Available online at www.sciencedirect.com
ScienceDirect

Games and Economic Behavior 57 (2006) 123–154

GAMES and Economic Behavior

www.elsevier.com/locate/geb

An experimental study of storable votes ☆

Alessandra Casella [a,b], Andrew Gelman [c], Thomas R. Palfrey [d,*]

[a] *Department of Economics, Columbia University*
[b] *GREQAM, NBER and CEPR*
[c] *Departments of Statistics and Political Science, Columbia University*
[d] *Division of Humanities and Social Sciences, California Institute of Technology*

Received 24 February 2004

Available online 12 June 2006

Abstract

The storable votes mechanism is a voting method for committees that meet periodically to consider a series of binary decisions. Each member is allocated a fixed budget of votes to be cast as desired over the sequence of decisions. This provides incentives for voters to spend more votes on those decisions that matter to them more, typically generating welfare gains over standard majority voting with non-storable votes. Equilibrium strategies have a very intuitive feature—the number of votes cast must be monotonic in the voter's intensity of preferences—but are otherwise difficult to calculate, raising questions of practical implementation. We present experimental data where realized efficiency levels were remarkably close to theoretical equilibrium predictions, while subjects adopted monotonic but off-equilibrium strategies. We are led to conclude that concerns about the complexity of the game may have limited practical relevance.

JEL classification: C92; D71; D72

Keywords: Voting; Experiments

☆ The journal is grateful to the authors for contributing this article in honor of Richard McKelvey to *Games and Economic Behavior*. It was submitted originally for publication in the special issue in honor of Richard D. McKelvey (Vol. 51, Issue 2, May 2005) and was under review when the special issue went to press. The authors are grateful for the financial support of the National Science Foundation (Grants MRI-9977244, SES-0084368 and SES-00214013) and the SSEL and CASSEL laboratories at Caltech and UCLA, respectively.

* Corresponding author.
E-mail address: trp@hss.caltech.edu (T.R. Palfrey).

0899-8256/$ – see front matter
doi:10.1016/j.geb.2006.04.004

1. Introduction

In binary decision problems, simple voting schemes where each voter has one vote to cast either for or against a proposal allow voters to express the direction of their preferences, but not their strength. This remains true if several binary decisions are taken in series, and voters are asked to cast their vote over each, independently of their other choices.[1] It is possible, however, to elicit voters' strength of preferences through voting mechanisms where voting choices are linked. Casella (2005) proposes a mechanism of this sort that has the advantage of being extremely simple: *storable votes* allow each voter to allocate freely a given total budget of votes over several consecutive decisions. Each decision is then taken in accordance with the majority of votes cast, but voters are allowed to cast multiple votes over the same decision, as long as they respect their budget constraint. Thus votes function as a kind of fiat money, playing a role similar to that of transfer payments in more familiar mechanism design problems: voters are induced to cast more votes over those decisions they care more about, increasing their probability of having their way exactly where it matters to them most. As stressed by Jackson and Sonnenschein (in press), this linkage principle has broad applicability to any voting context where the group is making more than one decision, with the potential for substantial efficiency gains.[2]

We are aware of no historical examples of pure storable votes institutions, but the storable votes mechanism is related to a scheme that is not uncommon: *cumulative voting*, where each voter distributes a fixed budget of votes across a field of candidates in a single multi-candidate election. Storable votes can be thought of as a version of cumulative voting applied not to a single election with multiple choices, but to a series of binary decisions taking place over time.[3]

In Casella (2005) voters receive no initial endowment but can accumulate a bank of votes by abstaining on the early votes. That paper discusses conditions under which such a form of storable votes increases efficiency relative to a standard sequence of simple (non-storable) votes. Here we modify the model by endowing voters with an initial stock of votes and derive theoretical results for the extended model. But we also address a more central question. Although the intuition behind storable votes is immediate, voters must play a complicated dynamic game, comparing the marginal effect of an extra vote on the probability of being pivotal today to its effect on the probability of being pivotal sometime in the future, a trade-off that must depend on the whole distribution of voting choices by all committee members. If storable votes were to

[1] If voting is optional and costly, then strength of preference is indirectly expressed through the choice to abstain or vote (for example, Börgers, 2004). But it can be ranked in two classes only—stronger or weaker than the cost of voting. In addition, biases will result if the cost of voting is correlated with voters' preferences (Campbell, 1999; Osborne et al., 2000).

[2] Jackson and Sonnenschein (in press) construct more complex rules that lead to first best efficiency as the number of decision problems becomes large. By a "voting context," we mean a public decision problem where direct side payments are not possible. The standard mechanism design approach (e.g. Crémer et al., 1990) allows unlimited side payments and assumes quasi-linear utility functions.

[3] The idea of cumulative voting has a long history (Dodgson, 1884), and has been promoted as a fair way to give voice to minorities (Guinier, 1994). Cumulative voting was the norm in the Illinois Lower House until 1980, it is commonly used in corporate board elections, and in recent years has been adopted as remedy for violations of fair representation in local elections. See Sawyer and MacRae (1962) for an early discussion of the experience in Illinois, Brams (1975) and Mueller (1989) for more theoretical surveys, Issacharoff et al. (2001) and Bowler et al. (2003) for descriptions of recent experiences. Other voting mechanisms that allow strength of preferences to affect outcomes are peremptory challenges in jury selection (Brams and Davis, 1978), voting by successive veto (Mueller, 1978 and Moulin, 1982), and, less formally, vote trading and log-rolling (Ferejohn, 1974; Philipson and Snyder, 1996; Piketty, 1994). For comparisons to storable votes, see the discussion in Casella (2005).

be used in practical decision-making, would voters be able to solve the problem well enough to achieve something resembling the theoretical properties of the voting mechanism? To address this point, we conducted a laboratory experiment.

Our most important finding is that the efficiency improvements predicted by the theory were observed in the data: the realized experimental efficiencies tracked the theoretical predictions almost perfectly across all treatments. The result is particularly remarkable in light of the fact that the actual choice behavior of the subjects did not track the theory nearly as closely. Equilibrium strategies in this game require voters to cast a number of votes that, for each decision, is increasing in the intensity of their preferences, and take the form of thresholds, or cutpoints, that determine how many votes to cast as a function of valuation. While monotonicity is very intuitive and characterizes all best response strategies with storable votes, calculating the equilibrium thresholds is much more complex. In our data, nearly all subjects adopted approximately monotonic strategies (with a small number of errors), but the same cannot be said of the thresholds equilibrium values: the best fitting cutpoint rules varied across subjects, and even the average estimated cutpoint rule was typically different from the equilibrium cutpoint.

The two observations—realized efficiency that matches the theory and choice behavior that does not—together suggest a robustness of the storable votes mechanism. Monotone voting strategies must be used in order to realize the efficiency gains, but monotone behavior is simple and intuitive, and in real committee settings, voters would be more experienced than our subjects. The potential usefulness of storable votes in practical applications is a difficult policy question, but our experiment provides an encouraging initial probe.

In order to account for some of the behavioral deviations from the theory, we estimated several different models of stochastic choice behavior. We find that logit equilibrium provides a close description of the subjects' strategies and easily outperforms a range of plausible alternative models. The model not only allows for stochastic choice, with the likelihood of errors negatively correlated to foregone expected payoff, but also endogenizes the equilibrium payoff in a way that is similar to Nash equilibrium.

The paper proceeds as follows. In the next section we describe the theoretical model and its main properties; Section 3 describes the experimental design and the theoretical predictions of the model with the parameter values used in the experiment; Section 4 presents the results; Section 5 concludes.

2. The model

A group of n individuals meets regularly over time to vote up or down each period t a proposal P_t, with $t = 1, \dots, T$. In the *Storable votes* mechanism, voters allocate a given initial stock of votes across the different proposals. In each period, each voter casts a single *regular vote* for or against proposal P_t, but in addition, voter i is endowed at time 0 with B_i^0 *bonus votes* and any bonus votes cast by i in period t are added to his regular vote in that period. The decision of how many votes to cast is made sequentially, period by period, after each voter observes his valuation for the current proposal. Thus in the first period voter i casts 1 regular vote, plus any number of bonus votes, b_{i1} between 1 and B_i^0, resulting in a total vote in the first period equal to $x_{i1} = 1 + b_{i1}$. In the second period, i casts his regular vote plus any number of bonus votes, b_{i2}, between 1 and $B_i^1 = B_i^0 - b_{i1}$, resulting in a total vote $x_{i2} = 1 + b_{i2}$, and so forth. Voters cast their votes simultaneously, but once the decision is taken the number of votes that each member has spent—and thus the number of votes remaining—is made public. Each period, the decision is taken according to a simple majority of votes cast, with ties broken randomly.

Individual i's valuation over proposal P_t is summarized by the variable v_{it}, drawn each period from a distribution $F_{it}(v)$ defined over a support $[\underline{v}, \bar{v}]$, with $\underline{v} < 0 < \bar{v}$. A negative realization of v_{it} indicates that individual i opposes proposal P_t. When a proposal is voted upon, each individual i receives utility u_i equal to $|v_{it}|$ if the vote goes in the desired direction, and 0 otherwise. In this paper, we make several assumptions about the distribution functions, F_{it}, that simplify both the theory and the laboratory environment:

(i) v_{it} is identically and independently distributed both across periods and across individuals;
(ii) the common distribution function $F(v)$, defined over $[-1, 1]$, is continuous, atomless and symmetric around 0;

F is common knowledge, and in each period each player observes his own current valuation, v_{it}. The realized valuations of the members of the group other than i are not known to i, nor are i's own future valuations.[4]

Individuals choose each period how many votes to cast to maximize the sum of expected utilities over all proposals. Given F, n, B^0, T the storable votes mechanism defines a multistage game of incomplete information, and we study the properties of the perfect Bayesian equilibria of this game.

Because valuations are not correlated over time, the direction of one's vote holds no information about the direction of future preferences and cannot be used to manipulate other players' future voting strategies. Assuming in addition that players do not use weakly dominated strategies, the direction of each individual vote will be chosen sincerely: all x_{it} votes are cast in favor of proposal P_t if $v_{it} > 0$ and all x_{it} votes are cast against proposal P_t if $v_{it} < 0$. The game however remains complex: it is a dynamic game where preferences are realized over time and the information over the stock of votes held by all other members is updated each period; and it is non-stationary because the horizon is finite. This said, the basic intuition is simple and strong: storable votes should allow voters to express the intensity of their preferences. This intuition is reflected in the formal properties that the equilibrium can be shown to possess.

We focus on strategies such that the number of votes each individual chooses to cast each period, x_{it}, depends only on the state of the game at t, which is the profile of bonus votes each voter has still available, $B = (B_1, \dots, B_n)$, and the number of remaining periods, $T - t$. Hence we refer to (B, t) as the state of the game and denote strategies by $x_{it}(v_{it}, B, t)$. Equilibrium strategies have the following properties:

1. Monotonicity. *We call a strategy monotonic if, at a given state, the number of votes cast is monotonically increasing in the intensity of preferences, $|v_{it}|$. For any number of voters n, horizon length T, and state (B, t), all best response strategies are monotonic.*

The proof follows closely the proof in Casella (2005), with a slightly different model, and is omitted here,[5] but monotonicity is very intuitive: for any number of votes cast by the other voters, the probability of obtaining one's favorite outcome must be increasing, if possibly weakly, in the number of votes one casts. Hence, everything else equal, if it is optimal to cast x votes when the

[4] As discussed in Casella (2005), the informational assumptions were inspired by the example that motivated the idea of storable votes: The Governing Council of the European Central Bank, meeting every month to decide whether to change interest rates or maintain the status quo.

[5] The proof can be found in Casella et al. (2003).

valuation attached to a given decision is $|v|$, it cannot be optimal to cast fewer votes than x when the valuation is higher than $|v|$.

2. Equilibrium. *The game has a perfect Bayesian equilibrium in pure strategies. All equilibrium strategies are monotone cutpoint strategies*: *at any state* (B, t) *and for any voter* i *with* $k_i = B_i + 1$ *available votes there exists a set of cutpoints* $\{c_{i1}(B,t), c_{i2}(B,t), \ldots, c_{ik_i}(B,t)\}$, $0 \leqslant c_{ix} \leqslant c_{ix+1} \leqslant 1$, *such that* i *will cast* x *votes if and only if* $|v_{it}| \in [c_{ix}, c_{ix+1}]$.

The existence of an equilibrium is proved in Casella (2005). Once that is established, the characterization of the equilibrium strategies follows directly from monotonicity. If all best response strategies are monotonic, equilibrium strategies must be monotonic, and given a fixed number of available votes and a continuum of possible valuations, they must take the form of monotonic cutpoints. The cutpoints depend on the state of the game and two or more may coincide (some feasible numbers of votes may never be cast in equilibrium at a given state). The equilibrium need not be unique.

