Electoral Engineering: One Man, One ~~Vote~~ Bid

Jacob K. Goeree and Jingjing Zhang*

March 18, 2013

Abstract

We compare two mechanisms to implement a simple binary choice, e.g. adopt one of two proposals. We show that when neither alternative is ex ante preferred, simple majority voting cannot implement the first best outcome. The fraction of the surplus lost *rises* with the number of voters and the total surplus loss *diverges* in the limit when the number of voters grows large. We introduce a simple bidding mechanism where votes can be bought at a quadratic cost and voters receive rebates equal to the average of others' payments. This mechanism is budget-balanced, individually rational, and fully efficient in the limit. Moreover, the mechanism redistributes from those that gain from the outcome to those that lose and *everyone* is better off under bidding compared to voting.

We test the two mechanisms in the lab using an environment with "moderate" and "extremist" voters. In the first part of the experiment, subjects gain experience with both bidding and voting. Then they collectively decide which mechanism applies in the second part. This endogenous choice of institutions provides clear evidence in favor of bidding: with groups of size eleven, 90% of the groups opt for the bidding mechanism. The observed efficiency losses under voting are close to theoretical predictions and significantly larger than under bidding. Because of redistribution, the efficiency gain from bidding benefits mostly the moderate voters. Observed behavior under the bidding mechanism deviates from theoretical predictions to some extent, which can be explained by a quantal response equilibrium model if we assume that voters overestimate the chance of being pivotal. These deviations hardly reduce any of the efficiency-improving effect of the mechanism, which further supports the adoption of this mechanism in the real world applications.

Keywords: *electoral design, information aggregation, experiments, probability weighting, quantal response equilibrium*

JEL Code: C92

*Chair for Organizational Design, Department of Economics, University of Zürich, Blümlisalpstrasse 10, CH-8006, Zürich, Switzerland. We gratefully acknowledge financial support from the Swiss National Science Foundation (SNSF 135135) and the European Research Council (ERC Advanced Investigator Grant, ESEI-249433). We are grateful to Nixon Li for his input and help with an earlier related project "Inefficient Voting." We would like to thank seminar participants at the Max Planck Institute in Bonn and at the Workshop for Behavioral Public Economics in Copenhagen for useful comments.

1. Introduction

Many countries use referenda to decide whether or not to undertake public projects, implement a policy change, or to accept a new law. Most US states, for example, add bond issues for schools or roads as "yes/no" items on the ballot of a general election. European countries have issued national referenda to decide whether to join the European Union, accept European lawmaking, and a common currency. Referenda are also commonly employed on a local level, e.g. whether or not to accept a change in zoning restrictions. This type of "direct democracy" is particularly prevalent in Switzerland where business hours, vacation duration, etc., are often settled by referendum.

In these examples, it is typically the case that some voters care more about which outcome prevails than others. Changes in zoning laws might preclude some families from sending their kids to a good public school and cause the value of their homes to drop. Other families may simply be unaffected. Likewise, relaxing opening and closing times in Zürich may be very valuable to store owners and a select group of shoppers while others may see little (or negative) value. It is an important question whether simple referenda can achieve the socially optimal outcome when voters differ not only in terms of which outcome they prefer but also in terms of their intensity, i.e. how much they care about either outcome.

There are several stylized cases where voting is easily seen to be optimal. Suppose each voter is more likely than not to favor increasing public spending on local schools, and each proponent values such an increase more than an opponent values the status quo level of spending. For a large enough electorate, simple majority voting will then result in the optimal outcome. Next, consider the case where each voter is less likely (say, 40% likely) to favor an increase, but proponents value such an increase three times more than those opposed value keeping current spending levels. Now, majority voting is no longer optimal but a super-majority rule is: one should require the status quo to receive at least 75% of all the votes since conditional on receiving more/less than 75% of the votes the expected total benefit of keeping the status quo is higher/lower. For a large electorate, the increase will be implemented as only 60% ($< 75\%$) favor the status quo – the outcome is again fully efficient in the limit.[1]

The previous situations are "easy" in that they have a clear ex ante preferred outcome, i.e. an outcome with a higher expected value. But when it is clear that one of the outcomes has higher value on average, and the electorate is sufficiently large, then why vote at all? It would suffice to simply command that the outcome with the higher expected value be implemented,

[1] The need for super-majority rules when preference intensities vary across alternatives is more systematically explored by Ledyard and Palfrey (2002).

unless voting serves only the purpose of ratification. In many instances, however, the main reason for calling a referendum is that it is *not* clear which of the two alternatives is ex ante preferred. In such cases, the referendum serves the important role of an information aggregation tool.

In this paper we consider binary-choice problems where voters' preferences are symmetrically distributed across two alternatives so that, ex ante, there is no preferred outcome. Of course, given the actual value realizations, one of the alternatives is preferred ex post and the question is what mechanism implements the first best outcome. Majority voting is not the answer because inefficiencies arise when the mean and median value are on opposite sides, which occurs with non-negligible probability. We show that the fraction of surplus realized under voting is strictly less than 100% by Jensen's inequality (Proposition 1). Importantly, the fraction of the surplus that is lost is *increasing* in the size of the electorate and the total surplus loss *diverges* in the limit.

We propose a novel electoral design that allows individuals to express an intensity, or bid, for their preferred outcome. In this bidding mechanism, individuals can buy votes at a quadratic cost and receive a rebate equal to the average of others' payments. The alternative that receives the largest total bid is selected. We show that the bidding mechanism is budget-balanced, individually rational,[2] and fully efficient in the limit (Proposition 2). The intuition is that, in a large electorate, the impact of a bid on the probability that the preferred outcome is selected is small, and, hence, approximately linear in the bid. As a result, the marginal benefit of raising one's bid is proportional to one's value. With a quadratic payment rule the marginal cost is linear in the bid. Equating the marginal cost and benefit yields optimal bids that are proportional to value. And when bids are proportional to value, the alternative with the largest total bid is also the one that is value maximizing: bidding is efficient.

As noted above, another important drawback of voting is that there is no compensation for those harmed by the outcome.[3] This "tyranny of the majority" seems untenable in many cases,

[2]We consider situations without a status quo, e.g. when one of two proposals will replace the current policy. With a status quo, the probability that a new proposal is accepted may vanish when the collective grows large even when it is common knowledge that adopting the proposal is welfare improving (Mailath and Postlewaite, 1990). This negative result stems from a clash between the requirement of efficiency, budget balancedness, and incentive compatibility on the one hand and voluntary participation (individual rationality) on the other hand. Without a status quo, individual rationality is satisfied. But even with a status quo, there are many situations where voluntary participation constraints do not play a role either because participants cannot "escape" the outcome of the mechanism or compliance is enforced. For example, a country's tax scheme to fund public goods is typically enforced by law with severe punishments for those who evade it. Likewise, homeowners cannot simply avoid the consequences of a change in zoning laws by not participating in the relevant referendum.

[3]Varian (1994) describes a procedure for inducing socially optimal choices by requiring participants to compensate those harmed by their decisions. This procedure has been tested with positive results in the lab (An-

e.g. when a change in zoning restrictions causes house values to drop it seems unreasonably harsh not to compensate its owners. Redistribution from those that benefit to those that are harmed is automatic in our bidding mechanism. Indeed, we show that under the bidding mechanism *everyone* is better off. In other words, the increase in overall welfare does not simply accrue to those with more "extreme" valuations but is evenly distributed among voters of all types (Proposition 2).

We test our bidding mechanism in the laboratory and compare its performance to majority voting. The experiments employ a setup with "moderate" (low-value) and "extremist" (high-value) voters who are equally likely to prefer one of two alternatives. In the first twenty periods of the experiment, subjects make choices under both the bidding and voting mechanism. Then they collective choose whether the bidding or voting mechanism applies in the final twenty periods of the experiment.[4] This endogenous choice of institutions provides perhaps the clearest evidence for the desirability of bidding versus voting: with electorates of size eleven, we find that 90% of the groups (18 out of 20) opt for the bidding mechanism.

An additional benefit of our experimental design is that subjects have ample opportunity to learn in the first part, which makes equilibrium behavior in the second part more likely. We find that voting is near-perfect in this second part: only 0.7% of all votes are "mistakes" that go against the preferred alternative. As a result, observed efficiency losses under voting are as theory predicts: 22% on average with groups of size three and 28% with groups of size eleven. Observed efficiency losses are much smaller under the bidding mechanism: 7% with a group size of three and 13% with a group size of eleven. Furthermore, the bidding mechanism results in significant levels of redistribution and it is the "moderate" voter who benefits the most.

In the experiments, bidding is more efficient than voting because those with higher values buy more votes than those with lower values. However, the observed ratio of high to low bids is less extreme than theoretically predicted. In particular, low-value voters tend to buy too many votes while high-value voters buy close to the equilibrium amount. As a result, bidding is not 100% efficient. Another reason is that bidding tends to be somewhat "noisy." To explore this further we estimate a quantal response equilibrium model assuming that voters overestimate

dreoni and Varian, 1999; Chen and Gazzale, 2004). Unlike the bidding mechanism that we propose, Varian's (1994) procedure works only when the costs and benefits are commonly known.

[4]This endogenous choice of institution is inspired by Wicksell's (1896) principle of unanimity and voluntary consent in taxation: "if a proposal yields a common outcome that applies to all members of a collective, and the outcome provides higher overall welfare, then there exists a compensation scheme whereby those that gain from the outcome compensate those that are harmed in a manner such that everyone is better off. As a result, the proposal will be accepted unanimously." In our experiments there is no status quo so it is unclear which of the two institutions would require unanimous consent. To not favor either bidding or voting, we let the choice of institution be determined by majority voting.

the chance of being pivotal. We find that this model explains the individual and aggregate data quite well.

1.1. Related Approaches

The need to allow for an expression of preference intensity to improve on voting outcomes can also be found in the work of Casella (2005) who introduced the idea of "storable votes." When voters face a series of binary choices and are endowed with a budget of votes, they can "store" votes when indifferent about an issue and allocate more votes to issues they feel strongly about. Casella (2005) shows that storable votes can be welfare improving relative to the case where each voter gets a single vote per issue, a prediction that has been corroborated by laboratory experiments (Casella, Gelman, and Palfrey, 2006).

