New Perspectives on Social Choice

Eric Pacuit, University of Maryland

Lecture 10: Storable Votes and Quadratic Voting

PHIL 808K

The Preference Intensity Problem

51	49
a	b
b	a

51% of the voters have a *slight* preference for a over b and 49% of the voters have a *strong* preference for b over a.

- Utilitarian considerations suggests that b should win.
- Majoritarian considerations suggests that a should win.

Systematic Minority

- If voters cast a single vote for a single candidate, the majority, no matter how slender, is guaranteed victory.

Systematic Minority

- If voters cast a single vote for a single candidate, the majority, no matter how slender, is guaranteed victory.
- When group barriers are permeable, the minority can occasionally belong to the winning side.

Systematic Minority

- If voters cast a single vote for a single candidate, the majority, no matter how slender, is guaranteed victory.
- When group barriers are permeable, the minority can occasionally belong to the winning side.
- When preferences are fully polarized and the power of a cohesive majority bloc is secure, the minority remains disenfranchised.

Systematic Minority

- If voters cast a single vote for a single candidate, the majority, no matter how slender, is guaranteed victory.
- When group barriers are permeable, the minority can occasionally belong to the winning side.
- When preferences are fully polarized and the power of a cohesive majority bloc is secure, the minority remains disenfranchised.
- Some solutions:
  - Ensure that the political districts are *fair*: `https://mggg.org/`
  - In some instances power-sharing is imposed directly, and the constitution grants executive positions to specific groups, typically on the basis of their ethnic or religious identity. The problem is that constitutional provisions of this type are difficult to enforce and heavy-handed.

Cumulative Voting

A possible remedy is the creation of larger, less arbitrary, multi-member districts, where multiple positions are at stake.

By itself multi-member districts do not alleviate the mechanical control granted to the majority, but they make it possible to adopt alternative voting systems that weaken majority control.

Cumulative Voting (CV): each voter has as many votes as there are open seats, and the candidates with more votes win, and each voter is allowed to distribute the votes freely among any number of candidates.

Cumulative Voting

Cumulative Voting (CV): each voter has as many votes as there are open seats, and the candidates with more votes win, and each voter is allowed to distribute the votes freely among any number of candidates.

Cumulative Voting

Cumulative Voting (CV): each voter has as many votes as there are open seats, and the candidates with more votes win, and each voter is allowed to distribute the votes freely among any number of candidates.

- First proposed by Charles Dodgson (aka Lewis Carroll) in 1884
- CV treats every voter equally; yet, a cohesive minority can ensure itself some victories by cumulating its vote.
- CV was used for more than 100 years, from 1870 to 1980, to elect representatives to the Illinois State House and is the rule now in tens of local jurisdictions, often as the remedy imposed by the courts in response to voting rights litigation.
- Outside local politics, it is used to elect corporate boards in approximately 10% of S&P 500 companies.

J. Sawyer and D. MacRae. *Game Theory and Cumulative Voting in Illinois: 1902-1954*. American Political Science Review, 56, 936-946, 1962.

D. Cooper and A. Zillante. *A Comparison of Cumulative Voting and Generalized Plurality Voting*. Public Choice 150, 363-83, 2012.

A. Casella, J. Guo, and M. Jiang. *Minority Turnout and Representation under Cumulative Voting. An Experiment*. NBER Working Paper, 2021.

Consider a two-agent society making a binary decision represented by $d \in \{a, b\}$.

- An agent's preferences are summarized by the difference in utilities between decisions a and b: $v_i = u_i(b) - u_i(a)$. For simplicity, assume that each v_i is independently and identically distributed and takes on values in $\{-2, -1, 1, 2\}$. with equal probability.

- Suppose that we wish to choose the decision that maximizes the sum of the utilities, and in the case of a tie we flip a coin.

M. O. Jackson and H. F. Sonnenschein. *Overcoming Incentive Constraints by Linking Decisions*. Econometrica, 2005.

Consider a two-agent society making a binary decision represented by $d \in \{a, b\}$. Use majority rule to make the group decision: have the agents vote over the alternatives, the alternative with the most votes wins, and in case of a tie, flip a coin

- The inefficiency is that we are not able to discover the agents' intensity of preference in the event of a tied vote. This is not an issue of inter-personal comparisons, but rather intra-personal comparisons.
- An agent would be better off if he or she won ties when having a type -2 or 2, at the cost of losing ties when of type -1 or 1. With just one decision, the agent would always pretend to be of higher type.

Consider a two-agent society making a binary decision represented by $d \in \{a, b\}$.

- If two decisions are linked, we could, for instance, ask the agents to declare that they are of a high type on just one of the two decisions. Essentially, by linking the decisions together, we can ask, "Which decision do you care more about?"
- Effectively, linking the decision problems changes the ex ante inefficiencies "I would like to make trades over my different possible future selves" to ex post inefficiencies "I now actually have different selves and would be happy to make trades across them".

