Vote Trading With and Without Party Leaders.[1]

Alessandra Casella[2] Thomas Palfrey[3] Sébastien Turban[4]

January 10, 2014

[1]For helpful comments, we thank participants at the ESA 2011 meeting in Tucson and at seminars at Caltech, NYU, the Paris School of Economics, and Ecole Polytechnique, three referees, and the co-Editor. We thank Kyle Gerry for careful research assistance. We gratefully acknowledge financial support from the National Science Foundation (SES-0617820, SES-0617934, and SES-0962802), the Gordon and Betty Moore Foundataion, the Center for Experimental Social Science at NYU, the Social Science Experimental Laboratory at Caltech, and the Straus Institute at NYU Law School. Casella thanks the Straus Institute and the Paris School of Economics for their hospitality, and the Alliance Program.

[2]Columbia University, NBER, CEPR, ac186@columbia.edu

[3]Caltech, NBER, trp@hss.caltech.edu

[4]Caltech, sebastien.turban@gmail.com

Abstract

Two groups of voters of known sizes disagree over a single binary decision to be taken by simple majority. Individuals have different, privately observed intensities of preferences and before voting can buy or sell votes among themselves for money. We study, theoretically and experimentally, the implication of such trading for outcomes and welfare when trades are coordinated by the two group leaders and when they take place anonymously in a competitive market. The theory has strong predictions. In both cases, trading falls short of full efficiency, but for opposite reasons: with group leaders, the minority wins too rarely; with market trades, the minority wins too often. As a result, with group leaders, vote trading improves over no-trade; with market trades, vote trading can be welfare reducing. The theoretical predictions are strongly supported by the experimental data.

1 Introduction

Consider a number of people collectively choosing between two alternatives through majority voting. The voters are divided into two groups, depending on which alternative they prefer. Suppose that before voting all votes can be freely traded for money: individuals feeling strongly about the decision can buy votes from those who are less concerned about the outcome. To concentrate on vote trading per se, suppose also that none of the voters is budget constrained so that they all can express the intensity of their preferences through the price they are willing to pay. In this setting, where inequality and credit constraints do not play a role, is vote trading a good idea?

In this paper, we address this question in two scenarios: when trades are coordinated by two group leaders, and when they take place anonymously in a competitive market. The theory has strong predictions. In both cases, trading falls short of full efficiency, but for opposite reasons: with group leaders, the minority wins too rarely; with market trades, the minority wins too often. As a result, with group leaders, vote trading improves over no-trade; with market trades, vote trading can be welfare reducing. We find that these predictions are strongly supported by experimental results.

There are at least three major reasons to study vote trading from a normative perspective. First, as is well-known, majority voting fails to account for the intensity of preferences. Second, economic theory teaches that markets typically work well in allocating goods to those who most value them. It is natural to ask whether this insight extends to votes. Third, corporate shares are traded in markets and come not only with rights to dividends and future profits, but also to votes. To what extent does the inherent trading of votes affect share prices and trades? It is difficult to answer this question without understanding the fundamental forces operating in a market for votes[1].

It is not surprising, then, that questions about vote markets, whether mediated by money or by promises of future support (log-rolling), intrigued the early scholars in modern political economy: Buchanan and Tullock (1962), Coleman (1966, 1967), Park (1967), Wilson (1969), Tullock (1970), Haefele (1971), Kadane (1972), Riker and Brams (1973), Mueller (1973), Bernholtz (1973, 1974)[2]. Writing in 1974, however, Ferejohn summarized the sad state of knowledge on the subject succinctly: "[W]e really know very little theoretically about vote trading. We cannot be sure about when it will occur, or how often, or what sort of bargains

[1] See for example, Demichelis and Ritzberger (2007) and Dhillon and Rossetto (2011), and the references they cite.

[2] The papers had different methodological approaches (for example, cooperative versus non-cooperative games; or log-rolling versus markets for votes) and often focused on specific examples. McKelvey and Ordeshook (1980) report a laboratory experiment that studies the Riker and Brams (1973) logrolling example.

will be made. We don't know if it has any desirable normative or efficiency properties" (p. 25).

The crux of the problem is that votes have characteristics that make them very different from typical goods. Votes are indivisible and intrinsically worthless; their value depends on the influence they provide on decision-making, and therefore on the holdings of votes by all other individuals. Thus, demands are interdependent, and payoffs discontinuous at the point at which a voter becomes pivotal. These unique features pose a major theoretical obstacle to understanding vote trading. Both in a market for votes and in log-rolling games, equilibrium and other stability concepts such as the core typically fail to exist. Ferejohn's early observation was echoed in later works (Schwartz (1977, 1981), Shubik and van der Heyden (1978), Weiss (1988), Philipson and Snyder (1996)), and with very few exceptions (Piketty (1994), Kultti and Salonen (2005)), the theoretical interest in voters trading votes among themselves effectively came to an end.

The literature shifted instead to modeling vote trading not as uncoordinated trades among vote holders, but as centralized agreements mediated either by a market-maker or by party leaders[3]. Coordinated vote trading is not only easier to study but a more promising route for efficiency gains, because it can address the externalities caused by individual trades on voters who are not part of the transaction. Studying models that incorporate some strong assumptions, both Koford (1982) and Philipson and Snyder (1996) conclude that vote trading through a market-maker improves welfare.[4]

In this paper, we go back to addressing these claims–both the lack of equilibrium in uncoordinated trading, and the scope for welfare gains when trading occurs through party leaders. To do so, we build on two existing contributions, one based on general equilibrium theory, and one based on mechanism design theory.

To overcome the problem of equilibrium existence in standard competitive models of vote markets, Casella et al. (2012) developed the concept of Ex Ante Competitive Equilibrium: a market price and (stochastic) demands such that each individual is maximizing his expected utility and the market clears in expectation. If realized demands do not clear the market exactly, a rationing rule determines which demands are satisfied.

That paper shows that an equilibrium exists in a symmetric environment where each

[3] A different literature studies vote-buying by either candidates or lobbyists: for example, Myerson (1993), Groseclose and Snyder (1996), Dal Bò (2007), Dekel, Jackson and Wolinsky (2008) and (2009). We focus instead on vote-buying *within* the committee (or the electorate). The agents buying or selling votes are the voters themselves, acting either independently or through their leaders.

[4] Philipson and Snyder assume that only trades that are unanimously preferred to no-trade by all members of the two parties are allowed to take place. Koford assumes that the two party leaders cooperate in maximizing their members' surplus. We refer to Philipson and Snyder for an eloquent discussion of the practical relevance of vote trading against a numeraire.

voter is expected to favor either alternative with equal probability. Casella and Turban (2012) extend the analysis to asymmetric scenarios where the two groups are of known and different sizes, and thus can study explicitly the effect of the market on the outcomes of an ex-ante majority group and minority group. The environment seems particularly relevant for applications: often sides are not equal-sized and are well-established by party labels, cultural and geopolitical characteristics, or historical voting patterns. It is this latter approach we adopt in this paper.

We characterize an ex ante equilibrium with trade for the parametrization implemented in the experiment. For the great majority of possible realizations of intensities of preferences, only two actions are observed in equilibrium: voters either offer their vote for sale, or demand a majority of votes; and only two voters demand votes with positive probability: the highest-intensity member of the majority and the highest-intensity member of the minority. The competition for votes becomes a competition for dictatorship between these two voters. The frequency of minority victories then reflects the intensity of preferences of the most intense minority member, without taking into account the smaller size of the minority and the aggregate group values. As a result, relative to utilitarian efficiency, the minority wins too often. For the parameters used in the experiment, the bias is strong enough that ex ante welfare is lower with a vote market than in the absence of trade.

Results are quite different when trading occurs through party leaders. Koford (1982) and Philipson and Snyder (1996) studied the benchmark case of benevolent and all-powerful party leaders who are fully informed about their members' preferences and internalize their party's aggregate utility. We show that even under these ideal conditions vote trading will generally fall short of full efficiency. The reason is that centralized vote trading closely resembles a bilateral bargain between the leaders of the two opposing groups, which enables us to use two standard results from the mechanism design literature on bilateral bargaining. When a majority exists, it "owns" the decision: Myerson and Satterthwaite (1983)'s seminal theorem thus implies that there is no incentive compatible mechanism that guarantees efficient trade and voluntary participation. Trade is then too rare, and the minority wins too infrequently. However, because the minority never wins in the absence of trade, we confirm the conclusion in the literature: trading through party leaders has higher ex ante welfare than no trade. When the two groups have the same size and ties are broken randomly, the two parties "own" half of the decision: Cramton, et al. (1987)'s model of efficient dissolution of an equal partnership then implies that decision power can be transferred efficiently with voluntary participation.

We conduct a series of laboratory experiments to explore both the decentralized and the centralized approaches, and in the latter case with groups of equal and unequal size.

We design the trades as a continuous auction. With party leaders, the experimental design comes to resemble well-known auction games from bilateral bargaining theory; with market trades, it is based on the widely accepted experimental design for competitive markets for goods and assets (Smith, 1965,1982; Forsythe, Palfrey and Plott 1982; Gray and Plott, 1990, and Davis and Holt, 1992)[5]

The theory generates two main hypotheses: (1) In a decentralized market the minority wins too often and welfare is lower than in the absence of trade; (2) Centralized vote trading coordinated by party leaders leads to efficiency gains relative to majority rule in the absence of vote trading. If the two groups have different sizes, however, the minority wins too rarely and welfare falls short of full efficiency. Both hypotheses are supported in the data.

There are two minor departures from the theory in terms of its specific quantitative predictions. First, while centralized trade with equal size groups leads to efficiency gains, full efficiency (first best) is not usually achieved with equal sized groups. Second, our vote markets exhibit some overpricing, although it declines with experience.

The next section describes the basic model, in the two specifications applying to groups leaders and to competitive trading, and derives the theoretical predictions; Section 3 describes the experimental design; Section 4 discusses the experimental results, starting with voting outcomes and welfare, and then proceeding to vote allocations and prices; Section 5 concludes.

2 The Model

A committee of size n must decide between two alternatives, X and Y, and is divided into two groups with opposite preferences: it is publicly known that M individuals prefer alternative X, and m prefer alternative Y, with $m = n - M \leq M$. We will use M and m to indicate not only the size of the two groups, but also the groups' names. While the *direction* of each individual's preference is known, the *intensity* of such preference is private information. Intensity is summarized by a value v_i representing the utility that individual i attaches to obtaining his preferred alternative, relative to the competing one: individual i's utility is v_i if his preferred alternative is chosen, and 0 if it is not. It is common knowledge that each v_i is drawn independently for each individual from a distribution $F_i(v)$, atomless and with support $[0, 1]$ for all i.

