Social Choice

Joseph Chuang-Chieh Lin

Dept. CSIE, Tamkang University, Taiwan

Outline

The Setting of Social Choice

Take voting scheme for example.

- A set O of outcomes (i.e., alternatives, candidates, etc.)
- A set A of agents s.t. each of them has a preference $\succ$ over the outcomes.
- The social choice function: a mapping from the profiles of the preferences to a particular outcome.

Outcomes & preferences

outcomes : a, b, c, d

$d \succ b \succ a \succ c$

preferences

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

Preferences

- A binary relation $\succ$ such that
  - for every $a, b \in O$, $a \neq b$, we have either $a \succ b$ or $b \succ a$ but NOT both.
  - for $a, b, c \in O$, if $a \succ b$ and $b \succ c$, then we have $a \succ c$.
- $\succeq$ can be defined similarly.
  - $\prec$: $\neg\succ$

Agents with preferences

- E.g., five agents (voters).
- Each agent has its preference over four candidates $\{a, b, c, d\}$.

outcomes : a, b, c, d

$d \succ b \succ a \succ c$

preferences

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

Agents with preferences

- E.g., three agents (voters).
- Each agent has its preference over four candidates $\{a, b, c, d\}$.

v_1	v_2	v_3
d	b	a
b	c	b
a	a	c
c	d	d

Plurality rule $\Rightarrow$ a

outcomes : a, b, c, d

$d \succ b \succ a \succ c$

preferences

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

- Plurality rule: each agent can only give score 1 to the most preferred one and 0 to the others.

Plurality rule (contd.)

v_1	v_2	v_3
d	***b***	***a***
b	c	b
a	a	c
c	d	d

- Plurality rule: each agent can only give score 1 to the most preferred one and 0 to the others.

Plurality rule (contd.)

v_1	v_2	v_3
$\mathbf{d}$	$\mathbf{b}$	$\mathbf{a}$
b	c	b
a	a	c
c	d	d

- Plurality rule:

Plurality rule (contd.)

v_1	v_2	v_3
d	***b***	***a***
b	c	b
a	a	c
c	d	d

- Plurality rule: depending on the tie-breaking rule.

Condorcet rule

v_1	v_2	v_3
d	b	a
b	c	b
a	a	c
c	d	d

- Condorcet rule:
  - a vs. b
  - a vs. c
  - a vs. d

Condorcet rule

v_1	v_2	v_3
d	b	a
b	c	b
a	a	c
c	d	d

- Condorcet rule:
  - a vs. $b \to b$
  - a vs. $c \to a$
  - a vs. $d \to a$

Condorcet rule

v_1	v_2	v_3
d	b	a
b	c	b
a	a	c
c	d	d

- Condorcet rule:
  - c vs. a
  - c vs. b
  - c vs. d

Condorcet rule

v_1	v_2	v_3
d	b	a
b	c	b
a	a	c
c	d	d

- Condorcet rule:
  - c vs. $a \rightarrow a$
  - c vs. $b \rightarrow b$
  - c vs. $d \rightarrow c$

Condorcet rule

v_1	v_2	v_3
d	b	a
b	c	b
a	a	c
c	d	d

- Condorcet rule:
  - b vs. a
  - b vs. c
  - b vs. d

Condorcet rule

v_1	v_2	v_3
d	b	a
b	c	b
a	a	c
c	d	d

- Condorcet rule:
  - b vs. $a \rightarrow b$
  - b vs. $c \rightarrow b$
  - b vs. $d \rightarrow b$

Condorcet rule

v_1	v_2	v_3
d	b	a
b	c	b
a	a	c
c	d	d

- Condorcet rule: b
  - b vs. $a \rightarrow b$
  - b vs. $c \rightarrow b$
  - b vs. $d \rightarrow b$

Borda rule

v_1		v_2		v_3	
d	**3**	b	**3**	a	**3**
b	**2**	c	**2**	b	**2**
a	**1**	a	**1**	c	**1**
c	**0**	d	**0**	d	**0**

- Borda count rule:

Borda rule

v_1		v_2		v_3	
d	**3**	b	**3**	a	**3**
b	**2**	c	**2**	b	**2**
a	**1**	a	**1**	c	**1**
c	**0**	d	**0**	d	**0**

- Borda count rule:
  - score of a: $1 + 1 + 3 = 5$.
  - score of b: $2 + 3 + 2 = 7$.
  - score of c: $0 + 2 + 1 = 3$.
  - score of d: $3 + 0 + 0 = 3$.