M. O. Jackson and H. F. Sonnenschein. *Overcoming Incentive Constraints by Linking Decisions*. Econometrica, 2005.

Qualitative Voting

Just as we contemplate the importance of the *willingness to pay* in the provision of public goods will increase efficiency, we would expect that taking into account the *willingness to influence* in a voting situation will increase efficiency.

Qualitative Voting (QV): In a setting with a closed agenda of N issues, allow voters to simultaneously and freely distribute a given number of votes among the issues.

- Voters have a broader set of strategies than the classical 'one person — one decision — one vote' but at the same time preserving equality, since all individuals are endowed with the same ex ante voting power.

R. Hortala-Vallve. *Qualitative voting*. Journal of Theoretical Politics 24(4), pp. 526-554, 2011.

Qualitative Voting

QV introduces two main improvements on the usual voting rules.

1. 'the problem of intensity' is solved by allowing strong minorities to decide over weak majorities.
2. voters are allowed to trade off their voting power, adding more weight to the issues they most care about, and so unlocks conflict-resolution situations.

R. Hortala-Vallve. *Qualitative voting*. Journal of Theoretical Politics 24(4), pp. 526-554, 2011.

Qualitative Voting

Two friend Anna ($i = 1$) and John ($i = 2$) are spending an evening together.

Above all they want to be together; however, they can't come to an agreement: Anna wants to see a horror film and would like to have dinner in a new Italian restaurant while John prefers a comedy film and eating sushi in a Japanese restaurant (i.e. we have two linked battle of the sexes games).

Qualitative Voting

If they vote on each of the issues nothing is decided and they have to stay at home (which we assume is not optimal for either of them).

In addition, suppose that Anna really cares about the restaurant decision while John cares more about the film. It seems sensible that, as good friends, each of them will give up on their least preferred option; they will both go to the Italian restaurant and the comedy film.

From a game theoretic perspective, they are both coordinating on the Pareto-optimal allocation that maximises the sum of utilities. QV is precisely a mechanism that allows voters to coordinate non-cooperatively on the only ex ante optimal outcome.

Storable Votes

In a setting with a finite number of binary issues, the Storable Votes mechanism (SV) grants a fixed number of total votes to each voter with the freedom to divide them as wished over the different issues, knowing that each issue will be decided by simple majority.

- SV allows the minority to prevail occasionally and yet is anonymous and treats everyone identically.
- SV can apply to direct democracy in large electorates, or to smaller groups, possibly legislatures or committees formed by voters' representatives.

A. Casella (2005). *Storable votes*. Games and Economic Behavior, 51(2), pp. 391 - 419.

A. Casella (2012). *Storable votes: Protecting the minority voice*. Oxford: Oxford University Press.

Although easy to describe, SV poses a challenging strategic problem: how should a voter best divide her votes over the different issues? Note a central ingredient of the strategic environment: the hide-and-seek nature of the game between majority and minority voters. If the majority spreads its votes evenly, then the minority can win some issues by concentrating its votes on them, but if the majority knows in advance which issues the minority is targeting, then the majority can win those too.

A. Casella, J.-F. Laslier, and A. Macé. *Democracy for Polarized Committees: The Tale of Blotto's Lieutenants*. Games and Economic Behavior, 106, pp. 239 - 225, 2017.

Colonel Blotto game

Two opposing military leaders with given army sizes must choose how many soldiers to deploy on each of several battlefields. Each battlefield is won by the army with the larger number of soldiers. Each colonel could win if he knew the opponent's plan. At equilibrium, choices must be random.

Colonel Blotto game

Two opposing military leaders with given army sizes must choose how many soldiers to deploy on each of several battlefields. Each battlefield is won by the army with the larger number of soldiers. Each colonel could win if he knew the opponent's plan. At equilibrium, choices must be random.

The SV's model can be phrased as in the classical Colonel Blotto scenario, with "issues" and "votes" instead of "battlefields" and "soldiers".

- The game is asymmetric— the majority has more votes
- The game is a decentralized Blotto game. Each voter, whether in the majority or in the minority, controls a number of votes, to be allocated to the different issues.

- N individuals must resolve $K \geq 2$ binary issues, where $\mathcal{K} = \{1, \ldots, K\}$.

- The same M individuals are in favor of all proposals and the remaining $N - M = m$ are opposed to all proposals with $m \leq M$.

- The specific direction of preferences is irrelevant, what matters is that the two groups are fully cohesive and fully opposed.

- Each individual receives utility 1 from any issue resolved in her preferred direction, and 0 otherwise. Thus each individual's goal is to maximize the fraction of issues resolved according to her—and her group's—preferences.

- Individuals are all endowed with K votes each. Voters are allowed to distribute their votes freely among the different issues. Each issue is then again decided according to the majority of votes cast. Voting on the K issues is contemporaneous, and all individuals vote simultaneously. Ties are resolved by a fair coin toss.

- p_m the expected fraction of minority victories.

- The parameters of the game are common knowledge, in particular each voter knows exactly the size of the two groups, and thus both her own and everyone else's preferences.

