Social Choice Theory and Machine Learning
Lecture 4

Eric Pacuit, University of Maryland

August 8, 2024

Plan for today

- ✓ A brief introduction to social choice theory
- ✓ A survey of voting methods
- ✓ Characterizing voting methods
- ✓ Splitting cycles and breaking ties
- ► Learning voting rules
- ► Strategic voting

Other approaches

Lirong Xia (2013). *Designing social choice mechanisms using machine learning*. In Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems, pp. 471 - 474.

Cem Anil and Xuchan Bao (2021). *Learning to Elect*. In Advances in Neural Information Processing Systems, pp. 8006 - 8017.

Dávid Burka, Clemens Puppe, László Szepesváry, and Attila Tasnádi (2022). *Voting: a machine learning approach*. European Journal of Operational Research, 299(3), pp. 1003 - 1017.

Inwon Kang, Qishen Han, and Lirong Xia (2023). *Learning to Explain Voting Rules*. In Proceedings of 6th AAAI/ACM Conference on AI, Ethics, and Society (AIES '23).

Other approaches

Lirong Xia (2013). *Designing social choice mechanisms using machine learning*. In Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems, pp. 471 - 474.

Cem Anil and Xuchan Bao (2021). *Learning to Elect*. In Advances in Neural Information Processing Systems, pp. 8006 - 8017.

Dávid Burka, Clemens Puppe, László Szepesváry, and Attila Tasnádi (2022). *Voting: a machine learning approach*. European Journal of Operational Research, 299(3), pp. 1003 - 1017.

Inwon Kang, Qishen Han, and Lirong Xia (2023). *Learning to Explain Voting Rules*. In Proceedings of 6th AAAI/ACM Conference on AI, Ethics, and Society (AIES '23).

Training Neural Networks

Which voting rule corresponds to the implicit selection mechanism employed by a trained neural network. By answering this question we hope to shed light on the 'conceptual' complexity (in a non-technical sense), or the salience of different voting rules.

Training Neural Networks

Which voting rule corresponds to the implicit selection mechanism employed by a trained neural network. By answering this question we hope to shed light on the ‘conceptual’ complexity (in a non-technical sense), or the salience of different voting rules.

Train a Multi-Layer Perceptron (MLP) to output the Condorcet winner, the Borda winners, and plurality winners, respectively, and statistically compare the chosen outcomes by the trained MLP.

Multi-Layer Perceptron

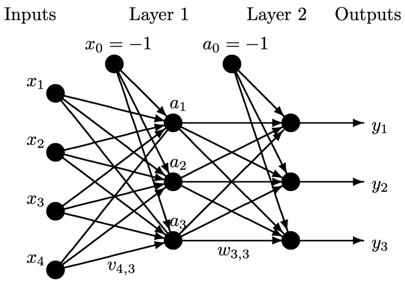

Figure 1: Structure of the MLP

It is well-known that MLPs are **universal function approximators** (Hecht-Nielsen 1987, Funahashi 1989), in particular an MLP will learn any voting rule on which it is trained with arbitrary accuracy provided that the size of the training sample is sufficiently large.

Training Neural Networks

What is the relation between the sample size and the accuracy by which the MLP learns different voting rules:

- Plurality
- Borda
- Copeland
- Kemeny-Young
- 2-Approval

Training Neural Networks

What is the relation between the sample size and the accuracy by which the MLP learns different voting rules:

- Plurality
- Borda
- Copeland
- Kemeny-Young
- 2-Approval

Conclusion: for limited, but still reasonably large sample sizes, the implicit voting rule employed by the MLP is most similar to the Borda rule and differs significantly from plurality rule; the Condorcet consistent methods such as Copeland and Kemeny-Young lie in between.

Training Neural Networks

- Randomly generated a set of profiles using the impartial culture (IC).
- 7, 9, or 11 voters and 3, 4, or 5 alternatives.
- Three scenarios:
  1. Trained on the subset of profiles with a (necessarily unique) Condorcet winner from the randomly generated profiles;
  2. Trained the Borda count as a set-valued function on the set of randomly generated profiles; and
  3. Trained the plurality rule as a set-valued function on the set of randomly generated profiles.
- For all scenarios, generate random training sets ranging from 100 to 3000 profiles.

Training Neural Networks

- Generated five random sample training seeds and took five random network seeds for the training procedure of the MLP.
- Selected one random testing seed pair for each training seed to generate test samples as well.
- For a given random training seed the five trained MLPs were each tested using the sample based on the respective testing seed pair. An alternative was selected as a winner on a test sample if it was selected by the majority of the five MLPs

Results

Consider the two-layered perceptron with a fixed sample size of 1000 profiles.

The table entries give the average percentages of those cases in which a trained MLP selects a winner of the method appearing in the respective column heading.

