

Democracy for Polarized Committees: The Tale of Blotto’s Lieutenants

Alessandra Casella, Jean-François Laslier and Antonin Macé

Abstract

In a polarized committee, majority voting disenfranchises the minority. By allowing voters to spend freely a fixed budget of votes over multiple issues, Storable Votes restores some minority power. We study a model of Storable Votes that highlights the hide-and-seek nature of the strategic game. With communication, the game replicates a classic Colonel Blotto game with asymmetric forces. We call the game without communication a decentralized Blotto game. We characterize theoretical results for this case and test both versions of the game in the laboratory. We find that, despite subjects deviating from equilibrium strategies, the minority wins as frequently as theory predicts. Because subjects understand the logic of the game — minority voters must concentrate votes unpredictably — the exact choices are of secondary importance. The result is an endorsement of the robustness of the voting rule.

Keywords: Storable Votes, Polarization, Colonel Blotto, Tyranny of the Majority, Committees

1 Introduction

How should political power be shared? Majoritarian democracy is desirable under many criteria [8, 25, 27] but in polarized societies, where the same group is on the losing side on all essential issues, it effectively disenfranchises the minority.¹ Polarization can exist in rich as well as poor countries, in old as well as new democracies, and can pre-date the democratic institutions or be generated by the institutions themselves. For instance, polarization can rest on the exogenous separation of the population into two main religions. But it can also result from electoral competition in a winner-take-all system, in otherwise very different countries.² Emerson [11, 12], referring to Northern Ireland, the Balkans, and other places plagued by civil wars, claims that majority rule is the problem, not a solution, and that more consensual rules must be implemented.

In modern democracies, the main tool for power-sharing is representation. The complexity of the political agenda, which unfolds over time and allows changing coalitions, logrolling, and compromises makes representation in Parliament valuable even to a minority. When group barriers are permeable, the minority can occasionally belong to the winning side. But

¹Political philosophy has long recognized that the *tyranny of the majority* poses a fundamental challenge to the legitimacy of majority voting [9].

²See Fiorina et al. [13], Jacobson [20] for the US case, or Eifert et al. [10], Kabre et al. [21], Reynal-Querol [28] for African cases.

when preferences are fully polarized and the power of a cohesive majority bloc is secure, the minority remains disenfranchised. In some instances, therefore, power-sharing is imposed directly, and the constitution grants executive positions to specific groups, typically on the basis of their ethnic or religious identity.³ The problem is that constitutional provisions of this type are difficult to enforce and heavy-handed, unsuited to changing realities. We argue that power-sharing in polarized societies could be achieved in a more subtle and more flexible manner via the design of appropriate voting rules.⁴

The *Storable Votes* mechanism (henceforth SV) does just that: it allows the minority to prevail occasionally and yet is anonymous and treats everyone identically [5]. In a setting with a finite number of binary issues, the SV mechanism grants a fixed number of total votes to each voter with the freedom to divide them as wished over the different issues, knowing that each issue will be decided by simple majority. SV can apply to direct democracy in large electorates, or to smaller groups, possibly legislatures or committees formed by voters' representatives, as in the model we study in this paper.

Although easy to describe, SV poses a challenging strategic problem: *how* should the votes best be divided? Testing whether voters are in fact able to use SV profitably is thus central to recommending its use in concrete applications. We study the SV mechanism in a model that highlights a central ingredient of the strategic environment: the *hide-and-seek* nature of the game between majority and minority voters. If the majority spreads its votes evenly, then the minority can win some issues by concentrating its votes on them, but if the majority knows in advance which issues the minority is targeting, then the majority can win those too.

Such strategic interaction is studied in the literature under the name of Colonel Blotto game [3, 15], with two armies (here: groups) allocating soldiers (votes) on several battlefields (issues), each battlefield being won by the army with the largest number of soldiers. The SV model differs however on one important dimension: the game is decentralized, as the votes of a given group (minority or majority) are controlled by several voters.

To our knowledge, the decentralized Blotto game has not been studied before. We first stress the difference between the decentralized game and its centralized counterpart, and we then provide theoretical results on the decentralized game. We experiment the game in the laboratory, under two conditions: in a first phase, subjects cannot communicate, and in a subsequent phase, subjects can communicate within their group before allocating their votes. The analysis of the decentralized game serves as reference for the game without communication, whereas the game with communication is well described by the centralized game studied in Hart [17]. In both treatments, the essential logic of the game — the minority needs to concentrate and randomize its votes — is immediately clear to minority players in the lab.

³For example, in Lebanon [26, 31], in Mauritius [4], and occasionally elsewhere [24].

⁴Note that neither vetoes or supermajority requirements, nor log-rolling can overcome the problem posed by a systematic minority. If on each issue there is a fixed majority of, say, 60 percent, versus a fixed minority of 40 percent, then vetoes and supermajorities stall all voting, and logrolling has no role because the majority is always winning.

The paper contrasts with previous models on SV [5, 6], which studied environments where the strategic problem faced by the voters is simpler. More precisely, existing models assume that voters have private information about their cardinal intensities of preferences, and that intensities are uncorrelated across voters. In such a scenario, a voter’s optimal strategy is to cast more votes on issues that the voter consider higher priorities (at a given state). Because the intensity of one’s own preferences is instinctively focal, it is natural to ask whether the good performance of the mechanism in the lab extends to more complicated strategic settings. If intensities are commonly known, or if it is known that intensities are correlated across groups, the minority’s need to be unpredictable returns and with it the more complex equilibrium with randomization.⁵ In this paper, we abstract from cardinal intensities and assume that each issue is judged equally important by all. The assumption can be read literally, reflecting a lack of clear priorities. But more generally, it is the modeling device we employ to give full weight to the strategic complexity of the hide-and-seek game. As a result of this modeling choice, one could argue that minority victories are not justified on normative utilitarian terms. Such a perspective, however, would be very narrow: in the absence of different intensities, the fairness requirement of *some* minority representation can be easily derived from a social welfare function that is concave in individual utilities, with the degree of concavity mirroring the strength of the social planner’s concern with equality [22, 23].⁶

Our experimental results also bear comparison to a small recent experimental literature on the asymmetric Colonel Blotto game. In line with Avrahami and Kareev [2] and Chowdhury et al. [7], we observe that the minority concedes some battlefields in order to win others. However, the key difference in our setting is the decentralization of decisions in the non-communication treatment, which renders the game more complex. Rogers [30] introduces some decentralization in a related game, whose payoffs differ from classical Blotto payoffs along several dimensions.⁷ One side consists of two players fighting against a single opponent, a structure that we examine in one of our treatments. Contrary to the conclusions of that paper, we observe that decentralization need not be detrimental to the divided side.

2 The Model

A committee of N individuals must resolve $K \geq 2$ binary issues: they must decide whether to pass or fail each of K independent proposals. The set of issues is denoted by $\mathcal{K} = \{1, \dots, K\}$. The same M individuals are in favor of all proposals, and the remaining $N - M = m$ are opposed to all, with $m \leq M$. We call M the *majority group*, and m the *minority group*, and we use the symbol M (m) to denote both the group and the number of individuals in

⁵In a related symmetric game, for example, Hortala-Vallve and Lorente-Saguer (2012) show that when priorities are known, pure strategy equilibria can exist only under very restrictive conditions.

⁶As pointed out in these papers, a normative basis for fairness also arises from individual utility functions which are concave with respect to the individual frequency of wins.

⁷Some battlefields are easier to win for one side, some for the other side; a bonus is added for the side winning a majority of battlefields; a bonus (resp. malus) is added for each winning (resp. losing) battlefield according to the margin of victory (resp. defeat).

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

Each voter is endowed with K votes, that she can distribute freely among the different issues. Each issue is decided according to the majority of *votes* cast — which, crucially, can differ from the majority of *voters*. Voting on the K issues is contemporaneous, and all individuals vote simultaneously. Ties are resolved by a fair coin toss. Each voter receives utility 1 from any issue resolved in her preferred direction, and 0 otherwise. Thus each voter’s goal is to maximize the fraction of issues resolved according to her — and her group’s — preferences.

With undominated strategies voters vote sincerely: they never cast a vote against their preferences. We simply assume that all m voters never vote in favor of a proposal and all M voters never vote against. We focus instead on each voter’s distribution of votes among the K issues.

The action space for each player is: $S(K) = \left\{ s = (s_1, \dots, s_K) \in \mathbb{N}^K \mid \sum_{k=1}^K s_k = K \right\}$, where s_k is the number of votes cast on issue k . Let the minority players be ordered from 1 to m . For each minority-profile $\mathbf{s} = (s^1, \dots, s^m) \in S(K)^m$, the number of votes allocated by the minority to issue k is: $v_k^m(\mathbf{s}) = \sum_{i=1}^m s_k^i$. We denote by $v^m(\mathbf{s}) = (v_k^m(\mathbf{s}))_{k \in \mathcal{K}} \in S(mK)$ the allocation of votes by the minority group. Similarly, for each majority-profile $\mathbf{t} = (t^1, \dots, t^M) \in S(K)^M$, the number of votes allocated by the majority to issue k is: $v_k^M(\mathbf{t}) = \sum_{i=1}^M t_k^i$ and we denote by $v^M(\mathbf{t}) = (v_k^M(\mathbf{t}))_{k \in \mathcal{K}} \in S(MK)$ the allocation of votes by the majority group. For a given profile $(\mathbf{s}, \mathbf{t}) \in S(K)^m \times S(K)^M$, the payoffs for each member of the two groups, called g_m and g_M , are given by

$$g_m(\mathbf{s}, \mathbf{t}) = \frac{1}{K} \sum_{k=1}^K \left(\mathbf{1}_{\{v_k^m(\mathbf{s}) > v_k^M(\mathbf{t})\}} + \frac{1}{2} \mathbf{1}_{\{v_k^m(\mathbf{s}) = v_k^M(\mathbf{t})\}} \right)$$

$$g_M(\mathbf{s}, \mathbf{t}) = \frac{1}{K} \sum_{k=1}^K \left(\mathbf{1}_{\{v_k^M(\mathbf{t}) > v_k^m(\mathbf{s})\}} + \frac{1}{2} \mathbf{1}_{\{v_k^M(\mathbf{t}) = v_k^m(\mathbf{s})\}} \right) = 1 - g_m(\mathbf{s}, \mathbf{t}),$$

where $\mathbf{1}$ is the indicator function.

