

Engelmann, Dirk; Grüner, Hans Peter; Hoffmann, Timo; Possajennikov, Alex

Conference Paper

Preventing the Tyranny of the Majority - Experimental Evidence on the Choice of Voting Thresholds in Bayesian Games

Beiträge zur Jahrestagung des Vereins für Socialpolitik 2017: Alternative Geld- und Finanzarchitekturen - Session: Experiments - Games IV, No. F01-V2

Provided in Cooperation with:

Verein für Socialpolitik / German Economic Association

Suggested Citation: Engelmann, Dirk; Grüner, Hans Peter; Hoffmann, Timo; Possajennikov, Alex (2017) : Preventing the Tyranny of the Majority - Experimental Evidence on the Choice of Voting Thresholds in Bayesian Games, Beiträge zur Jahrestagung des Vereins für Socialpolitik 2017: Alternative Geld- und Finanzarchitekturen - Session: Experiments - Games IV, No. F01-V2, ZBW - Deutsche Zentralbibliothek für Wirtschaftswissenschaften, Leibniz-Informationszentrum Wirtschaft, Kiel, Hamburg

This Version is available at:

<http://hdl.handle.net/10419/168295>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Preventing the Tyranny of the Majority - Experimental Evidence on the Choice of Voting Thresholds in Bayesian Games*

Dirk Engelmann[†] Hans Peter Grüner[‡] Timo Hoffmann[§]
Alex Possajennikov[¶]

February 2017

Abstract: In democracies, an absolute majority of the population may choose policies that are harmful to the rest of the population. A purpose of super-majority rules is to prevent this from happening. We study whether individuals optimally choose sub- or super-majority rules when the rights of minorities should be protected. In our Bayesian experiment individuals receive information about the distribution of valuations for a public project before knowing their own valuation. Subjects propose more extreme voting rules for more skewed distributions. However, we also find that rule choices are biased towards balanced rules, leading to under-protection of the minority that is associated with substantial welfare losses.

*We thank participants at the European ESA conference in Bergen 2016 for helpful comments. Financial support from the German research foundation (DFG), SFB 884 is gratefully acknowledged.

[†]Humboldt-Universität zu Berlin, Department of Economics, CESifo, Munich and CERGE-EI, Prague, Czech Republic. E-mail: dirk.engelmann@hu-berlin.de

[‡]University of Mannheim, Department of Economics and CEPR, London. E-mail: gruener@uni-mannheim.de.

[§]University Erlangen-Nuremberg, Department of Economics. E-mail: Timo.Hoffmann@fau.de

[¶]University of Nottingham, School of Economics. Email: alex.possajennikov@nottingham.ac.uk

“[We acknowledge] the principle that the majority must give the law.”

Thomas Jefferson to William Carmichael, 1788. ME 7:28

“[Sometimes] the minorities are too respectable, not to be entitled to some sacrifice of opinion, in the majority.”

Thomas Jefferson to James Madison, 1788. ME 7:184

1 Introduction

From its very beginning, democracy relies on voting (Thorley, 2005). Whereas voting is a common feature of democratic institutions, majority thresholds are diverse, ranging from sub- via super-majority rules to unanimity rule (Vermeule, 2004). Despite this vast diversity, simple majority rule plays the role of an important benchmark. It is actually applied so often that many people (wrongfully) identify democracy – the rule of the people – with the rule of a majority. However, this narrow interpretation of democracy has a downside. When a simple majority rules, it may choose policies that are very harmful to the rest of the population, a phenomenon labelled “the tyranny of the majority” (Adams, 1788, p. 291). It is a main purpose of super-majority rules is to prevent this from happening.

Still, simple majority rule is applied so frequently in practice, that one may wonder whether those who establish voting rules properly understand the risks of balanced rules. The present paper studies whether this is the case. We run a two stage voting experiment to find out whether individuals efficiently choose majority thresholds when the rights of minorities should be protected. In our experiment, individuals suggest voting rules at a stage when their own preferences have not yet realized.¹

In many practically relevant cases the voting thresholds have to be chosen (long) before stakeholders’ preferences materialize. Our experimental analysis focuses on these cases. In such cases, stronger majority requirements can effectively protect minorities in cases

¹It is a key politico-economic insight that rules governing collective decision making should ideally be chosen before individual preferences about outcomes have realized. Rawls (1972) builds his theory of justice on the view that a fair system should maximize expected utility under a veil of uncertainty, and Brennan and Buchanan (1985) argue that the establishment of such a system is more likely if the decision about the institution is taken *before* preferences have materialized. The present paper considers this case. There are other situations where informed voters pick rules. This case has been analyzed experimentally in Engelmann and Grüner (2013).

where preferences in favor of or against a decision may be particularly strong.² Efficient minority protection requires that these majority thresholds are optimally adjusted to the distribution of preference intensities. Our main finding is that while our subjects react to the distribution of preferences in the direction required for efficient minority protection, rule choices are biased towards balanced rules, leading to under-protection of the minority and substantial welfare losses.

