Logics for Social Choice Theory

Eric Pacuit, University of Maryland

Lecture 3

ESSLLI 2022

Social Choice Theory

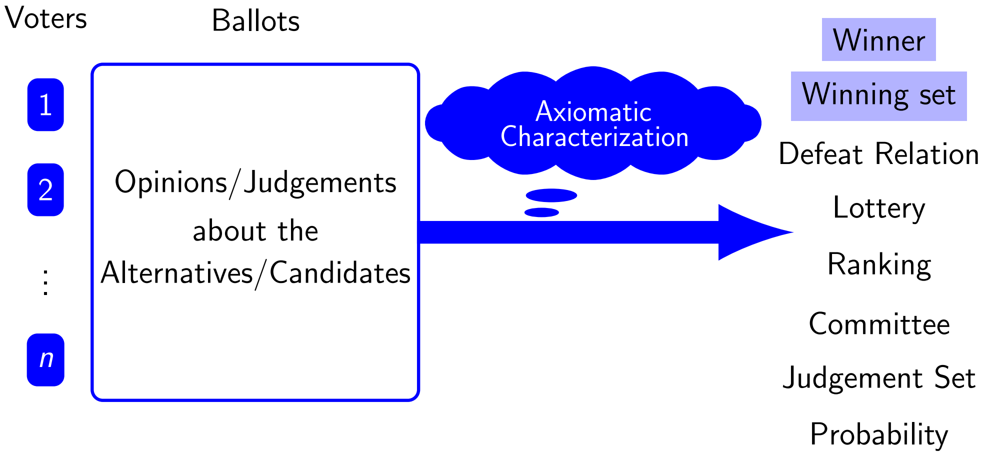

Anonymous profiles

40	35	25
t	r	k
k	k	r
r	t	t

Margin

Let P be a profile and $a, b \in X(\mathrm{P})$. Then the **margin of** a **over** b is:

$$Margin_{\mathrm{P}}(a, b) = |\{i \in V(\mathrm{P}) \mid a\mathrm{P}_i b\}| - |\{i \in V(\mathrm{P}) \mid b\mathrm{P}_i a\}|.$$

We say that a is **majority preferred** to b in P when $Margin_{\mathrm{P}}(a, b) > 0$.

Margin Graph

The **margin graph of** P, $\mathcal{M}(\mathrm{P})$, is the weighted directed graph whose set of nodes is $X(\mathrm{P})$ with an edge from a to b weighted by $Margin(a, b)$ when $Margin(a, b) > 0$. We write

$$a \xrightarrow{\alpha}_{\mathrm{P}} b \text{ if } \alpha = Margin_{\mathrm{P}}(a, b) > 0.$$

40	35	25
t	r	k
k	k	r
r	t	t

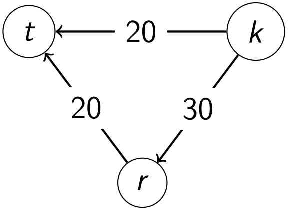

Margin Graph

A **margin graph** is a weighted directed graph $\mathcal{M}$ where all the weights have the same parity.

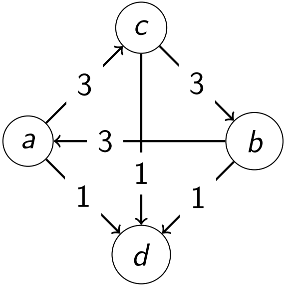

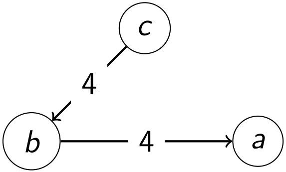

Theorem (Debord, 1987)

For any margin graph $\mathcal{M}$, there is a profile P *such that $\mathcal{M}$ is the margin graph of* P*.*

Social choice correspondence

A **voting method** is a function F on the domain of all profiles such that for any profile P, $\varnothing \neq F(\mathrm{P}) \subseteq X(\mathrm{P})$ (also called a **variable social choice correspondence** VSCC).

- A (V, X)-SCC is a social choice correspondence defined on (V, X)-profiles.
- A voting method F is **resolute** if for all P, $|F(\mathrm{P})| = 1$. Resolute SCCs are called **social choice functions**.

Social choice correspondence

A **voting method** is a function F on the domain of all profiles such that for any profile P, $\varnothing \neq F(\mathrm{P}) \subseteq X(\mathrm{P})$ (also called a **variable social choice correspondence** VSCC).

- A (V, X)-SCC is a social choice correspondence defined on (V, X)-profiles.
- A voting method F is **resolute** if for all P, $|F(\mathrm{P})| = 1$. Resolute SCCs are called **social choice functions**.

There are many examples of voting methods.

See `https://pref_voting.readthedocs.io` for a Python package that provides computational tools to study different voting methods.

