for women, or for English and for history majors). In this section, although not in the empirical analysis, we also assume that individual preferences are independent across elections, and thus voters' valuations over election k are drawn from some distribution $F_k(v)$. Each individual knows his or her own valuation over each proposal, but only the probability distribution of the others' valuations. There is no cost of voting.

Simple majority voting is designed to give weight to the *extent* of support for a candidate — the mass of voters who in election k prefer a to b . Storable votes allow voters to express not only the direction of their preferences but also, through the bonus vote, their *intensity*. How important accounting for intensity is depends crucially on the shape of the distributions $F_k(v)$.

Suppose first that the environment is fully symmetric, and nothing systematic distinguishes either the elections or the two sides in each election: $F_k(v) = F(v)$ for all k and $F(v)$ symmetric around 0. Then: (a) in equilibrium each voter casts the bonus vote in the election to which the voter attaches highest intensity; (b) with positive probability, one or more of the elections,

although not all, are won by a candidate supported by a minority of the electorate; and (c) ex ante expected utility with storable votes is higher than ex ante expected utility with simple majority voting.⁶

The results are clean, but the assumption of full symmetry is strong and in fact minimizes the importance of the voting rule. The welfare gains must vanish asymptotically as the number of voters becomes very large and the empirical frequencies of the preference draws approximate more and more precisely the theoretical distributions. In the limit, in each election the two candidates come to be supported by an equal mass of voters with equal distribution of intensity, and thus, from a welfare point of view, become interchangeable.

Allowing for asymmetries is important, but not only are asymmetries difficult to handle analytically, it is also not clear how best to model them. Suppose first that all asymmetries came from the extent of support: in each election, one candidate is more popular than the other, and the difference in popularity — the difference in the expected mass of supporters — in general varies across elections. Conditional on supporting either candidate, however, the distribution of preferences intensity is equal across the two groups of supporters.⁷ This is the environment to which majority voting is ideally suited. Then: (a) In equilibrium all voters cast their bonus vote in the election expected to be closest. (b) Storable votes and simple majority yield identical outcomes: the minority never wins. (c) Both voting rules are ex ante efficient.

Now suppose instead that the source of asymmetry is the intensity of support. In all elections, the two candidates are equally popular ex ante, but intensities are not equally distributed. For concreteness, suppose that in each election k supporters of candidate a_k have higher expected intensity.⁸ Then: (a) in equilibrium each voter casts the bonus vote in the election with highest valuation; (b) a minority candidate is expected to win occasionally with positive probability; (c) ex ante expected utility with storable votes is higher than expected utility with simple majority if a voter's highest valuation is expected to be on a candidate of type a ; and (d) the difference in ex ante utility does not disappear asymptotically. Condition (c) is satisfied automatically if, for example, the distribution of intensities among a_k supporters first-order stochastically dominates the distribution of intensities among b_k supporters. More generally, it depends on the shape of the distributions and on the number of elections K .

These results are interesting but come short of characterizing the properties

⁶To be precise, in the scenario described in the text the welfare superiority of storable votes holds only if the value of the bonus vote is not too high. In a slightly more general model, however, where the probability of supporting either candidate is not $1/2$ but is stochastic and distributed according to some distribution symmetric around $1/2$, the result holds for all values of the bonus vote. The assumption of symmetry across elections ($F_k(v) = F(v)$) simplifies the analysis but can be dropped fairly easily; what really matters is the symmetry across the two candidates ($F(v)$ symmetric around 0).

⁷Formally, call $G_{ak}(v)$ the distribution of intensity for supporters of candidate a in election k , and p_k the ex ante probability of favoring candidate a in election k . Then suppose $G_{ak}(v) = G_{bk}(v) = G(v)$ for all k , but $p_k \neq p_s$ for $k \neq s$.

⁸Suppose $p_k = p = 1/2$, but $G_{ak} = G_a \neq G_{bk} = G_b$ for all r . In particular suppose that G_a has higher mean than G_b : $Ev_a > Ev_b$.

of storable votes in realistic situations, where both extent and intensity of support differ across candidates and across elections. It is difficult to say anything general in such cases because equilibrium strategies will depend on the exact shape of the distributions of valuations. And in practice, would voters have the information and the rationality to solve as complex a problem? The limitations of a purely theoretical analysis are the motivation for this study: what shapes do the distributions of valuations take in practice? What strategies do voters follow? How do storable votes compare to majority voting then?

3 The Experiment

3.1 The Design

Several of Columbia’s schools hold elections in the spring, and the students’ organizations and the deans at the School of General Studies (GS) and Columbia College (CC) agreed to collaborate with us. In each school, voters elected representatives for multiple positions: GS students voted on a total of twenty different elections and CC students on twelve. We selected the subset of elections with two candidates or two mutually exclusive party lists only — three elections in GS (Board President, Alumni Representative and International Representative), and four elections in CC (Executive Board, Senator-Two Year Term, Senator-One Year Term, and Academic Affairs Representative)⁹.

All voting was electronic, and at the end of the ballot students were invited to participate in our survey. In the GS case, interested students clicked on a link and were redirected to a web page containing the survey. Students’ votes and their responses to the survey were matched and saved under anonymous identifiers and were later forwarded to us by the student bodies supervising the elections. In the CC case, the survey was on paper, because of logistical difficulties with the voting stations. At the end of their electronic ballot, CC students interested in answering the survey were given a number to be copied at the top of the paper questionnaire. The number allowed us to link their responses to their actual votes, again forwarded to us anonymously after the voting was concluded.

Understandably, the organizers of the elections were concerned with keeping our interference minimal, and the survey had to be short. We asked two sets of questions: first, how much the voter cared about the outcome of each of those elections, on a scale from 1 (not at all) to 10 (a lot); second, in which of these elections the voter would cast a single additional bonus vote in support of his or her favorite candidate, had one been available. The paper questionnaire for CC is reproduced in the Appendix; the electronic GS questionnaire was identical, with one exception — being electronic, and thus faster, we added a question about expected election outcomes.

⁹We excluded the elections for class presidents or representatives because they concern different subsets of the electorates. We also excluded two elections where one of the candidates was accused of irregularities and later disqualified.

