Designing binary social decisions

Margarita Kirneva and Matías Nuñez [1]

Abstract

We design a mechanism, Majority voting with random checks, that fully implements the majority rule for binary social decisions. After a simultaneous vote over the two options, the winner must be confirmed by some agent from a random sample of agents voting sequentially. The mechanism incentivizes agents to act truthfully as a lottery is held if no agent confirms the outcome. We extend our results to incomplete information and abstention and introduce additional implementation mechanisms based on the concept of network formation.

1 Introduction

Majority voting is the most popular voting rule for binary decisions. It is often a synonym for voting: in legislatures and committees, agents often express their opinion for one of two available options, the most popular one being elected. Early axiomatized with simple and intuitive axioms (see [19], and [15] for a recent contribution), it plays a central role in analyzing democratic institutions.

Nevertheless, a non-negligible share of voters regret their choice (their vote or their abstention) in different sorts of elections.[2] While many narratives could justify these results, we consider the theoretical insights by [21] and [24] as a reasonable justification: preventing such bad outcomes is not easy since voting mechanisms are plagued with a multiplicity of equilibria, and in some of them, the chosen option might be the minority preferred one. The former shows that majority rule fails to satisfy Bayesian monotonicity (a necessary condition for implementation via deterministic mechanisms in the presence of incomplete information) while the latter proves that only dictatorships are implementable through voting mechanisms.[3] The logic behind these negative results is related to the following observation. Whenever a voting profile contains too many votes for the same option, this profile is an equilibrium irrespective of preferences since no agent is pivotal.

These inefficiencies make appealing the design of alternative voting rules.[4] The primary solution considered in the literature[5] when a single election is at stake is as follows: shifting from simultaneous voting to the sequential majority mechanism or roll-call voting, where voters express their preference for one of the two options one after the other, and the most

[1]We thank Antonin Macé, Jeff Ely, Yukio Koriyama, Jean-François Laslier, Hervé Moulin, Tom Palfrey, Marzena Rostek, Fedor Sandomirskiy and Siyan Xiong as well as three anonymous reviewers for their useful remarks and comments as well as conference participants at the Asian School in Economic Theory (Singapore 2022) and Université Paris Dauphine. This research is supported by a grant of the French National Research Agency (ANR), "Investissements d'Avenir" (LabEx Ecodec/ANR-11-LABX-0047).

[2]See [6] and [8] for empirical evidence that regret matters in electoral settings. In particular, these authors point out that an "overwhelming majority of those who voted in an election feel, ex post, that they made the right decision, while non-voters are less certain about the correctness of their choice to abstain". This justifies our emphasis on voting mechanisms where abstention is possible.

[3]A voting mechanism allows each agent to vote for each option and is monotonic in the usual sense: if x wins at some profile and gets additional support from some agent, x remains the winner. For instance, delegation mechanisms (such as [12] and [13]) belong to the category of voting mechanisms.

[4]An optional approach is to modify the voters' rationality as for instance [21] who explore implementation under weakly undominated strategies. While this approach is sound, it would imply that voters never abstain in equilibrium since voting for one's favorite option weakly dominates abstention.

[5]See [14] for a review of the literature adopting a different perspective: namely, addressing alternatives to majority rule that take into account preference intensities such as voting rules with transfers (quadratic voting or markets for votes) or voting over several issues (storable votes).

popular option wins.[6] Roll-call voting subgame perfect implements the majority-preferred option as long as agents' reasoning relies upon backward induction. Recent experimental findings (see [9]) suggest that while theoretically different, sequential and simultaneous majority voting do not differ much; note that both systems co-exist in the U.S. state legislatures.[7] Moreover, as shown by [22], in the U.S. Senate, where roll-call voting is used, early movers tend to drop the party line to free-ride on late movers; this, in turn, can lead to failures in implementing the majority preferred option.

Do other mechanisms guarantee that the majority-preferred option is the only equilibrium winner with fewer steps? If one focuses on simultaneous mechanisms, the contribution of [18] ensures their existence, as majority voting fulfills the sufficient conditions for being Nash implementable. However, the main objection to such result is that the mechanism (integer game) achieving the implementation is hard to apply in practice. This objection is reinforced by the recent contribution of [24], as discussed above, since no known voting mechanism implements the majority option.

We contribute to this literature by providing a positive result: we design a mechanism that implements the majority option, the Majority voting with random checks. This mechanism relies on the use of lotteries off-equilibrium[8] to incentivize agents to vote for their preferred alternative. To deal with lotteries, we impose the mild restriction of stochastic dominance, virtually satisfied by all preference extensions. The timing of Majority voting with random checks is as follows. Initially, the $2p+1$ agents vote simultaneously for one of the two options. After the vote totals are revealed, a second stage takes place to determine whether the winner of the first round should be the outcome. Each agent (out of the $p+1$ randomly selected ones) sequentially declares whether she agrees with the winner. If any of these agents agrees, the winner of the first round is the outcome; otherwise, the winner is determined by a lottery by randomly picking one of the initial votes. This lottery's role is to give proper incentives to agents to reveal their preferred alternative. The Majority voting with random checks implements the majority rule under complete and incomplete information.[9] Observe that this mechanism requires at most $p+2$ stages (voting stage and $p+1$ voters). We also extend the mechanism to a setting where agents can abstain in any of the stages making the strategic problem richer.[10]

Our second set of contributions is the design of two alternative mechanisms achieving the same implementation goal: the Bloc formation mechanism and the Verification one. The Bloc formation mechanism is the first simultaneous voting-based mechanism, different than an integer game, that Nash implements majority rule. The mechanism is not a pure voting mechanism because agents do not have a strategy that consists exclusively of voting

[6] The introduction of supermajorities or qualified majority rules where one option needs more than half of the votes to win is often considered. The rationale behind such a solution is clear: to avoid events where both candidates get the same popular support and the winner is considered illegitimate, one should raise the threshold required to be called the winner. This class of rules does not escape from the impossibility described by [24] and thus also admits inefficient equilibria.

[7] In common value settings, roll-call voting has a main drawback: the dynamic structure might trigger informational cascades (see [1] for a review of the experimental literature on cascades).

[8] Mechanisms with off-equilibrium lotteries are known to be more permissive than deterministic mechanisms. See [4], [7] and [17] for recent contributions. [10] develop a similar idea to show that one can achieve Pareto improvements over random dictatorship through simultaneous mechanisms. In the related framework of the Condorcet jury theorem, [16] proposes the introduction of a "slightly randomized majority rule" to ensure that the unique equilibrium is informative. Our mechanism shares a similar spirit to this idea but without randomization in equilibrium. Likewise, [20] and [2] describe similar ideas for large populations of agents. See also [23] for a characterization of ordinal and onto choice rules that are subgame perfect implementable.

[9] We focus on implementation in sequential equilibria (see [3]). Implementation in extensive form games is also possible via other concepts such as Perfect bayesian (see [11]) or simply relying on belief restrictions (see [5]).

[10] We thank Omer Tamuz for suggesting this extension.

for one of the two options. Instead, the Bloc formation mechanism requires each agent to vote for one of the two options *and* to nominate p other agents.[11] This mechanism's definition actually allows us to see a voting profile as a directed network where vertices are the agents and the edges their nominations.[12] The outcome of the vote depends on the network structure generated by the voting profile. Finally, we design the Verification mechanism, a voting mechanism in the sense that voters are just required to vote for one of the options. More specifically, these mechanisms request agents to vote for an option in the first stage and, in case some agent disagrees with the first stage majority outcome, to confirm their choice to ensure no coordination problem exists. The mechanism also relies on the bloc formation as its simultaneous counterpart and uses lotteries to increase agents' pivotality. We show that this mechanism subgame perfect implements the majority option under complete and incomplete information.

This work is organized as follows. After laying out the model in Section 2, Section 3 analyzes the Majority voting with Random checks. Section 4 considers the Bloc formation and the Verification mechanisms. Section 5 concludes, and the Appendix contains some proofs.

2 Model

We consider a finite set $I = \{1, \ldots, n\}$ of agents, with generic element i, who need to choose one option out of the set $A = \{a, b\}$. Each agent has strict complete and reflexive preferences over A where aR_ib denotes that a is strictly preferred to b. We denote by $\Theta = \{\theta_a, \theta_b\}$ the type space for each agent. A vector $\theta^n \in \Theta^n$ denotes the profile of all agents' preferences. An agent i is of type θ_a (resp. θ_b) if and only if aR_ib (resp. bR_ia). A social choice correspondence (SCC) is a mapping $f : \Theta^n \rightrightarrows A$ that selects a non-empty set of options for each profile θ^n. The majority rule is the SCC that selects $Maj(\theta^n)$ the majority preferred option(s) of the profile θ^n. With $2p+1$ agents (p being a non-negative integer), the majority rule is uniquely defined since $Maj(\theta^n) = x$ if and only if $|\{i \in I \mid xR_iy\}| \geq p + 1$.[13]

We let Δ denote the set of lotteries over A with $\Delta = \{\beta : A \to [0, 1] \mid \sum_A \beta(a) = 1\}$. A *mechanism* is a function $g : M \to \Delta$ that assigns to every $m \in M$ a unique element of Δ, where $M = \prod_{i \in I} M_i$, and M_i is the strategy space of agent i.

We assume that preferences over lotteries satisfy stochastic dominance (**SD**). A lottery β stochastically dominates lottery η if and only if β yields at least as much expected utility as η for any von-Neumann Morgenstern utility representation consistent with the ordinal preferences.

