Democracy for Polarized Committees

The Tale of Blotto's Lieutenants*

Alessandra Casella† Jean-François Laslier‡ Antonin Macé§

October 19, 2017

Abstract

In polarized committees, majority voting disenfranchises the minority. Allowing voters to spend freely a fixed budget of votes over multiple issues restores some minority power. However, it also creates a complex strategic scenario: a hide-and-seek game between majority and minority voters that corresponds to a decentralized version of the Colonel Blotto game. We offer theoretical results and bring the game to the laboratory. The minority wins as frequently as theory predicts, despite subjects deviating from equilibrium strategies. Because subjects understand the logic of the game — minority voters must concentrate votes unpredictably — the exact choices are of secondary importance, a result that vouches for the robustness of the voting rule to strategic mistakes.

*We thank participants to seminars and conferences in Bath, Berkeley, Caen, Caltech, Columbia, Lund, Marseilles, Penn State, Stony Brook, St. Petersburg, UC Santa Barbara, Tokyo, and Toulouse, 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).

†Columbia University, NBER and CEPR, ac186@columbia.edu.

‡Paris School of Economics and CNRS, jean-francois.laslier@ens.fr.

§Aix-Marseille Université (Aix-Marseille School of Economics), CNRS and EHESS, antonin.mace@univ-amu.fr.

Keywords: Storable Votes, Polarization, Colonel Blotto, Tyranny of the Majority, Committees, Voting.
JEL classification: D71, C72, C92.

1 Introduction

How should political power be shared? Majoritarian democracy is desirable under many criteria (Condorcet, 1785; May, 1952; Rae, 1969), 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 predate the democratic institutions or be generated by the institutions themselves.[2] Referring to Northern Ireland, the Balkans, and other places plagued by civil wars, Emerson (1998, 1999) claims that in such situations 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 when preferences are fully polarized and the power of a cohesive majority bloc is secure — a scenario we refer to as a *systematic* minority — 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-

[1]Political philosophy has long recognized that the *tyranny of the majority* poses a fundamental challenge to the legitimacy of majority voting (Dahl, 1991).

[2]See Jacobson (2008); Fiorina et al. (2005) for the US case, or Reynal-Querol (2002); Eifert et al. (2010); Kabre et al. (2013) for African cases.

[3]For example, in Lebanon (Picard, 1994; Winslow, 2012), in Mauritus (Bunwaree and Kasenally, 2005), and occasionally elsewhere (Lijphart, 2004).

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 (Casella, 2005). 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. In fact, our arguments apply to virtually any divided group: for instance, SV could be used by the board of directors of a company.

Although easy to describe, SV poses a challenging strategic problem: *how* should a voter best divide her votes over the different issues? Note 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: in the original version of the game (Borel and Ville, 1938; Gross and Wagner, 1950), two opposite military leaders with given army sizes must choose how many soldiers to deploy on each of several battlefields. Each battlefield is won by the army with the larger number of soldiers. Each colonel could win if he knew the opponent's plan. At equilibrium, choices must be random.

The SV's model can be phrased as in the classical Colonel Blotto scenario, with "issues" and "votes" instead of "battlefields" and "soldiers". The game is asymmetric — the majority has more votes — and thus recalls

[4]Note that neither vetoes or supermajority requirements, nor logrolling 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.

Colonel Blotto analyses that allow for heterogeneous armies.[5] It differs however on one important dimension: it is a *decentralized* Blotto game. Each voter, whether in the majority or in the minority, controls a number of votes, to be allocated to the different issues. In its military analogue, it is as if multiple, individual lieutenant colonels in each of the two armies controlled their own battalions and chose how to distribute them over the different battlefields. Again, each battlefield is won by the army that deploys more soldiers.

To our knowledge, the decentralized Blotto game has not been studied before. In this game, although the interests of all lieutenants within each army are perfectly aligned, decentralizing the centralized solution is generally not possible: the centralized solution requires centralized randomization and thus cannot be replicated unless the randomization can be communicated, and communication is truthful and believed. The decentralized Blotto game can be of independent interest, beyond the specific application to SV's. From lobbying to campaign spending, from patent races to fighting criminal networks, traditional applications of the centralized Blotto games can be extended profitably to situations where one or both sides consist of multiple independent actors.

We start by studying the game in the absence of communication: we develop theoretical results, in particular results that will be useful for the experimental tests we describe in the second part of the paper. The game has many equilibria but, reverting to SV terminology, if the difference in size between the two groups is not too large, the minority is expected to win occasionally in all equilibria. We identify a class of simple strategies, neutral with respect to the issues and symmetric within each group, and characterize conditions under which profiles constructed with such strategies are equilibria. Strategies are such that each minority member concentrates her votes on a subset of issues, randomly chosen, and again induce a positive expected fraction of minority victories in equilibrium. In fact, the result is stronger and holds off equilibrium too: if each minority member concen-

[5] As in Roberson (2006) if soldiers can be deployed continuously, or Hart (2008) if soldiers are discrete.

trates her votes and does so randomly, the minority can guarantee itself a positive probability of victories, for *any* strategy by the majority, whether coordinated or not, and regardless of whether or not the minority voters all choose precisely the same strategy. When communication within each group is possible, the equilibria of the non-communication game continue to exist as chattering equilibria. However other equilibria exist, including the equilibria of the centralized Blotto game in which each of the two groups, the minority and the majority, acts as a single agent. Here we borrow from Hart (2008)'s results on discrete Blotto games and identify equilibria of the centralized Blotto game that hold for the parameter values we use in the laboratory. Again the theoretical prediction is a positive fraction of minority victories, in fact, interestingly, a very similar fraction to that predicted by our simple equilibrium strategies in the absence of effective communication, for the same parameter values.

We test the theoretical predictions in the laboratory in two treatments, one without and one with communication. 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. In contrast, majority subjects appear to alternate between exploiting their size advantage by covering all issues, and mimicking minority subjects. Be it with or without communication, the strategies of both groups deviate from the precise predictions of the theoretical equilibria, and yet the fraction of minority victories we observe is very close to equilibrium, varying 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. We read these findings as endorsement of the robustness of the voting rule to strategic mistakes. As in the off-equilibrium theoretical result described earlier, as long as minority voters recognize the importance of concentrating and randomizing their votes, their exact choices are of secondary importance: whether votes are concentrated on two or on only one issue, whether they are split equally or unequally, all this affects minority victories only marginally. This conclusion is the main result of the paper.

The robustness of SV to strategic mistakes has been noted before (Casella et al., 2006, 2008). Previous models, however, 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 she considers higher priorities (at a given state), and this behavior is observed in the lab. But the intensity of one's own preferences is instinctively focal, and the question then arises whether the good performance of the mechanism extends to more complicated settings.