Prior to the elections, students in both schools were informed about the survey through a school-wide email message, and through posters and fliers distributed widely throughout campus. Possible prizes from answering the survey were advertised: respondents from each school would take part in a lottery with iPods and \$20 gift certificates at Barnes & Noble awarded to the winners.¹⁰

3.2 The Data.

Out of a total of 1161 GS students, 476 voted in the GS elections, and of these 297 answered our survey; in the College, 2057 voted out of a potential electorate of 4073, and 644 answered the survey. After eliminating CC questionnaires that were either unreadable or unmatchable to actual voters, we cleaned the data according to the following criteria: (1) we assigned a score of 0 to any election in which a respondent abstained, a plausible option logically, and the only one available since we could not match the score to a voting choice; (2) we assigned a score of 1 to any election in which a respondent voted but left the ranking blank, and (3) we eliminated from the sample respondents who stated that they would cast the bonus vote in an election in which they in fact abstained. Because the CC questionnaires were on paper, they were missing automatic completeness checks that the electronic program forced instead on the GS students. In particular, nine CC respondents did not indicate where they would have cast the bonus vote. But choosing not to use the bonus vote could be a legitimate choice, and we left these respondents in the sample. After cleaning the data, we were left with 276 responses in the GS sample and 502 in the CC sample, or a valid response rate among voters of 58 percent in GS and 24 percent in CC.¹¹

3.3 The distributions of the scores.

Figure 1 reports the distributions of the scores assigned by respondents to each election, for both GS (Figure 1a), and CC (Figure 1b). Scores are labeled positive or negative, according to which of the two candidates the respondent voted for; for ease of reading, in each election we assign positive scores to the candidate commanding a majority in our sample. Consider for example the election for

¹⁰The prizes were 2 iPods and 8 gift certificates for the CC lottery (with a potential electorate of 4,073); and 1 iPod and 5 gift certificates for the GS lottery (with a potential electorate of 1,161 students).

¹¹Participation in the survey was voluntary and we cannot expect our samples to be representative of the population. For the electorate as a whole, we have data on the number of votes cast for either candidate in each election, and the number of abstentions. Table A1 in the Appendix tests the hypothesis that in each school both the rest of the electorate and our sample are random draws from a common population. When we look at abstention rates, the hypothesis is rejected at the 5 percent confidence level for all GS elections and for two of the four CC elections — predictably individuals who answered our questionnaires have significantly lower abstentions rates. When we look at support for the winning candidate among voters, the hypothesis is rejected in the President election in the GS sample and in the Academic Affairs election in the CC data — both elections were much closer in our sample than in the full electorate, a surprising finding for which we have no explanation.

Board President in the GS sample, at the top of Figure 1a. The scores assigned by Susannah’s supporters are plotted on the positive axis; those of Liz’s supporters on the negative axis, while the cell at zero, in black, reports abstentions. The histogram tells us that 23 of Liz’s supporters and 31 of Susannah’s assigned to the election a score of 10; 11 of Liz’s supporters and 13 of Susannah’s assigned it a score of 9, etc. The box next to the histogram summarizes the main data concerning this election: in our sample of 276 students, 142 voters supported Susannah, 129 supported Liz, and 5 abstained. Among our respondents, 208 stated that they would have cast the bonus vote in this election. Attributing the bonus votes according to each respondent’s actual candidate choice, Susannah again wins a majority of the bonus votes (110 against 98). Susannah did in fact win the election among all voters, with 241 votes in favor, versus 188 for Liz and 47 abstentions. The box reports also the average score assigned to the election by the supporters of each candidate: the average score is slightly higher among Susannah’s supporters (6.7 versus 6.3 for Liz’s supporters).

In the GS sample, the election for President was overwhelmingly the most salient: abstentions are less than one tenth of the second best attended election (International), and average scores are about seventy percent higher than in the other elections. Two of the three elections (President and Alumni Affairs) were close elections, and in these elections not only is the extent of support similar between the two candidates, but the distributions of voters’ scores are also approximately symmetrical across the two sides. The third election (International) was quite lopsided. If the elections had been held with bonus votes, none of the outcomes would have changed. In the President and in the International election, the majority winner commands a majority of the bonus votes; it does not in Alumni Affairs, but the difference is small (2 votes), and cannot override the difference in regular votes.

In the CC sample, there is again one election that voters in our sample considered most salient (Executive Board), with both low abstentions and higher average scores, but the difference with respect to the others is less pronounced than in GS. The Senate-Two Year election was a landslide, but the others, and in particular Senate-One Year and Academic Affairs were close elections. The Academic Affairs election, with scores distributed according to the histogram at the bottom of Figure 1b, is particularly interesting. Here Alidad won a majority of the votes (230 versus 195 for Ehizoje), but the average score is higher among Ehizoje’s supporters (6.1 versus 4.9 for Alidad). The distribution is not symmetric, with the majority of Ehizoje’s supporters attributing high importance to the outcome, while the scores given by Alidad’s supporters are concentrated in the middle range. Storable votes are designed to address situations of this type, and not surprisingly Ehizoje receives about 25 percent more bonus votes than Alidad (38 to 30), although the difference of 8 votes would not have been enough to counter the majority advantage. In the CC sample too bonus votes would not have changed any of the outcomes.

The scores reflect the importance attached by respondents to the different elections. We interpret them as measures of intensity of preferences. More precisely, and in line with the theoretical model, we read the scores as proxies

of the differential utility that each respondent attaches to having the preferred candidate win that specific election, as opposed to the opponent. The important information is the relative score assigned by each respondent to different contests: it is preferable to have one’s favorite candidate win an election rated as a 4 than an election rated as a 2. To give a measurable meaning to intensity of preferences, we must interpret the scores as not only ordinal but cardinal values: we use the simplest linear mapping, so that winning an election rated as a 4 is not only preferable but twice as valuable as winning an election rated as a 2. The distributions of the scores in Figure 1 are then the empirical counterpart of the distributions of preferences discussed in the theory.

If the scores measure intensity of preferences, we can construct measures of aggregate welfare in our samples. The most immediate utilitarian measure simply sums all scores over supporters of each candidate and calls *efficient* in any election the victory of the candidate whose supporters have higher aggregate scores. For example, in the Board President election in GS, Susannah’s supporters are both more numerous and have higher average score, guaranteeing a higher aggregate score (954, compared to 810 for Liz’s supporters), leading to the conclusion that efficiency favored Susannah’s victory. In fact, in the GS sample, with one landslide election and two closed elections with symmetrical distributions of scores, efficiency always favored the victory of the candidate preferred by a majority of respondents. In the Academic Affairs election in the CC sample, on the other hand, the conclusion is reversed. The aggregate score of Alidad’s supporters is 1132, against 1181 for Ehizoje’s supporters: contrary to the majority voting result, efficiency favored, marginally, a victory by Ehizoje. In all other CC elections, efficiency agrees with the majority selection.