The monotonicity of the equilibrium cutpoints supports our intuitive understanding of how storable votes might lead to welfare gains relative to non-storable votes (the reference case where each voter casts one vote each period). Very simply, by shifting votes from low to high realizations of $|v|$, a voter shifts the probability of obtaining the desired outcome towards decisions over which he feels more intensely. The effect appears clearly, and can be proved rigorously, in the transparent case of two voters for any arbitrary horizon length. With more than two voters, matters are more complicated because the number of asymmetric states—states where voters have different stocks of available votes—multiplies. In these states voting strategies are asymmetric, and in comparing across different voters, a larger number of votes cast need not always correlate with stronger preferences, breaking the link that underpins the expected welfare gains. As discussed in Casella (2005), and as intuition suggests, the welfare gains appear to be robust when the horizon is long enough, and the option value of a vote sufficiently important, or when the group is not too small. In fact, for any number of voters and proposals, there always exists a cooperative strategy such that the storable votes mechanism ex ante Pareto-dominates non-storable votes. In what follows we concentrate on symmetric games where all voters are given the same endowment of bonus votes ($B_i^0 = B^0$ for all i), and consider symmetric equilibria where voters play the same strategy when they are in the same state.

3. Welfare. *Call* $EV_0(B^0)$ *the expected value of the game at time* 0, *before the realization of any valuation, when votes are storable and each voter has a stock of* B^0 *bonus votes, and* EW_0 *the corresponding value with non-storable votes. Then*: (i) *For* $n = 2$, $EV_0(B^0) > EW_0$ *for all* $T > 1$, *with* $EV_0(B^0)/EW_0$ *monotonically increasing in* T. (ii) *Let the valuations' distribution* $F(v)$ *be uniform. Then for all* $n \geqslant 2$ *and* $B^0 > 1$, *there exists a monotonic cutpoint strategy profile* $x_t^C(v_{it}, B, t)$ *such that* $EV_0^C(B^0) > EW_0$ *for all* $T > 1$.

(The proofs are omitted here but can be found in Casella et al., 2003.)

With two voters, equilibrium strategies always result in expected welfare gains, relative to non-storable votes. Indeed, as shown in Casella et al., 2003, a stronger result holds: all strictly monotonic strategies, including simple rules-of-thumb, lead to expected welfare gains (for example, partitioning the interval of possible intensities [0, 1] into as many equally sized sub-intervals as the number of available votes). If the two voters choose the same rule-of-thumb, than the result holds for each voter individually; if instead they choose different but monotonic strategies, then the result holds in the aggregate: total expected welfare will be higher, although not necessarily each individual expected welfare. With more than two voters, ex ante welfare gains can be guaranteed by choosing cutpoints cooperatively.

3. Experimental design

All sessions of the experiment were run either at the Hacker SSEL laboratory at Caltech or at the CASSEL laboratory at UCLA, with enrolled students who were recruited from the whole campus through the laboratory web sites. No subject participated in more than one session. In each session, all subjects were initially allocated T bonus votes to spend over T successive proposals (in addition to one regular vote for each proposal): $B_i^0 = T$ for all i. There were two main treatment variables: the *number of voters in a group*, n; and the *number of proposals*, T. Overall we considered six different treatments: $n = 2, 3, 6$ for each of two possible horizons $T = 2, 3$.[6] The experimental design is given in Table 1.

After entering the computer laboratory, the subjects were seated randomly in booths separated by partitions and assigned ID numbers corresponding to their computer terminal[7]; when everyone was seated the experimenter read aloud the instructions, and questions were answered publicly.[8] The session then began. Subjects were matched randomly in groups of n each. Valuations were drawn randomly by the computer independently for each subject and could be any integer value between -100 and 100 (excluding 0),[9] with equal probability. Each subject was shown his valuation for the first proposal and asked to choose how many votes to cast in the first election. After everyone in a group had voted, the computer screen showed to each subject the result of the vote for the group, and the number of votes cast by all other group members.

Table 1
Experimental design

Session	n	T	Subject pool	# Subjects	Rounds
c1	2	2	Caltech	10	30
c2	2	2	Caltech	10	20
c3	2	3	Caltech	10	30
c4	2	3	Caltech	8	30
c5	3	2	Caltech	12	30
c6	3	3	Caltech	9	30
c7	6	2	Caltech	12	30
u1	2	2	UCLA	16	30
u2	2	3	UCLA	20	30
u3	3	2	UCLA	21	30
u4	3	3	UCLA	18	30
u5	6	2	UCLA	18	30
u6	6	3	UCLA	18	30

[6] In addition to our core treatments, we ran one session at UCLA with $n = 2$, $T = 2$ but $B_i^0 = 3$ for all i. We discuss it separately, later in the text.

[7] We used the Multistage Game software package developed jointly between the SSEL and CASSEL labs. This open source software can be downloaded from http://research.cassel.ucla.edu/software.htm.

[8] A sample of the instructions from one of the sessions is reproduced in Casella et al. (2003).

[9] In the first session, $c1$, one subject was assigned a valuation of 0 due to a programming error, which was corrected for later sessions. This observation is treated as missing data.

Table 2
Equilibrium and efficiency

$T = 2$ (2 *bonus votes*)

n	(c_{12}, c_{23})	surplus sv	surplus nsv	surplus sv^C	(c^C_{12}, c^C_{23})
2	(50,50)	37	0	44	(33,66)
3	$(35,67)^A$, $(50,50)^B$, $(0,100)^N$	82.6^A 80.4^B 85.7^N	85.7	89	(50,100)
4	(44,56)	83.3	73.2	84.5	(38,74)
5	$(50,50)^B$, $(0,100)^N$	86^B 88.3^N	88.3	90.8	(56,100)
6	(45,55)	86.5	80.8	88.5	(43,82)
7	$(50,50)^B$, $(0,100)^N$	87.5^B 88.6^N	88.6	91.1	(58,100)

$T = 3$ (3 *bonus votes*)

n	(c_{12}, c_{23}, c_{34})	surplus sv	surplus nsv	surplus sv^C	$(c^C_{12}, c^C_{23}, c^C_{34})$
2	(37,75,100)	52.6	0	58.2	(37,73,100)
3	$(47,66,84)^A$, $(48,66,84)^B$, $(47,69,84)^N$	$80.7^{A,B,N}$	85.7	89.3	(48,87,100)
4	(49,64,83)	84.6	73.2	88.3	(41,76,100)
5	$(53,65,77)^*$	85.7	88.3	91.8	(56,100,100)
6	(54,64,76)	86.1	80.8	91.1	(45,81,100)
7	$(56,65,73)^*$	86.3	88.6	92.3	(56,100,100)

* Although there are multiple equilibria in some of the second period subgames, the differences in first period cutpoints and expected welfare disappear when cutpoints are discretized. The experimental treatments are in bold.

Valuations over the second proposal were then drawn, and the session continued in the same fashion. After T proposals, subjects were rematched and a new sequence of T proposals occurred. Each session consisted of 30 such sequences (rounds).[10] Subjects were paid privately at the end of each session their cumulative valuations for all proposals resolved in their preferred direction, multiplied by a predetermined exchange rate. Average earnings were about \$25.00 an hour.

The parameters of the experiment mirrored closely the theoretical model, with the distribution $F(v)$ uniform over $[-100, 100]$. To construct a benchmark upper bound for efficiency, we calculated expected earnings if the decision were always resolved in favor of the side having the highest total valuation (i.e. expected ex post efficiency).[11] Because subjects in the experiments were not given the option of voting against their preferred decision, ex post efficiency and expected equilibrium payoffs are normalized using a lower bound, or "zero-efficiency" point that recognizes unanimity. We define this lower bound by the expected earnings if the decision is random whenever preferences are not unanimous.[12] We then measure surplus by the proportion of the difference between the lower bound and the upper bound that is achieved. The equilibrium cutpoints are shown in column 2 of Table 2, for different values of n (those corresponding to our experimental treatments are in bold; a few other values of n are also given for comparison). Theoretical surplus with simple majority rule (non-storable votes) is shown in column 4. The cooperative cutpoints are given in the last column, and the surplus at these cutpoints in the next to last column.

[10] With one exception: session $c2$ at Caltech had 20 rounds.

[11] See Appendix A.

[12] Other measures are possible—for example, expected earnings if the decision is always resolved in favor of the side having the lowest total valuation. The results are unchanged.

The first half of Table 2 reports the theoretical predictions in the case of two consecutive proposals and two bonus votes ($T = 2$). When the second (final) proposal is put up for a vote, all remaining votes are cast; thus the only strategic decision is how many votes to cast over the first proposal, from a minimum of one to a maximum of three. The equilibrium strategy is summarized by the equilibrium cutpoints in the second column, indicating at which valuations a voter switches from casting one vote to casting two (c_{12}) and from casting two to casting three (c_{23}): in the case of two voters, for example, the two cutpoints coincide at $|v| = 50$, indicating that casting two votes is never an equilibrium strategy. The normalized surplus measure in equilibrium is reported in the third (storable votes) and fourth (non-storable votes) columns. In the case of two voters, storable votes allow voters to appropriate 37 percent of the possible surplus over randomness, while non-storable votes lead to a random outcome whenever the two voters disagree. If voters chose their cutpoints cooperatively, the expected surplus share would rise to the number in the fifth column, and the cutpoints would be those reported in the last column (the superscript C stands for "cooperative"). With two voters, the cooperative strategy is to cast one vote for valuations smaller than 33, two votes between 33 and 66, and three votes above. This strategy is not an equilibrium, and individual deviations would be profitable, but if the two voters used this rule they could expect to appropriate 44 percent of possible surplus.

A few observations are instructive. When $n = 2$, the equilibrium is particularly simple: as said above, each voter should cast one vote if his realized valuation over the first proposal is below the mean, and three votes otherwise—he should never split his bonus votes. The equilibrium strategy is dominant: not only is the equilibrium unique, but each voter's best response strategy is not affected by the other voter's strategy.[13] Unique equilibria also hold when n equals 4 or 6, but now the equilibrium strategy is not dominant. When n is odd, the equilibrium is not unique. For the case $n = 3$, if voters 1 and 2 each cast two votes, the third voter is pivotal only if the other two voters have opposing valuations, but in this case he is pivotal regardless of the number of votes cast, and this is true in both periods. Thus always casting two votes ($c_{12} = 0$, $c_{23} = 100$) is an equilibrium, and the outcome then is identical to non-storable votes. One can see immediately that the result holds for all n odd, if $T = 2$. But other equilibria exist too, the more robust being $c_{12} = c_{23} = 50$ (never cast two votes, and switch from one to three votes at the mean valuation) which again can be shown to exist for all n odd.[14]

The efficiency difference between storable and non-storable votes is greatest in the case of two voters, although it remains positive, if declining, for all n even. When n is odd and small and the horizon is short, non-storable votes may yield a higher equilibrium surplus than storable votes. Typically, storable votes lead to states where the number of available votes differs across voters; in some cases, for example in the terminal period, the number of votes cast does not reflect the relative intensity of preferences across voters, and if a minority of the voters comes to dominate this last choice, the efficiency implications can be undesirable. And since non-storable votes perform quite well with n odd, the higher efficiency of storable votes in symmetrical states need not be enough to overcome the problem that can arise in asymmetrical states. However, although not shown in the table, as the number of voters increases, the storable votes mechanism eventually is more efficient than non-storable votes, whether n is odd or even.

[13] With $n = 2$ and $T = 2$, casting one more vote always increases a voter's probability of being pivotal at $t = 1$ by exactly the same amount it reduces it at $t = 2$. Thus the voter's optimal strategy is to spend all bonus votes in $t = 1$ if $|v_{i1}| > 50$ and save them all otherwise.

[14] Unlike in the two-player game, this strategy is not dominant.

The efficiency of storable votes can be improved by choosing the cutpoints cooperatively, in which case, as stated earlier, the storable votes mechanism always dominates non-storable votes in terms of ex ante welfare. In general, the cooperative strategy has voters casting two votes for a larger range of valuations than the equilibrium strategy, while casting three votes becomes less likely and never occurs when n is odd.[15]

The second half of Table 2 summarizes the theoretical predictions when three successive proposals are considered, and voters are given three initial bonus votes ($T = 3$). We report here the equilibrium and the cooperative cutpoints for the first proposal only, when voters can cast any number of votes between 1 and 4.[16] Two features of the equilibrium are worth noticing. First, although we know from the $T = 2$ case that some of the second period subgames have multiple equilibria, the equilibrium cutpoints induced in the first proposal election are empirically indistinguishable, once we constrain them to be discrete numbers (with the exception of the 3-voter case). The same conclusion applies to expected equilibrium payoffs: the second period multiplicity of equilibria does not translate into detectable multiplicity in expected surplus. Second, as one would expect when the horizon lengthens, the equilibrium cutpoints are now strongly asymmetrical, relative to the mean valuation: voters are always at least twice as likely to use no bonus votes than to use all three. Still, with the exception of the 2-voter case, there is a sizable range of valuations for which using all bonus votes is an equilibrium, in clear contrast to the cooperative strategy, where casting four votes should never occur. Indeed for $n = 5$ and $n = 7$, the expected payoff is maximized when voters cast at most on bonus vote on the first election.[17]

If votes are non-storable, lengthening the horizon has no effect on the expected surplus: valuations are independent and the one-period game repeats itself. The longer horizon has an effect when votes are storable. For $n = 2$, we have stated earlier that ex ante expected welfare cannot decrease in T, but in several of the other treatments Table 2 shows it declining, if only slightly. A similar result appears in Casella (2005) for the case $n = 3$, in a different specification of the storable votes mechanism. In that model the effect of a longer horizon is non-monotonic: after an initial decline with $T = 3$, expected welfare increases with the number of proposals. We have not solved the present model for more than three proposals, but the reader should not extrapolate from Table 2 that increasing the number of proposals must lead to lower expected welfare. As for the cooperative solution, moving from two to three proposals increases, in all treatments, the share of the efficient surplus that voters can expect to appropriate.[18]

We chose the experimental treatments according to two criteria. First, we selected cases with clearly different theoretical predictions, particularly in terms of efficiency. Second, we wanted to

---

[15] The intuition is clear for $n = 3$. Ruling out three votes at $t = 1$ rules out the possibility of the two states $(1, 1, 2)$ and $(1, 1, 3)$ at $t = T = 2$, the only states where a single voter can possibly override the opposition of the other two in a vote that does not reflect valuations (because T is the terminal period).