Finally, we denote by $\Sigma(K) = \Delta(S(K))$ the set of all probability measures on $S(K)$, i.e. the set of mixed strategies. Then the expected payoff to the minority $E[g_m]$ equals p_m , the expected fraction of minority victories, and is defined on $\Sigma(K)^m \times \Sigma(K)^M$ as the multi-linear extension of g_m . Two (mixed strategy) group profiles $(\boldsymbol{\sigma}, \boldsymbol{\tau}) \in \Sigma(K)^m \times \Sigma(K)^M$ naturally define two probability measures (V^m, V^M) on the minority and majority allocations of votes $(v^m, v^M) \in S(mK) \times S(MK)$. Formally, our scenario corresponds to a decentralized Blotto (DB) game, in contrast to the traditional, centralized Colonel Blotto (CB) game, in which the “minority” colonel directly chooses $v^m \in S(mK)$, while the “majority” colonel chooses $v^M \in S(MK)$.

3 Two Preliminary Remarks

With incentives fully aligned within each group, a natural question is whether the decentralized Blotto game actually differs from the centralized game. We provide a positive answer in our first remark. We say that an equilibrium of the CB game is *replicated* in the DB game if there exists an equilibrium of the DB game which induces the same distribution on the total minority and majority allocations of votes (v^m, v^M) . The most complete characterization of equilibria of the CB game with discrete allocations is due to Hart [17].⁸

Remark 1 *For any K and m , none of the equilibria of the CB game in Hart [17] can be replicated in the DB game if M is larger than a finite threshold $\bar{M}(K)$.*

Nevertheless, in many applications, the assumption of no communication may be too strong. Consider then a modification of the model above where, before casting votes, each voter can exchange messages freely with all other members of her group. The messages are costless and non-binding (they are cheap talk), and we impose no constraint on their content. With communication, the logic behind Remark 1 breaks down.

Remark 2 *With communication, any equilibrium of the centralized Colonel Blotto game can be replicated. Other equilibria exist, including chattering equilibria replicating the equilibria of the no-communication game.*⁹

In this paper we study two different versions of the game, without and with communication. The first version is new and we derive results for that game, that we then use as reference for our experimental treatment without communication. The equilibria of the CB game in Hart [17] provide the theoretical benchmark for the second treatment, with communication.

4 Theory: no communication

4.1 Equilibria

The game is a normal-form game with $m + M$ players and finite strategy spaces. Therefore, a Nash equilibrium always exists. In addition, it is easy to see that the voting rule fulfills its fundamental purpose: if the size of the two groups is not too different, the smaller one *must* win occasionally.

Theorem 1 *If $M < m + K$, the expected share of minority victories is strictly positive at any Nash equilibrium.*

⁸Hart [17] does not characterize optimal strategies for all parameter values. Roberson [29] provides general results for the CB game with continuous allocations. In our problem, we did not see obvious advantages from abandoning the more realistic case of discrete votes.

⁹Other types of equilibria exist too. For example, asymmetric equilibria in which communication is ignored by one group but not by the other, and thus one group coordinates its strategy while the other does not.

The coordination problem within each of the two groups results in many equilibria. We do not aim to characterize them all, we rather focus on equilibria that have a simple enough structure to provide a plausible theoretical reference for the experiment. We delegate to the appendix (Section A.3) a discussion on pure-strategy equilibria arising in that game.

If several minority members concentrate votes on a given issue, the minority may be able to win it. But only if the majority does not know which specific issue is being targeted. Thus, minority members need not only to concentrate their votes but also to randomly choose the issues on which the votes are concentrated. We introduce a family of simple strategies that achieve this goal. For any c factor of K , we define the strategy σ^c (noted τ^c for a majority player) as follows: choose randomly K/c issues,¹⁰ and allocate c votes to each of the selected issues. Suppose for example $K = 4$, a value we will use in the experiment. Then σ^4 corresponds to casting all four votes on one single issue, chosen randomly; σ^2 to casting two votes each on two random issues; σ^1 to casting one vote on each of the four issues. Note that, in this family, the parameter c can be interpreted as the degree of concentration of a player's votes. We denote by σ^c (resp. τ^c) the subgroup profile for which each minority (resp. majority) player plays σ^c (resp. τ^c).

Intuitively, we expect the minority to concentrate its votes, so as to achieve at least some successes, and the majority to spread its votes, because its larger size allows it to cover, and win, a larger fraction of issues. The intuition is confirmed by the following two propositions, characterizing parameter values for which strategy profiles with such features are supported as Nash equilibria: when the difference in size between the two groups is as small as possible — either nil or one member — or when it is very large.

Proposition 1 *Suppose K even and M is odd. Then (σ^2, τ^1) is an equilibrium if $M \leq m + 1$,¹¹ with*

$$p_m = \begin{cases} \frac{1}{2} & \text{if } M = m \\ \frac{1}{2} - \frac{1}{2^{m+1}} \binom{m}{m/2} & \text{if } M = m + 1. \end{cases}$$

When the difference in size between the two groups is larger, we expect minority members to concentrate their votes even further. Indeed, as the next result shows, at large M/m there exist equilibria in which each minority voter concentrates all of her votes on a single issue. Majority voters continue to spread their votes.

Proposition 2 *Suppose M is divisible by K . Then (σ^K, τ^1) is an equilibrium if and only if $M \geq \frac{mK}{2}$. In such an equilibrium:*

$$p_m = \begin{cases} \sum_{p=M/K+1}^m \binom{m}{p} \frac{(K-1)^{m-p}}{K^m} + \frac{1}{2} \binom{m}{M/K} \frac{(K-1)^{m-M/K}}{K^m} & \text{if } M \leq mK \\ 0 & \text{if } M > mK. \end{cases}$$

¹⁰I.e. choose each subset of K/c issues with equal probability $1/\binom{K}{K/c}$.

¹¹The strategies in the proposition are also an equilibrium if $M \geq 2m + K - 1$. This is a “trivial” equilibrium in which the majority's much larger size allows it to win all proposals ($p_m = 0$). For $K \geq 4$ and $M < 2m + K - 1$, one can show that (σ^2, τ^1) is an equilibrium *if and only if* $M \leq m + 1$.

4.2 Beyond equilibrium: positive minority payoff with concentration and randomization

The equilibrium strategies characterized in Propositions 2 and 3 combine features that appear very intuitive (concentration and randomization for minority voters; less concentration for majority voters) with others that are most likely difficult for players to identify (the exact number of issues to target, the exact division of votes over such issues), or to achieve in the absence of communication (the symmetry of strategies within each group). The question we ask in this section is how robust minority victories are to deviations from equilibrium behavior in these last two categories.

We introduce a definition of *neutrality* of a strategy to capture the randomization across issues. The notion of neutrality is appealing in this game because the issues are identical ex-ante.

Definition 1 *A strategy σ is said to be neutral if for any permutation of the issues π and any allocation $s \in S(K)$, we have: $\sigma(s) = \sigma(s_\pi)$, where $s_\pi = (s_{\pi(1)}, \dots, s_{\pi(K)})$.*¹²

We assume that each minority voter concentrates her votes on a subset of issues, chosen randomly and with equal probability. However, we do not specify the precise number of issues targeted, do not require that votes be divided equally over such issues, and do not impose symmetry within the minority group. In addition, we evaluate the probability of minority victories by allowing for a worst-case scenario in which the majority jointly best responds. We find that the probability of minority victories is surprisingly robust.

Proposition 3 *For all $M \leq mK$, there exists a number $\underline{k} \in \{1, \dots, K\}$ such that if every minority player's strategy: (i) is neutral, and (ii) allocates votes on no more than \underline{k} issues with probability 1, then for any strategy profile of the majority τ ,*

$$p_m(\sigma, \tau) > 0.$$

The result of Proposition 3 is important because it is very broad, and its wide scope makes us more optimistic about the voting rule's realistic chances of protecting the minority. The game is complex, and robustness to deviations from equilibrium behavior should be part of the evaluation of the voting rule's potential. The result will indeed play a role in explaining our experimental data. In this particular game, studying deviations from equilibrium is made easier by the intuitive salience of some aspects of the strategic decision (concentration and randomization), and the much more difficult fine-tuning required by optimal strategies (*how many issues? How many votes?*).¹³ To complement the result, we provide in the appendix (Section A.7) an assessment of the magnitude of the minority payoff, showing that it can be significant.

¹²Note that neutrality does not require that votes be cast in equal number on each issue.

¹³Note, for comparison, that Proposition 3 holds under the identical condition $M \leq mK$ for the centralized game (with both discrete and continuous allocations).

5 The Experiment

5.1 Protocol

We designed the experiment to focus on two treatment variables: the size of the two groups, m and M , and the possibility of communication within each group. Each experimental session¹⁴ consisted of 20 rounds with fixed values of m and M ; the first ten rounds without communication, and the second ten with communication.

At the start of each session, each subject was assigned a color, either Blue or Orange, corresponding to the two groups. Members of the two groups were then randomly matched to form several committees, each composed of m Orange members and M Blue members. Every committee played the following game. Each subject entered a round endowed with K balls of her own color. She was asked to distribute them as she saw fit among K urns, depicted on the computer screen, knowing that she would earn 100 points for each urn in her committee in which a majority of balls were of her color. In case of ties, the urn was allocated to either the Blue or the Orange group with equal probability. Figure 4 in the appendix reproduces the relevant computer screen in one of our treatments for a Blue voter who has already cast one ball.

After all subjects had cast their balls, the results appeared on the screen under each urn: the number of balls of each color in the urn, the tie-break result if there was a tie, and the subject's winnings from the urn (either 0 or 100). The session then proceeded to the next round. The first ten rounds were all identical to the one just described. Subjects kept their color across rounds, but committees were reshuffled randomly. After the first round, subjects could consult the history of past decisions before casting their balls. By clicking a History button, each subject accessed a screen summarizing ball allocations and outcomes in previous rounds, by urn, in the committee that in each round included her.

