Storable Votes: Giving Voice to Minority Preferences without Sacrificing Efficiency*

Alessandra Casella**,

Thomas Palfrey*** and

Raymond Riezman****

The principle of majority rule is the foundation of democratic constitutions, but provides an immediate and fundamental challenge to the legitimacy of any government that the constitution empowers: the risk of excluding minority groups from representation. At least since Madison, Mill, and Tocqueville, political thinkers have argued that a necessary condition for the legitimacy of a democratic system is for no group with socially acceptable goals to be disenfranchised. In the history of constitutional law, ensuring fair representation to each group is seen as the crucial second step in the evolution of democratic institutions, after granting the franchise: once all individuals are guaranteed the right to participate in the political process, the problem remains how to assign appropriate weights to each group's political interest. The core of the difficulty is that the two goals seem inherently contradictory.

One remedy is recourse to the judiciary system: basic rights can be guaranteed in the fundamental laws of the country, and the courts can be appealed to when such rights are imperiled. But protecting a political minority when its rights are threatened does not address the subtler problem of ensuring that its preferences are sufficiently represented. For this, the cor-rect design of political institutions is required. In our work, we approach the problem from the perspective of voting theory, and propose a simple voting mechanism that, without violating the basic principle of "one-person one-vote", allows the minority to win occasionally. The mechanism is not based on super-majorities, avoiding the costs of inertia and ineffi-ciency they can entail, nor on geographical parti-tions, with the inevitable arbitrariness and instability of redistricting. In addition, although the mecha-nism's main property is its ability to protect minori-ties, and thus to increase fairness and legitimacy, it does so without sacrificing efficiency.

A simple example will make our words more trans-parent, but precision is important and we must begin with some definitions. We define a *minority* as a clearly identifiable group characterized by two fea-tures: first, a small relative size; second, preferences that are systematically different from the prefer-ences of the majority. Thus, a minority in our discus-sion is a *political* minority, which may, but need not, correspond to a minority according to racial, ethnic, religious or other types of non-political group iden-tity. In terms of political decisions, what matters are the coherent and idiosyncratic policy preferences of the group, independent of the source of its identity. Consider then the following example.

A polity comprised of 100 citizens has two (political) groups, with 55 members in Group A and 45 in group B. Three proposals are being considered. All citizens in group A have identical preferences and strictly prefer to pass all proposals; all citizens in group B have identical preferences and strictly pre-fer the status quo on all three issues. Thus, group B fits our definition of a minority. Suppose that the utility each citizen receives from each alternative is as given in Table 1, with the utility of the less pre-ferred option normalized to 0.

Table 1

Issue	U_A(pass)	U_A(sq)	U_B(pass)	U_B(sq)
1	3	0	0	1
2	2	0	0	2
3	1	0	0	3

Note that the intensity of preferences varies across the issues. That is, some issues are "more important" to one group than to the other group – issue 1 is important to group A but not to group B, and issue 3 is important to group B but not to group A.

Now consider what would happen with simple majority rule when issues are decided independently: since group A has a majority, all three proposals pass. Indeed, as long as preferences are perfectly correlated within groups, then even if there were a million different issues, group A would always have a majority on all issues, so the B citizens are effectively disenfranchised – the outcome is exactly the same as it would be in a political system where only A citizens were allowed to vote.

Why is this situation undesirable? First, a formally open franchise is meaningless if the outcomes are equivalent to what would arise if political access were denied to one of the groups. If one considers outcomes as well as process, then equity considerations demand that political minorities be able to win on at least some issues. Moreover, from a purely utilitarian standpoint, the outcomes described above are socially inefficient according to widely accepted welfare criteria. In our example, if each individual is treated equally and decisions are evaluated ex ante, before membership into the groups is known, the status quo should prevail on issue 3. Thus letting the majority prevail on all issues has costs both in terms of equity and in terms of ex ante efficiency: the equity problem stems from the existence of a smaller group whose members' preferences are *systematically opposed* to the members of the larger group; the efficiency problem stems from differences in the *intensity* of preferences of the two groups. This failure of simple majority rule is often referred to as "the tyranny of the majority".[1]

How can the failure be avoided, or at least mitigated? An immediate answer may be that, in practice, vote trading and logrolling schemes produce outcomes responsive to different intensities of preferences and thereby improve efficiency: members of one group could trade their vote on one issue in exchange for votes on other issues. But in our simple example, vote trading and logrolling by itself will not

change the outcomes: citizens in group A already win on all issues, so B members have nothing to trade. An explicit institution "re-enfranchising" the minority is necessary.