Positional scoring rules

A *scoring vector* is a vector $\langle s_1, \ldots, s_n \rangle$ of numbers such that for each $m \in \{1, \ldots, n-1\}$, $s_m \geq s_{m+1}$.

Given a profile $\mathbf{P}$ with $|X(\mathbf{P})| = n$, $x \in X(\mathbf{P})$, a scoring vector $\vec{s}$ of length n, and $i \in V(\mathbf{P})$, define $\mathit{score}_{\vec{s}}(x, \mathbf{P}_i) = s_r$ where $r = \mathit{Rank}(x, \mathbf{P}_i)$.

Let $\mathit{score}_{\vec{s}}(x, \mathbf{P}) = \sum_{i \in V(\mathbf{P})} \mathit{score}_{\vec{s}}(x, \mathbf{P}_i)$. A voting method F is a **positional scoring rule** if there is a map $\mathcal{S}$ assigning to each natural number n a scoring vector of length n such that for any profile $\mathbf{P}$ with $|X(\mathbf{P})| = n$,

$$F(\mathbf{P}) = \mathrm{argmax}_{x \in X(\mathbf{P})} \mathit{score}_{\mathcal{S}(n)}(x, \mathbf{P}).$$

Examples

Borda:	$\mathcal{S}(n) = \langle n-1, n-2, \ldots, 1, 0\rangle$
Plurality:	$\mathcal{S}(n) = \langle 1, 0, \ldots, 0\rangle$
Anti-Plurality:	$\mathcal{S}(n) = \langle 1, 1, \ldots, 1, 0\rangle$

1	3	2	4
a	b	b	c
c	a	c	a
b	c	a	b

Borda winner	c
Plurality winner	b
Anti-Plurality winner	a

Iterative procedures: Instant Runoff

- If some alternative is ranked first by an absolute majority of voters, then it is declared the winner.
- Otherwise, the alternative ranked first be the fewest voters (the plurality loser) is eliminated.
- Votes for eliminated alternatives get transferred: delete the removed alternatives from the ballots and "shift" the rankings (e.g., if 1st place alternative is removed, then your 2nd place alternative becomes 1st).

Also known as Ranked-Choice, STV, Hare

How should you deal with ties? (e.g., multiple alternatives are plurality losers)

Iterative procedures

Variants:

- Plurality with runoff: remove all candidates except top two plurality score;
- Coombs: remove candidates with most last place votes;
- Baldwin: remove candidate with smallest Borda score;
- Nanson: remove candidates with below average Borda score

Example

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

Instant Runoff	$\{b\}$
Plurality with Runoff	$\{a, b\}$
Coombs	$\{d\}$
Baldwin	$\{a, b, d\}$
Strict Nanson	$\{a\}$

Condorcet criteria

The **Condorcet winner** in a profile P is a candidate $x \in X(\mathrm{P})$ that is the maximum of the majority ordering, i.e., for all $y \in X(\mathrm{P})$, if $x \neq y$, then $Margin_{\mathrm{P}}(x, y) > 0$.

The **Condorcet loser** in a profile P is a candidate $x \in X(\mathrm{P})$ that is the minimum of the majority ordering, i.e., for all $y \in X(\mathrm{P})$, if $x \neq y$, then $Margin_{\mathrm{P}}(y, x) < 0$.

A voting method F is **Condorcet consistent**, if for all P, if x is a Condorcet winner in P, then $F(\mathrm{P}) = \{x\}$.

A voting method F is susceptible to the **Condorcet loser paradox** (also known as *Borda's paradox*) if there is some P such that x is a Condorcet loser in P and $x \in F(\mathrm{P})$.

Condorcet paradox

n	n	n
a	b	c
b	c	a
c	a	b

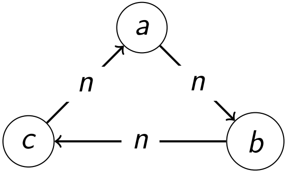

20	13	21	14	22	10
a	a	c	b	b	c
c	b	b	a	c	a
b	c	a	c	a	b

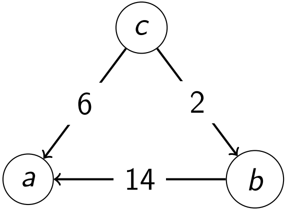

Condorcet winner: c

Instant Runoff winner: b

Plurality winner: b

Borda winner: b

Theorem (Smith 1973, Young 1974)

A voting method satisfies Anonymity, Neutrality and **Reinforcement** *if and only if F is a scoring rule.*

Saari's argument, Balinski and Laraki (2010, pg. 77); Zwicker (2016, Proposition 2.5): Multiple districts paradox, *f cancels properly*.

