New Perspectives on Social Choice

Eric Pacuit, University of Maryland

Lecture 9: Perpetual Voting and Storable Votes

PHIL 808K

60	40		65	35		45	55		30	70	
a_1	b_1		a_2	b_2		a_3	b_3		a_4	b_4	
b_1	a_1		b_2	a_2		b_3	a_3		b_4	a_4	$\cdots$

Winner: a_1 Winner: a_2 Winner: b_3 Winner: b_4

E.g., a group of friends are planning a party and meet and different times to decide when to have the party, what catering to use, how much money to spend on the catering, what drinks to get, etc.

60	40
a_1	b_1
b_1	a_1

Winner: a_1

65	35
a_2	b_2
b_2	a_2

Winner: a_2

45	55
a_3	b_3
b_3	a_3

Winner: b_3

30	70
a_4	b_4
b_4	a_4

Winner: b_4

$\cdots$

E.g., a group of friends are planning a party and meet and different times to decide when to have the party, what catering to use, how much money to spend on the catering, what drinks to get, etc.

- A sequence of different decision to be made by a fixed set of voters for which inclusion and participation is desirable.
- Can we guarantee a fair share of the decision power for each voter?
- Not a good model for large-scale elections and/or when there are extreme or dangerous opinions from the voters.

Fix a set of voters $N = \{1, \ldots, n\}$

A decision instance is a tuple (N, A, C) where

- C is a set of alternatives
- $A = (A(1), \ldots, A(n))$ is an approval profile, where for all $v \in N$, $A(v) \subseteq C$

Fix a set of voters $N = \{1, \ldots, n\}$

A decision instance is a tuple (N, A, C) where

- C is a set of alternatives
- $A = (A(1), \ldots, A(n))$ is an approval profile, where for all $v \in N$, $A(v) \subseteq C$

k-decision sequence: $(N, \overline{A}, \overline{C})$, where $\overline{A} = (A_1, \ldots, A_k)$ is a sequence of approval profiles and an associated k-tuple of alternatives $\overline{C} = (C_1, \ldots, C_k)$.

Fix a set of voters $N = \{1, \ldots, n\}$

A decision instance is a tuple (N, A, C) where

- C is a set of alternatives
- $A = (A(1), \ldots, A(n))$ is an approval profile, where for all $v \in N$, $A(v) \subseteq C$

k-decision sequence: $(N, \overline{A}, \overline{C})$, where $\overline{A} = (A_1, \ldots, A_k)$ is a sequence of approval profiles and an associated k-tuple of alternatives $\overline{C} = (C_1, \ldots, C_k)$.

k-decision history: $(N, \overline{A}, \overline{C}, \overline{w})$, where $(N, \overline{A}, \overline{C})$ is a k-decision sequence and for each $i \in \{1, \ldots, k\}$, $w_i \in C_i$ is the alternative chosen in round i.

Perpetual AV:

$$\alpha_{k+1}(v) = \alpha_k(v) = 1$$

Perpetual Unit Cost:

$$\alpha_{k+1}(v) = \begin{cases} \alpha_k(v) + 1 & \text{if } w_k \notin A_k(v) \\ \alpha_k(v) & \text{if } w_k \in A_k(v) \end{cases}$$

Perpetual Reset:

$$\alpha_{k+1}(v) = \begin{cases} \alpha_k(v) + 1 & \text{if } w_k \notin A_k(v) \\ 1 & \text{if } w_k \in A_k(v) \end{cases}$$

Example: Perpetual AV

Suppose that $N = \{1, 2, 3, 4\}$, the same four alternatives $\{a, b, c, d\}$ are used in each round, and alphabetic tie-breaking is used.

	1	2	3	4	winner
weights	1	1	1	1	
A_1	$\{a\}$	$\{a\}$	$\{b\}$	$\{c, d\}$	a
weights	1	1	1	1	
A_2	$\{a\}$	$\{a, b, c\}$	$\{d\}$	$\{c\}$	a
weights	1	1	1	1	
A_3	$\{a\}$	$\{b\}$	$\{a, c\}$	$\{b\}$	a
weights	1	1	1	1	
A_3	$\cdots$				

Example: Perpetual Unit Cost

Suppose that $N = \{1, 2, 3, 4\}$, the same four alternatives $\{a, b, c, d\}$ are used in each round, and alphabetic tie-breaking is used.