As a welfare measure, the sum of the reported scores has the important advantage of being closest to the questionnaire. The difficulty is that in general, evaluated over all elections, it will give different weights to different voters, reflecting differences in the total scores that each of them has assigned. An alternative is to construct and sum *normalized* scores: scores obtained by constraining all individuals to the same total score. Specifically, if the number of elections is K and s_{ik}^R is the reported score assigned by individual i to election k , normalized scores s_{ik} equal $s_{ik}^R / \sum_{r=1}^K s_{ik}^R$, and for each individual sum up to 1 and have mean $1/K$. A normalized utilitarian measure of welfare is then the sum of the normalized scores over supporters of each candidate.

The histograms describing the distributions of normalized scores are reported in the Appendix. In Figure 2 we summarize the properties of our data, using both reported and normalized scores. The figure reports margins in favor of the majority, in each election, in terms of the number of votes, aggregate, and average scores.¹² Figure 2a is based on reported scores, as was Figure 1, while Figure 2b uses normalized scores.

¹²If \mathbf{M} is the size of the majority, and \mathbf{m} the size of the minority, the margin of victory among voters is simply: $(\mathbf{M} - \mathbf{m})/n$. Using s to denote the reported score in Figure 2a and the normalized score in Figure 2b, the aggregate score margin in favor of the majority is calculated as: $(\sum_{i \in M} s_i - \sum_{i \in m} s_i) / \sum_{i=1}^n s_i$, and the average score margin is: $(\hat{s}_M - \hat{s}_m) / (\hat{s}_M + \hat{s}_m)$ where $\hat{s}_M = (\sum_{i \in M} s_i) / \mathbf{M}$ and $\hat{s}_m = (\sum_{i \in m} s_i) / \mathbf{m}$.

The two sets of figures are not identical, suggesting that the normalization does play some role: because it accounts for each individual’s overall scoring pattern, the normalization in general rescales differently scores from different individuals, and thus changes both average and total scores. However, the figures make clear that the qualitative conclusions are robust: if efficiency is measured by higher aggregate scores, the efficient outcome in each election is unchanged whether we refer to reported or normalized scores — the dark gray columns always have equal sign in Figures 2a and 2b. In the GS sample, efficiency always supports the majority choice. In the CC sample, efficiency supports the majority side on three of the four elections, but the majority’s total score margin is slightly negative, albeit with a margin so small as to be undetectable in the figure, in the fourth (Academic Affairs) where the larger size of the majority is countered by the stronger intensity of preferences of the minority. Again, the observation holds for both sets of scores.

4 The bonus vote decision

The analysis of the bonus vote choice is disciplined by the theory, but our data do not allow us to test the theory precisely. The theoretical equilibrium strategy depends on the *common knowledge* of the distributions of valuations *in the electorate at large*, because it is the choice of the electorate as a whole that any single voter is trying to influence. The distributions of valuations we construct from the survey refer to a sample only — the respondents of the survey — and we have no basis for extrapolating the shape of the distributions to the electorate as a whole. Nor do we have any evidence that preferences were common knowledge. In fact we have some evidence to the contrary. As mentioned earlier, in the GS questionnaire, electronic and thus faster, we added a question about expected election outcomes. In the election for President, more than 80 percent of the students filling the questionnaire answered the question, but about half of them predicted the wrong winner, and a quarter predicted that she would win by a large margin; in the other two elections, more than half of those filling the questionnaires chose not to answer the question, in line with the large abstention rate, but among those who did respond, a majority predicted the wrong winner in the Alumni Affairs election, half of them by a large margin, and one third did so in the International election.¹³ These answers may reflect something

¹³The exact numbers are the following. For the President election, 50 of the 276 voters in our sample declined to answer the question; 107 predicted that Liz would win, and of these 58 predicted a margin of victory larger than 10 percent; 119 predicted that Susannah would win, and of these 52 predicted a margin of victory larger than 10 percent (in fact, Susannah won with a margin of 12 percent). For Alumni Affairs, 157 of the 276 voters in our sample declined to answer the question; 72 predicted that Bob would win, and of these 36 predicted a margin of victory larger than 10 percent; 47 predicted that Maria would win, and of these 17 predicted a margin of victory larger than 10 percent (Maria won with a margin of 2 percent). Finally, for International Representative, 146 of the 276 voters in our sample declined to answer the question; 91 predicted that Makiko would win, and of these 36 predicted a margin of victory larger than 10 percent; 39 predicted that Liron would win, and of these 14 predicted a margin of victory larger than 10 percent (Makiko won with a margin of 28 percent).

other than rational calculations of expected outcomes and probably should not be taken too seriously, but cannot be read as support for common knowledge of the full distributions of valuations. In addition, the survey question did not state explicitly that in the thought experiment all other voters would also be casting their bonus vote. Our first priority was to keep the questionnaire as simple as possible, and we knew that, lacking the distributions of valuations for the electorate, we would not be able to test equilibrium behavior. In practice, given the poor information about others' preferences, we doubt that answers would have been different, but we cannot rule it out. Our exploration of the stated use of the bonus vote will need to be mostly descriptive.

In the GS sample, 89 percent of the voters cast their bonus vote in an election to which they assigned their highest score, and 81 percent did so in the CC sample. If voters knew little about all elections and were aware of their lack of information, this would be both the simplest and the optimal strategy. Which other criteria influenced their choice? Figures 3a for GS and 3b for CC show the relevant data, election by election. Each diagram reports, for all respondents who said they would cast the bonus vote on that specific election, the reported score assigned to that election, on the vertical axis, and to the highest ranked of the other elections, on the horizontal axis. If the election selected for the bonus vote is the highest score election, then the respondent is indicated by a point above the 45 degree line; if not, by a point below the 45 degree line.¹⁴

The most salient election in each data set (President in GS and Executive Board in CC) received the great majority of the bonus votes, and was also the election most respondents ranked as most important to them. In all elections, some bonus votes were cast by respondents who ranked a different contest higher. The number of such bonus votes is relatively small in both the GS President election (8 percent) and the CC Executive Board election (12 percent), but less so, in relative terms, in the other contests, reaching 50 percent in the Senate-One Year election in CC. Over the full electorate — and it is the full electorate that a voter would consider when deciding where to cast the bonus vote — Senate-One Year was the closest election in CC. If voters rationally weigh the probability of being pivotal, then we would expect the pattern seen in the CC sample. In the GS data set, on the other hand, the bonus vote choice does not seem to correlate in any transparent manner with the realized margin of victory: the closest election was Alumni Affairs, with a margin of 2 percent, but the fraction of voters who cast the bonus vote in that election while ranking a different one higher is smaller than the fraction in the International election, the most lopsided, with a margin of victory of 28 percent. The observation is surprising, but a natural explanation is the probable lack of information on the part of the voters. At least in GS, we doubt that voters' information was precise enough to allow sophisticated strategic considerations.