Our experiment is designed as follows. Subjects decide on a voting rule before knowing their own valuation for an alternative that changes the status quo, such as implementing a public project. Experimental participants decide in groups of five in a two-step procedure about whether to enact the change. In the first step, for given distribution of possible valuations they suggest a voting rule (such as simple majority, super-majority, or even an extreme minority one where one vote in favor is sufficient to enact the change). They make these decisions for 21 different distributions. In the second step the participants learn their own valuations and vote about the implementation of the change according to the voting rule randomly selected from the suggested ones.³ Thus, the experimental subjects vote on the outcome after receiving the information about their own valuation but they have to decide on the voting rule before the uncertainty is lifted.

We choose payoff distributions that vary in their skewness so that the total-payoff maximizing voting rule ranges across all possible thresholds from unanimity required for the public project to unanimity required against the public project. The total-payoff maximizing voting threshold for a decision increases in the expected cost for the opponents rises and decreases in the expected gains for the supporters. Our analysis shows that this is true qualitatively for the chosen voting rules. We find strong evidence for a monotonic relationship between the relative preference intensity of supporters and opponents and the chosen voting rule. This monotonicity is weaker than predicted, however, and subjects tend to shy away from choosing unanimity rules. While on average subjects respond in their rule choice to the underlying distribution of valuations by picking more extreme voting rules for more skewed distributions, fewer than half of the rule choices are for the total-payoff maximizing rule. The suggested rule choices imply that on average more than one third of the total surplus would be lost if all subjects voted selfishly in the second stage, which they typically do.

There is substantial heterogeneity across subjects. While most subjects show a positive correlation between the chosen and the payoff-maximizing rule, only one among 130 subjects chooses the total-payoff maximizing voting rule more than two thirds of the times (and

²In a binary voting decision with conditional expected monetary gains E^+ of winners and losses E^- of losers, a total payoff maximizing voting rule should specify that a decision is made if and only if more than a share of $E^- / (E^- + E^+)$ votes are in favor.

³Since we are interested mostly in the choice of the voting rules, only for three of the 21 distributions, randomly selected, valuation are drawn and votes are cast.

indeed does always do so). Experimental participants do not exhibit a particular preference for majority voting (only 4% choose the simple majority voting rule more than half of the time). They do, however, show a tendency towards moderate rules (73% of subjects never propose an extreme rule requiring unanimity for one of the outcomes, so that a single subject can determine the outcome), even though a unanimity rule is efficient in one third of the tasks. As a result, ex-ante rule choice is not generally sufficient to overcome the “tyranny of the majority”. Experimental participants lose most expected surplus when unanimity rules would be total surplus maximizing, more than twice the loss in payoff when majority voting is optimal.

There are many important real world situations in which ex-ante choices of rules are relevant. In our experiment, we consider an election with a small (five) number of voters. Hence, taken literally, our experiment evaluates individual’s ability to choose rules for clubs (sports clubs etc.), local trade unions, parent school boards, faculties and other “small” institutions. On the political level, these “small” decision making bodies also include several national and international organizations that take decisions in committees (including Parliamentary subcommittees, the United Nations, the IMF, Nato, the EU council of ministers, or central bank councils).

Many procedural and constitutional rules make use of the simple majority rule even in classes of situations where gains and losses from decisions are unlikely to be distributed symmetrically. Below are some examples of decision problems in which the optimal rule is likely to be biased towards one outcome in order to protect minority interests but the actual rule in parliament is the simple majority rule:

1. Decision on same sex marriage law. These decisions seem to affect the feelings of some people to some - perhaps rather limited - extent, but matter - potentially a lot - to those who are directly affected by them.
2. Decisions about public investment into research on specific rare diseases. Ex-ante it is very unlikely that one benefits from this kind of research but interim one may care a lot.
3. Rules regarding the consumption of specific types of food. Some people dislike it when others consume specific types of food (e.g. whales, horses, dogs). Those (few) who like that food may however care a lot about it.

While, for obvious reasons, our experiment considers a small election, the basic choice problem is similar to the one of the choice of voting rules in large elections where voting rules should also take the joint distribution of preferences into account. Constitutions also usually specify different majority requirements for different classes of actions. An example is the German constitution that for most issues requires a simple majority to make a decision in

parliament. However, for some issues such as constitutional changes the constitution requires a two third majority. Some rules cannot even be changed unanimously. Such majority thresholds in constitutions rarely change over time. Therefore when these thresholds were implemented, these were clearly choices under a veil of ignorance.

While there is a lot of empirical research about what determines individuals' voting behavior under a given rule, very little is known about how individuals choose rules that they want to apply in the future (with the exception of Engelmann and Grüner, 2013). More specifically, little is known about whether individuals are able to properly adjust the voting threshold to the underlying problem - i.e. the underlying distribution of voter preferences. A recent paper by Weber (2016) experimentally studies the choice of voting rules for representatives of homogenous groups, but does so in a setup with binary valuations of identical absolute size that does not permit to study how the absolute size and size-distribution of individual valuations react the choice of voting rules.