Borda rule

v_1		v_2		v_3	
d	**3**	b	**3**	a	**3**
b	**2**	c	**2**	b	**2**
a	**1**	a	**1**	c	**1**
c	**0**	d	**0**	d	**0**

- Borda count rule: b.
  - score of a: $1 + 1 + 3 = 5$.
  - score of b: $2 + 3 + 2 = 7$.
  - score of c: $0 + 2 + 1 = 3$.
  - score of d: $3 + 0 + 0 = 3$.

Inefficiency of Borda Count

v_1		v_2		v_3		v_4		v_5	
a	**2**	a	**2**	a	**2**	b	**2**	b	**2**
b	**1**	b	**1**	b	**1**	c	**1**	c	**1**
c	**0**	c	**0**	c	**0**	a	**0**	a	**0**

Inefficiency of Borda Count

Points	v_1	v_2	v_3	v_4	v_5
2	a	a	a	b	b
1	b	b	b	c	c
0	c	c	c	a	a

- Who is the winner by Borda counting?

Inefficiency of Borda Count

v_1		v_2		v_3		v_4		v_5	
a	**2**	a	**2**	a	**2**	b	**2**	b	**2**
b	**1**	b	**1**	b	**1**	c	**1**	c	**1**
c	**0**	c	**0**	c	**0**	a	**0**	a	**0**

- Who is the winner by Borda counting? *a*: 6, *b*: 7, *c*: 2.

Inefficiency of Borda Count

Points	v_1	v_2	v_3	v_4	v_5
2	a	a	a	b	b
1	b	b	b	c	c
0	c	c	c	a	a

- Who is the winner by Borda counting? *a*: 6, *b*: 7, *c*: 2.
- Condorcet principle follows?

Inefficiency of Borda Count

v_1		v_2		v_3		v_4		v_5	
a	**2**	a	**2**	a	**2**	b	**2**	b	**2**
b	**1**	b	**1**	b	**1**	c	**1**	c	**1**
c	**0**	c	**0**	c	**0**	a	**0**	a	**0**

- Who is the winner by Borda counting? a: 6, b: 7, c: 2.
- Condorcet principle follows? $a \succ b$, $a \succ c$.

Inefficiency of Borda Count

v_1		v_2		v_3		v_4		v_5	
a	**2**	a	**2**	a	**2**	b	**2**	b	**2**
b	**1**	b	**1**	b	**1**	c	**1**	c	**1**
c	**0**	c	**0**	c	**0**	a	**0**	a	**0**

- Who is the winner by Borda counting? a: 6, b: 7, c: 2.
- Condorcet principle follows? $a \succ b$, $a \succ c$.
- Who is the winner under the plurality rule?

Inefficiency of Borda Count

v_1		v_2		v_3		v_4		v_5	
a	**2**	a	**2**	a	**2**	b	**2**	b	**2**
b	**1**	b	**1**	b	**1**	c	**1**	c	**1**
c	**0**	c	**0**	c	**0**	a	**0**	a	**0**

- Who is the winner by Borda counting? a: 6, b: 7, c: 2.
- Condorcet principle follows? $a \succ b$, $a \succ c$.
- Who is the winner under the plurality rule? a.

Successive elimination

v_1	v_2	v_3
b	a	c
d	b	a
c	d	b
a	c	d

- Successive elimination with ordering $a \to b \to c \to d$:

Successive elimination

v_1	v_2	v_3
b	a	c
d	b	a
c	d	b
a	c	d

- Successive elimination with ordering $a \rightarrow \not{b} \rightarrow c \rightarrow d$:

Successive elimination

v_1	v_2	v_3
b	a	c
d	b	a
c	d	b
a	c	d

- Successive elimination with ordering $\not a \to \not b \to c \to d$:

Successive elimination

v_1	v_2	v_3
b	a	c
d	b	a
c	d	b
a	c	d

- Successive elimination with ordering $\not{a} \to \not{b} \to \not{c} \to d$:

Successive elimination

v_1	v_2	v_3
b	a	c
d	b	a
c	d	b
a	c	d

- Successive elimination with ordering $\not{a} \to \not{b} \to \not{c} \to d$: d

Successive elimination

v_1	v_2	v_3
b	a	c
d	b	a
c	d	b
a	c	d

- Successive elimination with ordering $\not a \rightarrow \not b \rightarrow \not c \rightarrow d$: d
  - The issue: all of the agents prefer b to d!