Condorcet Winners

Method	Cop	Kem-You	Borda	Plu	2-AV
q=3, n=7	96.12%	96.12%	94.36%	86.82%	80.94%
q=3, n=9	97.88%	97.88%	95.60%	88.28%	79.00%
q=3, n=11	97.58%	97.56%	94.50%	87.12%	77.80%
q=4, n=7	96.68%	93.58%	92.90%	80.80%	84.00%
q=4, n=9	92.54%	89.40%	92.30%	76.98%	82.62%
q=4, n=11	89.58%	86.38%	90.72%	75.76%	81.40%
q=5, n=7	84.30%	80.92%	87.12%	71.26%	78.86%
q=5, n=9	78.14%	74.48%	81.98%	66.22%	74.66%
q=5, n=11	77.02%	72.76%	80.16%	63.44%	72.78%

Table 1: Trained on Condorcet winners

Condorcet Winners

Method	Cop	Kem-You	Borda	Plu	2-AV
q=3, n=7	96.12%	96.12%	94.36%	86.82%	80.94%
q=3, n=9	97.88%	97.88%	95.60%	88.28%	79.00%
q=3, n=11	97.58%	97.56%	94.50%	87.12%	77.80%
q=4, n=7	96.68%	93.58%	92.90%	80.80%	84.00%
q=4, n=9	92.54%	89.40%	92.30%	76.98%	82.62%
q=4, n=11	89.58%	86.38%	90.72%	75.76%	81.40%
q=5, n=7	84.30%	80.92%	87.12%	71.26%	78.86%
q=5, n=9	78.14%	74.48%	81.98%	66.22%	74.66%
q=5, n=11	77.02%	72.76%	80.16%	63.44%	72.78%

Table 1: Trained on Condorcet winners

Condorcet Winners

Method	Cop	Kem-You	Borda	Plu	2-AV
q=3, n=7	96.12%	96.12%	94.36%	86.82%	80.94%
q=3, n=9	97.88%	97.88%	95.60%	88.28%	79.00%
q=3, n=11	97.58%	97.56%	94.50%	87.12%	77.80%
q=4, n=7	96.68%	93.58%	92.90%	80.80%	84.00%
q=4, n=9	92.54%	89.40%	92.30%	76.98%	82.62%
q=4, n=11	89.58%	86.38%	90.72%	75.76%	81.40%
q=5, n=7	84.30%	80.92%	87.12%	71.26%	78.86%
q=5, n=9	78.14%	74.48%	81.98%	66.22%	74.66%
q=5, n=11	77.02%	72.76%	80.16%	63.44%	72.78%

Table 1: Trained on Condorcet winners

Borda Winners

Method	Cop	Kem-You	Borda	Plu	2-AV
$q = 3, n = 7$	93.60%	93.60%	96.50%	86.08%	83.24%
$q = 3, n = 9$	93.36%	93.34%	98.20%	87.08%	81.82%
$q = 3, n = 11$	92.40%	92.32%	97.86%	86.48%	80.84%
$q = 4, n = 7$	88.84%	86.14%	95.38%	78.82%	85.62%
$q = 4, n = 9$	87.40%	84.30%	94.08%	75.86%	83.94%
$q = 4, n = 11$	87.32%	84.16%	92.78%	75.34%	82.86%
$q = 5, n = 7$	82.00%	78.70%	88.26%	69.88%	78.60%
$q = 5, n = 9$	78.42%	74.94%	84.20%	65.36%	76.26%
$q = 5, n = 11$	76.56%	72.60%	80.96%	62.88%	73.10%

Table 2: Trained on Borda winners

Borda Winners

Method	Cop	Kem-You	Borda	Plu	2-AV
$q = 3, n = 7$	93.60%	93.60%	96.50%	86.08%	83.24%
$q = 3, n = 9$	93.36%	93.34%	98.20%	87.08%	81.82%
$q = 3, n = 11$	92.40%	92.32%	97.86%	86.48%	80.84%
$q = 4, n = 7$	88.84%	86.14%	95.38%	78.82%	85.62%
$q = 4, n = 9$	87.40%	84.30%	94.08%	75.86%	83.94%
$q = 4, n = 11$	87.32%	84.16%	92.78%	75.34%	82.86%
$q = 5, n = 7$	82.00%	78.70%	88.26%	69.88%	78.60%
$q = 5, n = 9$	78.42%	74.94%	84.20%	65.36%	76.26%
$q = 5, n = 11$	76.56%	72.60%	80.96%	62.88%	73.10%

Table 2: Trained on Borda winners

Plurality Winners

Method	Cop	Kem-You	Borda	Plu	2-AV
$q = 3, n = 7$	94.44%	94.42%	95.64%	86.24%	82.54%
$q = 3, n = 9$	92.00%	91.82%	94.46%	86.16%	80.14%
$q = 3, n = 11$	89.76%	89.52%	92.62%	84.86%	78.48%
$q = 4, n = 7$	80.52%	78.10%	83.74%	73.16%	80.38%
$q = 4, n = 9$	76.48%	73.82%	79.56%	69.72%	75.54%
$q = 4, n = 11$	77.42%	74.32%	80.52%	69.58%	75.46%
$q = 5, n = 7$	66.62%	63.68%	70.12%	59.90%	67.04%
$q = 5, n = 9$	65.56%	62.72%	67.32%	57.02%	64.34%
$q = 5, n = 11$	60.78%	57.10%	61.74%	53.02%	58.80%

Table 3: Trained on plurality winners

Plurality Winners

Method	Cop	Kem-You	Borda	Plu	2-AV
$q = 3, n = 7$	94.44%	94.42%	95.64%	86.24%	82.54%
$q = 3, n = 9$	92.00%	91.82%	94.46%	86.16%	80.14%
$q = 3, n = 11$	89.76%	89.52%	92.62%	84.86%	78.48%
$q = 4, n = 7$	80.52%	78.10%	83.74%	73.16%	80.38%
$q = 4, n = 9$	76.48%	73.82%	79.56%	69.72%	75.54%
$q = 4, n = 11$	77.42%	74.32%	80.52%	69.58%	75.46%
$q = 5, n = 7$	66.62%	63.68%	70.12%	59.90%	67.04%
$q = 5, n = 9$	65.56%	62.72%	67.32%	57.02%	64.34%
$q = 5, n = 11$	60.78%	57.10%	61.74%	53.02%	58.80%