	1	2	3	4	winner
weights	1	1	1	1	
A_1	$\{a\}$	$\{a\}$	$\{b\}$	$\{c, d\}$	a
weights	1	1	2	2	
A_2	$\{a\}$	$\{a, b, c\}$	$\{d\}$	$\{c\}$	c
weights	2	1	3	2	
A_3	$\{a\}$	$\{b\}$	$\{a, c\}$	$\{b\}$	a
weights	2	2	3	3	
A_3	$\cdots$				

Example: Perpetual Reset

Suppose that $N = \{1, 2, 3, 4\}$, the same four alternatives $\{a, b, c, d\}$ are used in each round, and alphabetic tie-breaking is used.

	1	2	3	4	winner
weights	1	1	1	1	
A_1	$\{a\}$	$\{a\}$	$\{b\}$	$\{c, d\}$	a
weights	1	1	2	2	
A_2	$\{a\}$	$\{a, b, c\}$	$\{d\}$	$\{c\}$	c
weights	2	1	3	1	
A_3	$\{a\}$	$\{b\}$	$\{a, c\}$	$\{b\}$	a
weights	1	2	1	2	
A_3	$\cdots$				

Perpetual PAV:

$$\alpha_{k+1}(v) = \frac{1}{sat_k(v, \overline{w}) + 1} = \begin{cases} \alpha_k(v) & \text{if } w_k \notin A_k(v) \\ \frac{\alpha_k(v)}{\alpha_k(v)+1} & \text{if } w_k \in A_k(v) \end{cases}$$

Example: Perpetual PAV

Suppose that $N = \{1, 2, 3, 4\}$, the same four alternatives $\{a, b, c, d\}$ are used in each round, and alphabetic tie-breaking is used.

	1	2	3	4	winner
weights	1	1	1	1	
A_1	$\{a\}$	$\{a\}$	$\{b\}$	$\{c, d\}$	a
weights	$\frac{1}{2}$	$\frac{1}{2}$	1	1	
A_2	$\{a\}$	$\{a, b, c\}$	$\{d\}$	$\{c\}$	c
weights	$\frac{1}{2}$	$\frac{1}{3}$	1	$\frac{1}{2}$	
A_3	$\{a\}$	$\{b\}$	$\{a, c\}$	$\{b\}$	a
weights	$\frac{1}{3}$	$\frac{1}{3}$	$\frac{1}{2}$	$\frac{1}{2}$	
A_3	$\cdots$				

Perpetual Consensus:

- The weights of satisfied voters is reduced by a total of $n = |N|$ split equally among them.
- Some voters may have negative weights. Voters with negative weights are not taken into account when determining the winning alternative.
- After each decision, the weight of all voters is increased by 1.

$$N_k^+(c) = \{v \in N \mid c \in A_k(v) \text{ and } \alpha_k(v) > 0\}$$

$$\alpha_{k+1}(v) = \alpha_k(v) + 1 = \begin{cases} \alpha_k(v) & \text{if } v \notin N_k^+(w_k) \\ \alpha_k(v) + 1 - \frac{n}{|N_k^+(w_k)|} & \text{if } v \in N_k^+(w_k) \end{cases}$$

Example: Perpetual Consensus

Suppose that $N = \{1, 2, 3, 4\}$, the same four alternatives $\{a, b, c, d\}$ are used in each round, and alphabetic tie-breaking is used.

