New Perspectives on Social Choice

Eric Pacuit, University of Maryland

Lecture 4: Split Cycle

PHIL 808K

Let us say that a voting method F is regarded as a *pre-tiebreaking voting method* if one regards $F(\boldsymbol{P})$ as the set of undefeated candidates and regards any further narrowing of $F(\boldsymbol{P})$ as “tiebreaking.”

Let us say that a voting method F is regarded as a *pre-tiebreaking voting method* if one regards $F(\boldsymbol{P})$ as the set of undefeated candidates and regards any further narrowing of $F(\boldsymbol{P})$ as "tiebreaking."

The political significance of this distinction is that if $F(\boldsymbol{P})$ contains a single winner, then that winner may be viewed as having a stronger mandate from voters, as a result of a more unambiguous election, than a candidate who is among several undefeated candidates in $F(\boldsymbol{P})$ but wins by some further tiebreaking process.

Since there can be multiple undefeated candidates, the question arises of how to pick an ultimate winner from among the undefeated.

- non-anonymous tiebreaker (e.g., let the Chair decide)
- non-neutral tiebreaker (e.g., use seniority to decide among the undefeated candidates)
- non-deterministic tiebreaker (e.g., randomly choose an undefeated candidate)
- apply an anonymous, neutral, and deterministic tiebreaker before resorting to tiebreakers that violates one of the above properties.

- The Condorcet method as a pre-tiebreaking method, then Borda to break ties
- Top Cycle as a pre-tiebreaking method, then IRV as a tiebreaker
- Copeland as a pre-tiebreaking method, then *global*/*local* Borda as a tiebreaker
- Split Cycle as a pre-tiebreaking method, then *global*/*local* Plurality as a tiebreaker
- Split Cycle as a pre-tiebreaking method, then *global*/*local* Beat Path as a tiebreaker
- ...

What is a good anonymous, neutral and deterministic tiebreaker?

- No anonymous and neutral method is resolute, but the tiebreaker should be more resolute than the pre-tiebreaking method.
- Why don't you use the tiebreaker method from the beginning?
  - Choose a method for the tiebreaker that *refines* the pre-tiebreaking method.

Sizes of Winning Sets

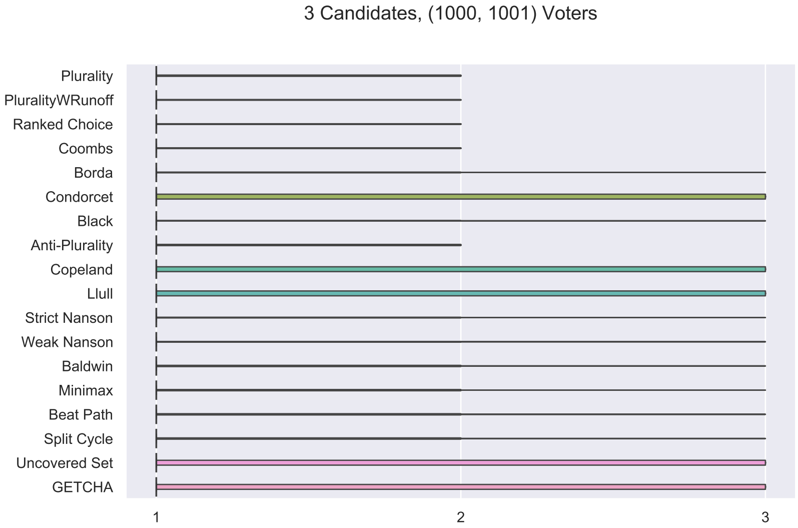

Sizes of Winning Sets

Sizes of Winning Sets

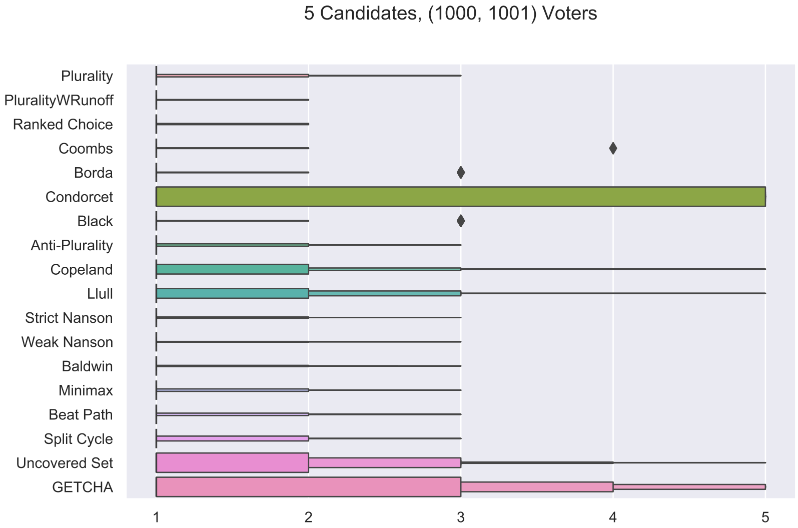

Sizes of Winning Sets

Sizes of Winning Sets

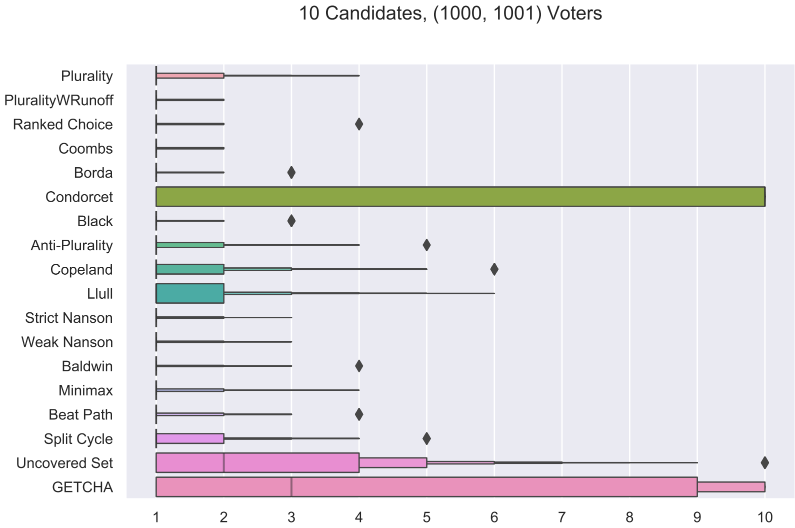

Sizes of Winning Sets

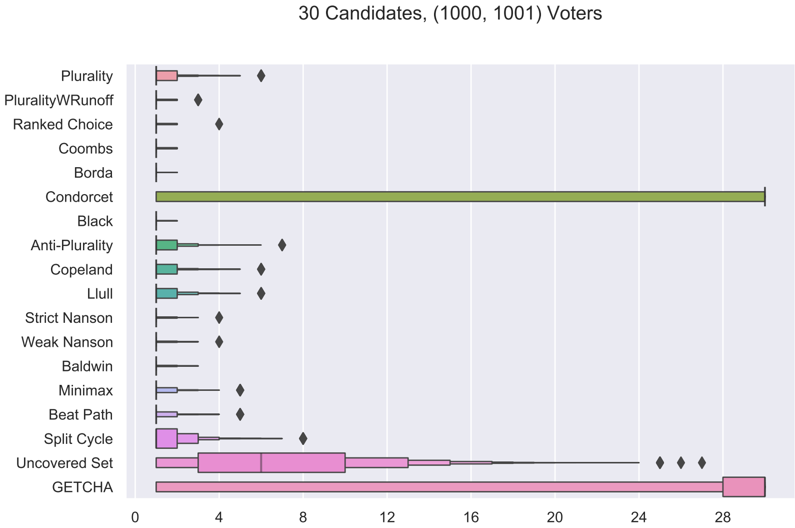

Any definition of defeat satisfying Anonymity and Neutrality will yield multiple undefeated candidates in some profiles.

Any definition of defeat satisfying Anonymity and Neutrality will yield multiple undefeated candidates in some profiles.

But some definitions of defeat are still more resolute than others...

Any definition of defeat satisfying Anonymity and Neutrality will yield multiple undefeated candidates in some profiles.

But some definitions of defeat are still more resolute than others...

Definition

A VCCR is **asymptotically resolvable** if the proportion of profiles with multiple undefeated candidates approaches 0 as the number of voters approaches ∞.

Any definition of defeat satisfying Anonymity and Neutrality will yield multiple undefeated candidates in some profiles.

But some definitions of defeat are still more resolute than others...

Definition

A VCCR is **asymptotically resolvable** if the proportion of profiles with multiple undefeated candidates approaches 0 as the number of voters approaches ∞.

Examples of asymptotically resolvable VCCRs are Plurality, Borda, and Beat Path.