Table 3: Trained on plurality winners

Sample Size: Condorcet Winners

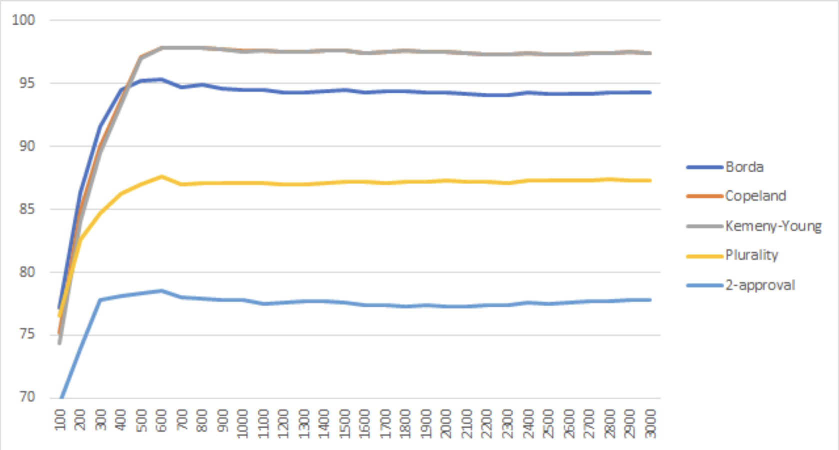

Figure 2: 2LP trained on Condorcet winners for $q = 3$

Sample Size: Condorcet Winners

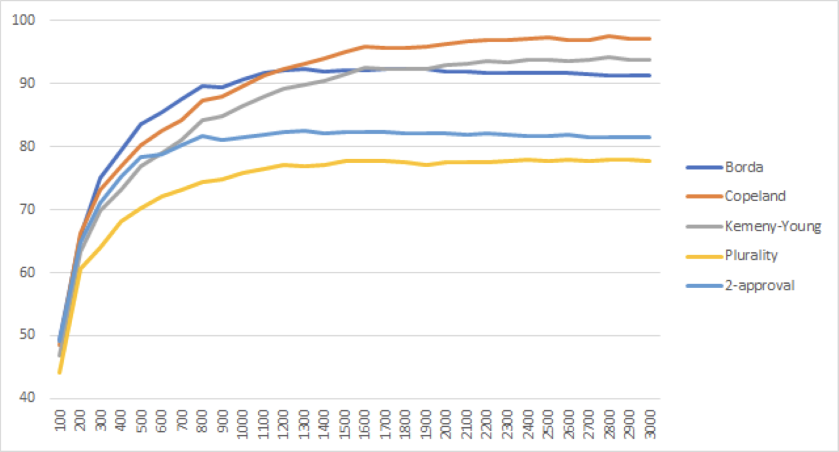

Figure 3: 2LP trained on Condorcet winners for $q = 4$

Sample Size: Condorcet Winners

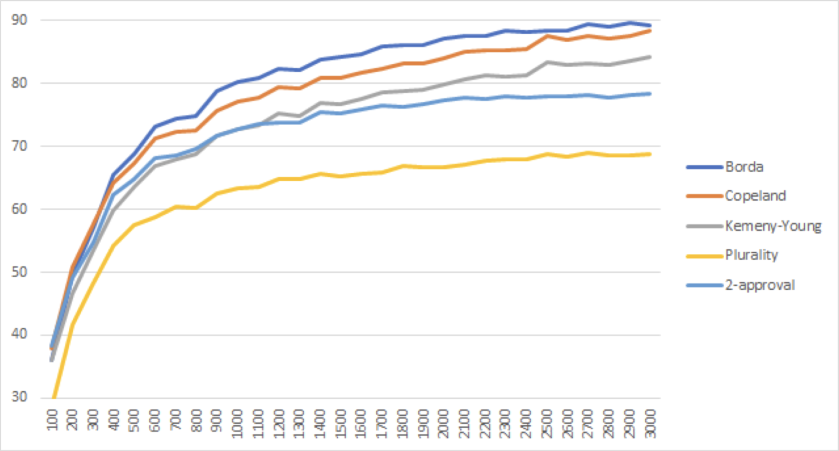

Figure 4: 2LP trained on Condorcet winners for $q = 5$

Sample Size: Plurality Winners

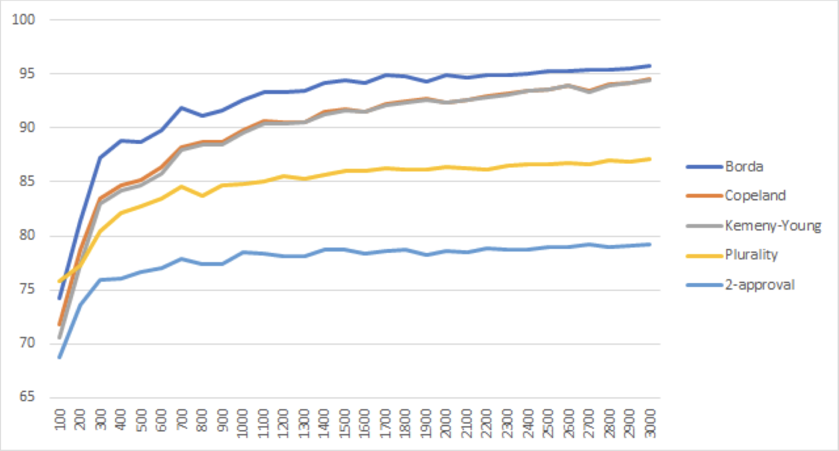

Figure 14: 2LP trained on the plurality rule for $q = 3$

Sample Size: Plurality Winners

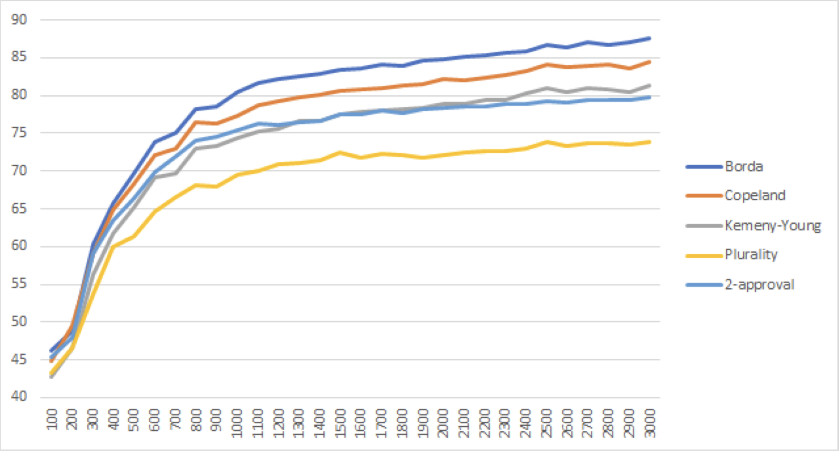

Figure 15: 2LP trained on the plurality rule for $q = 4$

Sample Size: Plurality Winners

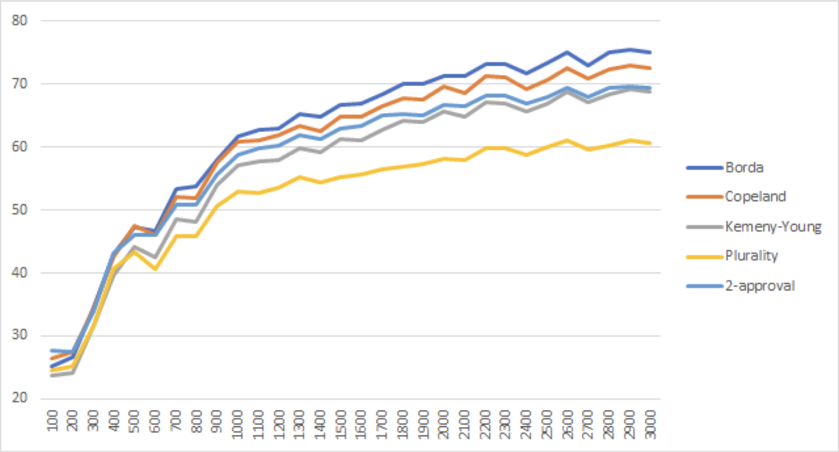

Figure 16: 2LP trained on the plurality rule for $q = 5$

Conclusions

- Theoretical Properties of MLPs:
  - Every voting rule can be learned by the trained network.
  - When trained on a rule different from the Borda count, the MLP's choices will converge to that rule.
  - The required training sample size can be very large for some rules.

Conclusions

- Theoretical Properties of MLPs:
  - Every voting rule can be learned by the trained network.
  - When trained on a rule different from the Borda count, the MLP's choices will converge to that rule.
  - The required training sample size can be very large for some rules.
- Borda Count:
  - Among popular voting rules, the Borda count stands out from a machine learning perspective.
  - The Borda count is the most salient voting method for limited sample sizes.
  - It best describes the behavior of a trained neural network in a voting environment.

Conclusions

- Complexity of Voting Rules:
  - Learning by neural networks can be seen as a way to select an appropriate degree of internal complexity.
  - Plurality rule and 2-approval are too simple.
  - The two investigated Condorcet consistent methods are too sophisticated.

Conclusions

- Complexity of Voting Rules:
  - Learning by neural networks can be seen as a way to select an appropriate degree of internal complexity.
  - Plurality rule and 2-approval are too simple.
  - The two investigated Condorcet consistent methods are too sophisticated.
- Human Intuition and Voting Methods:
  - It is conjectured that untrained humans might prefer the Borda count over other voting methods.
  - This hypothesis requires validation through carefully designed experiments with human subjects.

Strategic Voting

AI and Voting

ANALYSIS

How AI Puts Elections at Risk — And the Needed Safeguards

Widely accessible artificial intelligence tools could fuel the rampant spread of disinformation and create other hazards to democracy.