If intensities are commonly known, or if it is known that intensities are correlated across groups, the hide-and-seek nature of the game appears and with it the minority's need to randomize its strategy.[6] 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 quite 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 (Laslier, 2012; Koriyama et al., 2013).[7]

Our experimental results also bear comparison to a small recent experimental literature on the asymmetric Colonel Blotto game. In line with Avrahami and Kareev (2009) and Chowdhury et al. (2013), we observe that the minority concedes some battlefields (some issues) in order to win others. However, the key difference in our setting is the decentralization of

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

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

decisions, which renders the game more complex. Rogers (2015) introduces some decentralization in a related game, whose payoffs differ from classical Blotto payoffs along several dimensions.[8] 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.

Arad and Rubinstein (2012) identify several salient strategy dimensions in the Colonel Blotto game and argue that subjects use multi-dimensional hierarchical reasoning in deciding their behavior. Our environment with multiple heterogeneous players is more complex, but our experimental results are in line with the idea of multi-dimensional reasoning: our minority subjects, in particular, appear to identify easily the qualitative features of the equilibrium–concentrate votes and be unpredictable–but show substantial confusion on the exact number of votes cast on targeted issues.

The game we study can be seen as a “contest” between two teams. Dechenaux et al. (2015) provides a comprehensive survey of the experimental literature on contests, and in particular on Colonel Blotto games, seen as multi-battle contests. However, as recalled by Sheremeta (2015), the key question addressed in this literature is the efficiency —or inefficiency— that arises from the trade-off faced by participants between the cost of effort and the incentive to cooperate within their group. In our case, the efficiency question is moot because any outcome is Pareto optimal by design (contrary to Cason et al. (2012)), and there is neither an individual trade-off between cooperation and effort, nor intra-group punishment (as in Abbink et al. (2010)), nor the possibility of alliance between players with disjoint interests (as in Kovenock and Roberson (2012)). Note also that our decentralized game is conceptually different from team contests in which members of opposite teams fight each other in pairwise battles (as in Arad (2012), Rinott et al. (2012) or Fu et al. (2015)).

The paper is organized as follows. After the introduction, Section 2

[8] Some 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).

presents the model. Section 3 discusses two preliminary remarks on the distinction between centralized and decentralized games. The theory for the decentralized game is presented in Section 4. Section 5 describes the experimental protocol, and Section 6 presents the experimental results. Section 7 concludes. All proofs are in the appendix (Section A). A supplementary appendix contains a discussion of the evolution of the experiment over time, an analysis of the chat messages, and a copy of the experimental instructions.

2 The Model

A committee of N individuals must resolve $K \geq 2$ binary issues: the committee must decide whether to pass or fail each of K independent proposals. The set of issues is denoted by $\mathcal{K} = \{1, \ldots, K\}$. 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 the group. The specific direction of preferences is irrelevant, what matters is that the two groups are fully cohesive and fully opposed. We summarize these two features by calling m a *systematic minority.*

Each individual receives utility 1 from any issue resolved in her preferred direction, and 0 otherwise. Thus each individual's goal is to maximize the fraction of issues resolved according to her — and her group's — preferences.

Individuals are all endowed with K votes each, and each issue is decided according to the majority of votes cast. If each voter is constrained to cast one vote on each issue, M wins all proposals. This *tyranny of the majority* is our point of departure: with simple majority voting, a systematic minority is fully disenfranchised. The conclusion changes substantively if voters are allowed to distribute their votes freely among the different issues. Each issue is then again decided according to the majority of *votes* cast — which now, 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. The voting rule is then a specification of *Storable Votes*, with votes on all issues cast at the same time.[9]

A specific welfare criterion (a specific degree of concavity in the social welfare function) will capture the society's normative concern with minority representation. If we call p_m the expected fraction of minority victories, such a concern will translate into an optimal $p^*_m(M, m)$. Here we do not specify the welfare criterion and limit ourselves to measuring p_m.

We suppose that the parameters of the game are common knowledge, in particular each voter knows exactly the size of the two groups, and thus both her own and everyone else's preferences. Our framework is thus a one-stage, full information game.

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:

$$S(K) = \left\{ s = (s_1, \dots, s_K) \in \mathbb{N}^K \;\middle|\; \sum_{k=1}^{K} s_k = K \right\},$$

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 $\boldsymbol{s} = (s^1, \dots, s^m) \in S(K)^m$, where the bold font indicates a vector of allocations, the number of votes allocated by the minority to issue k is denoted by:

$$v^m_k(\boldsymbol{s}) = \sum_{i=1}^{m} s^i_k.$$

We denote by $v^m(\boldsymbol{s}) = (v^m_k(\boldsymbol{s}))_{k \in \mathcal{K}} \in S(mK)$ the allocation of votes by the minority side associated to the minority-profile $\boldsymbol{s}$.

Similarly, let the majority players be ordered from 1 to M. Denoting by $\boldsymbol{t} = (t^1, \dots, t^M) \in S(K)^M$, the majority profile, the number of votes

[9] As in chapters 5 and 6 in Casella (2012). See also Hortala-Vallve (2012).

allocated by the majority to issue k is denoted by:

$$v_k^M(\boldsymbol{t}) = \sum_{i=1}^{M} t_k^i,$$

and we denote by $v^M(\boldsymbol{t}) = (v_k^M(\boldsymbol{t}))_{k\in\mathcal{K}} \in S(MK)$ the allocation of votes by the majority side associated to the majority-profile $\boldsymbol{t}$.

For a given profile $(\boldsymbol{s}, \boldsymbol{t}) \in S(K)^m \times S(K)^M$, the payoffs for each member of the two groups, called g_m and g_M, are given by

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

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

Finally, we denote by $\Sigma(K) = \Delta\left(S(K)\right)$ the set of all probability measures on $S(K)$, i.e. the set of mixed strategies. Then the expected payoff to the minority $E\left[g_m\right]$ 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 $(v^m, v^M) \in S(mK) \times S(MK)$. Then we will also write, with abuse of notation, $p_m(V^m, V^M)$. Our goal is to study this game, both theoretically and experimentally. 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 (2008).[10]

Remark 1 *For any K and m, none of the equilibria of the CB game in Hart (2008) can be replicated in the DB game if M is larger than a finite threshold $\overline{M}(K)$.*

The intuition is straightforward: with the exception of knife-edge cases, equilibrium strategies in the centralized game must be such that the marginal allocation of forces on any given battlefield follows a uniform distribution. But the sum of independent variables cannot form a uniform distribution in general: unless the randomization is centralized, the strategy cannot be replicated.

In some applications, a precise description of the strategic environment would include a communication phase before the game is played. With preplay communication, subjects can send costless and non-binding messages (they are cheap talk), and a communication protocol describes who can send a message to whom. For any communication protocol, the equilibria of the decentralized game are still equilibria of the game with communication, but new equilibria can arise.