A simple statistical model provides a concise description of the patterns we see in the data.

¹⁴The only relevant information in the figure is the ordinal ranking of the elections, and reported scores show the difference across elections more clearly. Normalized scores yield a rescaled but otherwise identical figure.

4.1 A Simple Statistical Model

The model describes respondents' bonus vote choices through a set of simple behavioral criteria, each of which specifies how to cast the bonus vote. Each respondent's answer, given his or her scores, can then be written in terms of the probability of following the different criteria. We posit four mutually exclusive criteria: (1) cast the bonus vote in the election with highest score; (2) cast the bonus vote in the most salient election; (3) cast the bonus vote in the closest election; (4) some other criterion we ignore, and such that the choice appears to us fully random. We suppose that each criterion is followed with some probability, which we call $pmax$ for criterion 1, $psal$ for criterion 2, $pclose$ for criterion 3 and $prand$ for criterion 4, and we describe a respondent's choice through these probabilities. For example, consider a GS voter whose highest score is on the President election, and who indicates that he would cast the bonus vote on that election. Under our model, this behavior occurs with probability $pmax + psal + (1/3)prand$. If the voter assigns the highest score also to a second election, then the probability of the observed behavior becomes $(1/2)pmax + psal + (1/3)prand$, and correspondingly for the other cases. Assuming that respondents' choices are independent, the likelihood of observing the data set is simply the product of the probabilities of all individual choices. The probabilities $pmax$, $psal$, $pclose$, and $prand$ can then be estimated immediately through maximum likelihood or Bayesian methods.

We have estimated the probabilities on the two data sets separately, and in both cases maximum likelihood and Bayesian estimation yield identical results, summarized in Table 1, with standard errors are in parentheses.¹⁵

	GS	CC
$pmax$	0.55 (0.06)	0.52 (0.04)
$psal$	0.34 (0.06)	0.30 (0.04)
$pclose$	0.01 (0.01)	0.04 (0.01)

Table 1: Population frequencies of three behavioral criteria as estimated from a simple statistical model of observed bonus vote choices (with standard errors in parentheses).

In both data sets and with both estimation methods, $psal$, the probability of casting the bonus vote on the most salient election, is relatively large and significantly different from zero. At the same time, $pclose$, the probability of casting the bonus vote on the election with smallest margin of victory, is always small and in the GS sample insignificant. Both results are theoretically puzzling: the intensity of one's preferences should be fully captured by the score, and,

¹⁵In conducting the Bayesian inference, we assigned uniform prior distributions to each of $pmax$, $psal$ and $pclose$ and used the software Bugs (Spiegelhalter et al., 1994, 2002). The Bayesian approach works well here because with moderate sample sizes, maximum likelihood estimates can be at the boundary of parameter space (as discussed in general terms in chapter 4 of Gelman et al., 2004).

as mentioned, the use of the bonus vote should be higher in close elections. However, the two most salient elections were also reasonably close elections, the second closest in both data sets, and attracted much more debate and attention prior to voting than any of the others. It is possible that voters chose them disproportionately not because they were salient, but because in fact they knew them to be close contests, and had little information about the other elections.

The statistical model we have posited is extremely simple. In particular, the probabilities with which the different criteria are followed are treated as exogenous parameters constant across voters, as opposed to being the expression of each voter's strategies, dependent on the voter's full set of scores, as they would in a strategic model. We interpret the estimates in Table 1 not as test of a theoretical model, but as synthetic description of the data. We use this description as one of four plausible behavioral rules to check the robustness of the bootstrapping exercise that forms the rest of this paper.

5 Estimating the Probable Impact of the Bonus Vote

As shown in Figure 1 and stated earlier, in our samples no outcome would have changed with the addition of the bonus vote, given how respondents stated they would cast it. But this is not very informative: because the bonus vote links the outcomes of the bundle of elections over which the voters can choose to use it, we have only two independent data points, one for each school, too few to form an estimate of the bonus vote's probable impact. To obtain such an estimate, ideally we would want to replicate the same elections many times, with many different electorates whose preferences are all drawn from the same underlying distribution. We cannot rerun the elections, but we can approximate such iterations by bootstrapping our data.¹⁶ The objective is to estimate the impact of the bonus vote in a population for which our samples are representative.

The maintained assumption of the bootstrapping exercise is that preferences are independent across individuals, although not necessarily across elections for a single individual. We sample with replacement n individuals from our data sets, where $n = 276$ for GS and $n = 502$ for CC: for each individual, we sample the scores assigned to all elections and the choice of where to cast the bonus vote. We then replicate this procedure 10,000 times, where each replication generates a distribution of preferences over each of the three elections in GS and four elections in CC, and a bonus vote choice for each voter. Finally, for each replication, we calculate outcomes and measures of welfare if the elections were held with simple majority voting, and with four alternative rules governing the use of the bonus vote. Rule A is closest to the data: we use the bonus vote choice actually made by the individuals sampled in the bootstrapping; rule B applies the statistical model described above to each bootstrap sample of preferences;

¹⁶The classical references are Efron, 1979, and Efron and Tibshirani, 1993. For a recent treatment, see Davidson and MacKinnon, 2006.

rule C states that every individual casts the bonus vote in the election with highest score (randomizing with equal probability if more than one election has the highest score); rule D replicates rule C but excludes the two most lopsided elections in each sample, International in GS and Senate-Two year in CC. Each of the four rules has some argument in its support, but it is the consistency of the results across rules that gives us confidence in their robustness.

The bonus vote can give weight to the intensity of preferences, and thus, given the distributions of preferences in our data, the hypothesis is that with high probability storable votes should be equivalent to majority voting in the GS elections, but have more impact and better welfare properties in the CC elections. We have verified that the results are not sensitive to using either reported or normalized scores, as is to be expected given Figure 2. In what follows we describe the analysis based on normalized scores which seem theoretically somewhat superior.

5.1 The Frequency of Minority Victories

The first question is the frequency with which, using the bonus vote, at least one of the elections in each set is won by the minority candidate. As shown in Figure 4, in the GS bootstrap samples the frequency is stable across the different rules at around 15 percent. In the CC samples, the frequency increases monotonically as we move from rule A to rule D: from 11 to 15 to 23 to 29 percent.

The aggregate number, appropriate because of the linkage across elections imposed by the bonus vote, is the result of different impacts on the specific elections. Tables 2a for GS and 2b for CC report the frequency with which any individual election is won by the minority, for each bonus vote rule. The last column reports the aggregate frequency depicted in the figure. (In all cases, we assigned a weight of 1/2 to any sample where either simple majority or the bonus vote resulted in a tie).

