Computational Social Choice and Complexity Theory

Ronald de Haan
University of Amsterdam
https://staff.science.uva.nl/r.dehaan/esslli2018
me@ronalddehaan.eu

ESSLLI 2018 – **Day 3**

Recap

- Voting theory, SCFs, SWFs
- Plurality, Borda, IRV, Kemeny
- (Parameterized) complexity theory
- Complexity of winner determination for different voting rules
- Strategic manipulation in voting
- The "computational defense" against manipulation, and its caveats

What we'll do today

- Bribery
- Domain Restrictions
- (Nearly) Single-Peaked Elections
- Black's Theorem

Bribery

Example: Swap Bribery for Borda

P	1	2	3	4	5
#1	a	b	c	e	e
#2	c	a	d	d	d
#3	b	f	b	b	c
#4	f	c	a	f	b
#5	e	e	f	c	f
#6	d	d	e	a	a

- We use Borda voting
- How many swaps of adjacent candidates do we need to make d (one of) the Borda winner(s)?
- What are the current Borda scores?

  a: 11, b: 16, c: 15, d: 12, e: 12, f: 9

Example: Swap Bribery for Borda

P	1	2	3	4	5
#1	a	b	d	e	e
#2	c	a	c	d	d
#3	f	f	b	b	c
#4	e	c	a	f	b
#5	d	e	f	c	f
#6	b	d	e	a	a

- We use Borda voting
- How many swaps of adjacent candidates do we need to make d (one of) the Borda winner(s)? 4
- What are the new Borda scores?

  a: 11, b: 13, c: 14, d: 14, e: 13, f: 10

Swap Bribery

- Let $P = (\succ_1, \dots, \succ_n)$ be a profile
- Let $\{c^i_{a,b}\}_{i \in N, a,b \in A}$ be a collection of swap costs $c^i_{a,b}$, one for each voter $i \in N$ and for each two candidates $a, b \in A$, $a \neq b$
- Candidates a, b can be swapped in $\succ_i$ for cost $c^i_{a,b}$ if a and b are adjacent in $\succ_i$

Swap-Bribery(F)

Input: a set N of voters, a set C of candidates, a preference profile P (for N and C), a candidate $c^* \in C$, a budget $b \in \mathbb{N}$, and a collection $\{c^i_{a,b}\}_{i \in N, a,b \in A}$ of swap costs.

Output: Can we get from P to some P' with a sequence of swaps of total cost at most b such that $c^* \in F(P')$?

Complexity of Swap Bribery for Borda

Theorem (Elkind, Faliszewski, Slinko, 2009)

Swap-Bribery(Borda) *is NP-complete.*

E. Elkind, P. Faliszewski, and A. Slinko. Swap Bribery. In: Proceedings of the 2nd International Symposium on Algorithmic Game Theory (SAGT), 299–310, 2009.

Complexity of Swap Bribery for c-Approval

- Let $c \in \mathbb{N}$. c-Approval is the voting rule that, for each voter, gives the top c candidates on that voters' ballot one point each, and selects as winners the candidates with highest score.

Theorem (Elkind, Faliszewski, Slinko, 2009)

Let $c \geq 3$. Swap-Bribery(c-Approval) *is NP-complete.*

Theorem (Dorn, Schlotter, 2012)

Let c be a constant. Swap-Bribery(c-Approval) *is:*

- *W[1]-hard when parameterized by b*
- *FPT when parameterized by $|A|$*
- *FPT when parameterized by $|N|$*

B. Dorn, and I. Schlotter. Multivariate Complexity Analysis of Swap Bribery. Algorithmica, 64(1), 126–151, 2012.

Variants of Bribery

- Swap Bribery: only adjacent swaps, each swap has a cost $c^i_{a,b} \in \mathbb{N}$
- Shift Bribery: swap bribery, with the restriction that each swap must involve the distinguished candidate
- Bribery/$Bribery: entire ballots can be changed for cost c_i (for Bribery, $c_i = 1$ for all $i \in N$)
- Microbribery: non-adjacent swaps can be made (leading to possibly non-linear orders)
- Weighted elections: the voters have different weights
- Constructive/destructive: for the constructive variants, the goal is to make a distinguished candidate $c^* \in C$ win; for the destructive variants, the goal is to make c^* lose

Different Interpretation of Bribery

- Bribery as strength / robustness of victory:
  - If one would need a high budget to make a different alternative win, then the current winner's victory is strong,

    and it is unlikely that the election has been tampered with

- Bribery as campaign guidance:
  - Take the polls as current election and find out which voters you need to bribe to make your candidate win,

    use this information to guide where to invest time and money in the election campaign