Consider two lotteries $\beta, \eta \in \Delta$ and let x be the most preferred option of agent i. **SD** requires that the agent prefers β to η if and only if the probability with which β selects x is greater or equal than the probability that η selects x. In formal terms, we can write for any agent i that:

$$\beta \tilde{R}_i^{\mathbf{SD}} \eta \iff \beta(x) \geq \eta(x) \text{ and } \beta R_i^{\mathbf{SD}} \eta \iff \beta(x) > \eta(x),$$

where $\beta \tilde{R}_i^{\mathbf{SD}} \eta$ means that agent i weakly prefers β to η and $\beta R_i^{\mathbf{SD}} \eta$ implies that she strictly prefers the former to the latter.

[11] For any positive integer p, if n denotes the number of agents with $n = 2p + 1$, an agent nominates p agents. Note when an agent nominates p agents, a majority agent is always among the nominees since there at most p minority agents.

[12] For completeness, we also provide an equivalent definition of the Majority voting with random checks as a network formation game.

[13] The paper is written assuming an odd number of agents to ease notation. However, several of the results extend to the case with an even number of them. If there are $2p$ agents, the majority preferred candidate need not be unique since $Maj(\theta^n) = x$ if $|i \in I \mid xR_iy| \geq p + 1$ and $Maj(\theta^n) = \{x, y\}$ if $|i \in I \mid xR_iy| = p$.

A mechanism specifies a game-form: this means that, when the mechanism is coupled with preferences over options for each of the agents, it defines a normal-form or extensive-form game. For a mechanism g, let $\text{NE}^g(\theta^n)$ and $\text{SPNE}^g(\theta^n)$ respectively denote the set of Nash equilibria and subgame perfect equilibria when the preference profile is θ^n. A mechanism Nash (resp. subgame perfect) implements a social choice correspondence f if for any θ^n, the outcome of any member of $\text{NE}^g(\theta^n)$ (resp. $\text{SPE}^g(\theta^n)$) is an element of $f(\theta^n)$ and any element of $f(\theta^n)$ is the outcome of some member of $\text{NE}^g(\theta^n)$ (resp. $\text{SPE}^g(\theta^n)$). A similar idea applies to the incomplete information setting where the equilibrium notion upon which we rely is sequential equilibrium (a formal definition is skipped).

3 Majority voting with Random checks

This section introduces the main contribution of this work: Majority voting with Random checks (χ_{RC} in the sequel). Its formal definition follows.

Timing.
Voting stage: Each agent i votes for an option $v_i \in A$. The profile of votes v is publicly announced and the option which gets most votes in v is denoted the winner of the Voting stage.
Confirmation stage: An order of agents $\pi = (\pi_1, \ldots, \pi_n)$ is selected through a uniform draw. At each stage $t \in \{1, \ldots, p+1\}$, agent π_t announces Y or N.

At each stage t, the mechanism ends if the nominating agent π_t announces Y with the outcome being the winner (most announced option) of the Voting stage. Otherwise, the mechanism proceeds to stage $t+1$ and the next agent announces Y or N. If all agents in $\{\pi_1, \ldots, \pi_{p+1}\}$ announced N in the Confirmation stage the outcome is a lottery $\beta(v)$ determined by profile v. This lottery assigns to each option its share of first-stage votes, so that $\beta_a(v) = \frac{|\{i \in I | v_i = a\}|}{n}$ and $\beta_b(v) = 1 - \beta_a(v)$. The purpose of this lottery is to give incentives to agents to vote for their most preferred option.

Notice that the Confirmation stage announcements Y and N can be viewed as agreement and disagreement with the Voting stage outcome respectively. That is, if at least one agent in $\{\pi_1, \ldots, \pi_{p+1}\}$ agrees with the outcome being the majority winner of the Voting stage, this option is the outcome. On the other hand, if no one agrees, the outcome is the previously mentioned lottery.

We now establish the implementation under complete and incomplete information and discuss the extension to abstention.

3.1 Complete information

Before starting the formal argument with complete information, we present an example that illustrates the logic under this mechanism.

Example 1: Preferences are given by aR_ib for $i = 1, 2, 3$ and bR_ia for $i = 4, 5$. Consider any strategy profile in which $v_i = b$ for every agent i; the rules of the mechanism imply that b is the outcome, independently of the Confirmation stage. This profile is not an equilibrium since b is not majority preferred. More precisely, any agent with type θ_a has a profitable deviation. Consider, for instance, that agent 1 deviates and votes $v_1' = a$ rather than $v_1 = b$. The outcome after such deviation depends on the agents selected for the Confirmation stage. If either 4 or 5 is selected, b is the outcome since both agents prefer b to a and, thus, in any SPE of the Confirmation stage at least one of them announces Y. If neither 4 or 5 is in $\{\pi_1, \ldots, \pi_{p+1}\}$, the Confirmation stage involves only agents $\{1, 2, 3\}$. All of them strictly prefer to announce N if Y was not announced before. Thus, in the unique

SPE of the Confirmation stage, N is announced by all three agents and the outcome of the mechanism is a lottery that selects a with probability 1/5 (random dictatorship in the profile (a, b, b, b, b)). Table 1 illustrates this example. The left part represents the Voting stage profiles: unanimous in the first case, and after the deviation of agent 1 afterwards. The right part illustrates a possible SPE of the Confirmation stage given Voting stage profiles.

Voting stage					Confirmation stage					Outcome
1	2	3	4	5	1	2	3	4	5	
Unanimous vote for b, $\{\pi_1, \pi_2, \pi_3\} = \{1, 2, 3\}$										
b	b	b	b	b	N	N	N	-	-	b
Deviation to $v'_1 = a$, $\{\pi_1, \pi_2, \pi_3\} = \{1, 2, 5\}$										
a	b	b	b	b	N	N	-	-	Y	b
Deviation to $v'_1 = a$, $\{\pi_1, \pi_2, \pi_3\} = \{1, 2, 3\}$										
a	b	b	b	b	N	N	N	-	-	$1/5a + 4/5b$

Table 1: Majority voting with Random checks

Therefore, by deviating from $v_1 = b$ to $v'_1 = a$, agent 1 induces a lottery that assigns a a positive probability; by **SD**, agent 1 prefers to deviate showing that any strategy profile in which $v_i = b$ for every agent i cannot be an equilibrium.

A similar logic to the one described in the example shows that at least $p+1$ agents who prefer the majority option are sincere in the Voting stage of *any* equilibrium which leads to the implementation result, stated formally as follows.

Proposition 1. *The Majority voting with Random checks subgame perfect implements the majority rule.*

Proofs of all the results can be found in the Appendix.

3.2 Incomplete information

We now prove that the Majority voting with Random checks implements the majority rule when we relax the assumption of complete information. For simplicity of the argument we assume that each agent believes that the types of other agents are *i.i.d.* and assigns probability q_a (resp. $1 - q_a$) to each agent being θ_a (resp. θ_b). Later we show that the *i.i.d.* assumption can be relaxed without affecting the result. A strategy for an agent i is a mapping $\sigma_i = (\sigma_i^1, \sigma_i^2)$ where $\sigma_i^1 : \theta \to A$ and $\sigma_i^2 : \theta \times H_{i-1} \to \{Y, N\}$ stand for the strategies in each of the stages and H_{i-1} is the set of all possible histories before the Confirmation stage announcement of agent i. A first-stage vote of agent i, v_i^1, is revealing given strategy σ_i^1 if either $(\sigma_i^1)^{-1}(v_i^1) = \theta_a$ or $(\sigma_i^1)^{-1}(v_i^1) = \theta_b$, in other words, if after observing v_i^1 other agents learn the type of agent i.

Proposition 2. *Under incomplete information, Majority voting with Random checks implements the majority rule in sequential equilibrium.*

Notice that our initial assumption on prior beliefs being *i.i.d.* was unnecessarily demanding. If the prior beliefs satisfy the following weaker conditions, the result remains valid:

- Each agent assigns a positive probability to the event T_a (the event where p agents other than i have type θ_a and p agents have type θ_b);

- For any subset of agents $I' \subset I$ and any agent $j \in I \setminus I'$, for any $x, y \in A$, $x \neq y$, $Pr(\forall i \in I', i \text{ is of type } \theta_x \mid j \text{ is of type } \theta_y) > 0$.

3.3 Abstention

We now discuss a second extension of the Majority voting with Random checks mechanism where we allow the agents to abstain. The possibility of abstention makes the strategic problem richer. Indeed, the abstention of many agents of the majority type can induce the victory of the minority and make agents indifferent between abstaining or voting for any of the options.[14] In order to deal with abstention, we extend the definition of the mechanism as follows.

Timing.
Voting stage: Each agent i votes for an option or abstains, that is $v_i \in A \cup \emptyset$. The profile of votes v is publicly announced. We call the option which gets most votes in v the winner of the Voting stage. If no agent participates, i.e. $v_i = \emptyset$ for all $i \in I$ or there is a tie between the two alternatives, the outcome is a lottery which assigns probability 1/2 to each of the options.
Confirmation stage: An order of agents $\pi = (\pi_1, \ldots, \pi_n)$ is selected through an uniform draw (from all the agents independent of whether they participated or abstained in the Voting stage). At each stage $t \in \{1, \ldots, p+1\}$, agent π_t announces Y or N or abstains.