[16] The equilibrium cutpoints for the second proposal depend on the state. With the number of possible states at $t = 2$ equal to 4^n, we have chosen not to report the cutpoints in the paper. They are available from the authors. (See also the Appendix in Casella et al., 2003.)

[17] The cooperative strategy maximizes the expected payoff taking into account the whole path of the game. In this specific game, we can solve the problem backward, recognizing that voters will play the appropriate cooperative strategy in any future state (but for the last proposal, when all remaining votes are cast).

[18] Table 2 reports the results of the theoretical model where $F(v)$ is continuous. We have verified that all equilibria in Table 2 remain equilibria with the discrete distribution used in the experimental treatment (with no probability at 0), with the following minor corrections: (1) for $T = 2$, $n = 6$, the first cutpoint becomes 46; (2) for $T = 3$, $n = 2$, it is 36; (3) for $T = 3$, $n = 3$, all equilibria are $(48, 66, 85)$; (4) finally for $T = 3$, $n = 4$, the third cutpoint is 84. All cooperative strategies remain identical.

compare "simple" treatments, where equilibrium strategies are more straightforward, with more complex ones, where either the larger number of voters, or the longer horizon, or both, complicate calculations. Our efficiency criterion lead us to three choices: 2-voter treatments, where storable votes do much better than non-storable votes; 6-voter treatments where the efficiency differences are less pronounced, and 3-voter treatments, where again surplus differences between the two mechanisms are moderate, but now go in the opposite direction. As for the complexity of the game, the 2-voter 2-proposal case leads to the simplest equilibrium (a unique equilibrium in dominant strategies where bonus votes are never split), while increasing the number of voters from two to three to six increases the difficulty of computing equilibrium strategies. Even more so does increasing the number of proposals: calculating equilibrium strategies in the 3-proposal treatments, especially with three and six voters, is objectively difficult.

4. Results

Our motivation for investigating storable votes in the laboratory is that theory suggests it can produce efficiency gains over standard sequential voting. Given the complex structure of the equilibrium, and the highly strategic behavior predicted by game theory, it is not obvious that these gains will be achieved. Therefore, we begin by presenting our results on efficiency. Later we analyze individual choice behavior.

4.1. Efficiency

How do realized outcomes compare to the efficiency predictions of the theory? The short—and perhaps surprising—answer from our data is that realized payoffs match the theoretical predictions almost perfectly in all of our treatments, for all group sizes and number of proposals. The theory is highly successful in predicting both the welfare effects of storable votes and the mechanism's sensitivity to environmental parameters.

Figure 1(a) reports realized vs. predicted surplus levels in all sessions. The horizontal axis is the normalized surplus associated with equilibrium payoffs, where the equilibrium payoff is the (ex post) payoff that the subjects would have obtained had they all played the equilibrium strategies, given the actual valuation draws in the experiment. The vertical axis is the normalized surplus calculated from realized aggregate payoffs in the experiment.[19] There is one data point for each session, with the lighter points corresponding to UCLA sessions and the darker points to Caltech; the larger dots are 3-proposal sessions and the smaller dots 2-proposal sessions. Points on the 45° line represent sessions where the realized aggregate payoff equals the theoretical prediction, and points above (below) the line are sessions with realized payoffs above (below) the equilibrium payoff. Overall, the realized efficiencies align nicely with the theoretical predictions. The largest deviation from the equilibrium payoff is the 3-voter 3-proposal treatment at UCLA, but the subject pool does not appear to be an important factor for aggregate efficiency, nor is there evidence that the higher complexity of the 3-proposal game results in systematically lower payoffs.

To make a more formal statement of how well the data tracks the theory, we fit a regression line to the data in Fig. 1(a). The slope is not significantly different from one and the intercept is

[19] The ex post efficient and the random payoffs—the reference points of our surplus measure—are both calculated from the actual experimental valuations. Random payoffs are obtained from aggregating 50 percent of individuals' valuations whenever subjects are not unanimous.

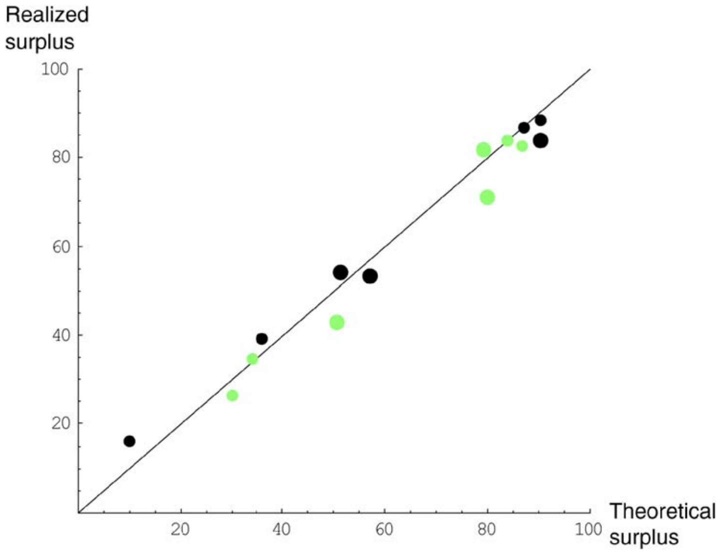

(a)

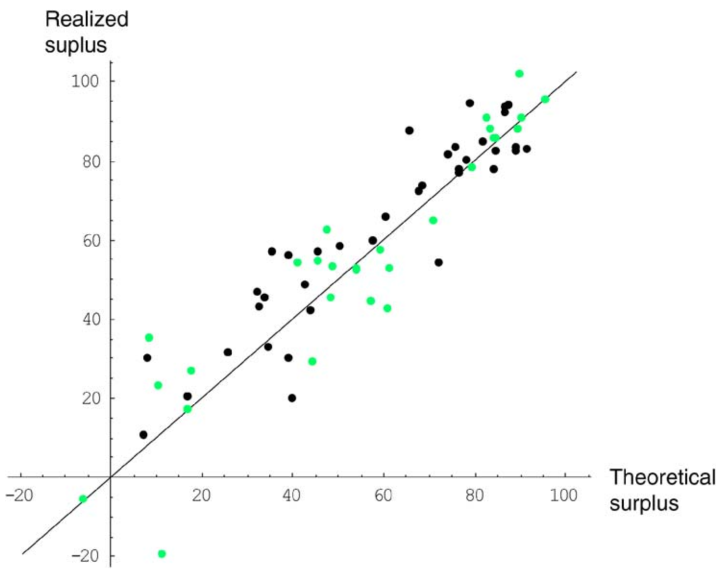

(b)

Fig. 1. Realized vs. theoretical surplus. All experiments. (a) Aggregate data by session. Caltech data in black; UCLA data in gray (green in the web version). The larger circles correspond to the $T = 3$ experiments. (b) Group data. Caltech data in black; UCLA data in gray (green in the web version).

not significantly different from zero, at the 5 percent level.[20] The R^2 is equal to 0.975. A second way of seeing how well the aggregate efficiency data track the theory is to look at comparative

[20] The estimated constant is 3.37 with a standard error of 2.88, and the estimated slope is 0.92 with a standard error of 0.043.

static predictions. Across all six treatments (in fact, each session of each treatment is actually a "subtreatment" since the random draws are different, implying slightly higher or lower theoretical predictions), the data mirror the theory well, whether the share of potential surplus that voters are expected to appropriate is in the lower range (below 60 percent, for all 2-voter treatments), or in the higher range (above 80 percent for the 3 and 6-voter treatments). The experimental data respond to changes in the environment as the theoretical model does.

A possibility that must be considered is whether perhaps the close fit of the aggregate payoffs to the theoretical payoffs is masking a large variance in more disaggregated data. Given our measure of surplus, the natural unit to be evaluated is the group—the committee voting over a sequence of proposals, and Fig. 1(b) replicates the previous figure at the group level (in the absence of systematic differences, no distinction is made here between 2 and 3-proposal treatments).[21] The smaller number of draws for each group translates into a larger range of outcomes, but the figure does not show an unexpected level of dispersion. If we fit a regression line to the group data, the R^2 equals 0.85, the slope is again not significantly different from one and the intercept barely positive, at the 5 percent level.[22]

From a practical standpoint, the interesting question is whether the storable votes mechanism is a desirable and workable voting mechanism, and more concretely whether it leads to better

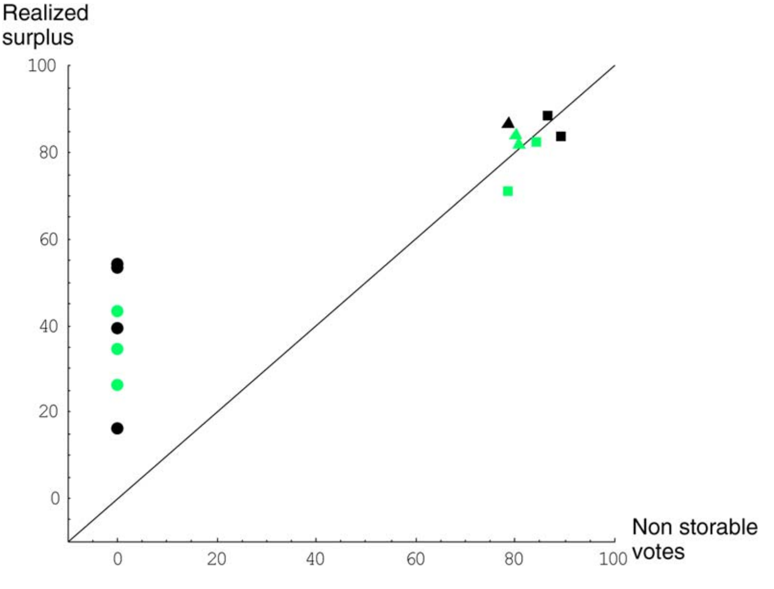

Fig. 2. Comparison to non-storable votes. Circles are $n = 2$ experiments, squares $n = 3$ and triangles $n = 6$. Caltech data in black; UCLA data in gray (green in the web version).

[21] Because individual voters are rematched after each series of votes, the group has no fixed identity—in any experiment the composition of a committee identified by the same label (for example committee 1) is unstable. Aggregating the payoff to a committee with given label in any single experiment is just an averaging device smoothing the noise in each individual round.

[22] The estimated constant is 6.62 with a standard error of 3.08, and the estimated slope is 0.93 with a standard error of 0.05.

outcomes than non-storable votes. In our experiments, subjects were able to extract on average 93 percent of the available welfare gain over non-storable votes. But we can also compare more directly the outcomes from the two mechanisms. Figure 2 plots realized aggregate surplus in each experiment versus aggregate surplus had subjects cast a single vote in all elections, with the realized valuations, distinguishing among treatments according to group size. Once again, the experimental results across different treatments match the theoretical predictions: the larger gains from storable votes appear in the 2-voter experiments, while the difference between the two mechanisms is much smaller for 3-voter and 6-voter sessions. Storable votes are associated with higher surplus than non-storable votes in all treatments, with the exception of three of the 3-voter sessions. The average surplus gain over non-storable votes is 38 percent in 2-voter sessions, 4 percent with 6 voters and −3 percent with 3 voters.[23]

4.2. *Individual behavior*

We turn now to individual behavior. Did the theory predict outcomes so well because individuals indeed followed the equilibrium strategies?

4.2.1. *Two-proposal sessions*

We begin by analyzing our data in the 2-proposal treatments ($T = 2$), because this is the simplest case: the only decision concerns the first proposal and there is only one relevant state—each voter has three available votes. Figure 3 displays the voting behavior of a sample of subjects from 2-proposal sessions. Each graph summarizes the behavior of one subject. The horizontal axis is the (absolute) valuation of proposal 1, and the vertical axis the number of votes cast in election 1. The dots correspond to the 30 rounds of decision/valuation pairs for that subject.

The figure is organized in three subclasses, according to how strictly subjects followed monotonic strategies. The first subject at the very top is perfectly monotonic: the number of votes cast at higher valuations is *always* at least as large as the number cast at lower valuations. This behavior is present in our data, but less frequent than the behavior shown in the second subclass: subjects who are *almost* perfectly monotonic: the minimum number of voting choices that would have to be changed to achieve perfect monotonicity is very small, one or two in the examples shown here. As we discuss in more detail below, almost perfect monotonicity is by far the most common pattern in the data. Finally, the last subject in the figure is a rare example of apparently erratic behavior.