	1	2	3	4	winner
weights	1	1	1	1	
A_1	$\{a\}$	$\{a\}$	$\{b\}$	$\{c, d\}$	a
weights	0	0	2	2	
A_2	$\{a\}$	$\{a, b, c\}$	$\{d\}$	$\{c\}$	c
weights	1	1	3	-1	
A_3	$\{a\}$	$\{b\}$	$\{a, c\}$	$\{b\}$	a
weights	0	2	2	0	
A_3	$\cdots$				

Perpetual Nash: Maximize Nash welfare, i.e., the product of voters' utilities, where the voters' utility is their satisfaction if their satisfaction is larger than 0; voters with a satisfaction of 0 have a utility of some small constant, e.g., $\epsilon = n^{-n}$

$$u_{k+1}(v, c) = \begin{cases} \max(sat_k(v, \overline{w}), \epsilon) & \text{if } c \notin A_{k+1}(v) \\ sat_k(v, \overline{w}) + 1 & \text{if } c \in A_{k+1}(v) \end{cases}$$

The Nash score of c is $nash_{k+1}(c) = \Pi_{v \in N} u_{k+1}(v, c)$. The alternative with the maximum Nash score is chosen.

Support, Quota

The support of a voter $v \in N$ in round j is defined as

$$supp_j(v) = \frac{1}{n} \max_{c \in A_j(v)} |\{u \in V \mid A_j(u) = c\}|$$

The quota of voter $v \in N$ in round j is

$$qu_j(v) = \sum_{i \leq j} supp_i(v)$$

The support of a voter in round j is the proportion of voters that can collectively agree on some alternative that v approves. The quota of voter v in round j is v's cumulative support from round 1 to j.

Simple Proportionality

We say that a k-decision sequence $(N, \overline{A}, \overline{C})$ is simple if $A_1 = \cdots = A_k$, $C_1 = \cdots = C_k$ and $|A_1(v)| = 1$ for all $v \in N$.

For any simple n-decision sequence $(N, \overline{A}, \overline{C})$ with $|N| = n$, we say that $\overline{w} \in \overline{C}$ is proportional if $sat_n(v, \overline{w}) = qu_n(v)$ for every $v \in N$.

A perpetual voting rule satisfies simple proportionality if for any simple $|N|$-decision sequence $(N, \overline{A}, \overline{C})$, the rule is proportional.

Proposition 1. AV, Perpetual Equality, Perpetual Reset, and Perpetual Unit-Cost fail simple proportionality.

Theorem 1. Perpetual PAV, Perpetual Consensus, Perpetual Nash, and Perpetual Quota satisfy simple proportionality.

Independent of Uncontroversial Decisions

An approval profile A is **uncontroversial due to** c if $\bigcap_{v \in N} A(v) = \{c\}$.

Given a k-tuple $L = (l_1, \ldots, l_k)$ and $i \in \{0, \ldots, k\}$, let $L \oplus_i x$ be the $(k+1)$-tuple $(l_1, \ldots, l_i, x, l_{i+1}, \ldots, l_k)$.

A perpetual voting rule $\mathcal{R}$ is *independent of uncontroversial decisions* if for any k-decision sequence $(N, \overline{A}, \overline{C})$, approval profile A for C that is uncontroversial due to c, and $i \in \{0, \ldots, k\}$, it holds that

$$\mathcal{R}(N, \overline{A} \oplus_i A, \overline{C} \oplus_i C) = \mathcal{R}(N, \overline{A}, \overline{C}) \oplus_i c$$

Proposition 2. Perpetual PAV, Perpetual Nash, and Perpetual Reset fail independence of uncontroversial decisions.

Theorem 2. AV, Perpetual Equality, Perpetual Quota, Perpetual Unit-Cost, and Perpetual Consensus satisfy independence of uncontroversial decisions

Dry Spells

Given a k-decision history $(N, \overline{A}, \overline{C}, \overline{w})$, we say that a voter $v \in N$ has a **dry spell of length** ℓ if there exists $t \leq k - \ell$ such that $sat_t(v, \overline{w}) = sat_{t+\ell}(v, \overline{w})$, i.e., voter v is not satisfied by any choice in rounds $t + 1, \ldots, t + \ell$.

Let $d : \mathbb{N} \to \mathbb{N}$. A perpetual voting rule has a **dry spell guarantee of** d if, for any decision sequence, no voter has a dry spell of length $d(|N|)$.

A perpetual voting rule $\mathcal{R}$ has **bounded dry spells** if $\mathcal{R}$ has a dry spell guarantee of some d.

Proposition 3. AV, Perpetual PAV, Perpetual Equality, Perpetual Quota, Perpetual Nash, and Perpetual Unit-Cost have unbounded dry spells.