Quasi-Resolute

Definition

A voting method F is *quasi-resolute* if for every uniquely-weighted $\boldsymbol{P} \in \mathrm{dom}(F)$, $|F(\boldsymbol{P})| = 1$.

Quasi-Resolute

Definition

A voting method F is *quasi-resolute* if for every uniquely-weighted $\boldsymbol{P} \in \mathrm{dom}(F)$, $|F(\boldsymbol{P})| = 1$.

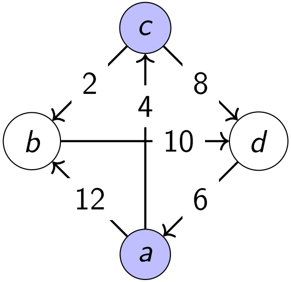

Both a and c are undefeated according to Split Cycle; Ranked Pairs picks only a; and Beat Path picks only c.

Violations of Quasi-Resoluteness

The known methods that satisfy Binary Expansion violate Asymptotic Resolvability/Quasi-Resoluteness.

Voting Method	3	4	5	6	7	8	9	10	20	30
Split Cycle	1	1.01	1.03	1.06	1.08	1.11	1.14	1.16	1.42	1.62
Uncovered Set	1.17	1.35	1.53	1.71	1.9	2.09	2.26	2.46	4.56	6.82
Top Cycle	1.17	1.44	1.8	2.21	2.72	3.31	3.94	4.68	13.55	22.94

Figure: Estimated **average sizes of winning sets** for profiles with a given number of candidates (top row) in the limit as the number of voters goes to infinity, obtained using the Monte Carlo simulation technique in M. Harrison Trainor, "An Analysis of Random Elections with Large Numbers of Voters," arXiv:2009.02979.

Stability for Winners

If  wins and 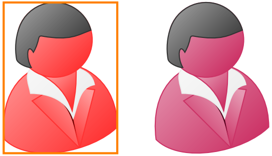 wins

Stability for Winners

If  wins and 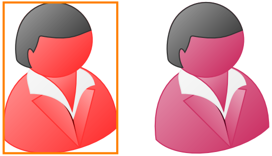 wins

then 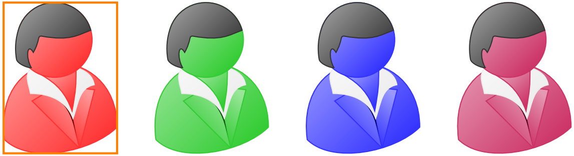 wins .

Stability for Winners

Definition

A VCCR satisfies **Stability for Winners** if for any profile $\boldsymbol{P}$ and $a, b \in X(\boldsymbol{P})$, if a is undefeated in $\boldsymbol{P}_{-b}$ and $Margin_{\boldsymbol{P}}(a, b) > 0$, then a is undefeated in $\boldsymbol{P}$.

Stability for Winners

Definition

A VCCR satisfies **Stability for Winners** if for any profile $\boldsymbol{P}$ and $a, b \in X(\boldsymbol{P})$, if a is undefeated in $\boldsymbol{P}_{-b}$ and $Margin_{\boldsymbol{P}}(a, b) > 0$, then a is undefeated in $\boldsymbol{P}$.

Example violations:

Stability for Winners

Definition

A VCCR satisfies **Stability for Winners** if for any profile $\boldsymbol{P}$ and $a, b \in X(\boldsymbol{P})$, if a is undefeated in $\boldsymbol{P}_{-b}$ and $Margin_{\boldsymbol{P}}(a, b) > 0$, then a is undefeated in $\boldsymbol{P}$.

Example violations:

- arguably the 2000 US Presidential Election in Florida, run with Plurality Voting, where $\boldsymbol{a}$ was Al Gore and $\boldsymbol{b}$ was Ralph Nader.

Stability for Winners

Definition

A VCCR satisfies **Stability for Winners** if for any profile $\boldsymbol{P}$ and $a, b \in X(\boldsymbol{P})$, if a is undefeated in $\boldsymbol{P}_{-b}$ and $Margin_{\boldsymbol{P}}(a, b) > 0$, then a is undefeated in $\boldsymbol{P}$.

Example violations:

- arguably the 2000 US Presidential Election in Florida, run with Plurality Voting, where $\boldsymbol{a}$ was Al Gore and $\boldsymbol{b}$ was Ralph Nader.
- definitely the 2009 mayoral election in Burlington, Vermont, run with Instant Runoff Voting, where $\boldsymbol{a}$ was the Democrat and $\boldsymbol{b}$ was the Republican.

Stability for Winners

Definition

A VCCR satisfies **Stability for Winners** if for any profile $\boldsymbol{P}$ and $a, b \in X(\boldsymbol{P})$, if a is undefeated in $\boldsymbol{P}_{-b}$ and $Margin_{\boldsymbol{P}}(a, b) > 0$, then a is undefeated in $\boldsymbol{P}$.

Example violations:

- arguably the 2000 US Presidential Election in Florida, run with Plurality Voting, where $\boldsymbol{a}$ was Al Gore and $\boldsymbol{b}$ was Ralph Nader.
- definitely the 2009 mayoral election in Burlington, Vermont, run with Instant Runoff Voting, where $\boldsymbol{a}$ was the Democrat and $\boldsymbol{b}$ was the Republican.
- there are also violations in profiles with no Condorcet winner.

Stability for Winners

Definition

A VCCR satisfies **Stability for Winners** if for any profile $\boldsymbol{P}$ and $a, b \in X(\boldsymbol{P})$, if a is undefeated in $\boldsymbol{P}_{-b}$ and $Margin_{\boldsymbol{P}}(a, b) > 0$, then a is undefeated in $\boldsymbol{P}$.

Example violations:

- arguably the 2000 US Presidential Election in Florida, run with Plurality Voting, where $\boldsymbol{a}$ was Al Gore and $\boldsymbol{b}$ was Ralph Nader.
- definitely the 2009 mayoral election in Burlington, Vermont, run with Instant Runoff Voting, where $\boldsymbol{a}$ was the Democrat and $\boldsymbol{b}$ was the Republican.
- there are also violations in profiles with no Condorcet winner.

Proposition

Anonymity, Neutrality, Monotonicity (for two-candidate profiles), and Coherent IIA together imply Stability for Winners.

The Cost of Quasi-Resoluteness

Theorem (W. Holliday, EP, and S. Zahedian)

There is no Anonymous and Neutral voting method that satisfies Binary Expansion and Quasi-Resoluteness.

Moral: Making room for tiebreaking (runoff, lottery, etc.) is necessary and sufficient to find voting methods that satisfy Binary Expansion.

Multiple claims based on stability

The basic problem is that inevitably there are profiles with *multiple* candidates who have the same kind of claim to winning based on stability for winners:

Multiple claims based on stability

The basic problem is that inevitably there are profiles with *multiple* candidates who have the same kind of claim to winning based on stability for winners:

wins

and

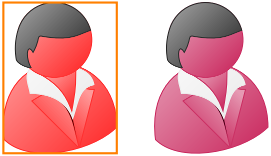

wins

Multiple claims based on stability

The basic problem is that inevitably there are profiles with *multiple* candidates who have the same kind of claim to winning based on stability for winners:

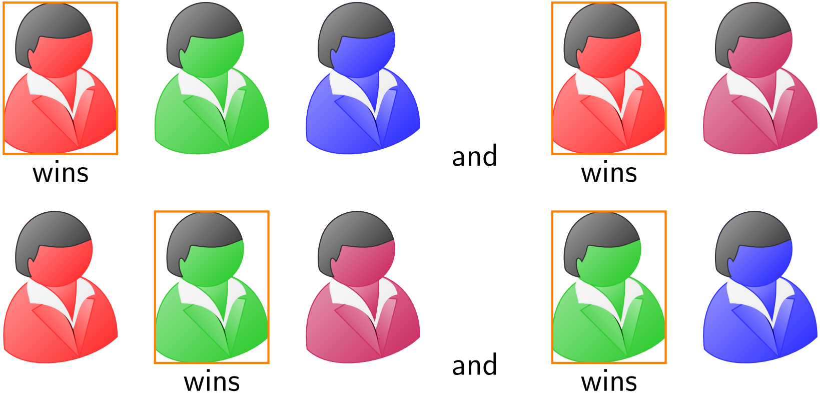

Multiple claims based on stability

The basic problem is that inevitably there are profiles with *multiple* candidates who have the same kind of claim to winning based on stability for winners:

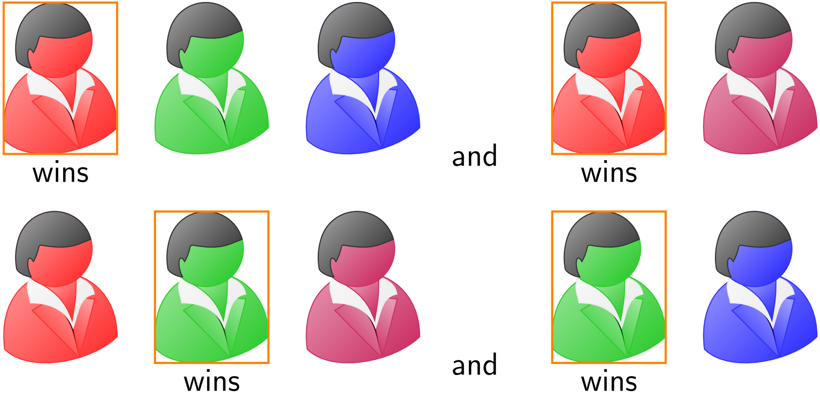

In such a situation—and only such a situation—it is legitimate to violate stability for winners for one of red or green in the name of tiebreaking between them.