2	2	2	1	2
a	b	c	a	b
b	c	a	b	a
c	a	b	c	c

- no Condorcet winner in the left profile
- b is the Condorcet winner in the right profile
- a is the Condorcet winner in the combined profiles

Condorcet consistent voting methods

- Minimax
- Copeland
- Beat Path
- Ranked Pairs
- Split Cycle

Minimax: For a profile P, The Minimax winners in P are:

$$\text{argmin}_{x \in X(\mathsf{P})} \text{max}\{\mathit{Margin}_{\mathsf{P}}(y, x) \mid y \in X(\mathsf{P})\}$$

Copeland/**Llull**: For $\alpha \in [0, 1]$, the Copeland$_\alpha$ score of a in P is the number of $b \in X(\mathsf{P})$ such that $\mathit{Margin}_{\mathsf{P}}(a, b) > 0$ plus α times the number of $b \in X(\mathsf{P})$ such that $\mathit{Margin}_{\mathsf{P}}(a, b) = 0$. *Copeland*(P) (resp. *Llull*(P)) is the set of candidates with maximal Copeland$_{1/2}$ (resp. $\mathit{Copeland}_1$) score in P.

Schulze Beat Path

For $a, b \in X(\mathrm{P})$, a *path from a to b in* P is a sequence $\rho = x_1, \ldots, x_n$ of distinct candidates in $X(\mathrm{P})$ with $x_1 = a$ and $x_n = b$ such that for $1 \leq k \leq n-1$, $Margin_{\mathrm{P}}(x_k, x_{k+1}) > 0$.

The *strength of* ρ is $\min\{Margin_{\mathrm{P}}(x_k, x_{k+1}) \mid 1 \leq k \leq n-1\}$.

Then a defeats b in P according to Beat Path if the strength of the strongest path from a to b is greater than the strength of the strongest path from b to a.

$BP(\mathrm{P})$ is the set of undefeated candidates.

Tideman Ranked Pairs, I

For a profile P and $T \in \mathcal{L}\big(\{(x,y) \mid x \neq y \text{ and } \mathit{Margin}_{\mathrm{P}}(x,y) \geq 0\}\big)$, called the *tie-breaking ordering*

A pair (x, y) of candidates has a *higher priority* than a pair (x', y') of candidates according to T when either $\mathit{Margin}_{\mathrm{P}}(x,y) > \mathit{Margin}_{\mathrm{P}}(x',y')$ or $\mathit{Margin}_{\mathrm{P}}(x,y) = \mathit{Margin}_{\mathrm{P}}(x',y')$ and $(x,y)\ T\ (x',y')$.

Tideman Ranked Pairs, II

We construct a *Ranked Pairs ranking* $\succ_{\mathbf{P},T} \in \mathcal{L}(X)$ as follows:

1. Initialize $\succ_{\mathbf{P},T}$ to $\varnothing$.
2. If all pairs (x, y) with $x \neq y$ and $Margin_{\mathbf{P}}(x, y) \geq 0$ have been considered, then return $\succ_{\mathbf{P},T}$. Otherwise let (a, b) be the pair with the highest priority among those with $a \neq b$ and $Margin_{\mathbf{P}}(a, b) \geq 0$ that have not been considered so far.
3. If $\succ_{\mathbf{P},T} \cup \{(a, b)\}$ is acyclic, then add (a, b) to $\succ_{\mathbf{P},T}$; otherwise, add (b, a) to $\succ_{\mathbf{P},T}$. Go to step 2.

When the procedure terminates, $\succ_{\mathbf{P},T}$ is a linear order.

The set $RP(\mathbf{P})$ of Ranked Pairs winners is the set of all $x \in X(\mathbf{P})$ such that x is the maximum of $\succ_{\mathbf{P},T}$ for some tie-breaking ordering T.

Split Cycle

Split Cycle defeat: a candidate a defeats a candidate b just in case

- the majority margin of a over b is greater than 0, and
- for every majority cycle containing a and b, the margin of a over b is greater than the smallest margin between consecutive candidates in the cycle.

The Split Cycle winners are the undefeated candidates.

An intuitive way defeat relation is as follows:

1. In each majority cycle, identify the wins with the smallest margin in that cycle.
2. After completing step 1 for all cycles, discard the identified wins. All remaining wins count as defeats.

Example

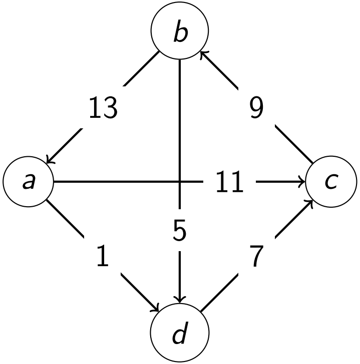

Minimax: $\{d\}$
Copeland: $\{a, b\}$
Beat Path: $\{d\}$
Ranked Pairs: $\{b\}$
Split Cycle: $\{b, d\}$

Key idea: Unequivocal increase in support for a candidate should not result in that candidate going from being a winner to being a loser.