Successive elimination

v_1	v_2	v_3
b	a	c
d	b	a
c	d	b
a	c	d

- Successive elimination with ordering $a \to b \to c \to d$: d
- Successive elimination with ordering $a \to c \to b \to d$:

Successive elimination

v_1	v_2	v_3
b	a	c
d	b	a
c	d	b
a	c	d

- Successive elimination with ordering $a \to b \to c \to d$: d
- Successive elimination with ordering $a \to c \to b \to d$: b

Successive elimination (sensitive to the agenda order)

v_1	v_2	v_3
b	a	c
d	b	a
c	d	b
a	c	d

- Successive elimination with ordering $a \to b \to c \to d$: d
- Successive elimination with ordering $a \to c \to b \to d$: b
- Successive elimination with ordering $b \to c \to a \to d$:

Successive elimination (sensitive to the agenda order)

v_1	v_2	v_3
b	a	c
d	b	a
c	d	b
a	c	d

- Successive elimination with ordering $a \to b \to c \to d$: d
- Successive elimination with ordering $a \to c \to b \to d$: b
- Successive elimination with ordering $b \to c \to a \to d$: a

Condorcet Winner vs. Plurality

- Let's say we have 1,000 agents each of which has a preference over three candidates A, B, C.
  - 499 agents for $A \succ B \succ C$.
  - 3 agents for $B \succ C \succ A$.
  - 498 agents for $C \succ B \succ A$.
- Who is the Condorcet winner?

Condorcet Winner vs. Plurality

- Let's say we have 1,000 agents each of which has a preference over three candidates A, B, C.
  - 499 agents for $A \succ B \succ C$.
  - 3 agents for $B \succ C \succ A$.
  - 498 agents for $C \succ B \succ A$.
- Who is the Condorcet winner? B.

Condorcet Winner vs. Plurality

- Let's say we have 1,000 agents each of which has a preference over three candidates A, B, C.
  - 499 agents for $A \succ B \succ C$.
  - 3 agents for $B \succ C \succ A$.
  - 498 agents for $C \succ B \succ A$.
- Who is the Condorcet winner? B.
- Who is the winner under the plurality rule?

Condorcet Winner vs. Plurality

- Let's say we have 1,000 agents each of which has a preference over three candidates A, B, C.
  - 499 agents for $A \succ B \succ C$.
  - 3 agents for $B \succ C \succ A$.
  - 498 agents for $C \succ B \succ A$.
- Who is the Condorcet winner? B.
- Who is the winner under the plurality rule? A.

Exercise

On Borda Count & Condorcet

We have five voters with the following preferences (ordering) over the outcomes A, B, C, and D.

- $B \succ C \succ A \succ D$.
- $B \succ D \succ C \succ A$.
- $D \succ C \succ A \succ B$.
- $A \succ D \succ B \succ C$.
- $A \succ D \succ C \succ B$.

Who is the winner by the Borda Count rule?
Who is the Condorcet winner?

Let's consider a practical application in MOOCs.

MOOCs

- MOOCs: Massive Online Open Courses
  - e.g., Coursera, EdX.

MOOCs

- MOOCs: Massive Online Open Courses
  - e.g., Coursera, EdX.
- Outscourcing the grading task to the students.

MOOCs

- MOOCs: Massive Online Open Courses
  - e.g., Coursera, EdX.
- Outscourcing the grading task to the students.
- They may have incentives to assign LOW scores to everybody else.

MOOCs

- MOOCs: Massive Online Open Courses
  - e.g., Coursera, EdX.
- Outscourcing the grading task to the students.
- They may have incentives to assign LOW scores to everybody else.
  - ▷ Ask each student to grade a SMALL number of her peers' assignments.

MOOCs

- MOOCs: Massive Online Open Courses
  - e.g., Coursera, EdX.
- Outscourcing the grading task to the students.
- They may have incentives to assign LOW scores to everybody else.
  - ▷ Ask each student to grade a SMALL number of her peers' assignments.
  - Then merge individual rankings into a global one.

Terminologies

- $\mathcal{A}$: universe of n elements (students).
- (n, k)-grading scheme:

  a collection $\mathcal{B}$ of size-k subsets (bundles) of $\mathcal{A}$, such that each element of $\mathcal{A}$ belongs to exactly k subsets of $\mathcal{B}$.
- The bundle graph:

  Represent the (n, k)-grading scheme with a bipartite graph.
- $\prec_b$: a ranking of the element b contains (partial order).

The aggregation rule

An aggregation rule:
profile of partial rankings $\mapsto$ complete ranking of all elements.