2 Experiment design

Our experiment considers the following two-stage decision setup. In the first stage, experimental participants are given a distribution of possible valuations (positive and negative) for alternative A, representing net payoffs from a public project, or from a change from the status quo. This distribution is the same for each of the members of the group of five individuals. The valuation for alternative B (representing the status quo) is always 0€ for all individuals. For example, a distribution of valuations for alternative A can be

- the valuation of A is 2€ with probability 2/3 and it is -5€ with probability 1/3.

For a given distribution, each individual chooses one of the threshold voting rules, specifying how many individuals in the group of five need to vote for alternative A for it to be adopted. Abstentions are not allowed in the voting stage, thus the threshold for alternative A automatically implies the threshold for alternative B. Therefore, for groups of five voters, the available voting rules are:

- Rule I. At least 1 vote for alternative A is required for A to be chosen, thus 5 votes for alternative B are required for B to be chosen (unanimity for B);
- Rule II. At least 2 votes for alternative A are required for A to be chosen, thus at least 4 votes for alternative B are required for B to be chosen (qualified majority for B);
- Rule III. At least 3 votes are required for either A or B to be chosen (simple majority), that is, whichever has more votes wins;

Rule IV. At least 4 votes for alternative A are required for A to be chosen (qualified majority for A), hence at least 2 votes for alternative B are required for B to be chosen;

Rule V. 5 votes for alternative A are required for A to be chosen (unanimity for A), thus at least 1 vote for alternative B is required for B to be chosen.

In the second stage of the experimental setup, one of the suggested rules by the five group members is randomly chosen to be the actual voting rule (a random dictator mechanism).⁴ The subjects are informed about which rule was chosen. The participants' valuations are then realized according to the given distribution and each participant learns his/her own valuation for alternative A. The participants then cast a binary vote (either for A or for B). The votes are tallied and the outcome (either A or B) is decided according to the chosen voting rule.

The two-stage procedure is designed to have individuals make decisions on the voting rule under the "veil of ignorance" (in the first stage, before they know their own valuation). The second step is the one more commonly tested in the experimental literature on voting, and we include it as a check on subjects' voting behavior and to make the rule-choice stage incentive compatible. However, our interest is mainly in the decision in the first step.

In the second stage of the procedure, it is a (weakly) dominant strategy to vote for A if your realized valuation for A is positive and vote for B if your realized valuation for A is negative. If voting in the second step is going to follow the dominant decisions, and if individuals maximize their expected payoff (or the expected payoff of the whole group), then which rule is optimal depends on the distribution. For example, for the distribution of valuations for A 2€ with probability 2/3 or -5€ with probability 1/3, the optimal rule Rule IV is skewed towards B: four votes for A are needed. This is because benefits from A are much lower than the losses that it causes. Even though players with losses are likely to be in minority (the probability of a negative value is only 1/3), they need to be protected from ex-ante point of view. For the optimal decision rule it is actually irrelevant how likely negative (or positive) values are, because the rule becomes relevant for a given size of the minority. Since in the first stage the individuals make a decision which rule to suggest before knowing their own valuation for alternative A (and thus they may end up being the players with valuation -5€), it is in their own interest to suggest such a protection.

In the experiment, the two stages did not immediately follow each other. In fact, in the first part of an experimental session, the participants made rule choice for 21 distributions of valuations for alternative A. These distributions are listed in Table 1. The order in which the distributions were shown to the subjects was randomly determined and thus varied between subjects; there was no feedback between rounds in the first part of the experiment.

⁴We chose a random dictator mechanism because it is incentive compatible and moreover easy for participants to understand to be incentive compatible.

No.	V1	V2	V3	Pr(V1)	Pr(V2)	Pr(V3)	Optimal Rule
1	-5	1		1/3	2/3		V
2	-4	1.5		1/3	2/3		IV
3	-2.5	2.5		1/2	1/2		III
4	-1.5	4		2/3	1/3		II
5	-1	5		2/3	1/3		I
6	-5	1		1/2	1/2		V
7	-4	1.5		1/2	1/2		IV
8	-2.5	2.5		1/3	2/3		III
9	-2.5	2.5		2/3	1/3		III
10	-1.5	4		1/2	1/2		II
11	-1	5		1/2	1/2		I
12	-5	0.5	1.5	1/3	1/3	1/3	V
13	-4	1	2	1/3	1/3	1/3	IV
14	-2.5	2	3.5	1/2	1/3	1/6	III
15	-1.5	3	5	2/3	1/6	1/6	II
16	-1	3.5	6.5	2/3	1/6	1/6	I
17	-1.5	0.5	5	1/3	1/3	1/3	II
18	-2	1	4	1/3	1/3	1/3	II
19	-3.5	-2	2.5	1/6	1/3	1/2	III
20	-5	-3	1.5	1/6	1/6	2/3	IV
21	-6.5	-3.5	1	1/6	1/6	2/3	V