Condorcetian candidates

Definition

Given a voting method F, profile $\boldsymbol{P}$, and $a \in X(\boldsymbol{P})$, we say that *a is Condorcetian for F in $\boldsymbol{P}$* if there is some $b \in X(\boldsymbol{P})$ such that $a \in F(\boldsymbol{P}_{-b})$ and $Margin_{\boldsymbol{P}}(a, b) > 0$.

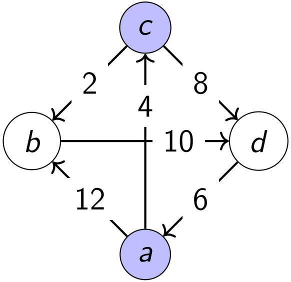

- There are two Condorcetian candidates a and c
- Beat Path elects c
- Ranked Pairs elects a

Stability for Winners with Tiebreaking

Definition

A voting method satisfies **Stability for Winners with Tiebreaking** if for any profile $\mathbf{P}$ and $a, b \in X(\mathbf{P})$, if a wins in $\mathbf{P}_{-b}$ and $Margin_{\mathbf{P}}(a, b) > 0$,

Stability for Winners with Tiebreaking

Definition

A voting method satisfies **Stability for Winners with Tiebreaking** if for any profile $\mathbf{P}$ and $a, b \in X(\mathbf{P})$, if a wins in $\mathbf{P}_{-b}$ and $\mathit{Margin}_{\mathbf{P}}(a, b) > 0$, then either

Stability for Winners with Tiebreaking

Definition

A voting method satisfies **Stability for Winners with Tiebreaking** if for any profile $\mathbf{P}$ and $a, b \in X(\mathbf{P})$, if a wins in $\mathbf{P}_{-b}$ and $Margin_{\mathbf{P}}(a, b) > 0$, then either

- a wins in $\mathbf{P}$ or

Stability for Winners with Tiebreaking

Definition

A voting method satisfies **Stability for Winners with Tiebreaking** if for any profile $\mathbf{P}$ and $a, b \in X(\mathbf{P})$, if a wins in $\mathbf{P}_{-b}$ and $Margin_{\mathbf{P}}(a, b) > 0$, then either

- a wins in $\mathbf{P}$ or
- there are $a', b' \in X(\mathbf{P})$ such that a' wins in $\mathbf{P}_{-b'}$, $Margin_{\mathbf{P}}(a', b') > 0$, and a' wins in $\mathbf{P}'$.

That is, all winners are Condorcetian.

Recursion to the Rescue: Stable Voting

Our proposed voting method is Stable Voting, defined *recursively* as follows:

Recursion to the Rescue: Stable Voting

Our proposed voting method is Stable Voting, defined *recursively* as follows:

- If only one candidate a appears on all ballots, then a wins.

Recursion to the Rescue: Stable Voting

Our proposed voting method is Stable Voting, defined *recursively* as follows:

- If only one candidate a appears on all ballots, then a wins.
- Otherwise list all head-to-head matches a vs. b, where a is undefeated according to Split Cycle, in order from the largest to the smallest margin of a vs. b.

Recursion to the Rescue: Stable Voting

Our proposed voting method is Stable Voting, defined *recursively* as follows:

- If only one candidate a appears on all ballots, then a wins.
- Otherwise list all head-to-head matches a vs. b, where a is undefeated according to Split Cycle, in order from the largest to the smallest margin of a vs. b.

  Find the first match such that a wins according to Stable Voting *after b is removed from all ballots*; this a is the winner for the original set of ballots.

Recursion to the Rescue: Stable Voting

Our proposed voting method is Stable Voting, defined *recursively* as follows:

- If only one candidate $\boldsymbol{a}$ appears on all ballots, then $\boldsymbol{a}$ wins.
- Otherwise list all head-to-head matches $\boldsymbol{a}$ vs. $\boldsymbol{b}$, where $\boldsymbol{a}$ is undefeated according to Split Cycle, in order from the largest to the smallest margin of $\boldsymbol{a}$ vs. $\boldsymbol{b}$.

  Find the first match such that $\boldsymbol{a}$ wins according to Stable Voting *after* $\boldsymbol{b}$ *is removed from all ballots*; this $\boldsymbol{a}$ is the winner for the original set of ballots.

W. Holliday and EP. *Stable Voting*. arXiv:2108.00542 [econ.TH].

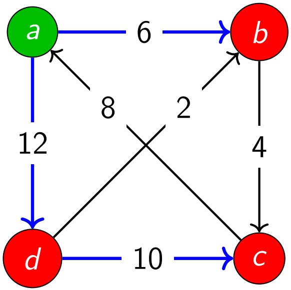

Stable Voting winner: a

Beat Path winner: a

Ranked Pairs winner: a

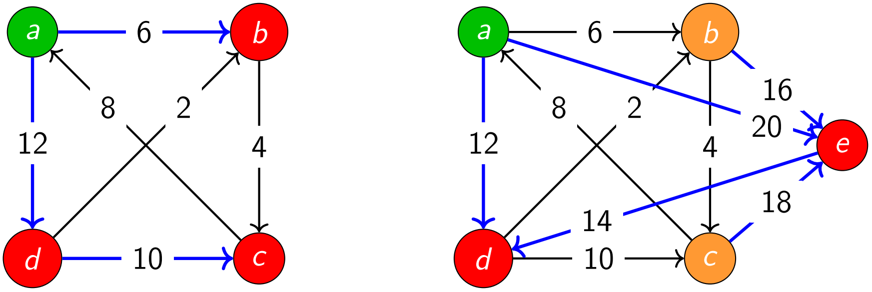

Stable Voting winner: *a*
Beat Path winner: *a*
Ranked Pairs winner: *a*

Stable Voting winner: *a*
Beat Path winner: *b*
Ranked Pairs winner: *c*

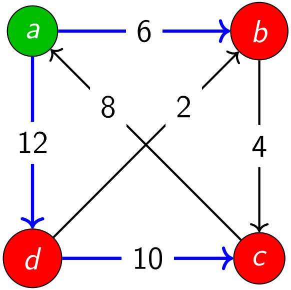

Stable Voting winner: a
Beat Path winner: a
Ranked Pairs winner: a

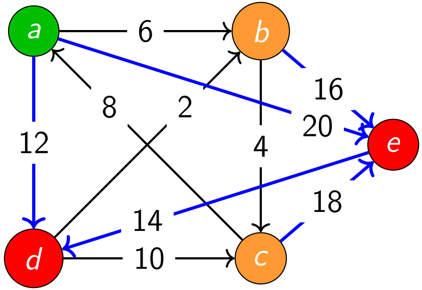

Stable Voting winner: a
Beat Path winner: b
Ranked Pairs winner: c

On the right, SV chooses the winner by going down the list of matches:

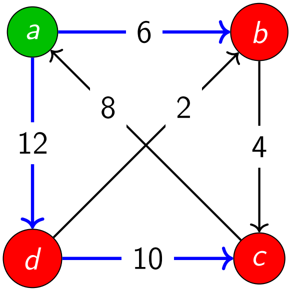

Stable Voting winner: *a*
Beat Path winner: *a*
Ranked Pairs winner: *a*

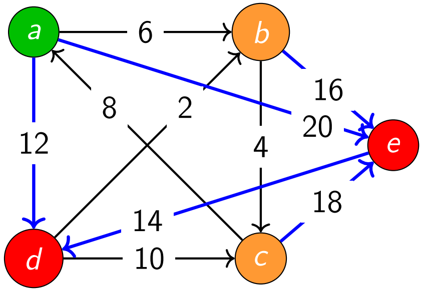

Stable Voting winner: *a*
Beat Path winner: *b*
Ranked Pairs winner: *c*

On the right, SV chooses the winner by going down the list of matches:

- ***a*** vs. ***e***: margin of 20.