Voters never cast a vote against their preferences. The focus is on each voter's distribution of votes among the K issues. The action space for each player is:

$$S(K) = \{s = (s_1, \ldots, s_K) \in \mathbb{N}^K \mid \sum_{k=1}^{K} s_k = K\}$$

where s_k is the number of votes cast on issue k.

Examples

- If $m = 4$, $M = 5$ and $K = 3$, there exists an equilibrium in which the minority wins two of the three issues.

- If $m = 1$ and $M = 2$ and $K = 4$, there exists no pure-strategy equilibrium.

If several minority members concentrate votes on a given issue, the minority may be able to win it. But only if the majority does not know which specific issue is being targeted. Thus, minority members need not only to concentrate their votes but also to randomly choose the issues on which the votes are concentrated. Mixed strategies allow them to do so.

If several minority members concentrate votes on a given issue, the minority may be able to win it. But only if the majority does not know which specific issue is being targeted. Thus, minority members need not only to concentrate their votes but also to randomly choose the issues on which the votes are concentrated. Mixed strategies allow them to do so.

For any c factor of K, define σ^c (note τ^c for a majority player) as follows: choose randomly K/c issues, and allocate c votes to each of the selected issues.

Intuitively, we expect the minority to concentrate its votes, so as to achieve at least some successes, and the majority to spread its votes, because its larger size allows it to cover, and win, a larger fraction

- Suppose $K = 4$ and $M = 10$, and $m \in \{1, \ldots, 10\}$.

- Minority voters adopt the σ^c strategies with $c = 2$ (each minority voter casts two votes each on half of the issues, chosen with equal probability) and $c = 4$ (each minority voter casts all votes on a single issue chosen randomly).

- Consider two rules for the majority: (1) each majority voter casts his votes randomly and independently over all issues (an upper bound on p_m) or (2) all majority voters together best respond to the minority (the lower bound).

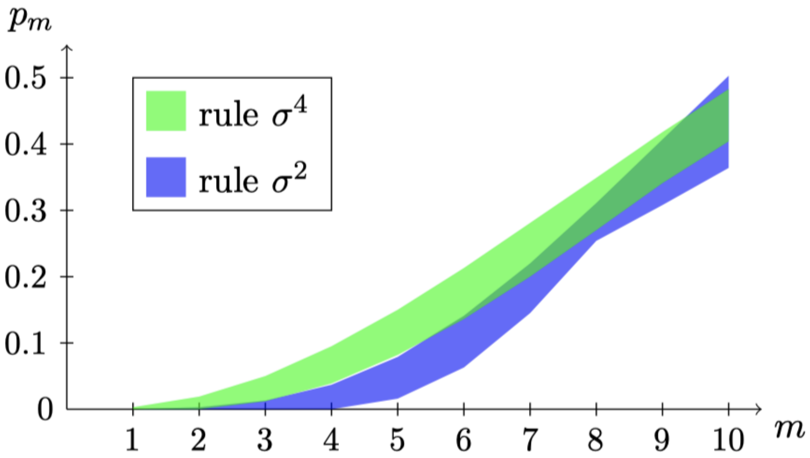

Figure 1: Minority payoffs for two minority rules ($M = 10$)

Experiments

We designed the experiment with two goals in mind.

1. We wanted to learn how substantive are minority victories in the lab and how well the theory predicts subjects' behavior.
2. We wanted to compare results with and without communication. Does communication helps or hinders the relative success of the minority?

Experiments

Each experimental session consisted of 20 rounds with fixed values of m and M; the first ten round without communication, and the second ten with group-communication.

Experiments

Each experimental session consisted of 20 rounds with fixed values of m and M; the first ten round without communication, and the second ten with group-communication.

At the start of each session, each subject was assigned a color, either Blue or Orange, corresponding to the two groups. Members of the two groups were then randomly matched to form several committees, each composed of m Orange members and M Blue members.

Experiments

Each experimental session consisted of 20 rounds with fixed values of m and M; the first ten round without communication, and the second ten with group-communication.

At the start of each session, each subject was assigned a color, either Blue or Orange, corresponding to the two groups. Members of the two groups were then randomly matched to form several committees, each composed of m Orange members and M Blue members.

Every committee played the following game. Each subject entered a round endowed with K balls of her own color. She was asked to distribute them as she saw fit among K urns, depicted on the computer screen, knowing that she would earn 100 points for each urn in her committee in which a majority of balls were of her color.