Figure 3 also illustrates a second important feature of our data: monotonicity is consistent with a wide range of individual behaviors. Cutpoints need not be interior (see the subject casting 3 votes at almost all valuations), and if they are interior, they need not replicate the theoretical equilibrium cutpoints ($c_{12} = c_{23} = 50$ is the dominant strategy in this treatment). The most common behavior we observed in the data is followed by the second subject in the figure: cast one vote at low valuations, two at intermediate valuations and three at high ones (with a few monotonicity violations), but the best-fitting cutpoints differed across subjects, and clearly differed from Nash equilibrium. Thus, while monotonicity on the whole is strongly supported by the data, Nash equilibrium behavior is not. In what follows, we systematically explore the extent to which these features are confirmed in the whole data set.

[23] The range of outcomes in 2-voter sessions reflect the empirical variability in the frequency of unanimous preferences, a likely scenario with two voters.

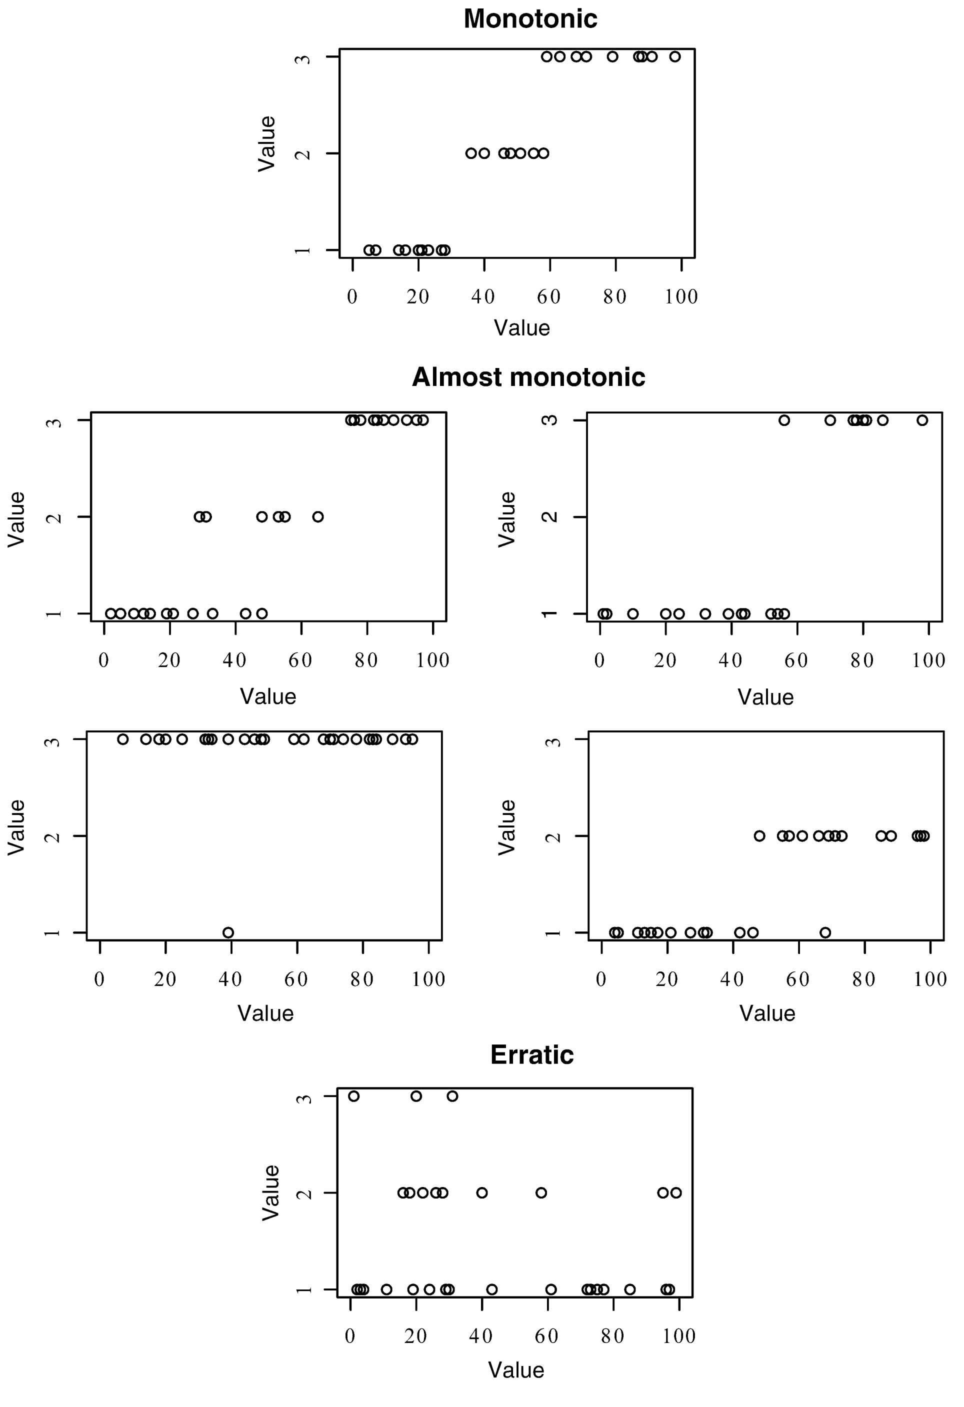

Fig. 3. Data from some example individuals ($n = 2$, $T = 2$). Each panel corresponds to an individual subject.

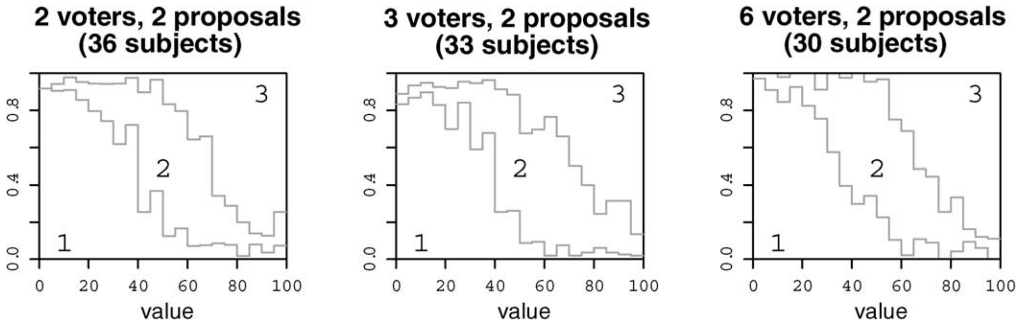

Fig. 4. Empirical frequency of votes for 2-proposal experiments.

Monotonicity Figure 4 displays the observed frequency of the three possible voting decisions (cast one, two, or three votes) as function of (absolute) valuation in the three treatments $n = 2, 3, 6$, aggregated over all sessions. We have partitioned all draws of (absolute) valuations into bins with an equal number of sample points in each bin, and plotted the observed frequency of the voting choice corresponding to each bin, so that the total equals one.

In all treatments, the observed frequency of "vote 1" decisions is close to one at very low valuations and approaches zero for valuations close to 100; the reverse is true for "vote 3" decisions, while two votes are mostly cast at intermediate valuations. The frequencies are not perfectly monotonic—for example, in the 2-voter game we observe a higher frequency of one's for valuations between 45 and 50 than between 40 and 45 (or between 95 and 100 than between 80 and 85). But the apparent violations are not per se very meaningful. Although the figure is an informative summary of the aggregate features of the data, it does not allow us to read individual behavior: a subject who always casts three votes, for example, follows a weakly monotonic strategy, but could induce an apparent violation of monotonicity in the figure if he happened not to draw intermediate valuations (while a subject who does violate monotonicity in his or her individual strategy need not induce an upward jump in these curves if that behavior is more than compensated by that of others).

Figure 5 shows the histograms of all individuals' error rates for each treatment—the minimum number of voting decisions that for each subject would have to be changed to achieve perfect monotonicity. As a comparison, the last histogram is obtained from a simulation where each voter casts one, two, or three votes with equal probability at all valuations (with 21 subjects and 30 rounds).[24] In the random simulation, only two voters have an error rate (just) below 40%; in contrast, in the actual data the number of subjects with error rates below 40% is 35 out of 36 in the 2-voter game, 30 out of 33 in the 3-voter game, and 30 out of 30 in the 6-voter game. In every session, more than half of the subjects had error rates below 10% (i.e., zero, one, or two violations of monotonicity out of 30 decisions).

A natural question is whether subjects are learning to employ monotonic strategies as they gain experience. We divided the data for each treatment into two subsamples—rounds 1–10 and 11–30—and calculated the error minimizing cutpoints separately for each subsample. The percentage of subjects with error rates below 5 percent increases in the later rounds, going from

[24] With three possible choices the error rate associated with random behavior tends to $2/3$ asymptotically, but is lower in the small sample simulation because of the ex post estimation of the best-fitting cutpoints.

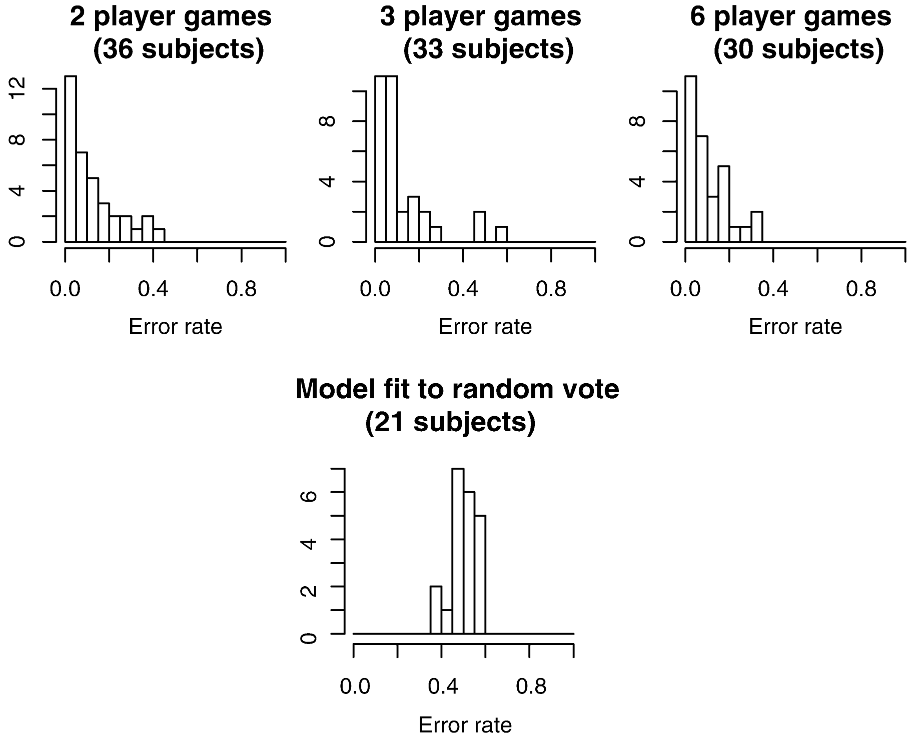

Fig. 5. Histograms of individual subjects' error rates for 2-proposal experiments (cutpoints estimated to minimize each subject's error rate). The vertical axis is the number of subjects; the scale in each panel is the relative frequency.

45 percent to above 60 percent on average over the three treatments, and reaching at least 50 percent for each treatment in the second subsample.[25]

Cutpoints The data support the hypothesis of monotonic strategies. A more difficult question is whether the estimated cutpoints are consistent with the theoretical equilibrium cutpoints. Figure 6 reports the estimated cutpoints for the three treatments. The horizontal axis measures the cutpoint at which a subject switches from one to two votes (c_{12}), and the vertical axis the cutpoint from two to three votes (c_{23}); points on the diagonal correspond to strategies such that the two cutpoints coincide (never cast two votes). Monotonicity is built into the definition of the cutpoints and the estimation method, and implies that all estimated cutpoints must lie on or above

[25] Error rates do not distinguish between violations with small impact on payoffs (hesitations over the correct voting strategy at intermediate valuations) and those revealing more fundamental confusion about the game (casting bonus votes at very low valuations and not casting any at very high valuations). A measure of the severity of the monotonicity violations is the minimum average error distance (that is, the average error distance that results from cutpoints estimated to minimize such a distance). We find some differences between the Caltech and UCLA subject pools (over all treatments, there are 2 Caltech subjects (out of 44) with average error distance larger than 5, but there are 9 UCLA subjects, out of 55). But combining the two subject pools, the results are very similar to those obtained from minimizing and counting error rates. We also estimated for each subject an ordered logit model of the number of votes cast against the (absolute) valuation. The model yields best-fit curves and estimated cutpoints, and identifies violations of monotonicity as voting "errors." The results are virtually identical to those reported in the text.