One protocol deserves particular attention in our game. As the two groups are fully opposed and fully cohesive, each may want to coordinate its voting, without making that information public to the other group. We call *group-communication* the protocol in which any subject can only send messages to all members of her group. With group-communication, a protocol that we test in the experiment, the logic behind Remark 1 breaks down.

[10] Hart (2008) does not characterize optimal strategies for all parameter values. Roberson (2006) 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.

It then becomes possible, and advantageous, for each group to randomize over the possible allocations at the central level, and then decentralize the realized allocations.

Remark 2 *With group-communication, the equilibria of the centralized Colonel Blotto game can be replicated.*

In the experimental part of this paper, we will use as theoretical benchmark of the treatment with group-communication the equilibria of the CB game in Hart (2008).[11]

Before turning to the experiment, we derive theoretical results for the decentralized Blotto game without communication, for which no analysis exists in the literature. We will use such results as the theoretical reference for the experimental treatment without communication.

4 Theory: the decentralized Blotto game without 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.*

The coordination problem within each of the two groups results in many equilibria. We do not aim to characterize them all; rather in this section we

[11] Other equilibria exist, including chattering equilibria in which communication is ignored, or asymmetric equilibria in which communication is ignored by one group but not by the other. The messages our subjects sent during the experiment, however, show that communication was used to coordinate actions, in line with the equilibria of the CB game. We discuss the experimental messages in the second part of the paper and in the supplementary appendix.

focus on equilibria that either stress the difference between the decentralized and the centralized version of the game, or that have a simple enough structure to provide a plausible theoretical reference for the experiment.

4.1.1 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 1 *If $M \geq m \geq 2$ and $M + m \geq (K+1)^2/K$, a pure-strategy equilibrium always exists.*

This result clearly indicates that the DB game differs from the CB game, in which pure-strategy equilibria generically fail to exist.[12] 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 $\mathcal{K}_m = \mathcal{K} \backslash \mathcal{K}_M$. 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.*

[12]In 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 (Hortala-Vallve and Llorente-Saguer, 2012).

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.*

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 1 require a large amount of coordination, both within and across groups. In addition, not only in those equilibria, but also in the "trivial" equilibrium with $M \geq m + K$ (where every voter casts one vote on each issue and the minority loses all decisions), each voter has only a weak incentive not to deviate. This seems particularly problematic when $M \geq m + K$: even non-strategic minority members seem likely to realize that some concentration is called for.

4.1.2 Symmetric equilibria in mixed strategies

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. Mixed strategies allow them to do so.

In this section, we focus on a family of simple strategies that treat each issue symmetrically and we assume that all voters within the same group play the same strategy. 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,[13] 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

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

the degree of concentration of a player's votes.[14] We denote by $\boldsymbol{\sigma^c}$ (resp. $\boldsymbol{\tau^c}$) the group 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 and belonging to the $(\boldsymbol{\sigma^c}, \boldsymbol{\tau^c})$ family 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 2 *Suppose K even and M odd. Then $(\boldsymbol{\sigma^2}, \boldsymbol{\tau^1})$ is an equilibrium if $M \leq m + 1$,*[15] *with*

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

What is remarkable in Proposition 2 is that when the difference in size between the two groups is as small as possible — at most a single member — equilibrium strategies for majority and minority voters can be quite different: while each majority voter simply casts one vote on each issue, each minority voter concentrates all votes on exactly half of the issues, chosen randomly, and casts two on each. Numerically, the expected frequency of minority victories is significant at this equilibrium, starting from $1/4$ when $(m, M) =$

[14] Arad and Rubinstein (2012) suggest that subjects faced with the Colonel Blotto game intuitively organize their strategy according to three dimensions, decided sequentially: (i) the number of targeted issues (ii) the apportionment of votes on targeted issues (iii) the choice of issues. The class of strategies $(\sigma^c)_{c \text{ factor of } K}$ is particularly easy to describe with respect to these three dimensions: (i) the number of targeted issues is $\frac{K}{c}$ (ii) the votes are equally split on all targeted issues (iii) the choice of targeted issues is random, with equal probability for each issue. This class of strategies has been independently introduced by Grosser and Giertz (2014), who refer to them as *pure balanced number strategies*.

[15] The 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 $(\boldsymbol{\sigma^2}, \boldsymbol{\tau^1})$ is an equilibrium *if and only if* $M \leq m + 1$.

$(2, 3)$ and converging to $1/2$ for large m and M.

When the difference in size between the two groups increases, the equilibrium breaks down because it becomes advantageous for minority voters to concentrate their votes even further. Consider the strategic problem of minority voter i, and suppose $M = m + 1$. If all other voters follow the strategy in the proposition, a random issue will receive M majority votes and, excluding i, a random even number of minority votes, between 0 and $2(m-1)$.[16] It is not difficult to verify that the most likely number of minority votes on any random issue, excluding i, is then either m or $m-2$, with equal probability, and thus the most likely difference in votes i wants to counter is either $M - m = 1$, or $M - m + 2 = 3$. Casting two votes on half of the issues, chosen randomly, is then a best reply[17] When $M - m > 1$, i can increase the minority's expected share of victories by cumulating more than two votes on a smaller, random subset of targeted issues. If $M - m = 3$, for example, casting three or four votes on individual issues, rather than two as dictated by Proposition 2, is a profitable deviation.

It is not surprising to see that minority voters' incentive to concentrate votes increases with the difference in size between the two groups. 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 3 *Suppose M is divisible by K. Then $(\boldsymbol{\sigma^K}, \boldsymbol{\tau^1})$ is an equilibrium if and only if $M \geq \frac{mK}{2}$. In such an equilibrium:*

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

Predictably, the minimum ratio M/m at which the equilibrium is supported must increase with K: recall that K is both the number of proposals

[16] Given the strategies described in Proposition 2, the probability that the remaining $m-1$ minority voters cast $2x$ votes on any issue k is given by $\binom{m-1}{x}(1/2)^{m-1}$.

[17] Casting four votes on a random quarter of the issues is also a best reply. But with $M = m + 1$ and M odd, it is not consistent with a simple equilibrium.

and the number of votes with which each voter is endowed; with majority voters spreading all their votes evenly, in equilibrium $v_M^k = M$ for all $k \in \mathcal{K}$, and thus, for given M/m, a minority voter's temptation to spread some of the votes increases at higher K.

Propositions 2 and 3 characterize p_m, the *expected* fraction of minority victories. But does the minority *always* win at least one of the issues, i.e. does it win at least one issue with probability one? And the majority? The following remark provides the answers.

Remark 3 *When the individuals use the equilibrium strategies identified in Propositions 2 and 3:*

- *the minority may win no proposal*
- *the majority always wins at least one proposal.*

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. For example the family of strategies $\{\sigma^c\}$ introduced in the previous section satisfies this property.