Domain Restrictions

The Assumption of Universal Domain

- In many of the core results that we saw so far, an implicit assumption is used
  - E.g., Arrow's Theorem, the Gibbard-Satterthwaite Theorem
- The assumption of universal domain:
  - We take into account all possible profiles P

The Condorcet Domain

- The Condorcet domain $\mathcal{D}_{\text{Condorcet}}$ consists of all profiles P that have a Condorcet winner
  - A Condorcet winner for P is a candidate $a \in A$ such that for all $b \in A$, $a \neq b$, it holds that a strict majority of voters in P prefer a to b

- There is a voting rule $F : \mathcal{D}_{\text{Condorcet}} \to A$ for the Condorcet domain that is resolute, anonymous, nonimposed and strategyproof
  - Take the rule F that always chooses the Condorcet winner

- *Downside:* being in $\mathcal{D}_{\text{Condorcet}}$ is a property of the profile, not of single ballots $\succ_i$

Single-Peakedness: Example

- Suppose you're going on a hike with your ESSLLI friends, and you want to vote on the distance of the hike

- $N = \{1, \ldots, n\}$: you and your friends (odd n);
  $A = \{1, \ldots, 20\}$: the possible distances (in km)

- Suppose also that each of you has an optimal distance d_i and that for each $d', d'' \in A$ it holds that:

  $$\text{if } d_i > d' > d'' \text{ or } d_i < d' < d'', \quad \text{then } d_i \succ_i d' \succ_i d''$$

  (such preferences are called single peaked)

- Then taking the median of all votes is a resolute, anonymous, nonimposed and strategyproof voting rule

Single-Peakedness: Example

P	1	2	3	4	5
#1	d_2	d_5	d_8	d_5	d_{10}
#2	d_5	d_8	d_{10}	d_2	d_8
#3	d_8	d_2	d_5	d_8	d_5
#4	d_{10}	d_{10}	d_2	d_{10}	d_2

- The median is d_5 (the median of $d_2, d_5, d_5, d_8, d_{10}$)
- Why is taking the median strategyproof?
  - Let d^* be the median
  - Take some voter i and take some $d' \succ_i d^*$
  - A majority of voters (not including i) prefers d^* to d'
  - So there is no insincere ballot $\succ_i'$ that i can cast to make a majority prefer d' to d^*

Single-Peakedness

- Let $>$ be a linear order over a set A of alternatives
- Let $\succ_i$ be a ballot and let a_i the top-ranked alternative in $\succ_i$.
- Then $\succ_i$ is single peaked w.r.t. $>$ if for each $a', a'' \in A$ it holds that:

$$\text{if } a_i > a' > a'' \text{ or } a_i < a' < a'', \quad \text{then } a_i \succ_i a' \succ_i a''$$

- A profile $P = (\succ_1, \ldots, \succ_n)$ is single peaked if there exists one linear order $>$ over A such that all $\succ_i$ in P are single peaked w.r.t. $>$

- *Possible motivation*: think of a political left-to-right spectrum; all voter and candidates are located on this spectrum; each voter ranks the candidates from close to far on the spectrum

Black's Theorem

- Let P be a single-peaked profile and let $>$ be the linear order that witnesses single-peakedness
- The median rule is the voting rule that selects the median (w.r.t. $>$) of the top-ranked candidates in the profile

Theorem (Black, 1958)

Let P be a single-peaked profile for odd n. Then P is in $\mathcal{D}_{\text{Condorcet}}$, and the median rule selects the Condorcet winner of P. Moreover, the median rule is resolute, anonymous, nonimposed and strategyproof for the single-peaked domain.

- So this is one way of getting around the Gibbard-Satterthwaite Theorem

Tractability Results for Condorcet Extensions

- Every single-peaked profile P is in $\mathcal{D}_{\text{Condorcet}}$
- Computing the Condorcet winner of a profile in $\mathcal{D}_{\text{Condorcet}}$ is polynomial-time solvable
- Let F be a Condorcet extension.

$\Rightarrow$ WinDet(F) is polynomial-time solvable when restricted to the single-peaked domain

Theorem

WinDet(Kemeny) *is polynomial-time solvable when restricted to the single-peaked domain*

Identifying Single-Peakedness

- How hard is it to check if a given election is single peaked?
- Polynomial-time algorithm:
  - Iteratively construct the ordering $>$ in rounds:
  - In each round, consider the (unassigned) alternatives that are bottom-ranked in the profile

    (these have to be assigned at the sides of the ordering)
  - If there are ≥ 3 of them, fail
  - If there are ≤ 2, try to find a way to put them at the sides of the partially constructed ordering
    - If this is not possible, fail
    - Otherwise, assign them, and continue

B. Escoffier, J. Lang, and M. Öztürk. Single-Peaked Consistency and Its Complexity. In Proceedings of the 8th European Conference on Artificial Intelligence (ECAI), 366–370, 2008.