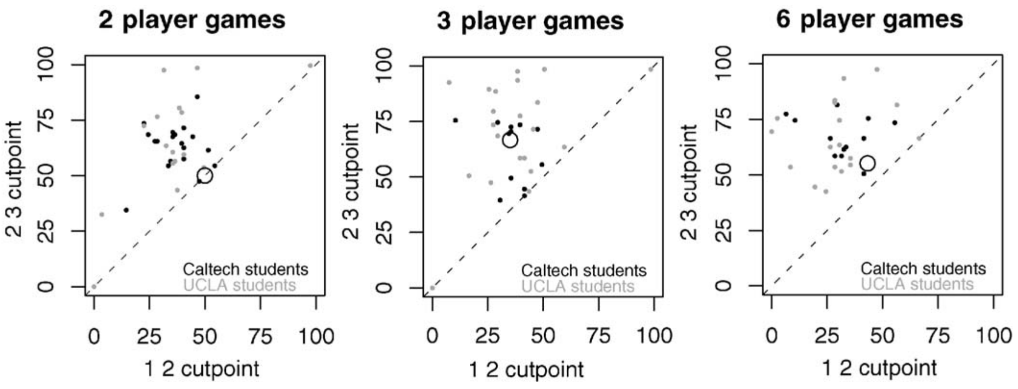

Fig. 6. Individuals' cutpoints, estimated to minimize error rates for 2-proposal experiments (circles show theoretical equilibrium values).

Table 3
Median cutpoints

Population	n	# Subjects	Median c_{12}, c_{23}	Equil. c_{12}, c_{23}	Coop. c_{12}, c_{23}
Caltech	2	10	36, 68	50, 50	33, 66
Caltech	2	10	38, 58	50, 50	33, 66
UCLA	2	16	37, 62	50, 50	33, 66
Caltech	3	12	36, 70	35, 67*	50, 100
UCLA	3	21	38, 73	35, 67*	50, 100
Caltech	6	12	32, 66	45, 55	43, 82
UCLA	6	18	30, 65	45, 55	43, 82

* Equilibrium cutpoints closest to the observed cutpoints.

the diagonal. The three corners of the box consistent with monotonicity correspond to weakly monotonic strategies: the axes origin corresponds to "always cast three votes," the top left corner to "always cast two votes," and the top right corner to "always cast one vote." The darker symbols refer to Caltech subjects, the lighter ones to UCLA. For both subsamples the figure shows large dispersion in the estimated cutpoint values around the equilibrium values.[26]

Table 3 summarizes the median cutpoints in each of the 2-proposal sessions (a statistic we use instead of the mean because of outliers), together with the equilibrium cutpoints and the cooperative cutpoints.

The most noticeable feature of the table is the similarity of the observed cutpoints across treatments, a similarity that contrasts with both equilibrium and cooperative predictions. The table can suggest strategies approaching cooperation in the case of two voters, or approaching the equilibrium predictions in the case of three (in both cases with added noise that is extraneous to the theory). But when the three treatments are observed together, a rule-of-thumb behavior, monotonic but noisy, seems a more convincing reading of the data—a point to which we return later when estimating alternative models.

[26] For some subjects, there are multiple cutpoints that minimize the number of monotonicity violations. The figure presents cutpoints estimated at the lowest value consistent with minimizing each subject's number of violations.

The similarity of the cutpoints across treatments remains true if we divide the data into early and late trials. The data suggest that learning, if it occurs, is weak: the distance between the estimated and the equilibrium cutpoints tends to decline in the later rounds, but remains large. For example, in the 2-voter treatments the distance falls in the late trials for 25 out of 36 subjects, but the median cutpoints move only from $(33, 61)$ to $(36, 63)$ (relative to equilibrium values of $(50, 50)$).

4.2.2. Three-proposal sessions

When individuals vote over three successive proposals, in the first election everybody has the same number of bonus votes. But in the second election the distribution of available bonus votes depends on the voting decisions at $t = 1$, and the number of possible states multiplies. Describing the equilibrium strategies in the second election is then complicated, and so in fact is describing the data, because each state has to be evaluated separately. We discuss here the data from the first election.[27]

The main features are very similar to those described in the 2-proposal sessions. With very few exceptions, subjects employed monotonic strategies, with a small number of errors: in all treatments, more than half of all subjects had error rates below 10 percent.[28]

But once again their strategies were more similar across treatments than theory suggests. Figure 7 depicts the aggregate frequency of the different voting choices in the three treatments (recall that the equilibrium cutpoints are presented in Table 2). Over a large range of valuations, subjects cast one vote with high probability, while they were clearly reluctant to cast four votes. To some extent, these choices match the theory: for example, in 2-voter sessions the equilibrium strategy has voters never casting four votes, and in 6-voter sessions it has them casting one vote for a majority of possible valuations. However, equilibrium strategies differ across treatments more than the data: in the 2-voter treatments, the frequency of voting one seems too high, and in the 6-voter treatments, the frequency of voting four too low.

Because we have set $B_0 = T$, by adding one proposal we have also added one bonus vote. To disentangle what the chosen strategies owe to the longer horizon per se, we have run a 2-proposal 2-voter session with three bonus votes. The frequency of the voting choices in the first election of this treatment is presented at the bottom of Fig. 7. The figure shows clearly that the length of the horizon does matter: as theory suggests, the propensity to cast three and four votes is much higher in the 2-proposal session.[29]

4.3. Relationship between individual behavior and efficiency results

Given the individual behavior shown by our subjects, the efficiency results presented earlier are surprisingly good. So much so, in fact, that one must wonder whether payoffs could be rather insensitive to the strategies played. To check this possibility we have simulated the payoffs

[27] The appendix of Casella et al. (2003) discusses the data from the second election in the $n = 2$ games, where the number of states remains small enough to be tractable.

[28] For comparison, we simulated a model with random voting (with 18 subjects and 30 rounds). The minimum error rate was 45 percent, and the median 65 percent. With four possible choices the error rate associated with random behavior tends to $3/4$ asymptotically (although it is lower in the small sample simulation).

[29] The session was run at UCLA with 20 subjects and 30 rounds. The dominant strategy is to cast one vote for valuations smaller or equal to 50, and four votes otherwise.

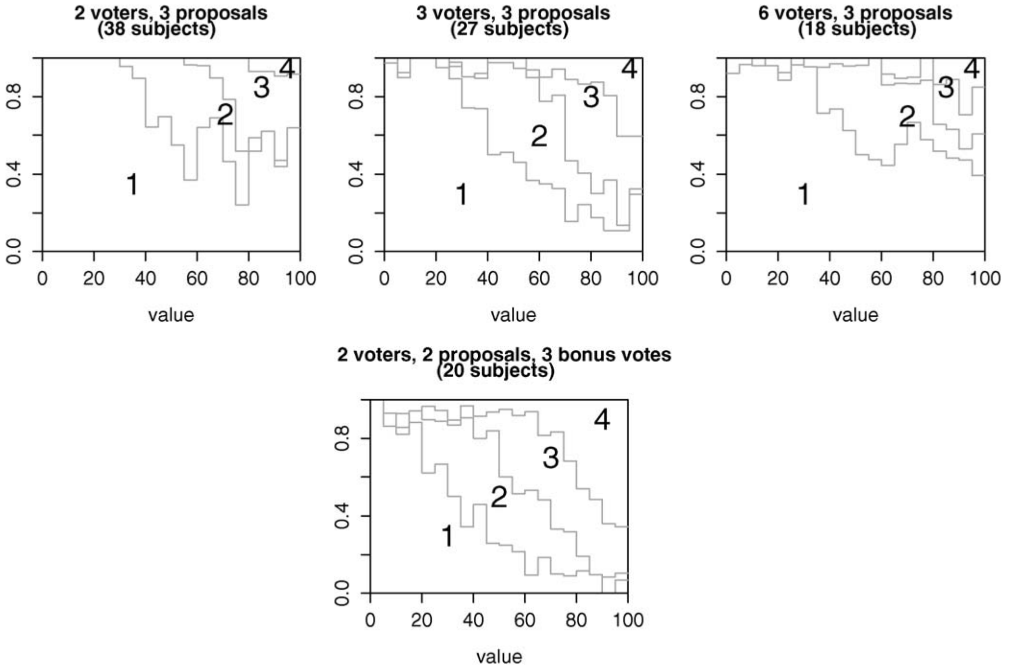

Fig. 7. Empirical frequencies of votes. 3-bonus votes.

that the subjects would have obtained in the experiments had they chosen the number of votes randomly, over their feasible alternatives.[30] In every treatment, the payoff would have been lower.

The result underscores the importance of monotonicity. Storable votes can lead to efficiency gains because voters can express the intensity of their preferences by casting more votes when their valuations are higher; even if the cutpoints are incorrect, as long as strategies are monotonic (even with a few errors), outcomes will reflect strength of preference and the essence of the mechanism is captured. And this is also why the efficiency results for the experiments are not systematically different for the 3-proposal treatments: the higher complexity of the equilibrium strategies has limited importance if subjects can approach the equilibrium payoffs by following the simpler criterion of monotonicity.

In the case of two voters, the conjecture can be supported by formal argument: as shown in Casella et al. (2003), the storable votes mechanism leads to higher expected aggregate payoffs, relative to non-storable votes, whenever the two voters follow monotonic strategies for any arbitrary value of the thresholds, and strictly higher if at least one of the thresholds is strictly interior. What we could not anticipate from the theory is how close to the equilibrium payoffs the experimental payoffs would be. Nor could we anticipate that the central role of monotonicity in the 2-voter case would extend to our other experimental treatments.

What the experimental results emphasize is that as long as strategies are monotonic, the payoff functions are rather flat at the top—the loss from not choosing the correct thresholds is small. Figure 8(a) illustrates this point in a graph; we have drawn it for the more transparent case of two

[30] The payoffs from randomization are the averages over 100 realizations (where each realization is a full 2 or 3-proposal game). Subjects always cast all their votes in the last election.

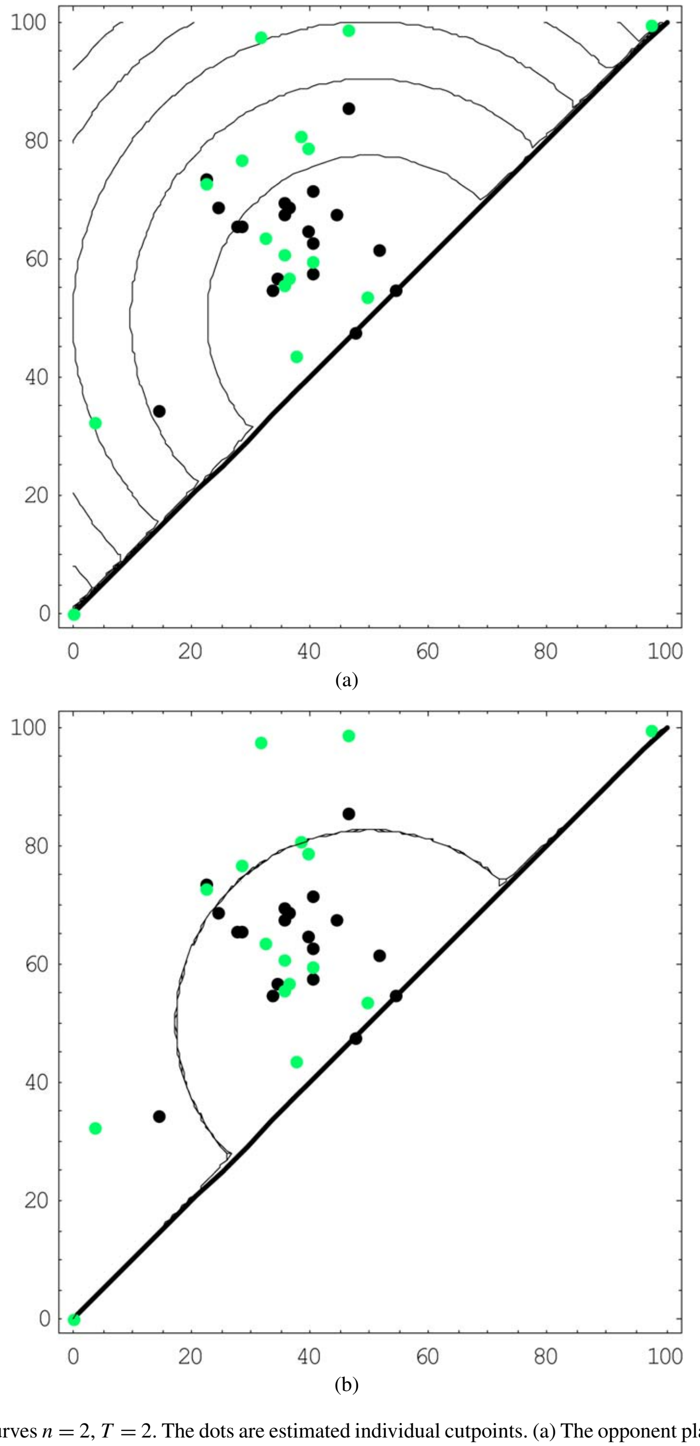

Fig. 8. Isopayoff curves $n = 2$, $T = 2$. The dots are estimated individual cutpoints. (a) The opponent plays the equilibrium strategy. Each contour is a loss of 7.5 percent of potential payoff. (b) The opponent plays the average estimated strategy. The isopayoff curve is drawn at the level of the equilibrium payoff.

voters (and two proposals), but its lessons apply to all treatments. The figure depicts individual expected isopayoff curves, when the other voter follows the equilibrium strategy. The horizontal axis is the first cutpoint (c_{12}), the vertical axis the second (c_{23}). Recall that the equilibrium strategy, and hence the highest payoff, corresponds to $c_{12} = c_{23} = 50$, the center of the square. Every isopayoff contour, moving away from the center, indicates a payoff loss of 7.5 percentage points, reaching down to the zero surplus associated with randomness (and in 2-voter treatments with non-storable votes) at the three corners, with a cumulative loss of 37 percent. The dots in the figure are the individual cutpoints estimated from the data and reported earlier in Fig. 6. The figure makes precise our observation about the flatness of the expected payoff function: the area within the first contour is large enough to encompass more than half of all of our data points.