But note that this institution cannot be a supermajority or unanimity requirement for passing any proposal: it would result in maintaining the status quo in all issues, an outcome that is worse than simple majority voting on both equity and efficiency grounds. Any solution must deviate from issue-by-issue simple majority voting.

Consider then endowing every voter with an initial stock of votes, and rather than requiring voters to cast exactly one vote on each issue, allowing them to lump their votes together, casting "heavier" votes on some issues and "lighter" votes on other issues. It is this voting mechanism, called *storable votes*, that we study in our work. If decisions are made according to the majority of *votes* cast (as opposed to the majority of *voters*), storable votes allow the minority to win some of the time, and in particular, to win when its preferences are most intense. And because the majority generally holds more votes, it is in a position to overrule the minority if it cares to do so: the minority can win only those issues over which its strength of preferences is high and, at the same time, the majority's preference intensity is weak. But these are exactly the issues where the minority "should" win from an efficiency viewpoint: *the equity gains resulting from the possibility of occasional minority's victory need not come at a cost to aggregate efficiency.* Nor does the representation of minority's preferences come at the expense of the equal treatment of all voters: all individuals are granted the same number of votes and all votes count equally.[2] Thus the scheme need not be redesigned if the size, or the very existence, of the minority changes.

In fact, even without systematically opposed preferences, the use of storable votes can increase efficiency in symmetrical voting environments, where no systematic minority exists.[3] When a voter bunches his votes on his highest intensity issues, the probability of obtaining his desired outcome shifts away from decisions that matter little to that voter and towards decisions that matter more, with positive welfare consequences. The value of a storable votes mecha-

---

[1] Nothing fundamental depends on all citizens in a group having the same intensity of preferences on every issue, or even on the direction of preferences within the group being perfectly correlated. The same problem exists with imperfect correlation and heterogeneity of intensities within groups.

[2] Many variations that we have not studied are also possible, including the granting of different numbers of votes to different individuals.
[3] Casella (2005).

nism in the presences of minorities is even more compelling because, in addition to the efficiency gains, it addresses fundamental considerations of equity and legitimacy.

An existing voting system that resembles storable votes is cumulative voting, a mechanism used in multi-candidate elections. It grants each voter a budget of votes, with the proviso that the votes can be spread or concentrated on as many or as few of the candidates as the voter wishes. The winner is the candidate who receives the most votes. In the United States, cumulative voting is used commonly in corporate elections, with the explicit goal of making it possible for minority shareholders to elect members to the board of directors (Williams 1951). In political elections, cumulative voting was used from 1870 to 1980 to elect representatives to the state House in Illinois; has been advocated more generally for the protection of minority rights (Guinier 1994) and has been imposed by the courts to redress violations of fair representation in local elections (Issacharoff, Karlan and Pildes 2002). There is evidence – theoretical (Cox 1990), experimental (Gerber, Morton and Rietz 1998), and empirical (Pildes and Donoghue 1995, Bowler, Donovan and Brockington 2003) – that cumulative voting does indeed help minorities.[4] The storable votes mechanism differs from cumulative voting because it applies to a collection of independent binary decisions – a series of proposals, each of which can pass or fail – as opposed to a single multi-candidate election, but the motivation is similar.

Experimental evidence

The desirable properties of storable votes are features of the equilibrium of the resulting voting game: they emerge if every voter chooses the correct number of votes, given what he rationally expects others to do. But could the outcome be much worse if voters made mistakes? This is an appropriate concern here because the storable votes game is quite complex: voters need to trade-off the different probabilities of casting the pivotal vote along the full logical tree of possible scenarios, a task further complicated by coordination problems within the two groups, and multiple equilibria. If actual voters were confronted with the problem, what type of decisions would they make?