Table 1: Distributions used in the experiment

The distributions are chosen to vary in their skewness so that the total-payoff maximizing voting rule ranges across all possible thresholds from unanimity required for alternative A (rule V) to unanimity required for alternative B (rule I). Distributions 1-5 are taken as the base; the rest of the distribution are derived from them. For example, distributions 6-11 are the variants of distributions 1-5, but with different probabilities of each value. Distributions 12-21 are variants of distributions 1-5 but with increased variance, with one of the outcomes in distribution 1-5 split into a probability distribution with the mean being the same outcome as in the corresponding base distribution.⁵

After the first part was finished, three of the 21 distributions were randomly selected for the second part of an experimental session. Therefore subjects in the same session voted on the same three distributions, potentially using a different voting rules in different groups, but subjects in different sessions typically voted on different distributions. In each round of the second part, valuations for the participants were drawn according to the distribution and the participants were informed about their own valuation. One voting rule among those suggested by the five group members for this distribution was randomly selected and the participants were informed about which rule is selected. The participants then voted for alternative A or alternative B, and the outcome of the voting was determined according to the voting rule. At the end of a round, the participants were informed about the outcome of the voting and their payoff.

The experiments were conducted at the Laboratory of Experimental Research Nuremberg (LERN) in December 2015. We ran 5 sessions, with the number of participants ranging between 15 and 30 in each. In total there were 130 participants. The experimental sessions were programmed with using zTree (Fischbacher, 2007) and the recruitment of the participants was done with ORSEE (Greiner, 2015). Each participant was given a starting budget of 15€. The valuations in Table 1 are in Euro; with the three distributions actually played out, the minimum amount a participant could earn was 3€ and the maximum amount was 27€.

Our experimental setup permits to test the following hypotheses:

1. Distribution matters: Voting rule choices take into account the skewness of the distributions towards larger positive or negative outcomes, thus reflecting which rules are optimal.
2. Preference for majority rule: Rule III is chosen more often than is warranted by it being theoretical optimal. Such a preference could result from a preference for democracy and a perception that majority voting best represents democracy.

⁵Following this rule, distribution #17 should have been -1.5 with probability 1/3, -0.5 with probability 1/3, 5 with probability 1/3, with the optimal rule being rule I. Due to a copying error, 0.5 was entered instead of -0.5, making II the optimal rule.

3. Asymmetry: There is a systematic bias towards rules IV and V as compared with rules I and II. This may reflect risk attitudes: since alternative A is a more risky than alternative B, a risk averse person would suggest rules IV and V more often. A maxmin person should always pick rule V.
4. Variance matters. Although distributions with one outcome replaced by a mean-preserving spread (for example, distributions 1 and 12 in Table 1) have the same theoretically optimal rule, decisions in the experiment may not reflect this. A distribution with higher variance is less attractive for a risk-averse person, who should hence choose more conservative rules for a mean-preserving spread.
5. Probabilities of various outcomes matter. Although distributions with the same outcomes but different probabilities of them (for example, distributions 1 and 6 in Table 1) have the same theoretically optimal rule, decisions in the experiment may not reflect this. Specifically, having a higher probability of a loss may make a distribution appear to be more risky and hence lead to more conservative rule choices by risk-averse subjects.
6. Extreme rules: A risk-neutral person who does not understand that there will be voting will pick rule V or rule I depending on the expected payoff of the simple lottery determining their own valuation.
7. Focality: Prominent numbers, or easy to calculate ratios of numbers or of probabilities may play a role in the decision process.

3 Experimental results

3.1 Rule Choices - aggregate data

As a first basic result, rule choices of the subjects are regressed on the optimal rule. Table 2 presents this for each session separately and over all sessions, where *OptRule* denotes the optimal rule variable.⁶

The correlations coefficients and the regression coefficients on *OptRule* are all significantly different from 0 and from 1. Thus the subjects appear to take into account which rule is optimal for a given distribution, as the first hypothesis in the previous section suggests, but clearly not as much as the optimal rule implies. For example, from the last line of the table, if *OptRule* = 1, then subjects' average choice is predicted to be 2.215 rather than 1; if *OptRule* = 4, the predicted choice is 3.511 rather than 4. Note also that there does

⁶Since we did not provide any feedback in the first part of the experiment, each of our subjects presents an independent observation.

Table 2: Regression of choice on optimal rule

	Obs	Correlation	Regression
Session 1	525	0.407***	$1.605^{***} + 0.459^{***} \text{OptRule}$ ($R^2 = 0.165$)
Session 2	630	0.520***	$1.408^{***} + 0.551^{***} \text{OptRule}$ ($R^2 = 0.270$)
Session 3	630	0.349***	$2.026^{***} + 0.362^{***} \text{OptRule}$ ($R^2 = 0.122$)
Session 4	630	0.342***	$2.091^{***} + 0.349^{***} \text{OptRule}$ ($R^2 = 0.117$)
Session 5	315	0.449***	$1.731^{***} + 0.455^{***} \text{OptRule}$ ($R^2 = 0.201$)
Overall	2730	0.409***	$1.783^{***} + 0.432^{***} \text{OptRule}$ ($R^2 = 0.167$)

not appear to be much difference between sessions. Hence there is qualitative support for Hypothesis 1, but subjects do not fully react to the optimal rule as theoretically predicted.