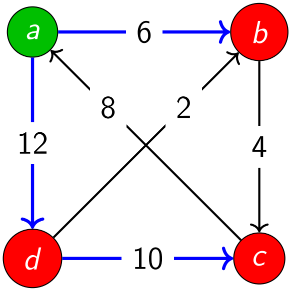

Stable Voting winner: *a*
Beat Path winner: *a*
Ranked Pairs winner: *a*

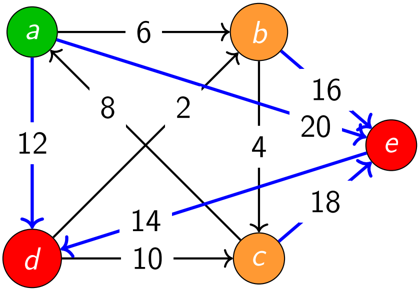

Stable Voting winner: *a*
Beat Path winner: *b*
Ranked Pairs winner: *c*

On the right, SV chooses the winner by going down the list of matches:

- ***a*** vs. ***e***: margin of 20.
  a wins after removing ***e***. Hence ***a* is elected**.

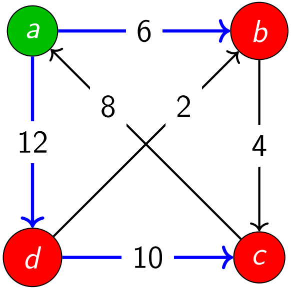

Stable Voting winner: *a*

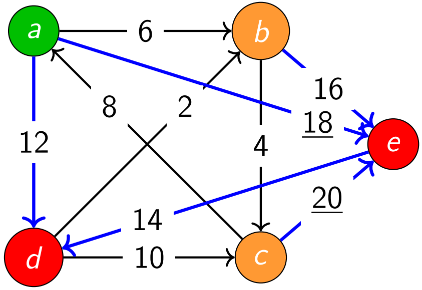

Stable Voting winner: *a*

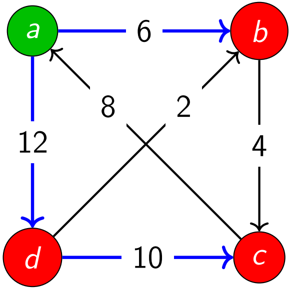

Stable Voting winner: a

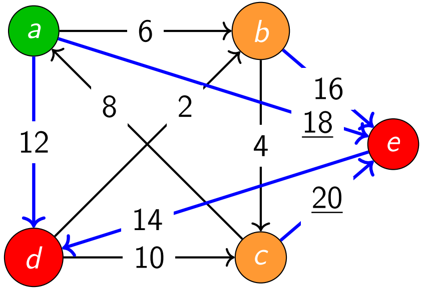

Stable Voting winner: a

On the right, SV chooses the winner by going down the list of matches:

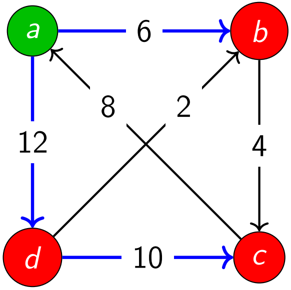

Stable Voting winner: *a*

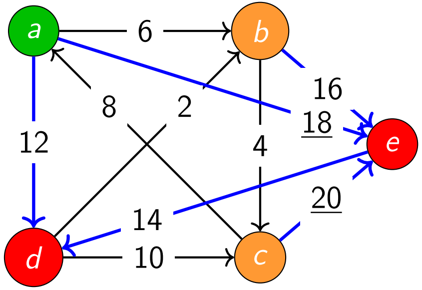

Stable Voting winner: *a*

On the right, SV chooses the winner by going down the list of matches:

- ***c*** vs. ***e***: margin of 20.

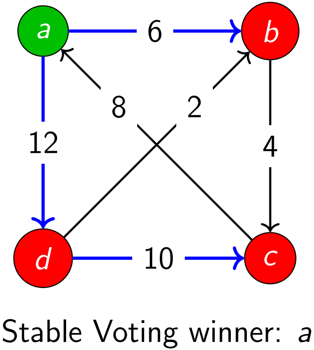

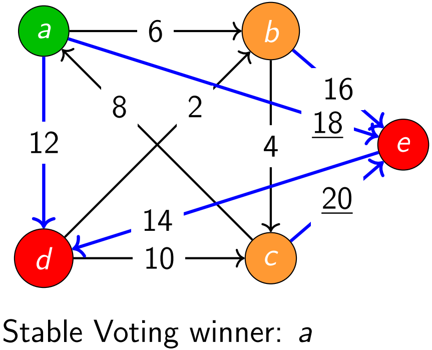

On the right, SV chooses the winner by going down the list of matches:

- c vs. e: margin of 20.
  c loses after removing e. Continue to next match:

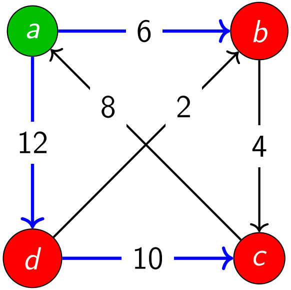

Stable Voting winner: a

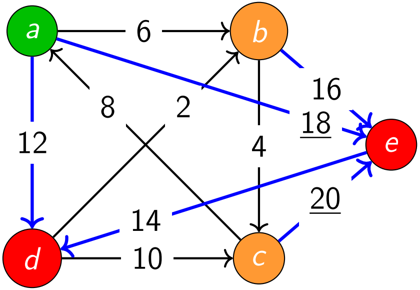

Stable Voting winner: a

On the right, SV chooses the winner by going down the list of matches:

- c vs. e: margin of 20.
  c loses after removing e. Continue to next match:
- a vs. e: margin of 18.

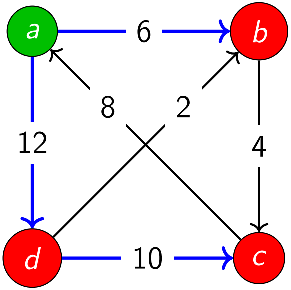

Stable Voting winner: *a*

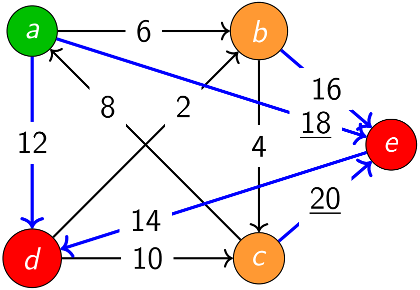

Stable Voting winner: *a*

On the right, SV chooses the winner by going down the list of matches:

- ***c*** vs. ***e***: margin of 20.
  c loses after removing ***e***. Continue to next match:
- ***a*** vs. ***e***: margin of 18.
  a wins after removing ***e***. Hence ***a* is elected**.

Stable Voting

Good news: Stable Voting satisfies **Stability for Winners with Tiebreaking** and **Quasi-resoluteness**.

Stable Voting

Good news: Stable Voting satisfies **Stability for Winners with Tiebreaking** and **Quasi-resoluteness**.

In fact, SV has a remarkable ability to avoid ties even in elections with small numbers of voters that can produce tied margins.

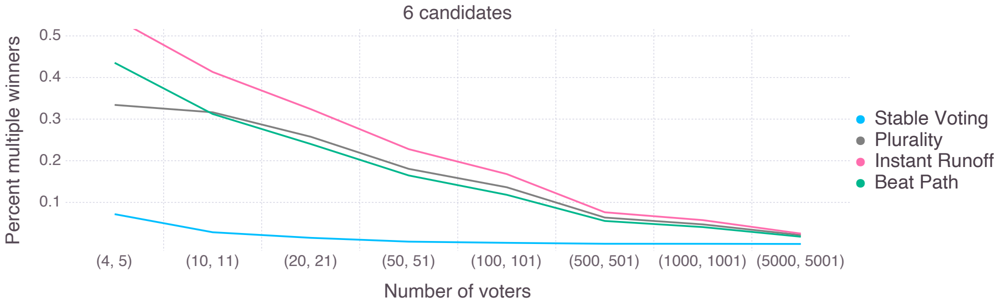
6 candidates
Percent multiple winners
0.5
0.4
0.3
0.2
0.1
(4, 5)
(10, 11)
(20, 21)
(50, 51)
(100, 101)
(500, 501)
(1000, 1001)
(5000, 5001)
Number of voters
Stable Voting
Plurality
Instant Runoff
Beat Path

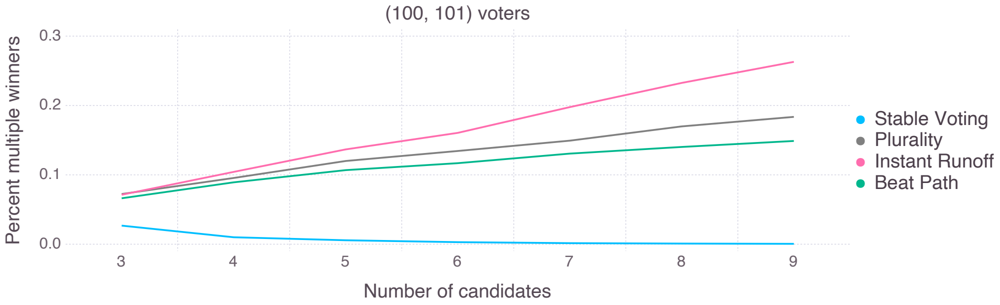
(100, 101) voters
Percent multiple winners
0.3
0.2
0.1
0.0
3
4
5
6
7
8
9
Number of candidates
Stable Voting
Plurality
Instant Runoff
Beat Path

Costs of Stable Voting

For truth in advertising, there are some costs of Stable Voting:

Costs of Stable Voting

For truth in advertising, there are some costs of Stable Voting:

1. Computing the SV winners using our current recursive implementation can be computationally expensive above 20 candidates.