Mekela Panditharatne

Noah Giansiracusa

Chris Burnett

AI and Voting

JUNE 7, 2023 | 6 MIN READ

How AI Could Take Over Elections--And Undermine Democracy

An AI-driven political campaign could be all things to all people

BY ARCHON FUNG, THE CONVERSATION US & LAWRENCE LESSIG

AI and Voting

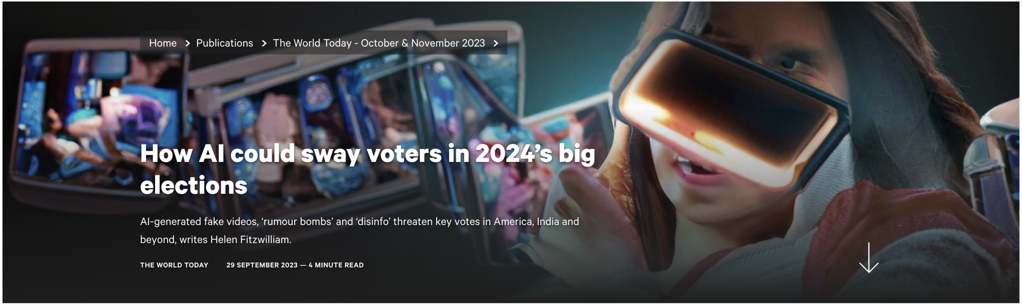

Home > Publications > The World Today - October & November 2023 >

How AI could sway voters in 2024's big elections

AI-generated fake videos, 'rumour bombs' and 'disinfo' threaten key votes in America, India and beyond, writes Helen Fitzwilliam.

THE WORLD TODAY 29 SEPTEMBER 2023 — 4 MINUTE READ

Helen Fitzwilliam

Journalist and filmmaker

Image — Image created with Adobe's Firefly generative AI.

AI and Voting

AI's War on Manipulation: Are We Winning?

Piotr Faliszewski and Ariel D. Procaccia

P. Faliszewski and A. Procaccia. *AI's War on Manipulation: Are We Winning?*. AI Magazine, 31:4, 2010.

Strategic Voting

While a voting method may satisfy a list of axioms that looks ideal, what happens if voters misrepresent their preferences in order to “game the system”?

Strategic Voting

While a voting method may satisfy a list of axioms that looks ideal, what happens if voters misrepresent their preferences in order to "game the system"?

It has long been noted that a voter can achieve a preferred election outcome by misrepresenting his or her actual preferences.

Strategic Voting

While a voting method may satisfy a list of axioms that looks ideal, what happens if voters misrepresent their preferences in order to "game the system"?

It has long been noted that a voter can achieve a preferred election outcome by misrepresenting his or her actual preferences.

C.L. Dodgson refers to a voters tendency to

> *"adopt a principle of voting which makes it a game of skill than a real test of the wishes of the elector."*

Strategic Voting

While a voting method may satisfy a list of axioms that looks ideal, what happens if voters misrepresent their preferences in order to "game the system"?

It has long been noted that a voter can achieve a preferred election outcome by misrepresenting his or her actual preferences.

C.L. Dodgson refers to a voters tendency to

> *"adopt a principle of voting which makes it a game of skill than a real test of the wishes of the elector."*

and that in his opinion

> *"it would be better for elections to be decided according to the wishes of the majority than of those who happen to be more skilled at the game."*

Strategic voting

By well-known theorems (the Gibbard-Sattherthwaite Theorem and its generalizations), any reasonable voting method sometimes gives voters an incentive to vote strategically.

Strategic voting

By well-known theorems (the Gibbard-Sattherthwaite Theorem and its generalizations), any reasonable voting method sometimes gives voters an incentive to vote strategically.

By reporting an insincere preference, a voter may obtain an election result that they prefer, according to their sincere preference, to the result they would obtain if there were to report their sincere preference.