Each individual has one vote, and the group decision is taken through majority voting, with ties broken with a coin toss. Prior to voting, however, individuals can purchase or sell votes among themselves for money: a trade is an actual transfer of the vote and of all rights

[5] We use a one-sided bid-only auction instead of a double auction.

to its use. Individual utility u_i is given by:

$$u_i = v_i I + t_i$$

where I equals 1 if i's preferred decision is chosen and 0 otherwise, and t_i is i's net monetary transfer, which can be positive, if i is a net seller of votes, or negative, if i is a net buyer. Each individual makes his trading and voting choices so as to maximize his expectation of u_i. In all that follows, we define as *efficient* the decision that maximizes the sum of realized utilities, or, equivalently, the decision preferred by the group of voters with higher total values.

With two alternatives and a single voting decision, voting sincerely is always a weakly dominant strategy, and we restrict our attention to sincere voting equilibria. Our focus is on the vote trading mechanism, and in particular, on two alternative trading institutions: a competitive spot market for votes, and bilateral bargaining between the two group leaders. We begin by studying the latter.

2.1 Trading through group leaders

The key problem with vote trading is the externality caused on individuals who are not part of the transaction. With this in mind, Koford (1982) and Philipson and Snyder (1996) study centralized vote trading: each group is represented by a leader who internalizes the values of all members of his group, and only the two leaders are authorized to buy or sell votes. We denote the leader's internalized value for the majority party and the minority party as v_M and v_m respectively. The value internalized by the opposite group's leader is not known, but it is known that the value is independently drawn from a distribution G. Efficiency is measured with respect to the internalized leaders' values v_M and v_m.[6]

The efficiency implications of trade through party leaders depends on the relative sizes of the two parties. If $M > m$, then the majority party "owns" the decision: in the absence of trade, it wins. The model is then isomorphic to Myerson and Satterthwaite's (1983) bargaining model, and the conclusion follows immediately: there is no mechanism that always guarantees ex post efficiency and satisfies incentive compatibility and interim individual rationality. The best mechanism has too little trade.

Alternatively, if $M = m$, then the vote is tied in the absence of trade. With the random

[6] A natural alternative is to allow for different distributions for the two parties, as would happen for example if the party value is the sum of the individual members' values, and parties' sizes differ. The theoretical results do not change substantively under this second specification. As we discuss in section 3, the assumption of identical distributions helped us to generate enough trade to keep experimental subjects engaged.

tie-break rule, each group leader expects to win with probability one-half. Thus, each leader "owns" half of the decision, and the model is isomorphic to the dissolution of an equal share partnership in a private good. We know from Cramton et al. (1987) that fully efficient trade is possible in this case.

We summarize these observations in the following remark:

Remark 1. (Myerson and Satterthwaite (1983), and Cramton et al. (1987)). *Suppose all vote trades occur through the two party leaders and the tie break rule is a coin toss. Then an efficient, incentive compatible, and interim individually rational trading mechanism exists if and only if $M = m$. If $M \neq m$, the most efficient, incentive compatible and individually rational mechanism has too little trade: the majority wins too often relative to efficiency.*

Beyond optimal mechanisms, our interest in this paper is in a specific institution, a market for votes, and in its properties. Different trading rules are plausible, but the experimental focus of the paper helps us restrict the theoretical models. The classic experiments on competitive goods markets are designed as a continuous open-book auction between buyers and sellers (for example, Smith, 1982, Plott, 1982, Plott and Smith, 1978). Remaining close to this trading institution is then both desirable per se – because it provides an immediate comparison between goods and votes markets–and has the added advantage of generating tractable auction models for the case of trading between group leaders.

When trade occurs exclusively through the groups leaders, the model has two agents only. The natural unit of trade is the minimum number of votes necessary to acquire decision power. We can normalize the object of trade to one vote without loss of generality.

2.1.1 Two equal-sized groups.

In the case of equal-sized groups, in the absence of trade, either alternative is chosen with probability 1/2. The continuous auction implemented in market experiments (and in our experimental design) is equivalent to one where the party with the highest bid wins, and pays the losing bid to the other party.[7]

Call b_j the bid submitted by party leader j, $j = m, M$ and focus on symmetric bidding strategies, such that $b_j = B(v_j)$. Thus j's payoff is given by:

$$u_j = \begin{cases} v_j - b_k & \text{if } b_j > b_k \\ b_j & \text{if } b_j < b_k \end{cases}$$

[7]In the continuous auction, the winner barely overbids the opponent, and thus effectively pays the losing bid.

and j maximizes:

$$\max_b \; [E[v_j - B(v_k)|v_k < B^{-1}(b)]G(B^{-1}(b)) + b[1 - G(B^{-1}(b))]]$$

where, in a symmetric equilibrium, $B^{-1}(b) = v_j$.

If G is uniform over $[0,1]$, as it will be in the experiment, standard derivations yield equilibrium bidding strategy:

$$B(v_j) = \frac{1 + 2v_j}{6}$$

Since B is strictly increasing, the mechanism is efficient. It is individually rational because j's expected utility is bounded below by $\frac{v_j}{2}$.

2.1.2 Two unequal-sized groups.

With $M > m$, the majority party wins if there is no trade. Thus the majority party leader assumes the role of the seller in a bargaining problem, and the minority leader the role of the buyer. In the continuous auction of the experiment, trade occurs only if party m bids high enough to exceed party M's value. The payoffs are:

$$u_M = \begin{cases} v_M & \text{if } b_M > b_m \\ b_m & \text{if } b_m > b_M \end{cases} \qquad u_m = \begin{cases} v_m - b_m & \text{if } b_M < b_m \\ 0 & \text{if } b_M > b_m \end{cases}$$

Again, standard derivations imply that the majority party leader's dominant strategy is to bid $B_M(v_M) = v_M$, and if G is uniform, the minority party leader will bid according to:

$$B_m(v_m) = \frac{v_m}{2}$$

The trading mechanism is individually rational but ex post inefficient: trade occurs only if the buyer's value is at least *twice* the seller's value.

To summarize, in our application to vote trading between group leaders, the theoretical prediction is unambiguous: vote trading is not fully efficient because the majority wins too often. However, it dominates majority voting without vote trading because it allows for *some* minority victories, when the disparity in values is sufficiently high.[8]

[8]The mechanism has lower trade and lower expected efficiency than the optimal Myerson and Satterthwaite mechanism, but the difference is not large. With G uniform over $[0,1]$, the optimal mechanism can sustain trade whenever $v_B \geq v_S + 1/4$; the probability that a trade is concluded is then 28 percent (versus 25 in our auction), and expected welfare 96 percent of full efficiency (versus 94).

2.2 Decentralized competitive vote trading

Characterizing the equilibrium and welfare properties of a competitive market for votes is more challenging. It is well-known that with standard models of competitive equilibrium there exists no price and allocation of votes such that the market clears[9]. To solve this problem, we borrow the concept of Ex Ante Competitive Equilibrium developed in Casella et al. (2012) and its extension to asymmetric scenarios with groups of different sizes in Casella and Turban (2012). The main findings of these papers are summarized in the following Remark:

Remark 2. *Consider a competitive market for votes. Under a weak condition on the ordering of realized values, there exists an Ex Ante Equilibrium such that trade occurs for any m and M. The highest value voter in each group demands $(n-1)/2$ votes with positive probability and sells otherwise; all other voters sell. A competitive market for votes resembles an auction for dictatorship.*

A striking feature of the equilibrium is that, given $v_{\overline{M}}$ and $v_{\overline{m}}$ (the realized highest values in the two groups), neither the price nor the traded quantities depend on the relative size of m. Because all other votes are offered for sale, the advantage of the majority evaporates. As a result, the theory predicts that the minority wins too frequently, relative to what efficiency dictates.

In the Appendix, we characterize equilibrium strategies and prices for the specific parametrization we use in the experiment. We find that the expected frequency of minority victories is so high that efficiency is lower than under no trade: competitive vote trading leads to a decline in expected welfare.

2.3 Summary of the predictions

We can summarize our predictions under the parametrization chosen in the experiment as follows:

1. Vote trading intermediated by leaders of different-sized groups falls short of efficiency, but increases welfare relative to the no-trade outcomes. The fraction of minority victories falls short of the efficient fraction but, by definition, is higher than under majority voting.

2. Vote trading intermediated by leaders of identically-sized groups is efficient, and thus welfare is higher than under majority voting with no trade.

[9]See, for example, Ferejohn (1974), Philipson and Snyder (1996), Piketty (1994)

3. Competitive vote-trading falls short of efficiency, and allowing trade decreases welfare relative to the no-trade outcomes. The fraction of minority victories is higher than the efficient fraction, and the market for votes is inferior to simple majority voting with no trading.

3 Experimental design

The experiment was run at the Center for Experimental Social Science at NYU (CESS), and at the Social Science Experimental Laboratory at the California Institute of Technology (SSEL), with enrolled students recruited from the whole campus through the laboratories' web sites. No subject participated in more than one session. After entering the computer laboratory, the students were seated randomly in booths separated by partitions and assigned ID numbers corresponding to their computer terminal; the experimenter then read aloud the instructions, projected views of the computer screens during the experiment, and answered all questions publicly. Each session of the experiment amounted to 25 paid rounds, preceded by one unpaid practice round. Each experimental session consisted of a single treatment.

In the experiment, the two groups were called X and Y, from the name of the preferred alternative. The sizes of the two groups were commonly known. Each round followed the same procedure. At the start of the round, subjects were matched randomly in committees and assigned either to group X or to group Y. Each subject i was told by the computer whether he belonged to group X or group Y, and the value v_i he would win if his preferred policy prevailed. Values were expressed in experimental points, and subjects knew that values were drawn randomly by the computer, independently and privately for each subject, and could assume any integer value between 1 and 100, with equal probability. After values were assigned, the market for votes opened. Any subject could post a bid specifying the price he was willing to pay for a vote; the bid appeared on all monitors, together with the name of the group the bidder belonged to, and a running tally of the votes belonging to each group.[10] If anyone accepted the bid, the transaction was concluded; if not, anyone could post a new bid, higher than the previous one. After each trade, a new bid could be posted, at any value. The market for votes was open for three minutes, during which as many transactions were concluded as there were accepted bids.[11] Once the market for votes closed, voting

[10]Each subject started the session with an initial endowment of 200 points, to be paid back at the end of the experiment. Posted prices had to be between 1 and 100, and no subject was allowed to post a bid higher than his current endowment of points. If a subject reached a 0 balance he would be excluded from bidding until the balance turned positive, but in the experiment all balances remained positive.