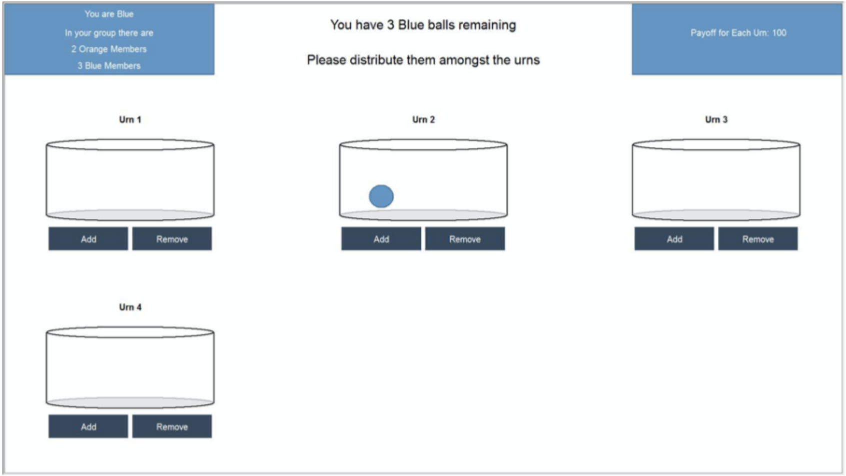

Figure 2: The Allocation screen

After all subjects had cast their balls, the results appeared on the screen under each urn: the number of balls of each color in the urn, the tie-break result if there was a tie, and the subject's winnings from the urn (either 0 or 100).

The session then proceeded to the next round. The first ten rounds were all identical to the one just described. Subjects kept their color across rounds, but committees were resuhffled randomly.

After the first round, subjects could consult the history of past decisions before casting their balls. By clicking a History button, each subject accessed a screen summarizing ball allocations and outcomes in previous rounds, by urn, in the committee that in each round included her.

After ten rounds, the session paused and new instructions were read for the second part.

Parameters and choices remained unchanged and subjects kept the same color, but now a chatting option was enabled: before casting their balls, subjects had two minutes to exchange messages with other members of their committee who shared their color. They could consult the history screen while chatting.

The second part of the session again lasted ten rounds, and again committees were reshuffled after each round but subjects kept the same color.

Thus each subject belonged to the same group, m or M, for the entire length of the session, a design choice we made to allow for as much experience as possible with a given role.

Results 1: No Communication

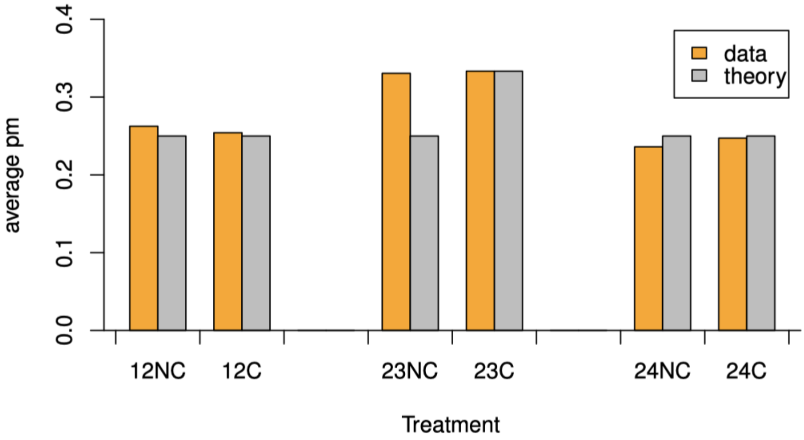

Figure 3: Fractions of minority victories

Results 2: No Communication

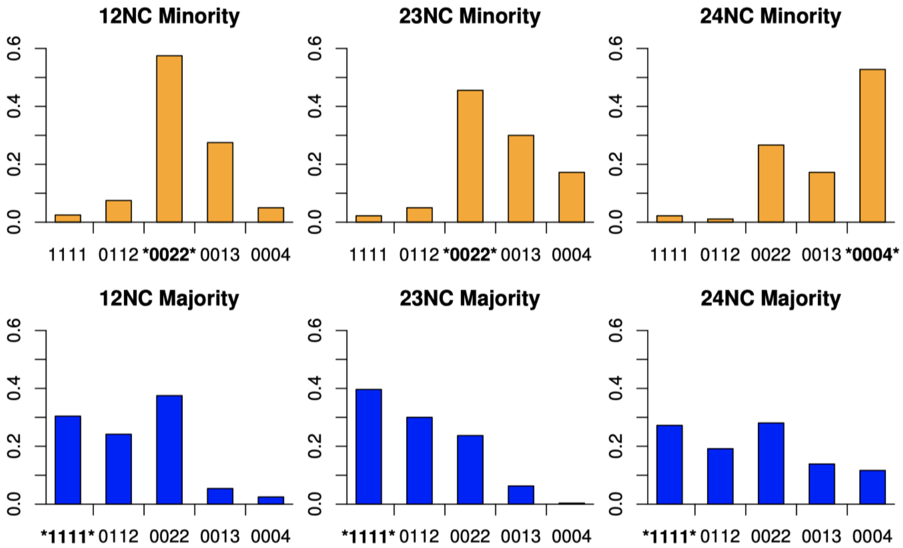

1. There is substantial deviation from equilibrium strategies: in all treatments and in both groups, at least forty percent of all individual allocations do not correspond to equilibrium strategies.

2. However, equilibrium predictions have some explanatory power for minority subjects. In all treatments, the most frequently observed allocation for minority subjects corresponds to the equilibrium strategy, a particularly clear result in treatment 12NC and 24NC, where more than half of all observed allocations correspond to the predictions.