Strategic Voting

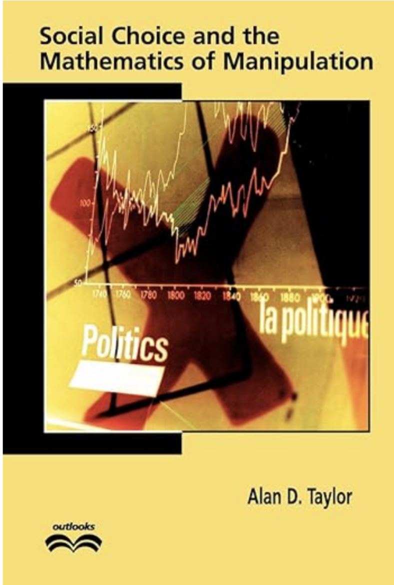

Strategic Voting: Example

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

Borda Winner: *c*

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

Borda Winner: *b*

Manipulation

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

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

and a voter i such that

Manipulation

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

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

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

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

and a voter i such that

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

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

Manipulation

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

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

and a voter i such that

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

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

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

Strategizing

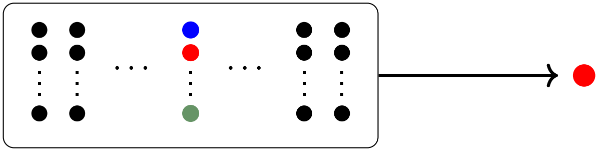

Strategizing

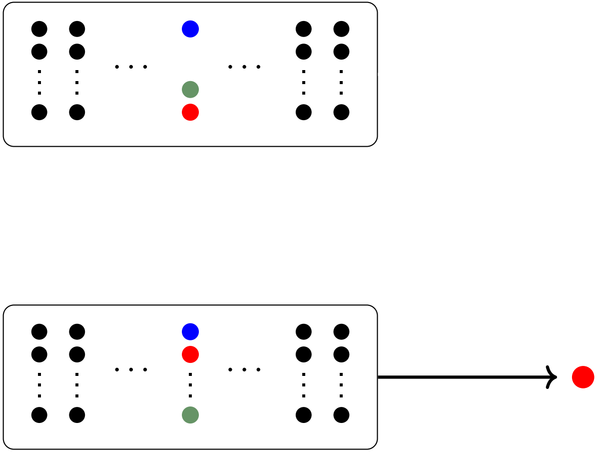

Strategizing

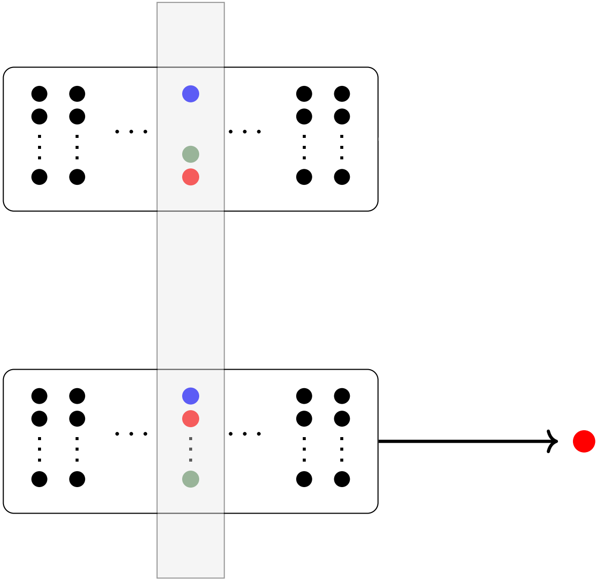

Strategizing

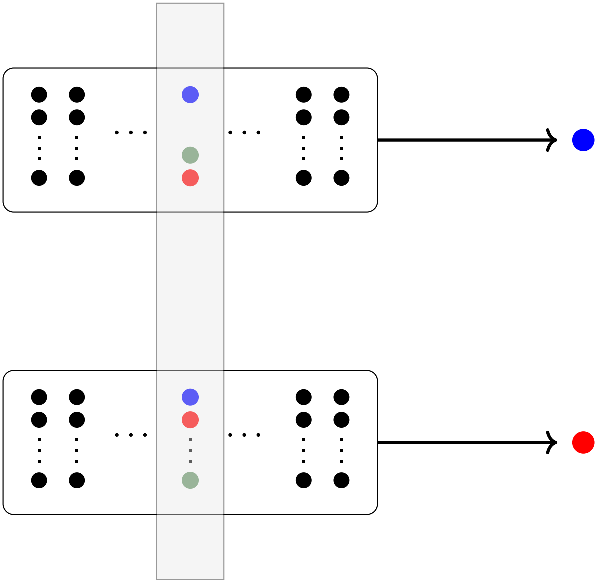

Strategizing

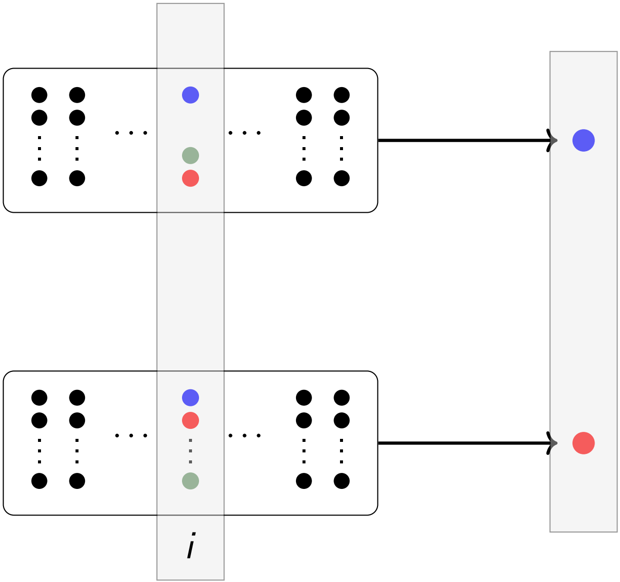

Strategizing

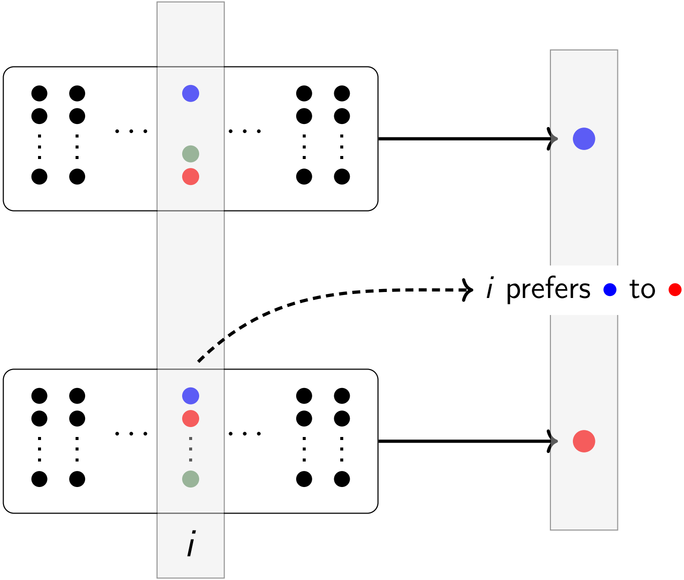

2	3	1	1
e	d	a	a
c	e	b	b
a	b	c	c
d	c	d	d
b	a	e	e

Copeland winning set: $\{e\}$

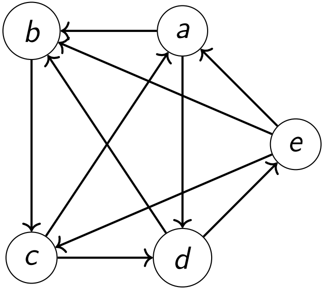

2	3	1	1
e	d	a	e
c	e	b	d
a	b	c	c
d	c	d	b
b	a	e	a

Copeland winning set: $\{d\}$

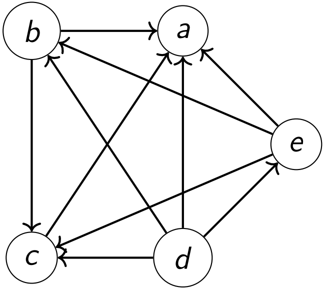

2	3	1	1
e	d	a	a
c	e	b	b
a	b	c	c
d	c	d	d
b	a	e	e

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

Strategic Voting is Unavoidable

Theorem (Gibbard-Satterthwaite Theorem)

Suppose that F is a resolute voting method for 3 or more candidates (and no restrictions on the domain). Then, F 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.