After ten rounds, the session paused and new instructions were read for the second part. Parameters and choices remained unchanged and subjects kept the same color, but now a chatting option was enabled: before casting their balls, subjects had two minutes to exchange messages with other members of their committee who shared their color. They could consult the history screen while chatting. The second part of the session again lasted ten rounds, and again committees were reshuffled after each round but subjects kept the same color.¹⁵ Thus each subject belonged to the same group, m or M , for the entire length of the session, a design choice we made to allow for as much experience as possible with a given role. Each session lasted about 75 minutes, and earnings ranged from \$18 to \$44,

¹⁴All sessions were run at the Columbia Experimental Laboratory for the Social Sciences (CELSS) in April and May 2015, with Columbia University students recruited from the whole campus through the laboratory's Orsee site [14]. No subject participated in more than one session. In the laboratory, the students were seated randomly in booths separated by partitions; the experimenter then read aloud the instructions, projected views of the relevant computer screens, and answered all questions publicly. Two unpaid practice rounds were run before starting data collection.

¹⁵In all sessions we ran first the ten rounds without the chat option to prevent subjects from learning a coordinated strategy in the first part of the session and then trying to replicate it in the second — in the absence of communication. Subjects used the two minutes available for chatting fully, but we have no indication that they found the time too short.

with an average of \$33 (including a \$10 show-up fee).

We designed the experiment with two goals in mind. First, we wanted to learn how substantive are minority victories in the lab and how well the theory predicts subjects' behavior. We chose three sets of parameters: $(m, M) \in \{(1, 2); (2, 3); (2, 4)\}$, to make vary the absolute and relative sizes of the two groups, while keeping the size of the committee small — to maintain the possibility of conscious strategic choices by inexperienced players. Second, we wanted to compare results with and without communication. We have thus six treatments, denoted by mMD without chat, and mMC with chat.

Sessions	m, M	# Subjects	# Committees	# Rounds (no chat, chat)
s1, s2, s3	1, 2	12×3	4×3	10, 10
s4, s5, s6	2, 3	15×3	3×3	10, 10
s7, s8, s9	2, 4	18×3	3×3	10, 10

Table 1: Experimental Design

5.2 Parameter values and theoretical predictions

The theoretical predictions for the treatments are summarized in Table 2. The predicted strategies are reported for the no-chat treatments; they follow¹⁶ the results of Section 4.1. The predictions for the chat treatments are based on the results of Hart [17]. The precise equilibrium strategies are described in the appendix (Section A.9).

Treatment	Simple symmetric equilibrium	p_m	Treatment	p_m
12D	(σ^2, τ^1)	1/4	12C	1/4
23D	(σ^2, τ^1)	1/4	23C	1/3
24D	(σ^4, τ^1)	1/4	24C	1/4

Table 2: Predictions for the experimental treatments

6 Experimental Results

6.1 Minority victories

We report in Figure 1 the realized fractions of minority victories in the six treatments — the percentage of urns won by an orange team. The orange columns correspond to the experimental data, and the grey columns to the theoretical equilibrium predictions.

¹⁶For our experimental parameters equilibria of the decentralized game exist in the family of simple profiles (σ^c, τ^d) , symmetric within groups, and within this family are unique. We select such equilibria as theoretical reference for the experiment because of their intuitive simplicity.

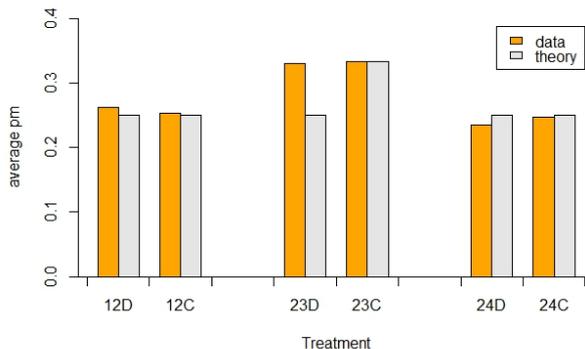


Figure 1: Fractions of minority victories

Whether with or without communication, the fraction of minority victories in the data is non-negligible, ranging from a minimum of 0.24 (in treatment 24D) to a maximum of 0.33 (in treatment 23C). Even more remarkable, realized values are very close to the theoretical predictions, although the difference is more sizable in treatment 23D.¹⁷

6.2 Strategies (no communication)

In the absence of communication, equilibrium strategies are defined at the individual level. Figure 2 reports the observed frequency of different ball allocations, across individual subjects, in the treatments without communication. The horizontal axis lists all possible allocations — with four balls and four urns there are five — and the vertical axis reports the frequency of subjects choosing the corresponding allocation, over all rounds, committees, and sessions of the relevant treatment.¹⁸ The panels are organized in two rows, corresponding to the two groups, with the minority in orange in the upper row, and the majority in blue in the lower row. The allocation denoted in bold and surrounded by two stars, on the horizontal axis, corresponds to the equilibrium strategy in Table 2.

The figure teaches three main lessons. First, there is substantial deviation from equilibrium strategies: in all treatments and in both groups, at least forty percent of all individual allocations do not correspond to equilibrium strategies. However — and this is the second lesson — equilibrium predictions have some explanatory power for minority subjects. In all treatments, the most frequently observed allocation for minority subjects corresponds to the equilibrium strategy, a particularly clear result in treatment 12D and 24D, where more than half of all observed allocations correspond to the predictions.¹⁹ Equilibrium predictions are

¹⁷The difference is not statistically significant. In treatment 23D there is an asymmetric equilibrium in which $p_m = 11/32 \approx 0.34$ (v/s 0.33 in the data): all m members play σ^4 , one M member plays τ^1 , and two play τ^2 . However, we do not see this equilibrium in the data. As mentioned above, random rematching at each round means that subjects in general cannot coordinate on an asymmetric equilibrium.

¹⁸Thus, for example, the column corresponding to “0112” reports the frequency of subjects casting two balls in one urn, and one ball each in two other urns.

¹⁹This need not be a best response, given the variability in the data and the more random behavior of

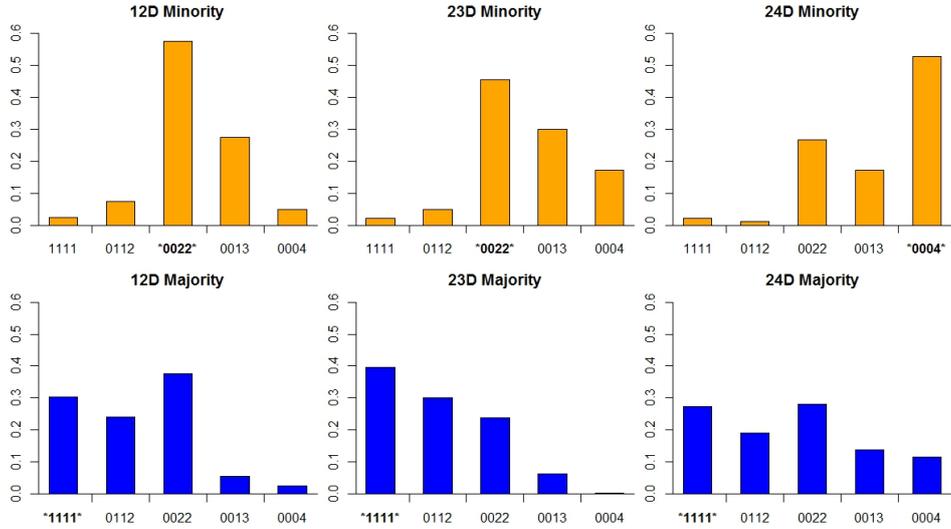


Figure 2: Frequency of individual ball allocations (no-chat treatments)

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

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

A finer analysis of voters' strategies at the individual level is provided in the appendix (Section A.10), it confirms the main observations: *most* minority subjects concentrate their votes *on all rounds*, and we observe more concentration in treatment 24D.

6.3 Further results

A description of group's allocations in the treatments with communication is given in the appendix (Section A.11). We observe that subjects do use communication to coordinate their strategies within a group. Moreover, groups' allocations are consistent with Nash equilibrium for small committees (treatment 12C), but this hypothesis is clearly rejected for larger committees (treatments 23C and 24C).

The main result of the theory (Proposition 3) states that minority victories can be guaranteed, if minority voters concentrate their votes, and do so unpredictably. Although the concentration of allocations can be directly observed, the unpredictability of these allocations is not easy to measure — with rematching after each round, the serial correlation

majority subjects.

of allocations is not a good proxy for predictability. In the appendix (Section A.12), we assess it indirectly, by computing each group’s guaranteed payoff at each session — this is the payoff a group would have obtained, had the opposite group best replied to its strategy (over the whole session). In our hide-and-seek game, a predictable allocation of the minority group has a guaranteed payoff of 0. On the contrary, we observe in each treatment that the minority guaranteed a positive and significant payoff, ranging from 0.16 in treatment 24D to 0.28 in treatment 23C. This confirms that minority strategies were indeed unpredictable.

7 Conclusions

We have investigated the ability of the SV mechanism to protect a minority in a fully polarized committee. Both in theory and in a laboratory experiment, we find that the mechanism is effective: in line with equilibrium predictions, the fraction of minority victories observed in the experiment varied from 25 percent in treatments in which the minority is half the size of the majority, to 33 percent, when the minority’s relative size increases to two thirds. Allowing voters to communicate before casting their votes does not alter our conclusions.