[11]The market was open for two minutes in treatment $32C$ described below. In all treatments, the market closed early if there was no activity for 30 seconds.

took place. All votes were automatically cast for the preferred option of the post-market owner of the vote, with ties broken randomly.[12] The session then proceeded to the following round, where subjects were randomly regrouped into new committees. At the end of each session, subjects were paid their cumulative earnings from all rounds, summing payoffs from obtaining their preferred committee decisions and net transfers from the market for votes, multiplied by a pre-announced exchange rate, plus a fixed show-up fee. Each session lasted about 90 minutes, and average earnings were around $33. A sample of the instructions from one of the sessions is reproduced in the on-line Appendix.[13]

Our treatment variables are the relative size of the two groups, m and M, and whether trade takes place through the group leaders or through the market. The first treatment, called $1, 1$, captures vote trading through group leaders when the groups have equal size: each group is represented by a single subject, with opposite preferences, and each subject enters the vote market with a single vote. The second treatment, called $3, 2C$, captures vote trading through group leaders when the groups have different sizes: each group is again represented by a single subject, with opposite preferences, but the two subjects enter the market endowed with three and two votes, respectively. In line with the assumption in section 2.1, we implemented this treatment by generating each subject's value as a single random draw from 1 to 100, assuming each integer value with equal probability. Alternatively, we could have generated each leader's value as the sum of multiple independent draws (three and two, respectively, for the majority and minority leader). We chose our design because it generates more frequent opportunities for trade: when a leader's value is a single random draw from a uniform distribution, the theory predicts a frequency of trade of 25 percent; when it is generated as sum of multiple independent draws, the predicted frequency of trade in equilibrium falls to 10 percent.[14] Maintaining the interests of the subjects is crucial, and here they play no other role but trading. Note that the design allows us a direct comparison between $1, 1$ and $3, 2C$ treatments. The third treatment is the market treatment: the two groups, with opposite preferences, are formed by three and two subjects respectively, and each individual subject is free to trade, independently of the other members of his group. Each individual value is an independent random draw, assuming any integer value between

[12] In one of the $1, 1$ sessions run at Caltech (session s5 in Table 1 below), in case of tie each subject received 50 percent of his value with probability 1. We changed the design to check whether the uncertainty of the outcome in case of a tie affected the results, but the session is indistinguishable from the others.

[13] We used the Multistage Game software package developed jointly between the SSEL and CASSEL labs. This open-source software can be downloaded from http://software.ssel.caltech.edu/

[14] Going back to section 2.1, if v_M equals the sum of three independent draws from a uniform distribution, the minority leader's problem yields $b_m \leq 3v_m/4$. Since $b_M = v_M$, the frequency of trade is bounded above by the frequency with which $3v_m/4 \geq v_M$. When v_m equals the sum of two independent draws from a uniform, the latter frequency is 10.4 percent.

Session	n	Treatment	Subject pool	# Subjects
s1	2	$1,1$	NYU	12
s2	2	$1,1$	NYU	8
s3	2	$1,1$	NYU	10
s4	2	$1,1$	NYU	16
s5	2	$1,1$	CIT	12
s6	2	$1,1$	CIT	10
s7	2	$3,2C$	NYU	12
s8	2	$3,2C$	NYU	12
s9	2	$3,2C$	CIT	10
s10	2	$3,2C$	CIT	10
s11	5	$3,2$	NYU	15
s12	5	$3,2$	NYU	20
s13	5	$3,2$	NYU	10
s14	5	$3,2$	NYU	10
s15	5	$3,2$	CIT	15
s16	5	$3,2$	CIT	15

Table 1: Experimental Design

1 and 100 with equal probability, and each group value then is the sum of either two or three independent draws. We call this treatment $3,2$. Table 1 reports the experimental design.

Table 2 summarizes the theoretical predictions discussed in the previous section. Columns 2 and 3 report the expected frequency of minority victories in equilibrium and under full efficiency; columns 4 and 5 report, respectively, expected ex ante utility in equilibrium and in the absence of trade, expressed as share of expected ex ante utility with full efficiency.

	Min victs %	Eff. min victs %	E(eff share) %	E(eff share, maj rule) %
1,1			100	75
3,2C	25	50	94	75
3,2	52.5	22.5	84.2	95

Table 2: Theoretical predictions. Uniform distribution

4 Experimental Results.

We begin by evaluating whether the strong theoretical predictions on efficiency and frequency of minority victories are confirmed by the experimental data. We will later describe in detail the transaction and price data. Because the latter show evidence of learning, we focus our discussion on the second half of the experimental rounds, i.e. rounds 11 to 25. All qualitative results are unchanged if we consider all rounds.[15]

[15]Casella, Palfrey and Turban (2012) compare price and allocation data from early and late rounds.

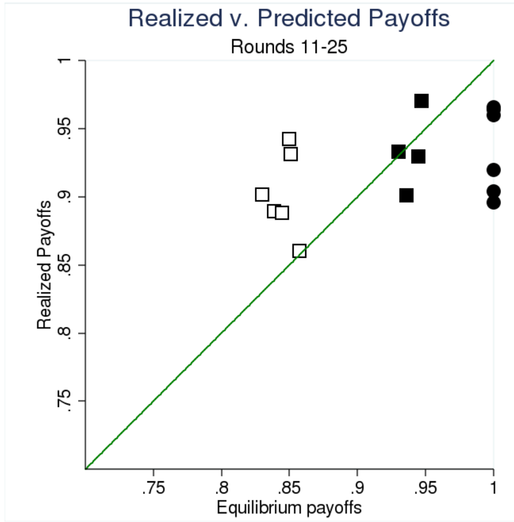

Figure 1: Experimental payoff per session versus equilibrium (as share of efficiency). Rounds 11-25. Solid dots correspond to $1,1$ sessions; solid squares to $3,2C$ sessions, and empty squares to $3,2$.

4.1 Welfare and minority victories

4.1.1 Comparison to equilibrium and to efficiency

Figure 1 shows the aggregate payoff in each experimental session, on the vertical axis, versus the corresponding payoff if all subjects had played the equilibrium strategies, on the horizontal axis. Both payoffs are expressed as share of the efficient payoff (the maximal aggregate payoff) and all predictions are calculated on the basis of the realized experimental valuation draws. The symbols distinguish the different treatments: circles are $1,1$ sessions; squares are sessions with groups of unequal size, with solid squares corresponding to $3,2C$ sessions and empty squares to $3,2$ sessions. The diagonal line is the 45 degree line; thus points above the line indicate experimental payoffs in excess of the theoretical prediction, and points below payoffs that fall short of equilibrium payoffs.

Compared to equilibrium predictions, experimental payoffs are higher in the $3,2$ market sessions, they are comparable in the $3,2C$ sessions and lower in the $1,1$ sessions. The more noticeable feature of the figure is the disparity between the clear efficiency rankings of the theory and the more uniform level of the experimental payoffs: while the different treatments are distinctly organized along the horizontal axis, experimental payoffs as share of efficiency are similar.

Consider first the experimental payoffs in $1,1$ sessions. In itself, the observation that they fall short of the theoretical prediction is not very surprising. Since theory predicts full efficiency, any noise in behavior must result in lower payoffs than expected. More surprising

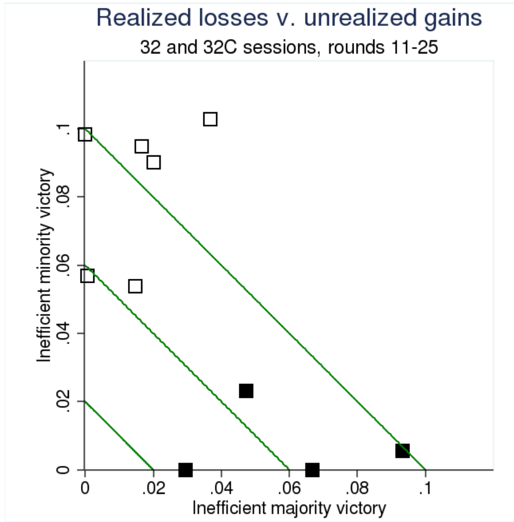

Figure 2: Realized losses versus unrealized gains, both as share of full efficiency. Rounds 11-25. The diagonal lines are iso-efficiency loss lines. Solid squares are $3,2C$ sessions; empty squares $3,2$.

however is the lack of differentiation in payoffs relative to $3,2C$ sessions. In part, this can again be rationalized as a mechanical result of noisy choices: because $3,2C$ payoffs are not expected to be fully efficient, noise in $3,2C$ sessions can move experimental payoffs both above and below the equilibrium prediction. If experimental subjects make unbiased errors, the result must be closer experimental payoffs in $1,1$ and $3,2C$ treatments than theory predicts. That said, the close similarity of the observed payoffs violates the unambiguous theoretical result on the higher efficiency of bargaining when the two groups have equal size.

The similarity in payoffs between $3,2$ and $3,2C$ sessions, on the other hand, hides systematically different efficiency deviations that can be well explained by the theory.

Figure 2 disaggregates the source of efficiency losses for these two treatments. The diagonal lines in the figure are iso-efficiency loss lines. The origin, at (0,0), denotes full efficiency, or the maximum possible aggregate payoff; moving up and to the right, the efficiency losses increase, with the three iso-loss lines corresponding to losses of 2, 6 and 10 percent respectively, relative to the efficient payoffs. The horizontal axis measures losses from *missed trading* opportunities: instances where the aggregate experimental values of the minority group were higher than the majority's and yet the majority won. The vertical axis measures losses from *inefficient trading*: instances where the aggregate experimental values of the minority group were lower than the majority's and yet the minority won. The symbols are as in Figure 1: solid squares are $3,2C$ sessions, and empty squares are $3,2$ sessions.

The predictions are well borne out by the data. We did not add a 45 degree line to keep the figure clean, but all empty square lie above it, and all solid squares lie below it: in $3,2C$

sessions, the source of losses is almost exclusively missed trading opportunities, while in $3, 2$ sessions losses come primarily from too much trade and inefficient minority victories. In particular, in two of the $3, 2C$ sessions there is not a single instance of inefficient minority victories, and in two of the $3, 2$ sessions not a single instance of inefficient majority victories.

Figure 3 tests directly how the realized fraction of minority victories in $3, 2C$ and $3, 2$ data compares to equilibrium predictions and to efficiency outcomes. The first panel plots the realized frequency of minority victories in each experimental session, on the vertical axis, and the frequency predicted by the theoretical models, given the realized experimental draws, on the horizontal axis. In $3, 2C$ sessions, the minority won slightly more than the theory predicts while in $3, 2$ sessions the minority won less frequently than theory predicts. Nevertheless, in *all* $3, 2$ sessions, realized minority victories were more frequent than in *any* $3, 2C$ session.

The second panel does a similar comparison to the efficient fraction of minority victories, again calculated for each session according to the realized experimental draws. The theory is strongly supported. Every point representing a $3, 2$ session is above the 45 degree line, and every point representing a $3, 2C$ session is below: as expected, the minority wins too much in $3, 2$ sessions, and too little in $3, 2C$ sessions.[16]

The conclusion is very robust and remains true at the level of the individual groups, even with the inevitable added noise. In Figure 4, we replicate the two panels of Figure 3 at the group level: each dot now corresponds to the fraction of minority victories for a *group label,* in the two treatments.[17] In the first panel, the symbols tend to align themselves along the 45 degree line, implying that realized minority victories are around the predicted values, although again we observe that they tend to be more frequent than predicted in the $3, 2C$ sessions and less frequent in the $3, 2$ sessions. Comparing across treatments, realized minority victories remain more frequent in $3, 2$ sessions than in $3, 2C$ sessions. In the second panel, again we find that realized minority victories are too frequent relative to efficiency in $3, 2$ sessions, but less frequent in $3, 2C$ sessions.