In contrast to Casella (2005), we study a single binary choice for which the introduction of a budget of votes has no effect (i.e. the entire budget of votes would be used). Moreover, the bidding mechanism we propose is fully efficient and redistributes from those that benefit to those that are harmed. Unlike the storable votes approach, our bidding mechanism could be applied to classic hold out problems such as land assemblages for commercial development (e.g. shopping malls) or public projects (e.g. highways or power lines).

Another related approach is Smith's (1977) "compensation election," which was tested in laboratory experiments by Oprea, Smith, and Winn (2007). In this mechanism, voters submit bids and the alternative with the largest total bid is implemented. Those that win pay their bids and those that lose get paid their (own) bids. Unlike our proposed bidding mechanism, the compensation election is thus not budget balanced. Furthermore, as pointed out by Oprea, Smith, and Winn (2007), no equilibrium model exists to predict behavior in the compensation election. Finally, the experimental setup of Oprea, Smith, and Winn (2007) does not allow them to test whether the compensation election improves efficiency compared to voting, which is the main point of this paper.

1.2. Organization

The paper is organized as follows. In the next section we show that majority voting is inefficient when neither alternative is ex ante preferred. We start with an example of normally distributed values for which welfare under voting and the optimal welfare can be computed exactly. We then use the central limit theorem to show that voting is inefficient for arbitrary value distributions. In Section 3 we introduce our bidding mechanism and show that it is fully efficient in the limit when the electorate grows large. We test both the voting and bidding mechanisms in

the laboratory. Section 4 describes the design of the experiments, the procedures, and the hypotheses. The experimental results are discussed in Section 5. Section 6 concludes. The instructions for the experiment can be found in the Appendix.

2. Inefficient Voting

Suppose there are $2n+1$ voters whose values are symmetrically distributed according to a normal distribution, $\Phi(v)$, with mean zero and variance one. Here a negative value corresponds to a preference for alternative A while a positive value reflects a preference for alternative B. A voter with value v receives $|v|$ if her preferred alternative is implemented and $-|v|$ otherwise. Thus $|v|$ captures the intensity of the preference.

Under majority voting, a voter with value v generates a surplus equal to $|v|$ when $k \geq n$ others prefer the same alternative and $-|v|$ when $k < n$ others prefer the same alternative. Since the event where k others prefer the same alternative is equally likely as the event where $2n-k$ others prefer the same alternative, all terms cancel except the one for $k = n$. The ex ante expected surplus, or expected welfare, under simple majority voting is thus

$$W_{voting} \;=\; (2n+1)\sqrt{\frac{2}{\pi}}\binom{2n}{n}\Big(\frac{1}{2}\Big)^{2n}$$

where we used that $E(|v|) = \sqrt{2/\pi}$ when values are normally distributed.

Under an efficient mechanism, a voter with value v generates a surplus equal to v when the sum of others' values is greater than $-v$, which occurs with probability $\Phi(v/\sqrt{2n})$, and she generates a surplus equal to $-v$ with complementary probability. Expected welfare is therefore

$$W_{optimal} \;=\; (2n+1)\int_{-\infty}^{\infty} v\Big(2\Phi\big(\frac{v}{\sqrt{2n}}\big)-1\Big)d\Phi(v) \;=\; \sqrt{\frac{2}{\pi}(2n+1)}$$

The open orange circles in the left panel of Figure 1 show the fraction of the surplus realized by majority voting for various electorate sizes. Note that efficiency $W_{voting}/W_{optimal}$ is strictly *decreasing* in the size of the electorate and falls from 100% when the size of the electorate is one to its limit value

$$\lim_{n\to\infty}\frac{W_{voting}}{W_{optimal}} \;=\; \sqrt{\frac{2}{\pi}}$$

as indicated by the dashed line. (The filled blue circles show the fraction of total surplus realized by the bidding mechanism, which is discussed in the next section.) Since the fraction

of the surplus lost grows with the size of the electorate the total surplus loss $W_{optimal} - W_{voting}$ diverges to $+\infty$ in the limit.

Remark 1. One misconception is that since neither alternative is ex ante preferred, flipping a coin would do as well as voting. The mistake in this reasoning is that even though each voter is equally likely to favor either alternative ex ante, one of the alternatives is preferred by the majority ex post. Indeed, the expected difference in the total support for either alternative diverges when the electorate grows large. Voting serves as an information aggregation tool by picking out the correct alternative preferred by the majority. As a result, total expected welfare under voting diverges to $+\infty$ as the number of voters grows large, while flipping a coin would result in an expected welfare of 0. □

Remark 2. Another misconception is that while voting is known to be inefficient in small groups the inefficiency loss disappears for large electorates. As Figure 1 shows, voting gets *worse* as the size of the electorate grows. The fraction of the surplus realized by voting falls with the size of the electorate and the total surplus loss diverges. □

The assumption of normally distributed values is helpful in that the convolution of normal distributions is normal. As a result, an exact expression for the optimal welfare can be derived for all electorate sizes. For a general symmetric distribution, $F(v)$, a similar expression can be derived for large electorates using the central limit theorem. The probability that the sum of others' values is greater than $-v$ limits to $\Phi(v/\sqrt{2n}\sigma)$ where $\sigma^2 = E(v^2)$ denotes the variance. Hence, expected welfare limits to

$$W_{optimal} \approx (2n+1)\int_{-\infty}^{\infty} v\Big(2\Phi\big(\frac{v}{\sqrt{2n}\sigma}\big) - 1\Big)dF(v) \approx \sqrt{\frac{2}{\pi}(2n+1)E(v^2)}$$

where the far right expression follows from a Taylor approximation as $n \to \infty$.[5] The expression for expected welfare under voting derived above applies to any symmetric distribution if we replace $\sqrt{2/\pi}$ by $E(|v|)$. To summarize:

Proposition 1. *For any symmetric value distribution, the fraction of the total surplus realized by majority voting falls from 1 when the size of the electorate is one to*

$$\lim_{n\to\infty} \frac{W_{voting}}{W_{optimal}} = \frac{E(\sqrt{v^2})}{\sqrt{E(v^2)}}$$

which is less than 1 (for non-degenerate distributions) by Jensen's inequality. The total surplus loss $W_{optimal} - W_{voting}$ diverges in the limit.

[5]To first order $\Phi(v/\sqrt{2n}\sigma) \approx \Phi(0) + \Phi'(0)v/\sqrt{2n}\sigma = \frac{1}{2}(1 + v/\sqrt{n\pi}\sigma)$.

Our results contrast with those of Ledyard and Palfrey (2002) who show that voting is fully efficient in the limit when one of the alternatives is ex ante preferred. In this case, a super-majority rule can be used to ensure that, for large electorates, the alternative with the higher expected value wins with probability one. In fact, there are many super-majority rules that are efficient, including a command mechanism that simply implements the ex ante preferred alternative.[6] In this sense, when one alternative is ex ante preferred, the election does not need to aggregate the information that is dispersed in the electorate.

3. Efficient Bidding

An early approach to the efficient provision of public goods is due to Clarke (1971). The "Clarke Tax" mechanism is a member of the Vickrey-Clarke-Groves (VCG) family for which truth telling is a dominant strategy and agents are taxed the negative externality their choice imposes on others. The Clarke Tax mechanism is also known as the "Pivotal" mechanism since agents only pay if their choice changes the outcome, i.e. when they are pivotal.

To illustrate, suppose two voters, labeled 1 and 2, indicate a preferred alternative and their values for it. If both voters favor the same alternative, neither is pivotal and, hence, neither pays anything. If they prefer different alternatives and $|v_2| > |v_1|$ then voter 2 is pivotal and is taxed $2|v_1|$ (since her vote changes voter 1's payoff from $|v_1|$ to $-|v_1|$). Note that voter 2's payment is independent of her report and that it is a weakly dominant strategy for her to report truthfully. If she reports a different value $\hat{v}_2$ with $|\hat{v}_2| > |v_1|$ then the outcome and the tax are the same and if $|\hat{v}_2| < |v_1|$ then her payoff changes from $|v_2| - 2|v_1|$ to $-|v_2|$, which is worse.

An important shortcoming of the Clarke Tax mechanism is that it is not budget balanced, i.e. excess tax contributions are discarded. In the above example, voter 2's payment does not accrue to voter 1 but is "burned" in order to maintain incentive compatibility.[7] It is well known, however, that whenever VCG runs a surplus, an efficient, Bayesian incentive compatible mechanism that is individually rational and balances the budget exists (Krishna and Perry,

[6] Suppose, for instance, that values are distributed according to a normal distribution, $\Phi(v-\mu)$, with mean $\mu > 0$ and variance one. Then the optimal welfare is given by

$$W_{optimal} \;=\; (2n+1)\int_{-\infty}^{\infty} v\Big(2\Phi\big(\frac{v+2n\mu}{\sqrt{2n}}\big)-1\Big)d\Phi(v-\mu) \;\approx\; (2n+1)\mu + \sqrt{\frac{2}{\pi}(2n+1)}e^{-(2n+1)\mu^2}$$

A command mechanism yields $W_{command} = (2n+1)\mu$. So when $\mu > 0$ there is no efficiency loss in the limit when the electorate grows large (Ledyard and Palfrey, 2002). However, this argument breaks down when $\mu = 0$.

[7] Results from laboratory experiments by Attiyeh et al. (2000) confirm that compared to simple majority voting the Clarke Tax mechanism is inefficient because excess taxes must be discarded.

1998). This mechanism was proposed by Arrow (1979) and independently by d'Aspremont and Gérard-Varet (1979a,b) and is known as the AGV or "expected externality" mechanism. The basic idea is that voters receive rebates that depend on others' tax payments so that the budget can be balanced without affecting incentives.