Identifying Single-Peakedness

P	1	2
#1	a_1	a_6
#2	a_2	a_5
#3	a_3	a_4
#4	a_4	a_3
#5	a_5	a_2
#6	a_6	a_1

Identifying Single-Peakedness

P	1	2
#1	a_1	a_6
#2	a_2	a_5
#3	a_3	a_4
#4	a_4	a_3
#5	a_5	a_2
#6	a_6	a_1

$a_1 >$ $> a_6$

Identifying Single-Peakedness

P	1	2
#1	a_1	a_6
#2	a_2	a_5
#3	a_3	a_4
#4	a_4	a_3
#5	a_5	a_2
#6	a_6	a_1

$a_1 > a_5 > \qquad > a_2 > a_6$

Identifying Single-Peakedness

P	1	2
#1	a_1	a_6
#2	a_2	a_5
#3	a_3	a_4
#4	a_4	a_3
#5	a_5	a_2
#6	a_6	a_1

$a_1 > a_5 > \qquad > a_2 > a_6$

Identifying Single-Peakedness

P	1	2
#1	a_1	a_6
#2	a_2	a_5
#3	a_3	a_4
#4	a_4	a_3
#5	a_5	a_2
#6	a_6	a_1

$a_1 > a_2 >$ $\qquad$ $> a_5 > a_6$

Identifying Single-Peakedness

P	1	2
#1	a_1	a_6
#2	a_2	a_5
#3	a_3	a_4
#4	a_4	a_3
#5	a_5	a_2
#6	a_6	a_1

$$a_1 > a_2 > a_3 > a_4 > a_5 > a_6$$

Identifying Single-Peakedness

P	1	2	3
#1	a_1	a_6	a_1
#2	a_2	a_5	a_5
#3	a_3	a_4	a_3
#4	a_4	a_3	a_4
#5	a_5	a_2	a_2
#6	a_6	a_1	a_6

$$a_1 > a_2 > \qquad\qquad > a_5 > a_6$$

Identifying Single-Peakedness

P	1	2	3
#1	a_1	a_6	a_1
#2	a_2	a_5	a_5
#3	a_3	a_4	a_3
#4	a_4	a_3	a_4
#5	a_5	a_2	a_2
#6	a_6	a_1	a_6

$a_1 > a_2 > \qquad > a_5 > a_6$

Identifying Single-Peakedness

P	1	2	3
#1	a_1	a_6	a_1
#2	a_2	a_5	a_5
#3	a_3	a_4	a_3
#4	a_4	a_3	a_4
#5	a_5	a_2	a_2
#6	a_6	a_1	a_6

fail

Single-Peakedness in Practice

- Single-peakedness is not robust, in the following sense:
- You can have a profile of 1 million voters whose preferences are single peaked w.r.t. the same ordering $>$

  and one 'maverick' voter whose preference makes the profile not single-peaked

Nearly Single-Peakedness

- Various notions of nearly single-peakedness have been considered:
  - single-peakedness after deleting a small # of voters
  - single-peakedness after deleting a small # of alternatives
  - single-peakedness after contracting alternatives into small intervals
  - etc.

Nearly Single-Peakedness: Example

P	1	2	3
#1	a	c	f
#2	b	a	e
#3	c	b	d
#4	d	d	b
#5	e	e	c
#6	f	f	a

- Is this profile single peaked? No, no Condorcet winner
- What if we merge a, b, c?

Nearly Single-Peakedness: Example

P	1	2	3
#1	u	u	f
#2	d	d	e
#3	e	e	d
#4	f	f	u

- Is **this** profile single peaked? Yes

SP Width

- The Single-Peaked (SP) Width of a profile P is one way to measure how close P is to being single peaked
- A subset $I \subseteq A$ of alternatives is an interval if the alternatives in I are consecutive in each ballot in P
- A contraction of intervals of P is the profile obtained by partitioning A into intervals $I_1, \ldots, I_u$ and replacing each interval I_j by some candidate $a_j \in I_j$
- The SP width of P is the smallest value of $\max_i |I_i| - 1$ over all contractions of intervals that are single-peaked

D. Cornaz, L. Galand, and O. Spanjaard. Kemeny Elections with Bounded Single-Peaked or Single-Crossing Width. In Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI), 76–82, 2013.