At each stage t, the mechanism ends if the confirming agent π_t announces Y for some $t \in \{1, \ldots, p+1\}$ with the outcome being the winner of the Voting stage. If the two options are tied in the Voting stage the outcome is the lottery which assigns probability 1/2 to each option. If all agents in $\{\pi_1, \ldots, \pi_{p+1}\}$ announced N or abstained in Confirmation stage the outcome is a lottery $\beta(v)$, which assigns to each option its share of first-stage votes, so that $\beta_a(v) = \frac{|\{i \in I | v_i = a\}|}{n}$ and $\beta_b(v) = 1 - \beta_a(v)$.

Proposition 3. *The Majority voting with Random checks subgame perfect implements the majority rule in the presence of abstention.*

The proof of the result can be found in the Appendix. As a final comment on this mechanism, observe the existence of the following equilibrium. Consider a strategy profile with only two votes for A (the rest being abstentions) in the Voting stage and all agents announcing Y in the Confirmation stage. This is a subgame perfect equilibrium for any preference profile where A is the majority preferred option. Indeed, no deviation is possible in the Voting stage since the winner is not altered by adding or substracting one vote. In the Confirmation stage, by definition, there is some agent preferring A among the $p+1$. Since this agent's best response is to vote Y , the outcome is A independently of the rest of the votes in the Confirmation stage. Thus, only two agents participating in the voting stage suffice to guarantee that the majority option is elected.

4 Alternative implementing mechanisms

In this section, we introduce two mechanisms, beyond Majority voting with random checks, that implement the majority rule. The first one, the Bloc formation mechanism, is the first simultaneous mechanism (beyond integer games) that implements the majority rule. The second one, the Verification mechanism, allows agents to verify their vote after an initial vote. Note that both mechanisms are based on the idea of network formation.

[14]If one considers the simultaneous majority mechanism, any strategy profile in which all majority agents abstain and two minority agents vote sincerely is a Nash equilibrium irrespective of preference.

4.1 Bloc formation mechanism

4.1.1 Simultaneous blocs

In the Bloc formation mechanism, the message m_i of agent i consists of (1.) a vote for an option v_i and (2.) a nomination of p agents excluding herself (denoted c_i). Formally, the mechanism is denoted $\chi_{BF} : M \to \Delta$ with, for all $i \in I$, $M_i := A \times 2_p^{-i}$ where 2_p^{-i} denotes the sets of p agents different from i where $|I| = 2p + 1$.

The central notion of this mechanism is the idea of a bloc of agents. A bloc in favor of option x is a majority group of agents, denoted B, such that each agent votes for x while only nominating agents in B. This can be defined as follows.

Definition 1. *A set B of agents forms a **bloc** in favor of option x in the profile m if:*

1. *$|B| \geq p + 1$ (majority group),*
2. *$v_i = x$ for each $i \in B$ (only votes for x),*
3. *$c_i \subset B$ for each $i \in B$ (only nominations in B).*

The outcome of the Bloc formation mechanism χ_{BF} depends on whether a bloc forms in the message profile or not. Denote by B^m the set of blocs formed in profile m. By definition, all blocs in a profile (if any) favor the same option since each bloc contains a majority of agents. Therefore, for any profile m in which there is a bloc in favor of option x, $\chi_{BF}(m) = x$.

If the profile m does not contain a bloc, the outcome is a lottery over A that depends on m. For each option $x \in A$, the weight of x associated to m, denoted $\eta^x(m)$, equals:

$$\eta^x(m) = \sum_{i \in I} \eta_i(m) \mathbb{1}\{v_i = x\} \text{ with } \eta_i(m) = \frac{|\{j \in I \setminus \{i\} \mid i \in c_j\}|}{np}.$$

To see the logic behind this formula, we let $\eta_i(m)$ be the weight of agent i, that is the share of nominations of i in the total nominations np. By construction, $\sum_{i \in I} \eta_i(m) = 1$ for any $m \in M$. When all the other agents nominate i, agent i has the maximal possible weight of $\eta_i(m) = \frac{n-1}{np}$ whereas $\eta_i(m) = 0$ when none of the other agents nominate i.

We thus interpret $\eta^x(m)$ as the sum of the weights of the agents who vote for x so that, by construction, $\eta^a(m) + \eta^b(m) = 1$. Notice that the weight $\eta^x(m)$ is strictly increasing in the number of nominations for x-agents and, thus, in the number of x-agents among nominated agents.

The previous rules of the mechanism can be summarized as follows. For each message profile m, the outcome of the Bloc formation mechanism χ_{BF} coincides with:

$$\chi_{BF}(m) = \begin{cases} a & \text{if } m \text{ admits a bloc in favor of } a, \\ b & \text{if } m \text{ admits a bloc in favor of } b, \\ \eta(m) & \text{otherwise.} \end{cases}$$

A noteworthy comment on the Bloc formation mechanism deals with its strategic implications. Deciding for which option to vote should be easy. This is the case as we show next. For any agent i with xR_iy, any nomination $c_i \in 2_p^{-i}$ and any message m_{-i} of the agents different from i, agent i weakly prefers to vote x, that is to vote for her most preferred option since:

$$\chi_{BF}(x, c_i, m_{-i}) \tilde{R}_i^{SD} \chi_{BF}(y, c_i, m_{-i}).$$

To see why there is a preference of voting honestly, observe that for a given announcement profile of the rest of the agents, the following cases arise. In the case that the outcome is a lottery (i.e. no bloc), this is immediate since she strictly prefers to be sincere. In the case that the agent is part of a bloc, it is always better for an agent to vote for her preferred option (it can create a bloc, prevent its formation or have no impact). Finally, if the profile admits a bloc independently of the agent's behavior, agent i's vote is irrelevant. Similarly, there is a weak preference for nominating agents who voted for one's preferred option.

Notice that for any message profile $m = (v, c)$ the nomination profile c creates a directed graph in which vertices are the agents and the edges - their nominations. In the Appendix we provide an alternative formulation of the mechanism which regards blocs from the perspective of graph theory.

4.1.2 Nash implementation

The main result of this section is as follows.

Proposition 4. *The Bloc formation mechanism Nash implements the majority rule.*

The formal proof of Proposition 4 is included in the Appendix, we give some intuition for it in the next lines. Let a be the majority preferred option and b the minority one (assume that the majority is uniquely defined for simplicity).

First, it is rather intuitive that an equilibrium selecting a exists. By definition, there are at least $p + 1$ such agents. If these agents vote in favor of a and nominate each other, this constitutes a bloc in favor of it and hence an equilibrium. Indeed, no agent within the bloc wants to deviate (she obtains her most preferred outcome) and no agent outside the bloc can alter the outcome (by definition).

Second, suppose that the outcome is a lottery with full support. Since the weight with which each option wins is strictly increasing in the votes it obtains from the agents with positive weight, the agents find it optimal to behave sincerely and vote for their most preferred option. This, in turn, leads to the formation of a bloc in favor of a.

Finally, we argue that no bloc can be formed in favor of b, the minority option. Assume, by contradiction, that such bloc exists. Then, as shown in 7, an effective bloc exists and includes some majority agent. Then, by Proposition 8 a majority agent - member of the effective bloc can break all the blocs in favor of b in the profile leading to a lottery being the outcome. Such deviation is profitable for a majority agent, contradicting the existence of an equilibrium in which a bloc in favor of b is formed. The next example illustrates the logic of the mechanism on this precise point.

Example 2: The preferences of the agents ($N = \{1, 2, 3, 4, 5\}$) are such that aR_ib for $i = 1, 2, 3$ and bR_ia for $i = 4, 5$ so that a is majority-preferred. Remark that no equilibrium profile admits a bloc in favor of b. Indeed, let $m = (c, v)$ be the profile where each agent votes b (i.e. $v_i = b \forall i$) and nominations are as follows: $c_1 = 4, 5$, $c_2 = 1, 3$, $c_3 = 1, 2$, $c_4 = 1, 5$ and $c_5 = 1, 4$. The profile m admits two blocs: $\{1, 4, 5\}$ and $\{1, 2, 3, 4, 5\}$. The bloc $\{1, 4, 5\}$ is the effective one since one cannot find a smaller group that nominate each other while voting b. If any agent $i \in \{1, 4, 5\}$ deviates to $m'_i = (a, c_i)$, the profile (m'_i, m_{-i}) admits no bloc and the outcome is a lottery between a and b. Since agent 1 prefers a to b, she has a profitable deviation and thus the profile m is not an equilibrium.

More generally, take any profile in which a bloc is formed in favor of b. Any bloc consists of at least $p+1$ agents, thus, it includes at least one majority agent. Since the effective bloc is included in the rest of the blocs, we can always find a majority agent with a profitable deviation. Thus no bloc in favor of the minority preferred option is possible.

4.2 Verification mechanism

This final section presents the Verification mechanism. This mechanism, that we denote χ_{DV}, incentivizes the formation of blocs in favor of the majority option and it leads to a lottery in the absence of blocs. The main advantage is that agents are treated symmetrically (all agents vote either once or twice) its main flaw is that other equilibria might involve more than n stages.

In the first stage, agents vote simultaneously for one of the two options. The votes are announced at the end of the stage. In the second stage, agents are asked to ratify the result. If all ratify the result, the outcome is the majority winner of the first stage. If some agent fails to ratify, a verification stage starts. The verification stage asks agents to vote again for one of the two options, this time in a sequentially in an exogenous commonly known order. If at least $p+1$ agents of the ones who voted for the majority winner of the first stage confirm their votes, the outcome is the winner of the first stage. Otherwise (namely many agents do not confirm their vote), the outcome is a lottery with endogenously determined weights. As we show in the sequel, the χ_{DV} mechanism admits several equilibria, all of which select the majority preferred option. Among its main advantages, it admits equilibria without verification, that is two-stage equilibria.