The mechanism we propose is inspired by the AGV mechanism but simpler. In our bidding mechanism, votes for either alternative can be bought at a quadratic cost and voters receive rebates equal to the average of others' payments. The alternative with the largest total bid is implemented. The intuition behind our proposal is as follows. For a large electorate, the impact of anyone's bid on the probability that the preferred outcome is selected is small and approximately linear in the bid. As a result, the marginal benefit of raising one's bid is proportional to one's value. With a quadratic payment rule the marginal cost is linear in the bid. Hence, optimal bids are proportional to value and selecting the alternative with the largest total bid is value maximizing.

To formalize this intuition, let $G(b_i)$ denote the expected probability that alternative B is selected from voter i's point of view if she buys b_i votes:

$$G(b_i) \;=\; \text{Prob}\big(b_i + \sum_{j \neq i} b_j \;>\; 0\big)$$

let $C(b_i)$ denote the cost of buying b_i votes

$$C(b_i) \;=\; \frac{1}{2\sqrt{n\pi}\sigma}\, b_i^2$$

and let $R(b_{-i})$ denote the rebate voter i receives

$$R(b_{-i}) \;=\; \frac{1}{2n} \sum_{j \neq i} C(b_j)$$

which is independent of voter i's own bid.

Proposition 2. *The bidding mechanism is budget balanced and individually rational. For large electorates, truthful bidding constitutes a Bayes-Nash equilibrium and*

$$\lim_{n \to \infty} \frac{W_{bidding}}{W_{optimal}} \;=\; 1$$

Furthermore, voters of (almost) all types are better off under bidding compared to voting.

Proof. It is readily verified that $\sum_i R_i = \sum_i b_i$, i.e. the bidding mechanism is budget balanced. Consider voter i and suppose all others bid truthfully, i.e. $b_j = v_j$ for $j \neq i$. Voter i's expected payoff when her value is v_i and she bids b_i is given by

$$\pi_i(v_i, b_i) \;=\; v_i\big(2G(b_i) - 1\big) - C(b_i) + \frac{\sigma}{2\sqrt{n\pi}}$$

where the final term on the right side is the expected rebate when others bid truthfully. Since $G(0) = \frac{1}{2}$, voter i's payoff when she bids zero is equal to the rebate, which is non-negative, so the bidding mechanism is individually rational. Moreover, for a large electorate, the central limit theorem implies that $G(b_i)$ limits to $\Phi(b_i/\sqrt{2n}\sigma)$, which can be approximated by $\Phi(b_i/\sqrt{2n}\sigma) \approx \frac{1}{2}(1+b_i/\sqrt{n\pi}\sigma)$. Hence, voter i's profit is maximized at $b_i = v_i$. In other words, truthful bidding constitutes a Bayes-Nash equilibrium. The bidding mechanism is efficient because it selects the alternative with the largest total bid, which is the one with the largest total value.

Finally, under truthful bidding, the expected payoff of a voter with value v limits to

$$\pi(v) \;=\; \frac{1}{2\sqrt{n\pi}\sigma}(v^2 + \sigma^2)$$

while under voting the expected payoff limits to

$$\pi(v) \;=\; \frac{1}{\sqrt{n\pi}}\,|v|$$

The difference is proportional to $(v - \text{sign}(v)\sigma)^2$, which is strictly positive for all $|v| \neq \sigma$. $\square$

The full efficiency obtained under bidding contrasts with the inefficiencies that arise under majority voting (Proposition 1). In addition, the gain in welfare that results with bidding does not simply accrue to those that care more about the outcome. The redistribution that is inherent to the bidding mechanism ensures that voters of all types are better off.

We next show that bidding works not only in the limit but also in smaller groups. For example, for an electorate of size two the equilibrium bids can be computed exactly, e.g.

$$B(v) \;=\; \text{sign}(v)\,\sqrt{2(1 - \exp(-\tfrac{1}{2}v^2))}$$

when values are (standard) normally distributed.[8] Since the equilibrium bidding function is

[8] Assuming the other voter uses the proposed bidding function, a voter with value v who makes a bid b has expected payoffs

$$\pi(v, b) \;=\; v\big(2\Phi(B^{-1}(b)) - 1\big) - b^2/\sqrt{2\pi} + (\sqrt{2} - 1)/\sqrt{\pi}$$

where the final term on the right side is the expected rebate and $B^{-1}(b) = \text{sign}(b)\,\sqrt{\log(1/(1 - \frac{1}{2}b^2)^2)}$ is the inverse bidding function. Optimizing with respect to the bid yields $b = B(v)$.

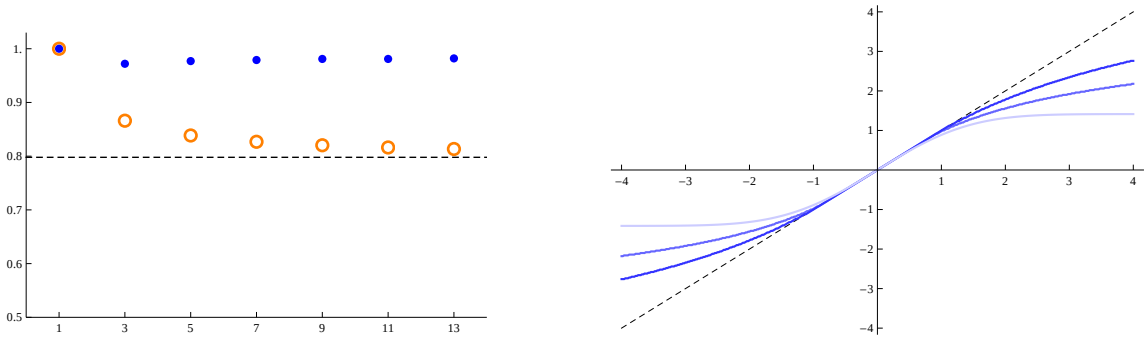

Figure 1. The left panel shows efficiency levels under voting (open orange circles) and bidding (filled blue circles) for various electorate sizes with normally distributed values. Voting is efficient with one voter but efficiency falls quickly for larger electorate sizes and limits to $\sqrt{2/\pi}$ as shown by the dashed line. In contrast, bidding is close to 100% efficient for all electorate sizes. The right panel shows equilibrium bids for electorates of size two (light blue), five (medium blue), and eleven (dark blue).

increasing, the bidding mechanism is fully efficient. Moreover, voters of *all* types are better off under bidding compared to voting when the electorate size is two.[9]

While we have no analytic characterization of the bidding functions for other finite electorate sizes, it is straightforward to determine them numerically.[10] The right panel in Figure 1 shows the equilibrium bidding functions for electorates of size two (light blue curve), five (medium blue curve), and eleven (dark blue curve). For the latter case, bids are close to the 45-degree line for values within plus or minus two standard deviations. In other words, for 95% of the values, bidding is close to truthful, resulting in efficient outcomes. The filled blue circles in the left panel of Figure 1 show the predicted efficiency levels under bidding by electorate size. The bidding mechanism is fully efficient with one (or two) voters, there is only a slight drop in efficiency for intermediate electorate sizes, and full efficiency obtains in the limit.

There are two main points to be taken away from our analysis. One is that bidding improves on voting because it allows for an expression of voter intensity. For intermediate electorate sizes, bids do not fully reveal values and, hence, bidding is not fully efficient, but it does improve substantially on voting. Second, the predicted efficiency levels for bidding and voting are close to their limit values for moderate electorate sizes, say, an electorate of size eleven (see the left panel of Figure 1). This is a useful feature when designing an experiment.

[9]The equilibrium payoff under voting, $\pi(v) = \frac{1}{2}|v|$, is strictly less than the equilibrium payoff under bidding, $\pi(v) = v(2\Phi(v) - 1) + 2\phi(v) - 1/\sqrt{\pi}$, for all v.

[10]We consider values between -4 and $+4$ (recall that the standard deviation is 1 so this captures 99.994% of all values) on a discrete grid with 201 possible values: $v \in \{-4.00, -3.98, \ldots, 0, \ldots, 3.98, 4.00\}$. The numerical process starts with some initial bidding function, say $B(v) = v$, and then computes the best response against this initial bidding function for each value type. This yields a new bidding function against which the best response is computed, etc., until there is no more change in the computed bids.

4. Experimental Design, Procedures, and Hypotheses

Our interest is in testing bidding and voting when the limit results of Propositions 1 and 2 apply. An obvious difficulty is that we cannot study very large electorates because of physical constraints (lab size) as well as financial constraints. Fortunately, moderately small electorate sizes suffice to test the alternative mechanisms for two reasons. First, recall that efficiency under voting falls with the number of voters. So testing majority voting with a smaller group puts an *upper* bound on how well voting does with large electorates. Furthermore, a finite electorate size can be chosen such that the predicted efficiency level only marginally exceeds its limit value for an infinite electorate size. Second, it is possible to design environments with a finite number of bidders for which the logic behind limit results of Proposition 2 applies and bidding is fully efficient making everyone better off.

4.1. Design

We used a two-point value distribution (instead of normally distributed values) since it (i) is much easier to understand for subjects, (ii) results in more separation between bidding and voting, (iii) makes redistribution easier to analyze, and (iv) allows us to estimate a quantal response equilibrium model using the individual bid data. In particular, voters are equally likely to prefer one of two alternatives, labeled A and B, and their values for the preferred alternative are low, $v_L = 10$, with probability $p = 0.75$ and high, $v_H = 70$, with complementary probability. We refer to the low-value and high-value voters as "moderates" and "extremists" respectively. For this value distribution, the open orange circles in the left panel of Figure 2 show the predicted efficiency levels under majority voting for various (odd) electorate sizes. (This figure parallels Figure 1, which pertains to normally distributed values.) Voting is 100% efficient when there is only a single voter but efficiency falls as the electorate size increases. For the experiments, we chose a group size of eleven for which predicted efficiency is 71.1%, which is within two percentage points of the limit efficiency (69.3%) indicated by the dashed line in Figure 2. Since we are interested in how electorate size affects performance, we also considered groups of size three for which the predicted efficiency is at least somewhat different (78.0%).[11]

A second departure from the theoretical setup discussed above is that we restricted possible bids to the set $\{0, 1, \ldots, 10\}$ to avoid the possibility of erroneously high bids that could lead to bankruptcies. To ensure that the equilibrium bids are within this set we scaled up the cost function by a factor of 10 so that, in large electorates, the optimal bid is not equal to value, as

[11] Predicted efficiencies for other group sizes are rather close to that for $n = 11$ (or $n = \infty$), which would make any treatment effect unlikely. The exception is a group of size 1 but this would be too trivial an experiment.