To test formally whether the frequencies of minority victories observed in the experiment are significantly different across the two treatments, we need to account for the different value draws. We can use the efficient frequency of minority victories as a means of normalizing the value draws, and test whether, across groups, the ratios of realized to efficient minority victories observed in the two treatments could be drawn from the same sample.

[16] In the second panel, the solid square located most to the right hides a second solid square with almost identical values.

[17] Recall that groups are formed randomly at each round; thus a given group label does not correspond to a fixed set of subjects. Aggregating at the group level is an averaging device, smoothing some individual noise while allowing for more variation than at the aggregate session level.

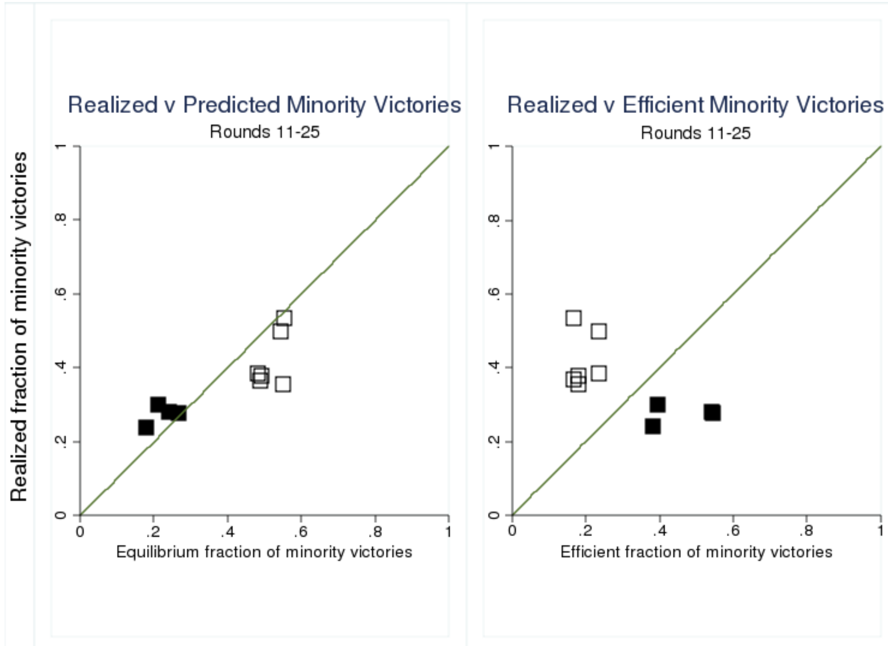

Figure 3: Realized frequency of minority victories, versus equilibrium frequency (panel 1) and efficient frequency (panel 2). Rounds 11-25. Solid squares are $3, 2C$ sessions; empty squares $3, 2$.

The hypothesis is strongly rejected by a two-sided Kolmogorov-Smirnov test: the D statistic is 0.96, with p-value 0.000.

4.1.2 Comparison to majority voting

The theory has sharp predictions on the efficiency of vote trading relative to majority voting with no trade in the different treatments. Recall that centralized trading sessions are predicted to perform better, while decentralized market sessions should do worse.

Figure 5 plots the aggregate experimental payoffs per session, on the vertical axis, versus the aggregate session payoffs in the absence of trade, on the horizontal axis. Both measures are normalized as a share of the efficient payoffs.

The results strongly support the theory. As expected, payoffs were systematically higher than in the absence of trade in the two treatments with leaders' trading, and consistently lower in the treatment with market trading. The theoretical prediction is confirmed in every single experimental session.

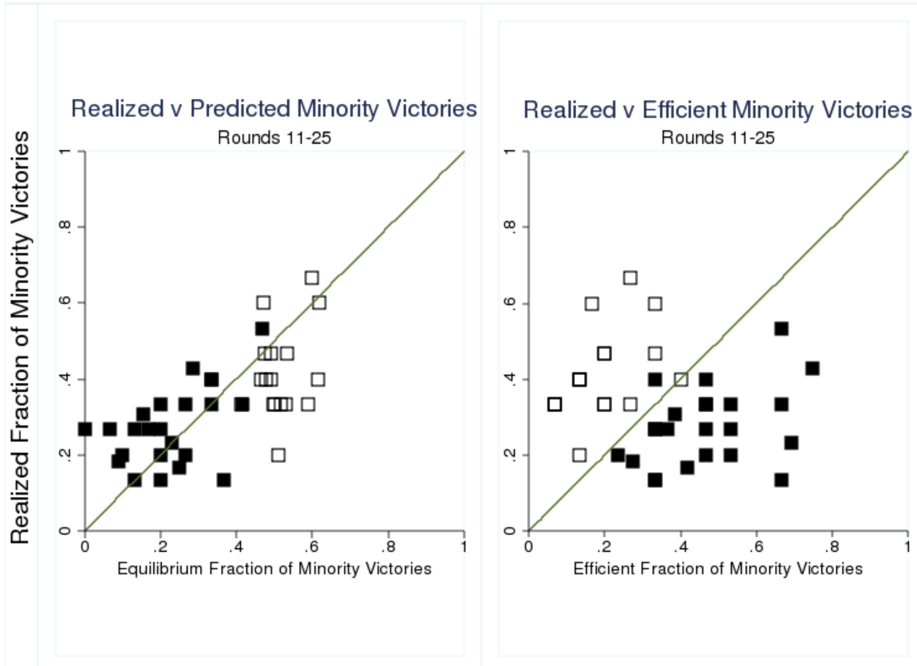

Figure 4: Realized frequency of minority victories by group versus equilibrium (panel 1) and efficient frequency (panel 2). Rounds 11-25. Solid squares are $3, 2C$ sessions; empty squares are $3, 2$.

4.2 Trading outcomes

The theoretical predictions on the efficiency of the different treatments result from predictions of specific trade patterns. In this subsection, we evaluate whether such patterns are indeed observed in the data.

In $1, 1$ sessions, the theory states that there should always be trade and it should always be efficient: the higher value voter should buy out the lower value voter. Figure 6 reports realized and unrealized gains from trade. The first panel displays all instances in which trading happened. The vertical axis is the buyer's value, and the horizontal axis the seller's value, so that points above the diagonal are efficient trades, and points below the diagonal are inefficient trades. The second panel displays all instances in which trade did not happen. The vertical axis is the higher value and the horizontal axis the lower value for the two experimental subjects in those instances, so that all points lie above the diagonal by construction.

The frequency of trade–the fraction of all points that lie in the first panel–is 79 percent. Conditional on trade, the frequency of efficient transactions–the fraction of points in the first panel that lie above the diagonal–is 84 percent. On the whole, then, transactions in $1, 1$ sessions were less efficient than the theory predicts, both because of lower trade and of

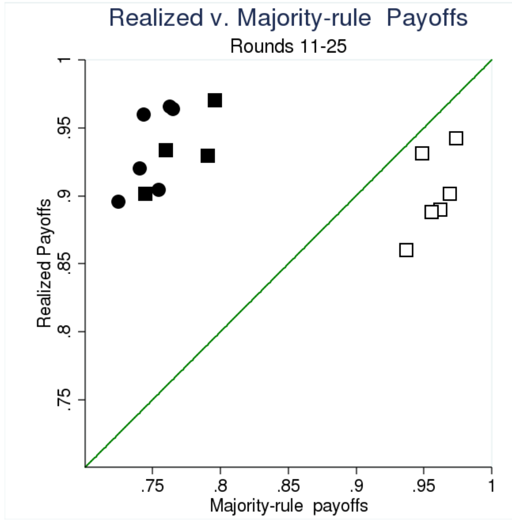

Figure 5: Realized payoff per session versus no-trade payoff (as share of efficiency). Rounds 11-25. Solid dots are $1,1$ sessions; solid squares $3,2C$ sessions, and empty squares $3,2$.

mistakes in the direction of trade.

In $3,2C$ sessions, the theory states that trade, when it occurs, should be efficient, but some efficient opportunities will be missed. Figure 7 reports the same results as Figure 6 for these sessions, although the axes differ. The first panel displays all instances where trading happened and the minority won. The vertical axis is the minority value and the horizontal axis the majority value. The second panel plots the instances in which a minority victory would have been efficient but did not take place[18]. Again by construction, all points above the diagonal represent efficient trades. The steeper line is the theoretical boundary for trade: our auction model predicts that trade will occur for points above the line, but not below.

As predicted, when trade happens, it is usually efficient: the fraction of points below the diagonal in the first panel is very small[19]. The prediction that some efficient trading opportunities will be missed is also validated. Of all these opportunities, i.e. of all points above the diagonal, the fraction in the first panel is 57 percent. The number is above 50 percent, the predicted share, in line with the over-trading observed in bargaining experiments with similar design (Radner and Schotter (1989), Valley et al. (2002)), but not significantly[20].

As expected, experimental subjects concluded fewer trades in the $3,2C$ treatment, relative to the $1,1$ treatment.[21] The theory, however, tells us more: the disparity between the two

[18]The figures do not report values for which the absence of trade was efficient and was observed.

[19]4 of 91 in rounds 11-25 (4.4. percent).

[20]The Pearson χ^2 value is 1.45, which is insignificant at the 5 percent level.

[21]In $1,1$ sessions, the fraction of realized efficient trades is 66 percent (0.79 x 0.84), v/s the 57 percent observed in $3,2C$ sessions. According to a χ^2 test, the difference is statistically significant at the 5 percent level.

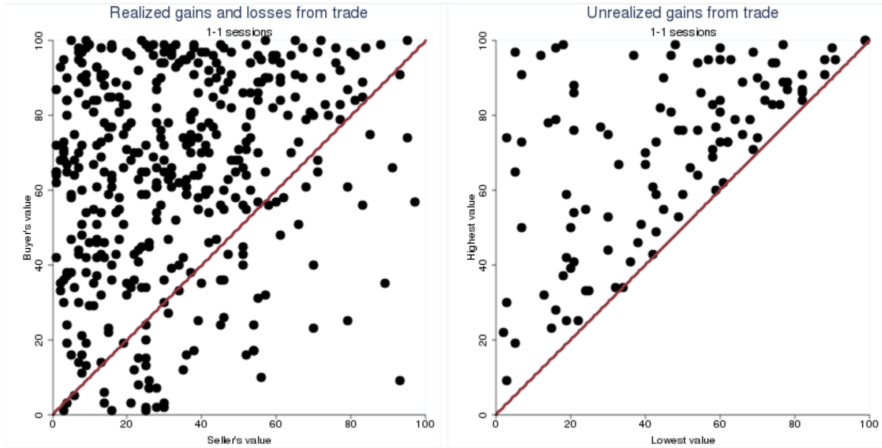

Figure 6: Trade in $1, 1$ sessions. Rounds 11-25. On the left panel, each dot is an individual transaction. Equilibrium predicts trade above the diagonal. On the right panel, each dot is a missed transaction. Equilibrium predicts no dots.

treatments should be concentrated on realized pairs of valuations for which the buyer's (the minority's) value is higher than the seller's (the majority's) but less than twice as high. The figures show quite clearly that in both treatments a large fraction of missed trading opportunities corresponded to realized draws close to the diagonal, as intuition suggests. But is the concentration significantly higher in $3, 2C$ sessions?