A surprising aspect of our results is that experimental outcomes closely replicate the theoretical predictions even though subjects often deviate from equilibrium strategies. The reason is that the fundamental logic of the game — its *hide-and-seek* nature, requiring minority voters to concentrate their votes and to do so unpredictably — seems to be immediately clear to the experimental subjects. We see this in the experimental results, and we establish it theoretically by studying the robustness of predicted outcomes to plausible off-equilibrium behavior: as long as each minority voter concentrates her votes sufficiently and randomizes the target issues, minority victories are guaranteed (in expectations). The conclusion holds even if the number of target issues is not optimal, even if other minority voters choose different degrees of concentration, and even if majority voters coordinate their strategies and best-respond. We interpret this result as an encouraging check on the robustness of the voting mechanism and on its potential to overcome the tyranny of the majority in realistic applications. SV treat all individuals equally, avoid the inertia and obstruction of supermajority rules and vetoes, and yet ensure that the minority voice is heard, even in the difficult strategic environment studied here.

A Appendix

A.1 Proof of Remark 1

With the parameters of our model, the number of votes of each group is a multiple of the number of issues, K . In that case, optimal strategies for the majority are identified in Hart [17] if K is even and/or M is odd (by combination of his Theorem 4 and Proposition 6). They are such that the marginal distribution of majority votes on each issue is uniform over a set of consecutive odd integers: $\forall k \in \mathcal{K}, V_k^M \sim \mathcal{U}(\{1, 3, \dots, 2M - 1\})$.

Let us assume that this strategy is replicated by M independent lieutenants. We denote by S^i the allocation of lieutenant i on issue 1. We have:

$$\sum_{i=1}^M S^i = V_1^M \sim \mathcal{U}(\{1, 3, \dots, 2M - 1\}).$$

As we have $\forall i = 1 \dots M, 0 \leq S^i \leq K$ and $\mathbb{E}[V_1^M] = M$, we obtain by Hoeffding's inequality [18]:

$$\mathbb{P}(V_1^M - M \geq M - 1) = \frac{1}{M} \leq \exp\left(\frac{-2(M - 1)^2}{MK^2}\right).$$

This inequality can be written $Me^{-\frac{2(M-1)^2}{MK^2}} \geq 1$, which is equivalent to

$$K \geq \frac{\sqrt{2}(M - 1)}{\sqrt{M \log(M)}} := \bar{K}(M).$$

Hence, we get a contradiction if $K < \bar{K}(M)$. As we have $\frac{\partial \bar{K}}{\partial M} > 0$, the function \bar{K} is one-to-one, and we denote its inverse by $\bar{M}(K) := \bar{K}^{-1}(K)$. As \bar{M} is increasing, we have a contradiction if $M > \bar{M}(K)$. \square

A.2 Proof of Theorem 1

Consider a profile of (possibly mixed) strategies such that the majority wins all the decisions with probability one. Consider any pure-strategy profile (\mathbf{s}, \mathbf{t}) played with positive probability. Consider any minority player i . For each issue $k \in \mathcal{K}$, let $b_k = v_k^M(\mathbf{t}) - v_k^m(\mathbf{s})$ be the margin (bias) by which the majority beats the minority on issue k , and let s_k^i be the number of votes allocated by i to issue k . As the average of the $(b_k)_{k \in \mathcal{K}}$ is $M - m$, while the average of the $(s_k^i)_{k \in \mathcal{K}}$ is one, it follows that the average of the numbers $(b_k + s_k^i)_{k \in \mathcal{K}}$ is $M - m + 1$. There must be an issue $k' \in \mathcal{K}$ such that:

$$b_{k'} + s_{k'}^i \leq M - m + 1.$$

Subtracting K from both sides:

$$b_{k'} - (K - s_{k'}^i) \leq M - m + 1 - K.$$

The term $(K - s_{k'}^i)$ captures the amount by which i 's votes on k' fall short of the maximum possible, K . The left-hand side of the inequality equals the majority's vote margin on k' when i allocates all her votes to k' . But if $M < m + K$, $M - m + 1 - K \leq 0$, and the majority cannot be winning with probability one. Either $s_{k'}^i = K$, and we have obtained a contradiction. Or $s_{k'}^i < K$, and i has a profitable deviation; but then the initial profile is not an equilibrium. \square

A.3 Equilibria in pure strategies

We begin by remarking that the condition in Theorem 1 is tight: if $M \geq m + K$, the profile of strategies such that every player allocates one vote per issue is an equilibrium, and the expected share of minority victories is zero. This same profile of strategies is also an equilibrium if $M = m$, in which case $p_m = 1/2$. More generally, we establish the existence of an equilibrium in pure strategies when the committee is large enough.

Proposition 4 *If $M \geq m \geq 2$ and $M + m \geq (K + 1)^2/K$, a pure-strategy equilibrium always exists.*

Proof. Assume that $M \geq m \geq 2$ and $M + m \geq (K + 1)^2/K$. We construct a pure-strategy equilibrium for the DB game, based on a partition of the set of issues $\mathcal{K} = \mathcal{K}_m \cup \mathcal{K}_M$. We note $K_m = \#\mathcal{K}_m$ and $K_M = \#\mathcal{K}_M = K - K_m$.

Step 1. There exists a partition of the set of issues $\mathcal{K} = \mathcal{K}_m \cup \mathcal{K}_M$ satisfying $K_m \in \left[\max \left(K - \left\lfloor \frac{MK}{K+1} \right\rfloor, 1 \right), \min \left(\left\lfloor \frac{mK}{K+1} \right\rfloor, K - 1 \right) \right]$.

As $M \geq m \geq 2$, it is immediate that

$$\left\lfloor \frac{mK}{K+1} \right\rfloor \geq 1 \quad \text{and} \quad K - \left\lfloor \frac{MK}{K+1} \right\rfloor \leq K - 1.$$

As $M + m \geq (K + 1)^2/K$, we get $\frac{mK}{K+1} \geq K + 1 - \frac{MK}{K+1}$, and therefore

$$\left\lfloor \frac{mK}{K+1} \right\rfloor \geq K - \left\lfloor \frac{MK}{K+1} \right\rfloor.$$

Step 2. For any such partition, any pure-strategy profile (\mathbf{s}, \mathbf{t}) for which

$$\begin{cases} v_k^m(\mathbf{s}) \geq \left\lfloor \frac{mK}{K_m} \right\rfloor & \text{if } k \in \mathcal{K}_m \\ v_k^m(\mathbf{s}) = 0 & \text{otherwise} \end{cases} \quad \begin{cases} v_k^M(\mathbf{t}) \geq \left\lfloor \frac{MK}{K_M} \right\rfloor & \text{if } k \in \mathcal{K}_M \\ v_k^M(\mathbf{t}) = 0 & \text{otherwise} \end{cases}$$

is an equilibrium. In such equilibria, $p_m = K_m/K$.

As $K_m \leq \left\lfloor \frac{mK}{K+1} \right\rfloor \leq \frac{mK}{K+1}$, we have $K \leq \frac{mK}{K_m} - 1$, which leads to

$$K < \left\lfloor \frac{mK}{K_m} \right\rfloor.$$

We conclude that a majority player cannot upset the outcome of an issue in \mathcal{K}_m : she has no profitable deviation.

As $K_m \geq K - \left\lfloor \frac{MK}{K+1} \right\rfloor$, we have $K_M \leq \left\lfloor \frac{MK}{K+1} \right\rfloor$. We conclude as before that no minority player has a profitable deviation. ■

This result clearly indicates that the DB game differs from the CB game, in which pure-strategy equilibria generically fail to exist.²⁰ The equilibria we construct are such that the two groups target different issues: the majority only votes on a subset \mathcal{K}_M of issues, while the minority votes on the remaining subset $\mathcal{K}_m = \mathcal{K} \setminus \mathcal{K}_M$. As each voter is small in a large committee, no voter can upset the outcome of any given issue, and thus gain from deviating.

We note one surprising effect of decentralization: in these equilibria, it is possible for the minority to win *more frequently* than the majority, whereas no such outcome exists in the CB game.

Example 1 *If $m = 4$, $M = 5$ and $K = 3$, there exists an equilibrium in which the minority wins two of the three issues.*

We also note that pure-strategy equilibria may not exist for small committees. The following example describes a parametrization we use in the experiment.

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

Proof. Note first that since $3 < 25/4$, Proposition 4 does not apply. Consider an arbitrary pure-strategy profile (\mathbf{s}, \mathbf{t}) . For each issue $k \in \mathcal{K}$, let $b_k = v_k^M(\mathbf{t}) - v_k^m(\mathbf{s})$, so that $\frac{1}{4} \sum_{k=1}^4 b_k = 1$. Assume for simplicity that $b_1 \leq b_2 \leq b_3 \leq b_4$.

We first remark that, if there is a tie, a majority player deviates. Assume that for some issue k , $b_k = 0$. Then, there must be some issue j for which $b_j \geq 2$. At least one majority player can withdraw a vote from issue j and allocate it to issue k . This is a profitable deviation.

We distinguish four cases:

- (a) $p_m = 0$. The average number of majority votes per proposal is 2. Thus, the minority player can win a decision by allocating all her votes to an issue with no more than 2 majority votes. This is a profitable deviation.

²⁰In the CB game, the profile for which every player allocates one vote per issue is an equilibrium only when $M = m = 1$ or $M > mK$. Beyond these special cases, if $K > 2$, the CB game has no equilibria in pure strategies. A pure-strategy equilibrium may exist in a non-zero sum variant in which the two sides attribute heterogeneous and asymmetric values to the different issues [19].

- (b) $p_m = 1/4$ and $b_1 \leq -3$. The average number of majority votes on issues 2,3,4 is at most $8/3$. Thus, one of these issues (say $k = 2$) receives no more than 2 majority votes. The minority player can withdraw 2 votes from issue 1, and allocate them to issue 2 to obtain a tie. This is a profitable deviation.
- (c) $p_m = 1/4$ and $b_1 \geq -2$. The average of the $(b_k)_{k=2\dots4}$ on the issues won by the majority is at least 2. This means that one of the two majority players can withdraw 2 votes from issues 2,3,4 at no cost. By allocating these 2 votes on issue 1, she obtains at least a tie. This is a profitable deviation.
- (d) $p_m \geq 1/2$. The minority wins issues 1 and 2. The average of the $(b_k)_{k=3\dots4}$ is at least 3. A majority player can withdraw 2 votes from issues 3 and 4, and cast the 2 votes on the issue with the lowest number of minority votes among 1 and 2. On this issue, there cannot be more than 2 minority votes, so the majority player obtains at least a tie. This is a profitable deviation.