To answer this question, in 2004 and 2005 we ran a number of laboratory experiments at Caltech, UCLA and Princeton University where subjects recruited from campus were asked to vote over a sequence of elections in five-person committees, with storable votes. In each experimental session, a committee faced two consecutive proposals. Three members of the committee (the *majority* group) earned a monetary reward whenever a proposal failed; two others (the *minority* group) if it passed. The composition of the two groups was constant over the two proposals, mimicking a systematic minority that would always lose with simple majority voting.

Immediately before each proposal was voted on, each voter was assigned a random value between 1 and 100 which translated directly into the monetary reward, inducing a voter's intensity of preference, as in the theoretical model.[5] Although group membership was known to everyone in the group, the intensities were private information. We considered two treatments regarding the distribution of preference intensity within a group, representing stronger or weaker group cohesion.

In the first case, as in the example in Table 1, all members of a group had identical intensity and agreed not only on the preferred direction of the proposal, but also on priorities across the two proposals. We called this treatment *C*, as in "correlated intensities". In the second case, intensities were drawn independently, and members of a group agreed on the direction of preferences but not necessarily on the strength of their preferences. We called this treatment *B*, the "base" treatment. In both cases, intensities were independent across groups, so members of one group were uncertain about the intensities in the other group. Everyone was always informed of the statistical process by which intensities were assigned.

Each subject had one standard vote to cast over each proposal and a total of two bonus votes to spend as desired, either dividing them over the two proposals or cumulating both over a single proposal. The outcome was determined by majority rule, with ties broken randomly. After each round of two proposals, another round started with two new proposals and a new endowment of bonus votes, and intensities were reassigned. This was repeated for

---

[4] There is even a blog on cumulative voting: http://www.fairvote.org.

[5] See Smith (1976) for an explanation of induced value theory.

15 to 30 rounds. We ran 11 experimental sessions with a total of 167 students.[6]

In both treatments, theory predicts that minority members should always cumulate their bonus votes on a single proposal: on the first proposal if the intensity attached to it is at least 50, and on the second otherwise. The threshold is 50 because 50 is the expected intensity over the second proposal. The same is true for a majority member in treatment B. However, in treatment C the equilibrium has the majority collectively casting 5 votes on the first proposal if its intensity is below 50, and 7 otherwise. Since neither 5 nor 7 is divisible by 3, there is no simple symmetric individual strategy that produces this total: different majority voters must use different strategies, a difficult coordination problem. To see whether the coordination problem affected the results, we designed two variations of the C treatment: in one we allowed subjects to chat electronically with other members of their own group; in the second, we let a single subject represent an entire group, and thus rephrased the game as taking place between two voters only, with asymmetrical voting power and preferences.

The key intuitive property of rational voting behavior with storable votes is evident: the number of votes cast on a proposal increases monotonically in intensity. This applies to each individual voter in case B of our experiment, and to the group as a whole in case C. As we said earlier, this property is quite general and is the main reason why storable votes have good efficiency properties: by casting more votes when they care more, voters are more likely to have their way when it matters most.

Table 2 reports the equilibrium outcomes predicted by theory with the parameterization we used in the experiment. The first two rows report the equilibrium expected frequency of minority victories, and the fully efficient frequency, respectively. The third and fourth rows give the expected share of payoffs for minority versus majority members, in equilibrium and with full efficiency respectively.[7] The last two rows describe theoretical gain from storable votes over simply majority rule.[8]

---

[6] In each session, multiple committees were run in parallel, with random rematching of subjects into committees after each 2-proposal round.
[7] Full efficiency is defined here as deciding in favor of the group with higher total intensity. It guarantees to the committee as a whole the maximum aggregate payoff the committee could achieve.
[8] The percentage surplus measures are normalized so that full efficiency is 100 and the surplus from random committee decisions is 0.