Table 2a: GS

	President	Alumni	International	Aggregate
Rule A	5%	11%	0	16%
Rule B	6%	9%	0	14%
Rule C	8%	9%	0	16%
Rule D	7%	8%	0	15%

Table 2b: CC

	Exec Board	Senate-Two	Senate-One	Academic	Aggregate
Rule A	0	0	2%	9%	11%
Rule B	0	0	1%	14%	15%
Rule C	0	0	1%	23%	23%
Rule D	0	0	1%	28%	29%

Table 2: Frequency of minority victories in 10,000 bootstrap samples, based on four different assignment rules for each of the elections in our study.

Predictably, the numbers reflect the distributions of preferences over each election. The most lopsided elections (International Representative in GS, and Executive Board and Senate-Two Year in CC) are never won by the minority candidate. The election most susceptible to the impact of the bonus vote is Academic Affairs in CC; moving from rule A to rule D the number of bonus votes cast on Academic Affairs progressively increases; so does the frequency with which the minority candidate wins, and so does the aggregate frequency.¹⁷ In the GS samples, the impact of the bonus vote reflects mostly how close the different contests are. In the CC samples, again the bonus vote affects exclusively the two closest elections, but the higher intensity of the minority supporters in Academic Affairs plays a clear additional role in tilting the results in the minority’s favor when the bonus vote is available.

We used 10,000 bootstrap samples to guarantee that the experimental error (the binomial standard error of the observed frequencies due to the finite size of the sample) is negligible. As a test, we recalculated the frequencies in tables 2a and 2b with 20,000 independent bootstrap samples, confirming that all frequencies changed by less than half of one percent.

Our first conclusion then is that in both sets of data and for all rules, although particularly in CC for rules C and D, the bonus vote allows the minority to win with substantial probability.

5.2 The Impact of Minority Victories on Aggregate Welfare

How costly are minority victories in terms of aggregate welfare? Perhaps surprisingly, our second result is that minority victories not only are not costly, but in fact typically come with aggregate welfare gains.

In each election, realized welfare is defined as the sum of all (normalized) scores of all voters who supported the winner. With the bonus vote, the relevant unit is again the full set of elections in each school, for each bootstrap sample. More precisely, if M_k is the set of voters whose favorite candidate commands a majority of regular votes in election k , then welfare with majority voting, W_M is defined as $W_M = \sum_k \sum_{i \in M_k} s_{ik}$ where s_{ik} is the score assigned to election k by voter i . Similarly, welfare with storable votes W_{SV} is defined as $W_{SV} = \sum_k \sum_{i \in \mathcal{M}_k} s_{ik}$ where \mathcal{M}_k is the set of voters whose favorite candidate commands a majority of all votes, including bonus votes, in elections k . We select all bootstrap samples where the minority wins at least one election: in each set of elections there exists at least one k for which $\mathcal{M}_k \neq M_k$.¹⁸ Among those samples only, the upper panel of Figure 5 shows the frequency with which $W_{SV} > W_M$, and the lower panel shows the mean percentage welfare change

¹⁷In CC, rule B assigns some bonus votes to the closest election. Taking into account the bonus votes, Senate-1 is the closest election in 26 percent of the bootstrap samples, and Academic Affairs in 72 percent. The remaining 2 percent of bootstrap samples are ambiguous — either election can be closest, depending on how the bonus votes are cast. The ambiguity does not affect the frequency reported in Figure 4.

¹⁸In the case of ties, we assign a probability of 50 percent to the victory of either side.

(mean $(W_{SV} - W_M)/W_M$).

For all four rules in the CC data and for three of the four in the GS data the mean effect on welfare is positive. In the CC data, the result is strong: among all bootstrap samples where the minority wins at least one election, an increase in aggregate welfare occurs from a minimum of 86 percent of the time (rule A) to a maximum of 96 percent (rule C). The minority must be winning when its preferences are more intense than the majority preferences. In the GS data, the mean welfare change is always smaller than in CC, and in one case, when it is smallest, the sign is negative (rule A); with rules B, C and D the expected impact is positive, but the frequency with which aggregate welfare rises reaches a maximum at 72 percent (rule D).

In addition to generating the summary results reproduced in Figure 5, the bootstrap exercise provides us with the full distribution of the impact of the bonus vote on aggregate welfare across all bootstrap samples. Figures 6a for GS and 6b for CC present the histograms of the percentage difference in aggregate welfare between each of the four bonus vote rules and majority voting, in all samples where, using the bonus vote, the minority wins at least one election in each set. The histograms show clearly the concentration of the probability mass around the positive mean welfare change in the CC samples, while the distribution is more dispersed in the case of GS. The higher variability of the GS results is consistent with approximately symmetric preferences across the two sides in elections affected by the bonus vote. Summary measures of dispersion — standard errors and bootstrap confidence intervals — are obtained easily; we omit them here, preferring to present the full distributions. Because our focus is on samples where the minority wins at least one election, the number of samples analyzed here is between 1,400 and 1,600 for GS and between 1,000 and 3,000 in CC, depending on the bonus vote rule (see Table 2). We have compared the results reported here to those obtained from a different set of 10,000 bootstrap samples, confirming that the experimental error remains negligible.

5.3 Inequality

Minorities should win occasionally not only because aggregate efficiency is higher when intensity of preferences is recognized, but also because preferences can be correlated across elections, and individuals who find themselves repeatedly on the minority side will feel unrepresented by the political system. If it comes at no cost to aggregate welfare, a more equal distribution of decision-making influence seems desirable in itself. Our data allow us to construct the full distribution of ex post utility across voters in each of the bootstrap samples, where individual utility is defined as the sum of the voter’s scores over all elections won by the voter’s preferred candidate. Formally, recalling that \mathcal{M}_k is the set of voters whose favorite candidate commands a majority of all votes, including bonus votes, in elections k , ex post utility of voter i , U_i , is defined as $U_i = \sum_{k|i \in \mathcal{M}_k} s_{ik}$. The distribution across voters thus captures the different frequencies with which each voter’s preferred candidate wins, weighted by the relative importance the voter assigns to each election. If instead of summing over all elections for which

i belongs to \mathcal{M}_k , we sum over all elections for which i belongs to M_k — the set of voters whose favorite candidate commands a majority of regular votes in election k — we obtain for comparison ex post utility with majority voting. We can then calculate the impact of the bonus vote on voters’ ex post inequality.