Note that although the coefficient of *OptRule* is highly significant, R^2 is not high. There is a lot of variance in subject's choices that is not explained by which rule is optimal for a given distribution. Table 3 shows the distribution of actual rule choices by optimal rule.

Table 3: Rule choices

OptRule	Obs	Median	Mean (St.Dev.)	Rule choices		
				I or II	III	IV or V
I	390	2	2.254 (1.330)	64%	17%	18%
II	650	2	2.594 (1.261)	51%	24%	25%
III	650	3	3.080 (1.189)	30%	35%	34%
IV	520	4	3.598 (1.297)	21%	22%	57%
V	520	4	3.896 (1.348)	18%	13%	69%

The table shows just how much noise there is in the choice of rules: even for the distributions where an extreme rule is optimal (1 or 5), 18% of suggested rules are on the wrong side of the simple majority rule III.⁷

There is not much evidence in favor of Hypothesis 2 that subjects have a preference for majority rule in the sense that rule III is chosen most often. Rule III is not chosen much more often than other rules. Even for the distributions where III is the optimal rule, it is chosen only about 1/3 of the time.

Hypothesis 3 predicted asymmetry, i.e. rules I and II are chosen less often than rules IV and V. Although there is a slight shift to the rules with larger numbers (evident from the

⁷Note also that standard deviation of suggested rules is lower for those distribution whose optimal rule is III and increases as the optimal rule moves to the extremes. It is not clear though if this increase is anything more than simply allowing larger distances from the mean.

averages and the percentages), it does not appear to be large, providing at best very weak support for Hypothesis 3.

Our distributions were chosen so that some of them are variants of others but with different variance and different probabilities of the same values. Theoretically, changing variance or probabilities of values does not change which rule is optimal. Table 4 lists means and standard deviations of suggested rules for the distributions, organized by the optimal rule (recall that distribution #17 was supposed to be a variant of distribution #5 but was not correctly implemented and is thus omitted from the table):

Table 4: Effects of changes in distributions

OptRule	Base	Var 1	Var 2	Prob 1	Prob 2
I	D#5 2.508 (1.342)	D#16 2.208 (1.322)		D#11 2.046 (1.293)	
II	D#4 2.892 (1.277)	D#15 2.677 (1.331)	D#18 3.085 (1.168)	D#10 2.208 (1.132)	
III	D#3 2.962 (0.901)	D#14 2.746 (1.095)	D#19 3.500 (1.087)	D#8 2.446 (1.114)	D#9 3.746 (1.235)
IV	D#2 2.454 (1.283)	D#13 3.308 (1.167)	D#20 3.708 (1.349)		D#7 3.923 (1.310)
V	D#1 3.854 (1.330)	D#12 3.807 (1.365)	D#21 3.992 (1.327)		D#6 3.931 (1.376)

Hypothesis 4 states that the variance of distribution matters. From Table 4, it does not appear that variance per se matters. Var 1 distributions had an increased variance of positive values compared with the Base ones while Var 2 distributions had an increased variance of negative values. The means for Var 1 distributions are lower than for the Base distributions and those of Var 2 are higher. It appears that what matters perhaps is the magnitude of values (larger numbers loom larger, reflecting Hypothesis 7 to some effect) but not their variance per se.

Hypothesis 5 states that probabilities of values matter although they should not in theory. Table 4 lends some support to this hypothesis. Prob 1 distributions had an increased probability of positive values compared with the Base distributions and Prob 2 distributions

had an increased probability of negative values. This seems to have an effect: an increased probability of positive value lowered the average suggested rule (thus fewer votes in favor of alternative A would be needed) while an increased probability of negative values led to higher average suggested rules (this can be most clearly seen for distributions with optimal rule III, where changes in probability moved the choices closer to II and IV respectively).

Hypothesis 7 states that prominent numbers play a role. Depending on how this hypothesis is interpreted, it may have some support. Distributions with higher variances meant that numbers were “stretched”, so that e.g. 4 becomes {3 with probability 1/2 and 5 with probability 1/2}. Then 5 perhaps played a more prominent role than 4 in realizing that the optimal rule should involve a low threshold (as indeed happens since High Var 1 have on average lower suggested rule than Base) (although the presence of two positive values rather than one may also play this role).

Finally, Hypothesis 6 that someone risk-neutral who does not understand that there will be voting will pick rule V or rule I depending on the expected payoff of the lottery determining their own value. There may be some subjects who behave in a similar (though less extreme) way (e.g. Subjects #3, #5, #11, #17 in session 1, Subjects #18, #20, #21, #23 in session 2, ...) but this is left for the analysis of individual types rather than of aggregate patterns.

To summarize the rule choice results,

Result 1: Rule Choices Rule choices follow the optimality of rules to some extent although there is a lot of heterogeneity and possible errors. There is little evidence of preference for majority rule III or of risk-aversion (or asymmetry). Variance in a distribution plays little role, but shifts in probability and in magnitude (or in the number) of values seem to play a role.