In Table 3 we regress the fraction of efficient trades realized in the two treatments in each round on a constant, the round number[22], and three indicator variables. The first indicator captures whether data were generated in $1, 1$ or $3, 2C$ sessions; the second whether the valuation draws were in the critical area $v_b \in [v_s, 2v_s]$ (where we indicate by b the buyer and by s the seller); the third indicator is the interaction term selecting instances where the valuation draws were in the critical area in $1, 1$ sessions. We report both logit and probit estimations.

As expected, when the valuations are in the critical area, the frequency of realized trades is significantly smaller. Both estimations show that the effect is larger in $3, 2C$ sessions, in line with theoretical predictions, although the interaction term is only significant at the 10 percent level. Note that all the difference between $1, 1$ and $3, 2C$ sessions is concentrated in this area: over the remaining range of value realizations, the coefficient of the indicator variable for $1, 1$ sessions is not significantly different from zero. The constant, capturing the frequency of trade common to both treatments when the buyer's value is more than twice the seller's, is predictably positive and large.

In the $3, 2$ treatment, the market design induces multiple trades within each group. The

[22]Round 11, the first round of the second half rounds, is transformed as Round 1 and so on.

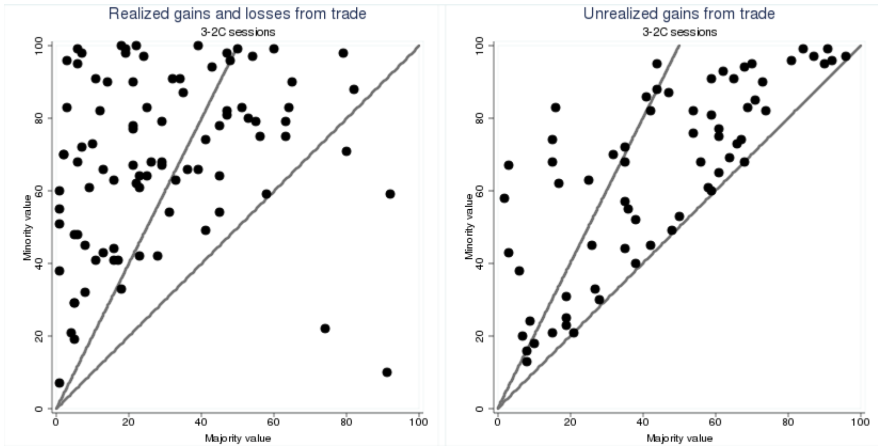

Figure 7: Trade in $3, 2C$ sessions.Rounds 11-25. On the left panel, each dot is a realized trade (minority victory). Equilibrium predicts trade above the steep line. On the right panel, each dot is a missed opportunity for trade (minority victory). Equilibrium predicts all dots between the diagonal and the steeper line.

validity of the theory should thus be tested on the basis of the final allocation of votes, when the market closes. Individual strategies depend on the realized vector of valuations, and, in line with the detailed equilibrium characterization in the Appendix, Figure 8 distinguishes five possible scenarios, depending on the group membership of the voters with highest values overall. In the figure, values are in increasing order from left to right. The black bars are the realized allocations of votes when the market closes, averaged over all groups and rounds with the relevant order of value realizations; the grey bars are the predicted expected allocations, given the experimental value draws.

The theory predicts that in all scenarios with trade, the largest purchases of votes should come from the highest value voter in each group, regardless of the exact position that those two voters occupy in the overall ranking of values. It is only when all three majority voters have values higher than any minority voter (as in the last panel) that the equilibrium with trade characterized in the Appendix does not exist, and we conjecture that the no-trade equilibrium will be focal. The grey bars reflect these predictions.

The hypothesis that the realized allocation of votes displayed by the black bars is generated by the theoretical distribution represented by the grey bars is rejected at the 5 percent level in each panel[23]. Yet, the more aggressive strategies of the highest-value minority and majority voters appear in the data: in all four of the top panels, the highest black bars correspond to those two voters, even when theirs are not the two highest values overall, as in the panels in the second row. The result is worth remarking because neither others'

[23]The Kolmogorov-Smirnov test needs to be adjusted here for the discreteness of the distributions. We ran a Cramér-von Mises test (Choulakian, Lockhart and Stephens, 1994).

VARIABLES	(1) probit	(2) logit
$1,1$-sessions	0.00247	0.00311
	(0.181)	(0.313)
$v_b \in [v_s, 2v_s]$	-1.122***	-1.827***
	(0.216)	(0.363)
$v_b \in [v_s, 2v_s]$ interacted with $1,1$-sessions	0.476*	0.764*
	(0.246)	(0.412)
Experience (Round)	0.00909	0.0154
	(0.0118)	(0.0197)
Constant	0.688***	1.124***
	(0.185)	(0.315)
Observations	672	672

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 3: Probability of realizing an efficient trade

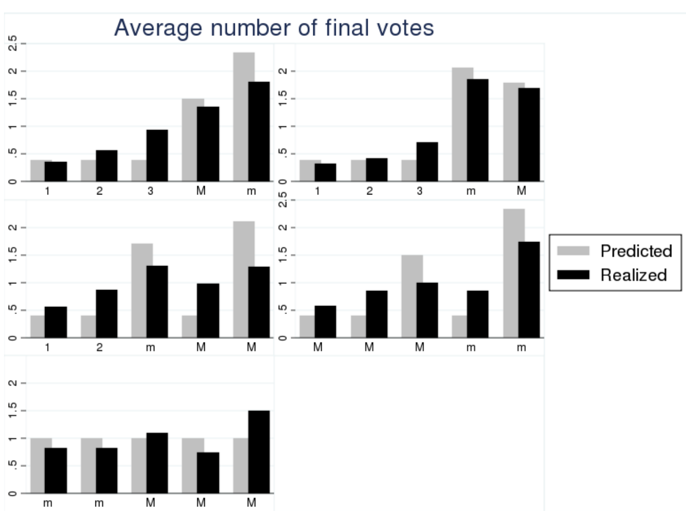

Figure 8: Trade in $3,2$ sessions. Rounds 11-25. Average votes allocations by profile of value ranks. Experimental results (black) and equilibrium predictions (grey).

realized values nor one's own position in the overall ranking were known to the subjects, and the experimental design was such that values changed every round. And yet, as theory predicts, the voters' relative rankings were revealed through trading. It is noticeable too that the concentration of votes in the hands of highest-value voter in each group is always observed in the data, with the exception of the last panel, when the no-trade equilibrium seems particularly plausible. Although the data show that some trade did take place, the dispersion in votes holdings across members of the group is much less pronounced than in all other scenarios.

4.3 Prices

For each of the three treatments, the theory also has precise predictions about equilibrium prices. We thus evaluate our theory further by analyzing whether it is a good predictor of the prices realized in individual trades. Note that in our experiment, price discovery is particularly difficult because new values are randomly realized at each round, and thus the equilibrium price change at each round.

Figure 9 plots for each treatment the average percentage difference between realized and predicted equilibrium prices in each round. For each group in each session, we calculate the predicted equilibrium price given the realized values in each round; we then compare it to the price at which trade occurred in the experimental data. For each round, the resulting percentage difference is averaged over all groups and all sessions. If there are more than one trade, we calculated both the average and the last traded price. In $1, 1$ and $3, 2C$ treatments, the theory predicts a single trade; multiple trades are occasionally observed in the data but the results are indistinguishable whether we use average or last traded price. In the market treatment $3, 2$, multiple trades are expected and observed, and the results are sensitive to the price measure. The average price is less noisy and it is the measure we report in the figure.[24]

The figure shows two main regularities. First, in $1, 1$ and $3, 2C$ treatments, there is tendency towards overpricing: almost all bars are above zero, indicating that experimental prices are above equilibrium prices. Second, $1, 1$ and $3, 2$ sessions show evidence of convergence to equilibrium; no trend appears for $3, 2C$.[25]

A regression of the percentage difference between the realized price and the equilibrium price on the round number and a constant confirms what the figure shows and is reported in

[24]The data on last traded prices are presented in Casella, Palfrey and Turban (2012).

[25]By looking at rounds 11-25, we ask whether realized prices approached the equilibrium price at the end of an experimental session. The full range of dynamic convergence appears more clearly when we consider all rounds–see the discussion in Casella, Palfrey, and Turban (2012).

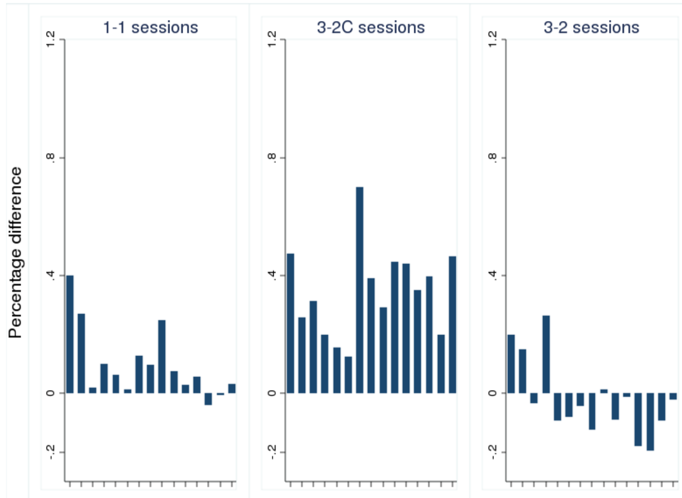

Figure 9: Average percentage difference between realized and equilibrium prices in each round. Rounds 11-25.

Table 4. The percentage difference falls significantly with the round number in treatments $1,1$ and $3,2$, but not in $3,2C$. The constant term is positive and statistically significant in $1,1$ and $3,2C$ treatments, indicating convergence from above in $1,1$ and systematic overpricing for $3,2C$. The evolution across rounds explains a fair fraction of the observed price variability in $1,1$ and $3,2$ sessions, but none for $3,2C$.[26]

The systematic overpricing in $3,2C$ sessions replicates the findings of Radner and Schotter (1989). Both the $3,2C$ treatment in this paper and the experiment in Radner and Schotter study two-person trading mechanisms, but the experimental designs differ: Radner and Schotter analyze a sealed-bid double auction, while our experiment has open, on-going bids. The result appears robust to changing the frame of the auction.