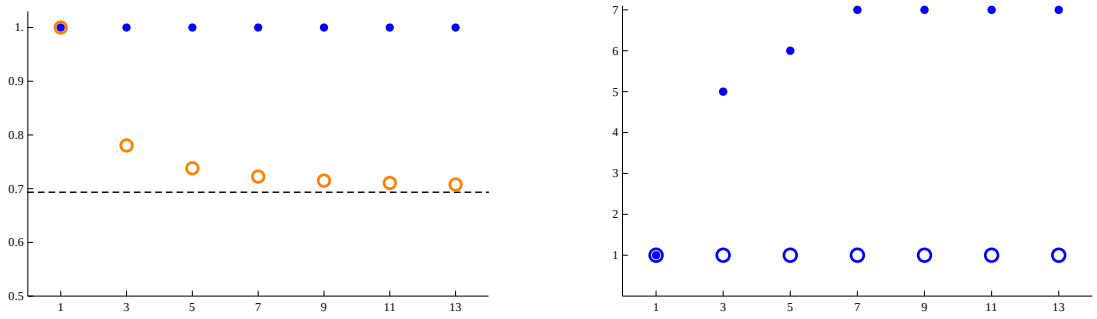

Figure 2. The left panel shows efficiency under voting (open orange circles) and bidding (filled blue circles) by electorate size for the value distribution used in the experiment. Voting is efficient with a single voter but efficiency falls quickly for larger electorate sizes. The experiments use electorates of size three for which efficiency is 78.0% and eleven for which efficiency is 71.1%, which is close to the limit value of 69.3% (as indicated by the dashed line). Bidding is predicted to be 100% efficient for all electorate sizes. The right panel shows the equilibrium bids by electorate size for moderates (open blue circles) and extremists (filled blue circles).

in Proposition 2, but rather to one-tenth of the value (1 for a moderate and 7 for an extremist). This normalization results in a cost function $C(b) = b^2/3$ for an electorate of size eleven.[12] As we explain in the Hypotheses section below, restricting the highest possible bid to 10 has no consequence for the theoretical predictions (but precludes very costly mistakes).

Increasing the group size from three to eleven roughly halves the chance of being pivotal. For this reason the costs of buying votes in the bidding mechanism should be twice as high in groups of three compared to groups of eleven. Hence, we implemented a cost function $C(b) = 2b^2/3$ for groups of size three. Of course, turnout in real elections (and previous voting experiments) suggests that voters' perceptions of being pivotal may be biased. For instance, there is ample evidence that subjects tend to overestimate small probabilities (as in models of probability weighting). To address this possibility, we also ran a treatment with groups of size eleven and high costs $C(b) = 2b^2/3$, thus ignoring the (theoretically predicted) drop in the chance of being pivotal when increasing the group size from three to eleven.

The three treatments are labeled "3H," "11H," and "11L," where the number indicates group size and the letter indicates whether costs are low or high (see Table 1). All three treatments used two parts of twenty periods each, for a total of forty periods. In each of these forty periods, subjects were randomly assigned a preferred alternative, A or B, and a private value for that alternative using the two-point value distribution described above.[13] In the first

[12] The exact cost function is $C(b) = 10/(2\sqrt{5\pi}\sigma)b^2$ where $\sigma = \sqrt{\frac{3}{4}(1)^2 + \frac{1}{4}(7)^2}$, which yields $C(b) = 0.349b^2$. With a discrete bid space, equilibrium bids are unaffected by small changes in the cost function and we decided to use the cost function $C(b) = b^2/3$ to simplify the instructions.

[13] Subjects knew their own preferred outcome and their own value, but they knew only the distribution of others' preferred outcomes and values. We used the same value draws in treatments 11H and 11L.

Treatment Name	Values $\{v_L, v_H\}$	Cost Function	Group Size	Number of Groups	Part 1 Bid and Vote	Part 2 Bid or Vote	Optimal Bids $\{b_L^*, b_H^*\}$
11L	$\{10, 70\}$	$b^2/3$	11	10	20 periods	20 periods	$\{1, 7\}$
3H	$\{10, 70\}$	$2b^2/3$	3	10	20 periods	20 periods	$\{1, 5\}$
11H	$\{10, 70\}$	$2b^2/3$	11	10	20 periods	20 periods	$\{1, 5\}$

Table 1. The experimental design consists of three treatments that vary group size and the cost of buying votes. In the first 20 periods, subjects make choices under both bidding and voting. Then they collective decide by majority vote which mechanism applies in the final 20 periods. Subjects' values are low, $v_L = 10$, with probability $p = 3/4$ and high, $v_H = 70$, with complementary probability.

twenty periods, subjects made choices under both the bidding and voting mechanism (subjects' screens were split in two parts and for each subject it was randomly determined whether voting or bidding appeared on the left part of the screen for the entire first part). Then they decided by majority vote which mechanism would apply in the final twenty periods.[14] For completeness we summarize the rules of both mechanisms.

Voting Mechanism: subjects chose to vote for either A or B. The outcome with the most votes was selected as the group outcome (ties are not possible with odd-sized groups). A subject received her value if the group outcome matched her preferred outcome, otherwise she received the negative of her value.

Bidding Mechanism: subjects chose to bid for either A or B and they chose a bid amount. Integer amounts from 0 to 10 were allowed. The outcome with the largest total bid amount was selected as the group outcome (in the event of a tie, each outcome had an equal chance of being selected). A subject received her value if the group outcome matched her preferred outcome, otherwise she received the negative of her value. Bidding is costly and irreversible. Irrespective of the group outcome, a subject who bid b paid $C(b) = 2b^2/3$ in the high cost treatments (3H and 11H) and $C(b) = b^2/3$ in the low cost treatment (11L). In addition, each subject received a rebate equal to the average of others' payments.

4.2. Procedures

We used ten groups per treatment for a total of 250 subjects who were recruited at the University of Zürich and the neighboring ETH. The experiment was computerized using zTree (Fischbacher, 2007). Each session started with two unpaid practice periods followed by 20 paid periods in part 1 and another 20 paid periods in part 2. Subjects were not informed about

[14]This decision could not be made by bidding since subjects did not know their values for the second part. To not favor voting or bidding we used a simple majority rule rather than a super-majority or unanimity rule.

part 2 until they had completed part 1. In the first part, subjects made decisions under both bidding and voting and received the sum of their earnings from both mechanisms. After finishing the first part, they collectively decided by majority vote which mechanism would apply in the second part.[15] In this second part, subjects only made decisions for, and received earnings from, the selected mechanism. Sessions lasted about 60-75 minutes, including instructions and payment. The conversion rate from experimental points to Swiss Francs was 30:1 for groups of size three and 15:1 for groups of size eleven. Average earnings were 43.4 Swiss Francs including a 15 Franc show-up fee.

4.3. Hypotheses

We have not yet discussed in detail the predictions for the bidding mechanism. The left panel of Figure 2 shows full efficiency under bidding for all electorate sizes. This result obtains when bidding costs are scaled such that they roughly fall with the square root of the electorate size.[16] One of the treatments (11H) purposefully deviates from this scaling law to explore the (behavioral) assumption that subjects may overestimate the probability of being pivotal. As a result, the equilibrium bids for this treatment are not equal to one-tenth of the value but instead are $b_L^* = 1$ for a moderate voter and $b_H^* = 5$ for an extremist voter. These bids are also the equilibrium bids for treatment 3H while for treatment 11L the equilibrium bids are $b_L^* = 1$ and $b_H^* = 7$, see the final column in Table 1.

It is easy to see how the equilibrium bids for treatment 3H guarantee full efficiency. Whenever two moderates are on one side but an extremist voter is on the other side, bidding yields the efficient outcome, since $5 > 1+1$, while voting is inefficient, since $1 < 1+1$. Other scenarios can be checked similarly. In a sense, it is relatively easy for a bidding mechanism to improve on voting in small groups of size three.

The fact that bidding is fully efficient in treatment 11L is for an entirely different reason, i.e. because the influence of any bid is small enough such $2G(b) - 1$ is approximately linear over the range of feasible bids. This claim is illustrated in Figure 3. The left panel shows $2G(b) - 1$ for bids between -70 and 70, where negative (positive) numbers correspond to bids for alternative A (B). Since there are ten others who potentially could each bid 7 for either alternative, $2G(b) - 1$ is equal to 1 (-1) only when a bid of 70 (-70) or more (less) is made. The right panel zooms in and shows $2G(b) - 1$ for the range of feasible bids, were it is well

[15] After they cast their vote, but prior to knowing the majority choice, subjects had an opportunity to briefly state a reason for their choice.

[16] In particular, the cost function used to generate Figure 2 is $C(b) = c_{2n+1}b^2$ for an electorate of size $2n+1$, where $c_3 = \frac{2}{3}$, $c_5 = \frac{1}{2}$, $c_7 = \frac{2}{5}$, $c_9 = \frac{7}{20}$, $c_{11} = \frac{1}{3}$, and $c_{13} = \frac{3}{10}$. Other parameter values that are close and roughly follow the same scaling result in the same optimal bids and full efficiency.