1. *monotonicity*: if a candidate x is a winner given a preference profile P, and P′ is obtained from P by one voter moving x up in their ranking, then x should still be a winner given P′.
   (fixed-electorate axiom)

2. *positive involvement*: if a candidate x is a winner given P, and P^* is obtained from P by adding a new voter who ranks x in first place, then x should still be a winner given P^*.
   (variable-electorate axiom)

More-is-Less Paradox: Instant Runoff

6	5	4	2
a	c	b	b
b	a	c	a
c	b	a	c

6	5	4	2
a	c	b	a
b	a	c	b
c	b	a	c

More-is-Less Paradox: Instant Runoff

6	5	4	2
a	c	b	b
b	a	c	a
c	b	a	c

6	5	4	2
a	c	b	a
b	a	c	b
c	b	a	c

More-is-Less Paradox: Instant Runoff

6	5	4	2
a	c	b	b
b	a	c	a
c	b	a	c

Instant Runoff Winner: a

6	5	4	2
a	c	b	a
b	a	c	b
c	b	a	c

More-is-Less Paradox: Instant Runoff

6	5	4	2
a	*c*	*b*	*b*
b	*a*	*c*	*a*
c	*b*	*a*	*c*

Instant Runoff Winner: *a*

6	5	4	2
a	*c*	*b*	*a*
b	*a*	*c*	*b*
c	*b*	*a*	*c*

Instant Runoff Winner: *c*

More-is-Less Paradox: Instant Runoff

6	5	4	2
a	c	b	b
b	a	c	a
c	b	a	c

Instant Runoff Winner: a

6	5	4	2
a	c	b	a
b	a	c	b
c	b	a	c

Instant Runoff Winner: c

Any failure of monotonicity for a resolute voting rule F represents an opportunity for a voter to manipulate F in a particular way: via a simple drop or simple lift.

Manipulation

Suppose that F is a resolute voting rule
F is **manipulable** provided there are two profiles

$$\mathrm{P} = (P_1, \ldots, P_i, \ldots, P_n) \text{ and } \mathrm{P}' = (P'_1, \ldots, P'_i, \ldots, P'_n)$$

and a voter i such that

Manipulation

Suppose that F is a resolute voting rule
F is **manipulable** provided there are two profiles

$$\mathrm{P} = (P_1, \ldots, P_i, \ldots, P_n) \text{ and } \mathrm{P}' = (P'_1, \ldots, P'_i, \ldots, P'_n)$$

and a voter i such that

$P_j = P'_j$ for all $j \neq i$, and

Manipulation

Suppose that F is a resolute voting rule
F is **manipulable** provided there are two profiles

$$\mathrm{P} = (P_1, \ldots, P_i, \ldots, P_n) \text{ and } \mathrm{P}' = (P'_1, \ldots, P'_i, \ldots, P'_n)$$

and a voter i such that

$P_j = P'_j$ for all $j \neq i$, and

i strictly prefers the winner under P' to the winner under P:
aP_ib where $F(\mathrm{P}') = \{a\}$ and $F(\mathrm{P}) = \{b\}$.

Manipulation

Suppose that F is a resolute voting rule
F is **manipulable** provided there are two profiles

$$\mathrm{P} = (P_1, \ldots, P_i, \ldots, P_n) \text{ and } \mathrm{P}' = (P'_1, \ldots, P'_i, \ldots, P'_n)$$

and a voter i such that

$P_j = P'_j$ for all $j \neq i$, and

i strictly prefers the winner under P' to the winner under P:
aP_ib where $F(\mathrm{P}') = \{a\}$ and $F(\mathrm{P}) = \{b\}$.

Intuition: P_i is voter i's "true preference".