Theorem 4. Perpetual Consensus has a dry spell guarantee of at least $\frac{n^2+3n}{4}$.

Perpetual Lower Quota

A measure to which degree voters receive at least a “fair share” of favorable choices, and thus have a fair say in the decision process.

Let $(N, \overline{A}, \overline{C})$ be a k-decision sequence. A k-choice sequence $\overline{w} \in \overline{C}$ satisfies **perpetual lower quota** if for every voter $v \in N$, it holds that

$$sat_k(v, \overline{w}) \geq \lfloor qu_k(v) \rfloor$$

Perpetual Lower Quota

A measure to which degree voters receive at least a “fair share” of favorable choices, and thus have a fair say in the decision process.

Let $(N, \overline{A}, \overline{C})$ be a k-decision sequence. A k-choice sequence $\overline{w} \in \overline{C}$ satisfies **perpetual lower quota** if for every voter $v \in N$, it holds that

$$sat_k(v, \overline{w}) \geq \lfloor qu_k(v) \rfloor$$

Proposition 4. There are decision sequences for which no choice sequence exists that satisfies perpetual lower quota.

	1	2	3	$\cdots$	2^n
A_1	$\{0\}$	$\{0\}$	$\{0\}$	$\cdots$	$\{1\}$
A_2	$\{0\}$	$\{1\}$	$\{0\}$	$\cdots$	$\{1\}$
A_3	$\{0\}$	$\{0\}$	$\{1\}$	$\cdots$	$\{1\}$
$\vdots$					
A_n	$\{0\}$	$\{0\}$	$\{0\}$	$\cdots$	$\{1\}$

- The outcome is a 0-1 sequence of length n.
- Each voter agrees with $\frac{n}{2}$ other voters
- For each possible outcome, there is a voter that has satisfaction 0.

Let $(N, \overline{A}, \overline{C}, \overline{w})$ be a k-decision history. The **perpetual lower quota compliance** of $\overline{w}$, $compl(\overline{w})$ is the average proportion of voters in each round that have their perpetual lower quota satisfied:

$$compl(\overline{w}) = \frac{1}{nk} \sum_{i=1}^{k} |\{v \in N \mid sat_i(v, \overline{w}) \geq \lfloor qu_i(v) \rfloor\}|$$

- 20 voters which decide upon 20 decision instances, i.e., we have 20-decision sequences. For each decision 5 alternatives are available—these differ from round to round.

- 20 voters which decide upon 20 decision instances, i.e., we have 20-decision sequences. For each decision 5 alternatives are available—these differ from round to round.
- Generate voters and alternatives in a two-dimensional Euclidean space:
  - Voters are split in two groups and are placed on the $2d$ plane by a bivariate normal distribution. For the first group (6 voters) both x- and y-coordinates are independently drawn from $\mathcal{N}(-0.5, 0.2)$; for the second group (14 voters) for the second group (14 voters) x- and y-coordinates are independently drawn from $\mathcal{N}(0.5, 0.2)$.

- 20 voters which decide upon 20 decision instances, i.e., we have 20-decision sequences. For each decision 5 alternatives are available—these differ from round to round.
- Generate voters and alternatives in a two-dimensional Euclidean space:
  - Voters are split in two groups and are placed on the $2d$ plane by a bivariate normal distribution. For the first group (6 voters) both x- and y-coordinates are independently drawn from $\mathcal{N}(-0.5, 0.2)$; for the second group (14 voters) for the second group (14 voters) x- and y-coordinates are independently drawn from $\mathcal{N}(0.5, 0.2)$.
  - Alternatives are distributed uniformly in the rectangle $[-1, 1] \times [-1, 1]$

- 20 voters which decide upon 20 decision instances, i.e., we have 20-decision sequences. For each decision 5 alternatives are available—these differ from round to round.
- Generate voters and alternatives in a two-dimensional Euclidean space:
    - Voters are split in two groups and are placed on the $2d$ plane by a bivariate normal distribution. For the first group (6 voters) both x- and y-coordinates are independently drawn from $\mathcal{N}(-0.5, 0.2)$; for the second group (14 voters) for the second group (14 voters) x- and y-coordinates are independently drawn from $\mathcal{N}(0.5, 0.2)$.
    - Alternatives are distributed uniformly in the rectangle $[-1, 1] \times [-1, 1]$
    - Voters approve all alternatives that have a distance of at most 1.5 times the distance to the closest alternative. This yields approval sets of size 1.8 on average.