■

The fact that, unexpectedly, pure strategy equilibria may exist is interesting. How empirically plausible they are, however, is open to question. The equilibria obtained in Proposition 4 require a large extent of coordination, both within and across groups. In addition, not only in those equilibria, but also in the “trivial” equilibrium with $M \geq m + K$, each voter has only a weak incentive not to deviate. This seems particularly problematic when $M \geq m + K$ and the minority loses all decisions, under the equilibrium profile in which each player allocates one vote per decision. Even non-strategic minority members seem likely to realize that some concentration is called for.

A.4 Proof of Proposition 1

Deviations for an M -player We consider the point of view of an M -player, denoted by i . On each issue $k \in \mathcal{K}$, the total number of votes cast by the other M -players is $v_k^{M-1} = M - 1$. The total number of votes cast by the minority is denoted by v_k^m . The random variable $v_k^m/2$ follows a binomial distribution of parameters m and $\frac{1}{2}$.

Let a_k^i be the number of votes cast by voter i on issue k , and $p_k^i(a_k^i)$ the payoff of i on this issue:²¹

$$p_k^i(a_k^i) = \mathbb{P}(v_k^{M-1} + a_k^i > v_k^m) + \frac{1}{2}\mathbb{P}(v_k^{M-1} + a_k^i = v_k^m).$$

In what follows, we omit to mention the subscript k in the computations, as all the strategies

²¹By convention, the payoff on issue k can take values between 0 and 1, and the overall payoff is the mean of the payoffs over all the issues.

are symmetric across decisions. As M is an odd number, we have for all $a \in \{1, \dots, K\}$:

$$\begin{aligned}
p^i(a) - p^i(a-1) &= \frac{1}{2}\mathbb{P}(v^m = M-1+a) + \frac{1}{2}\mathbb{P}(v^m = M-2+a) \\
&= \frac{1}{2}\mathbb{P}\left(\frac{v^m}{2} = \frac{M-1}{2} + \frac{a}{2}\right) + \frac{1}{2}\mathbb{P}\left(\frac{v^m}{2} = \frac{M-1}{2} + \frac{a-1}{2}\right) \\
&= \frac{1}{2}\mathbb{P}\left(\frac{v^m}{2} = \frac{M-1}{2} + \left\lfloor \frac{a}{2} \right\rfloor\right) \\
&= \frac{1}{2^{m+1}} \binom{m}{\frac{M-1}{2} + \lfloor \frac{a}{2} \rfloor} \mathbf{1}_{\{\frac{M-1}{2} + \lfloor \frac{a}{2} \rfloor \leq m\}}.
\end{aligned}$$

For any $a \in \{1, \dots, K\}$, we have $\frac{M-1}{2} + \lfloor \frac{a}{2} \rfloor \geq \frac{M-1}{2} \geq \frac{m}{2} - \frac{1}{2}$, this implies:

$$\binom{m}{\frac{M-1}{2} + \lfloor \frac{a}{2} \rfloor} \mathbf{1}_{\{\frac{M-1}{2} + \lfloor \frac{a}{2} \rfloor \leq m\}} \leq \binom{m}{\frac{M-1}{2}} \mathbf{1}_{\{\frac{M-1}{2} \leq m\}}.$$

Therefore $p^i(a) - p^i(a-1) \leq p^i(1) - p^i(0)$. It follows that τ^1 is a best reply for player i .

Deviations for an m -player. We consider a player j in team m . On a given decision, the payoff of j , playing $a \in \{0, \dots, K\}$ is:

$$p^j(a) = \mathbb{P}(v^{m-1} + a > v^M) + \frac{1}{2}\mathbb{P}(v^{m-1} + a = v^M)$$

where $v^M = M$ and $v^{m-1}/2$ is a random variable following a binomial distribution of parameters $(m-1)$ and $\frac{1}{2}$. As M is an odd number, we have:

$$\begin{aligned}
p^j(a) - p^j(a-1) &= \frac{1}{2}\mathbb{P}(v^{m-1} = M-a) + \frac{1}{2}\mathbb{P}(v^{m-1} = M+1-a) \\
&= \frac{1}{2}\mathbb{P}\left(\frac{v^{m-1}}{2} = \frac{M-1}{2} - \left\lfloor \frac{a-1}{2} \right\rfloor\right) \\
&= \frac{1}{2^m} \binom{m-1}{\frac{M-1}{2} - \lfloor \frac{a-1}{2} \rfloor} \mathbf{1}_{\{0 \leq \frac{M-1}{2} - \lfloor \frac{a-1}{2} \rfloor \leq m-1\}}.
\end{aligned}$$

In particular, $p^j(2) - p^j(1) = p^j(1) - p^j(0) = \frac{1}{2^m} \binom{m-1}{\frac{M-1}{2}} \mathbf{1}_{\{M \leq 2m-1\}}$.

As $M \leq m+1$, for any $a \in \{3, \dots, K\}$, we have $\frac{M-1}{2} - \lfloor \frac{a-1}{2} \rfloor \leq \frac{M-1}{2} \leq \frac{m-1}{2} + \frac{1}{2}$. Therefore, $p^j(a) - p^j(a-1) \leq p^j(2) - p^j(1) = p^j(1) - p^j(0)$. As a result, σ^2 is a best reply for j . \square

A.5 Proof of Proposition 2

We assume that $M \leq mK$. Indeed, if $M > mK$, the profile (σ^K, τ^1) is trivially an equilibrium in which the majority wins all the decisions.

Deviations for an M -player We write, as before, for any $a \in \{0, \dots, K\}$:

$$p^i(a) = \mathbb{P}(M-1+a > v^m) + \frac{1}{2}\mathbb{P}(M-1+a = v^m),$$

where v^m/K follows a binomial distribution of parameters m and $1/K$. We get:

$$p^i(a) - p^i(a-1) = \frac{1}{2}\mathbb{P}(v^m = M-1+a) + \frac{1}{2}\mathbb{P}(v^m = M-2+a).$$

As M is a multiple of K , it is the only one in the set $\{M-2, \dots, M-1+K\}$. As v^m must be a multiple of K , we obtain $p^i(2) - p^i(1) = p^i(1) - p^i(0) = \frac{1}{2}\mathbb{P}(v^m = M)$ and for all $a \in \{3, \dots, K\}$, $p^i(a) - p^i(a-1) = 0$. We conclude that τ^1 is a best reply for player i .

Deviations for an m -player We write as before, for $a \in \{0, \dots, K\}$:

$$p^j(a) = \mathbb{P}(v^{m-1} + a > M) + \frac{1}{2}\mathbb{P}(v^{m-1} + a = M),$$

where v^{m-1}/K follows a binomial distribution of parameters $(m-1)$ and $1/K$. We get:

$$p^j(a) - p^j(a-1) = \frac{1}{2}\mathbb{P}(v^{m-1} = M-a) + \frac{1}{2}\mathbb{P}(v^{m-1} = M+1-a).$$

There are two multiples of K in $\{M-K, \dots, M\}$, namely $M-K$ and M . We obtain:

$$\begin{aligned} p^j(1) - p^j(0) &= \frac{1}{2}\mathbb{P}(v^{m-1} = M) \\ \forall a \in \{2, \dots, K-1\}, \quad p^j(a) - p^j(a-1) &= 0 \\ p^j(K) - p^j(K-1) &= \frac{1}{2}\mathbb{P}(v^{m-1} = M-K). \end{aligned}$$

There are only two candidates for the best reply of voter j : playing one vote on every issue or playing K votes on a single issue. It follows that the strategy σ^K is a best reply for player j if and only if

$$\begin{aligned} p^j(K) + (K-1)p^j(0) \geq Kp^j(1) &\Leftrightarrow p^j(K) - p^j(1) \geq (K-1)(p^j(1) - p^j(0)) \\ &\Leftrightarrow \mathbb{P}(v^{m-1} = M-K) \geq (K-1)\mathbb{P}(v^{m-1} = M). \end{aligned}$$

We know that $v^{m-1} = M$ (resp $v^{m-1} = M-K$) if exactly M/K m -players (resp. exactly $M/K-1$ m -players) play K on the considered issue. Thus:

$$\begin{aligned} \mathbb{P}(v^{m-1} = M) &= \binom{m-1}{M/K} \left(\frac{1}{K}\right)^{M/K} \left(\frac{K-1}{K}\right)^{m-1-M/K} \\ \mathbb{P}(v^{m-1} = M-K) &= \binom{m-1}{M/K-1} \left(\frac{1}{K}\right)^{M/K-1} \left(\frac{K-1}{K}\right)^{m-M/K}. \end{aligned}$$

We obtain

$$\frac{\mathbb{P}(v^{m-1} = M-K)}{(K-1)\mathbb{P}(v^{m-1} = M)} = \frac{M/K}{m - M/K}.$$

The strategy σ^K is a best reply for player j if and only if this ratio is larger than or equal to 1, or equivalently $M \geq mK/2$. \square

A.6 Proof of Proposition 3

Assume that $M \leq mK$, and define $\underline{k} \equiv \lfloor \frac{Km}{M} \rfloor$. Note that $\underline{k} \in \{1, \dots, K\}$.

Let σ be a minority profile satisfying the two conditions of the proposition. For each player, and each allocation played with positive probability, there is at least one issue receiving at least $\frac{K}{\underline{k}}$ votes from this player. By symmetry across issues, each player allocates with positive probability at least $\frac{K}{\underline{k}}$ votes on each issue. As a result, each issue receives at least $\frac{mK}{\underline{k}}$ votes from the minority with positive probability.