Strategizing

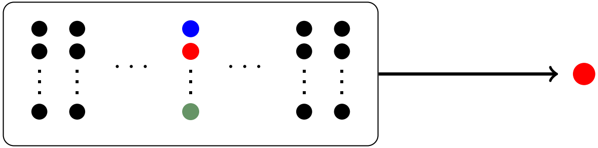

Strategizing

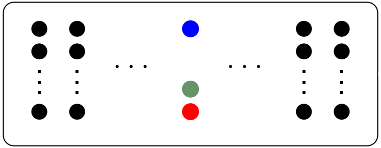

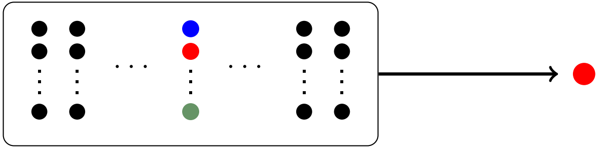

Strategizing

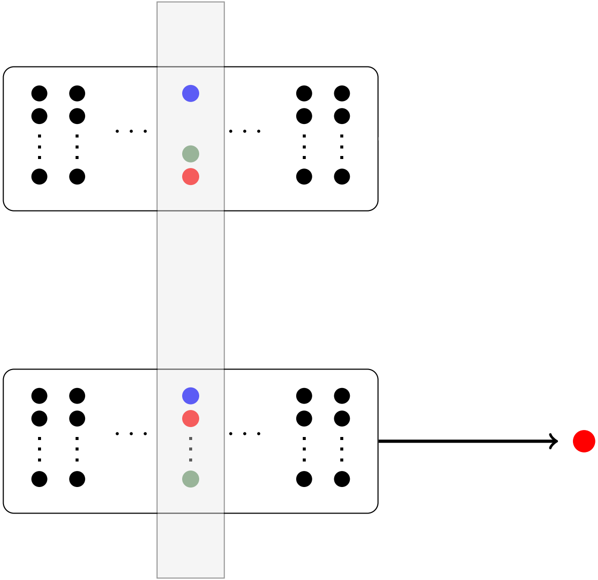

Strategizing

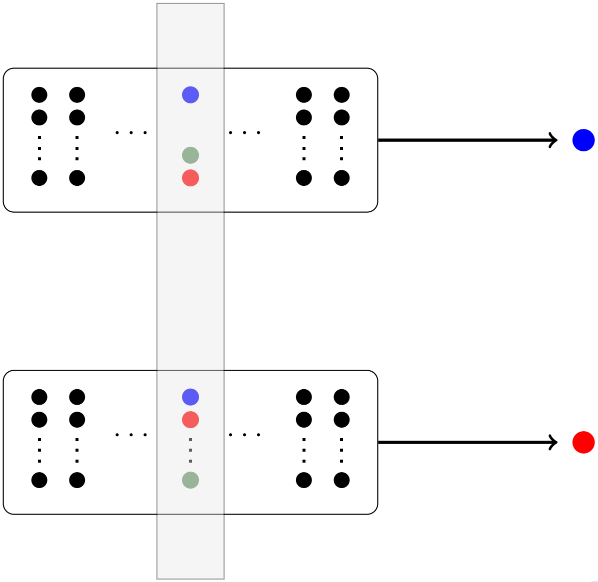

Strategizing

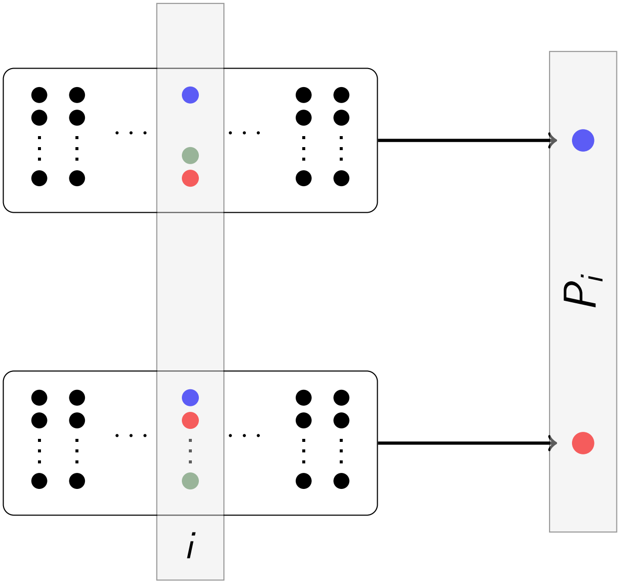

2	3	1	1
e	*d*	*a*	*a*
c	*e*	*b*	*b*
a	*b*	*c*	*c*
d	*c*	*d*	*d*
b	*a*	*e*	*e*

Copeland winning set: $\{e\}$

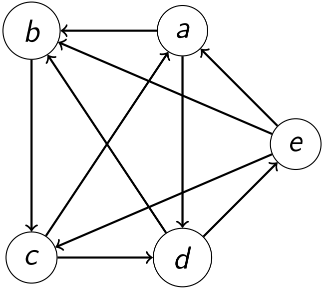

2	3	1	1
e	*d*	*a*	*e*
c	*e*	*b*	*d*
a	*b*	*c*	*c*
d	*c*	*d*	*b*
b	*a*	*e*	*a*

Copeland winning set: $\{d\}$

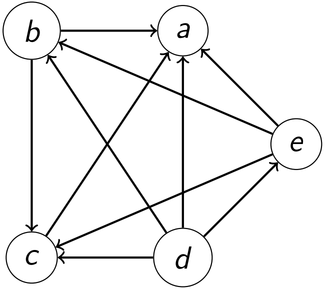

2	3	1	1
e	d	a	a
c	e	b	b
a	b	c	c
d	c	d	d
b	a	e	e

Borda winning set: $\{e\}$

Borda scores:
a: 12
b: 12
c: 13
d: 16
e: 17

2	3	1	1
e	d	a	d
c	e	b	a
a	b	c	b
d	c	d	c
b	a	e	e

Borda winning set: $\{d\}$

Borda scores:
a: 11
b: 11
c: 12
d: 19
e: 17

Monotonicity Properties

Strategyproofness if for all profiles P, if $\mathrm{P}' = \mathrm{P}[P_i/Q_i]$, then *not* $F(\mathrm{P}')P_iF(\mathrm{P})$