- 20 voters which decide upon 20 decision instances, i.e., we have 20-decision sequences. For each decision 5 alternatives are available—these differ from round to round.
- Generate voters and alternatives in a two-dimensional Euclidean space:
  - Voters are split in two groups and are placed on the $2d$ plane by a bivariate normal distribution. For the first group (6 voters) both x- and y-coordinates are independently drawn from $\mathcal{N}(-0.5, 0.2)$; for the second group (14 voters) for the second group (14 voters) x- and y-coordinates are independently drawn from $\mathcal{N}(0.5, 0.2)$.
  - Alternatives are distributed uniformly in the rectangle $[-1, 1] \times [-1, 1]$
  - Voters approve all alternatives that have a distance of at most 1.5 times the distance to the closest alternative. This yields approval sets of size 1.8 on average.
- Results are based on $10,000$ instances. For each instance and each voting rule, we compute the perpetual lower quota compliance.

Random Serial Dictatorship (RSD)

In each round, a permutation of voters is selected uniformly at random.

We maintain a set X that starts as the set of all alternatives (in this round). One voter after the other can shrink X further to include only approved alternatives; the set X remains unchanged by voters whose approval set has an empty intersection with X.

As soon as X has cardinality 1, this alternative is chosen. If $|X| > 1$ after all voters are considered, one alternative in X is chosen at random.

Figure 1: Perpetual lower quota compliance (values on top of the diagram are medians)

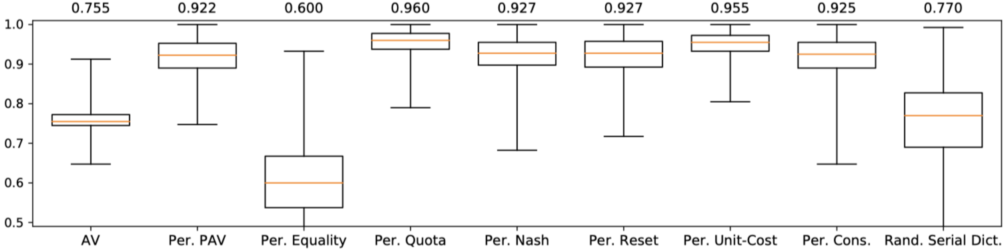

Gini Coefficient

The **Gini coefficient** is a metric of inequality (often used for income distributions); it is 0 for completely equal distributions and 1 for maximally unequal distributions. We use the Gini coefficient to capture inequality in voters' influence on the decision process.

We define the influence of a voter on a given choice as 1 divided by the number of voters supporting this choice. For example, a voter has an influence of 1 on a choice that everyone but her disagrees with; if a choice is supported by all n agents, then their (individual) influence is $\frac{1}{n}$.

Gini Influence Coefficient

Let $(N, \overline{A}, \overline{C}, \overline{w})$ be a k-decision history. The influence of voter $v \in N$ on the choice sequence $\overline{w}$ is

$$inf_k(v, \overline{w}) = \sum_{j \in [k]} \frac{\mathbb{I}_{w_j \in A_j(v)}}{|\{u \in V \mid A_j(u) = w_j\}|}$$

Let a be the average influence $a = \frac{1}{|N|} \sum_{v \in N} infl_k(v)$

The **Gini influence coefficient** of $\overline{w}$ is defined as the Gini coefficient of the sequence $(infl_k(v, \overline{w}))_{v \in N}$.

$$gini_k(\overline{w}) = \frac{1}{2a|N|^2} \sum_{u \in N} \sum_{v \in N} |infl_k(u, \overline{w}) - infl_k(v, \overline{w})|$$

Figure 2: Gini influence coefficient (values on top of the diagram are medians)