Table 2

| Treatment	C	B
% min wins, sv	25	19
% min wins, efficiency	33	22
% (min/maj) payoff, sv	38	26
% (min/maj) payoff, efficiency	52	35
% surplus sv	60	71
% surplus majority voting	53	75

In treatment C in equilibrium the minority is expected to win one quarter of the time; this is less than efficiency recommends, but obviously much more than with simple majority voting (when, by definition, the minority always loses). Similarly, again in treatment C, in equilibrium a minority member is expected to have a payoff that is just below 40 percent of a majority member's payoff; efficiency recommends a larger share (just above 50 percent), but again, the payoff is zero with majority voting. In treatment C, storable votes are not only more responsive to minorities but sensibly more efficient; in treatment B, storable votes may come with some efficiency losses, as they do in this case, but the losses are small in magnitude.

How closely did laboratory behavior conform to these theoretical predictions? The following figure summarizes the results. The three panels match the rows in Table 2. The left panel reports the realized share of minority victories in the experiments (vertical axis), against the share that would have been observed, given the experimental intensity draws, if all subjects had played equilibrium strategies (horizontal axis). Each point corresponds to an experimental session; the light blue points refer to treatment B; the smaller purple points to treatment C, the larger purple points to treatment C with intra-group communication, and the yellow points to treatment C when a single subject represented an entire group. If all points were on the 45 degree line, observed outcomes would match the theory perfectly; distance from the 45 degree line represents deviations from the theory. The exact coordinates of the different points allow us to evaluate the results quantitatively, relative to Table 2, and to identify any treatment effects. The center panel reports the share of the aggregate minority payoff to the aggregate majority payoff, again plotting the laboratory outcomes (vertical axis) against equilibrium predictions (horizontal axis). The right panel graphs the total payoff in each experimental session against the total payoff under simple majority rule (i.e., when majority group
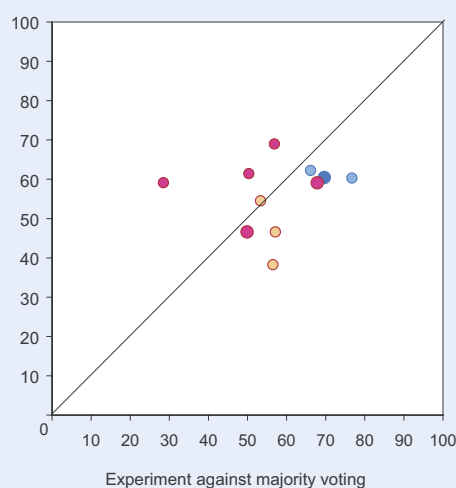
SHARE OF MINORITY VICTORIES

Experiment against equilibrium
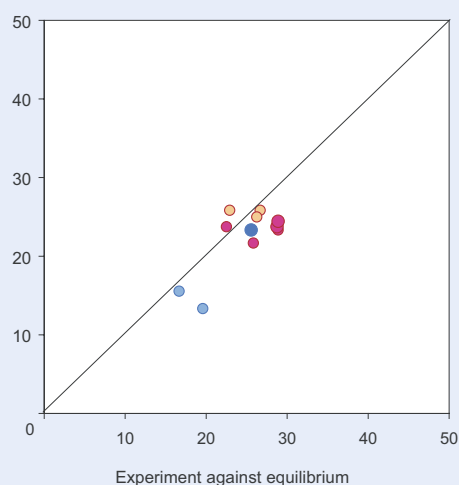
SHARE OF MINORITY PAYOFF

Experiment against equilibrium

SHARE OF SURPLUS

Experiment against majority voting
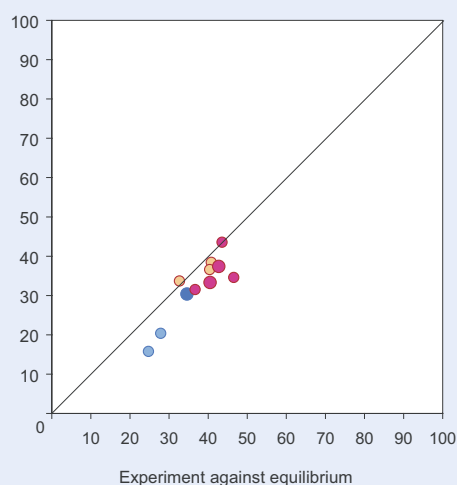
always wins), both expressed as percentage of potential surplus.[9]