3. Equilibrium predictions are noticeably less useful for majority subjects. We are not sure why. We can speculate that the difference may be due to the higher complexity of the majority members' problem: Should they spread their votes, or try to second-guess the minority?

4. The theory's qualitative predictions are mostly satisfied, both across treatments and between the two groups. We have ordered the five possible ball allocations with concentration increasing progressively from left to right. In all treatments, the distribution of minority allocations is shifted to the right, relative to the majority distribution: predictably, and in line with the theory, minority members tend to concentrate balls more than majority members do.

To what extent does communication influence the groups' allocations? We find that all subjects actively participate in the chats, and the messages are very relevant: 91 percent of exchanges 30 mention at least one ball allocation, 36 percent refer to the opposite group, and 84 percent include an explicit agreement.

Results: Communication

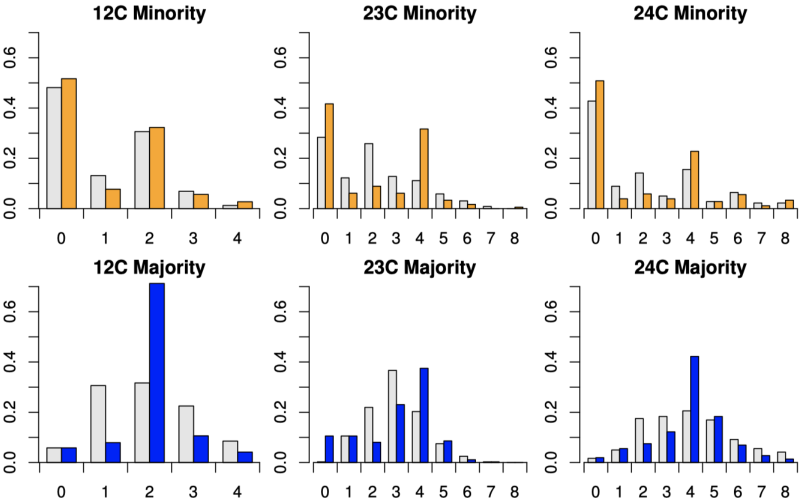

Figure 6: Frequency of group marginal allocations of balls

A. Casella, S. Turban and G. Wawro. *Storable votes and judicial nominations in the US Senate*. Journal of Theoretical Politics, 29:2, pp. 243 - 272, 2016.

This paper starts from the premise that the minority has a legitimate, important role in confirming nominations.

The expression of intense sentiment by the minority once figured prominently in filibuster battles, and its expression was valued by the majority because it provided an informative signal about public opinion. Yet, the power of the minority should not trump the majority's right to govern; it should consist in the institutional recognition of principled support or opposition to speicific nominees.

The puzzle is how to design transparent, formal institutions that balance the minority's right to be heard with the majority's right to rule.

The proposal is to use Storable Votes confirm or reject judicial nominees on slates submitted to the chamber.

Allows the parties a mechanism to reveal the salience of their preferences, and grants the minority the power to prevail on some nominations, but only on those that the minority considers a higher priority than the majority does.

The procedural innovation that we explore shares the same spirit: blocking a nomination should be costly, and the willingness to bear that cost measures the intensity with which the defeat of a nominee is desired.

Our simulations show that a higher correlation of intensities within parties—higher agreement on which nominations are most important—results in more coordinated voting and favors the minority, whose smaller numerical size makes coordination essential.

Our simulations show that a higher correlation of intensities within parties—higher agreement on which nominations are most important—results in more coordinated voting and favors the minority, whose smaller numerical size makes coordination essential.

On the other hand, stronger correlation in intensities across parties—when the nominees the president's party most wants to confirm are those the opposition most wants to block—favors the majority, because the larger party tends to win when the two parties prioritize the same nominees.

Our simulations show that a higher correlation of intensities within parties—higher agreement on which nominations are most important—results in more coordinated voting and favors the minority, whose smaller numerical size makes coordination essential.

On the other hand, stronger correlation in intensities across parties—when the nominees the president's party most wants to confirm are those the opposition most wants to block—favors the majority, because the larger party tends to win when the two parties prioritize the same nominees.

When both types of correlations are high—the case we consider most realistic for today's Senate—our results show that a minority of 45 senators can prevail on about 35 percent of nominations on any given slate.

Can a president name a nominee so objectionable to the opposition that all its votes are concentrated on defeating him, guaranteeing that the other nominations are confirmed? We find that, indeed, placing on the slate one or more “decoy” nominees can be advantageous.

Can a president name a nominee so objectionable to the opposition that all its votes are concentrated on defeating him, guaranteeing that the other nominations are confirmed? We find that, indeed, placing on the slate one or more "decoy" nominees can be advantageous.

Decoys work, then, but only if the remaining nominees are less polarizing. This is a direct effect of storable votes: because the number of votes cast depend on priorities, a decoy nominee can concentrate the votes of the opposition only if the others nominations are on the whole acceptable. As a result, storable votes exercise a moderating effect on the list of nominees.