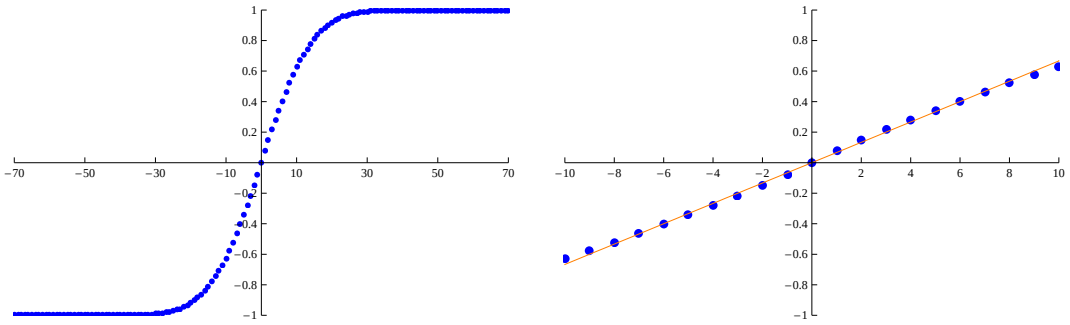

Figure 3. The left panel shows $2G(b) - 1$ for bids in the range $-70 \leq b \leq 70$ in treatment 11L assuming others bid $b_L^* = 1$ when they are moderate and $b_H^* = 7$ when they are extremist. The right panel zooms in on the set of feasible bids $-10 \leq b \leq 10$ for which $2G(b) - 1$ is well approximated by the linear function $b/15$ as shown by the orange straight line.

approximated by the orange line that represents the linear function $b/15$. Hence, in treatment 11L, a voter's expected payoff can be written as

$$\pi(v, b) \;=\; v\big(\frac{b}{15}\big) - \frac{1}{3}b^2$$

and the first-order condition for profit maximization yields $B(v) = v/10$.

From a theoretical viewpoint, restricting the set of feasible bids to $\{0, 1, \ldots, 10\}$ can be done without loss of generality. Suppose, for instance, that the voter's value is positive and consider any positive bid. The left panel of Figure 3 shows that the win probability is concave for any positive bid and, hence, so is the voter's profit function. As a result, there is a unique maximizer, $b_L^* = 1$ or $b_H^* = 7$, which is contained within the feasible range. Bids outside this range are never optimal and, hence, can be discarded without affecting theoretical predictions. Behaviorally, setting a maximum bid has the advantage of ruling out very costly mistakes.

To summarize, while full efficiency in treatment 3H is due to a "small numbers effect," treatment 11L effectively implements a setting for which the limit results of Proposition 2 apply. Truthful bidding is not in equilibrium in treatment 11H, but the optimal bids of 1 and 5 are such that predicted efficiency is still very high (99.7%). We can now state our hypotheses.

Hypothesis 1. In all treatments, bidding is more efficient than voting.

Hypothesis 2. In all treatments, moderates and extremists are better off under bidding.

Hypothesis 3. In all treatments, bidding is more often selected for part 2 of the experiment.

Hypothesis 4. Moderates bid the same in all treatments. Extremists bid the same in treatments 3H and 11H and more in treatment 11L. Hence, the bid ratio (extremist's bid divided

by moderate's bid) is the same in treatments 3H and 11H and higher in treatment 11L.

These hypotheses are based on Bayes-Nash equilibrium predictions. We also state one behavioral hypothesis that motivates treatment 11H and provides an alternative to Hypothesis 4. Suppose subjects ignore the effect of group size on the chance of being pivotal and that instead their behavior is driven by the total cost they are willing to incur (as a fraction of their value).

Hypothesis 4-Alt. Moderates bid the same in treatments 3H and 11H and more in treatment 11L. The same is true for extremists. The bid ratio is the same in all treatments.

5. Experimental Results

Our two-part design has two important advantages: it allows us to study the preference of groups over institutions (voting or bidding) and it provides subjects with the opportunity to learn in the first part making equilibrium behavior in the second part more likely. Subjects indeed learned over time. The percentage of votes (bids) against the preferred alternative dropped from 1.04% to 0.56% (1.88% to 1.16%) from the first 10 periods to the second 10 periods in part 1. These mistakes were made by very few subjects. Only 2.4% subjects voted against their preferred alternative more than once. 90% subjects never made a single mistake.

Figure 4 reports the observed vs. predicted efficiency calculated at the group level, averaged across the second 10 periods in part 1. The horizontal axis is the predicted efficiency under voting, given the actual value draws in the experiment. The predicted bidding efficiency is always 100%. The size of the marker corresponds to the number of observations in each data point. Typically there is one data point for each group. A larger marker has two groups and the largest has three groups in one data point. Points on the 45 degree line represent groups where the observed efficiency equals the predicted efficiency under voting, and points above (below) the line are groups performed better (worse) than prediction. In all treatments, the observed voting efficiencies align nicely with the theoretical predictions (orange circle). The observed efficiencies under bidding (blue dot) are higher than under voting most of the time and the differences are significant at 1%. [17] Figure 5 shows the aggregate data at the treatment level.

Given that bidding is significantly more efficient than voting, did subjects collectively choose to apply bidding in part 2? We analyze the group choice for either voting or bidding that

[17] A Wilcoxon signed-rank test that compares the efficiency levels ($p = 0.0095$ for all treatments, $p = 0.051$ for 3H and $p = 0.093$ for 11H). All reported tests are two-sided and based on group level data (unless otherwise stated).

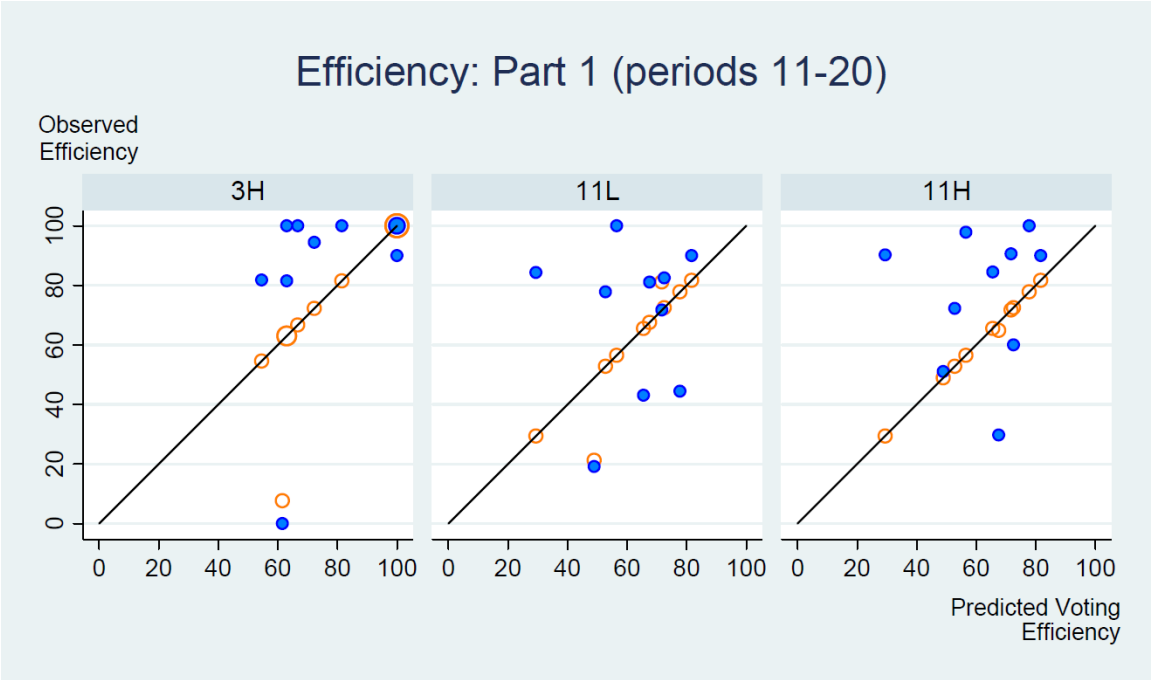

Figure 4. Observed vs. predicted efficiency at the group level for periods 11-20 by treatment. The orange circle is the observed voting efficiency and blue dot is the observed bidding efficiency. The size of the markers corresponds to the number of observations. The horizontal axis depicts the predicted voting efficiency while the predicted bidding efficiency is always 100%.

immediately followed the first part. A total of 30 groups participated in the experiments (see Table 1) and 20 of them made more money, as a group, under bidding ($p = 0.098$) in part 1. Furthermore, 22 groups preferred bidding while only 8 groups preferred voting for the second part ($p = 0.016$). We thus find strong support for Hypothesis 3, especially with large electorates.

> **Result 1.** The bidding institution is selected more frequently for part 2 of the experiment. In particular, with electorates of size eleven, 18 out of 20 groups chose bidding, and with an electorate of size three, 4 out of 10 groups chose bidding.

Besides the higher earnings accrued under bidding in the first part, subjects mentioned that they preferred bidding because it "provided them with more influence" while with voting "chance decides." In addition, subjects positively commented on the redistribution that occurs under bidding in the form of rebates. Subjects criticized the free-riding that is possible with bidding (e.g. by submitting a zero bid), which has a bigger impact in small electorates. In 3H, extremists never free ride in part 1 in the four groups that later chose the bidding mechanism in part 2. In the other 6 groups, 7.4% bids submitted by the extremists were zero and these groups selected the voting mechanism for part 2. Also in groups of three, 24.50% of the time the same preferred

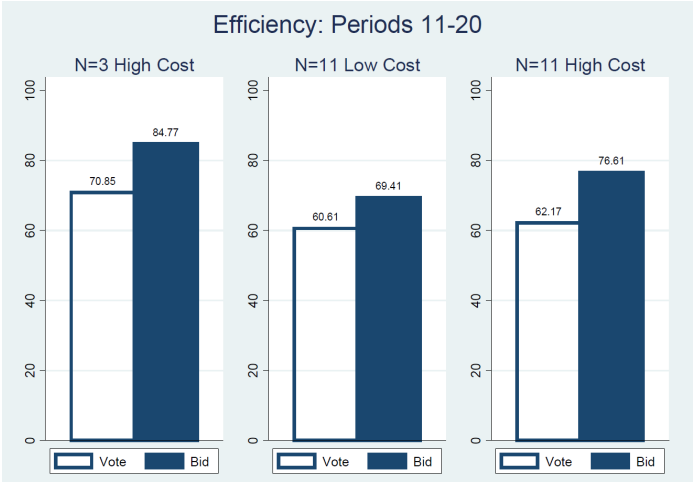

Figure 5. Observed efficiencies under voting (white bars) and under bidding (blue bars) averaged across periods and groups, by treatment. The predicted efficiencies under voting are 76.24%, 62.44% and 62.44% for treatment 3H, 11L and 11H respectively.

alternative was assigned to everyone in the group. This only occurred 0.5% in groups of eleven. These differences may explain why bidding was less frequently selected with groups of size three in part 2.