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 $s_\pi = (s_{\pi(1)}, \ldots, s_{\pi(K)})$.*[18]

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

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 4 *For all $M \leq mK$, there exists a number $\underline{k} \in \{1, \ldots, K\}$ 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 $\boldsymbol{\tau}$,*

$$p_m(\boldsymbol{\sigma}, \boldsymbol{\tau}) > 0.$$

The result of Proposition 4 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, if applications are considered seriously, 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?).[19]

Proposition 4 allows us to conclude that with randomization and sufficient concentration, the minority can expect to win some of the time, even off equilibrium. But how frequently?

Clearly the answer depends on the rules followed by each minority and majority voter. To assess the magnitude of minority payoffs off equilibrium, consider the following numerical example. Suppose $K = 4$, $M = 10$, and

[19] Note, for comparison, that Proposition 4 holds under the identical condition $M \leq mK$ for the centralized game (with both discrete and continuous allocations).

$m \in \{1, .., 10\}$. Minority voters adopt the σ^c strategies described in the previous section, with c assuming one of two values: $c = 2$ (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). Consider two plausible rules for the majority, corresponding to plausible bounds on the frequency of minority victories: 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).[20] Figure 1 reports such bounds for each value of m (on the horizontal axis) under minority rules σ^2 (in blue) and σ^4 (in green).

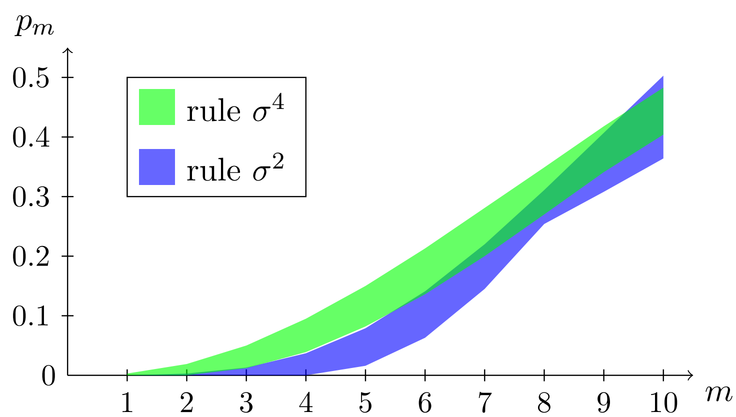

Figure 1: Minority payoffs for two minority rules ($M = 10$)

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 $M \leq mK$ 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,

[20] We 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.

when the minority is either 60 or 70 percent of the majority).

Note that the condition $M \leq mK$ in Proposition 4 is tight. The remaining case $M > mK$ refers to a committee of extreme asymmetry, in which the average number of votes of the majority per issue (M) is larger than the total amount of votes of the minority (mK). In this case, it is natural for majority players to spread their votes, and we should expect no minority victories: for any minority-profile $\boldsymbol{\sigma}$, $p_m(\boldsymbol{\sigma}, \boldsymbol{\tau}^{\mathbf{1}}) = 0$.

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 session consisted of 20 rounds with fixed values of m and M; the first ten rounds without communication, and the second ten with group-communication.

All 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 (Greiner, 2015). 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.

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 2 reproduces the relevant computer screen in one of our treatments for a Blue voter who has already cast one ball.

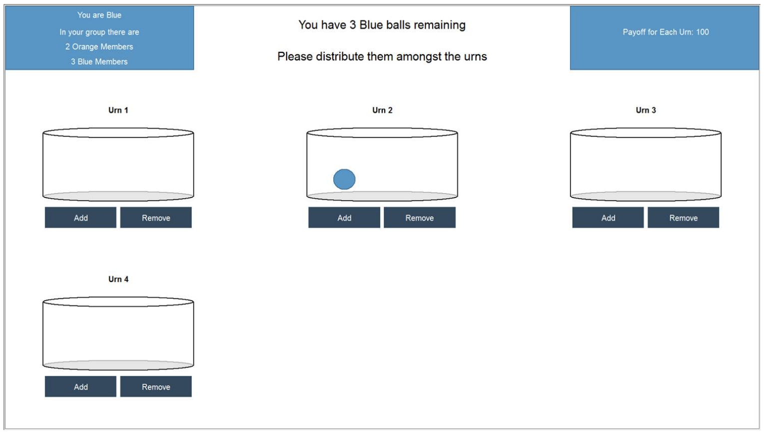

Figure 2: The Allocation screen

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. 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.

In 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. As we discuss below, subjects used the chat option intensely, with the explicit goal of coordinating their strategies.[21] The order of the treatments did not induce them to replicate the chattering equilibrium when communication was made available.

Each session lasted about 75 minutes, and earnings ranged from \$18 to \$44, with an average of \$33 (including a \$10 show-up fee). The experiment was programmed in ZTree (Fischbacher, 2007), and a copy of the instructions for a representative treatment is reproduced in the third section of the supplementary appendix.

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. Second, we wanted to compare results with and without communication. Does communication helps or hinders the relative success of the minority? As summarized in Table 1, we ran the experiment with and without the chat option for three sets of m, M values. We have thus six treatments, denoted by $mMNC$ without chat, and mMC with chat.

5.2 Parameter values and theoretical predictions

We chose the values for m and M according to three criteria. First, given the complexity of the game, we kept the size of the committee small enough to maintain the possibility of conscious strategic choices by inexperienced players. Second, we chose group sizes so as to have variation in the relative minority size m/M, keeping constant the absolute difference $M-m$ (sessions

[21] Subjects used the two minutes available for chatting fully, but we have no indication that they found the time too short.

Table 1: Experimental Design

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

s1-s3 and s4-s6), and to have variation in the absolute difference $M - m$, keeping constant the relative size m/M, (sessions s1-s3 and s7-s9). Finally, we chose parameter values such that equilibria of the decentralized game exist in the family of simple profiles $(\boldsymbol{\sigma}^{\boldsymbol{c}}, \boldsymbol{\tau}^{\boldsymbol{d}})$, symmetric within groups, and within this family are unique. We select such equilibria as theoretical reference for the experiment because of their intuitive simplicity. We know that asymmetric equilibria exist for some of the experimental parameters, and we do not rule out other symmetric equilibria with more complex mixing strategies, but their emergence seems unlikely in our experimental environment, with random rematching and inexperienced subjects.[22]

The theoretical predictions for our design are summarized in Table 2 and Table 3. Table 2 refers to the decentralized game: in both treatments 12NC and 23NC, $(\boldsymbol{\sigma}^2, \boldsymbol{\tau}^1)$ is an equilibrium; in treatment 24NC, the symmetric equilibrium is $(\boldsymbol{\sigma}^4, \boldsymbol{\tau}^1)$.[23] In all three treatments, the expected fraction of minority victories is $1/4$.

As we noticed in Remark 2, with group-communication, coordination around the equilibria of the centralized Blotto game (the CB game) becomes possible. The equilibrium strategies chracterized by Hart (2008) provide the theoretical reference for the experimental treatments with group-communication.