Costs of Stable Voting

For truth in advertising, there are some costs of Stable Voting:

1. Computing the SV winners using our current recursive implementation can be computationally expensive above 20 candidates.

2. There are some violations—in an extremely small fraction of profiles—of voting criteria satisfied by some other voting methods, such as *monotonicity*.

Costs of Stable Voting

For truth in advertising, there are some costs of Stable Voting:

1. Computing the SV winners using our current recursive implementation can be computationally expensive above 20 candidates.
2. There are some violations—in an extremely small fraction of profiles—of voting criteria satisfied by some other voting methods, such as *monotonicity*.

Re 1, we can handle larger profiles that are uniquely weighted with up to 20 candidates in the "Smith set." This covers many voting contexts.

Costs of Stable Voting

For truth in advertising, there are some costs of Stable Voting:

1. Computing the SV winners using our current recursive implementation can be computationally expensive above 20 candidates.
2. There are some violations—in an extremely small fraction of profiles—of voting criteria satisfied by some other voting methods, such as *monotonicity*.

Re 1, we can handle larger profiles that are uniquely weighted with up to 20 candidates in the "Smith set." This covers many voting contexts.

Re 2, the frequency with which Stable Voting violates monotonicity is minuscule compared to the frequency for Instant Runoff (in use in the Bay Area and NYC).

Monotonicity

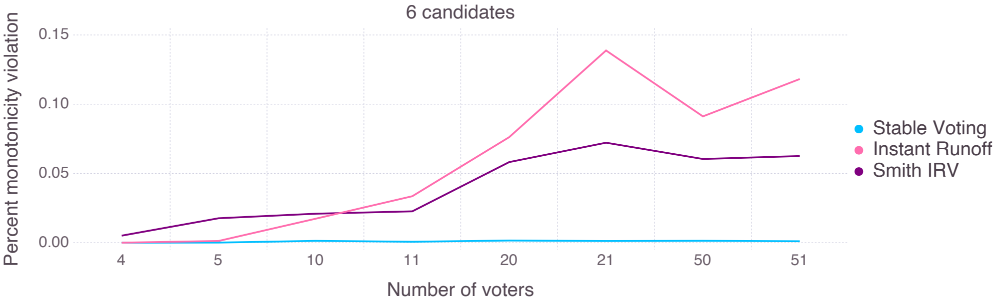

Demo

stablevoting.org

Since there can be multiple undefeated candidates, the question arises of how to pick an ultimate winner from among the undefeated.

- non-anonymous tiebreaker (e.g., let the Chair decide)
- non-neutral tiebreaker (e.g., use seniority to decide among the undefeated candidates)

$\Rightarrow$ non-deterministic tiebreaker (e.g., randomly choose an undefeated candidate)

✓ apply an anonymous, neutral, and deterministic tiebreaker before resorting to tiebreakers that violates one of the above properties.

Even-Chance Tiebreaking

All anonymous and neutral voting methods F may select more than one winner in a profile $\mathbf{P}$.

Even-Chance Tiebreaking

All anonymous and neutral voting methods F may select more than one winner in a profile $\boldsymbol{P}$.

When $F(\boldsymbol{P})$ is not a singleton set, one option is to use an even-chance lottery on $F(\boldsymbol{P})$ to break the tie and select a unique ultimate winner.

Definition

For a voting method F, let F^{eve} be the probabilistic voting method such that for any profile $\boldsymbol{P}$ and $x \in X(\boldsymbol{P})$, $F^{\text{eve}}(\boldsymbol{P})(x) = 1/|F(\boldsymbol{P})|$ if $x \in F(\boldsymbol{P})$ and $F^{\text{eve}}(\boldsymbol{P})(x) = 0$ otherwise.

Probabilistic Social Choice

Voters	Rankings
1	$a\ b\ c\ d$
2	$b\ a\ d\ c$
3	$b\ d\ a\ c$
4	$d\ c\ a\ b$

Aggregation Method

Winner

Winning set

Defeat Relation

Committee

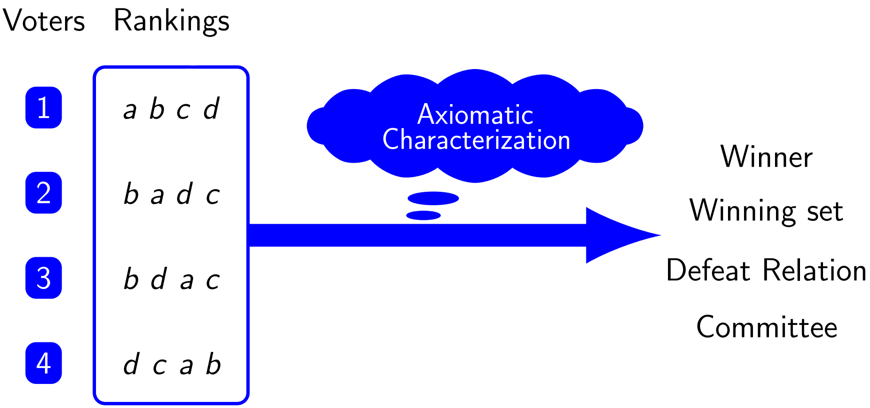
Voters
Rankings
1
2
3
4
a b c d
b a d c
b d a c
d c a b
Axiomatic Characterization
Winner
Winning set
Defeat Relation
Committee

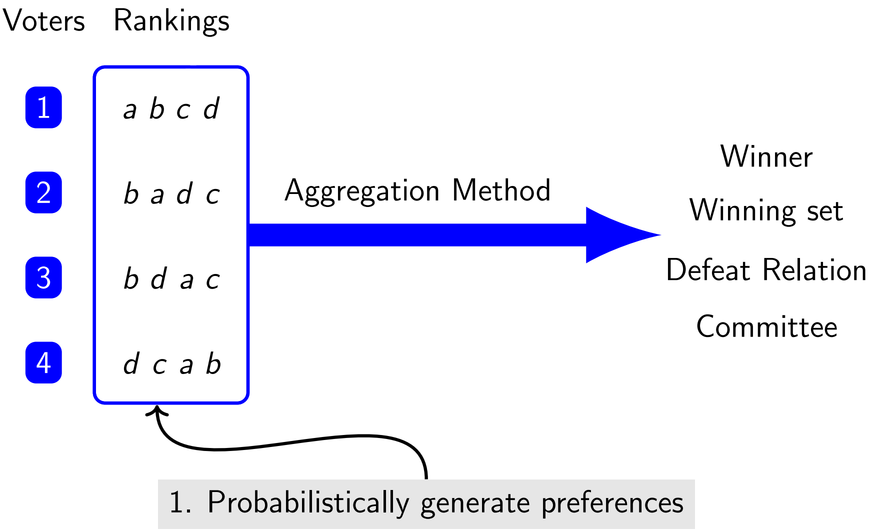
Voters
Rankings
1
2
3
4
a b c d
b a d c
b d a c
d c a b
Aggregation Method
Winner
Winning set
Defeat Relation
Committee
1. Probabilistically generate preferences