Monotonicity Properties

Strategyproofness if for all profiles P, if $\mathrm{P}' = \mathrm{P}[P_i/Q_i]$, then *not* $F(\mathrm{P}')P_iF(\mathrm{P})$

Maskin monotonicity if for all profiles P, if $\mathrm{P}' = \mathrm{P}[P_i/Q_i]$ and for all y, $F(\mathrm{P})P_iy$ implies $F(\mathrm{P}')Q_iy$, then $F(\mathrm{P}) = F(\mathrm{P}')$

The Gibbard-Satterthwaite Theorem

Gibbard-Satterthwaite Theorem. Consider a resolute voting rule F that is defined for some number m of alternatives with $m \geq 3$, with no restrictions on the preference domain. Then, this rule must be at least one of the following:

1. dictatorial: there exists a single fixed voter whose most-preferred alternative is chosen for every profile;
2. imposing: there is at least one alternative that does not win under any profile;
3. manipulable (i.e., not strategy-proof).

M. A. Satterthwaite. *Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions*. Journal of Economic Theory, 10(2):187-217, 1975.

A. Gibbard. *Manipulation of voting schemes: A general result*. Econometrica, 41(4):587-601, 1973.

Theorem 13.1 For $n = 3$ voters and $m > 3$ alternatives, no (resolute) voting rule satisfies both strategyproofness and the majority criterion.

Theorem 13.1 For $n = 3$ voters and $m > 3$ alternatives, no (resolute) voting rule satisfies both strategyproofness and the majority criterion.

Lemma 1. Let $m = 3$ and $n = 3$. There is no resolute voting rule F satisfying strategyproofness and the majority criterion

Lemma 2. Let $m \geq 3$ and $n = 3$. If F is a resolute voting rule satisfying strategyproofness and the majority criterion for $m + 1$ alternatives, then there exists a voting rule F' for m alternatives with the same properties.

Christian Geist and Dominik Peters. *Computer-aided Methods for Social Choice Theory*. Trends in Computational Social Choice, chapter 13, pages 249–267. AI Access, 2017.

Theorem (Muller-Satterthwaite) Assume that there are more than 3 candidates. Any resolute voting method satisfying surjectivity and Maskin monotonicity is dictatorial.

Key idea: Unequivocal increase in support for a candidate should not result in that candidate going from being a winner to being a loser.

1. *monotonicity*: if a candidate x is a winner given a preference profile P, and P′ is obtained from P by one voter moving x up in their ranking, then x should still be a winner given P′.
   (fixed-electorate axiom)

2. *positive involvement*: if a candidate x is a winner given P, and P^* is obtained from P by adding a new voter who ranks x in first place, then x should still be a winner given P^*.
   (variable-electorate axiom)

Violating Positive Involvement: Coombs

2	2	1	1	2	1	1
c	b	d	d	c	a	b
a	a	c	a	b	d	d
b	c	b	c	d	b	a
d	d	a	b	a	c	c

Coombs winner: $\{b\}$

(the order of elimination is d, c)

2	2	1	1	2	1	1	1
c	b	d	d	c	a	b	b
a	a	c	a	b	d	d	d
b	c	b	c	d	b	a	c
d	d	a	b	a	c	c	a

Coombs winner: $\{c\}$

(a and d are tied for the most last place votes)

Breaking Ties

There are many tiebreaking rules: non-anonymous, non-neutral, random

Parallel universe tiebreaking: x is a winner if x wins according to some tiebreaking rule.

S. Obraztsova, E. Elkind and N. Hazon. *Ties Matter: Complexity of Voting Manipulation Revisited*. Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence.

J. Wang, S. Sikdar, T. Shepherd, Z. Zhao, C. Jiang and L. Xia. *Practical Algorithms for Multi-Stage Voting Rules with Parallel Universes Tiebreaking*. Proceedings of AAAI, 2019.

Violating Positive Involvement: Coombs PUT

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

Coombs winner: $\{a, b\}$

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

Coombs winner: $\{b, d\}$

No Show Paradox

The term "No Show Paradox" was introduced by Fishburn and Brams for violations of what is now called *negative involvement*: Adding a new voter who ranks a candidate last should not result in the candidate going from being a loser to a winner.

P. Fishburn and S. Brams. *Paradoxes of Preferential Voting*. Mathematics Magazine, 56(4), pp. 207 - 214, 1983.

D. Saari. *Basic Geometry of Voting*. Springer, 1995.

No Show Paradox

Moulin changed the meaning of "No Show Paradox" to refer to violations of participation: A resolute voting method satisfies participation if adding a new voter who ranks x above y cannot result in a change from x being the unique winner to y being the unique winner.