"Voter i **strictly prefers** $F(\mathbf{P}')$ **to** $F(\mathbf{P})$."

"Voter i **strictly prefers** $F(\mathbf{P}')$ **to** $F(\mathbf{P})$."

If $F(\mathbf{P})$ and $F(\mathbf{P}')$ are singletons, then "i **prefers** $F(\mathbf{P}')$ **to** $F(\mathbf{P})$" means $F(\mathbf{P}')\ P_i\ F(\mathbf{P})$

"Voter i **strictly prefers** $F(\mathbf{P}')$ **to** $F(\mathbf{P})$."

If $F(\mathbf{P})$ and $F(\mathbf{P}')$ are singletons, then "i **prefers** $F(\mathbf{P}')$ **to** $F(\mathbf{P})$" means $F(\mathbf{P}')\ P_i\ F(\mathbf{P})$

What happens if $F(\mathbf{P})$ and $F(\mathbf{P}')$ are not singletons?

Suppose that $F(\mathbf{P}) = Y$ and $F(\mathbf{P}') = Z$ are not singletons

- Z **weakly dominates** Y for i provided

$$\text{for all } z \in Z \text{ and } y \in Y \quad z \text{ is weakly preferred to } y \text{ by } i \quad \text{and}$$

$$\text{there exists } z \in Z \text{ and } y \in Y \text{ such that } z \; P_i \; y$$

Suppose that $F(\mathbf{P}) = Y$ and $F(\mathbf{P}') = Z$ are not singletons

- Z **weakly dominates** Y for i provided

  for all $z \in Z$ and $y \in Y$ z is weakly preferred to y by i and

  there exists $z \in Z$ and $y \in Y$ such that $z \, P_i \, y$

- Z is preferred by an **optimist** to Y: $\max_i(Z) \; P_i \; \max_i(Y)$

Suppose that $F(\mathbf{P}) = Y$ and $F(\mathbf{P}') = Z$ are not singletons

- Z **weakly dominates** Y for i provided

$$\text{for all } z \in Z \text{ and } y \in Y \quad z \text{ is weakly preferred to } y \text{ by } i \quad \text{and}$$

$$\text{there exists } z \in Z \text{ and } y \in Y \text{ such that } z \; P_i \; y$$

- Z is preferred by an **optimist** to Y: $\max_i(Z) \; P_i \; \max_i(Y)$

- Z is preferred by a **pessimist** to Y: $\min_i(Z) \; P_i \; \min_i(Y)$

Fact. Plurality rules is weak dominance manipulable, but is never single-winner manipulable.

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

Plurality Winner: c

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

Plurality Winners: $\{b, c\}$

Fact. Condorcet rule is manipulable by optimists (and also by pessimists), but is never weak dominance manipulable.

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

Condorcet Winner: c

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

Condorcet Winners: $\{a, b, c\}$

The Duggan-Schwartz Theorem

i is a **nominator** if for all profiles $\mathbf{P}$, then $Top(P_i) \in F(\mathbf{P})$.

Non-Imposed: For all $a \in X$ there exists a profile $\mathbf{P}$ such that $F(\mathbf{P}) = \{a\}$.

Manipulated by Optimist/Pessimist

F can be manipulated by an **optimist** if there is a profile $\mathbf{P}$ an ordering $Q_i \in O(X)$ such that $Q_i \neq P_i$ and

$$\exists x \in F(\mathbf{P}[P_i/Q_i]), \forall y \in F(\mathbf{P}), xP_iy$$

F can be manipulated by a **pessimist** if there is a profile $\mathbf{P}$ an ordering $Q_i \in O(X)$ such that $Q_i \neq P_i$ and

$$\forall x \in F(\mathbf{P}[P_i/Q_i]), \exists y \in F(\mathbf{P}), xP_iy$$

The Duggan-Schwartz Theorem

Theorem (Duggan-Schwartz Theorem)

Suppose that there are at least three candidates. Any voting method that is non-imposed and cannot be manipulated by an optimist or a pessimist has a nominator.

J. Duggan and T. Schwartz. *Strategic manipulability without resoluteness or shared beliefs: Gibbard-Satterthwaite generalized*. Social Choice and Welfare, 17, pp. 85 - 93, 2000.