3.2 Rule Choices - individual data

We also looked at the individual behavior and tried to identify whether the heterogeneity we observe in the aggregate data is generated by a few subjects or whether choices of subjects are typically noisy. To provide some overview of the general performance of subjects, Table 5 shows how many subjects made less than 5, between 5 and 10, between 11 and 15 and more than 15 optimal rule choices. About half of the subjects made between 5 and 10 rule choices that correspond to the optimal rule. Remarkably the one subject that choose an optimal rule more than 15 times choose the optimal rule in all 21 distributions.

To see which implications these choices have on the surplus extracted, we calculate for each individual the expected missed surplus using all 21 rule choices an individual made and based on the assumption that subjects would have voted rational and selfish (i.e. for alternative A if valuation is positive and for alternative B if valuation is negative).

Table 5: Overview: Efficient rules choices by individuals

NumOfOptRuleChoices	Frequency
≤ 5	43
$5 < x \leq 10$	63
$10 < x \leq 15$	23
> 15	1

Table 6 states the expected missed surplus per session, as well as the expected missed surplus by the best (“min”) and worst (“max”) individual in each session. Thus this analysis takes the size of the error by not choosing the most efficient rule into account. Note that these are absolute numbers **per group**, e.g. 90.08 means that a group of five subjects would have in expectation earned 90 Euro less than with the optimal rule choices. As a reference, always choosing the most efficient rule yields a group surplus of 63.89 Euro (= 12.8 Euro/subject). Therefore “missing” 90 Euros means, that on average the group would have lost money.

Table 6: Expected surplus missed by session

Session #	mean	sd	min	max
1	25.04	22.69	4.53	90.08
2	20.01	14.39	0.00	53.45
3	28.12	18.98	5.67	84.10
4	26.22	17.77	5.14	69.10
5	22.05	10.97	5.99	45.36
Total	24.52	17.79	0.00	90.08

On average 25 Euro of the achievable surplus is missed. For the five sessions the expected missed surplus varies between 20 and 28 Euro, however, the individual variation is much larger. Session 1 stands out with regard to the standard variation and this is driven by three subjects that missed 60-90 Euro of the expected surplus.⁸

Figure 1 is a scatter plot with the expected missed surplus of each individual on the y-axis and the number of efficient rule choices on the x-axis. While the number of efficient

⁸Subjects in session 1 did not answer understanding questions after reading the instructions. In all other sessions subjects had to answer several questions regarding the resulting payoffs given various distributions and voting rules to ensure that participants understood the relationship between choosing a rule in stage 1 and the effect of these decisions rules for stage 2.

rules choices clearly drives the total surplus missed, the scatter plot also demonstrates that the variation is substantial. Especially for subjects that selected the optimal rule less than 10 times, the missed surplus varies a lot.⁹

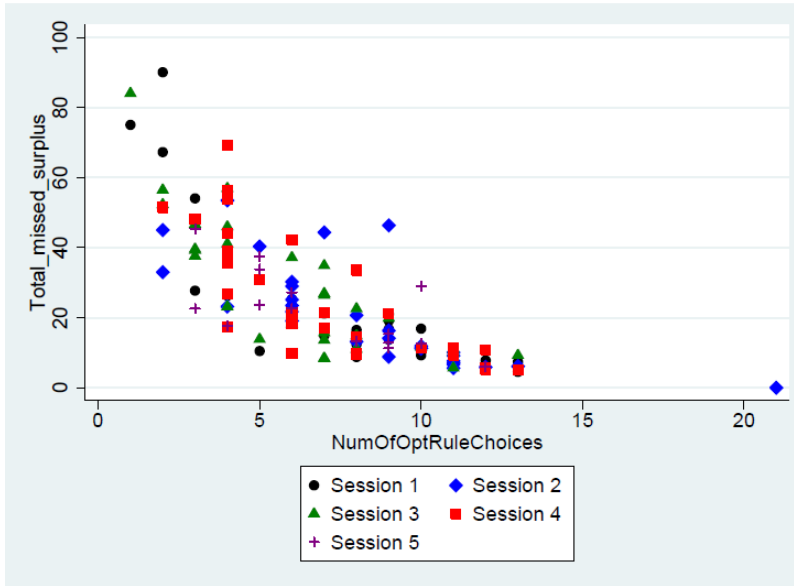


Figure 1: Scatter plot - missed surplus

We divided subjects into types according to the correlation between the chosen rules and the optimal rules. There are 4 different classifications based on the individual correlation between the choices and the efficient rule. Two subjects always choose the same rule (one always rule I, the other always rule V), therefore these two are not classified.

For most subjects the correlation is positive and for more than half it is above 0.5. These numbers indicate that a majority behaves as least in the predicted direction. Subjects seem to take the expected payoff into account when choosing the group decision rule. However, about 30% of subjects have a correlation below 0.25, which indicates that their choices are not in line with the efficient rule very often. In order to understand especially this third of the subject population we looked for various other “types”, most of them based on our hypotheses.