H. Moulin. *Condorcet's Principle Implies the No Show Paradox*. Journal of Economic Theory 45(1), pp. 53 - 64, 1988.

No Show Paradox

Peréz concludes that the Strong No Show Paradox is a common flaw of many *Condorcet consistent* voting methods, which are methods that always pick a Condorcet winner—a candidate who is majority preferred to every other candidate—if one exists.

J. Pérez. *The Strong No Show Paradoxes are a common flaw in Condorcet voting correspondences*. Social Choice and Welfare 18(3), pp. 601 - 616, 2001.

Violating Positive Involvement: Copeland

2	1	1
e	c	a
c	b	d
b	a	b
a	d	e
d	e	c

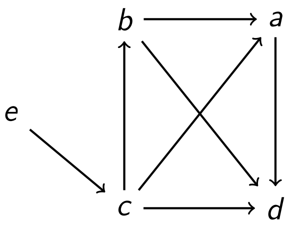

Copeland winners: $\{c\}$

2	1	1	
e	c	a	c
c	b	d	e
b	a	b	d
a	d	e	c
d	e	c	a

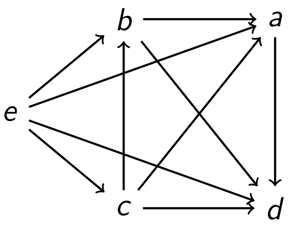

Copeland winners: $\{e\}$

Violating Positive Involvement: Beat Path

1	1	1	1	2	1	1	1	1	1
a	d	c	c	a	b	a	d	d	b
d	c	a	b	c	d	b	c	b	a
c	b	b	d	b	c	d	a	c	d
b	a	d	a	d	a	c	b	a	c

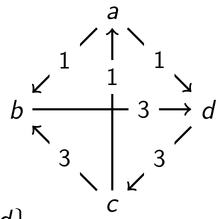

Beat Path winners: $\{a, b, c, d\}$

1	1	1	1	2	1	1	1	1	1	1
a	d	c	c	a	b	a	d	d	b	b
d	c	a	b	c	d	b	c	b	a	a
c	b	b	d	b	c	d	a	c	d	c
b	a	d	a	d	a	c	b	a	c	d

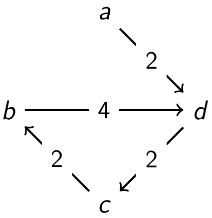

Beat Path winners: $\{a\}$

A logic for resolute social choice correspondences

G. Ciná and U. Endriss. *Proving classical theorems of social choice theory in modal logic*. Autonomous Agents and Multi-Agent Systems, 30, pp. 963 - 989, 2016.

N. Troquard, W. van der Hoek, and M. Wooldridge. *Reasoning about social choice functions*. Journal of Philosophical Logic 40(4), 473 - 498 (2011).

T. Agotnes, W. van der Hoek, and M. Wooldridge. *On the logic of preference and judgment aggregation*. Journal of Autonomous Agents and Multiagent Systems 22(1), 4 - 30 (2011).

Language

Atomic Propositions:

- $\mathit{Pref}[V, X] := \{p^i_{x \succeq y} \mid i \in V, x, y \in X\}$ is the set of preference atomic propositions, where $p^i_{x \preceq y}$ means i prefers y to x.
- Each $x \in X$ is an atomic proposition.

Modality:

- $\Diamond_C \varphi$: C can *ensure* the truth of φ.

Language

Atomic Propositions:

- $Pref[V, X] := \{p^i_{x \succeq y} \mid i \in V, x, y \in X\}$ is the set of preference atomic propositions, where $p^i_{x \preceq y}$ means i prefers y to x.
- Each $x \in X$ is an atomic proposition.

Modality:

- $\Diamond_C \varphi$: C can *ensure* the truth of φ.

$$p \mid \neg\varphi \mid \varphi \wedge \psi \mid \Diamond_C \varphi$$

Model

A **model** is a triple $M = \langle N, X, F \rangle$, consisting of a finite set of agents N (with $n = |N|$), a finite set of alternatives X, and a resolute SCC $F : \mathcal{L}(X)^V \to X$.

A **world** is a profile $(\mathrm{P}_1, \ldots, \mathrm{P}_n)$

Truth

Let $w = (\mathrm{P}_1, \ldots, \mathrm{P}_n)$

- $M, w \models p^i_{x \succeq y}$ iff $x\mathrm{P}_i y$
- $M, w \models x$ if and only if $F(\mathrm{P}_1, \ldots, \mathrm{P}_n) = x$
- $M, w \models \neg\varphi$ if and only if $M, w \not\models \varphi$
- $M, w \models \varphi \wedge \psi$ if and only if $M, w \models \varphi$ and $M, w \models \psi$
- $M, w \models \Diamond_C \varphi$ if and only if $M, w' \models \varphi$ for some $w' = (\mathrm{P}'_1, \ldots, \mathrm{P}'_n)$ with $\mathrm{P}_j = \mathrm{P}'_j$ for all $j \in N - C$.