As we remarked for the allocation of votes, we find the convergence towards the equilibrium price in the $3,2$ data particularly striking, because the equilibrium depends on the realization of the full vector of values. We attribute the result not to the conscious calculation of the equilibrium price, but to the underlying forces of competitive market exchange–gains and losses experienced at the traded prices. In exchanging votes as in exchanging goods, the

[26] An unexpected finding in $3,2C$ sessions is the presence of redundant trades: trades that do not change ownership of the decision power. Typically, they take the form of a low bid by the subject representing the majority party, low bid that seems intended to stimulate a counteroffer by the opponent but is instead accepted. Redundant trades are pure transfers with no allocative or efficiency effect, and we do not include them in the $3,2C$ data analysis reported in the text. All results are qualitatively unchanged if redundant trades are included, with the only exception that, since they tend to occur at low prices, the evidence for overpricing in $3,2C$ is reduced. The constant terms in the $3,2C$ regressions in Table 4 and, below, in Table 5 become quantitatively small and statistically insignificant.

	1-1 sessions	3-2C sessions	3-2 sessions
Rounds	-0.0163**	0.00603	-0.0189**
	(0.00664)	(0.00761)	(0.00654)
Constant	0.229***	0.298***	0.128*
	(0.0727)	(0.0813)	(0.0648)
Observations	15	15	15
R-squared	0.371	0.032	0.409

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 4: Regression of the percentage difference between realized and predicted prices on the round number. Average realized prices; Rounds 11-25.

experimental data appear to support the fundamental intuition at the heart of our competitive market theory.[27]

The evidence above refers to average realized prices. But are the disaggregated prices also consistent with the theory? Figure 10 presents scatter plots of the average transacted price, for each group and round, on the vertical axis, plotted against the equilibrium price for that group and round, on the horizontal axis. Each panel in the figure corresponds to a different treatment. The panels also show the best fit line.

The dispersion in realized prices is evident in the figure, but so is the positive correlation between realized prices and equilibrium prices. The figure makes visible one factor that may contribute to the observed overpricing: the upper bounds on equilibrium prices are a fraction of possible realized valuations. While valuations vary between 1 and 100, maximal equilibrium prices are 50, for $1, 1$ and $3, 2C$, and 33 for $3, 2$. If there is a diffuse random error in realized prices, with support over the full range of valuations, the result is systematic overpricing.[28]

Table 5 tests whether the regression lines are significantly different from the 45 degree line. The standard errors are clustered at the session level.

As the last line in the table shows, the slope is never significantly different from 1 at any conventional significance level, in all three treatments. In treatment $1, 1$ we cannot reject a zero constant; we can, at the 5 percent level, in $3, 2C$ and $3, 2$ treatments. Finally, the large

[27]Casella, et al. (2012) report similar convergence towards equilibrium price in the symmetrical scenario where each voter has equal probability of favoring either alternative.

[28]In $1, 1$ sessions, the equilibrium price is $p^* = (100 + 2v_s)/6$, and thus has a minimum at $100/6$. The figure also shows that, in later rounds especially, realized prices appear to lie above a linear function of the equilibrium price, with slope higher than 1 and negative intercept. The reason is that sellers only sell if $p \geq v_s/2$, since the default is a tie. With $p^* = (100 + 2v_s)/6$, $p \geq 3/2p^* - 25$.

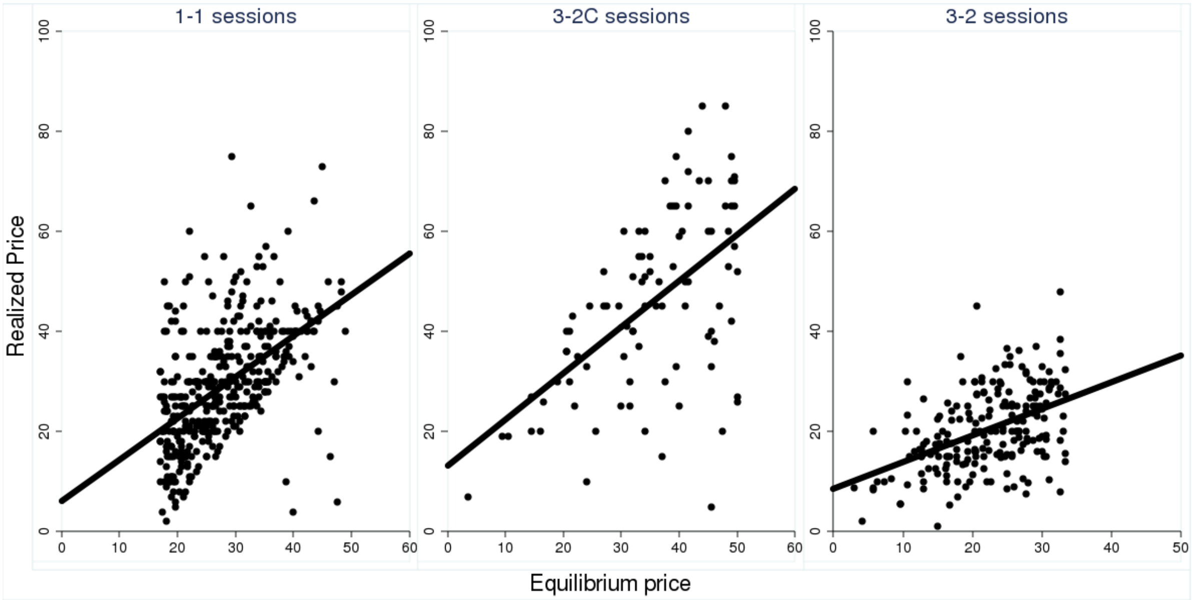

Figure 10: Traded prices versus equilibrium prices, per group and round, and linear regression lines. Rounds 11-25.

unexplained noise in realized prices is reflected in the low R^2's.

Summarizing then, we find that the disaggregated prices are also correlated with the theoretical predictions, and we cannot reject a linear, one-to-one relation. However, the data also show some evidence of systematic overpricing in the realizations relative to predictions.[29]

5 Conclusions

The objective of this paper is a better understanding of vote trading in committees and legislatures that operate under simple majority rule.

On the theoretical side, we show that standard economic models of bargaining and exchange can be reinterpreted to provide tractable equilibrium models of vote trading. If vote trading is centralized, in the sense that there is a single representative of the interests of each side of an issue–for example, a party leader–then results from the mechanism design approach to bargaining theory translate directly to voting environments. If the two parties have equal size–if there is no minority–then vote trading can theoretically lead to a first best outcome. This follows from the main result in Cramton et al. (1987). In the presence

[29] All results on prices and on allocations, and for all treatments, are unchanged if the data are disaggregated by subject pool, with a single exception. In the convergence regressions for realized prices in $3, 2C$ sessions, reported in Table 4, NYU subjects show no evidence of systematic overpricing for the average price per round: the constant term in the regression becomes insignificant. However, the lack of overpricing is not confirmed by the disaggregated price regressions in Table 5, suggesting the possible role of outliers in lowering average prices per round.

	1-1 sessions	3-2C sessions	3-2 sessions
Eq Price	0.825***	0.918**	0.865***
	(0.156)	(0.162)	(0.144)
Constant	6.182	13.38**	8.622**
	(4.681)	(3.730)	(3.230)
Obs.	409	93	208
R^2	0.265	0.305	0.187
p-val (coef=1)	0.315	0.648	0.390

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 5: Price regressions; rounds 11-25.

of a minority party, however, vote trading cannot lead to a first best outcome, although it improves over majority rule without trade. There is too little trading: the majority wins too often. This follows from the main result in Myerson and Satterthwaite (1983).

If vote trading is decentralized, in the sense that all trading takes place between individual party members rather than being coordinated by party leaders, then it can be studied by adapting the standard general equilibrium model of competitive markets. For this purpose, we apply the concept of Ex Ante Competitive Equilibrium developed by Casella et al. (2012) and its generalization by Casella and Turban (2012) to an asymmetric environment where the size of the two opposing groups is known to differ. We characterize the equilibrium for the parametrization used in the experiment, and show that an ex ante equilibrium exists and exhibits a significant volume of trade. Because the competition for votes between the two groups depends on the relative intensities of preferences of the highest intensity majority voter and highest intensity minority voter only, regardless of the size of the two groups, there is too much trade. Relative to utilitarian welfare or ex ante efficiency, the minority wins too often, and the theory predicts efficiency *losses* relative to the no-trade voting outcome.

We conduct laboratory experiments to explore the extent to which the actual outcomes in committees correspond to the equilibrium outcomes of the theoretical models of exchange. In line with the theoretical predictions, we observe efficiency gains to vote trading only when trading is centralized through party leaders. However, the efficiency gains with equal sized committees fall short of the first best. We observe efficiency losses in the experimental committees that engage in decentralized trade. Again in line with theory, in every single experimental session we observe too few minority victories, relative to first best efficiency, if trading occurs through party leaders, and too many if it occurs through the market. Prices converge towards the theoretical equilibrium prices when there is decentralized trade and when there is centralized trading between equal sized parties, but not with centralized trading between unequal sized parties. This latter observation confirms one of the conclusions in Radner and Schotter (1989).

Our theoretical results can be extended in a number of directions. First, with or without party leaders, the model should allow committees to consider more than one issue, and thus introduce the possibility of log-rolling, or vote trading *across* issues. This variety of vote trading is believed to be common practice in real committees, and could be accomplished with or without the use of a numeraire commodity. Second, it would be interesting to study more general specifications of preferences, in particular the spatial representation of preferences that has become the standard model for theoretical and empirical work in political science. Finally, our model does not address the complex strategic issues related to agenda setting and proposal power. We have taken as exogenous the proposal to be voted upon. In practice, votes are taken only after a proposal has been made, and proposal-making itself would need to consider the possibility of vote trading that can take place between the proposal stage and the voting stage. One might conjecture that vote trading could dilute the proposal power of the agenda setter.

If our theoretical results are robust to these generalizations, they suggest interesting lines of thought for empirical work. If party discipline translates into more control by party leaders and more centralized trading, according to our analysis stronger party discipline will also imply fewer vote trades and fewer minority victories. In principle, this could be tested.

References

[1] Allen, B. and J. S. Jordan (1998). "The existence of rational expectations equilibrium: a retrospective," Staff Report 252, Federal Reserve Bank of Minneapolis.

[2] Arrow, K. and F. Hahn (1971). *General Competitive Analysis.* San Francisco: Holden-Day.

[3] Bernholz, P. (1973), "Logrolling, Arrow Paradox and Cyclical Preferences", *Public Choice*, 15, 87–96.

[4] Bernholz, P. (1974), "Logrolling, Arrow Paradox and Decision Rules: A Generalization", *Kyklos,* 27, 49–62.

[5] Brams, S.J. and W.H. Riker (1973). "The Paradox of Vote Trading", *American Political Science Review*, 67, 1235–1247.

[6] Buchanan, J.M. and G. Tullock (1962). *The Calculus of Consent.* Ann Arbor: University of Michigan Press.

[7] Casella, A, A. Llorente-Saguer, and T. Palfrey (2012). "Competitive Equilibrium in Market for Votes", *Journal of Political Economy*,120, 593-658.

[8] Casella, A., T. Palfrey and S. Turban (2012). "Vote Trading With and Without Party Leaders", Social Science Working Paper 1359, California Institute of Technology, January 2013.