A formal definition of the χ_{DV} mechanism follows. For each arbitrary order π of agents, let $\chi_{DV} : M \to \Delta$ be a mechanism with $m_i = (v_i^1, \omega_i, v_i^2) \in M_i : A \times \{R, V\} \times A$ being the message of agent i. While an order of agents is required in stage 3, the equilibrium outcome is order-independent. The game stages are as follows.

Timing.
Voting stage: Each agent simultaneously announces $v_i^1 \in A$. Profile v^1 is publicly announced at the end of the stage.
Ratification stage: Each agent i simultaneously announces $\omega_i \in \{R, V\}$ to state whether she ratifies the outcome (R) or wants a verification (V). If all ratify, the outcome is the majority winner in the profile v^1. Profile ω is publicly announced at the end of the stage. If some agent votes V, stage 3 starts.
Verification stage: Each agent i sequentially announces $v_i^2 \in A$ according to the order π knowing announcements $v_{\pi_1}^2, v_{\pi_2}^2, \ldots, v_{\pi_{i-1}}^2$ of the predecessors.

The outcome of the Verification mechanism χ_{DV} for each message profile m is as follows.
A. If all agents ratify the voting outcome in stage 2, the outcome is the majority winner of the voting profile v^1.
B. If some agent asks for a verification in stage 2, the outcome depends on the votes in Stage 3.
B.1. If there is a bloc in favor of some option x, then x is selected. A bloc in favor of option x is formed in m if there is a set J of at least $p+1$ agents with $v_j^1 = x$ and $v_j^2 = x$ $\forall j \in J$. That is, a bloc in favor of x consists of a majority of agents voting twice for x.
B.2. In the absence of a bloc in the profile m, the outcome is a lottery β over A. For each profile m, let $s_a(m) = |\{j \in I \mid v_j^1 = a\}|$ be the votes for a in stage 1 and $\phi_a(m) = |\{j \in I \mid v_j^2 = a\}|$ be the votes for a in stage 3. In this case, the outcome of the profile m is:

1. a if $s_a(m) = n$ or $\phi_a(m) = n$ and $s_a(m) > 0$,
2. b if $s_b(m) = n$ or $\phi_b(m) = n$ and $s_b(m) > 0$,
3. otherwise is a lottery $\beta(m)$ with probabilities:

$$\beta_a(m) = \frac{\phi_a(m)\frac{s_a(m)}{s_b(m)}}{\phi_a(m)\frac{s_a(m)}{s_b(m)} + \phi_b(m)\frac{s_b(m)}{s_a(m)}} \text{ and } \beta_b(m) = 1 - \beta_a(m). \tag{1}$$

Remark that the probability $\beta_a(m)$ is strictly increasing in $s_a(m)$ and in $\phi_a(m)$ and strictly decreasing in $s_b(m)$ and in $\phi_b(m)$. This probability being increasing plays a key role in the sequel since, whenever the outcome is a lottery, an agent has a strict incentive to vote sincerely.

Our results consider majority rule implementation through χ_{DV} under complete and incomplete information.

Proposition 5. *For any order of the agents in the verification stage, the Verification mechanism subgame perfect implements the majority rule and admits equilibria where no verification is required.*

We now extend the implementation result of the Verification mechanism to an incomplete information setting. We assume that each agent believes that the types of other agents are *i.i.d.* and assigns probability q_a (resp. $1 - q_a$) to each agent being θ_a (resp. θ_b). Notice that the *i.i.d.* assumption can be relaxed for the same reasons as discussed in the analysis of the Majority voting with Random checks. A strategy for an agent i is a mapping $\sigma_i = (\sigma_i^1, \sigma_i^2)$ where $\sigma_i^1 : \theta \to A$ and $\sigma_i^2 : \theta \times H_{i-1} \to A$ stand for the strategies in each of the steps and H_{i-1} is defined as in the complete information case. We denote the message profile by $m = (v^1, v^2) \in M = A^{|I|} \times A^{|I|}$. A first-stage message of agent i, v_i^1, is revealing given strategy σ_i^1 if either $(\sigma_i^1)^{-1}(v_i^1) = \theta_a$ or $(\sigma_i^1)^{-1}(v_i^1) = \theta_b$, in other words, if after observing v_i^1 other agents learn the type of agent i.

Proposition 6. *Under incomplete information, the Verification mechanism implements the majority rule in sequential equilibria.*

5 Conclusion

We have considered the implementation of the majority rule in a preference aggregation setting through several mechanisms. The key idea of this work is that introducing lotteries in the mechanisms removes the undesirable equilibria widespread in the voting literature and helps design shorter mechanisms. Lotteries allow agents to be pivotal more often, while, in all of the considered mechanisms, lotteries do not arise in equilibrium so that the socially desirable option is always selected. We leave it to future empirical research to determine to which extent the proposed mechanisms outperform the classic voting procedures. Finally, some of our proposals consider voting as a network formation game; this deserves further investigation in the context of implementation and coalition formation.

References

[1] Lisa R Anderson and Charles A Holt. Information cascade experiments. Handbook of experimental economics results, 1:335–343, 2008.

[2] Eduardo M Azevedo and Eric Budish. Strategy-proofness in the large. The Review of Economic Studies, 86(1):81–116, 2019.

[3] Sandeep Baliga. Implementation in economic environments with incomplete information: the use of multi-stage games. Games and Economic Behavior, 27(2):173–183, 1999.

[4] Jean-Pierre Benoît and Efe A Ok. Nash implementation without no-veto power. Games and Economic Behavior, 64(1):51–67, 2008.

[5] James Bergin and Arunava Sen. Extensive form implementation in incomplete information environments. Journal of Economic Theory, 80(2):222–256, 1998.

[6] André Blais, Fernando Feitosa, and Semra Sevi. Was my decision to vote (or abstain) the right one? Party Politics, 25(3):382–389, 2017.

[7] Olivier Bochet. Nash implementation with lottery mechanisms. Social Choice and Welfare, 28(1):111–125, 2007.

[8] Damien Bol, André Blais, and Jean-François Laslier. A mixed-utility theory of vote choice regret. Public Choice, 176:461–478, 2018.

[9] Friedel Bolle and Philipp E Otto. Voting behavior under outside pressure: promoting true majorities with sequential voting? Social Choice and Welfare, 58(4):711–740, 2022.

[10] Tilman Börgers and Doug Smith. Robust mechanism design and dominant strategy voting rules. Theoretical Economics, 9(2):339–360, 2014.

[11] Sandro Brusco. Perfect bayesian implementation. Economic Theory, 5:419–444, 1995.

[12] Joseph Campbell, Alessandra Casella, Lucas de Lara, Victoria R Mooers, and Dilip Ravindran. Liquid democracy. two experiments on delegation in voting. Technical report, National Bureau of Economic Research, 2022.

[13] Amrita Dhillon, Grammateia Kotsialou, and Dimitris Xefteris. Information aggregation with delegation of votes. Technical report, 2021.

[14] Jacob K Goeree, Philippos Louis, and Jingjing Zhang. Improving on simple majority voting by alternative voting mechanisms. Oxford Research Encyclopedia of Economics and Finance, 2020.

[15] Sean Horan, Martin J Osborne, and M Remzi Sanver. Positively responsive collective choice rules and majority rule: a generalization of may's theorem to many alternatives. International Economic Review, 60(4):1489–1504, 2019.

[16] Jean-François Laslier and Jörgen W Weibull. An incentive-compatible condorcet jury theorem. The Scandinavian Journal of Economics, 115(1):84–108, 2013.

[17] Jean-François Laslier, Matías Núñez, and M. Remzi Sanver. A solution to the two-person implementation problem. Journal of Economic Theory, 194:105261, 2021.

[18] Eric Maskin. Nash equilibrium and welfare optimality. The Review of Economic Studies, 66(1):23–38, 1999.

[19] Kenneth O May. A set of independent necessary and sufficient conditions for simple majority decision. Econometrica, pages 680–684, 1952.

[20] Matías Núñez and Marcus Pivato. Truth-revealing voting rules for large populations. Games and Economic Behavior, 113:285–305, 2019.

[21] Thomas R Palfrey and Sanjay Srivastava. Mechanism design with incomplete information: A solution to the implementation problem. Journal of Political Economy, 97(3): 668–691, 1989.

[22] Jörg L. Spenkuch, Pablo Montagnes, and Daniel B. Magleby. Backward induction in the wild? evidence from sequential voting in the us senate. American Economic Review, 108(7):1971–2013, 2018.

[23] Hannu Vartiainen. Subgame perfect implementation of voting rules via randomized mechanisms. Social Choice and Welfare, 29:353–367, 2007.

[24] Siyang Xiong. Designing referenda: An economist's pessimistic perspective. Journal of Economic Theory, 191:105133, 2021.

Margarita Kirneva
CREST & Ecole Polytechnique
Palaiseau, France
Email: margarita.kirneva@polytechnique.edu

Matías Nuñez
CREST, CNRS & Ecole Polytechnique
Palaiseau, France
Email: matias.nunez@polytechnique.edu

A Proof of Proposition 1

We solve the game backwards starting from Confirmation stage. Notice first that in case the profile v is unanimous, Confirmation stage does not affect the outcome. Assume that v is not unanimous and denote by x the winning option of the Voting stage and by y the remaining one.