B. Laurence and I. Sher. *Ethical considerations on quadratic voting*. Public Choice, 172: pp. 195 - 222 2017.

E. Posner and G. Weyl. *Voting squared: Quadratic voting in democratic politics*. Vanderbilt Law Review, 68, pp. 441 - 500, 2015.

S. Lalley and G. Weyl. *Quadratic voting*. Unpublished manuscript, June 2016.

Quadratic Voting

The thought animating QV is that if we put a price on votes, then voters could express their preference intensity—more passionate voters are willing to pay more—and we solve the problem preference intensity problem.

Specifically in QV, it costs v^2 dollars to purchase v votes, which can then be used to vote for either alternative.

The alternative with more votes wins.

There is a systematic divergence between utility and willingness to pay: if a rich person and a poor person care about a decision equally—the decision has the same impact on their utility—the rich person will be willing to pay more than the poor person for that decision to be made.

There is a systematic divergence between utility and willingness to pay: if a rich person and a poor person care about a decision equally—the decision has the same impact on their utility—the rich person will be willing to pay more than the poor person for that decision to be made.

This is because the rich person will have to give up only certain luxuries if she spends $y, whereas the poor person will have to give up more basic wants or necessities for the same expenditure.

So under QV, the preferences of the rich will be overrepresented relative to their true ethical weight.

This could mean that majority voting is better than QV from a utilitarian point of view.

If preferences are independent of wealth, in the sense that overall poor voters have the same distribution of utilities for outcomes as rich voters, then QV will be optimal from a utilitarian perspective. The divergence between willingness to pay and utility will not cause any distortion in the outcome.

If preferences are independent of wealth, in the sense that overall poor voters have the same distribution of utilities for outcomes as rich voters, then QV will be optimal from a utilitarian perspective. The divergence between willingness to pay and utility will not cause any distortion in the outcome.

If, on the other hand, preferences are polarized by wealth, so that the poor prefer one thing and the rich another, then the louder voice of the rich under QV may win out even when the poor care more.

If there are more poor than rich, then under majority voting the poor win, as the utilitarian would have liked.

If preferences are independent of wealth, in the sense that overall poor voters have the same distribution of utilities for outcomes as rich voters, then QV will be optimal from a utilitarian perspective. The divergence between willingness to pay and utility will not cause any distortion in the outcome.

If, on the other hand, preferences are polarized by wealth, so that the poor prefer one thing and the rich another, then the louder voice of the rich under QV may win out even when the poor care more.

If there are more poor than rich, then under majority voting the poor win, as the utilitarian would have liked.

It is fundamentally ambiguous whether QV or majority voting is better from a utilitarian standpoint.

$x \in \{0, 1\}$ where 0 represents decision 0 and 1 represents decision 1.

$N = \{1, \ldots, n\}$ is the set of agents.

Each $i \in N$ has a wealth endowment $w_i^e \in \mathbb{R}_+$ representing i's initial wealth.

An outcome for i is a pair (x, w_i) where x is the public decision and w_i is i's final wealth holdings.

For each i, there is a quantity of money $\hat{u}_i$ such that i is indifferent between $(0, w_i^e)$ and $(1, w_i^e - \hat{u}_i)$.

Assuming i likes money, $\hat{u}_i > 0$ if i prefers decision 1 and $\hat{u}_i < 0$ if i prefers decision 0.

So, $|\hat{u}_i|$ is i's willingness to pay for decision 1 to replace decision 0 if i prefers 1, and is the minimum payment i is willing to accept for 1 to replace 0 if i prefers 0.

i's preference over outcomes (x, w_i) are represented by the quasilinear utility function

$$\hat{U}_i(x, w_i) = \hat{u}_i x + w_i$$

If $x = 0$, then $\hat{u}_i x = 0$ and if $x = 1$ then $\hat{u}_i x = \hat{u}_i$. So we can think of i's utility for decision 0 as being normalized to zero and i's utility for decision 1 is $\hat{u}_i$.

Let $v \in \mathbb{R}$ be a quantity of votes.

If v is positive, then v represents $|v|$ votes for decision 1 and if v is negative, then v represents $|v|$ votes for decision 0.

Let $c : \mathbb{R} \to \mathbb{R}_+$ be a costly voting rule that maps votes into dollars. Assume that c is even, which means that $c(v) = c(-v)$.

So we can think of c as a function of the absolute value $|v|$. Assume that c is differentiable, convex, and strictly increasing in $|v|$, and that $c(0) = 0$.

The interpretation is that it requires \$$c(|v|)$ to purchase $|v|$ votes, which can then be cast for either alternative, decision 0 or decision 1.

The alternative that receives the most votes wins.

Election proceeds are refunded (approximately) equal to all citizens. Let v_j be the votes purchased by voter j.

Then voter i is refunded

$$\$ \frac{\sum_{j \in N \setminus i} c(v_j)}{n-1}$$

This is is the average payment made by voters other than i.

Voter i's refund is independent of how i votes, and even whether i votes at all, so as not to affect i's incentive to vote.