SP Width: Example

P	1	2	3
#1	a	c	f
#2	b	a	e
#3	c	b	d
#4	d	d	b
#5	e	e	c
#6	f	f	a

- This profile P has SP width 2
- $I_1 = \{a, b, c\}$, $I_2 = \{d\}$, $I_3 = \{e\}$, $I_4 = \{f\}$

or

- $I_1 = \{a, b, c\}$, $I_2 = \{d, e, f\}$

Kemeny Voting for Small SP Width

Theorem (Cornaz, Galand, Spanjaard, 2013)

The SP width of a profile P, and a corresponding contraction of intervals, can be computed in polynomial time.

Theorem (Cornaz, Galand, Spanjaard, 2013)

Winner Determination for the Kemeny voting rule is fixed-parameter tractable when parameterized by the SP width w of the profile.

- *Idea:*
  - compute a suitable contraction of intervals
  - take the Condorcet winner of the reduced profile
  - consider all $\leq (w+1)!$ rankings of the winning interval

Some Other Topics in Voting

Other Types of Ballots

- So far, we looked at linear orders as ballots:

$$a_1 \succ a_2 \succ \cdots \succ a_m$$

- In some settings, one would like to elicit preferences from voters in a different form
  - For example, as weak orders (transitive, complete, reflexive):

e.g., $a_1 \succ a_2 \sim a_3 \succ a_4 \sim a_4 \succ \cdots \succ a_m$

e.g., $a_1 \succ \cdots \succ a_\ell \succ a_{\ell+1} \sim \cdots \sim a_m$

  - A special case of weak orders are dichotomous preferences: (two equivalence classes, ordered by $\succ$)

e.g., $a_1 \sim \cdots \sim a_\ell \succ a_{\ell+1} \sim \cdots \sim a_m$

- An interesting application using dichotomous preferences is approval voting:
  - For each vote, each alternative in the top equivalence class gets a point
  - Alternatives with highest score win

Control

- Control refers the act of influencing the election (with a certain goal) by an outside agent, by means of changing the structure of the election
- For example, consider the following Plurality election:

P	9x	8x	2x
#1	a	b	c
#2	b	c	b
#3	c	a	a

  Removing candidate c from this election changes the outcome
- Various types of control have been studied, e.g.:
  - Deleting voters or alternatives
  - Adding voters or alternatives
  - Gerrymandering

Incomplete Information

- Several problems related to voting often make more sense in the setting of incomplete information
  - E.g., strategic manipulation, bribery, etc.
  - Also: after having received a subset of the votes, do we have enough information to decide the outcome of the election?

- Often modelled by taking ballots to be preorders $\succeq$ (transitive, reflexive relations over A)

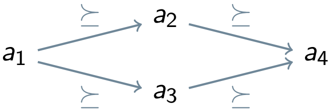

Interpretation: consider all linear orders consistent with this preorder as possibilities

Combinatorial Domains

- Suppose the ESSLLI organizing committee wanted to vote on what to offer at the coffee breaks, e.g., on whether or not there should be:
  - chocolate cookies
  - vanilla cookies
  - hazelnut cookies
  - apples
  - bananas
  - oranges
- Preferences over these topics are often not independent, e.g.:
  - "If we have chocolate cookies, I don't want vanilla cookies."
  - "If we have hazelnut cookies, I prefer apples over bananas."
- Many challenges arise in this setting:
  - E.g., for u issues, how can we efficiently represent preferences over the 2^u different combinations?

Iterated Voting

- We considered strategic manipulation as a one-shot move, where other voters have no chance to react
- Iterated voting studies the setting where each agent at every point is able to change their vote
  - Think of an online tool where you can hold an election, where you can see others' votes, and where you can change your vote
- In this setting, different questions arise, e.g.:
  - When voters always change their vote in their best (immediate) interest, does the process converge?
  - Does the result of this convergence lead to a result that is socially preferred over the truthful outcome?
  - etc.

Multiwinner / Committee Elections

- In many cases, it is desirable to select multiple winners
- For example, when selecting a committee

- Other constraints and desiderata could play a role:
  - Diversity constraints on the selected committee
  - Proportional representation

Recap

- Bribery
- Domain Restrictions
- (Nearly) Single-Peaked Elections
- Black's Theorem

Homework exercise

- Consider the following profile P

P	1	2	3
#1	g	c	b
#2	e	e	a
#3	d	d	d
#4	c	b	c
#5	b	a	e
#6	a	g	g
#7	f	f	f

- Show that P is not single peaked
- Find the SP width of P
- Compute the Kemeny winners of P