We only mention very few types in detail here, since many specified types cannot be identified very often. There are only 12 subjects (9%) who select the efficient rule in at least

⁹A possible effect of not asking understanding question in session 1 can be seen in the upper left: the three black dots indicate that three subjects in this session missed a lot of expected surplus. So while the session average is not very different from the others, maybe some individuals could be heavily affected.

Table 7: Correlation with efficient rule choice

Correlation	Freq.	Percent
< 0	25	19.53
≥ 0 and < 0.25	11	8.59
≥ 0.25 and < 0.5	22	17.19
≥ 0.5	70	54.69
Total	128	100.00

12 rounds. Even fewer subjects have a strong preferences for rule III. Only 5 out of the 130 subjects voted for rule III in at least 12 rounds (4%). The same number of subjects voted only for the extreme rules (rule I or V). Clearly rule III does not seem to be special.

Another finding is that 73% (95 subjects) **never** suggested rule I or V and therefore only used rules II-IV (even though an extreme rule was the optimal rule in 7 of 21 distributions). When generating types using our hypotheses that rule choices are affected by changes in variance or probabilities the individual choices confirm the results from the aggregate data. Many subjects change their chosen rule when either the variance or the probabilities change. However, we did not find one type that is very prominent. It looks like both variations lead to different chosen rules, but no single reason “dominates” the others.

The self classification of risk aversion does not explain a lot. Splitting the sample for subjects that stated to be “risk averse” or “somewhat risk averse” in a questionnaire administered at the end of the experiment (89 subjects, 68%) shows that these subjects tend to choose “higher” than optimal rules a little more often, but the differences are small (7.5 times compared to 6.9 for the remaining 32%). Just conditioning on those that state that they are “risk averse” yields somewhat bigger differences (9.6 vs. 7.1 times), but this regards only a small fraction of the population (10 subjects). So risk aversion “works” in the expected direction, but (at least self-assessed) risk aversion does not seem to be a major driver of the results.

Result 2: Short summary on types There is quite some heterogeneity in the rule choice behavior of subjects. Looking at the classification using the correlation between the rule choice and the optimal rule shows that about 1/3 of subjects behaves rather noisily. The other 2/3 react to the optimality of different rules, but to a smaller extend then it would be optimal.

3.3 Surplus extraction rates by rule

To identify the surplus extraction rates we used the “expected missed surplus” measure. Table 8 below states for the five rules how much surplus was lost on average if the given rule was the optimal rule. Since these numbers are absolute numbers, they do not take into account how much surplus was actually possible with the optimal rule. Therefore the column “maximal surplus” states the mean expected surplus with the efficient rule.

Table 8: SER rule summary table

OptRule	mean	sd	maximal surplus
I	1.61	2.37	7.16
II	1.22	1.37	4.41
III	0.65	0.71	2.55
IV	1.12	1.28	1.37
V	1.47	2.35	0.53
Total	1.17	1.69	3.04

Looking at the average of surplus lost, a U-shape is clearly visible. If an extreme rule is optimal, the surplus lost is clearly larger than with the intermediate rules. Especially with rule III, the surplus loss is the smallest. However, as mentioned these numbers are not directly comparable in the sense that only the absolute values are shown. One aspect that plays a big role is the fact that for many inefficient rules, the generated surplus is negative. This is especially true for the case that rule V is optimal.

It appear that under the veil of ignorance, subjects are on average inclined to pick balanced voting rules even in unbalanced situations. The fact that many subjects never chose an extreme voting rule might indicate that many subjects are just drawn to the middle of their choice set. Together with the fact that rule III does not stick out, a possible explanation would be that subjects react to the optimal rule, but often stay away from extreme rules.

3.4 Voting behavior

Although our main interest is in the choices of voting rules in the first stage of the procedure, we also looked at the voting choices in the second stage. Table 9 summarizes how consistent voting was with the realized values, by session and in aggregate (Consistent choice means voting for A if one’s value is positive and voting for B if one’s value is negative):

Voting in Session 1 was less consistent than in Sessions 2, 3, 4 but not that different from Session 5. Voting was more consistent when values were negative which may be just by

Table 9: Consistency in voting

Session	Consistent	If positive value	If negative value
1	56/75 (75%)	30/45 (67%)	26/30 (87%)
2	80/90 (89%)	45/51 (88%)	35/39 (90%)
3	81/90 (90%)	37/43 (86%)	44/47 (94%)
4	83/90 (92%)	43/47 (91%)	40/43 (93%)
5	32/45 (71%)	20/32 (63%)	12/13 (92%)
Total	332/390 (85%)	175/218 (80%)	157/172 (91%)

chance because the randomly selected distributions to be voted on in the second part more often had small positive and large negative values than the other way round (the average of realized positive values was 2.07 and the average of realized negative values was -2.65). The “inconsistency” in voting may thus reflect a concern for efficiency (especially relevant if the voting rule used is not the optimal one) or may be because of not wanting to enforce a loss on others (concern for equity).