Figure 8(a) ignores a second possible reason for the high off-equilibrium efficiency of the data—the possibility that the observed strategies are in fact closer to the cooperative choices. Figure 8(b) represents the individual expected isopayoff curve at the equilibrium level, again in the case of two voters and two proposals, when the other voter is playing the (estimated) average strategy. The lowest expected payoff is at the three corner points, and again corresponds to the expected payoff with non-storable votes. The strategy $c_{12} = c_{23} = 50$ is dominant and thus leads once more to the highest expected payoff, but now there is a whole region of (non-equilibrium) cutpoints that yields higher than equilibrium expected payoff. Two thirds of our estimated cutpoints (25 out of 36) belong to the region.

The bottom line of the analysis is clear. The efficiency gains from storable votes appear to be robust with respect to deviations from equilibrium strategies, provided that subjects are using monotone cutpoints. The next section explores more deeply the deviations from equilibrium behavior by investigating several alternative models of aggregate behavior that allow for random variations from monotone cutpoint strategies.

4.4. Quantal response equilibrium

The results above show that the basic monotone structure of strategies is reflected in the data, but there are some clear violations of the theoretical predictions. We see some nonmonotonicities for nearly every subject in every treatment, most estimated cutpoints differ from their equilibrium values, and there is little support even for the dominant strategy in the simplest treatment (the 2-proposal, 2-voter scenario where bonus votes should never be split). All these features are inconsistent with the perfect Bayesian equilibrium of the game.

In this section we estimate a stochastic choice model of behavior. While a standard procedure such as logit or probit is a reasonable first step (and we have used it to check the robustness of the cutpoint estimates reported so far), it is not completely adequate in the context of strategic games. The reason is that stochastic behavior by one player changes the other players' expected payoffs from different strategies (even with dominant strategies), and therefore can change the equilibrium. Moreover, even if stochastic behavior does not change players' best responses, in standard models of stochastic choice (such as logit) it will still affect the predicted stochastic choices of the other players, because it changes their expected payoffs. Only in a stochastic choice model where choice probabilities are unresponsive to payoffs would this interaction effect not be present. Therefore, what is needed is a more elaborate model that incorporates not only stochastic choices, but also the endogenous equilibrium effects.

Quantal response equilibrium (McKelvey and Palfrey, 1995, 1996, 1998) is a model that embodies stochastic choice into the standard noncooperative game approach. It solves the problem of stochastic choice interactive effects by looking at an equilibrium in which players' choices react stochastically to expected payoffs, while (in equilibrium) the expected payoffs are themselves a function of the stochastic choice behavior of the other players. This results in a generalization of Nash equilibrium to allow for stochastic choice.

For a finite n-player game, let K_i be the number of strategies available to player i. Let $\sigma = (\sigma_1, \ldots, \sigma_n) \in \Sigma$ be a mixed strategy profile, and let $u_i : S \to \mathbb{R}$ be i's payoff function. Denote by $U_{ik}(\sigma)$ the expected payoff to player i from using strategy k, when the other players are using profile σ_{-i} and let $U_i(\sigma) = (U_{i1}(\sigma), \ldots, U_{iK_i}(\sigma))$. We define a *quantal response function* as a mapping from utilities into choice probabilities, that is a function that maps $U_i(\sigma)$ into a K_i-vector of choice probabilities for player i. As is typical in applications we require such a function, $Q_i(U_i) = (Q_{i1}(U_i), \ldots, Q_{iK_i}(U_i))$, to be *interior*, *continuous*, and *payoff responsive* (see McKelvey and Palfrey, 1995, for details).[31] Interiority requires $Q_{ij}(U_i) > 0$ for all i, j, $U_i \in \mathbb{R}^{K_i}$. Continuity requires $Q_i(U_i)$ to be continuous for all i, $U_i \in \mathbb{R}^{K_i}$. Payoff responsiveness requires: (1) $U_{ij} > U_{ik} \to Q_{ij}(U_i) > Q_{ik}(U_i)$ for all i, j, $U_i \in \mathbb{R}^{K_i}$; and (2) $Q_{ij}(U_i)$ is weakly increasing in U_{ij} for all i, j. A *quantal response equilibrium* (QRE) is a strategy profile $\sigma^* = (\sigma_1^*, \ldots, \sigma_n^*) \in \Sigma$ such that $Q_i(U_i(\sigma^*)) = \sigma^*$ for all i.

4.4.1. Logit equilibrium

For estimation, we use a parametric version of QRE, logit equilibrium, which is the extension of the standard logit choice model to multiperson strategic choice problems. A *logit equilibrium* is a quantal response equilibrium in which the quantal response function is given by the standard logit response function below:

$$Q_{ij}(U_i) = \frac{\exp(\lambda U_{ij}(\sigma))}{\sum_{k=1}^{K_i} \exp(\lambda U_{ik}(\sigma))}, \tag{1}$$

where the parameter λ governs the degree of payoff responsiveness. When $\lambda = 0$, strategies are completely unresponsive to payoffs, and player i simply chooses each strategy with probability $1/K_i$. When $\lambda = \infty$, players choose best responses, and the logit equilibrium converges to the Nash equilibrium.[32] Therefore, we can write a logit equilibrium as any strategy profile $\sigma^* = (\sigma_1^*, \ldots, \sigma_n^*) \in \Sigma$ such that:

$$\sigma_{ij}^* = \frac{\exp(\lambda U_{ij}(\sigma^*))}{\sum_{k=1}^{K_i} \exp(\lambda U_{ik}(\sigma^*))} \quad \text{for all } i, j.$$

As λ is varied over $[0, \infty)$, one traces out the logit equilibrium correspondence, that is, the set of solutions to (1). This correspondence is upper hemicontinuous and its limit points, as λ tend to ∞, are Nash equilibria. In this paper, we consider the logit equilibrium correspondence of the storable votes game. Because of computational difficulties, we apply the logit equilibrium model only to the 2-proposal treatments.[33]

[31] Such response functions can be rationalized as a Bayesian equilibrium of a game of incomplete information with privately observed i.i.d. payoff perturbations.

[32] In most applications, it is assumed that λ is identical for all players, but this is not necessary. Heterogeneity with respect to λ has been explored in McKelvey et al. (2000).

[33] Quantal response equilibrium (and logit equilibrium) has also been defined for Bayesian games with continuous types (McKelvey and Palfrey, 1996), games with continuous strategy spaces (Goeree et al., 1998), and games in extensive form (McKelvey and Palfrey, 1998). The storable votes game described in the theoretical section combines all of these elements. In the experiment however, strategies and types are finite, and we use the standard model in the estimation that follows.

We study two representations of the logit equilibrium, corresponding to two different models of strategy choice in the storable votes game. In one representation, strategies are behavior strategies: a player must consider how many bonus votes to use conditional on his valuation (and, in later stages of the game, on the history of voting on past proposals). In this case, we apply the logit model to each discrete choice (0, 1, or 2 bonus votes) conditional on the player's (absolute) valuation. Each player's strategy is characterized by 100 probability distributions over 0, 1, or 2 bonus votes, strategies are cutpoint strategies: one for each possible (absolute) valuation. In the second representation, we suppose players are choosing ex ante, among the set of (weakly) monotone cutpoint strategies, before drawing valuations.[34] A monotone cutpoint strategy is a pair: given 100 possible valuations, there are 5050 distinct monotone strategies, and a logit equilibrium will be represented as a probability distribution over all of these cutpoint strategies. While the Nash equilibria are identical for the two representations, the logit equilibrium correspondences are quite different. Moreover, any logit equilibrium will imply a specific probability distribution over actions (i.e., number of bonus votes used in the first proposal) as a function of absolute valuation. As we see below, these probability distributions differ quite a bit depending on the representation of strategies we use.

4.4.2. QRE estimation

Because the logit equilibrium implies a probability distribution over actions, we can use it as a model to fit the data, by estimating the response parameter, λ, through standard maximum likelihood estimation. The derivation of the likelihood function is described in Appendix B. Table 4 presents the results of the estimation: the estimates for the behavior strategy model and the cutpoint strategy model are reported in columns 4 and 6, respectively, and the corresponding values of the log-likelihood function at the estimated value of λ are reported in columns 5 and 7, respectively.

Several observations can be made. First, the data are generally noisier in the UCLA subject pool, a fact reflected in both the value of the likelihood function and the estimated λ under both models, but especially under the cutpoint model. Second, the cutpoint model generally fits the data better in both subject pools, although the differences are not always significant. Third, $\hat{\lambda}_{\text{cut}} < \hat{\lambda}_{\text{beh}}$ in every session, a sensible result given that $\hat{\lambda}_{\text{cut}}$ has some additional "rationality" built into it: even with $\lambda_{\text{cut}} = 0$, players are still using monotone cutpoints (randomizing over all such monotone cutpoint strategies with equal probability), and hence the predicted choice behavior is highly responsive to valuation, as in the data. In contrast, if $\lambda_{\text{beh}} = 0$, choice behavior is completely random, and independent of valuation. Fourth, the estimated value of λ for both models is increasing in n. We do not have a good explanation for this, but it is an interesting and persistent finding.

Figures 9 and 10 show the implications of the QRE model for the probability distribution of votes as a function of (absolute) valuation, for the behavior strategy model as well as the cutpoint strategy model. Figure 9 shows, for session c1, the expected number of votes using the estimated value of $\hat{\lambda}_{\text{beh}}$ (0.46) and $\hat{\lambda}_{\text{cut}}$ (0.25). The darker curve corresponds to behavior strategies, the lighter one to cutpoint strategies. The data are superimposed, and each dot represents the empirical average number of votes, as function of absolute valuation. From this graph, it is clear that there is not much difference between the models in terms of expected number of votes cast.

[34] Therefore the cutpoint model is a logit equilibrium of the game with a restricted set of strategies.

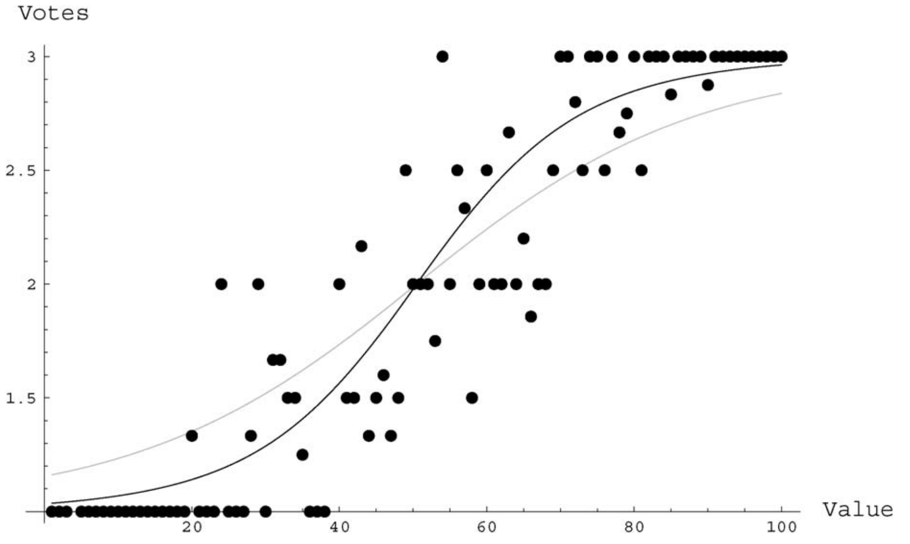

Fig. 9. Expected number of votes in the two QRE models session c1. The darker curve is estimated from the behavior strategies model, the lighter curve from the cutpoint strategies model.

However, the estimated distribution of votes cast is quite different in the two models. This is shown in Fig. 10, which displays the relative frequencies of casting 1, 2, or 3 votes as function of valuation for each session and for both QRE models. For each session, the graph on the left is for behavior strategies, and the one on the right is for cutpoint strategies. In each graph, the horizontal axis is absolute valuation, which ranges between 1 and 100, and the vertical axis is choice probability, and ranges from 0 to 1. For each valuation, the two curves in the graph partition the [0, 1] interval into three subintervals, with the size of these subintervals corresponding to the probability of casting exactly 1, 2, or 3 votes, respectively. Each graph is, for each session, the estimated version of Fig. 4.