The main conclusions from the experiment are clear. First, observed outcomes were close to equilibrium predictions: although not to the full extent predicted by theory, the minority nevertheless wins a substantial fraction of the times and a substantial share of the payoff.[10] The experiment confirms the potential of storable votes to empower minorities. Second, the results reflected treatment effects consistent with theory: in particular, the two light blue circles are below the other points in the diagrams, suggesting that cohesive minorities with correlated intensities will benefit the most from storable votes. Similarly, efficiency relative to majority voting rose in the simple C treat-

ments (the three purple dots above the 45 degree line in the third panel), and fell slightly in the B treatments (light blue). Third, contrary to our expectations, efficiency fell in the two modifications of the C treatment that we had designed to improve coordination, the yellow dots and the two large purple dots.

Finally, an examination of individual voting behavior reveals that voters did use responsive strategies, casting more votes when intensity was higher.[11] While bonus-vote choices were generally monotonic in intensity, observed behavior was not perfectly consistent with the equilibrium strategies: minority members did not always respect the threshold of 50, and occasionally split their bonus votes over the two proposals for intermediate intensities; majority members were particularly likely to split their votes, even in treatment B, and predictably found the group equilibrium strategy in treatment C very difficult. Nevertheless, the monotonicity of voting behavior was sufficient to produce committee outcomes very similar to equilibrium outcomes. Other experiments with storable votes in symmetric environments (Casella, Gelman and Palfrey 2006) corroborate this finding, indicating a robustness to small departures from equilibrium behavior that we

---

[9] The point denoted by a darker and larger blue circle refers to a session of the B treatment run, as a robustness check, with nine voters, with group sizes 5 and 4. As theory predicts, the observed minority win rate and the relative minority payoff were greater than in the 3–2 committees.

[10] The points in the first two panels are disproportionately below the 45 degree line, probably because strategic mistakes are much more costly for the minority than the majority.

[11] A full discussion of subjects' strategies can be found in Casella, Palfrey and Riezman (2007).

see as an encouraging sign of the practical viability of the mechanism.

Conclusion

Our theoretical work suggests that storable votes can be an effective and reasonably efficient way to enfranchise minority voters. Laboratory experiments confirm that the mechanism works in practice. Many questions remain open about the applicability of the mechanism to less stylized environments, but at this stage the idea appears promising.

References

Bowler, S., T. Donovan and D. Brockington (2003), *Electoral Reform and Minority Representation: Local Experiments with Alternative Elections,* Ohio State University Press, Columbus.

Casella, A. (2005), "Storable Votes", *Games and Economic Behavior* 51, 391–419.

Casella, A., A. Gelman and T.R. Palfrey (2006), "An Experimental Study of Storable Votes", *Games and Economic Behavior* 57, 123–54.

Casella, A., T. R. Palfrey and R. Riezman (2007), "Minorities and Storable Votes", *Working Paper,* California Institute of Technology, Pasadena.

Cox, G. (1990), "Centripetal and Centrifugal Incentives in Electoral Systems", *American Journal of Political Science* 34, 903–35.

Gerber, E. R., R. B. Morton and T. A. Rietz (1998), "Minority Representation in Multimember Districts", *American Political Science Review* 92, 127–44.

Guinier, L. (1994), *The Tyranny of the Majority,* Free Press, New York.

Issacharoff, S., P. Karlan and R. Pildes (2002), *The Law of Democracy: Legal Structure and the Political Process,* Foundation Press, Westbury, N.Y. (2nd edition).

Pildes, R. and K. Donoghue (1995), "Cumulative Voting in the United States", *The University of Chicago Legal Forum,* 241–313.

Smith, V. (1976), "Experimental Economics: Induced Value Theory", *American Economic Review: Papers and Proceeding* 66(2), 274–79.

Williams, C. (1951), *Cumulative Voting for Directors,* Graduate School of Business Administration, Harvard University, Boston.