Comparing to part 1, observed efficiencies under bidding were higher in the second part for all three treatments.[18] In the subsections below, we use the data from the second part to compare bidding and voting in terms of efficiency (Section 5.1) and the distribution of earnings (Section 5.2). In Section 5.3 we analyze the bid data in more detail.

5.1. Efficiency

Figure 6 plots the observed vs. predicted efficiency at the group level in part 2. For electorates of size eleven, we have only two groups that selected voting (orange circle) and eighteen groups that selected bidding (blue dots), which makes it difficult to compare bidding and voting using non-parametric tests. For this reason, we compare the observed efficiencies under bidding with the *predicted* efficiencies under voting assuming perfectly optimal behavior in the voting mechanism. In other words, we give voting its best chance. The assumption of perfect behavior under voting is reasonable from an empirical viewpoint since only 0.7% of all votes were against the alternative in part 2. Also, for the six groups that selected voting with a group size of three, the observed efficiency is exactly as predicted (perfectly aligned on the 45 degree line), i.e. if there were any mistakes they were not pivotal. For electorates of size three, it thus does not

[18] A Wilcoxon signed-rank test that compares the efficiency levels in parts 1 and 2 for the 22 groups that selected bidding yields a significant difference ($p = 0.088$).

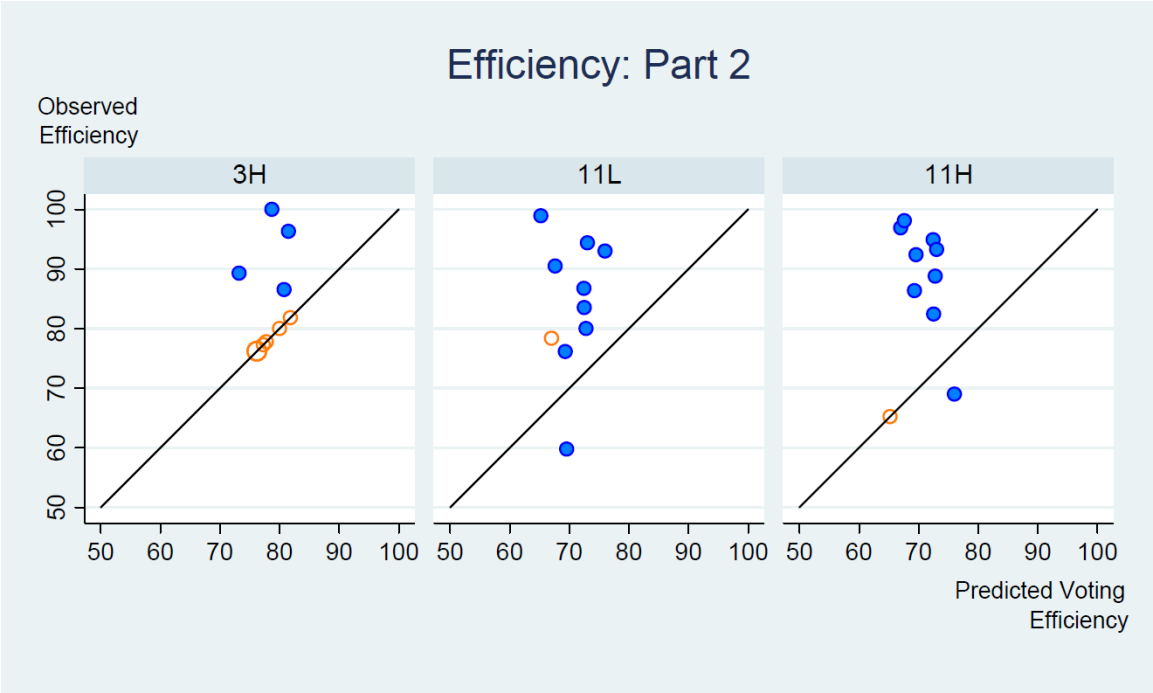

Figure 6. Observed vs. predicted efficiency at the group level given the endogenous group choice of applying voting or bidding mechanism in part 2. by treatment. The orange circle is the observed voting efficiency and blue dot is the observed bidding efficiency. The size of the markers corresponds to the number of observations. The horizontal axis depicts the predicted voting efficiency while the predicted bidding efficiency is always 100%.

matter whether we compare efficiency under bidding with the predicted or observed voting efficiencies. Overall, only two groups had lower efficiencies under bidding than the predicted voting efficiency (below the line).

Figure 7 shows the observed efficiencies under bidding (blue bars) and the associated predicted efficiencies under voting (white bars) by treatment. The numbers represent averages over all twenty periods and across groups that selected bidding for the second part. Predicted efficiencies under voting are 78.6% in 3H, 71.0% in 11L, and 71.2% in 11H.[19] Observed efficiencies under bidding are much higher: 93.0% in 3H, 84.8% in 11L, and 89.1% in 11H.[20] We thus find strong support for Hypothesis 1.

Result 2. Bidding is significantly and substantially more efficient than voting.

Support. For electorates of size three, we can compare the observed efficiencies for the six groups that selected voting with the observed efficiencies for the four groups that chose bid-

[19] Close to the ex ante expected efficiencies of 78.0% for 3H and 71.1% for 11H and 11L.

[20] A Kruskal-Wallis test shows no difference among the bidding efficiencies for the three treatments ($p = 0.39$).

Figure 7. Predicted efficiencies under voting (white bars) and observed efficiencies under bidding (blue bars) for the draws used in the experiment and averaged across periods and groups. The left panel is based on four groups that selected bidding in 3H, the middle panel is based on nine groups that selected bidding in 11L, and the right panel is based on nine groups that selected bidding in 11H.

ding. A Mann-Whitney rank-sum test shows that efficiency under bidding is significantly higher ($p = 0.01$).[21] Alternatively, we can compare the four observed bidding efficiencies with the predicted efficiencies under (perfect) voting. A Wilcoxon matched-pair signed-rank test confirms that bidding is significantly more efficient ($p = 0.069$).[22] For electorates of size eleven, we compare the observed efficiencies for the nine groups that selected bidding with the predicted efficiencies under (perfect) voting. A Wilcoxon matched-pair signed-rank test shows that bidding is significantly more efficient in 11H ($p = 0.01$) and 11L ($p = 0.02$).[23] □

To glean some intuition for why voting performs worse, consider those cases where the mean and median value are misaligned. For example, suppose there is a majority of moderate voters that prefer A and a few extremists that prefer B. In this case, outcome A will be selected under voting even though it might be optimal to implement outcome B. Figure 8 shows observed efficiency levels under bidding when the mean and median value are on opposite sides

[21] The six observed efficiency levels under voting are (77.8%, 76.2%, 80.0%, 81.8%, 76.2%, 77.3%) and the four observed efficiency levels under bidding are (89.3%, 100.0%, 96.3%, 86.5%).

[22] The predicted efficiency levels under voting are (73.2%, 78.7%, 81.5%, 80.8%).

[23] The nine observed efficiency levels under bidding in 11H are (88.8%, 69.0%, 96.9%, 94.9%, 82.4%, 98.1%, 86.4%, 92.4%, 93.3%) and for 11L they are (80.0%, 93.0%, 86.7%, 83.5%, 98.9%, 90.5%, 76.1%, 94.4%, 59.8%). We used the same value draws in treatments 11H and 11L so the predicted vote efficiencies are the same. However, the group that did not select bidding in 11H used different values than the one that did not select bidding in 11L. Hence, we list the predicted vote efficiencies for all ten groups (72.8%, 76.0%, 67.0%, 72.5%, 72.5%, 65.2%, 67.6%, 69.3%, 69.6%, 73.0%). In 11H, group 6 did not select bidding, and in 11L, group 3 did not select bidding.

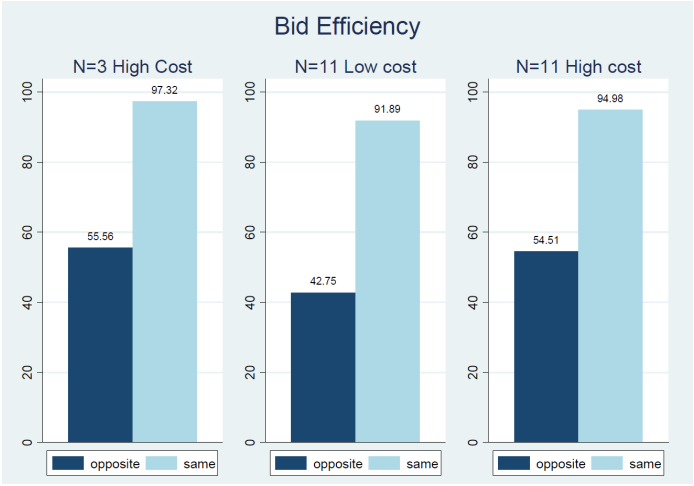

Figure 8. Observed efficiency levels under bidding when the mean and median value are on opposite sides (dark bars) or on the same side (light bars). Under voting, predicted efficiency is -100% when the mean and median value are on opposite sides and 100% when they are on the same side.

(dark bars) or on the same side (light bars). Bidding efficiency is near perfect ($> 90\%$) for the "easy" case when the mean and median value are on the same side but less so when they are on opposite sides (55% with high costs and 43% with low costs). This performance, though, is much better than predicted under voting even if we assume no mistakes are made: when the mean and median value are on the same side voting is predicted to be 100% efficient, but voting is -100% efficient when they are on opposite sides.[24]

5.2. Who Benefits?

Figure 9 shows the predicted payoffs, in Francs, of moderates and extremists under voting (white bars) and bidding (blue bars). The predicted payoffs are based on the value draws for those groups that selected bidding in part 2. Comparing the white and dark blue bars illustrates the effects of redistribution under bidding and confirms that voters of all types are predicted to be better off (Proposition 2). The green bars in Figure 9 correspond to the observed payoffs under bidding. We find partial support for Hypothesis 2.