[9] Casella, A. and S. Turban (2012). "Democracy Undone. Systematic Minority Advantage in Competitive Markets for Votes", NBER W.P. 18573.

[10] Choulakian, V., R. A. Lockhart and M. A. Stephens (1994), "Cramér-von Mises Statistics fro Discrete Distributions", *Canadian Journal of Statistics*, 22, 125-137.

[11] Coleman, J. (1966). "The possibility of a social welfare function", *American Economic Review*, 56, 1105–1122.

[12] Coleman, J. (1967). "Reply", *American Economic Review*, 57, 1311–1317

[13] Cramton, P., M. Gibbons and P. Klemperer (1987). "Dividing a Partnership Efficiently", *Econometrica*, 55, 615-632.

[14] Dal Bò, E. (2007). "Bribing Voters", *American Journal of Political Science*, 51, 789-803.

[15] Davis, D. and C.Holt (1992). *Experimental Economics*, Princeton, N.J.:Princeton University Press.

[16] Dekel, E., M.O. Jackson and A. Wolinsky (2008). “Vote Buying: General Elections”, *Journal of Political Economy*, 116, 351-380.

[17] Dekel, E., M.O. Jackson and A. Wolinsky (2009). “Vote Buying: Legislatures and Lobbying,” *Quarterly Journal of Political Science*, 4, 103-128.

[18] Demichelis, S. and K. Ritzberger (2007). ”Corporate Control and the Stock Market,” *Carlo Alberto Notebooks* 60, Collegio Carlo Alberto, Torino.

[19] Dhillon, A. and S. Rossetto (2011). ”The Role of Voting in the Ownership Structure,” unpublished manuscript, University of Warwick.

[20] Ferejohn, J. (1974). “Sour Notes on the Theory of Vote Trading”, *Social Science Working Paper #41*, California Institute of Technology, Pasadena, California.

[21] Forsythe, R., T. Palfrey, and C. Plott (1982). ”Asset Valuation in an Experimental Market”, *Econometrica*, 50, 537-68.

[22] Friedman, D. (1984). ”On the Efficiency of Experimental Double Auction Markets”, *American Economic Review*, 74, 60-72.

[23] Friedman, D. (1993). ”The double auction market institution: A survey” in D. Friedman and J. Rust (eds.), *The Double Auction Market: Institutions, Theories and Evidence*, Santa Fe’ Institute Studies in the Science of Complexity, New York: Addison-Wesley

[24] Gray, P., and C. Plott (1990). “The Multiple Unit Double Auction,” *Journal of Economic Behavior and Organization*, 13, 245-258.

[25] Groseclose, T. J., and J. M. Snyder Jr. (1996). “Buying Supermajorities”, *American Political Science Review*, 90, 303–15.

[26] Haefele, E. (1971). “A Utility Theory of Representative Government”, *American Economic Review*, 61, 350–365.

[27] Kadane, J.B. (1972). “On Division of the Question”, *Public Choice*, 13, 47–54.

[28] Kagel, J. and A. Roth (1993). *Handbook of Experimental Economics*, Princeton, N.J.:Princeton University Press.

[29] Koford, K. (1982). “Centralized vote trading”, *Public Choice*, 39, 245–68.

[30] Kultti, K. and H. Salonen (2005). "Market for Votes", *Homo Oeconomicus*, 22, 323-332.

[31] Mailath, G. and A. Postlewaite (1990). "Asymmetric Information Bargaining Problems with Many Agents," *Review of Economic Studies*, 57, 351-67.

[32] McKelvey, R. D., and Ordeshook, P. C. (1980). "Vote Trading: An Experimental Study", *Public Choice*, 35, 151–184.

[33] Mueller, D.C. (1973). "Constitutional Democracy and Social Welfare", *Quarterly Journal of Economics*, 87, 61–79.

[34] Myerson, R. (1993). "Incentives to Cultivate Favorite Minorities under Alternative Voting Systems", *American Political Science Review*, 87, 856-869.

[35] Myerson, R and M. Satterthwaite (1983). "Efficient Mechanisms for Bilateral Trading", *Journal of Economic Theory*, 29, 265-281.

[36] Park, R.E. (1967). "The Possibility of a Social Welfare Function: Comment", *American Economic Review*, 57, 1300–1304.

[37] Philipson, T. and J. Snyder (1996). "Equilibrium and Efficiency in an Organized Vote Market", *Public Choice*, 89, 245-265.

[38] Piketty, T. (1994). "Information Aggregation through Voting and vote trading", unpublished, available at: http://www.jourdan.ens.fr/piketty/fichiers/public/Piketty1994c.pdf.

[39] Plott, C. (1982). "Industrial Organization Theory and Experimental Economics", *Journal of Economic Literature*, 20, 1485-1527.

[40] Plott, C. and V. Smith (1978). "An Experimental Examination of Two Exchange Institutions", *Review of Economic Studies*, 45, 133-153.

[41] Radner, R. and A. Schotter (1989). "The Sealed-Bid Mechanism: An Experimental Study", *Journal of Economc Theory*, 48, 179-220.

[42] Riker, W.H. and S.J. Brams (1973). "The Paradox of Vote Trading", *American Political Science Review*, 67, 1235–1247.

[43] Schwartz, T. (1977). "Collective Choice, Separation of Issues and Vote Trading", *American Political Science Review*, 71, 999–1010.

[44] Schwartz, T. (1981). "The Universal Instability Theorem", *Public Choice*, 37, 487–501.

[45] Shubik, M. and L. Van der Heyden (1978). “Logrolling and Budget Allocation Games”, *International Journal of Game Theory*, 7, 151–162.

[46] Smith, V. (1965). ”Experimental Auction Markets and the Walrasian Hypothesis”, *Journal of Political Economy*, 73, 387-93.

[47] Smith, V. (1982). ”Microeconomic Systems as an Experimental Science”, *American Economic Review*, 72, 923-55.

[48] Tullock, G. (1970). ”A Simple Algebraic Logrolling Model”, *American Economic Review*, 60, 419-426.

[49] Valley, K., L. Thompson, R. Gibbons and M. H. Bazerman (2002), ”How Communication Improves Efficiency in Bargaining Games”, *Games and Economic Behavior*, 38, 127-155.

[50] Starr, R.M. (1969). “Quasi-Equilibria in Markets with Non-Convex Preferences”, *Econometrica*, 17, 25-38.

[51] Weiss, J.H. (1988). “Is Vote-Selling Desirable?”, *Public Choice*, 59, 117–194.

[52] Wilson, R. (1969). ”An Axiomatic Model of Logrolling”, *American Economic Review*, 59, 331-341.

Appendix A. Existence and Characterization of Ex Ante Competitive Equilibrium ($M = 3$ $m = 2$)

In an ex ante equilibrium, demands are allowed to be stochastic, the market clears in expectation, and a rationing rule determines the ex post allocation of votes. Call δ_i i's mixed demand: δ_i is a discrete probability distribution over support $[-1, 0, 1, .., n-1]$.[30] Call δ_{-i} the profile of demands excluding δ_i, d_i the realization of δ_i, and R a rationing rule establishing how votes are allocated if $\sum_i d_i \neq 0$.

Definition. A price p, a profile of demands δ, and a rationing rule R constitute *an ex ante competitive equilibrium* if $Eu_i(\delta_i, \delta_{-i}, p, R) \geq Eu_i(\widetilde{\delta_i}, \delta_{-i}, p, R)$ for all $\widetilde{\delta_i}$, for all i, and $\sum \delta_i = 0$.[31]

As in Casella et al. (2012), we use the rationing rule inspired by *All-Or-Nothing* orders in financial markets: either a voter fulfills his demand completely or is excluded from trade. More precisely, any individual with positive demand is considered with equal probability; in case his demand cannot be satisfied, the voter is left with his initial endowment, and the process goes on to another randomly selected voter with positive demand. We call this rationing rule $R1$.[32]

Denote by $\overline{v}_M$ ($\underline{v}_M$) and $\overline{v}_m$ ($\underline{v}_m$) the maximal (minimal) realized value of a voter in group M and m, respectively Then:

Proposition *Let $M = 3$ and $m = 2$. The rationing rule is $R1$. Then for all realizations of $\{v_1, .., v_5\}$ such that $\overline{v}_m \geq \underline{v}_M$, a fully revealing ex ante equilibrium with trade exists. The equilibrium price is always such that* $\min(\overline{v}_m, \overline{v}_M)$ *is indifferent between demanding two votes and selling his vote.*[33]

In the spirit of rational expectations models (Allen and Jordan, 1998), we call an equilibrium *fully revealing* if either: (1) the equilibrium price, together with the set of others' equilibrium strategies and market equilibrium, fully convey to voter i the direction of preferences associated to each demand; or (2) the information conveyed is partial but voter i has a unique best response, identical to his best response under full information.

We prove the Proposition by construction. Call g' the group such that $\overline{v}_{g'} \geq \overline{v}_g$, and call

[30] A demand of -1 corresponds to selling one's vote.

[31] As in the analysis of competitive equilibrium with externalities (e.g., Arrow and Hahn, 1971, pp. 132-6), the definition of the equilibrium requires voters to best reply not only to the price but also to the demands of other voters. Optimal demands are interrelated.

[32] For details see Casella et al. (2012). For clarity, note that in the experiment we do not impose any particualr rationing rule but let the market find its own equilibrium.

[33] A trivial no-trade equilibrium always exists, as well, because if all other voters are neither buying nor selling, then being inactive is a best response. Our interest is in nontrivial equilibrium with trade.

$v_{(2)g'}$ the second highest value in group g'. The equilibrium is characterized in the following three lemmas.

Lemma 1 *If $\overline{v}_g \geq (2/7)v_{(2)g'}$, then there exists a fully revealing ex ante equilibrium such that voters $\overline{v}_{g'}$ and $\overline{v}_g$ randomize between demanding two votes and selling their vote (with probabilities $\sigma_{\overline{v}_{g'}}$,and $\sigma_{\overline{v}_g}$) and all other voters offer to sell their vote. The randomization probabilities and the equilibrium price depend on the value realizations. In particular: (a) If $\overline{v}_m \in [(2/7)v_{(2)M}, (3/5)\overline{v}_M]$, then $\sigma_{\overline{v}_M} = 0$, $\sigma_{\overline{v}_m} = 1/3$, and $p = \overline{v}_m/3$; (b) If $\overline{v}_m \in (3/5\overline{v}_M, 5/6\overline{v}_M)$, then $\sigma_{\overline{v}_M}$, $\sigma_{\overline{v}_m}$, and p are solutions to the system:*

$$1 = 3(\sigma_{\overline{v}_M} + \sigma_{\overline{v}_m})$$
$$p = \overline{v}_M\left(\frac{1-\sigma_{\overline{v}_m}}{3+\sigma_{\overline{v}_m}}\right)$$
$$p = \overline{v}_m\left(\frac{1+\sigma_{\overline{v}_M}}{3+\sigma_{\overline{v}_M}}\right)$$

(c) If $\overline{v}_M \in [(2/7)v_{(2)m}, (6/5)\overline{v}_m]$, then $\sigma_{\overline{v}_M} = 1/3$, $\sigma_{\overline{v}_m} = 0$, and $p = \overline{v}_M/3$.