Fishburn set extension

Suppose i is a voter with a preference ordering R_i that expects the ties in the voting rule to be broken according to some linear tie-breaking order; however, i does not know which order will be used:

For sets of candidates X and Y, we have $X \; {P_i}^F \; Y$ provided that

1. $x \; P_i \; y$ for all $x \in X \setminus Y$ and $y \in X \cap Y$
2. $y \; P_i \; z$ for all $y \in X \cap Y$ and $z \in Y \setminus X$
3. $x \; P_i \; z$ for all $x \in X \setminus Y$ and $z \in Y \setminus X$

F is a C1 voting method provided that for all profiles $\mathbf{P}$ and $\mathbf{P}'$, if $M(\mathbf{P}) = M(\mathbf{P}')$, then $F(\mathbf{P}) = F(\mathbf{P}')$

Theorem (Brandt and Geist, 2016)

There is no C1 voting method that satisfying Neutrality, Pareto and Fishburn-strategyproofness for $m \geq 5$ candidates and $n \geq 7$ voters.

F. Brandt and C. Geist (2016). *Finding strategyproof social choice functions via SAT solving*. Journal of Artificial Intelligence Research, 55, pp. 565 - 602.

candidates	unlabeled weak tournaments	6-voter ANECs	6-voter anonymous profiles	6-voter profiles
3	7	83	462	46,656
4	42	19,941	475,020	191,102,976
5	582	39,096,565	4,690,625,500	2,985,984,000,000

Strategic Voting under Partial Information

The classic approach to strategic voting assumes that voters know how the other voters will vote (and assumes that the other voters are not manipulating).

Strategic Voting under Partial Information

The classic approach to strategic voting assumes that voters know how the other voters will vote (and assumes that the other voters are not manipulating).

Can a voter can successfully manipulate in an election under limited information about how other voters will vote?

Strategic Voting under Partial Information

The classic approach to strategic voting assumes that voters know how the other voters will vote (and assumes that the other voters are not manipulating).

Can a voter can successfully manipulate in an election under limited information about how other voters will vote?

The information of the manipulator is typically represented by (probability over) a set of preference profiles, all of which agree on (i) the manipulator's own ranking and (ii) some other *partial information* (e.g., poll information such as the current winners, or the number of first-place votes for each candidate).

Strategic Voting under Partial Information

V. Conitzer, T. Walsh, and L. Xia (2011). *Dominating Manipulations in Voting with Partial Information*. In Proceedings of the 25th National Conference on Artificial Intelligence (AAAI-11).

S. Chopra, E. Pacuit, and R. Parikh (2004). *Knowledge-Theoretic Properties of Strategic Voting*. In Logics in Artificial Intelligence. JELIA 2004.

A. Reijngoud and U. Endriss (2012). *Voter Response to Iterated Poll Information*. In Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012).

R. Myerson and R. Weber (1993). *A theory of voting equilibria*. The American Political Science Review.

Learning to Manipulate under Limited Information

We use machine learning to gauge how resistant a preferential voting method is to manipulation under limited information about how other voters will vote.

Wesley Holliday, Alexander Kristoffersen, Eric Pacuit. *Learning to Manipulate under Limited Information*. arxiv.org/abs/2401.16412, 1st Workshop on Social Choice and Learning Algorithms (SCaLA 2024).

How to manipulate

v_1	v_2	v_3	v_4	v_5
c	c	c	d	d
b	b	b	b	b
a	d	d	a	a
d	a	a	c	c

↓

$$Borda(\mathbf{P}) = \{b\}$$

Winners

How to manipulate

Rankings

$R \in \{a\ b\ c\ d,\ \ldots,\ d\ c\ b\ a\}$

v_1	v_2	v_3	v_4	v_5
c	c	c	d	d
b	b	b	b	b
a	d	d	a	a
d	a	a	c	c

$Borda(\mathbf{P}) = \{b\}$

Winners

v_1	v_2	v_3	v_4	v_5
c	c	c	d	d
a	b	b	b	b
d	d	d	a	a
b	a	a	c	c

$Borda(R, \mathbf{P}_{-v_1}) = \{c, d\}$

Winners

How to manipulate

Rankings

$R \in \{a\ b\ c\ d,\ \ldots,\ d\ c\ b\ a\}$

v_1	v_2	v_3	v_4	v_5
c	c	c		
b	b	b		
a	d	d		
d	a	a	c	c

$Borda(\mathbf{P}) = \{b\}$

Winners

Which ranking R *should* v_1 submit?

v_1	v_2	v_3	v_4	v_5
		c	d	d
		b	b	b
		d	a	a
b	a	a	c	c

$Borda(R, \mathbf{P}_{-v_1}) = \{c, d\}$

Winners

Profitable manipulations

Given a profile of utilities for each voters, we can define the profile of rankings submitted by each voter, where alternative a is ranked above alternative b when the utility of a is greater than the utility of b:

Voters	a	b	c	d
v_1	0.1	0.65	0.9	0.08
v_2	0.7	0.9	1.0	0.8
v_3	0.01	0.03	0.5	0.02
v_4	0.1	0.5	0	0.9
v_5	0.1	0.2	0.05	1.0

U

v_1	v_2	v_3	v_4	v_5
c	c	c	d	d
b	b	b	b	b
a	d	d	a	a
d	a	a	c	c

P

Profitable manipulations

Given a profile of utilities for each voters, we can define the profile of rankings submitted by each voter, where alternative a is ranked above alternative b when the utility of a is greater than the utility of b:

Voters	a	b	c	d
v_1	0.1	0.65	0.9	0.08
v_2	0.7	0.9	1.0	0.8
v_3	0.01	0.03	0.5	0.02
v_4	0.1	0.5	0	0.9
v_5	0.1	0.2	0.05	1.0

U

⟹

v_1	v_2	v_3	v_4	v_5
c	c	c	d	d
b	b	b	b	b
a	d	d	a	a
d	a	a	c	c

P