To check the efficiency concern, we calculated the realized average payoffs (net of 15 Euros) in each session (Actual), as well as what payoffs would have been if all votes were consistent with values (but the voting rule is the one actually used, not necessarily optimal) (Sincere) and if in addition to consistent voting the rule were optimal (Optimal)

Table 10: Voting and realized net actual payoff

Session	Actual	Sincere	Optimal
1	0.367	0.433	0.607
2	0.533	0.583	0.672
3	0.156	0.072	0.283
4	-0.044	-0.033	0.361
5	0.144	0.322	0.578
Total	0.236	0.264	0.487

As can be seen from Table 10, voting “insincerely” did not improve the average payoff (except in Session 3). “Insincere” voting did not make the average payoff much worse though. The main source of inefficiency appears to be the choice of voting rules: the payoff with the optimal rule (and consistent voting) would have been about twice as much as the actual payoff (or as the payoff if voting were sincere with actual rules used). Note also that from the analysis of rule choices, Sessions 4 and 3 are the ones where the correlation between

actually chosen rule and the optimal rule is the weakest, and those sessions appear to lose most from not using the optimal rule.

Result 3: Voting Voting is mostly consistent with values; where it was inconsistent, it did not noticeably changed the obtained payoff. Not choosing the optimal rule appears to be the main source of lower average payoff.

4 Conclusion

Many collective decisions are governed by institutions that rely on voting procedures to aggregate stakeholders' preferences. Often, the very same institution, including as diverse institutions as faculty boards or the US Senate, relies on several different voting procedures to decide on different kinds of issues. While, in some cases, voting outcomes are determined by a simple majority of participants, other decisions require support of a supermajority or even unanimous support. Clubs frequently decide on time and place of their next assembly by simple majority, whereas the decision to accept a new club member often requires more widespread support (see Grüner and Tröger, 2017 for a list of examples).

We analyzed the behavior of experimental subjects in a situation where they can decide on voting rules under “the veil of ignorance” – i.e. before knowing their own valuations for possible alternatives, even though they know the distribution of possible valuations (and they know that they may be actually playing this situation out, voting on the alternatives in the second stage of the experiment). For distributions skewed towards (possibly unlikely) high positive values, or towards very low negative values, a voting rule that is optimal in such a situation involves clear departures from the simple majority rule.

We find that subjects on average adjust the voting rule to the distribution. However, they fail to adjust the rules strongly enough, missing, in expected terms, quite a substantial proportion of available surplus. Thus, minorities with strong preferences are systematically under-protected against decisions made by a majority of voters in our experiment. There are several possible reasons why this may be the case. One reason could be that simple majority rule is frequently associated with the mere concept of democracy (i.e. the majority decides). Nevertheless, we do not see a clear preference for the simple majority rule. Another possible reason is that it is difficult to calculate which of the rules different from simple majority rule is optimal for a given distribution. Subjects thus may decide to avoid the extreme rules (such as rules I and V in our setting).

Our experiments suggest a way to explore designing voting rules and constitutions more generally. While it may be the case that rules can be designed to protect minorities, it remains challenging for decision-makers to determine how much protection there should be even in situations where uncertainty is relatively transparent and quantifiable. It is even

more difficult to determine the necessary protection and thus optimal voting rules in the real world, but the analysis in this paper sheds a light on how such decisions can possibly be improved.

References

- [1] Adams, J (1788) “A Defence of the Constitutions of Government of the United States of America”, Vol. 3, London.
- [2] Brennan, G. and J.M. Buchanan (1985). “The Reason of Rules: Constitutional Political Economy.” Cambridge: Cambridge University Press.
- [3] Engelmann, D. and V. Grimm (2012). “Mechanisms for Efficient Voting with Private Information about Preferences.” *Economic Journal* 122, 1010–1041.
- [4] Engelmann, D. and H.P. Grüner (2013). “Tailored Bayesian Mechanisms: Experimental Evidence from Two-Stage Voting Games.” CEPR Discussion Paper No. 9544.
- [5] Fischbacher, U. (2007). “z-Tree: Zurich Toolbox for Ready-made Economic experiments.” *Experimental Economics* 10, 171–178.
- [6] Greiner, B. (2015). “Subject Pool Recruitment Procedures: Organizing Experiments with ORSEE.” *Journal of the Economic Science Association* 1, 114–125.
- [7] Grüner, H.P. and T.Tröger (2017) “Optimal costly voting”, CEPR.
- [8] Jackson, M.O. and H.F. Sonnenschein (2007). “Overcoming Incentive Constraints by Linking Decisions.” *Econometrica* 75, 241–258.
- [9] Rawls. J (1972). “A Theory of Justice.” Oxford: Clarendon Press.
- [10] Thorley, J. (2005) “Athenian Democracy”, Routledge.
- [11] Vermeule, A. (2004) “Submajority Rules: Forcing Accountability upon Majorities”, U of Chicago, Public Law Working Paper No. 54.
- [12] Weber, M. (2016). “Choosing the Rules: Preferences over Voting Systems in Assemblies of Representatives.” mimeo.