Assume that no agent from $\{\pi_1, \ldots, \pi_p\}$ announced Y and consider agent π_{p+1}. If $xR_{\pi_{p+1}}y$ the unique best response is Y and it is N otherwise.

Consider now agent π_p and assume no Y was announced before. Then the best response of π_p is:
- Y if $xR_{\pi_p}y$ and $yR_{\pi_{p+1}}x$;
- N if $yR_{\pi_p}x$ and $yR_{\pi_{p+1}}x$;
- $\{N, Y\}$ otherwise.

This logic can be extended to earlier agents in π in the following way. For any agent π_i with $i \leq p+1$ the best response at Stage 2 is:
- Y if $xR_{\pi_i}y$ and $yR_{\pi_j}x$ for all $i < j \leq p+1$;
- N if $yR_{\pi_j}x$ for all $i \leq j \leq p+1$;
-$\{N, Y\}$ otherwise.

Then the SPE outcome of Stage 2 is the following one.

Lemma 1. *For any non-unanimous profile v of the Voting stage, the SPE outcome of the Confirmation stage is:*
- x (the winner of Voting stage) if xR_iy for some $i \in \{\pi_1, \ldots, \pi_{p+1}\}$;
- the lottery $\beta(v)$ otherwise.

Consider now the Voting stage where we let a denote the majority preferred option and b the minority preferred one. Notice first, that if a is the winner of the Voting stage, the equilibrium outcome is a. This follows from Lemma 1 and the fact that for any order π, the subset $\{\pi_1, \ldots, \pi_{p+1}\}$ includes some majority agent.

We claim that there are at least $p+1$ votes in favor of a in the profile v in any equilibrium. By contradiction, assume that this is not the case. Then there are only two possible cases in which b is selected as an outcome with positive probability.

Case 1: v is unanimous in favor of b. Consider some agent i with type θ_a. If she deviates to $v'_i = a$, then if $\{\pi_1, \ldots, \pi_{p+1}\}$ contains only majority agents, the outcome is a lottery according to Lemma 1 and it assigns positive probability to a. Thus, such deviation is profitable for agent i. It follows that the unanimous profile v in favor of b is not an equilibrium.

Case 2: There are less than $p+1$ votes for a in the profile v. It follows that there is some majority agent i with $v_i = b$. Then, the deviation to $v_i' = a$ is profitable. Indeed, if after such deviation a is the winner of the Voting stage, a is the equilibrium outcome. If this is not the case, b is the winner of $(v' - i, v_{-i})$. Thus, if some minority agent is in $\{\pi_1, \ldots, \pi_{p+1}\}$, b is the outcome by Lemma 1. However, if only majority agents are in $\{\pi_1, \ldots, \pi_{p+1}\}$, the outcome is the lottery $\beta(v_i', v_{-i})$. Deviation by i to $v_i' = a$ increases the probability of a in such lottery: assuming that the number of votes for a in v is n_a the probabilities are the following:

$$\beta_a(v) = \frac{n_a}{n} < \frac{n_a + 1}{n} = \beta_a(v_i', v_{-i}).$$

Thus, we have eliminated all profiles v in which less than $p+1$ agents vote a as potential equilibria and this completes the proof.

B Proof of Proposition 2

Fix any strategy profile σ^1, vote profile v and order π of agents. Denote by x the majority winner of the Voting stage given v and by y the remaining option.

Consider agent π_{p+1} and assume that all Confirmation stage votes prior to her were N. Then, the best response of π_{p+1} is to vote Y if $xR_{\pi_{p+1}}y$ and N otherwise.

Consider now some agent π_t and assume that all Confirmation stage votes prior to her were N. Her best response is the following:

- Y if $xR_{\pi_t}y$ and agent i assigns positive probability to the event that all agents in $\{\pi_{t+1}, \ldots, \pi_{p+1}\}$ prefer y to x;
- N if $yR_{\pi_t}x$ and agent i assigns positive probability to the event that all agents in $\{\pi_{t+1}, \ldots, \pi_{p+1}\}$ prefer y to x;
- $\{Y, N\}$ otherwise.

Note that based on our assumption on prior beliefs, if some agent upon observing v assigns probability 0 to the event that all agents in $\{\pi_{t+1}, \ldots, \pi_{p+1}\}$ prefer y to x then all agents do so. More precisely, this happens only if some agent in $\{\pi_{t+1}, \ldots, \pi_{p+1}\}$ prefers x to y and her Voting stage strategy was revealing. It follows, that in case of non-revealing strategies for agents in $\{\pi_{t+1}, \ldots, \pi_{p+1}\}$ agent π_t strictly prefers to be truthful, i.e. to announce Y if the winner of v is her preferred option and N otherwise. Thus, we can summarize the outcome of the Confirmation stage as follows.

Lemma 2. *For any strategy profile σ^1 and Voting stage profile v, the sequential equilibrium outcome of the Confirmation stage is:*

- x if $v_i = x$ for all $i \in I$,
- x if some agent in $\{\pi_1, \ldots, \pi_{p+1}\}$ prefers the winner at v,
-the lottery $\beta(v)$ otherwise.

Consider the Voting stage of the mechanism, some agent $i \in I$ and some profile σ_{-i}^1 of strategies of other agents. We now show that i strictly prefers to vote for her most preferred option.

Assume w.l.o.g. aR_ib. Notice first that there is no v_{-i} for which i strictly prefers to vote $v_i = b$. However, voting $v_i = b$ may be a best response if i is indifferent between voting a or voting b for all possible realizations of types θ_{-i}^n and votes v_{-i} of other agents given strategies σ_{-i}^1.

Consider some realization θ_{-i}^n and v_{-i} such that the majority preferred option (among all agents including i) is not the one getting the most votes in v_{-i}. In this case, i strictly prefers to vote $v_i = a$ independently of whether a or b is the majority preferred option.

Indeed, when i's vote is the $p+1^{th}$ in favor of a, by voting a rather than b, she induces the outcome to be a rather than a lottery (if a is the majority preferred option) or a lottery with a being selected with positive probability rather than b (if b is the majority preferred option). Otherwise, the outcome is the lottery $\beta(v)$ for any vote of i, so that she strictly prefers to vote a.

If in some strategy profile σ^1_{-i}, the probability of such event is 0, that is for all possible realizations θ^n_{-i} and v_{-i} the majority preferred option gets the most votes in v_{-i}, voting b is a best response. We now show that no such strategy profile exists.

Consider some profile θ^n_{-i} such that $p-1$ agents in $I\setminus\{i\}$ are of type θ_a and the remaining $p+1$ agents are of type θ_b. This event occurs with positive probability. By assumption, given σ^1_{-i}, any realization of v_{-i} is such that the majority of agents in $I\setminus i$ vote for b (since b is the majority preferred option). Now consider a different profile $\theta^{n'}_{-i}$ such that $\theta^n_h = \theta^{n'}_h$ for all $h \in I \setminus \{i,j\}$ with agent j being θ_b in θ^n_{-i} and θ_a in $\theta^{n'}_{-i}$. That is, the profile $\theta^{n'}_{-i}$ is such that a is the majority preferred option (since aR_ib) and the only difference with θ^n_{-i} is the preference of agent j. By assumption, in the profile $\theta^{n'}_{-i}$ for any realization v_{-i} the majority of agents in $I\setminus\{i\}$ vote a. Notice, however, that each agent can condition her strategy only on her type since this is the information available to agents at the time of the vote. Thus, for all agents in $I\setminus\{i,j\}$ the probability to vote for a or b remains the same when moving from θ^n_{-i} to $\theta^{n'}_{-i}$. Thus, the only change in votes occurs for agent j. Assume that either when j is of type θ_b (as in θ_{-i}) or of type θ_a (as in θ'_{-i}) agent j randomizes, i.e. votes for a and for b with positive probabilities. In this case there is some profile v_{-i} which occurs with positive probability under θ^n_{-i} and $\theta^{n'}_{-i}$. However, this contradicts the assumption that for any realization of v_{-i} the majority preferred option obtains the majority of the votes. Thus, it must be that agent j votes a when she is of type θ_a and b when she is of type θ_b: she votes sincerely. Notice that agent j was random so that the same logic applies to any agent in $I\setminus\{i\}$. Thus, i is indifferent between voting a or b only if all other agents are truthful.

Consider now profile θ^n_{-i} such that exactly p agents are of type θ_a and p remaining agents are of type θ_b. Since they are truthful there are p votes for a and p votes for b. In this case if $v_i = a$ the outcome is a as prescribed by Lemma 2 whereas if $v_i = b$, the outcome is a lottery which assigns positive probability to b. Thus, i strictly prefers to be truthful. This completes the proof.

C Proof of Proposition 3

First, notice that in the Confirmation stage all agents in $\{\pi_1,\ldots,\pi_{p+1}\}$ are indifferent between announcing N or abstaining. Indeed, by construction the mechanism treats equally these announcements and, in the Confirmation stage, the best response does not depend on the previous announcements.

Assume then that the Voting stage admits a unique winner, i.e. the two options are not tied in v. In this case, the Confirmation stage outcome coincides with the one presented in Lemma 1.

Assume now that both a and b are tied in the profile v. Thus, independently of announcements of the Confirmation stage the outcome is a lottery which assigns probability of 1/2 to each of the options being selected. Thus, all agents at the Confirmation stage are indifferent between all 3 possible announcements. Them, the counterpart of Lemma 1 can be formulated as follows.