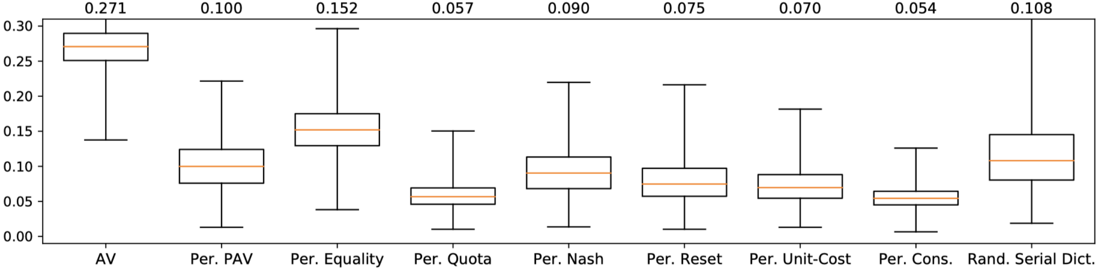

1. The issue of fluctuating voters was not addressed. How should weights be adapted if voters abstain some decisions or enter the decision process at a later stage?

1. The issue of fluctuating voters was not addressed. How should weights be adapted if voters abstain some decisions or enter the decision process at a later stage?

2. In perpetual voting one may encounter the freerider effect: if an alternative is guaranteed to be chosen, it is beneficial for voters to misrepresent their preferences (or abstain) so as not to pay the "price of winning".

1. The issue of fluctuating voters was not addressed. How should weights be adapted if voters abstain some decisions or enter the decision process at a later stage?

2. In perpetual voting one may encounter the freerider effect: if an alternative is guaranteed to be chosen, it is beneficial for voters to misrepresent their preferences (or abstain) so as not to pay the “price of winning”.

3. In a long-term decision process, compromise becomes a powerful concept. For example, if agents assign different importance to individual decisions, compromise can be found by deciding in favor of agents that consider the issue at hand critical, while assigning a higher priority for future decisions to agents that “yielded”.

1. The issue of fluctuating voters was not addressed. How should weights be adapted if voters abstain some decisions or enter the decision process at a later stage?

2. In perpetual voting one may encounter the freerider effect: if an alternative is guaranteed to be chosen, it is beneficial for voters to misrepresent their preferences (or abstain) so as not to pay the "price of winning".

3. In a long-term decision process, compromise becomes a powerful concept. For example, if agents assign different importance to individual decisions, compromise can be found by deciding in favor of agents that consider the issue at hand critical, while assigning a higher priority for future decisions to agents that "yielded".

4. Agents can strategize by controlling the agenda. The outcome will differ if the order of the decision problems faced by the voters is changed.

L. Bulteau, N. Hazon, R. Page, A. Rosenfeld, and N. Talmon. *Justified Representation for Perpetual Voting*. IEEE Access, 2021.

The Preference Intensity Problem

51	49
a	b
b	a

51% of the voters have a *slight* preference for a over b and 49% of the voters have a *strong* preference for b over a.

- Utilitarian considerations suggests that b should win.
- Majoritarian considerations suggests that a should win.

Related Work

P. Harrenstein, M.-L. Lackner and M. Lackner. *A Mathematical Analysis of an Election System Proposed by Gottlob Frege*. Erkenntnis, 2020.

A. Casella (2005). *Storable votes*. Games and Economic Behavior, 51(2), pp. 391 - 419.

A. Casella (2012). *Storable votes: Protecting the minority voice*. Oxford: Oxford University Press.

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. The problem is that constitutional provisions of this type are difficult to enforce and heavy-handed, unsuited to changing realities.

The Storable Votes mechanism (SV): 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 allows the minority to prevail occasionally and yet is anonymous and treats everyone identically.
- SV can apply to direct democracy in large electorates, or to smaller groups, possibly legislatures or committees formed by voters' representatives.
- 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.

Colonel Blotto game

In the original version of the game, 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.

Colonel Blotto game

In the original version of the game, 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
- The game 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.

- N individuals must resolve $K \geq 2$ binary issues, where $\mathcal{K} = \{1, \ldots, K\}$.