> **Result 3.** Moderates are better off under bidding in all treatments. Extremists are better off under bidding in groups of size three but not in groups of size eleven.

Support. For treatment 3H, comparing observed payoffs in the six groups that chose voting to the observed payoffs in the four groups that selected bidding shows that moderates and

[24]Of course, the event where the mean and median value are on the same side is more likely, which is why the overall predicted efficiency for the voting mechanism is positive.

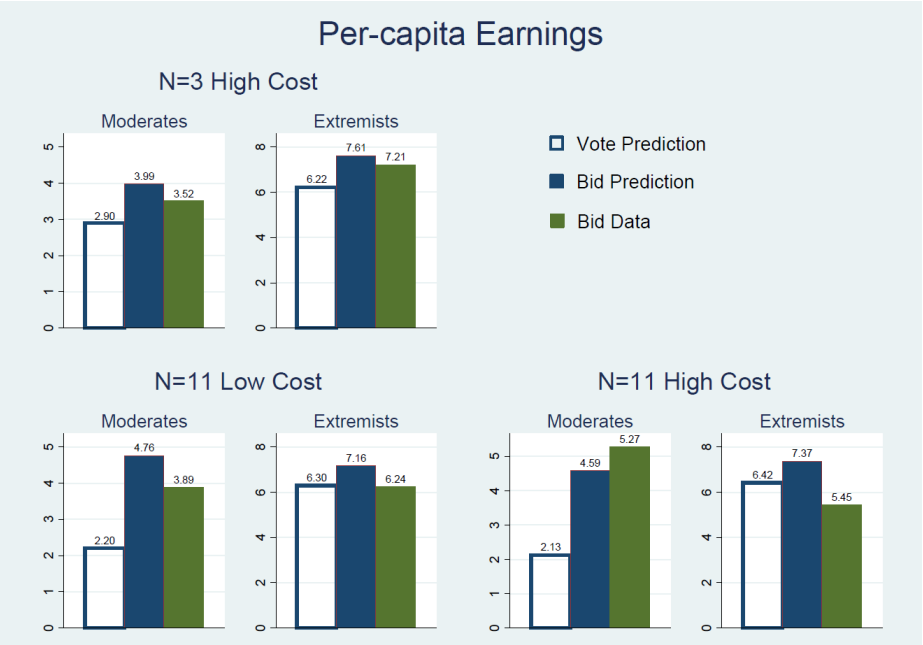

Figure 9. Predicted per-capita earnings for moderates and extremists under voting (white bars) and bidding (blue bars) as well as observed per-capita earnings under bidding (green bars). Earnings are in Swiss Francs and are based on groups that selected bidding in part 2.

extremists are better off.[25]. Additional support can be obtained by comparing observed extremists' payoffs under bidding to their predicted payoffs under (perfect) voting for the four groups that selected bidding in the second part, which shows that they are better off under bidding ($p = 0.068$). Comparing observed payoffs under bidding to predicted payoffs under voting shows that moderates are better off in treatments 11L ($p = 0.008$) and 11H ($p = 0.008$), but extremists are equally well off in treatment 11L ($p = 0.95$) and worse off in treatment 11H ($p = 0.038$). □

The theoretical prediction that moderates benefit more than extremists in all three treatments is confirmed. The comparative statics predictions of Bayes-Nash theory are thus corroborated by the data, but less so its point predictions.

5.3. Bidding Behavior

Figure 10 shows the observed average bids by treatment for the extremists (lower right panel) and moderates (lower left panel) as well their ratio (top panel). The data favor behavioral Hypothesis 4-Alt, which assumes that bidding behavior is mainly driven by costs, over Hypothesis

[25]Mann-Whitney test, one-sided $p = 0.068$ for moderates and $p = 0.067$ for extremists; Robust Rank Order test, one-sided $p = 0.038$ for moderates and $p = 0.043$ for extremists

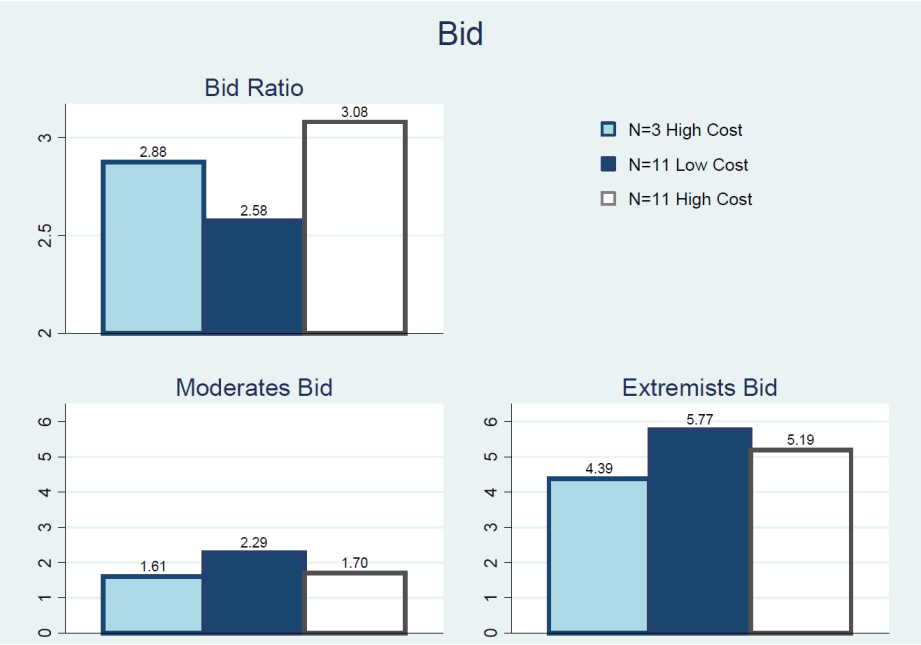

Figure 10. The top panel shows the observed bid ratio (extremist bid/moderate bid) by treatment. The bottom panels show the average bids of moderates (left panel) and extremists (right panel) by treatment. Averages are across groups and periods.

4, which is based on Bayes-Nash theory.

> **Result 4.** Moderates bid the same in treatments 3H and 11H and they bid higher in treatment 11L. The same is true for extremists. The bid ratio is the same in treatments 3H and 11H and lower in treatment 11L.

Support. Comparing treatments 3H and 11H shows no differences between the average bids submitted by either moderates (Mann-Whitney test, $p = 0.64$) or extremists ($p = 0.22$).[26] Moderates' bids in treatment 11L are significantly higher than those in treatment 3H ($p = 0.02$) and treatment 11H ($p = 0.05$).[27] Likewise, extremists' bids in treatment 11L are significantly higher than those in treatment 3H ($p = 0.04$) and treatment 11H ($p < 0.001$). Finally, the bid ratio is the same in treatments 3H and 11H ($p = 0.28$), the same in treatments 3H and 11L ($p = 1.00$), and lower in treatment 11L than in treatment 11H ($p = 0.058$). □

The observed bid ratios differ substantially form the theoretically predicted ratios: 5 in the

[26]The average bids in treatment 3H are (1.53, 1.62, 1.13, 2.14) for moderates and (4.53, 5.93, 3.75, 3.33) for extremists while in treatment 11H the averages are (1.99, 1.57, 1.87, 1.81, 1.85, 1.40, 1.52, 1.76, 1.56) for moderates and (4.93, 4.15, 5.22, 5.84, 5.45, 5.55, 5.19, 5.22, 5.16) for extremists.

[27]The average bids in treatment 11L are (2.51, 2.16, 2.56, 2.35, 1.98, 1.95, 2.42, 2.76, 1.92) for moderates and (4.23, 5.65, 6.21, 6.57, 6.29, 5.67, 6.20, 4.69, 6.38) for extremists.

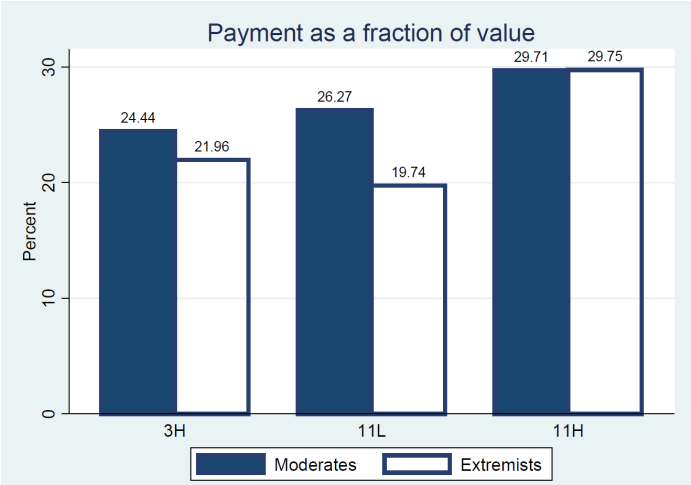

Figure 11. This figure plots the average cost moderates and extremists incurred to pay for their bids as a fraction of their value in each treatment. The fraction is calculated as $\frac{1}{3}b^2/v$ in treatment 11L and $\frac{2}{3}b^2/v$ in treatments 3H and 11H. Averages are calculated first over 20 periods per subject and then across 10 groups by treatment.

high-cost treatments and 7 in the low-cost treatment.[28] The reason is that moderates bid significantly more than 1 in all treatments while extremists bid less than 5 in 3H, bid 5 in treatment 11H, and they less than 7 in treatment 11L.

The bid ratio is an important determinant of efficiency. For example, if the bid ratio is 1, bidding is no more efficient than voting. But even though observed bid ratios are less extreme than predicted, they are substantial enough to result in highly efficient outcomes. For instance, in treatment 3H, if all moderates bid 1.61 and all extremists bid 4.39 (which are the observed averages) then predicted efficiency would be 100%. Similarly, for treatments 11L and 11H, the assumption of homogenous behavior results in predicted efficiencies of 96% or more when using the observed average bids.