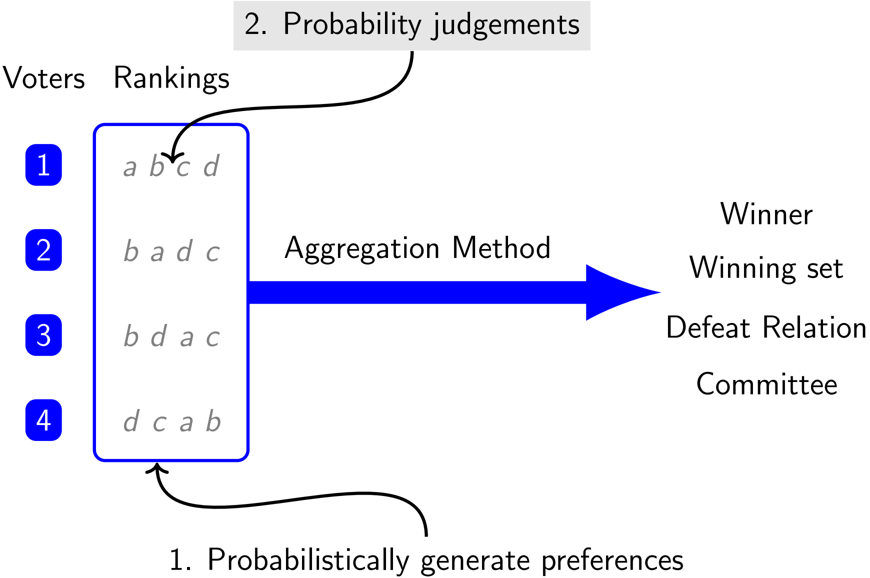

2. Probability judgements
Voters
Rankings
1
2
3
4
a b c d
b a d c
b d a c
d c a b
Aggregation Method
Winner
Winning set
Defeat Relation
Committee
1. Probabilistically generate preferences

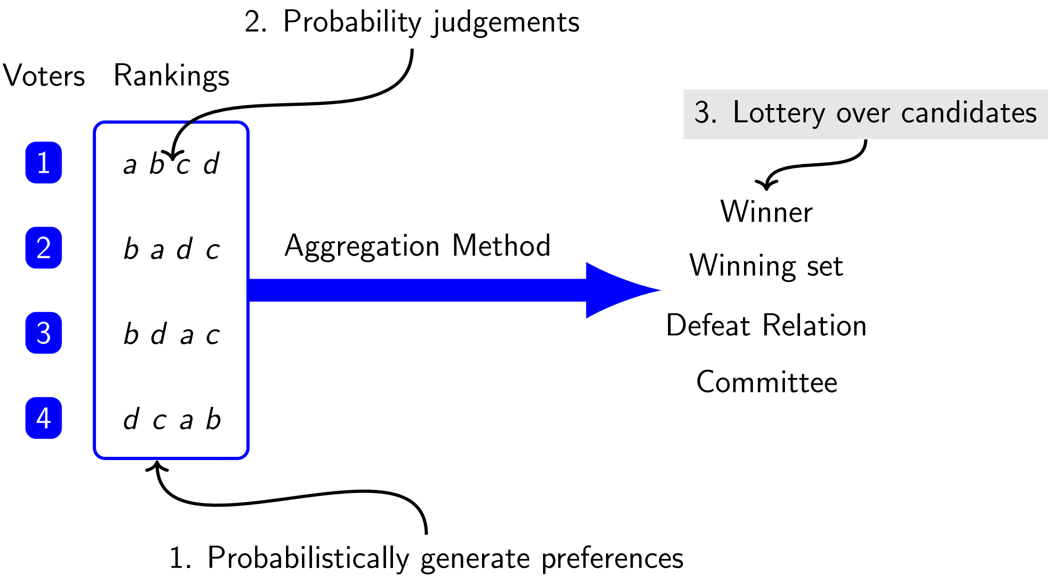
2. Probability judgements
Voters
Rankings
1
2
3
4
a b c d
b a d c
b d a c
d c a b
Aggregation Method
3. Lottery over candidates
Winner
Winning set
Defeat Relation
Committee
1. Probabilistically generate preferences

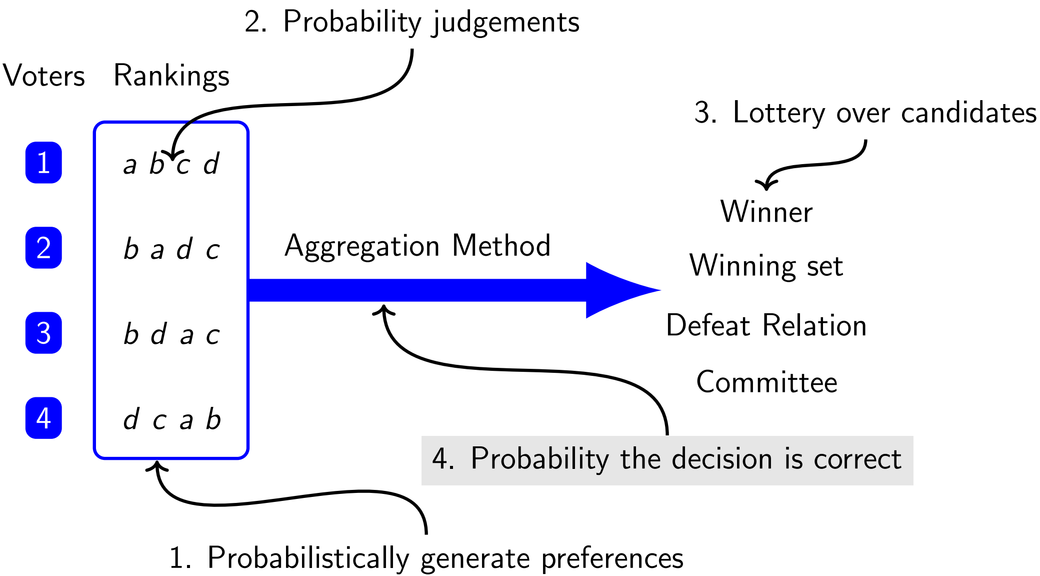
2. Probability judgements
Voters
Rankings
1
2
3
4
a b c d
b a d c
b d a c
d c a b
Aggregation Method
3. Lottery over candidates
Winner
Winning set
Defeat Relation
Committee
4. Probability the decision is correct
1. Probabilistically generate preferences

F. Brandt. *Rolling the Dice: Recent Results in Probabilistic Social Choice*. Handbook of Computational Social Choice, 2016.

Let V be a set of voters, X a set of m alternatives.

The set of all lotteries over X is:

$$\Delta(X) = \{p \in \mathbb{R}^X \mid p(x) \geq 0 \text{ for all } x \in X \text{ and } \sum_{x \in X} p(x) = 1\}$$

For $p \in \Delta(X)$, $supp(p) = \{x \mid p(x) > 0\}$

p is **degenerate** when $|supp(p)| = 1$.

We write lotteries as convex combinations of alternatives, e.g., the uniform lottery on $\{a, b\}$ where $p(a) = p(b) = 1/2$ is denoted as $1/2a + 1/2b$.

Probabilistic Voting Methods

A **probabilistic social choice function (PSCF)** is a map $F : O(X)^V \to \wp(\Delta(A)) \setminus \emptyset$ such that for all $\boldsymbol{P}$, $F(\boldsymbol{P})$ is a convex set of lotteries.

Anonymity and neutrality can be defined as usual.

Random (Serial) Dictator

Random dictatorship: A voter is picked uniformly at random and this voter's most-preferred alternative is selected. Thus, the probabilities assigned by RD are directly proportional to the number of agents who top-rank a given alternative (or, in other words, the alternative's plurality score).

Random serial dictatorship (RSD): RSD selects a permutation of the agents uniformly at random and then sequentially allows agents in the order of the permutation to narrow down the set of alternatives to their most preferred of the remaining ones.

Proportional Borda

Proportional Borda: Assign probabilities to the alternatives that are proportional to their Borda scores.

3	2
a	b
b	c
c	a

- $\mathit{Plurality}^{eve}$ lottery is a
- RD lottery is $\frac{3}{5}a + \frac{2}{5}b$
- Borda^{eve} lottery is b
- Borda_{pro} lottery is $\frac{6}{15}a + \frac{7}{15}b + \frac{2}{15}c$

Margin Matrix/Graph

49	48	3
c	a	b
a	b	c
b	c	a

$$\begin{array}{c c} & \begin{array}{ccc} a & b & c \end{array} \\ \begin{array}{c} a \\ b \\ c \end{array} & \begin{pmatrix} 0 & 94 & -4 \\ -94 & 0 & 2 \\ 4 & -2 & 0 \end{pmatrix} \end{array}$$

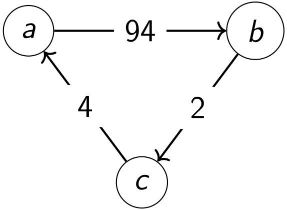

Margin Matrix/Graph

49	48	3
c	a	b
a	b	c
b	c	a

$$\begin{array}{c c} & \begin{array}{ccc} a & \quad b & \quad c \end{array} \\ \begin{array}{c} a \\ b \\ c \end{array} & \begin{pmatrix} 0 & 94 & -4 \\ -94 & 0 & 2 \\ 4 & -2 & 0 \end{pmatrix} \end{array}$$

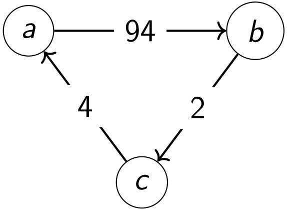

If the output of a neutral PSCF F only depends on the margin matrix/graph M, F is called **pairwise**. An advantage of pairwise PSCFs is that they are applicable even when individual preferences are incomplete or intransitive.