Average realized utility distributions are shown in Figure 7. The bins on the horizontal axis are ordered in increasing levels of utilities, each bin corresponding to a 5 percent range, from lowest (0) to highest (1). The height of each bin is the fraction of voters with utility falling into the bin, averaged over all samples where, with the bonus vote, the minority wins at least one election. The black line corresponds to the bonus vote (rule A), and the gray line to simple majority. Here rule A is fully representative of the results obtained with the other bonus vote rules. The higher polarization of the GS distributions reflects the smaller number of contests (3, as opposed to 4 in CC) and the overwhelming priority assigned to a single election (President). The impact of the bonus vote on the distribution is hardly detectable in the GS samples but is clear in the CC sample, where the frequency of realizations at the two ends of the distribution is reduced in favor of larger mass in the center.

Figure 8 reports the impact of the bonus vote on voters’ inequality as measured by the Gini coefficient. In the CC data, inequality declines unambiguously: over all bootstrap samples where the bonus vote allows the minority to win at least one election, inequality declines from a minimum of 90 percent of the samples (rule A) to a maximum of 98 percent (rule D). Across all bonus vote rules, the mean percentage decrease in the Gini coefficient is 12–13 percent. In the GS data, on the other hand, the bonus vote’s effect on inequality is of inconsistent sign and small magnitude. The Gini coefficient declines barely more than 50 percent of the times with rules B, C and D, and barely less with rule A. The effects are quantitatively small: the mean percentage change in the coefficient is an increase of 2 percent.

The different results reflect the different roles of the bonus vote in the two data sets. In the GS data, the elections affected by the bonus vote (President and Alumni Affairs) are particularly close, with little difference in volume and intensity of support between the two candidates. The side benefiting from the bonus vote is not consistent across bootstrap samples, and thus the average impact of the bonus vote on equality of representation is small and similarly non consistent. In the CC data, on the other hand, the main role of the bonus vote is to overthrow the majority victory in the Academic Affairs election, on the strength of the more intense minority preferences. The distribution of preferences is sufficiently asymmetric to ensure that in the bootstrap samples the bonus vote consistently helps the same candidate. And because the minority in the Academic Affairs election is also disproportionately on the minority side in the other elections, the final result is a positive, sizable improvement in equality.¹⁹

¹⁹The precise numbers reflected in Figure 8 are reported in the Appendix, together with 95 percent confidence intervals. An alternative measure of inequality — the ratio of the average utility of individuals at the bottom 20 percent of the utility distribution to average utility — yields identical conclusions, and is also reported in the Appendix.

6 Conclusions


Storable votes are a simple mechanism with the potential to improve fairness without costs, in fact with some gains, to efficiency. In our data, they perform well: when the margins of victory are small, they allow the minority to win occasionally, with no substantive effects on welfare and inequality if the two sides have preferences of similar intensity, but with improvements in both if minority preferences are particularly intense. Storable votes should be tested in higher stake contests, particularly in referendums, for which they are expressly designed. Using a new voting system in real elections is, appropriately, difficult to do, and assembling an experimental data set rich enough to evaluate the system's performance statistically is probably impossible. Survey methods, coupled with bootstrap resampling, are a more practical route. One of the goals of this work is to suggest that the bootstrap methodology is ready to be exploited, easily and cheaply.

7 References

1. Bowler, S., T. Donovan, and D. Brockington, 2003, *Electoral Reform and Minority Representation: Local Experiments with Alternative Elections*, Ohio State University Press, Columbus.
2. Casella, A., 2005, "Storable Votes", *Games and Economic Behavior*, 51, 391-419.
3. Casella, A., A. Gelman and T. Palfrey, 2006, "An Experimental Study of Storable Votes", *Games and Economic Behavior*, 57, 123-154.
4. Casella, A. and A. Gelman, forthcoming, "A Simple Scheme to Improve the Efficiency of Referenda", *Journal of Public Economics*.
5. Chwe, M., 1999, "Minority Voting Rights Can Maximize Majority Welfare", *American Political Science Review*, 93, 85-97.
6. Cox, G., 1990, "Centripetal and Centrifugal Incentives in Electoral Systems", *American Journal of Political Science*, 34, 903-935.
7. Davidson, R. and J. G. MacKinnon, 2006, "Bootstrap Methods in Econometrics", in K. Patterson and T. C. Mills (eds.), *Palgrave Handbook of Econometrics: Vol. 1 Econometric Theory*, Palgrave Macmillan Ltd, Houndmills, Basingstoke, Hampshire RG61 6XS, UK.
8. Efron, B., 1979, "Bootstrap Methods: Another Look at the Jackknife", *Annals of Statistics*, 7, 1-26.
9. Efron B. and R. J. Tibshirani, 1993, *An Introduction to the Bootstrap*, Chapman and Hall: New York.


10. Gelman, A., Carlin J., Stern H. S. and Rubin D. B., 2004, *Bayesian Data Analysis* CRC Press, London.
11. Gerber, E. R., 1999, *The Populist Paradox: Interest Group Influence and the Promise of Direct Legislation*, Princeton University Press.
12. Gerber, E. R., R. B. Morton and T. A. Rietz, 1998, "Minority Representation in Multimember Districts", *American Political Science Review*, 92, 127-144.
13. Guinier, L., 1994, *The Tyranny of the Majority*, Free Press: New York.
14. Hortala-Vallve, R., 2005, "Qualitative Voting", unpublished, Oxford University, Nuffield College.
15. Issacharoff, S., P. Karlan and R. H. Pildes, 2002, *The Law of Democracy: Legal Structure and the Political Process*, Foundation Press: New York, 2nd edition.
16. Jackson, M. and H. Sonnenschein, 2007, "Linking Decisions", *Econometrica*, 75, 241-57.
17. Matsusaka, J. G., 2004, *For The Many or The Few: The Initiative Process, Public Policy, and American Democracy*, Chicago: University of Chicago Press.
18. Pildes, R. H. and K. A. Donoghue, 1995, "Cumulative Voting in the United States", *The University of Chicago Legal Forum*, 241-313.
19. Spiegelhalter, D., Thomas, A., Best, N., Gilks, W., and Lunn, D., 1994 and 2002, *BUGS: Bayesian inference using Gibbs sampling*. MRC Biostatistics Unit, Cambridge, England. www.mrc-bsu.cam.ac.uk/bugs

Figure 1a: GS




Susannah >0
Votes: Susannah 142; Liz 129; Abs 5
Bonus Votes: Susannah 110; Liz 98
Average Score: Susannah 6.7; Liz 6.3

ELECTION RESULTS
Susannah 241; Liz 188; Abs 47



Bob >0
Votes: Bob 106; Maria 97; Abs 73
Bonus Votes: Bob 17; Maria 19
Average Score: Bob 5; Maria 4.8


ELECTION RESULTS
Maria 151; Bob 144; Abs 181



Makiko >0
Votes: Makiko 136; Liron 73; Abs 67
Bonus Votes: Makiko 21; Liron 11
Average Score: Makiko 4.5; Liron 4.7