To better understand what drives the bidding behavior, we examine the amount moderates and extremists paid for their bids as a fraction of their value. The fraction is calculated as $\frac{1}{3}b^2/v$ in treatment 11L and $\frac{2}{3}b^2/v$ in treatments 3H and 11H on individual level. Figure 11 reports the group average of these fractions by treatment. In treatments 3H and 11H, moderates and extremists on average spent roughly the same amount proportional to their value to bid for their preferred alternative.[29] In treatment 11L, moderates incurred significantly higher cost

[28]Not surprisingly, these differences are statistically significant. A sign test yields $p = 0.068$ in treatment 3H, $p = 0.008$ in treatment 11H, and $p = 0.008$ in treatment 11L.

[29]Mann Whitney tests comparing moderates with extremists report no significant difference in 3H (p = 1.00) and in 11H (p = 0.89).

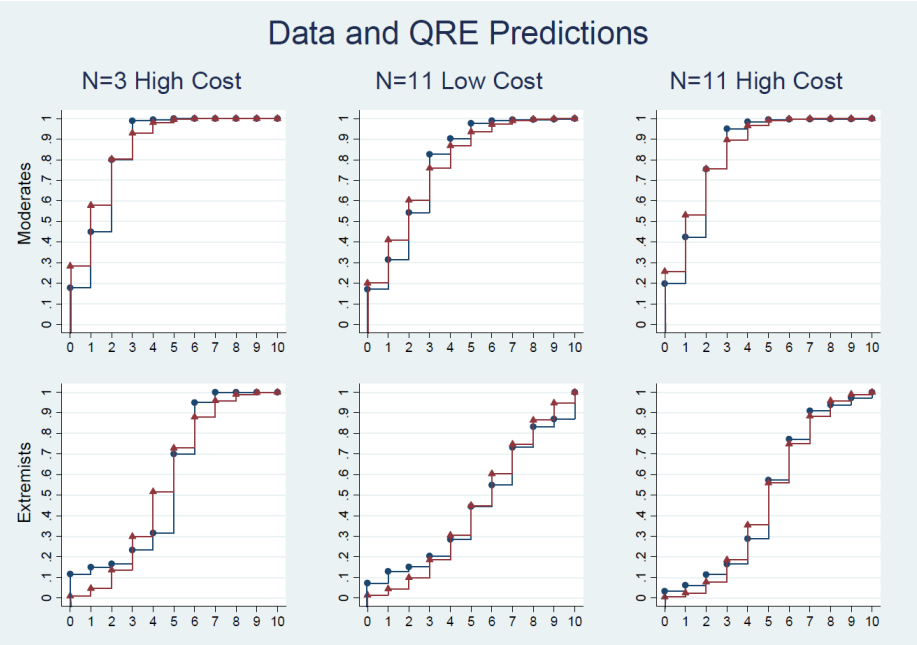

Figure 12. This figure compares the actual cumulative bid distributions (blue lines) with the predicted quantal response distributions (red lines) for moderates (top panel) and extremists (bottom panel) by treatment. The Nash equilibrium bids for moderates and extremists in high cost treatments are 1 and 5 and in low cost treatment are 1 and 7 respectively.

relative to their value than extremists (p = 0.038). This again indicates that bidders mainly consider the cost of bidding instead of the chance of being pivotal in forming their bidding strategy. Moreover, pairwise comparison across three treatments indicates no difference for the relative amounts paid by moderates. For extremists, the only difference was observed when comparing 11H with 11L.[30]

Another reason that may explain the small observed efficiency loss under bidding is that there is substantial heterogeneity, or "noise," in decision making. To this end we estimate a quantal response model where we allow subjects to overestimate the chance of being pivotal. Figure 12 compares the actual cumulative bid distributions with the quantal response distributions. Using a KolmogorovSmirnov test, the null hypothesis that the data are generated by the theoretical distribution, cannot be rejected at the 5% level.

[30]Mann Whitney tests comparing 3H with 11L, 3H with 11H and 11L and 11H report p = 0.44, p = 0.22, p =0.31 for moderates, respectively and p = 0.88, p = 0.12, p = 0.002 for extremists.

6. Conclusion

In this paper we have demonstrated the inefficiency of elections, i.e. their inability to achieve socially optimal outcomes in situations without a clear ex ante winner (Proposition 1). We have also shown that a simple bidding mechanism yields higher overall welfare and better outcomes for *all* voters. These bidding mechanisms therefore represent practical implementations of Wicksell's idea of "unanimity and voluntary consent in taxation," i.e. the insight that if a proposal yields a common outcome with a positive net value to society then there exists a compensation scheme by which those that gain from the proposal can compensate those that lose and the proposal can be approved unanimously.[31]

A common objection to bidding mechanisms vis-à-vis a simple one man - one vote mechanism is that those that care a lot more may try to"buy the election" at the expense of more moderate voters. One reason to doubt this argument is that buying the election is typically very costly and the payments required imply substantial redistributions that benefit all. Indeed, Proposition 2 shows that *all* voters may benefit when voting is replaced by a bidding mechanism. And the laboratory data show that moderate voters benefit most, especially under the simpler bidding mechanism. Another reason is that in many prominent elections, those that care a lot *are* trying to buy the election via excessive campaign contributions. And, unlike the budget balanced mechanism studied in this paper where monetary bids are redistributed to the general public, these campaign contributions often represent wasteful rent-seeking of little or no social value (e.g. buying TV time for negative and uninformative advertisements about an opponent or proposition). In this sense, our results about the inefficiency of elections would be further strengthened in a model that incorporates the possibility of vote buying through rent-seeking activities.

Our theoretical predictions about the inefficiency of elections are corroborated by results from controlled experiments. The laboratory data provide empirical proof that the proposed bidding mechanisms are not just of theoretical interest but are simple enough to work in practice. Observed behavior in the bidding mechanisms mimics equilibrium behavior closely enough to ensure that they outperform the voting mechanism. In the bidding mechanism, for instance, most choices are truthful resulting in a 15-20% increase in efficiency.

The one man-one bid mechanism applies more broadly than just to political elections (the electorate). For example, the one share-one vote principle governs the shareholders democracy. The connection between shareholder voting and the bidding mechanism is emerging as an

[31]The connection to Wicksell's original work is highlighted by Smith (1977) and Oprea, Smith, and Winn (2007).

importation issue in the study of corporate governance and corporate control (see Burkart and Lee, 2010).

We are currently exploring a number of possible extensions that could enhance the results of this paper in interesting ways. Regarding the laboratory experiments we plan to replicate the experiments with bigger groups, and implement asymmetric treatments where more voters favor one outcome but the minority group cares more (keeping the expected value of both outcomes the same). In this case, the optimal voting rule is a super-majority rule but, as in the symmetric case, it is incapable of achieving full efficiency.

Another possible extension is to introduce uncertainty about the distribution of voters' preferences. Suppose, for instance, that there are "left" and "right" voters and that, in expectation, both groups are of the same size. But there is uncertainty about which group cares more – say, it is equally likely that the left care more or the right care more. In this case, it is easy to see that no (threshold) voting can achieve optimality whereas a bidding mechanism can be used to discover which group is more intense. We leave this and other explorations for future work.

References

Andreoni, J. and H. Varian (1999) "Pre-Play Contracting in the Prisoner's Dilemma," *Proceedings of the National Academy of Sciences*, **96**, 10933-10938.

Arrow, K. (1979) "The Property Rights Doctrine and Demand revelation under Incomplete Information," in Boskin (ed.), *Economics and Human Welfare*, New York, NY: Academic Press, 23-39.

Attiyeh, G., R. Franciosi and R. M. Isaac (2000) "Experiments with the Pivot Process for Providing Public Goods," *Public Choice*, **102**, 95-114.

Burkart, M. and Lee, Samuel (2008) "One Share - One Vote: the Theory," *Review of Finance*, **12**, 1-49.

Carroll, G. (2011) "A Quantitative Approach to Incentives: Application to Voting Rules," working paper, MIT.

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

Casella, A., A. Gelman, and T. R. Palfrey (2006) "An Experimental Study of Storable Votes," *Games and Economic Behavior*, **57**, 123-154.

Chen, Y. and R. Gazzale (2004) "When Does Learning Generate Convergence to Nash Equilibria? The Role of Supermodularity in an Experimental Setting," *American Economic Review*, **94**, 1505-1535.

Clarke, E. (1971) "Multi-part Pricing of Public Goods," *Public Choice*, **11**, 17-33.

d'Aspremont, C., and L. A. Gérard-Varet (1979a) "Incentives and Incomplete Information," *Journal of Public Economics*, **11**, 25-45.

d'Aspremont, C., and L. A. Gérard-Varet (1979b) "On Bayesian Incentive Compatible Mechanisms," in J.-J. Laffont (ed.), *Aggregation and Revelation of Preferences*, Amsterdam: North-Holland, 269-288.

Krishna, V., and M. Perry (1998) "Efficient Mechanism Design," Working Paper, Penn State University.

Ledyard, J. O. and T. R. Palfrey (2002) "The Approximation of Efficient Public Good Mechanisms by Simple Voting Schemes," *Journal of Public Economics*, **83**, 153-171.

Mailath, G. and A. Postlewaite (1990) "Asymmetric Information Bargaining Problems with Many Agents," *Review of Economic Studies*, **57**, 351-367.

Oprea, R. D., V. L. Smith, and A. M. Winn (2007) "A Compensation Election for Binary Social Choice," *Proceedings of the National Academy of Sciences*, **104(3)**, 1093.

Smith, V. L. (1977) "The Principle of Unanimity and Voluntary Consent in Social Choice," *Journal of Political Economy*, **85**, 1125-1139.

Varian, H. (1994) "A Solution to the Problem of Ecternalities When Agents are Well-Informed," *American Economic Review*, **84**, 1278-1293.

Weyl, E. G. (2012) "Quadratic Vote Buying," working paper, University of Chicago.

Wicksell, K. (1896) in Musgrave RA (ed.), *Classics in the Theory of Public Finance*, Peacock AT (St Martin's, New York), 72-118.