Let τ be a majority profile and let v^M be a majority allocation played with positive probability. There exists at least an issue k receiving no more than M votes from the majority. Since $\underline{k} \leq \frac{Km}{M}$, it follows that $\frac{mK}{\underline{k}} \geq M$. Hence the minority wins the issue k with positive probability: $p_m(\sigma, \tau) > 0$. \square

A.7 Magnitude of the minority payoff

We can assess the magnitude of the minority payoff through simulations, under different assumptions over the rules followed by each minority and majority voter. As an example, we report here results obtained if the minority adopts the neutral σ^c strategies described in the previous section. We set $K = 4$, $M = 10$, and $m \in \{1, \dots, 10\}$, and consider two cases, with increasing concentration: $c = 2$ (each minority voter casts two votes each on half of the issues, chosen with equal probability), and $c = 4$ (each minority voter casts all votes on a single issue, again chosen randomly with equal probability). To establish plausible bounds on the frequency of minority victories, we consider two rules for the majority: either each majority voter casts his votes randomly and independently over all issues (an upper bound on p_m) or all majority voters together best respond to the minority rule (the lower bound).²² Figure 3 reports such bounds for each value of m (on the horizontal axis) under minority rules σ^2 (in blue) and σ^4 (in green).

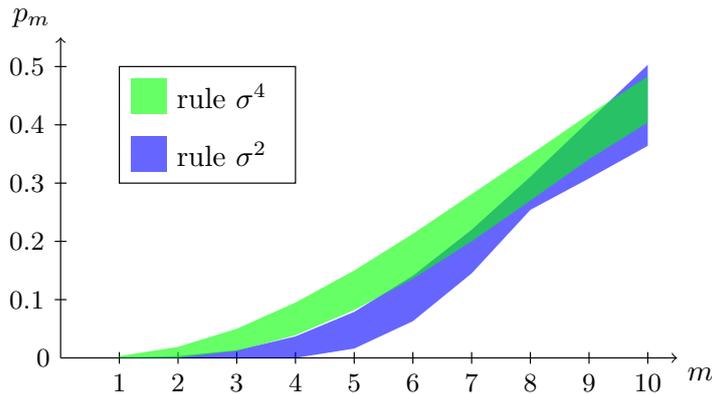


Figure 3: Minority payoffs for two minority rules ($M = 10$)

²²We compute p_m when the majority jointly best responds by considering all possible allocations of the MK majority votes, and then selecting the minimum p_m .

As expected, p_m increases with m . In addition, strategy σ^4 , allocating all votes on a single issue, outperforms σ^2 for all values of $m < M$. As long as $m > 2$ (a threshold that corresponds to the condition $M < mK$ in the proposition), σ^4 always results into a positive frequency of minority victories. Even for relatively large differences in size between the two groups, the expected fraction of minority victories is significant: in a range between 0.14 and 0.21 when $m = 6$, and between 0.20 and 0.28 when $m = 7$ (that is, when the minority is either 60 or 70 percent of the majority).

A.8 A screenshot

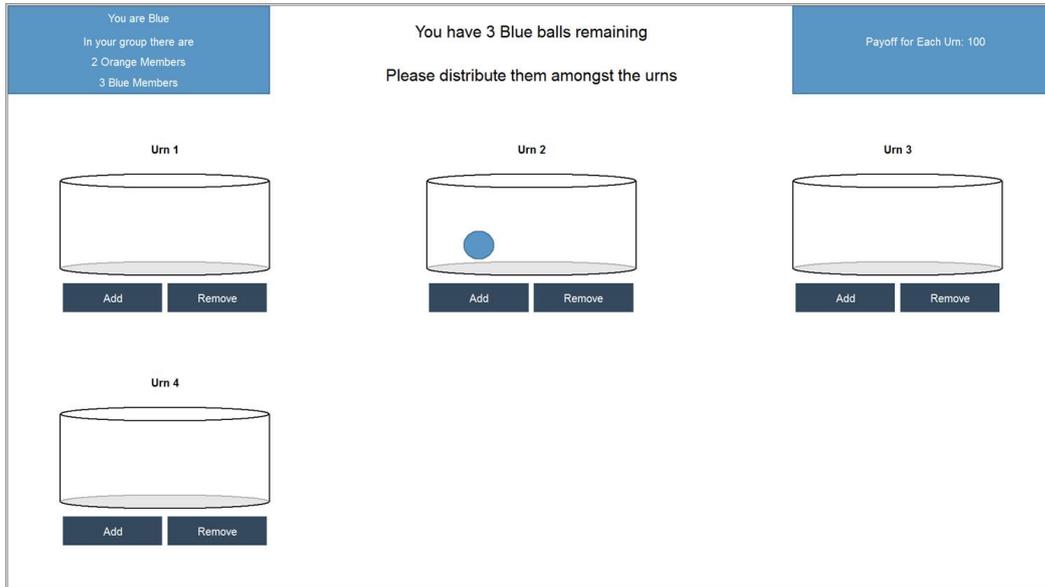


Figure 4: The Allocation screen

A.9 Predictions for the communication treatments

With communication within each group, we look at the equilibria of the centralized Blotto game. As established by Hart [17], with discrete allocations the value of the Blotto game (and thus p_m at equilibrium) is unique, but the optimal strategies are not, even in the special cases of our experimental parameters. And yet such strategies share a common intuitive structure. In the continuous Blotto game, where allocations need not be integer numbers, optimal strategies must be such that the marginal distribution of forces allocated to any one battlefield is uniform: M allocates to any urn a number drawn from a uniform distribution over $[0, 2M]$; m allocates to any urn either no balls, with probability $(1 - m/M)$, or a number of balls drawn from the uniform distribution on $[0, 2M]$ [29]. With integer numbers, the uniform requirement cannot be matched exactly, but is approximated. Using Hart's notation, we define as U_o^μ the uniform distribution over odd numbers with mean μ (i.e.

over $\{1, 3, \dots, 2\mu - 1\}$, U_e^μ the uniform distribution over even numbers with mean μ (i.e. over $\{0, 2, \dots, 2\mu\}$), and $U_{o/e}^\mu$ the convex hull of U_o^μ and U_e^μ (i.e. the set $\lambda U_o^\mu + (1 - \lambda)U_e^\mu$, for all $\lambda \in [0, 1]$). Table 3 reports the marginal allocations (on each urn) associated to Hart's optimal strategies for our experimental parameters, as well as p_m . Note that the optimal strategies in Hart [17] may not be unique; for example we identified new ones in the treatment 12C.²³

Treatment	Optimal strategies: marginal allocations	p_m
12C	m : $1/2\{0\} + 1/2(U_{o/e}^2)$; $1/2\{0\} + 1/2\{2\}$; any combination M : U_o^2 ; $\{2\}$; any combination	1/4
23C	m : $1/3\{0\} + 2/3(U_{o/e}^3)$ M : U_o^3	1/3
24C	m : $1/2\{0\} + 1/2(U_{o/e}^4)$ M : U_o^4	1/4

Table 3: Equilibria of the centralized game

The strategies can be implemented in different ways, as long as the equal probability restriction embodied by the marginal distribution is satisfied.²⁴ The important point of the table is that optimal strategies are such that the marginal distributions on the targeted urns must be uniform distributions or combinations of uniform distributions, for both groups, a relatively easy requirement to check on the experimental data.

A.10 Experimental strategies by individual subjects (no communication)

Our theoretical results establish that the minority can guarantee itself a positive expected fraction of victories, even when individual minority members follow different strategies, as long as each concentrates her votes on a sufficiently small subset of urns and casts them randomly. According to Proposition 3, on not more than \underline{k} urns, where $\underline{k} = 2$ in all our experimental treatments. We look in more detail at the subjects' behavior in the lab, keeping this result in mind.

Figure 5 plots individual subjects' average ball allocations in the three treatments with no communication. The vertical axis in the figure is the average largest number of balls cast in any one urn, a number that we denote by x_4 and that ranges from 1 to 4; the horizontal axis is the average second largest number, denoted by x_3 and ranging from 0 to 2. Each dot in the figure is a single subject's average ball allocation over the 10 rounds

²³The strategies involving $\{2\}$ in treatment 12C are not identified by Hart because they are not optimal strategies of the General Lotto game. See Hart [17].

²⁴For example, the majority strategy in 23C must correspond to mixing uniformly over $\{1, 3, 5\}$ for each urn, satisfying the budget constraint: in terms of specific allocations per urn, and keeping in mind that each urn is chosen with equal probability, one such strategy is $(1/3)(3, 3, 3, 3) + (2/3)(1, 1, 5, 5)$; another is $(2/3)(1, 3, 3, 5) + (1/3)(1, 1, 5, 5)$; in fact any combination of these two strategies also satisfies the requirement.

played, summarized by the subject's average x_4 and x_3 .²⁵ Orange dots denote members of the minority, and Blue dots members of the majority.

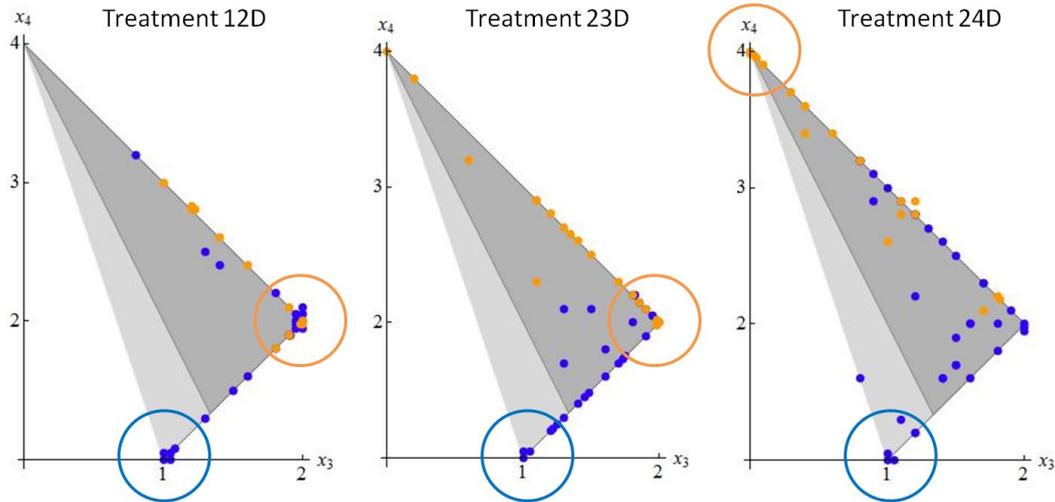


Figure 5: Individual subjects' average ball allocations (no-chat treatments)

The vertices of each triangle in the figure correspond to three feasible allocations: $(0, 4)$, at the upper end, corresponds to casting all balls in a single urn; $(1, 1)$, at the lower end, corresponds to casting one ball in each urn, and $(2, 2)$, at the right end, corresponds to dividing the balls equally over two urns.²⁶ In all three panels, the equilibrium strategy for majority subjects is the $(1, 1)$ vertex (marked by the large blue circle); for minority subjects it is the $(2, 2)$ vertex in the first two panels and the $(0, 4)$ one in the third (marked by the large orange circle).