Lemma 3. *For any non-unanimous profile v of the Voting stage, the SPE outcome of the Confirmation stage is:*

- x if xR_iy for some $i \in \{\pi_1, \dots, \pi_{p+1}\}$,

- a lottery $\beta(v)$ if there is a unique winner in v,

- a lottery which assigns equal probabilities to both options if v does not admit a unique winner.

Consider now the Voting stage of the mechanism. We show that there is no equilibrium which selects b (the minority preferred option) with positive probability. By contradiction, assume that such equilibrium exists.

Case 1: The outcome is deterministic and selects b with probability 1 for all orders π. In this case, given Lemma 3, one of the following statements holds:

- all participating majority agents vote b. If any of these agents deviates to $v_i' = a$, this is a profitable deviation since there is positive probability that only majority agents play at the Confirmation stage and, by Lemma 3, the outcome in this case is a lottery;

- no majority agent participates. Then for any majority agent i with $v_i = \emptyset$, deviating to $v_i' = a$ is profitable since it leads to a lottery as an outcome with positive probability.

Case 2: The outcome is b with positive probability. Notice, that if a is the winner of the Voting stage, a is the outcome for all possible orders π since some majority agent is among the first $p+1$ agents at the Confirmation stage. Thus, if b is selected with positive probability, it must be the winner of v, or that $v_j = \emptyset$ for all $j \in I$. If b is the majority winner of v then there is some majority agent who either abstains or votes for b in the Voting stage.

- Assume $v_j = \emptyset$ for all $j \in I$. In this case any agent has incentives to enter and vote for her favorite option since this option will be the outcome with only one agent present at the Voting stage.

- Assume b is the majority winner and there is some majority agent $i \in I$ with $v_i = b$. Then the deviation to $v_i' = a$ is profitable. Indeed, if after this deviation a is the majority winner of v, a is the outcome of the mechanism with probability 1. Otherwise, the outcome is the lottery based on the Voting stage profile for any π. Deviation by v to $v_i' = a$ increases the probability of a in such lottery.

- Assume b is the majority winner and there is some majority agent $i \in I$ with $v_i = \emptyset$. Then the deviation to $v_i' = a$ is profitable. Indeed, if after such deviation a is the winner of the Voting stage, a is the equilibrium outcome. If this is not the case, b is the winner of $(v' - i, v_{-i})$. Thus, if some minority agent is in $\{\pi_1, \dots, \pi_{p+1}\}$, b is the outcome by Lemma 3. However, if only majority agents are in $\{\pi_1, \dots, \pi_{p+1}\}$, the outcome is the lottery $\beta(v_i', v_{-i})$. Deviation by v to $v_i' = a$ increases the probability of a in such lottery: assuming that the number of votes for a in v is n_a and the total number of the Voting stage participants is n the probabilities are the following:

$$\beta_a(v) = \frac{n_a}{n} < \frac{n_a + 1}{n + 1} = \beta_a(v_i', v_{-i}).$$

This concludes the proof.

D Proof of Proposition 4

W.l.o.g. assume that any agent i in $\{1, \dots, p+1\}$ is such that aR_ib so that a is the majority-preferred option and b the minority-preferred one. Any agent i with aR_ib is a majority agent. We need to prove that (A.) there is an equilibrium implementing a and that (B.) any equilibrium selects a.

A. Existence of an equilibrium selecting a.

Consider the set $J = \{1, \dots, p+1\}$ that consists only of majority agents. Take the strategy profile m where for each $i \in J$, $v_i = a$ and $c_i \subset J \setminus \{i\}$ so that coalition J forms a bloc in favor of a. It follows that $\chi_{BF}(m) = a$. To see why m is an equilibrium, remark that each agent in J prefers a to b (and a to any lottery with both a and b in its support by **SD**) and hence does not want to deviate. Each agent outside J cannot affect the outcome since the bloc formed by J is formed independently of the deviation of any agent outside J. This shows the existence of an equilibrium selecting a.

B. Any equilibrium implements a.

For the sake of clarity, we divide this part of the proof in two sections. In section B.1, we show that there is no bloc in favor of b in equilibrium. In section B.2, we show that any strategy profile that leads to a full-support lottery cannot be an equilibrium, concluding the proof.

B.1. No bloc in favor of b in equilibrium.

Take any profile m with a bloc B in favor of b; hence $\chi_{BF}(m) = b$. The definition of a bloc means that at least $p+1$ agents vote for b and nominate only agents in B. Consider the effective bloc B^* which exists and is unique according to Proposition 7. Since a is the majority option, there is some agent $i \in B^*$ with $v_i = b$ in the profile m and aR_ib.

Assume that m is an equilibrium. Suppose that agent i deviates from $m_i = (b, c_i)$ to $m'_i = (a, c_i)$. This means that B^* is not anymore an effective bloc in favor of b in the profile (m'_i, m_{-i}). Moreover, since $B^* = \cap_{B \in B^m} B$, there is no other remaining bloc in the profile (m'_i, m_{-i}) as shown by Proposition 7; thus the outcome $\chi_{BF}(m'_i, m_{-i})$ is a lottery with support a and b with a being selected with positive probability since $\eta_i(m) > 0$ and thus $\eta_i(m'_i, m_{-i}) > 0$ (i was nominated by some other agent in m, being part of B^*). Thus, by **SD**, m'_i is a profitable deviation for i since it increases the probability of a being selected, proving that m is not an equilibrium.

B.2. There is no equilibrium which selects b with positive probability.

Assume that there is some equilibrium m where the outcome is a full-support lottery.

Notice that the following two statements hold for any equilibrium profile m with the outcome being a lottery:

(1) any agent i who is nominated ($\eta_i(m) > 0$) is sincere.

(2) any agent nominates the largest number of agents who announce her preferred option. In other words, if aR_ib then $|\{j \mid v_j = a \text{ and } j \in c_i\}| = \min\{p, \{h \in I \mid v_h = a\}\}$.

Indeed, (1) holds since with $\eta_i(m) > 0$ the vote of agent i affects the final outcome, thus, voting sincerely increases the probability of i's favorite option being selected. Statement (2) holds since the weight $\eta^x(\cdot)$ is increasing in the sum of the weights of x-agents and each agent's weight strictly increases on the number of votes that she receives.

Given that (1) and (2) hold since m is an equilibrium and that $B^m = \emptyset$, there is some majority agent which votes b and is not nominated. Indeed, assume this is not the case and such agent does not exist. According to (1) all nominated agents vote sincerely. It follows from (2) then that all majority agents nominate only other majority agents who are also sincere. This means that a bloc in favor of a exists contradicting $B^m = \emptyset$. Consider then some minority agent j, i.e. bR_ja. Since (1) holds, c_j does not include any majority agent who votes b, that is $|\{h \in c_j \mid v_h = b\}| < p$. Then since $\eta^b(m)$ is increasing in the number of nominations of b-agents, agent j has a profitable deviation: to nominate agent i in c_j rather than some a-agent. Formally, $m'_j = (b, c'_j)$ with $c'_j = (c_j \setminus \{h\}) \cup \{i\}$ for some h with $v_h = a$. This contradicts m being an equilibrium, and concludes the proof.

E Proof of Proposition 5

Lemma 4 shows that the unique SPE outcome of the Verification stage is the outcome of truth-telling. We now consider the earlier stages of the mechanism to complete the proof of the proposition.

A. Ratification stage

At this stage, an agent's decision is rather simple. She strictly prefers verifying if she anticipates that the outcome is better than at the Voting stage. Note that if the outcome after verification remains unchanged, the agent is indifferent at the Ratification stage.

B. Voting stage

We now consider the first stage of the mechanism, letting a denote the majority preferred option. There are 2 possible cases: either less than $p+1$ majority agents vote for a in the first stage (B.1) or and at least $p+1$ majority agents do so (B.2).

Case B.1: If the profile v^1 is unanimous in favor of b, the outcome is b. If the profile v^1 is not unanimous in favor of b, as established in A., the second-stage SPE outcome is equivalent to the one under truth-telling. Thus the SPE outcome is a lottery with all majority agents announcing a and all minority agents announcing b. Consider now a deviation $\hat{v}_i^1 = a$ by a majority agent who voted b in the first stage before ($v_i^1 = b$). Given the SPE of the second stage, after such deviation there is either a bloc in favor of a if this agent is the $p+1^{th}$ majority agent to vote a in the first stage, or a lottery where the probability of a strictly increases according to equation (1). Thus, this deviation is strictly profitable. Therefore, the first-stage profile in which less than $p+1$ majority agents vote a can not be part of a SPE equilibrium.

Case B.2: There are at least $p+1$ majority agents who vote for a in the first stage. Assume first that the profile v^1 is unanimous in favor of a. If $v_i^1 = a$ for all $i \in I$, it is an equilibrium. Indeed, by deviating to $\hat{v}_j^1 = b$ any minority agent j is not changing the outcome given that the outcome of the verification stage is equivalent to the one under truth-telling and the remaining agents $(n-1)$ vote for a in the first stage. If the profile v^1 is not unanimous, the second-stage SPE outcome is a with a bloc of majority agents being formed in favor of a. Each of the majority agents has no profitable deviation since a is their preferred option. No minority agent's deviation can prevent the bloc in favor of a. Thus, the described profile is an equilibrium and a is the only equilibrium outcome.

F Proof of Proposition 6

We now show that every agent strictly prefers to be sincere in the first stage.