[22]Note that the pure-strategy equilibria identified in Proposition 1 do not appear in our experimental treatments as $(K+1)^2/K = 25/4 > 6$.

[23]Proposition 2 applies to M odd, and thus does not cover treatment 12NC. However, one can verify immediately that (σ^2, τ^1) is an equilibrium for treatment 12NC when $K = 4$. In fact, if $K = 4$, Proposition 2 extends to all M even.

Table 2: Symmetric equilibria of the decentralized game

Treatment	Simple symmetric equilibrium	p_m
12NC	$(\boldsymbol{\sigma^2}, \boldsymbol{\tau^1})$	1/4
23NC	$(\boldsymbol{\sigma^2}, \boldsymbol{\tau^1})$	1/4
24NC	$(\boldsymbol{\sigma^4}, \boldsymbol{\tau^1})$	1/4

As established by Hart, with discrete allocations the value of the CB 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. The intuition is easier to grasp if we start from the continuous CB game, where allocations need not be integer numbers. In such a game, 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 $[0, 2M]$ (Roberson, 2006). 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 $\{1, 3, .., 2\mu - 1\}$), U_e^μ the uniform distribution over even numbers with mean μ (i.e. over $\{0, 2, .., 2\mu\}$), 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.[24]

The strategies can be implemented in different ways, as long as the equal probability restriction embodied by the marginal distribution is satisfied. For example, the majority strategy in $23C$ must correspond to mixing uni-

[24] Note that the optimal strategies identified by Hart may not be unique. For example, the strategies involving $\{2\}$ in treatment 12C do not appear in Hart (2008) because they are not optimal strategies of the General Lotto game. See the discussion in Hart (2008).

Table 3: Equilibria of the centralized game

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

formly 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. 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.

6 Experimental Results

As we document in the first section of the supplementary appendix, we observe little evidence of learning in the data, either in terms of strategies or outcomes, and thus report the results below aggregating over all rounds of the same treatment.

6.1 Minority victories

Is the minority able to exploit the opportunity provided by the voting system? This is the main question of the paper, and thus we begin our analysis of the experimental data by addressing it. Figure 3 plots 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.

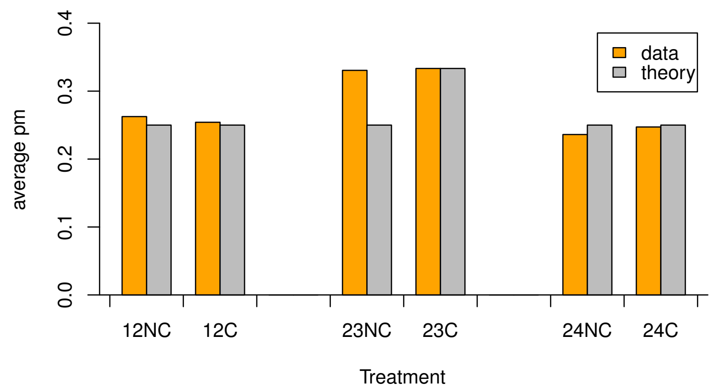

Figure 3: Fractions of minority victories

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 24NC) 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 23NC.[25]

Are the experimental subjects really adopting the rather sophisticated strategies suggested by the theory?

[25] The difference is not statistically significant. In treatment 23NC 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.

6.2 Strategies

6.2.1 No communication: ball allocations

In the absence of communication, equilibrium strategies are defined at the individual level. Figure 4 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.[26] 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.

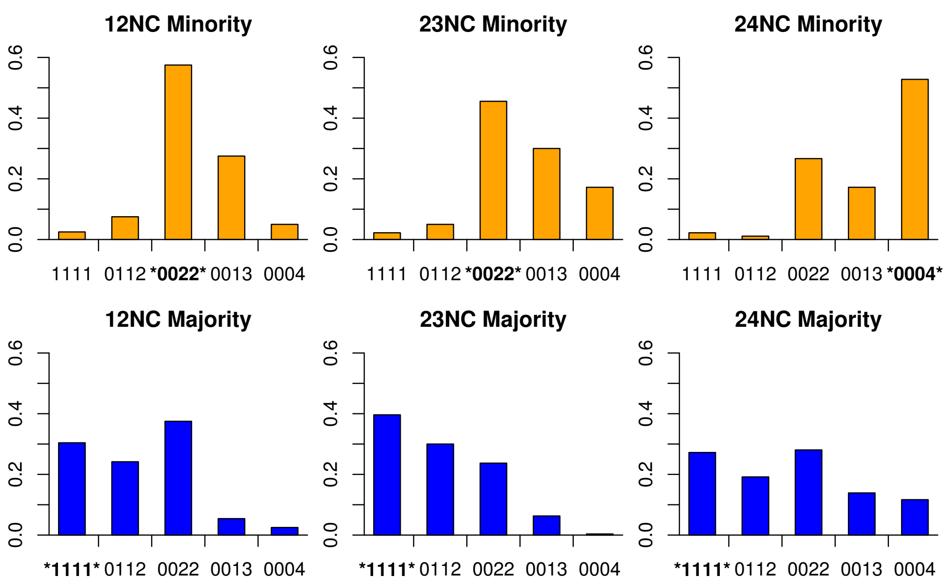

Figure 4: Frequency of individual ball allocations (no-chat treatments)

The figure teaches three main lessons. First, there is substantial devi-

[26]Thus, 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.

ation 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 12NC and 24NC, where more than half of all observed allocations correspond to the predictions.[27] Equilibrium predictions are 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. We have ordered the five possible ball allocations with concentration increasing progressively from left to right. In all treatments, the distribution of minority allocations is shifted to the right, relative to the majority distribution: predictably, and in line with the theory, minority members tend to concentrate balls more than majority members do. In all treatments, the fraction of minority subjects casting one ball in each urn, the left-most column in each panel, is negligible: the need to concentrate the number of balls cast is clear to all minority subjects since the very beginning of the game. Similarly, the fraction of majority members casting all balls in a single urn, the right-most column in each panel, is negligible in treatments 12NC and 23NC, although it surprisingly rises to 12 percent in treatment 24NC. Focusing on minority subjects, a shift to the right in the distribution of allocations is also evident as we move from treatment 12NC to 23NC, and finally to 24NC. The shift between 12NC, and 24NC is again in line with the theory, as the equilibrium strategy shifts from σ^2 to σ^4; the distribution in 23NC appears intermediate between these two cases. For majority subjects, on the other hand, the change in distribution across treatments is difficult to rationalize on the basis of the theory.

[27]This need not be a best response, given the variability in the data and the more random behavior of majority subjects.