- The same M individuals are in favor of all proposals and the remaining $N - M = m$ are opposed to all proposals with $m \leq M$.

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

- 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.
  - *Tyranny of the majority* : If each voter is constrained to cast one vote on each issue, M wins all proposals: 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.

p_m the expected fraction of minority victories.

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.

Voters never cast a vote against their preferences. The focus is on each voter's distribution of votes among the K issues. The action space for each player is:

$$S(K) = \{s = (s_1, \ldots, s_K) \in \mathbb{N}^K \mid \sum_{k=1}^{K} s_k = K\}$$

where s_k is the number of votes cast on issue k.

Given $\boldsymbol{s} = (s^1, \dots, s^m) \in S(K)^m$, the number of votes allocated by the minority to issue k is:

$$v_k^m(\boldsymbol{s}) = \sum_{i=1}^{m} s_k^i$$

Given a 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:

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

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

$\Sigma(K) = \Delta(S(K))$ the set of probability distributions on $S(K)$.

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)

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

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.

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.

We call *group-communication* the protocol in which any subject can only send messages to all members of her group.

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

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

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

Examples

- If $m = 4$, $M = 5$ and $K = 3$, there exists an equilibrium in which the minority wins two of the three issues.
- If $m = 1$ and $M = 2$ and $K = 4$, there exists no pure-strategy equilibrium.

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.

For any c factor of K, define σ^c (note τ^c for a majority player) as follows: choose randomly K/c issues, and allocate c votes to each of the selected issues.

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

Proposition 2. Suppose K is even and M is odd. Then $(\boldsymbol{\sigma}^2, \boldsymbol{\tau}^1)$ is an equilibrium if $M \leq m+1$ with

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

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

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 $2(m-1)$.
- 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.
- 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.

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:

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

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?

- 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.

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?

- 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.

- 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).

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 $s_\pi = (s_{\pi(1)}, \ldots, s_{\pi(K)})$.

E.g., σ^c for any c is neutral.

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 $s_\pi = (s_{\pi(1)}, \ldots, s_{\pi(K)})$.

E.g., σ^c for any c is neutral.

Proposition 4. For all $M \leq mK$, there exist a number $\underline{k} \in \{1, \ldots, K\}$ such that if every minority player's strategy: (i) is neutral, and (ii) allocated votes on no more than $\underline{k}$ issues with probability 1, then for any strategy profile of the majority $\boldsymbol{\tau}$,

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

- 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.

- 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?).

- Suppose $K = 4$ and $M = 10$, and $m \in \{1, \ldots, 10\}$.
- Minority voters adopt the σ^c strategies with $c = 2$ (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 chosen randomly).
- Consider two rules for the majority: (1) each majority voter casts his votes randomly and independently over all issues (an upper bound on p_m) or (2) all majority voters together best respond to the minority (the lower bound).

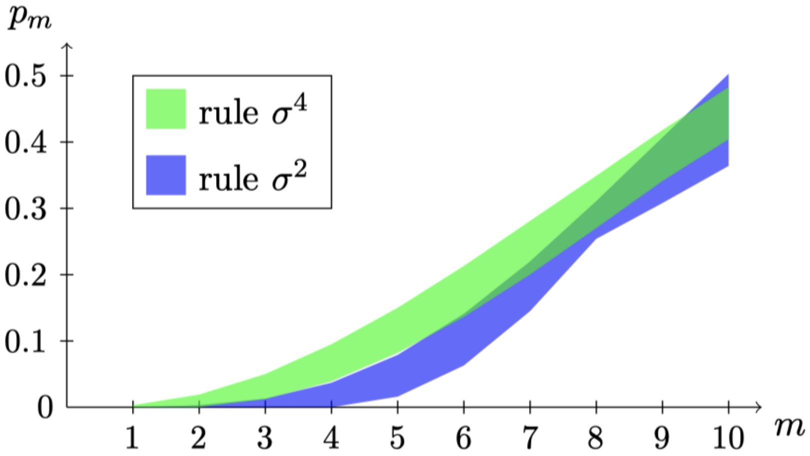

Figure 1: Minority payoffs for two minority rules ($M = 10$)