W.l.o.g. take agent 1 with type θ_a and let σ_{-1} denote the strategy of her opponents. Recall that agent 1 weakly prefers to be sincere in the first stage. It suffices to prove that there is some event where agent 1 strictly prefers to be sincere and that this event arises with positive probability. If this is the case, then **SD** implies that agent 1 strictly prefers to be sincere in the first stage.

Let T_a be the event where p agents have type θ_a and p agents have type θ_b, all different from agent 1. Agent 1 assigns to the event T_a probability $\pi_a > 0$ where $\pi_a = C_p^{n-1}(q_a)^p(1-q_a)^p$.

Moreover, if the event T_a arises, no bloc can be formed without agent 1 being part of it. Indeed, as shown by Lemma 5, the outcome of the verification stage is equivalent to the one under truth-telling. Moreover, by definition, the formation of a bloc requires at least $p+1$ agents voting equally in the voting and in the verification stage. However, exactly p of them have type θ_a and exactly p of them have type θ_b. W.l.o.g we can assume that the former agents vote a in the verification stage while the latter ones vote b due to Lemma 5.

Therefore, we have two cases to describe the vote of agent 1: either agent 1 is part of a bloc in favor of a or no bloc is formed.

In the first case, agent 1 strictly prefers to report her type ensuring that her best option is implemented. In the second case, no bloc is formed so that the outcome is a lottery: again, she strictly prefers to report her type since the weight of a is strictly increasing in the number of votes obtained in the first stage.

G Voting profile as a directed graph

It is useful to consider blocs in terms of the graph theory. Notice that for any message profile $m = (v, c)$ the nomination profile c creates a directed graph in which the vertices are the agents and the edges their nominations. Formally, for each message profile $m = (v, c)$ denote by $G_m = (I, C)$ the directed graph formed by c where the set I of agents coincides with the set of vertices, and C is the adjacency matrix such that $C_{ij} = 1$ if $j \in c_i$ and $C_{ij} = 0$ otherwise.

We can formulate an option definition of bloc using the adjacency matrix.

Definition 2. *A set B of agents forms a **bloc** in favor of option x in profile m if:*
1. $|B| \geq p + 1$
2. $v_i = x$ $\forall i \in B$ and
3. the restriction of C to set B of agents, denoted C_B, is such that $\sum_{h=1}^{|B|} C_{ih} = p$ $\forall i \in B$.

It follows from Definition 2 that if a set B of agents forms a bloc in favor of x in m, then there is no path from any agent $i \in B$ to any agent $j \in I \setminus B$ in the associated graph G_m. It is straightforward that Definition 2 is equivalent to Definition 1 in which all agents in B vote for the same option and nominate other agents in J exclusively. Firstly, both definitions require at least $p + 1$ agents to vote for the same option. To show that $\sum_{h=1}^{|B|} C_{ih} = p$ $\forall i \in B$ is equivalent to agents in B voting for other agents in B exclusively, observe that the row i of the adjacency matrix C gives the nominations of agent i. Thus, $\sum_{h=1}^{n} C_{ih} = p$ by definition of the mechanism. Since, according to Definition 2, $\sum_{h=1}^{|B|} C_{ih} = p$ $\forall i \in B$, $C_{ih} = 0$ for all $h \in I \setminus B$, that is no agent $i \in B$ is voting for any agent outside B which proves the equivalence.

We need some additional definitions to formulate the main results regarding blocs.

We say that G_J is a subgraph of a graph G induced by the set $J \subseteq I$ of vertices if it includes all vertices in J and its adjacency matrix C_J is the restriction of C to J (i.e includes only rows and columns corresponding to vertices in J).

Definition 3. *A subgraph G_J for some $J \subseteq I$ of a graph G is strongly connected if there exists a path in each direction between any pair i, j of vertices with $i, j \in J$.*

Definition 4. *A bloc $B \subseteq I$ in favor of x is effective iff it is strongly connected.*

Figure 1 gives two examples of the directed graphs created by the Bloc formation mechanism. The vertices represent the agents, the letters their votes and the arrows show nominations. Figure 1a presents a profile with no blocs. Indeed, the only possible bloc could involve agents $\{1, 2, 3\}$ since all vote for a and while the rest vote for b. However, agent 1 (resp. agent 3) nominates agent 5 (resp. agent 4), thus, violating the conditions to form a bloc. In Figure 1b, the profile admits two blocs: $\{1, 4, 5\}$ and $\{1, 2, 3, 4, 5\}$. Indeed, in any of these subsets all agents vote for b and nominate only agents in the bloc. Moreover, only the bloc $\{1, 4, 5\}$ is strongly connected, thus, effective. Notice that there is no path from agent 4 to agent 2, which prevents $\{1, 2, 3, 4, 5\}$ to be an effective bloc.

The next proposition defines some important properties of an effective bloc.

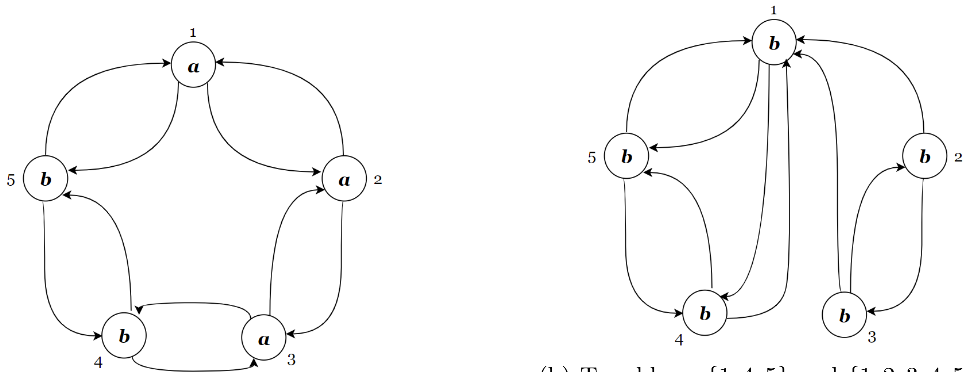

(a) No blocs.

(b) Two blocs, $\{1, 4, 5\}$ and $\{1, 2, 3, 4, 5\}$, in favor of b.

Figure 1: Voting profiles formed by the Bloc formation mechanism.

Proposition 7. *Any profile m of the Bloc formation mechanism admitting a bloc also admits an effective bloc B^*. Moreover, the effective bloc B^* is unique and satisfies $B^* = \cap_{B \in B^m} B$.*

Proof. **Existence.** Let m be some profile with $B^m \neq \emptyset$ and consider w.l.o.g. that all blocs are in favor of x. Assume by contradiction that there is no effective bloc B^* in m. Therefore, any bloc $B \in B^m$ is not effective, thus, not strongly connected. It follows that there are 2 vertices $i, j \in B$ such that there is no path in G_m either from i to j or from j to i, or both. W.l.o.g. assume that there is no path from i to j. It follows that there is a group $B' \subseteq B$ of agents with $i \in B'$ such that there is no path from any agent in B' to agent j. Since B is a bloc, it follows that $\forall h \in B'$, $c_h \subset B'$. Thus, B' is a bloc. It follows that if B is a bloc which is not strongly connected, it contains another bloc of smaller size. Thus, since the minimal size of a bloc is $p + 1$, for each bloc B which fails to be strongly connected, there is a bloc contained in B which is strongly connected.

Uniqueness. Assume by contradiction that for some profile m there are two non-identical effective blocs B^* and $B^{'*}$. Since each bloc consists of at least $p + 1$ agents, $B^* \cap B^{'*} \neq \emptyset$. Thus, there is some agent i such that $i \in B^* \cap B^{'*}$. By definition of a bloc, there is no path from i to any $j \in B^* \setminus (B^* \cap B^{'*})$ since $i \in B^{'*}$. Likewise, there is no path from i to any $h \in B^{'*} \setminus (B^* \cap B^{'*})$. By assumption, blocs B^* and $B^{'*}$ are effective and, thus, strongly connected. It follows that there is a path between any two vertices of an effective bloc, reaching the desired contradiction.

$\boldsymbol{B^* = \cap_{B \in B^m} B}$**.** We have shown that each bloc which is not effective must include an effective bloc. We have also shown that the effective bloc is unique. The claim follows directly from the two observations. □

The previous proposition shows that in any profile m with blocs, the intersection of the blocs is non-empty and is a bloc itself. Moreover, this intersection is strongly connected meaning that there is no agent which can be removed from it in such way that the profile still admits a bloc. This property has an important implication on the strategic behavior, as summarized by the next result: for any profile m admitting a bloc, any agent in the effective bloc has a strategy m'_i that allows her to break all blocs in m (i.e. no bloc in (m'_i, m_{-i})).

Proposition 8. *For any profile m admitting a bloc, any agent i in the effective bloc B^* has a strategy m'_i such that $B^{(m'_i, m_{-i})} = \emptyset$.*

Proof. Take some m with $B^m \neq \emptyset$. W.l.o.g. assume that all blocs in B^m are in favor of x and consider some agent i in the effective bloc B^*. Notice that since $|B^*| \geq p+1$ and B^* is strictly connected, for any $B \in B^m$ and any $j \in B$ there is a path from j to i. Consider a deviation $m'_i = (y, c_i)$, that is let agent i to switch and vote for y instead of x while keeping her nominations unchanged. After such deviation, i cannot be a part of a bloc in favor of x since she votes for y.