6.2.2 No communication: individual subjects

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 (not more than $\underline{k}$, according to Proposition 4, where $\underline{k} = 2$ in all our experimental treatments) and casts them randomly. 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 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 second largest number, denoted by x_3 and ranging from 0 to 2. Each dot in the figure is a single subject's ball allocation averaged over the 10 rounds played, summarized by the subject's average x_4 and x_3.[28] Orange dots denote members of the minority, and Blue dots members of the majority.

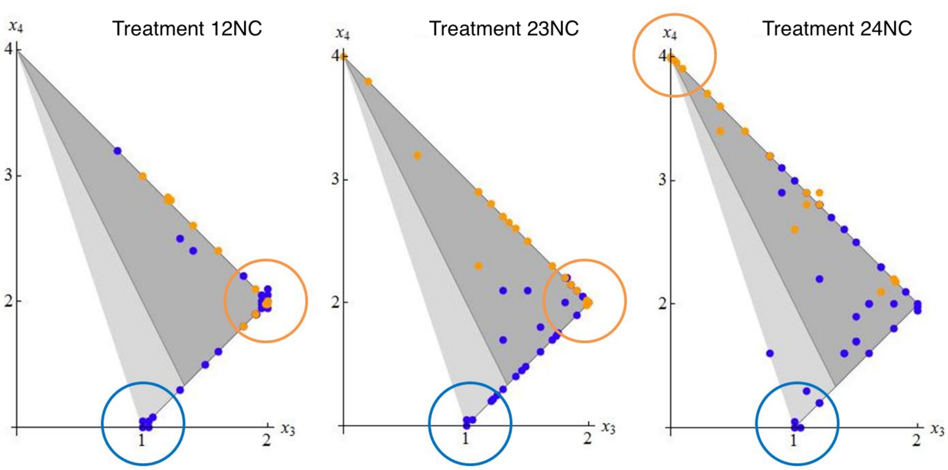

Figure 5: Individual subjects' average ball allocations (no-chat treatments)

The vertices of each triangle in the figure correspond to three feasible

[28]For instance, if a subject played 0022 on half of the rounds, and 0004 on the other half, her average allocation would be represented with $x_4 = 3$ and $x_3 = 1$.

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.[29] 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 $x_4 + 2\ x_3 = 4$, 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 12NC, 2 out of 18 in 23NC, and 3 out of 18 in 24NC *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 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-

[29]The 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. We have added a small amount of noise to the data to reduce overlaying.

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 24NC: 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 24NC is intuitive and could be the trigger for the majority subjects' own more frequent concentration.

6.2.3 Communication

To what extent does communication influence the groups' allocations? We analyze the content of the experimental chats in the second section of the supplementary appendix. We find that all subjects actively participate, and the messages are very relevant: 91 percent of exchanges[30] mention at least one ball allocation, 36 percent refer to the opposite group, and 84 percent include an explicit agreement.[31] The messages indicate an explicit effort at coordination and suggest that the equilibrium strategies of the CB game, summarized in Table 3, are indeed the appropriate theoretical reference for the experimental treatments with communication. 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),

[30]We call *exchange* the list of all messages within a group in a given committee at a given round.

[31]For 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.

averaged over all sessions, groups, and chat rounds of the same treatment.[32] The figure includes in gray, as a matter of comparison, the same frequencies computed for each of the no-chat treatments.

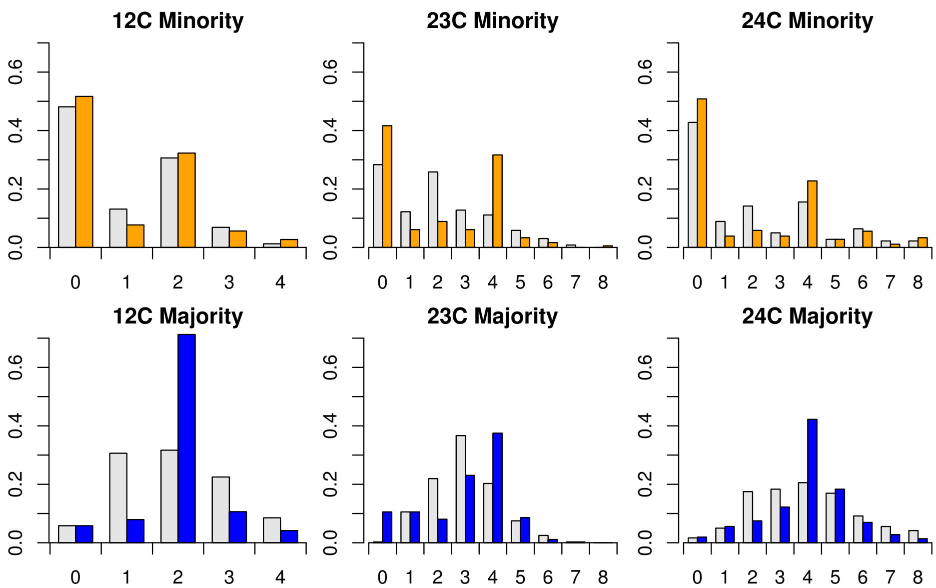

Figure 6: Frequency of group marginal allocations of balls

As in the no-chat 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 again reflected in the data.

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. Sim-

[32] In 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).

ilar 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.[33] 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. As we mentioned earlier, although not always optimal, the strategy matches well the hide-and-seek nature of the game.

6.3 Unpredictability and best replies

According to the theory, not only should the minority concentrate its balls on a subset of urns, but the targeted urns should be unpredictable. In our experimental design, with rematching groups, there are two complementary sources of unpredictability. First, a group allocation can be made unpredictable by the choices of its subjects if, at each round, they randomize over their targeted urns. Second, the composition of the group itself is random, and this reinforces the unpredictability of the group allocation. In this section, rather than separating these two forces, we assess globally the unpredictability of the observed allocations. We perform this analysis for both no-chat and chat treatments.

As a preliminary remark, we note that the spatial distribution of balls

[33]Note 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).

in the experiment is quite even: overall, each urn received between 24.4% (urn 4) and 25.5% (urn 2) of the balls. For all treatments, each of the four urns received between 20% and 30% of the balls. We see no systematic bias in favor of, or against, any particular urn.

We evaluate the unpredictability of one's group allocations (say, the minority) over a session by assessing how much the opposite group (the majority) could have extracted from the knowledge of the distribution of these allocations. That is, we measure the payoff gains available to the majority, had it best responded to the minority's experimental actions. A fully predictable minority strategy, for example, means that there exists a majority best response that translates into zero minority victories.

Precisely, for treatment T and session S, we fix the observed distribution of minority group's allocations $V^m_{T,S}$, 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 $BR^M(V^m_{T,S})$ assuming that majority members could coordinate, again distinguishing across urns. The corresponding *guaranteed payoff* $p_m(V^m_{T,S}, BR^M(V^m_{T,S}))$ is the minimal payoff that the minority could obtain by playing statistically as in the experiment.[34] 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 12NC, with average $p_m = 0.26$, slightly above the equilibrium prediction of 0.25.