Voting is a market in which voters purchase influence. This is the price-taking model.

A collective decision problem is a tuple $\{N, S, \overline{u}\}$ where N is the collection of agents, S, a positive real number, is the supply of influence, and $\overline{u} = (\hat{u}_i : i \in N) \in \mathbb{R}^N$ is the profile of utilities for decision 1.

A price-taking equilibrium of a collective decision problem under voting rule c is an influence vector $\overline{\boldsymbol{I}}^* = (I_i^* : i \in N) \in \mathbb{R}^N$, a price $p^* \in \mathbb{R}_+$ and a decision $x^* \in \{0, 1\}$. such that

- I_i^* maximizes $\hat{u}_i I_i - c(p^* I_i)$ over all $I_i \in \mathbb{R}$
- $\sum_{i \in N} |I_i^*| = S$
- $x^* = 1$ iff $\sum_{i \in N} p^* I_i^* \geq 0$

Let Q_i^* be the probability that decision 1 would win the election if i did not purchase any votes.

By buying votes, voter i purchases influence.

Let I_i be the additional probability that i's favorite alternative wins given the number of votes that i purchases.

Given the influence that i obtains by purchasing votes, the probability that the outcome is decision 1 is $Q_i^* + I_i$.

$I_i > 0$ if i votes for decision 1 and $I_i < 0$ if i votes for decision 0.

Acquiring l_i units of influence $v_i = p^* l_i$ votes. Thus, p^* is the "price of influence".

Equivalently, we can think of $\frac{1}{p^*}$ as a voter's marginal pivotality, that is, the additional probability of being pivotal she purchases with an additional vote.

For what follows, it is not essential that the is conversion from votes to pivotality probabilities be linear; it matters only that each voter perceives (approximately) the same marginal pivotality.

The voter's problem is

$$\max_{I_i} \left[\underbrace{\hat{u}_i \times (I_i + Q_i^*)}_{\text{expected value of decision}} - \underbrace{c(p^* I_i)}_{\text{vote cost}} \right]$$

QV is the costly voting rule $c(v) = kv^2$ for some $k > 0$. For simplicity assume $k = 1$ so that $c(v) = v^2$.

Under QV, the voter's optimization problem is

$$\max_{l_i} \hat{u}_i l_i - (p^* l_i)^2$$

Under QV, the voter's optimization problem is

$$\max_{I_i} \hat{u}_i I_i - (p^* I_i)^2$$

The first-order condition for an optimum is $\hat{u}_i = 2(p^*)^2 I_i^*$. Equivalently, I's optimal influence is $I_i^* = \frac{\hat{u}_i}{2(p^*)^2}$. Converting influence into votes (using $v_i^* = p^* I_i^*$), i purchases

$$v_i^* = \frac{\hat{u}_i}{2p^*}$$

Under QV, the voter's optimization problem is

$$\max_{l_i} \hat{u}_i l_i - (p^* l_i)^2$$

The first-order condition for an optimum is $\hat{u}_i = 2(p^*)^2 l_i^*$. Equivalently, l's optimal influence is $l_i^* = \frac{\hat{u}_i}{2(p^*)^2}$. Converting influence into votes (using $v_i^* = p^* l_i^*$), i purchases

$$v_i^* = \frac{\hat{u}_i}{2p^*}$$

So, under QV,

$$x^* = 1 \text{ iff } \sum_{i=1}^{n} \hat{u}_i \geq 0$$

Each costly voting rule c and collective decision problem $\{N, S, \overline{u}\}$ determine a unique price-taking equilibrium.

c is **robustly efficient** if for all collective decision problems, this unique price-taking equilibrium satisfies:

$$x^* = 1 \text{ iff } \sum_{i=1}^{n} \hat{u}_i \geq 0$$

Proposition 1 (Lalley and Weyl 2015) A costly voting rule c is robustly efficient if and only if $c(v) = kv^2$ for some $k > 0$.

QV vs. Majority Rule

Suppose there are three voters: Ann, Bob, and Carol.

Ann and Bob each prefer decision 0, and each would be willing to pay \$1 to cause decision 0 to replace decision 1.

Carol prefers decision 1 and willing to pay \$3 for her preference.

Under majority voting, there are two votes for decision 0, and one vote for decision 1. So decision 0 wins.

Yet $\hat{u}_{\mathsf{Ann}} + \hat{u}_{\mathsf{Bob}} + \hat{u}_{\mathsf{Carol}} = -1 + -1 + 3 = 1 > 0$.

So decision 0 wins while the aggregate willingness to pay for decision 1 to replace decision 0 is positive.

As we have seen above, under QV, decision 1 wins.

1. From a utilitarian point of view, decision 1 (the outcome of QV) is better than decision 0 (the outcome of majority voting) because the sum of utilities for decision 1 is greater.

2. If we enact decision 0, then everyone could be made better off: we could instead enact decision 1, and Carol could make a payment of $\$\frac{4}{3}$ to Ann and a payment of $\$\frac{4}{3}$ to Bob. The utility of each of the three agents would rise by $\frac{1}{3}$. So we should not enact decision 1. Majority voting, unlike QV, enacts an outcome against which there is a decisive argument.