The upper edge of the triangle, uniting $(0, 4)$ and $(2, 2)$, is the line segment described by $x_4 + x_3 = 4$, conditional on $x_4 \geq x_3$: all dots lying along this line represent subjects who in every round divided their balls over at most two urns. Dots lying to the interior of the line, on the other hand, represent subjects who in at least some rounds cast balls in more than two urns. The boundary between the two grey areas corresponds to the line segment $x_4 + 2x_3 = 4$, again conditional on $x_4 \geq x_3$. Dots below that line correspond to subjects who must have cast balls in all four urns in at least some rounds.

Figure 5 can now be read at a glance and reveals several regularities. First, in all three treatments, minority subjects almost unanimously concentrate balls in only two urns. Only 2 out of 12 minority subjects in treatment 12D, 2 out of 18 in 23D, and 3 out of 18 in 24D *ever* cast balls in more than two urns, and in 4 of these 7 cases the dots are close to the upper edge, implying that this occurred in a small number of rounds. Not only do minority

²⁵For instance, if a subject plays 0022 on half of the rounds, and 0004 on the other half, her average allocation will be represented with $x_4 = 3$ and $x_3 = 1$.

²⁶The other two possible allocations, 0013 and 0112, correspond to points $(1, 3)$ and $(1, 2)$ in the figure, and are, respectively, along the upper edge of the triangle, and along the line dividing the dark and light grey areas.

subjects follow the intuitive prescription of concentrating balls in a subset of urns; they also target not more than two urns. Second, there is much more variability in the number of target urns among majority subjects. In all treatments, a non-negligible number of subjects casts balls in all urns, but an equally large number casts balls in two or three urns only. A possible reading is that majority subjects are divided between exploiting their larger size by covering all urns (as equilibrium predicts), and second-guessing the minority, in the logic of the hide-and-seek game. The role of this latter motivation is supported by the right-most panel in Figure 5, and this is our third observation. Members of both groups tend to concentrate their balls more in treatment 24D: although again there is large variability, especially among majority subjects, the dots in the third panel tend to be shifted upward along the outer edge, relative to the dots in the first two panels, indicating that, among the two targeted urns, one is receiving an increasingly disproportionate share of balls. Minority members' incentive to concentrate their allocations more in treatment 24D is intuitive and could be the trigger for the majority subjects' own more frequent concentration.

A.11 Groups' allocations (communication)

To what extent do allocations change when communication is allowed? The messages exchanged while chatting show clearly that the opportunity to communicate is actively exploited: the subjects are very involved in the game, they send relevant messages on how to coordinate their actions, and then follow through.²⁷ Thus we ignore the possibility of chattering equilibria and compare the groups' actions to the equilibrium strategies of the CB game, summarized in Table 3. Note that only group-wide strategies are identified.

Figure 6 reports, for each treatment, the frequency of urns holding different numbers of Orange and Blue balls (in orange and blue in the figure), averaged over all sessions, groups and chat rounds of the same treatment.²⁸ The figure includes in gray, as a matter of comparison, the same frequencies computed for each of the no-chat treatments.

As in the no-communication treatments, the minority does concentrate its balls on a fraction of urns, and does so more than the majority. In all treatments more than 40 percent of urns receive no minority balls, while less than 10 percent receive no majority balls. The intuitive observation that a small budget demands concentration is reflected in the data.

²⁷For example, here is an edited but representative exchange between two minority members in round 13 of session 6, treatment 23C (using italics to distinguish one individual): “*2200 for me. We can do 4400.*”; “Or i could do 1030.”; “*2200. So we can do 4400*”; “And what the blues tend to be doing is just putting 3 in each.”; “i was checking the history”; “*2200*”; “hi hi.” At the same time, the majority members in the same committee were saying: “*So, even for now. lets see what happens. if they get smarter we will change next round.*”; “i think theyve figured out they needa concentrate their balls since they have fewer players.”; “*do even distribution. orange members not smart enough to do 2 urns 4 balls.*” Indeed, in this round and group the minority group played 4400, the majority played 3333, and the minority won two urns.

²⁸In principle, urns can contain up to MK majority balls in each treatment. However, by truncating the figure at mK balls (the upper bound for the minority), we still report 99 percent of all majority data (for chat and no-chat treatments), while making the figure much more readable. Note that casting more than $mK + 1$ balls in one urn is a strictly dominated strategy for the majority (and we observe it exactly once, out of a total of 1,200 urn allocations over the three chat treatments).

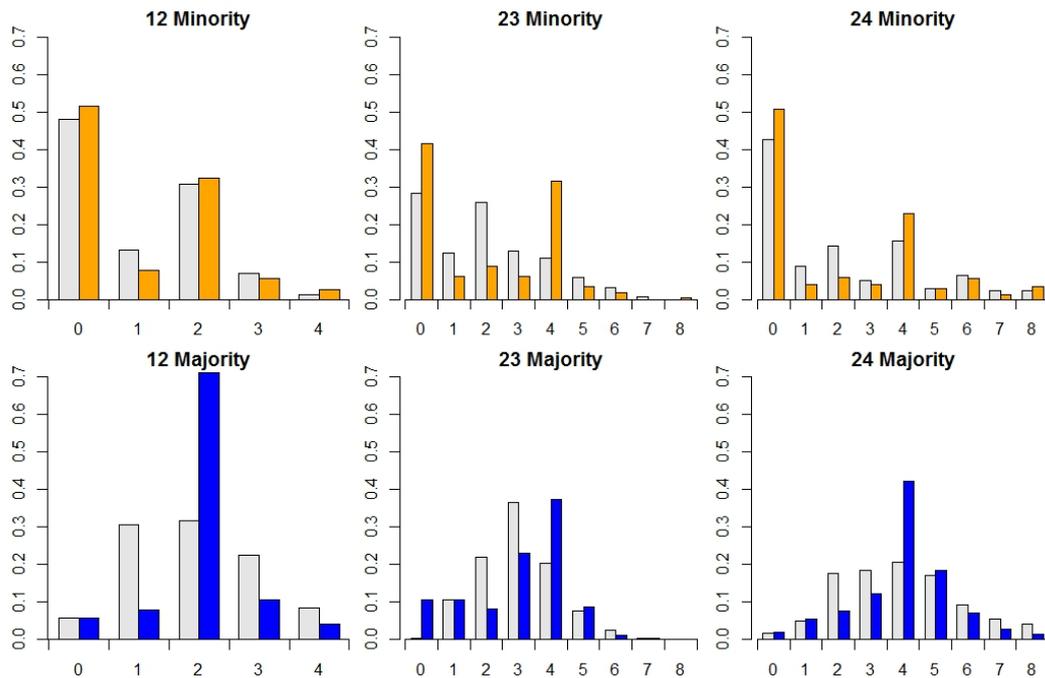


Figure 6: Frequency of group marginal allocations of balls

More precisely, in treatment 12C, the data are consistent with centralized equilibrium behavior. The minority targets 48 percent of the urns; it casts two balls in two thirds of the targeted urns, and one or three balls with very similar frequency, in line with the predictions of Table 3. Similar observations hold for the majority: the frequency of 2-ball urns is 71 percent and the frequency of 1 and 3-ball urns is very similar, again in line with Table 3.²⁹ Such consistency with equilibrium predictions, however, is not observed in the other two treatments. In 23C and 24C, according to the optimal strategies in Table 3, the majority should never cast an even number of balls, while the minority should cast two, four, and six balls with the same frequency. For both groups, on the other hand, the data show a peak at four balls.

In fact, in all three treatments, the modal number of balls cast by either group is $2m$. This coincides with optimal strategies in 12C, but does not in 23C and 24C. One plausible conjecture is that the minority tends to target two urns, and the majority mimics the minority. Although not always optimal, the strategy matches well the hide-and-seek nature of the game.

²⁹Note that in this treatment, for both groups the individual equilibrium strategies of the decentralized game add up to a team equilibrium strategy of the centralized game. Thus the comparison to the no-chat results is instructive. Communication is relevant only for the majority group, and the minority strategy remains mostly unchanged. For the majority, however, communication brings a clear change: the team plays 2222 in 65 percent of all rounds (versus 13 percent with no-chat), and the frequency of 2-ball urns more than doubles (from 32 to 71 percent).

A.12 Unpredictability and best replies

The theory makes predictions not only on how balls should be allocated across urns but also on the randomness of the minority’s allocation. In our experimental design, with rematching groups, an individual subject can maintain unpredictability while making balls allocations that are correlated over time. To a more limited extent this remains true at the group level, since there are multiple rematching committees in each experimental session.

An alternative approach is to evaluate unpredictability indirectly, by measuring the payoff gains available to each group, had it best responded to the opposite group’s experimental actions. A fully predictable minority strategy, for example, means that there exists a majority best response that translates into zero minority victories.

By focusing on best responses, the approach we take here also allows us to quantify the answer to a natural question: if the experimental subjects did not play equilibrium strategies, how far were they from playing optimally? Note that the answer can be read through two main perspectives. First, it tells us whether a group replied accurately to the other: this is the *best reply* perspective. Second, it reveals what payoff a group *guaranteed* to itself, by considering the worst-case scenario in which the opposite group best replies — as in in Proposition 3.