[34] For 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 $p_m = 1/12$.

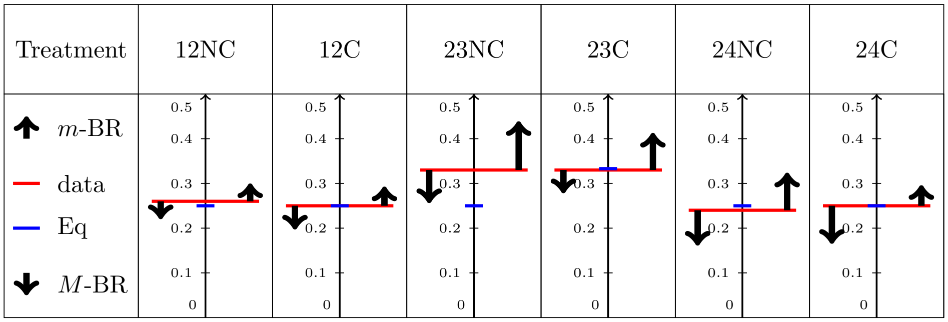

Figure 7: Best-reply minority payoffs, computed for each group and treatment, averaged across sessions

Given what the minority statistically played, p_m could have been as low as 0.22, had the majority best replied. This means that the minority guaranteed itself a significant payoff, way above 0 and not far from the equilibrium payoff of 0.25. Conversely, given what the majority statistically played, p_m could have been as high as 0.30[35].

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 24NC 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. The result is particularly remarkable for the no-chat treatments, in which minority subjects cannot coordinate, and provides the experimental counterpart of Proposition 4.

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 23NC to 0.29 in 12C. Taken together, these observations give us confidence on the robustness of the payoffs found in the experiment: although, on the whole, subjects

[35] Alternatively, the length of a group's arrow can be read as a measure of the 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.

did not play equilibrium strategies, both groups secured worst-case payoffs that were close to actual payoffs.[36] The similarity of experimental and theoretical payoffs observed in Figure 3 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.

7 Conclusions

We have investigated the ability of the SV mechanism to protect the minority group 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. Whether minority subjects concentrate votes on the correct number of target issues, and whether majority voters are able to best-respond to minority strategies, these finer strategic points are of secondary importance. 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

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

are guaranteed (in expectation). 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.

From a theoretical perspective, this paper has contributed a new version of the classic Colonel Blotto game: a decentralized game where the allocation of resources is deferred to multiple individual lieutenants within each army. Although incentives are perfectly aligned, in the absence of communication the decentralized game cannot replicate the equilibria of the centralized Blotto game (because randomization needs to be centralized). Thus the paper can be of interest beyond the specific application to SV, and opens the study of different problems as decentralized Blotto games. Possible applications include patent races with multiple intra-firm research teams; campaign spending in the US, with aligned and opposed political action committees (PAC's); or the fight against terrorism, with limited communication across terrorist cells and more or less coordination among international police forces.

A Mathematical 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 (2008) 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: $\forall k \in \mathcal{K}, \quad V_k^M \sim \mathcal{U}\left(\{1, 3, \ldots, 2M-1\}\right)$.

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:

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

As we have $\forall i = 1 \ldots M,\ 0 \leq S^i \leq K$ and $\mathbb{E}[V_1^M] = M$, we obtain by Hoeffding's inequality (Hoeffding, 1963):

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

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

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

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

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 $(\boldsymbol{s}, \boldsymbol{t})$ played with positive probability. Consider any minority player i. For each issue $k \in \mathcal{K}$, let $b_k = v_k^M(\boldsymbol{t}) - v_k^m(\boldsymbol{s})$ 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:

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

The term $(K - s^i_{k'})$ 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 $s^i_{k'} = K$, and we have obtained a contradiction. Or $s^i_{k'} < K$, and i has a profitable deviation; but then the initial profile is not an equilibrium. □

A.3 Proof of Proposition 1

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 $\dfrac{mK}{K+1} \geq K + 1 - \dfrac{MK}{K+1}$, 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 $(\boldsymbol{s}, \boldsymbol{t})$ for which

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

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. □

A.4 Proof of Example 1

With $K = 3$ and $(m, M) = (4, 5)$, we have

$$\max\left(K - \left\lfloor \frac{MK}{K+1} \right\rfloor, 1\right) = \max\left(3 - \left\lfloor \frac{15}{4} \right\rfloor, 1\right) = 1,$$
$$\min\left(\left\lfloor \frac{mK}{K+1} \right\rfloor, K - 1\right) = \min\left(\left\lfloor \frac{12}{4} \right\rfloor, 2\right) = 2,$$

so that we can choose $K_m = 2$ in the previous proof and obtain a pure strategy equilibrium for which the minority wins two of the three issues (with certainty).

A.5 Proof of Example 2

Note first that since $3 < 25/4$, Proposition 1 does not apply. Consider an arbitrary pure-strategy profile $(\boldsymbol{s}, \boldsymbol{t})$. For each issue $k \in \mathcal{K}$, let $b_k = v_k^M(\boldsymbol{t}) - v_k^m(\boldsymbol{s})$, 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.

(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 $(b_k)_{k=2\ldots4}$ 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 $(b_k)_{k=3\ldots4}$ 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. □

A.6 Proof of Proposition 2

Let K be even and M be odd.[37] We consider potential individual deviations from the profile $(\boldsymbol{\sigma^2}, \boldsymbol{\tau^1})$.

[37] We require K even, so that the strategy σ^2 is well-defined. Moreover, the reasoning we offer does not apply for M even. For instance, if $M = m + 1$, the most likely difference in votes a m-player wants to counter is $M - m + 1 = 2$, and allocating 3 votes on a subset of issues is a profitable deviation.

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:[38]

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

In what follows, we omit to mention the subscript k in the computations, as all the strategies are symmetric across decisions. As M is an odd number, we have for all $a \in \{1, \dots, K\}$:

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

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

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

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:

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

[38] By 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.

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:

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

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

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

A.7 Proof of Proposition 3

Let M be divisible by K,[39] and assume that $M \le mK$. This last assumption is made without loss of generality as, if $M > mK$, the profile $(\boldsymbol{\sigma}^K, \boldsymbol{\tau}^1)$ is trivially an equilibrium in which the majority wins all the decisions.

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

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

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

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

As M is a multiple of K, it is the only one in the set $\{M-2, \ldots, M-1+K\}$.

[39] Note that the result is valid only when M is divisible by K. If $M = aK + b$ with $2 \le b \le K-1$, any given majority player is useless playing σ^1 (any issue receives $aK+b-1$ votes from the other majority players, and a multiple of K votes from minority players), and consolidating a subset of issues is a profitable deviation. If $M = aK + 1$, a minority player i is decisive on an issue only when the other minority players allocate aK votes or $(a+1)K$ votes. As the first case is much more probable than the second one, one can show that i is better off playing σ^1 or σ^2, rather than σ^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 $a \in \{3, \ldots, K\}$, $p^i(a) - p^i(a-1) = 0$. We conclude that τ^1 is a best reply for player i.