This figure illustrates quite clearly the implications of the estimates in Table 4. First, observe that the third graph in the second column corresponds to the cutpoint model estimates for $n = 2$, $T = 2$ when $\lambda = 0$.[35] These are the predicted frequencies if behavior is monotone, but no additional rationality is assumed—players randomize over all monotonic cutpoints. Behavior remains regular, with the probability of casting 1 vote approaching 1 for low valuations, the probability of

Table 4
Results of logit equilibrium estimation

Population	n	# Obs.	$\hat{\lambda}_{\text{beh}}$	$-\ln L_{\text{beh}}$	$\hat{\lambda}_{\text{cut}}$	$-\ln L_{\text{cut}}$
Caltech	2	299	0.46	181	0.25	162
Caltech	2	200	0.52	108	0.32	102
UCLA	2	480	0.19	437	0.00	438
Caltech	3	360	0.81	248	0.37	248
UCLA	3	630	0.56	520	0.01	521
Caltech	6	360	1.66	203	0.84	200
UCLA	6	540	1.10	390	0.30	380

[35] The fifth graph in the same column is very similar, since the estimated value of λ is 0.01.

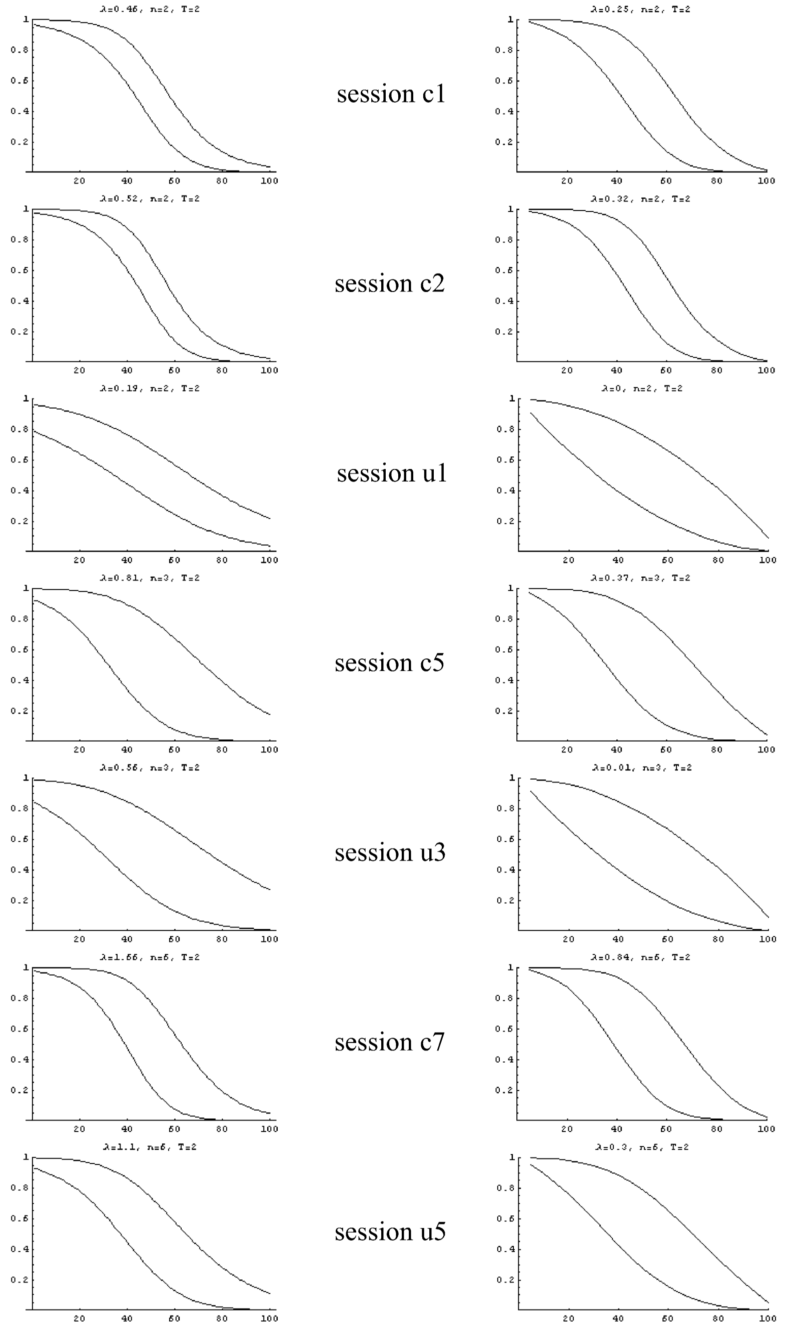

Fig. 10. QRE estimated choice frequencies. Behavior strategies on the left. Cutpoint strategies on the right.

casting 3 votes approaching 1 for high valuations, and the probability of casting 2 votes increasing as valuations approach 50 (from either direction). Second, observe that in all cases the curves for UCLA sessions are flatter than the corresponding curves for Caltech sessions, reflecting the fact that $\hat{\lambda}_{UCLA} < \hat{\lambda}_{Caltech}$ in all sessions and for both models. Third, for intermediate ranges of v, the probability of casting exactly 2 votes (the vertical distance between the two curves) is higher in the cutpoint model than in the behavior strategy model. This is one reason for the better fit of the cutpoint model.

4.4.3. Non-equilibrium models of stochastic choice

How well do our two logit equilibrium models fit the data compared to other plausible models? In addition to allowing us to better evaluate the results just described, the relative fit of the different models will help us to understand better the properties of the data, and hence the behavior of the subjects. Although we must acknowledge the heterogeneity revealed by our earlier description of the data, for consistency with the QRE estimations we limit ourselves to aggregate models that assume homogeneous behavior on the part of the players.[36]

We consider four alternative stochastic choice models, all of which are based on monotone behavior and all of which have at least one free parameter to estimate. In the first model, which we refer to as *aggregate best fit* (ABF), all players are assumed to use monotone strategies with an error rate that is not payoff dependent, but is assumed to be random over the alternative strategies. The cutpoints are constrained to be the same for all players and are estimated from the data. While the model is almost completely atheoretical (except to the extent that it assumes monotone behavior), it provides both the most natural benchmark and a particularly challenging comparison to QRE, which is instead based on a theoretical structure of equilibrium behavior. For each session, we estimate three parameters: two cutpoints, c_{12}, c_{23}, and the error rate, ϵ (reported in Appendix B). The second model is a variation on the Nash equilibrium that allows for errors, but (like the cutpoint models) assumes that errors are unrelated to equilibrium expected payoffs: individuals are assumed to choose their Nash equilibrium strategy with some probability, $1 - \epsilon$ and randomize over non-equilibrium actions otherwise. In contrast to QRE, the randomization is not taken into account by the other players, so the model is not quite an equilibrium model.[37] We call this the *noisy Nash equilibrium* (NNE) model. It has been investigated in other contexts and typically fits data better than Nash equilibrium but worse than logit equilibrium. In line with our earlier discussion of cooperation, the third model we study is a noisy version of cooperative behavior: as in the NNE case, individuals are assumed to choose their cooperative cutpoints with some probability, and randomize otherwise. We call this third model *noisy cooperative behavior* (NCB). Finally, we consider a constrained ABF model where subjects are all assumed to use, with some error, what we call "uniform cutpoint strategies." That is, they are assumed to adopt monotone cutpoint strategies that follow a simple rule of thumb: the range of possible valuations is divided in intervals of equal size so that each possible strategy has the same probability ex ante. In the case of $T = 2$, this corresponds to $c_{12} = 33$ and $c_{12} = 66$. Thus with some probability subjects vote according to these cutpoints, and with some probability they randomize.[38] The likelihood function for the ABF model is given in Appendix B, and the likelihood functions for the other two models are derived in a similar way. The (negative) log-likelihoods of these alternative models are presented in Table 5.

---

[36] Both QRE models we estimated assume that all subjects' (mixed) strategies are identical and all subjects share the same value of λ.

[37] Actually, for the $N = T = 2$ it is an equilibrium model with error. To see this, recall that in this treatment $(50, 50)$ is a dominant strategy. Therefore, in this case NNE is an equilibrium model with errors, much like QRE, but without the assumption that choice probabilities are strictly monotone in expected payoffs.

[38] 33,66 are also the expected cutpoints if subjects choose cutpoints monotonically but randomly.

Table 5
Alternative models. Log-likelihoods

Population	n	Obs.	*ABF*	QRE_{cut}	QRE_{beh}	*NNE*	*NCB*	33/66
Caltech	2	299	191	162	181	253	215	215
Caltech	2	200	108	102	108	154	128	128
UCLA	2	480	416	438	437	478	433	433
Caltech	3	360	267	248	248	323	374	287
UCLA	3	360	520	521	520	621	621	569
Caltech	6	360	228	199	203	293	260	238
UCLA	6	540	411	380	390	517	488	420

4.4.4. Discussion of estimation results

Table 5 indicates that both QRE models do much better than NNE—NNE is easily rejected. This is not really surprising, in light of the earlier description of the data, and the estimation confirms our reading of the data. NCB also fares worse than either QRE model (with the exception of the UCLA $n = 2$ session, which all models fit poorly, due to erratic behavior of a few subjects); but it does almost uniformly better than NNE (again with a single exception). By construction, the 33/66 model is identical to NCB in the 2-voter treatments, but would seem a more promising option in the other treatments, given the relative invariance of the observed cutpoints to the number of voters. Indeed, 33/66 fits the data better than NCB when the two models differ and, as expected, much better than NNE. More surprising is the inferior performance of 33/66 with respect to either QRE model (but for session $u1$). Finally, both QRE models do much better than ABF in the Caltech data; ABF does better in one UCLA session, but again it is the one session where no aggregate model fits well. This is the most surprising result: it would seem that ABF must outperform nearly any aggregate model, since, by definition, it estimates the cutpoints that best fit the data. However, while the QRE estimated cutpoints place an additional constraint on behavior (equilibrium), QRE, unlike ABF, allows errors to be correlated with expected payoffs rather than just assuming all errors are equally likely. The QRE parameter, λ, is, loosely speaking, an indicator of this correlation. Thus, for example, if $v_{i1}^k \in (0, \hat{c}_{12}]$, and thus no bonus votes should be cast, the logit equilibrium likelihood function will assign a lower likelihood to $b_{i1}^k = 2$ than it assigns to $b_{i1}^k = 1$, since $b_{i1}^k = 2$ yields lower expected utility than $b_{i1}^k = 1$. In contrast, the ABF likelihood function assigns the same likelihood to both of these observations. In the data, it is clearly the case that errors are related to expected payoffs.[39]

All of these models are aggregate models in the sense that individuals are treated as representative agents. Using the approach of the earlier section, where one obtains a best fit separately for each subject, one could improve significantly over all of the representative agent models.

Finally, we note that all of the models considered fit Caltech sessions better than UCLA sessions. There is more unexplained variation in the UCLA data, in part at least because of the presence of more outliers in the UCLA data.

[39] In addition, the QRE_{cut} model implicitly allows for some heterogeneity in the data, since different players are using different cutpoints in each round. The QRE_{beh} model does not have a simple interpretation in terms of heterogeneity of cutpoints.

5. Conclusions

The results of the experiment suggest several conclusions. First, the efficiency calculations based on the perfect Bayesian equilibrium model of behavior predicts almost perfectly the aggregate surplus for all treatments. This conclusion holds true across subject pools, across time, and across the various parametric environments.

Second, monotone cutpoint strategies, with some random deviation, appear to be used by the vast majority of subjects, again in all treatments and in both subject pools. Monotone behavior characterizes all best response strategies of the storable votes mechanism, but it is also highly intuitive: "use more bonus votes if your preferences over the current proposal are more intense." Moreover, it is monotone behavior that leads to the welfare findings. The efficiency gains from storable votes, compared to non-storable votes, derive from the ability of voters to shift the probability of obtaining their desired outcome towards those decisions that weigh more in their utility. This happens precisely because subjects use monotone strategies.

Third, the cutpoint strategies, while typically monotone, are significantly different from the perfect Bayesian equilibrium strategies, a fact we read as evidence of stochastic choice. We fit a logit equilibrium model to the data for all treatments, with two alternative representations of the subjects' strategies, one that allows for all possible behavior strategies, and one that assumes monotone cutpoint strategies. We compare their fit to four alternative models: a noisy Nash model where Nash equilibrium strategies obtain with random errors; a noisy cooperative model where subjects either behave cooperatively or make random errors; a uniform strategy model where monotone cutpoints are such that feasible voting choices are all assigned equal ex ante frequency, again with random errors, and a purely statistical model where aggregate cutpoints are directly estimated from the data by minimizing monotonicity violations. The qualitative and quantitative features of the data are best organized by the logit equilibrium model, underscoring the importance of modeling errors as negatively correlated with foregone payoffs.

The experimental treatment replicated the informational assumptions maintained in the theoretical model of storable votes proposed by Casella (2005). In particular, individuals do not know the valuations of the other voters, or their own future valuation, and valuations are independent across voters and proposals. There are obviously a wide variety of extensions to this model that can be obtained by changing these assumptions in different ways. Thus, a question of considerable merit is whether the basic intuition behind storable votes is robust: i.e., it often produces welfare improvements over standard majority rule because it allows voters to express preference intensities. While a definitive answer requires a great deal of work, we close with a simple example indicating why we think the intuition is robust.