Lemma 2 *If $\overline{v}_g \in [(1/14)v_{(2)g'}, (2/7)v_{(2)g'}]$, then there exists a fully revealing ex ante equilibrium such that voter $\overline{v}_{g'}$ demands two votes, $v_{(2)g'}$ randomizes between demanding one vote and offering his vote for sale (with probability $\sigma_{v_{(2)g'}}$),$\overline{v}_g$ randomizes between demanding two votes and selling his vote (with probability $\sigma_{\overline{v}_g}$), and all other voters offer to sell their vote. The randomization probabilities and the equilibrium price are solutions to the system:*

$$3 = 2\sigma_{v_{(2)g'}} + 3\sigma_{\overline{v}_g}$$
$$p = v_{(2)g'}\left(\frac{1-\sigma_{\overline{v}_g}}{6+3\sigma_{\overline{v}_g}}\right)$$
$$p = \overline{v}_g\left(\frac{2+\sigma_{v_{(2)g'}}}{10-\sigma_{v_{(2)g'}}}\right)$$

Lemma 3 *If $\overline{v}_g \leq (1/14)v_{(2)g'}$, then for all realizations of $\{v_1, .., v_5\}$ such that $\overline{v}_m \geq \underline{v}_M$ there exists a fully revealing ex ante equilibrium such that voters $\overline{v}_{g'}$ and $v_{(2)g'}$ demand one vote, $\overline{v}_g$ randomizes between demanding two votes and selling his vote (with probability $\sigma_{\overline{v}_g} = 2/3$) and all other voters offer to sell their vote. The equilibrium price is $\overline{v}_g/4$.*

We reproduce below the proof of Lemma 1. The proofs of Lemmas 2 and 3 are similar and the details are available as an online supplementary material. (See Appendix 5) of this manuscript version).

Proof of Lemma 1. If voters' preferred alternative is known, establishing that the candidate strategies and price are an equilibrium follows immediately from comparing the expected utilities of different actions, given others' strategies. Call $EU_{v_{ig}}A$ the expected utility of voter with value v_i belonging to group g from action $A \in \{S, 0, D1, D2\}$, with obvious notation. In this case, allowing for both $\sigma_{\overline{v}_M} > 0$ *and* $\sigma_{\overline{v}_m} > 0$:

$$\begin{aligned}
EU_{\overline{v}_M}D2 &= (\overline{v}_M - 2p)(1 + \sigma_{\overline{v}_m})/2 \\
EU_{\overline{v}_M}D1 &= \sigma_{\overline{v}_m}\overline{v}_M - p \\
EU_{\overline{v}_M}0 &= \sigma_{\overline{v}_m}\overline{v}_M \\
EU_{\overline{v}_M}S &= \sigma_{\overline{v}_m}\overline{v}_M + (1 - \sigma_{\overline{v}_m})p/2
\end{aligned}$$

$$\begin{aligned}
EU_{\overline{v}_m}D2 &= (\overline{v}_m - 2p)(1 + \sigma_{\overline{v}_M})/2 \\
EU_{\overline{v}_m}D1 &= \sigma_{\overline{v}_M}3\overline{v}_m/4 - p \\
EU_{\overline{v}_m}0 &= 0 \\
EU_{\overline{v}_m}S &= (1 - \sigma_{\overline{v}_M})p/2
\end{aligned}$$

$$\begin{aligned}
EU_{v_{iM}}D2 &= \sigma_{\overline{v}_M}[\sigma_{\overline{v}_m}(v_{iM} - 2p) + (1 - \sigma_{\overline{v}_m})(v_{iM}/2 - p)]+ \\
&\quad + (1 - \sigma_{\overline{v}_M})[\sigma_{\overline{v}_m}(v_{iM} - p) + 2(1 - \sigma_{\overline{v}_m})(v_{iM} - p)/3] \\
EU_{v_{iM}}D1 &= \sigma_{\overline{v}_M}[\sigma_{\overline{v}_m}(v_{iM} - p) - (1 - \sigma_{\overline{v}_m})p]+ \\
&\quad + (1 - \sigma_{\overline{v}_M})[\sigma_{\overline{v}_m}(v_{iM} - p) + (1 - \sigma_{\overline{v}_m})(2v_{iM} - p)/3] \\
EU_{v_{iM}}0 &= \sigma_{\overline{v}_m}[\sigma_{\overline{v}_M}v_{iM} + (1 - \sigma_{\overline{v}_M})v_{iM}] + (1 - \sigma_{\overline{v}_m})(1 - \sigma_{\overline{v}_M})v_{iM}/2 \\
EU_{v_{iM}}S &= \sigma_{\overline{v}_M}[\sigma_{\overline{v}_m}v_{iM} + (1 - \sigma_{\overline{v}_m})p/2]+ \\
&\quad + (1 - \sigma_{\overline{v}_M})[\sigma_{\overline{v}_m}(v_{iM} + p/2) + (1 - \sigma_{\overline{v}_m})(v_{iM}/2 + 2p/3)]
\end{aligned} \tag{S1}$$

$$\begin{aligned}
EU_{v_{im}}D2 &= \sigma_{\overline{v}_M}[\sigma_{\overline{v}_m}(v_i - 2p) + (1 - \sigma_{\overline{v}_m})(v_{im} - p)]+ \\
&\quad + (1 - \sigma_{\overline{v}_M})[\sigma_{\overline{v}_m}(v_{im}/2 - p) + 2(1 - \sigma_{\overline{v}_m})(v_{im} - p)/3] \\
EU_{v_{im}}D1 &= \sigma_{\overline{v}_M}[\sigma_{\overline{v}_m}(3v_{im}/4 - p) + (1 - \sigma_{\overline{v}_m})(v_{im} - p)]+ \\
&\quad + (1 - \sigma_{\overline{v}_M})[\sigma_{\overline{v}_m}(-p) + (1 - \sigma_{\overline{v}_m})(2v_{im} - p)/3] \\
EU_{v_{im}}0 &= (1 - \sigma_{\overline{v}_m})[\sigma_{\overline{v}_M}v_{im} + (1 - \sigma_{\overline{v}_M})v_{im}/2] \\
EU_{v_{im}}S &= (1 - \sigma_{\overline{v}_m})\sigma_{\overline{v}_M}(v_{im} + p/2) + (1 - \sigma_{\overline{v}_M})[\sigma_{\overline{v}_m}p/2 + (1 - \sigma_{\overline{v}_m})(v_{im}/2 + 2p/3)]
\end{aligned}$$

where $v_{iM} \leq \overline{v}_M$, and $v_{im} \leq \overline{v}_m$. Expected market balance requires $\sum \delta_i = 0$, where δ_i is individual i's expected demand, or $2(1 - \sigma_{\overline{v}_M}) + 2(1 - \sigma_{\overline{v}_m}) = 3 + \sigma_{\overline{v}_M} + \sigma_{\overline{v}_m}$, or $\sigma_{\overline{v}_M} + \sigma_{\overline{v}_m} = 1/3$. Given the equations in System (S1), it follows immediately that $\overline{v}_M$ and

$\overline{v}_m$ are both indifferent between $D2$ and S if:

$$p = \overline{v}_M \left(\frac{1 - \sigma_{\overline{v}_m}}{3 + \sigma_{\overline{v}_m}} \right)$$
$$p = \overline{v}_m \left(\frac{1 + \sigma_{\overline{v}_M}}{3 + \sigma_{\overline{v}_M}} \right)$$

It is not difficult to verify that expected market balance and the two indifference conditions can be satisfied simultaneously at $\sigma_{\overline{v}_M} \in (0, 1]$, $\sigma_{\overline{v}_m} \in (0, 1]$ only if $\overline{v}_m \in (3/5\overline{v}_M,\ 5/6\overline{v}_M)$. If $\overline{v}_m \geq (5/6)\overline{v}_M$, the price that makes $\overline{v}_M$ indifferent between $D2$ and S is too low to induce $\overline{v}_m$ to sell with positive probability: the equilibrium must then have $\sigma_{\overline{v}_m} = 0$, $\sigma_{\overline{v}_M} = 1/3$, and $p = \overline{v}_M/3$. If $\overline{v}_M \geq (5/3)\overline{v}_m$, the price that makes $\overline{v}_m$ indifferent between $D2$ and S is too low to induce $\overline{v}_M$ to sell with positive probability: the equilibrium must have $\sigma_{\overline{v}_M} = 0$, $\sigma_{\overline{v}_m} = 1/3$, and $p = \overline{v}_m/3$. Establishing that the stated strategies are best responses to each other is trivial, given p, $\sigma_{\overline{v}_m}$ and $\sigma_{\overline{v}_M}$ and the equations in System (S1). In addition, if $\overline{v}_m > \overline{v}_M$, the condition $v_{im} \leq (7/2)\overline{v}_M$ is required to prevent the profitable deviation of voter v_{im} to demanding a positive number of votes; similarly if $\overline{v}_M > \overline{v}_m$, the condition $v_{iM} \leq (7/2)\overline{v}_m$ is required to guarantee that selling is a best response for voter v_{iM}. The intuition is straightforward: if $v_{im} > (7/2)\overline{v}_M$, the price that makes $\overline{v}_M$ indifferent between selling and demanding two votes is too low to induce v_{im} to sell, as this equilibrium prescribes; and similarly if $v_{iM} > (7/2)\overline{v}_m$.

Finally, we need to show that the equilibrium is fully revealing. First notice that there can be no equilibrium with trade where both m members offer to sell their vote with probability one–because no M member would have an incentive to buy. Hence $\overline{v}_M$ knows that in equilibrium the other voter with positive expected demand belongs to group m; of the sellers, two must belong to M and one to m. Consider now the problem from the point of view of $\overline{v}_m$. Given others' equilibrium strategies, expected market balance requires $\overline{v}_m$ to demand a positive number of votes. It is not difficult to verify that at p if the voter mixing between $D2$ and S with probability $\sigma_{\overline{v}_M}$ belonged to m, $\overline{v}_m$'s best response is S. However, S does not satisfy expected market balance. Hence $\overline{v}_m$ knows that in equilibrium the other voter with positive expected demand belongs to group M; again, of the sellers, two must belong to M and one to m. As for the sellers, market balance requires each of them to sell with probability one. Among them, each M member knows that the two voters with positive expected demand cannot both belong to group m, by the argument above; not can they both belong to group M, because in equilibrium at least one m member must demand votes with positive probability. Similarly, the m seller knows that the other m member cannot also be selling with probability one. Hence, each seller knows that one but not both of the voters with positive demands must belong to his own group; the seller cannot know which one, but is indifferent: the unique best response is to sell. Thus the equilibrium is fully revealing.$\square$ $\blacksquare$