Maximal Lotteries

G. Kreweras. *Aggregation of preference orderings*. In Mathematics and Social Sciences I: Proceedings of the seminars of Menthon-Saint-Bernard, France (1–27 July 1960) and of Gösing, Austria (3 - 27 July 1962), pages 73 - 79, 1965.

P. C. Fishburn. *Probabilistic social choice based on simple voting comparisons*. Review of Economic Studies, 51(4):683-692, 1984.

Maximal Lotteries

H. Aziz, F. Brandl, F. Brandt, and M. Brill. *On the tradeoff between efficiency and strategyproofness*. Games and Economic Behavior, 110:1 - 18, 2018.

F.. Brandl, F. Brandt, and H. G. Seedig. *Consistent probabilistic social choice*. Econometrica, 84(5):1839 - 1880, 2016.

F. Brandl, F. Brandt, M. Eberl, and C. Geist. *Proving the incompatibility of efficiency and strategyproofness via SMT solving*. Journal of the ACM, 65(2):1 - 28, 2018.

F. Brandl, F. Brandt, and C. Stricker. *An Analytical and Experimental Comparison of Maximal Lottery Schemes* . forthcoming Social Choice and Welfare, 2021.

Maximal Lotteries

M admits a (weak) Condorcet winner if M contains a nonnegative row, i.e., there is a standard unit vector v such that

$$v^T M \geq 0$$

Maximal Lotteries

M admits a (weak) Condorcet winner if M contains a nonnegative row, i.e., there is a standard unit vector v such that

$$v^T M \geq 0$$

1	1	1
a	b	c
b	a	a
c	c	b

Maximal Lotteries

M admits a (weak) Condorcet winner if M contains a nonnegative row, i.e., there is a standard unit vector v such that

$$v^T M \geq 0$$

1	1	1
a	b	c
b	a	a
c	c	b

$$\begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}^T \begin{pmatrix} 0 & 1 & 1 \\ -1 & 0 & 1 \\ -1 & -1 & 0 \end{pmatrix} = \begin{pmatrix} 0 & 1 & 1 \end{pmatrix} \geq 0$$

Maximal Lotteries

A lottery p is maximal if $p^T M \geq 0$:

- randomized Condorcet winner
- p is "at least as good" as any other lottery:

  the expected number of agents who prefer the alternative returned by p to that returned by q is at least as large as the expected number of agents who prefer the outcome returned by q to that returned by p

Maximal Lotteries

A lottery p is maximal if $p^T M \geq 0$:

1	1	1
a	b	c
b	c	a
c	a	b

$$\begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}^T \begin{pmatrix} 0 & 1 & -1 \\ -1 & 0 & 1 \\ 1 & -1 & 0 \end{pmatrix} = \begin{pmatrix} 0 & 1 & -1 \end{pmatrix} < 0$$

Maximal Lotteries

A lottery p is maximal if $p^T M \geq 0$:

1	1	1
a	b	c
b	c	a
c	a	b

$$\begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}^T \begin{pmatrix} 0 & 1 & -1 \\ -1 & 0 & 1 \\ 1 & -1 & 0 \end{pmatrix} = \begin{pmatrix} -1 & 0 & 1 \end{pmatrix} < 0$$

Maximal Lotteries

A lottery p is maximal if $p^T M \geq 0$:

1	1	1
a	b	c
b	c	a
c	a	b

$$\begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}^T \begin{pmatrix} 0 & 1 & -1 \\ -1 & 0 & 1 \\ 1 & -1 & 0 \end{pmatrix} = \begin{pmatrix} 1 & -1 & 0 \end{pmatrix} < 0$$

Maximal Lotteries

A lottery p is maximal if $p^T M \geq 0$:

1	1	1
a	b	c
b	c	a
c	a	b

$$\begin{pmatrix} \frac{1}{3} \\ \frac{1}{3} \\ \frac{1}{3} \end{pmatrix}^T \begin{pmatrix} 0 & 1 & -1 \\ -1 & 0 & 1 \\ 1 & -1 & 0 \end{pmatrix} = \begin{pmatrix} 0 & 0 & 0 \end{pmatrix} \geq 0$$

Maximal Lotteries

- always exist due to the Von Neumann Minimax Theorem.
- almost always unique
- set of profiles with multiple maximal lotteries has measure zero
- always unique for odd number of voters with strict preferences
- does not require asymmetry, completeness, or even transitivity of individual preferences

20	20	60
c	a	b
a	b	c
b	c	a

$$\begin{array}{c} \\ a \\ b \\ c \end{array}\begin{array}{c} \begin{array}{ccc} a & b & c \end{array} \\ \begin{pmatrix} 0 & -40 & -60 \\ 40 & 0 & 60 \\ 60 & -60 & 0 \end{pmatrix} \end{array}$$

$$\begin{tikzcd}
a & & b \arrow[ll, "40"'] \arrow[dl, "60"] \\
& c \arrow[ul, "60"] &
\end{tikzcd}$$

$$\begin{array}{c|ccc|c}
 & a & b & c & \\
a & 0,0 & -40,40 & -60,60 & -60 \\
b & 40,-40 & 0,0 & 60,-60 & 40 \\
c & 60,-60 & -60,60 & 0,0 & -60 \\
 & -60 & 40 & -60 &
\end{array}$$

Nash equilibrium: (b, b)

49	48	3
c	a	b
a	b	c
b	c	a

$$\begin{array}{c c} & \begin{array}{ccc} a & b & c \end{array} \\ \begin{array}{c} a \\ b \\ c \end{array} & \begin{pmatrix} 0 & 94 & -4 \\ -94 & 0 & 2 \\ 4 & -2 & 0 \end{pmatrix} \end{array}$$

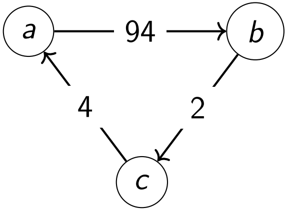

$$\begin{array}{c c} & \begin{array}{ccc} a & b & c \end{array} \\ \begin{array}{c} a \\ b \\ c \end{array} & \begin{pmatrix} 0,0 & 94,-94 & -4,4 \\ -94,94 & 0,0 & 2,-2 \\ 4,-4 & -2,2 & 0,0 \end{pmatrix} \end{array}$$

There is no *pure strategy Nash equilibrium*

There is a mixed Nash equilibrium: $2/100a + 4/100b + 94/100c$.

Maximal Lotteries can be efficiently computed via linear programming

https://voting.ml

Maximal Lottery Schemes

Every **maximal lottery scheme** is based on an odd and monotone function $\tau : \mathbb{Z} \to \mathbb{R}$ with $\tau(1) = 1$.

$$ML^\tau(\boldsymbol{R}) = \{p \in \Delta(X) \mid \sum_{x,y \in X} p(x)q(y)\tau(m_{xy}) \geq 0 \text{ for all } q \in \Delta(X)\}$$

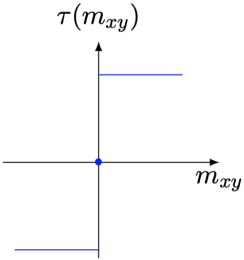

Figure 2: Different examples for τ (extended to $\mathbb{R}$): sign function, identity function, and cubic function.

$C1$-ML: ML schemes based on the sign function.

$C2$-ML: ML schemes based on the identity function.

$C1$-ML: ML schemes based on the sign function.

$C2$-ML: ML schemes based on the identity function.

49	48	3
c	a	b
a	b	c
b	c	a

$$\begin{array}{c c} & \begin{array}{ccc} a & b & c \end{array} \\ \begin{array}{c} a \\ b \\ c \end{array} & \begin{pmatrix} 0 & 94 & -4 \\ -94 & 0 & 2 \\ 4 & -2 & 0 \end{pmatrix} \end{array}$$

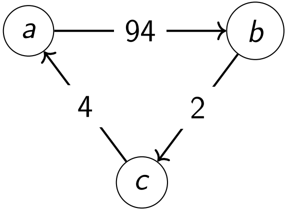

$$C1\text{-}ML\text{: } 1/3a + 1/3b + 1/3c$$

$$C2\text{-}ML\text{: } 2/100a + 4/100b + 94/100c$$

Theorem (Brandl, Brandt and Stricker) For any pair of ML schemes ML^{τ} and ML^{σ}, there is a preference profile $\boldsymbol{R}$ such that $ML^{\tau}(\boldsymbol{R}) = \{p\}$ and $ML^{\sigma}(\boldsymbol{R}) = \{q\}$ and $supp(p) \cap supp(q) = \varnothing$.