ELECTION RESULTS
Makiko 203; Liron 113; Abs 160

Figure 1b: CC



Open Columbia >0


Votes: Open Columbia 294; Evolution 199; Abs. 9

Bonus Votes: Open Columbia 205; Evolution 105

Average Score: Open Columbia 6.5; Evolution 5.9

ELECTION RESULTS

Open Columbia 1107; Evolution 819; Abs 131



Tiffany >0


Votes: Tiffany 324; Gerry 120; Abs 58

Bonus Votes: Tiffany 48 ; Gerry 9

Average Score: Tiffany 5.2; Gerry 4.3

ELECTION RESULTS

Tiffany 1233; Gerry 533; Abs 291



David >0


Votes: David 256; Yihe 208; Abs 38

Bonus Votes: David 29; Yihe 29

Average Score: David 4.5; Yihe 4.9

ELECTION RESULTS

David 1015; Yihe 809; Abs 233



Alidad >0

Votes: Alidad 230; Ehizoje 195; Abs. 77

Bonus Votes: Alidad 30 ; Ehizoje 38

Average Score: Alidad 4.9; Ehizoje 6.1

ELECTION RESULTS

Alidad 1060; Ehizoje 634; Abs 363

Figure 1: Histograms of recorded scores.

Figure 2a: Recorded scores




Figure 2b: Normalized scores




Figure 2. Margins in favor of the majority, per election. The first series (dotted) is the margin of victory among voters in our sample, ignoring the bonus vote. The second series (dark grey) is the aggregate score margin in favor of the majority. The third series (light grey) is the average intensity margin in favor of the majority.

Figure 3a: GS




Figure 3b: CC

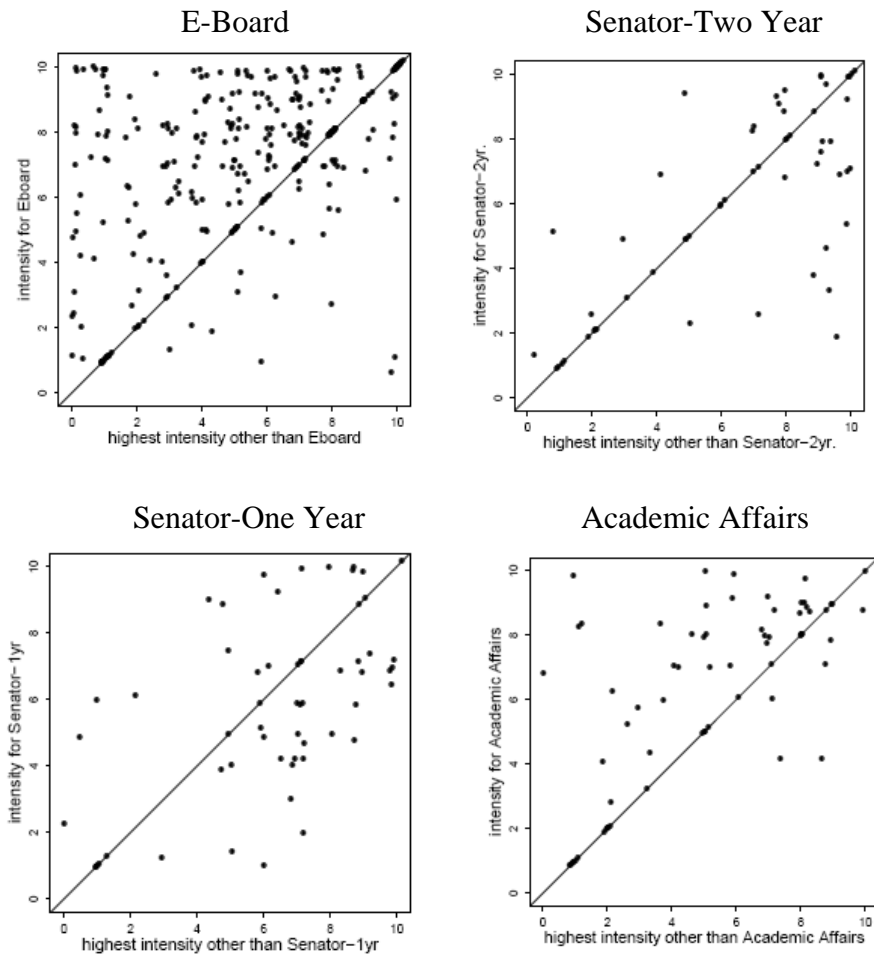


Figure 3: Relative raw score of the election receiving the bonus vote.




Figure 4. Percentage frequency with which at least one election in each set is won by the minority, out of 10,000 bootstrap samples, when voters cast their bonus vote according to each of four different rules.




Figure 5. Percentage frequency of aggregate welfare increases and mean percentage welfare change, relative to majority voting, when the minority wins at least one election.




Figure 6: Histograms of percentage welfare differences between the bonus vote scheme and majority voting in all samples where, with the bonus vote, the minority wins at least one election. The vertical axis is the percentage of such samples falling into each bin.

Figure 7a: GS




Figure 7b: CC




Figure 7: Realized utility distribution, simple majority (grey) and bonus vote, rule A, averaged over all samples in which the minority wins at least one election.




Figure 8. Percentage frequency of more equalitarian distributions of realized utility, and mean percentage decline in the Gini coefficient, relative to majority voting.

Appendix

TABLE A1
Table A1a: GS

	Abstentions			Share of votes for winner		
	non-sample	sample	diff.	non-sample	sample	diff.
President	0.21	0.02	0.19*	0.63	0.53	0.10*
Alumni Affairs	0.54	0.26	0.28*	0.59	0.48	0.11
Int'l Repr.	0.46	0.23	0.24*	0.63	0.65	-0.02

Table A1b: CC

	Abstentions			Share of votes for winner		
	non-sample	sample	diff.	non-sample	sample	diff.
Exec Board	0.08	0.02	0.06*	0.57	0.60	-0.03
Senate-1	0.15	0.12	0.03	0.69	0.73	-0.04
Senate-2	0.13	0.08	0.05*	0.56	0.55	0.01
Academic Affs.	0.18	0.15	0.03	0.65	0.54	0.11*

Table A1: Comparison of abstention rates and margins of victory in and outside the samples. Both in the electorate as a whole and in our samples: (i) abstention rates are calculated for each election for students who voted at least once, and thus do not reflect the fraction of students who did not take part in voting at all; (ii) the share of votes for the winner is calculated among voters in each election, and thus ignores abstentions.