We make an ethical assumption on utility that is widespread in the utilitarian tradition, namely, the assumption that the marginal utility of a dollar is decreasing in wealth. Under this assumption, if we had one dollar, and could either give it to a rich person or a poor person, then, other things being equal, it would be ethically better to give it to the poor person. Intuitively, the dollar makes more of a difference to the poor person.

U_i is the *ethical utility* and $\hat{U}_i$ from above is the value function and $\hat{u}_i$ is the willingness to pay.

$$U_i(x, w_i) = u_i x + g(w_i)$$

$x \in \{0, 1\}$ is the public decision, w_i is the final wealth of i.

u_i can be positive or negative and is i's utility of public decision 1 (and the utility of public decision 0 is 0).

The utility of wealth function $g(w_i)$ measures the utility an agent receives from having wealth w_i. For simplicity, g is the same for all agents i. We assume that g is concave to capture the diminishing marginal utility of wealth.

For simplicity, we assume that u_i is not a function of wealth w_i.

Let the agents be N^0.

The utilitarian objective is:

$$\sum_{i \in N^0} U_i(x, w_i)$$

Now that utility is understood in an ethical sense, we can see that QV's refund is ethically significant because it amounts to a transfer among voters with potentially different marginal utilities of wealth. But is the refund's effect on total utility positive or negative?

- Imagine that voters are split into two groups A and B. Voters in A care passionately about the election and voters in B are relatively indifferent. So voters in A purchase a large number of votes while voters in B purchase a small number of votes.

- As proceeds are refunded equally to all voters independently of voting behavior, the refund then amounts to a net transfer from A to B. B could be large or small in comparison to A and voters in B could be wealthy or poor in comparison to voters in A.

- Depending on how these parameters are resolved, the transfer could be from rich to poor or from poor to rich and its effect may be larger or smaller. So, the refund could have a positive or negative effect on aggregate utility.

A. Casella and L. Sanchez. *Democracy and Intensity of Preferences: A Test of Storable Votes and Quadratic Voting on Four California Propositions*. The Journal of Politics, 84:1, 2022.

Four proposition on the November 2016 California ballot:

1. Bilingual education (BE): re-instate the possibility of bilingual classes in public schools. (The proposition was included in the November 2016 ballot and passed.)
2. Immigration (IM): require all state law enforcement officials to verify immigration status in case of an infraction and report undocumented immigrants to federal authorities. (The proposition was not included in the final ballot.)
3. Teachers' tenure (TT): increase required pre-tenure experience for teachers from two to five years. (The proposition was not included in the ballot.)
4. Public Vote on Bonds (PB): require voters' approval for all public infrastructure projects of more than $2 billion. (The proposition was included in the ballot and failed.)

We then recruited 647 California subjects via Amazon Mechanical Turk (MTurk).

We first asked each subject how (s)he would vote on each of the four propositions, presented in random order, allowing for the option to abstain.

We then elicited measures of intensity of preferences. Each subject was asked to distribute 100 points among the four propositions, with the number of points used as scale of the importance attributed to each proposal ("How important is this issue to you?"). We used examples to clarify that importance is independent of whether the respondent is in favor or against a proposition, and summarized responses in terms of priorities, allowing for revisions and asking for a final confirmation.

After this first part of the survey, common to all respondents, subjects were randomly assigned either to the SV treatment (324 subjects; 306 after data cleaning) or to the QV treatment (323 subjects; 313 after cleaning).

In the SV treatment, the subjects were told that each was granted one extra vote, in addition to the regular votes cast earlier, and were asked to choose the proposition in which to use it. The vote was cast in the direction indicated in the first part of the survey, and the final outcome under SV was calculated summing regular and bonus votes.

The design of the QV scheme required some innovation.

We asked respondents to choose one of four classes of votes, distinguished by color and weight.

1. Blue votes are regular votes, four in number; a person choosing blue votes casts one vote on each proposition.
2. Three green votes, each worth more than a regular blue vote.
3. Two yellow votes, each stronger than a green vote.
4. One red vote, stronger than a yellow vote.

The weights we assigned to the different votes are 1 for blue votes, 1.2 for green votes, 1.5 for yellow votes, and 2 for the red vote. A subject who chooses green/yellow/red votes casts votes on only three/two/one proposition(s).

The simple four-class classification respects the convex cost of concentrating votes at the heart of QV. A voter casting votes on all four propositions—choosing blue votes—has a total weight of 4, but the total weight declines as votes are concentrated: the total weight corresponding to the three green votes is 3.6, to the two yellow votes is 3, and to the single red vote is 2. The decline is increasing with concentration, and increasing at an increasing rate, capturing the core feature of QV.

Our study confirms the theoretical promise of the two voting schemes. SV and QV allow for occasional minority victories on those issues over which the minority's intensity of preferences is sufficiently stronger than the majority's to make a minority victory desirable.