For treatment T and session S , we fix the observed distribution of minority group’s allocations $V_{T,S}^m$, distinguishing across urns (with one observation per group and per round): this is the “statistical strategy” of the minority. Then, we compute the best reply of the majority $BR^M(V_{T,S}^m)$ assuming that majority members could coordinate, again distinguishing across urns. The corresponding *guaranteed payoff* $p_m(V_{T,S}^m, BR^M(V_{T,S}^m))$ is the minimal payoff that the minority can obtain by playing statistically as in the experiment.³⁰ We do the same exercise for both groups.

Figure 7 summarizes the results, reported in terms of p_m . Because we observe little variation across sessions, for each treatment the results in the figure are averaged across sessions.

The different panels correspond to the different treatments; the red lines indicate the observed average frequency of minority victories in the data, and the blue traits the predicted equilibrium frequency. In each panel, the arrow on the left side indicates the value of p_m when the majority best replies, and the arrow on the right side when the minority best replies.

How should the figure be read? Consider for example treatment 12D, with average $p_m = 0.26$, slightly above the equilibrium prediction of 0.25. Given what the minority statistically played, p_m could have been as low as 0.22, had the majority best replied. Conversely, given what the majority statistically played, p_m could have been as high as 0.30. According to the *best reply* perspective, the length of a group’s arrow measures the

³⁰For example, the three observations $\{(2, 1, 1, 0), (2, 2, 0, 0), (2, 0, 2, 0)\}$ in treatment 1, 2 would correspond to the following statistical strategy: urn 1: 2 balls with probability 1, urn 2: 0,1, or 2 balls, each with probability 1/3; urn 3: 0,1, or 2 balls, each with probability 1/3; urn 4: 0 balls with probability 1. The majority’s best response is $(3, 2, 2, 1)$, implying $p_m = 1/12$.

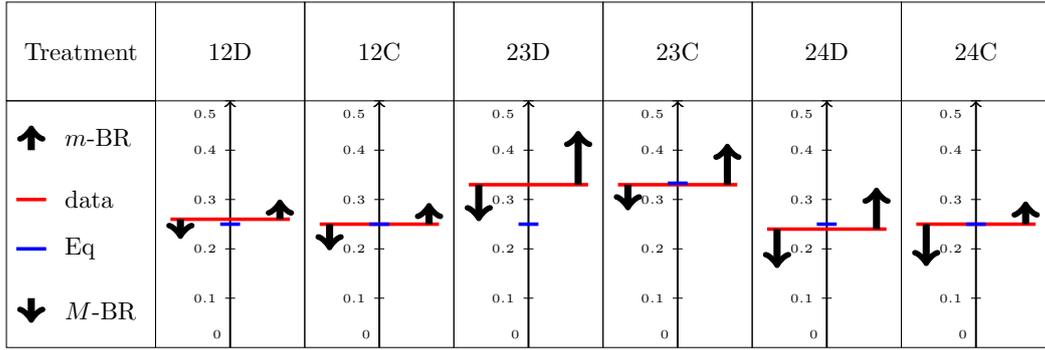


Figure 7: Best-reply minority payoffs, computed for each group and treatment, averaged across sessions

distance to the best reply. We see that the two groups were quite effective in maximizing their payoffs, with each group falling short of its best achievable payoff by an amount of 0.04. Note that when the two arrows collapse, the profile is an equilibrium of the centralized game. Thus, the length of the two arrows is also a measure of the distance from such an equilibrium.

From the *guaranteed payoff* perspective, the effectiveness of a group’s play is measured by the value corresponding to the *opposite* group’s best reply. The minority guaranteed itself a payoff of 0.22, way above 0 and not far from the equilibrium payoff of 0.25; while the majority guaranteed that p_m would stay below 0.30.

Reading the figure across all treatments, we are led to three main conclusions. First, in all treatments the minority was able to guarantee itself a significant fraction of victories, ranging from a minimum of 0.16 in treatment 24D to a maximum of 0.28 in 23C. Note that this observation does not depend on experimental majority allocations; rather, it reflects the fact that the minority was able to make its actions sufficiently unpredictable.

Second, the majority was also able to limit its losses, guaranteeing itself an upper bound on p_m that ranged between 0.43 in treatment 23D to 0.29 in 12C. These observations give us confidence on the robustness of the payoffs found in the experiment: although, on the whole, subjects did not play equilibrium strategies, both groups secured worst-case payoffs that were close to actual payoffs.³¹ The similarity of experimental and theoretical payoffs observed in Figure 1 did not occur by chance: in precisely defined payoff terms, experimental strategies were “close” to equilibrium.

Finally, for each minority and majority size, communication makes very little difference not only to observed payoffs, but also to guaranteed payoffs. In our experimental data, any difference between the two groups in the ability to communicate effectively and coordinate is not reflected in payoffs.

³¹The largest difference appears for the majority in treatment 23D, where $p_m = 0.33$ but could have been 0.44, had the minority best replied.

References

- [1] ARAD, A. AND A. RUBINSTEIN (2012): “Multi-dimensional iterative reasoning in action: The case of the Colonel Blotto game,” *Journal of Economic Behavior & Organization*, 84, 571–585.
- [2] AVRAHAMI, J. AND Y. KAREEV (2009): “Do the weak stand a chance? Distribution of resources in a competitive environment,” *Cognitive Science*, 33, 940–950.
- [3] BOREL, E. AND J. VILLE (1938): *Applications de la théorie des probabilités aux jeux de hasard*, Jacques Gabay.
- [4] BUNWAREE, S. AND R. KASENALLY (2005): *Political parties and democracy in Mauritius*, Johannesburg: EISA.
- [5] CASELLA, A. (2005): “Storable votes,” *Games and Economic Behavior*, 51, 391–419.
- [6] CASELLA, A., T. PALFREY, AND R. RIEZMAN (2008): “Minorities and Storable Votes,” *Quarterly Journal of Political Science*, 3, 165–200.
- [7] CHOWDHURY, S. M., D. KOVENOCK, AND R. M. SHEREMETA (2013): “An experimental investigation of Colonel Blotto games,” *Economic Theory*, 52, 833–861.
- [8] CONDORCET (1785): *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*, Paris: Imprimerie Royale.
- [9] DAHL, R. A. (1991): *Democracy and its Critics*, Yale University Press.
- [10] EIFERT, B., E. MIGUEL, AND D. N. POSNER (2010): “Political competition and ethnic identification in Africa,” *American Journal of Political Science*, 54, 494–510.
- [11] EMERSON, P. (1998): *Beyond the Tyranny of the Majority*, The de Borda Institute.
- [12] EMERSON, P. J. (1999): *From Belfast to the Balkans: Was Democracy Part of the Problem?*, The de Borda Institute.
- [13] FIORINA, M. P., S. J. ABRAMS, AND J. C. POPE (2005): *Culture war?*, Pearson Longman New York.
- [14] GREINER, B. (2015): “Subject pool recruitment procedures: organizing experiments with ORSEE,” *Journal of the Economic Science Association*, 1, 114–125.
- [15] GROSS, O. AND R. WAGNER (1950): “A continuous Colonel Blotto game,” Tech. rep., DTIC Document.
- [16] GROSSER, J. AND T. GIERTZ (2014): “Pork barrel politics, voter turnout, and inequality: an experimental study,” *University of Cologne working paper*.
- [17] HART, S. (2008): “Discrete Colonel Blotto and general lotto games,” *International Journal of Game Theory*, 36, 441–460.
- [18] Hoeffding, W. (1963): “Probability inequalities for sums of bounded random variables,” *Journal of the American statistical association*, 58, 13–30.
- [19] HORTALA-VALVE, R. AND A. LLORENTE-SAGUER (2012): “Pure strategy Nash equilibria in non-zero sum colonel Blotto games,” *International Journal of Game Theory*, 41, 331–343.

- [20] JACOBSON, G. C. (2008): *A Divider, Not a Uniter: George W. Bush and the American People: The 2006 Election and Beyond*, Longman Publishing Group.
- [21] KABRE, P. A., J.-F. LASLIER, K. VAN DER STRAETEN, AND L. WANTCHEKON (2013): “About political polarization in Africa: An experiment on Approval Voting in Benin,” *Mimeo*.
- [22] KORIYAMA, Y., A. MACÉ, R. TREIBICH, AND J.-F. LASLIER (2013): “Optimal apportionment,” *Journal of Political Economy*, 121, 584–608.
- [23] LASLIER, J.-F. (2012): “Why not proportional?” *Mathematical Social Sciences*, 63, 90–93.
- [24] LIJPHART, A. (2004): “Constitutional design for divided societies,” *Journal of democracy*, 15, 96–109.
- [25] MAY, K. O. (1952): “A set of independent necessary and sufficient conditions for simple majority decision,” *Econometrica*, 680–684.
- [26] PICARD, E. (1994): “Les habits neufs du communautarisme libanais,” *Cultures et conflits*, 49–70.
- [27] RAE, D. W. (1969): “Decision-rules and individual values in constitutional choice,” *American Political Science Review*, 63, 40–56.
- [28] REYNAL-QUEROL, M. (2002): “Ethnicity, political systems, and civil wars,” *Journal of Conflict Resolution*, 46, 29–54.
- [29] ROBERSON, B. (2006): “The Colonel Blotto game,” *Economic Theory*, 29, 1–24.
- [30] ROGERS, J. (2015): “An Experimental Investigation of Lobbying Strategies,” *Mimeo*.
- [31] WINSLOW, C. (2012): *Lebanon: war and politics in a fragmented society*, Routledge.

Acknowledgments

We thank participants to seminars and conferences in Berkeley, Caltech, Marseilles, St. Petersburg, UC Santa Barbara, and Tokyo, and in particular Yann Bramoullé, Matias Nunez and Rafael Treibich, for useful suggestions and comments. We thank Manuel Puente and Liu Yongfeng for their research assistance, and the National Science Foundation for financial support (grant SES-0617934).

Alessandra Casella
 Columbia University
 New York City, NY, USA
 Email: ac186@columbia.edu

Jean-François Laslier
 Paris School of Economics
 Paris, France
 Email: jean-francois.laslier@ens.fr

Antonin Macé
 Aix-Marseille School of Economics
 Marseille, France
 Email: antonin.mace@univ-amu.fr