By definition of a bloc, for any bloc B there is no path from the members of the bloc to the agents in $I \setminus B$. However, as stated before, since $i \in B^*$, there is a path to i from any member of any bloc in B^m. Thus, there is no bloc in favor of x in profile (m'_i, m_{-i}).

Notice also that since $B^m \neq \emptyset$ and since c_i includes only agents voting for x, there can be no bloc in favor of y in (m'_i, m_{-i}). Thus, $B^{(m'_i, m_{-i})} = \emptyset$, ending the proof. □

H Verification mechanism

Lemma 4. *For any profile $v^1 \in A^n$ of the Verification mechanism, the subgame-perfect equilibrium outcome of the verification stage is unique and coincides with the one in which agents report their true type.*

Proof. We prove this result by induction. W.l.o.g. we assume that the order of play is 1,2,..,n. Given this order of play, we denote by $\Gamma(h_{i-1})$ the subgame starting from agent i given history h_{i-1}. For any history h_i, we denote by $B(h_i)$ the set of blocs formed, i.e. $B(h_i) = \{J \subset \{1, ..., i\} \mid |J| \geq p + 1 \text{ and } v^1_j = v^2_j \; \forall j \in J\}$. We denote by $\tilde{v}^2_k$ the profile of votes $(v^2_k, \ldots, v^2_n)$ in which each of the agents from k to n vote truthfully for her preferred option. Likewise, we denote by $B(h_i \cup \{\tilde{v}^2_k\})$ the set of blocs formed by agents $1, \ldots, i$ and agents $k, \ldots, n$ for some $k > i$ given that each agent starting from k votes truthfully for her preferred option. $BR_i(h_{i-1})$ is the best response of agent i at verification stage after history h_{i-1} was realized. Notice that we abuse slightly the notation and do not include the preference of agent i in the notation.

Step 1: Agent n.

We consider the decision of agent n in the subgame $\Gamma(h_{n-1})$. Let x be agent n's preferred option.

If h_{n-1} is such that $B(h_{n-1}) \neq \emptyset$, agent n cannot affect the outcome, and any v^2_n is a best response, that is $BR_n(h_{n-1}) = A$.

If h_{n-1} is such that $B(h_{n-1}) = \emptyset$, that is no bloc was formed by the first $n - 1$ agents, truth-telling is the unique best response of agent n, i.e. $BR_n(h_{n-1}) = x$.

Notice that any SPE outcome of the subgame $\Gamma(h_{n-1})$ is equivalent to the outcome if she votes truthfully.

Step 2: Agent $n - 1$.

We now consider the decision of agent $n - 1$ in any subgame $\Gamma(h_{n-2})$.

If h_{n-2} is such that $B(h_{n-2}) \neq \emptyset$, then a bloc is formed by the first $n - 2$ agents. Thus, agent $n - 1$ cannot affect the outcome so her best response equals $BR_{n-1}(h_{n-2}) = A$.

If $B(h_{n-2}) = \emptyset$, by Step 1, the equilibrium outcome of any subgame $\Gamma(h_{n-1})$ coincides with the one obtained if agent n votes truthfully. Thus, it is without loss of generality to assume that agent n votes for her preferred option. There are two possible cases for agent $n - 1$:

If $B(h_{n-2} \cup \tilde{v}^2_n) \neq \emptyset$, then agent $n - 1$ is indifferent, since independently of her choice agent n will complete the bloc. Thus $BR_{n-1}(h_{n-2}) = A$.

If $B(h_{n-2} \cup \tilde{v}^2_n) = \emptyset$, then no bloc is going to be formed without the vote of agent $n - 1$. Then, agent $n - 1$ is strictly better off voting truthfully, so $BR_{n-1}(h_{n-2}) = y$ with y being her preferred option.

Notice that any SPE outcome of the subgame $\Gamma(h_{n-2})$ is again equivalent to the outcome if agents $n - 1$ and n vote truthfully.

We have shown that truth-telling is a SPE for all possible subgames $\Gamma(h_{n-2}), \Gamma(h_{n-1})$ and preferences (R_{n-1}, R_n) and any SPE of these subgames is outcome-equivalent to the truth-telling SPE. We now establish the induction argument.

Step 3: Induction argument.

Take some agent i. Assume that any SPE outcome of any subgame $\Gamma(h_i)$ is equivalent to the outcome under truth-telling.

If $B(h_{i-1} \cup \tilde{v}^2_{i+1}) \neq \emptyset$, then the bloc in favor of one of the options will be formed independently of the vote of agent i. Thus, agent i is indifferent, so $BR_i(h_{i-1}) = A$.

If $B(h_{i-1} \cup \tilde{v}^2_{i+1}) = \emptyset$, then no bloc can be formed without agent i. In this case the outcome is either a lottery, or agent i completes some bloc. In this case the best response of agent i is strict and it is to vote truthfully. Thus, we obtain that in the verification-stage, truth-telling is a SPE and, moreover, the outcomes of any other SPEs coincide with the truth-telling one. □

Lemma 5. *Under incomplete information, for any profile v^1, the unique sequential equilibrium outcome of the verification stage is the one associated to sincere behavior.*

Proof. For each agent i, we denote the profile of sincere announcements following agent i by $\tilde{v}^2_{i+1} = (v^2_{i+1}, v^2_{i+2} \dots, v^2_n)$ with $v^2_j = x$ iff xR_jy for all $j \geq i+1$.

Assume first that the strategy v^1 of the Voting stage is revealing for all agents, that is the preference profile is revealed before the Stage 2. In this case, the logic of the complete information case applies and the sincere voting outcome is the unique sequential equilibrium outcome of the verification stage.

Assume now that there is some agent j for whom v^1_j is not revealing, so that $(\sigma^1_j)^{-1}(v^1_j) \in \text{int}(\Delta(\Theta))$. It follows that the rest of agents upon observing v^1_j are uncertain about j's true preference.

The rest of the proof proceeds by induction.

Step 1. Consider agent n, i.e. the last agent to vote at the Verification stage.

- If $B(h_{n-1}) \neq \emptyset$ then agent n is indifferent, so $BR_n(h_{n-1}) = A$;
- If $B(h_{n-1}) = \emptyset$ then agent n strictly prefers to vote sincerely.

Thus, the outcome given any history h_{n-1} is the same as in the case when agent n votes sincerely in the second stage.

Step 2. Consider now agent $n-1$.

- If $B(h_{n-2}) \neq \emptyset$ then agent $n-1$ is indifferent, so $BR_{n-1}(h_{n-2}) = \Delta$;
- If $B(h_{n-2}) = \emptyset$ then there are several cases:

Step 2.1. v^1_n is revealing. Thus, preferences of agent n are known to agent $n-1$. Thus, agent $n-1$'s reasoning is identical to the complete information case.

Step 2.2. v^1_n is not revealing, so agent $n-1$ is not certain of the preferences of agent n. By the definition of a bloc, we cannot have that exactly one agent is needed to complete a bloc in favor of a and exactly one agent can complete a bloc in favor of b for any history h_{n-2} (since $n = 2p+1$ and each bloc requires at least $p+1$ agents). Thus, there is a positive probability assigned to the preference of agent n being such that $B(h_{n-2} \bigcup \tilde{v}^2_n) = \emptyset$. In this case agent $n-1$ is strictly better-off by voting truthfully since the outcome is a lottery. Likewise, agent $n-1$ also assigns positive probability to the event $B(h_{n-2} \cup \tilde{v}^2_n) \neq \emptyset$ and, in this case, a bloc is formed independently of the vote of agent $n-1$. Thus, she is indifferent and thus $BR_{n-1}(h_{n-2}) = A$. Thus, by **SD**, agent $n-1$ strictly prefers to vote for her favorite option.

All in all, the outcome given any history h_{n-2} and any message v^1_n coincides with the outcome in the case when agent $n-1$ votes truthfully.

Step 3. Assume now that given any history h_i the sequential equilibrium outcome in the subgame $\Gamma(h_i)$ is equivalent to the one under sincere behavior. Consider some history h_{i-1} and agent i who best responds to such history.

- If $B(h_{i-1}) \neq \emptyset$ then agent i is indifferent, so $BR_i(h_{i-2}) = \Delta$;

- If $B(h_{i-1}) = \emptyset$ then denote by J the set of agents in $\{i+1, \ldots, n\}$ whose first-stage messages are revealing. That is, for each $j \in J$, $(\sigma_j^1)^{-1}(v_j^1) \notin \text{int}(\Delta(\Theta))$.

Step 3.1. If $B(h_{i-1} \bigcup(\bigcup_{j\in J} \tilde{v}_j^1)) \neq \emptyset$, then agent i is certain that a bloc is formed independently of her vote, thus, she is indifferent and $BR_i(h_{i-1}) = A$.

Step 3.2. If $B(h_{i-1} \bigcup(\bigcup_{j\in J} \tilde{v}_j^1)) = \emptyset$, then $Pr(B(h_{i-1} \bigcup(\bigcup_{k\in\{i+1,\ldots,I\}} (\tilde{v}_k^1)) = \emptyset) > 0$, that is agent i assigns strictly positive probability to her verification stage vote being pivotal. With the remaining probability agent i is not decisive, so the outcome is independent of her vote. Thus, agent i strictly prefers to vote truthfully due to **SD**.

Thus, given that the outcome of a subgame starting from agent $i+1$ is as if all the agents vote truthfully, the outcome of the subgame starting from agent i is the same as in the case of truthful voting as well, which completes the induction argument. □