(cf. B. Dutta and J.-F. Laslier. Comparison functions and choice correspondences. Social Choice and Welfare, 16(4):513-532, 1999)

An SDS is **homogenous** when replacing every voter with a fixed number of identical clones (i.e., voters with the same preferences) does not change the outcome.

Theorem (Brandl, Brandt and Stricker) ML^{τ} is homogenous if and only if τ is based on τ' where there is a $t \geq 0$ such that $\tau'(k) = k^t$ for all $k \in \mathbb{N}$.

Efficiency

Pareto Efficiency: no voter can be made better off without making another voter worse off.

To define this, we need assumptions about how voters rank lotteries.

Suppose that $\succeq_i$ is voter i's weak preference relation.

$p \succeq_i^{DD'} q$ if and only if $x \succ_i y$ for all $x \in supp(p)$ and $y \in supp(q)$

$p \succeq_i^{DD} q$ if and only if $x \succeq_i y$ for all $x \in supp(p)$ and $y \in supp(q)$

$p \succeq_i^{ST} q$ if and only if
$(supp(p) \setminus supp(q)) \succ_i (supp(p) \cap supp(q)) \succ_i (supp(q) \setminus supp(p))$
and $p(x) = q(x)$ for all $x \in supp(p) \cap supp(q)$

Suppose that $\succeq_i$ is voter i's weak preference relation.

$p \succeq_i^{DD'} q$ if and only if $x \succ_i y$ for all $x \in supp(p)$ and $y \in supp(q)$

$p \succeq_i^{DD} q$ if and only if $x \succeq_i y$ for all $x \in supp(p)$ and $y \in supp(q)$

$p \succeq_i^{ST} q$ if and only if
$(supp(p) \setminus supp(q)) \succ_i (supp(p) \cap supp(q)) \succ_i (supp(q) \setminus supp(p))$
and $p(x) = q(x)$ for all $x \in supp(p) \cap supp(q)$

Suppose that $a \succ_i b \succ_i c$. Then,
$2/3a + 1/3b \succ_i^{DD'} c$
$2/3a + 1/3b \succ_i^{DD} 1/2b + 1/2c$
$1/2a + 1/2b \succ_i^{ST} 1/2b + 1/2c$

Bilinear Dominance

$p \succeq_i^{BD} q$ if and only if $p(x)q(y) \geq p(y)q(x)$ for all $x, y \in X$ with $x \succ_i y$

Bilinear Dominance

$p \succsim_i^{BD} q$ if and only if $p(x)q(y) \geq p(y)q(x)$ for all $x, y \in X$ with $x \succ_i y$

Suppose that $a \succ_i b \succ_i c$. Then,
$1/2a + 1/2b \succ_i^{BD} 1/3a + 1/3b + 1/3c$

(Fishburn 1984): p bilinearly dominates q iff p is preferable to q for every skew-symmetric bilinear (SSB) utility function consistent with $\succsim_i$

Stochastic Dominance

$p \succeq_i^{SD} q$ if and only if $\sum_{y \succeq_i x} p(y) \geq \sum_{y \succeq_i x} q(y)$ for all $x \in X$

Stochastic Dominance

$p \succeq_i^{SD} q$ if and only if $\sum_{y \succeq_i x} p(y) \geq \sum_{y \succeq_i x} q(y)$ for all $x \in X$

Suppose that $a \succ_i b \succ_i c$. Then,
$1/2a + 1/2c \succ_i^{SD} 1/2b + 1/2c$

p stochastically dominates q iff p is preferable to q for every von Neumann-Morgenstern utility function consistent with $\succeq_i$

$$\succeq^{DD'} \subseteq \succeq^{DD} \subseteq \succeq^{BD} \qquad \succeq^{DD'} \subseteq \succeq^{ST} \subseteq \succeq^{BD} \qquad \succeq^{BD} \subseteq \succeq^{SD} \subseteq \succeq^{PC}$$

A lottery p is *SD*-**efficient** for a preference profile $\boldsymbol{R}$ if there is no lottery $q \in \Delta(X)$ such that $q \succeq_i^{SD} p$ for all i and $q \succ_j^{SD} p$ for some j. (Similar definitions for other preferences over lotteries).

Theorem (Fishburn, 1984) Every ML^τ is *ex post efficient*: whenever there are alternatives x and y such that $x \succeq_i y$ for all i and $x \succ_j y$ for some j, then y should receive probability 0.

Theorem (Brandl, Brandt and Stricker). Every $C2$-ML schemes is SD-efficient. No other ML scheme is SD-efficient for all numbers of voters and candidates.

Theorem (Brandl, Brandt and Stricker). Suppose that m is the number of candidates and n is the number of voters. Every *majoritarian* and neutral $SPSC$ violates SD-efficiency for $m \geq 9$ (and $n = 5$, $n - 7$ or $n \geq 9$), even when preferences are strict.

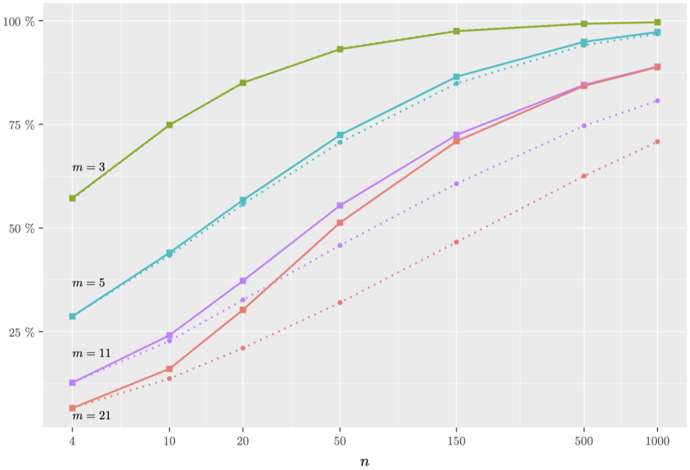

Figure 7: Frequencies of profiles R for which $|C2\text{-}ML(R)| = 1$ (solid lines) and $|C1\text{-}ML(R)| = 1$ (dotted lines), respectively. Each point is based on 100 000 samples according to the IAC model.

Figure 8: Distributions of the support sizes of *C1-ML* (left) and *C2-ML* (right) based on 100 000 samples according to the IAC model for every combination of parameters. The bars represent the frequency of a support with size $1, 3, 5,$ and 7 or more stacked from bottom to top. Frequencies lower than 4% are not labeled.

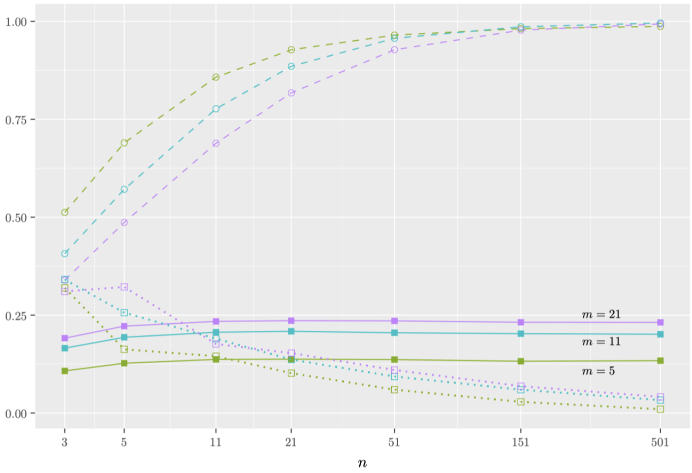

Figure 10: Average normalized Shannon entropies of *C2-ML* (solid line) compared to random dictatorship (dashed line) and the plurality rule (dotted line). Each point is based on 100 000 samples according to the IAC model.

F. Brandl and F. Brandt. *A Natural Adaptive Process for Collective Decision-Making*. manuscript, 2021.

Consider an urn filled with balls, each labeled with one of several possible collective decisions. Now, draw two balls from the urn, let a random voter pick her more preferred as the alternative, relabel the losing ball with the collective decision, put both balls back into the urn, and repeat. In order to prevent the permanent disappearance of some types of balls, once in a while, a randomly drawn ball is labeled with a random alternative.

Consider an urn filled with balls, each labeled with one of several possible collective decisions. Now, draw two balls from the urn, let a random voter pick her more preferred as the alternative, relabel the losing ball with the collective decision, put both balls back into the urn, and repeat. In order to prevent the permanent disappearance of some types of balls, once in a while, a randomly drawn ball is labeled with a random alternative.

Brandl and Brandt prove that this process will almost surely converge to the $C2$-ML.

300	300	300
*A*1	*A*1	*A*2
*A*2	*A*3	*A*3
*A*3	*A*2	*A*1

300	300	300
*A*1	*A*2	*A*3
*A*2	*A*3	*A*1
*A*3	*A*1	*A*2