TABLE A2
MEASURES OF INEQUALITY WHEN OUTCOMES DIFFER

Table A2a: General Study

Gini Coefficient

	Maj	Bv A	Bv B	Bv C	Bv D
Mean	0.347	0.353	0.353	0.353	0.354
95% CI	(0.20, 0.40)	(0.21, 0.40)	(0.21, 0.40)	(0.22, 0.40)	(0.22, 0.39)
Prob bv lower (%)		47.9	52.8	58.2	58

Ratio of average utility of bottom 20% to sample average

	Maj	Bv A	Bv B	Bv C	Bv D
Mean (%)	9.08	7.89	7.44	7.14	6.78
95% CI	(0, 47.1)	(0, 45.0)	(0, 44.9)	(0, 43.7)	(0, 43.1)
Prob bv higher (%)		52.5	51.2	53.5	49.4

Table A2b: Columbia College

Gini Coefficient

	Maj	Bv A	Bv B	Bv C	Bv D
Mean	0.280	0.246	0.243	0.243	0.243
95% CI	(0.24, 0.30)	(0.22, 0.29)	(0.22, 0.28)	(0.22, 0.27)	(0.22, 0.27)
Prob bv lower (%)		90.5	95.3	97.1	97.7

Ratio of average utility of bottom 20% to sample average

	Maj	Bv A	Bv B	Bv C	Bv D
Mean (%)	27.27	33.88	34.42	34.42	34.38
95% CI	(22.5, 37.7)	(24.9, 41.8)	(25.9, 41.5)	(27.2, 41.0)	(27.4, 41.2)
Prob bv higher (%)		88.6	94.0	96.6	97.0

Table A2: Measures on inequality when the minority wins at least one election. Because we are focussing on bootstrap samples where majority voting and the bonus vote lead to different outcomes, to each bonus vote rule corresponds a slightly different data set. The values reported in the table for majority rule correspond to bootstrap samples where majority and Rule A lead to different outcomes (the values in the other cases are almost identical).

Thank you for choosing to participate in this experiment. By entering your number below, and attaching your post-it to the questionnaire you acknowledge that you have read, and understand the consent form (attached to the ballot box). Please enter your number here: _____

A lottery will be conducted among all participants who have followed the instructions of the experiment and written the number above legibly. The winner will receive an iPod.

This experiment asks you how much you care about the outcome of some of the elections.

Each election can be scored on a scale of 1-10, with 10 meaning "care very much" and 1 meaning "care not at all". In order to indicate a box, please cross out the relevant number with an X like this:



E-Board

1	2	3	4	5	6	7	8	9	10
---	---	---	---	---	---	---	---	---	----

Senate (1 year)

1	2	3	4	5	6	7	8	9	10
---	---	---	---	---	---	---	---	---	----

Senate (2 year)

1	2	3	4	5	6	7	8	9	10
---	---	---	---	---	---	---	---	---	----

Academic Affairs

1	2	3	4	5	6	7	8	9	10
---	---	---	---	---	---	---	---	---	----

Suppose now that, in addition to the regular votes you cast earlier, you had 1 additional vote to cast in favor of your candidate in one of the elections. You can choose any of these four elections as you see fit. Which one would you choose? (NOTE: This is for experimental purposes only. Your answer will **not** change the outcome of the actual elections in any way).

Please check the box under the election you choose. Please check **only one** box.

E-Board

Evolution
Open Columbia

Senate (1 year)

David Ali
Yihe (Eric) Wang

Senate (2 year)

Tiffany Davis
Gerry Rodriguez

Academic Affairs

Ehizoje Azeke
Alidad Damooei

Thank you for participating. And good luck with the lottery!

Figure A1: The written questionnaire submitted to CC students.

Figure A2a: GS




Figure A2b: CC

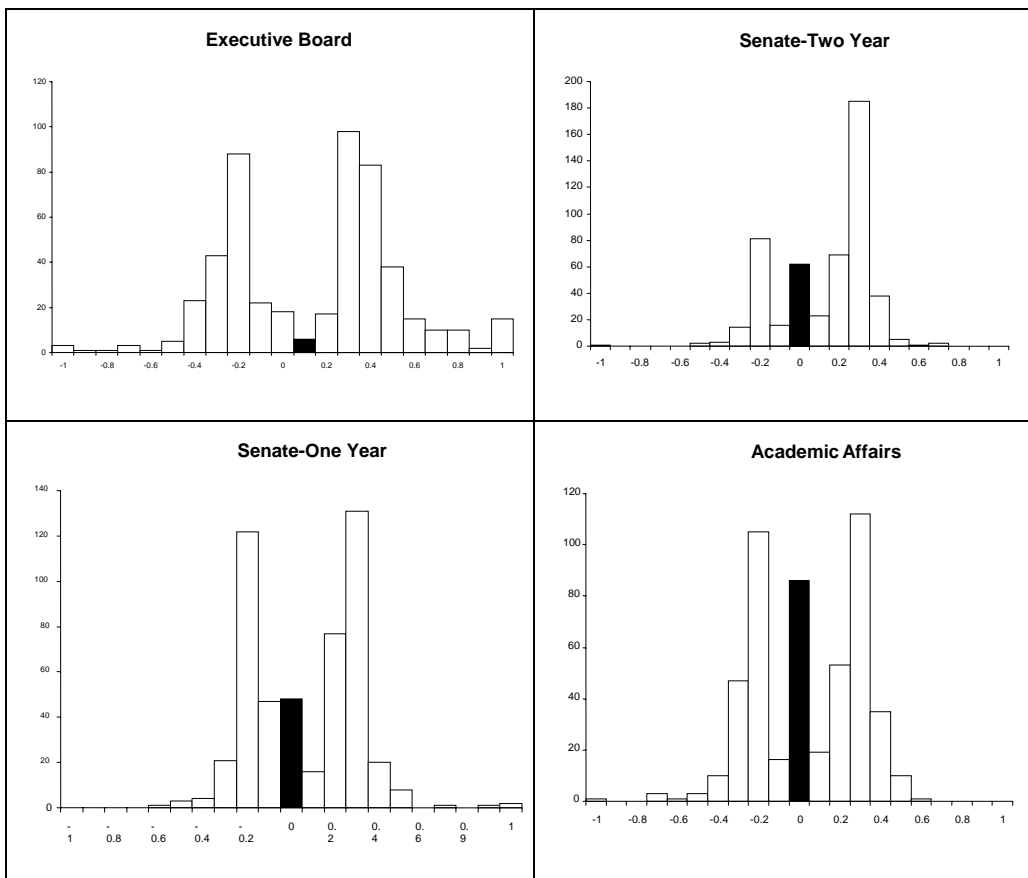


Figure A2: Histograms of normalized scores.