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

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

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

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

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

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

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:

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

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:

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

We obtain

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

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 $M \geq mK/2$. □

A.8 Proof of Remark 3

Under the equilibrium of Proposition 2, when $M = m + 1$, we have by assumption m even. As each minority player allocates 2 votes on any targeted issue, and as the average number of votes of the minority group per issue is m, the scenario in which the minority group allocates exactly m balls on each issue realizes with positive probability. In this scenario, the minority wins no decision.

Under the equilibrium of Proposition 3, the number of majority votes per urn is equal to M, and it is divisible by K, the number of votes that each minority player allocates on her chosen issue. As the total number of votes of the majority exceeds the total number of votes of the minority, a possible scenario is one where the minority and the majority are tied on a given number of issues, while the other issues receive a majority of majority votes. If all ties are resolved in favor of the majority, the minority wins no decision.

As the majority group has a larger amount of votes than the minority, there must always be an issue with more votes from the majority than from the minority. Therefore, the minority can never win all decisions. □

A.9 Proof of Proposition 4

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

Let $\boldsymbol{\sigma}$ 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 $\boldsymbol{\tau}$ 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: $p_m(\boldsymbol{\sigma}, \boldsymbol{\tau}) > 0$. □

References

Abbink, K., J. Brandts, B. Herrmann, and H. Orzen (2010): "Intergroup conflict and intra-group punishment in an experimental contest game," *American Economic Review*, 100, 420–447.

Arad, A. (2012): "The Tennis Coach Problem: A Game-Theoretic and Experimental Study," *The BE Journal of Theoretical Economics*, 12, 1–43.

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.

Avrahami, J. and Y. Kareev (2009): "Do the weak stand a chance? Distribution of resources in a competitive environment," *Cognitive Science*, 33, 940–950.

Borel, E. and J. Ville (1938): *Applications de la théorie des probabilités aux jeux de hasard*, Jacques Gabay.

Bunwaree, S. and R. Kasenally (2005): *Political parties and democracy in Mauritius*, Johanesbourg: EISA.

Casella, A. (2005): "Storable votes," *Games and Economic Behavior*, 51, 391–419.

——— (2012): *Storable votes: protecting the minority voice*, Oxford University Press.

Casella, A., A. Gelman, and T. R. Palfrey (2006): "An experimental study of storable votes," *Games and Economic Behavior*, 57, 123–154.

Casella, A., T. Palfrey, and R. Riezman (2008): "Minorities and Storable Votes," *Quarterly Journal of Political Science*, 3, 165–200.

Cason, T. N., R. M. Sheremeta, and J. Zhang (2012): "Communication and efficiency in competitive coordination games," *Games and Economic Behavior*, 76, 26–43.

Chowdhury, S. M., D. Kovenock, and R. M. Sheremeta (2013): "An experimental investigation of Colonel Blotto games," *Economic Theory*, 52, 833–861.

Condorcet (1785): *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*, Paris: Imprimerie Royale.

Dahl, R. A. (1991): *Democracy and its Critics*, Yale University Press.

Dechenaux, E., D. Kovenock, and R. M. Sheremeta (2015): "A survey of experimental research on contests, all-pay auctions and tournaments," *Experimental Economics*, 18, 609–669.

Eifert, B., E. Miguel, and D. N. Posner (2010): "Political competition and ethnic identification in Africa," *American Journal of Political Science*, 54, 494–510.

Emerson, P. J. (1998): *Beyond the Tyranny of the Majority*, The de Borda Institute.

——— (1999): *From Belfast to the Balkans: Was Democracy Part of the Problem?*, The de Borda Institute.

Fiorina, M. P., S. J. Abrams, and J. C. Pope (2005): *Culture war?*, Pearson Longman New York.

Fischbacher, U. (2007): "z-Tree: Zurich toolbox for ready-made economic experiments," *Experimental economics*, 10, 171–178.

Fu, Q., J. Lu, and Y. Pan (2015): "Team contests with multiple pairwise battles," *The American Economic Review*, 105, 2120–2140.

Greiner, B. (2015): "Subject pool recruitment procedures: organizing experiments with ORSEE," *Journal of the Economic Science Association*, 1, 114–125.

Gross, O. and R. Wagner (1950): "A continuous Colonel Blotto game," Tech. rep., DTIC Document.

Grosser, J. and T. Giertz (2014): "Pork barrel politics, voter turnout, and inequality: an experimental study," *University of Cologne working paper.*

Hart, S. (2008): "Discrete Colonel Blotto and general lotto games," *International Journal of Game Theory*, 36, 441–460.

Hoeffding, W. (1963): "Probability inequalities for sums of bounded random variables," *Journal of the American statistical association*, 58, 13–30.

Hortala-Vallve, R. (2012): "Qualitative voting," *Journal of Theoretical Politics*, 24, 526–554.

Hortala-Vallve, 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.

Jacobson, G. C. (2008): *A Divider, Not a Uniter: George W. Bush and the American People: The 2006 Election and Beyond*, Longman Publishing Group.

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.*

Koriyama, Y., A. Macé, R. Treibich, and J.-F. Laslier (2013): "Optimal apportionment," *Journal of Political Economy*, 121, 584–608.

Kovenock, D. and B. Roberson (2012): "Coalitional Colonel Blotto games with application to the economics of alliances," *Journal of Public Economic Theory*, 14, 653–676.

Laslier, J.-F. (2012): "Why not proportional?" *Mathematical Social Sciences*, 63, 90–93.

Lijphart, A. (2004): "Constitutional design for divided societies," *Journal of democracy*, 15, 96–109.

May, K. O. (1952): "A set of independent necessary and sufficient conditions for simple majority decision," *Econometrica*, 680–684.

Picard, E. (1994): "Les habits neufs du communautarisme libanais," *Cultures et conflits*, 49–70.

Rae, D. W. (1969): "Decision-rules and individual values in constitutional choice," *American Political Science Review*, 63, 40–56.

Reynal-Querol, M. (2002): "Ethnicity, political systems, and civil wars," *Journal of Conflict Resolution*, 46, 29–54.

Rinott, Y., M. Scarsini, and Y. Yu (2012): "A Colonel Blotto gladiator game," *Mathematics of Operations Research*, 37, 574–590.

Roberson, B. (2006): "The Colonel Blotto game," *Economic Theory*, 29, 1–24.

Rogers, J. (2015): "An Experimental Investigation of Lobbying Strategies," *Mimeo*.

Sheremeta, R. M. (2015): "Behavior in group contests: A review of experimental research," *Mimeo*.

Winslow, C. (2012): *Lebanon: war and politics in a fragmented society*, Routledge.