(1) $p^i_{x \succeq x}$

(2) $p^i_{x \succeq y} \leftrightarrow \neg p^i_{y \succeq x}$ for $x \neq y$

(3) $p^i_{x \succeq y} \wedge p^i_{y \succeq y} \rightarrow p^i_{x \succeq z}$

(1) $p^i_{x \succeq x}$

(2) $p^i_{x \succeq y} \leftrightarrow \neg p^i_{y \succeq x}$ for $x \neq y$

(3) $p^i_{x \succeq y} \wedge p^i_{y \succeq y} \rightarrow p^i_{x \succeq z}$

$$ballot_i(w) \; = \; p^i_{x_1 \succeq x_2} \wedge \cdots \wedge p^i_{x_{m-1} \succeq x_m}$$

$$profile(w) \; = \; ballot_1(w) \wedge \cdots \wedge ballot_n(w)$$

(4) all propositional tautologies

(5) $\Box_i(\varphi \to \psi) \to (\Box_i\varphi \to \Box_i\psi)$ (K(i))

(6) $\Box_i\varphi \to \varphi$ (T(i))

(7) $\varphi \to \Box_i\Diamond_i\varphi$ (B(i))

(8) $\Diamond_i\Box_j\varphi \leftrightarrow \Box_j\Diamond_i\varphi$ (confluence)

(9) $\Box_{C_1}\Box_{C_2}\varphi \leftrightarrow \Box_{C_1\cup C_2}\varphi$ (union)

(10) $\Box_\emptyset\varphi \leftrightarrow \varphi$ (empty coalition)

(11) $(\Diamond_i p \wedge \Diamond_i\neg p) \to (\Box_j p \vee \Box_j\neg p)$, where $i \neq j$ (exclusiveness)

(12) $\Diamond_i ballot_i(w)$ (ballot)

(13) $\Diamond_{C_1}\delta_1 \wedge \Diamond_{C_2}\delta_2 \to \Diamond_{C_1\cup C_2}(\delta_1 \wedge \delta_2)$ (cooperation)

(14) $\bigvee_{x\in X}(x \wedge \bigwedge_{y\in X\setminus\{x\}}\neg y)$ (resoluteness)

(15) $(profile(w) \wedge \varphi) \to \Box_N(profile(w) \to \varphi)$ (functionality)

Theorem (Cinà and Endriss) The logic $L[V, X]$ is sound and complete w.r.t. the class of models of resolute social choice correpsondences.

Pareto

$$Par := \bigwedge_{x \in X} \bigwedge_{y \in X - \{x\}} \left[\left(\bigwedge_{i \in N} p^i_{x \succeq y} \right) \rightarrow \neg y \right]$$

IIA

$$IIA \;:=\; \bigwedge_{w \in \mathcal{L}(X)^n} \; \bigwedge_{x \in X} \; \bigwedge_{y \in X - \{x\}} [\Diamond_V (\mathit{profile}(w) \wedge x) \rightarrow (\mathit{profile}(w)(x,y) \rightarrow \neg y)]$$

- $N^{w}_{x \succeq y} = \bigwedge \{ p^{i}_{x \succeq y} \mid x \mathrm{P}_i y \text{ in } w \}$
- $\mathit{profile}(w)(x,y) := N^{w}_{x \succeq y} \wedge N^{w}_{y \succeq x}$

Dictatorship

$$Dic \; := \; \bigvee_{i \in N} \bigwedge_{x \in X} \bigwedge_{y \in X - \{x\}} (p^i_{x \succeq y} \rightarrow \neg y)$$

Arrow's Theorem

Theorem (Ciná and Endriss) Consider a logic $L[V, X]$ with a language parameterised by X such that $|X| > 3$. Then we have:

$$\vdash \mathit{Par} \wedge \mathit{IIA} \rightarrow \mathit{Dic}$$

Strong Monotonicity

$$SM := \bigwedge_{w \in \mathcal{L}(X)^n} \bigwedge_{x \in X} \left[\Diamond_V (profile(w) \wedge x) \wedge \left(\bigwedge_{y \in X \setminus \{x\}} N^w_{x \succeq y} \right) \rightarrow x \right]$$

Surjectivity

$$Sur := \bigwedge_{x \in X} \bigwedge_{w \in \mathcal{L}(X)^V} \Diamond_V(profile(w) \wedge x)$$

Theorem (Ciná and Endriss) Consider a logic $L[V, X]$ with a language parameterised by X such that $|X| \geq 3$. Then we have:

$$\vdash SM \wedge Sur \rightarrow Dic$$