Consider a 2-voter 2-proposal world but with different informational assumptions. In particular suppose that each individual knows both his valuations. To make it interesting, assume that preferences of the voters are opposed, so whenever one voter's valuation for a proposal is positive, the other's is negative and vice versa. This game has a unique pure strategy equilibrium, where voters use all their bonus votes on the proposal they care about most (the one with the highest absolute valuation), and this remains true regardless of whether the other's valuations are known or not; regardless of whether voting over the two proposals is simultaneous or not; regardless of the joint distribution from which valuations are drawn. Except in the case where one player's absolute valuations are higher than the other player's on both issues, this equilibrium implements the utilitarian optimum.[40] In contrast, simple majority rule produces purely random allocations, an outcome that is inferior (strictly inferior, as long as the correlation between valuations, across players or across proposals for each player, is less than perfect).

We think this very simple example captures the essence of storable votes, without any of the added features we used in the present paper. The key to the example is, as in our theoretical model and in our experiments, *monotonicity* of the voting strategies. As long as monotonicity of the equilibrium strategies is preserved, the outcomes will reflect intensity of preference, so it is hard to imagine how or why the positive welfare implications we obtain here might somehow become reversed as we vary the information structure and timing. We leave a deeper theoretical study of the robustness of the mechanism to future research.

Acknowledgments

We thank the editor and the referee for their constructive comments. We also benefitted from conversations with Avinash Dixit and with participants at the 2002 meetings of the Economic Science Association in Boston, the Association for Public Economic Theory in Paris, the 2003 Economic Theory Conference in Rhodos, and at seminars at the University of Bologna, CORE, GREQAM, Princeton, Paris, SMU, Stanford, the University of Pennsylvania and Washington University. We thank Nobuyuki Hanaki and Brian Rogers for excellent research assistance.

Appendix A. Efficient yardstick

We define as *ex post efficient* a decision-making mechanism that resolves each election in favor of the side with higher aggregate valuations. Consider for example the case $n = 2$. For each proposal, half of the time the two voters agree, and the expected valuation is 50; half of the times they disagree, and in the ex post-efficient resolution a voter wins only if his valuation is higher than the other player's. Hence for each proposal each voter's expected utility from the efficient allocation is:

$$EU^*(2) = \frac{1}{2}50 + \frac{1}{2}\left(\int_0^{100}\int_{v_j}^{100} v_i \frac{\mathrm{d}v_i}{100}\frac{\mathrm{d}v_j}{100}\right) = 41.7. \tag{A.1}$$

We can calculate expected utility under ex post efficiency in a similar manner for different numbers of voters, keeping in mind that the density function of the sum w of n random variables, each independently distributed uniformly over [0, 100], is given by:

$$P_n(w) = \frac{1}{200(n-1)!}\sum_{k=0}^{n}(-1)^k\binom{n}{k}\left(\frac{w-100k}{100}\right)^{n-1}\operatorname{sign}(w-100k), \tag{A.2}$$

where $\operatorname{sign}(x)$ is the sign of x. For arbitrary n, expected utility from the efficient allocation is then:

$$EU^*(n) = \left(\frac{1}{2}\right)^{n-1}(50) + \sum_{k=0}^{n-2}\binom{n-1}{k}\left(\frac{1}{2}\right)^{n-1}$$

[40] The exceptional case is one where the ideal solution is dictatorship, not voting. This case would not arise if we normalized both players' utilities so that the valuation on the issue they prefer least is equal to 0.

$$\times\left[50\int_{0}^{100k} P_{n-1}(w)\,\mathrm{d}w+\int_{100k}^{100(k+1)}\left(\int_{w-100k}^{100} v\frac{\mathrm{d}v}{100}\right)P_{n-1}(w)\,\mathrm{d}w\right]. \tag{A.3}$$

In contrast to storable votes, the expected efficient payoff has no temporal dimension: given the assumed independent and identical distributions of the valuations, the expected utility under ex post efficiency is the same for all proposals, and the expected utility from a sequence of T proposals is simply the sum of T one-proposal expected utilities.

The zero-efficiency reference point is the expected individual payoff if the decision is random unless all agree. With n voters, the probability of unanimity is $(1/2)^{n-1}$. Thus the expected payoff equals $50(1/2)^{n-1}+25[1-(1/2)^{n-1}]=25[1+(1/2)^{n-1}]$, and again is constant across all proposals.

The expected utility from non-storable votes is equally constant across proposals. It must equal the probability of having the proposal decided in one's preferred direction multiplied by the expected utility from such an outcome, before the valuation has been observed. Or:

$$EU^{ns}(n)=50\sum_{k=(n+I-1)/2}^{n+I-1}\binom{n+I-1}{k}\left(\frac{1}{2}\right)^{n+I-1}, \tag{A.4}$$

where

$$I=\begin{cases}1 & \text{if } n \text{ is even,}\\ 0 & \text{if } n \text{ is odd.}\end{cases}$$

Appendix B. Alternative models: estimation

Table B.1
Cutpoints and error rates

Population	n	Obs.	*ABF*			*NNE*	*CB*	33/66
			c_{12}	c_{23}	$\hat{\epsilon}$	$\hat{\epsilon}$	$\hat{\epsilon}$	$\hat{\epsilon}$
Caltech	2	299	38	70	0.20	0.32	0.24	0.24
Caltech	2	200	40	65	0.16	0.27	0.20	0.20
UCLA	2	480	39	63	0.33	0.45	0.36	0.36
Caltech	3	360	38	68	0.25	0.36	0.50	0.29
UCLA	3	630	39	74	0.30	0.43	0.43	0.36
Caltech	6	360	41	67	0.20	0.30	0.24	0.21
UCLA	6	540	32	65	0.26	0.40	0.36	0.27

B.2. Likelihood functions

B.2.1. Logit equilibrium likelihood function

For each model (behavior strategy or cutpoint strategy), we compute the logit equilibrium correspondence as a function of the precision parameter, λ. To obtain the maximum likelihood estimate of λ given a dataset (from a session) we can compute the likelihood of observing the data for each value of $\tilde{\lambda}$, for each model, under the maintained hypothesis that the model is correct and $\tilde{\lambda}$ is the true value of λ. This gives us a likelihood function (as a function of λ) that is directly implied by the structure of the model.

Consider some set of data points from a session with some fixed value of n and $T=2$. Denote by $\sigma^{*\hat{\lambda}}$ the quantal response equilibrium of the game for this value of n and $\lambda=\hat{\lambda}$. For each individual, i, in the session we have a collection of K observations of proposal 1 valuations and bonus vote choices, denoted $y_i = (v_{i1}^1, b_{i1}^1, \dots, v_{i1}^K, b_{i1}^K)$. For any round, k, and any individual, i, the equilibrium specifies probabilities that individual i uses 0, 1, or 2 bonus votes for the first proposal as a function of i's first proposal value in that round. Denote these by $\sigma_0^{*\hat{\lambda}}(v_{i1}^k)$, $\sigma_1^{*\hat{\lambda}}(v_{i1}^k)$, and $\sigma_2^{*\hat{\lambda}}(v_{i1}^k)$. (These probabilities are assumed to be independent across individuals and across rounds.) Letting $y = (y_1, \dots, y_i, \dots, y_I)$, the log likelihood function, given $\hat{\lambda}$, is:

$$L\big(y \mid \hat{\lambda}\big) = \sum_{i=1}^{n}\sum_{k=1}^{K} \ln pr\big(b_{i1}^k \mid v_{i1}^k, \hat{\lambda}\big) = \sum_{i=1}^{n}\sum_{k=1}^{K} \ln\big(\sigma_{b_{i1}^k}^{*\hat{\lambda}}(v_{i1}^k) \mid \hat{\lambda}\big).$$

The maximum likelihood estimate, $\hat{\lambda}$, is the value of λ that maximizes this likelihood function for the given data set. This is computed for each session and for each model.

B.2.2. ABF likelihood function

The estimation of the ABF model involves the estimation of three parameters: two cutpoints, c_{12}, c_{23}, and an error rate, ϵ. The likelihood function is constructed in the following way. For each individual, i, we have a collection of K observations of proposal 1 valuations and bonus vote choices, denoted $y_i = (v_{i1}^1, b_{i1}^1, \dots, v_{i1}^K, b_{i1}^K)$: The log likelihood of this observation, given cutpoint and error parameter estimates, $\hat{c}_{12}$, $\hat{c}_{23}$, and $\hat{\epsilon}$, is given by:

$$L\big(y_i \mid \hat{c}_{12}, \hat{c}_{23}, \hat{\epsilon}\big) = \sum_{k=1}^{K} \ln pr\big(b_{i1}^k \mid v_{i1}^k, \hat{c}_{12}, \hat{c}_{23}, \hat{\epsilon}\big),$$

where

$$\Pr\big(b_{i1}^k \mid v_{i1}^k, \hat{c}_{12}, \hat{c}_{23}, \hat{\epsilon}\big) = 1 - \hat{\epsilon} \quad \text{if}$$

$$\begin{aligned} &v_{i1}^k \in (0, \hat{c}_{12}] \quad \text{and} \quad b_{i1}^k = 0 \quad \text{or} \\ &v_{i1}^k \in [\hat{c}_{12}, \hat{c}_{23}] \quad \text{and} \quad b_{i1}^k = 1 \quad \text{or} \\ &v_{i1}^k \in [\hat{c}_{23}, 1] \quad \text{and} \quad b_{i1}^k = 2, \end{aligned}$$

and

$$\Pr\big(v_{i1}^k, b_{i1}^k \mid \hat{c}_{12}, \hat{c}_{23}, \hat{\epsilon}\big) = \frac{\hat{\epsilon}}{2} \quad \text{if}$$

$$\begin{aligned} &v_{i1}^k \in (0, c_{12}] \quad \text{and} \quad b_{i1}^k = 1 \quad \text{or} \\ &v_{i1}^k \in (0, \hat{c}_{12}] \quad \text{and} \quad b_{i1}^k = 2 \quad \text{or} \\ &v_{i1}^k \in [\hat{c}_{12}, \hat{c}_{23}] \quad \text{and} \quad b_{i1}^k = 0 \quad \text{or} \\ &v_{i1}^k \in [\hat{c}_{12}, \hat{c}_{23}] \quad \text{and} \quad b_{i1}^k = 2 \quad \text{or} \\ &v_{i1}^k \in [\hat{c}_{23}, 1] \quad \text{and} \quad b_{i1}^k = 0 \quad \text{or} \\ &v_{i1}^k \in [\hat{c}_{23}, 1] \quad \text{and} \quad b_{i1}^k = 1. \end{aligned}$$

References

Börgers, T., 2004. Costly voting. Amer. Econ. Rev. 94, 57–66.

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

Brams, S., 1975. Game Theory and Politics. Free Press, New York.

Brams, S., Davis, M., 1978. Optimal jury selection: A game-theoretic model for the exercise of peremptory challenges. Math. Operations Res. 26, 966–991.

Campbell, C., 1999. Large electorates and decisive majorities. J. Polit. Economy 107, 1199–1217.

Casella, A., 2005. Storable votes. Games Econ. Behav. 51, 391–419.

Casella, A., Gelman, A., Palfrey, T.R., 2003. An experimental study of storable votes. Working paper No. 9982. National Bureau of Economic Research.

Crémer, J., d'Aspremont, C., Gérard-Varet, L.A., 1990. Incentives and the existence of Pareto-optimal revelation mechanisms. J. Econ. Theory 51, 233–254.

Dodgson, C., 1884. The Principles of Parliamentary Representation. Harrison and Sons, London (Supplement, 1885).

Ferejohn, J., 1974. Sour notes on the theory of vote trading. Caltech.

Goeree, J., Anderson, S., Holt, C., 1998. The all-pay auction: Equilibrium with bounded rationality. J. Polit. Economy 106, 828–853.

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

Issacharoff, S., Karlan, P., Pildes, R., 2001. The Law of Democracy: Legal Structure and the Political Process, second ed. Foundation Press, New York.

Jackson, M., Sonnenschein, H., in press. Overcoming incentive constraints by linking decisions, Econometrica.

McKelvey, R.D., Palfrey, T.R., 1995. Quantal response equilibria for normal form games. Games Econ. Behav. 10, 6–38.

McKelvey, R.D., Palfrey, T.R., 1996. A statistical theory of equilibrium in games. Japanese Econ. Rev. 47, 186–209.

McKelvey, R.D., Palfrey, T.R., 1998. Quantal response equilibria for extensive form games. Exper. Econ. 1, 9–41.

McKelvey, R.D., Palfrey, T.R., Weber, R., 2000. The effects of payoff magnitude and heterogeneity on behavior in 2×2 games with unique mixed strategy equilibria. J. Econ. Behav. Organ. 42, 523–548.

Moulin, H., 1982. Voting with proportional veto power. Econometrica 50, 145–162.

Mueller, D.C., 1978. Voting by veto. J. Public Econ. 10, 57–75.

Mueller, D.C., 1989. Public Choice II. Cambridge Univ. Press, Cambridge.

Osborne, M., Rosenthal, J., Turner, M., 2000. Meetings with costly participation. Amer. Econ. Rev. 90, 927–943.

Philipson, T., Snyder, J., 1996. Equilibrium and efficiency in an organized vote market. Public Choice 89, 245–265.

Piketty, T., 1994. Information aggregation through voting and vote-trading. MIT.

Sawyer, J., MacRae, D., 1962. Game theory and cumulative voting in Illinois: 1902–1954. Amer. Polit. Sci. Rev. 56, 936–946.