- Borda:

SPRING FEAST 2016 BALLOT

a	LE BLE D'OR	5
b	CRYSTAL SPOON	3
c	Bei Yuan Restaurant	1
d	Tasty Steak	2
e	Capricciosa	4

SPRING FEAST 2016 BALLOT

a	LE BLE D'OR	5
b	CRYSTAL SPOON	4
c	Bei Yuan Restaurant	2
d	Tasty Steak	1
e	Capricciosa	3

SPRING FEAST 2016 BALLOT

a	LE BLE D'OR	4
b	CRYSTAL SPOON	5
c	Bei Yuan Restaurant	1
d	Tasty Steak	3
e	Capricciosa	2

- a: 14; b: 12; c: 4; d: 6; e: 9.

$$a \prec b \prec e \prec d \prec c.$$

Order-revealing grading scheme

An aggregation rule in peer grading (Borda):

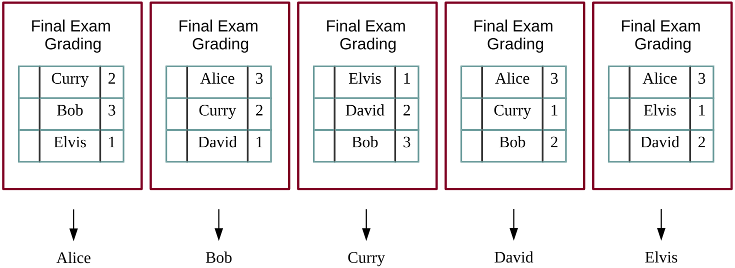

- Alice: 9; Bob: 8; Curry: 5; David: 5; Elvis: 3.

  Alice $\prec$ Bob $\prec$ Curry $\prec$ David $\prec$ Elvis.

Assumption (perfect grading)

Each student grades the assignments in her bundle consistently to the ground truth.

Order-revealing grading scheme (contd.)

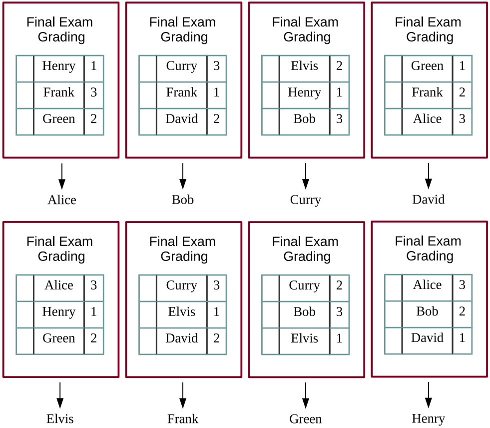

- Alice: 9; Bob: 8; Curry: 8; David: 5; Elvis: 4; Frank: 6; Green: 5; Henry: 3.
  Alice $\prec$ Bob $\prec$ Curry $\prec$ Frank $\prec$ David $\prec$ Green $\prec$ Elvis $\prec$ Henry.

Order-revealing grading scheme (contd.)

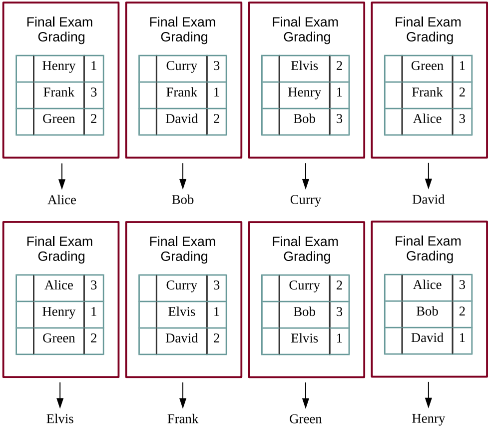

- Alice: 9; Bob: 8; Curry: 8; David: 5; Elvis: 4; Frank: 6; Green: 5; Henry: 3.

  Alice $\prec$ Bob $\prec$ Curry $\prec$ Frank $\prec$ David $\prec$ Green $\prec$ Elvis $\prec$ Henry.

The bundle graph

The bundle graph:

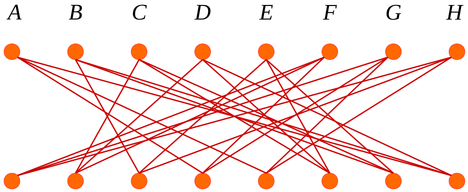

The bundle graph

The bundle graph:

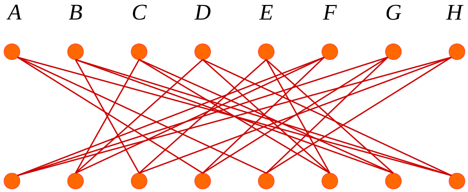

- A random k-regular graph:

  A complete bipartite $K_{n,n} \mapsto$ removing edges $\{v, v\}$, $\forall v \mapsto$

  repeat

  > *"draw a perfect matching uniformly at random among all perfect matchings of the remaining graph"*

  for k times.

The limitation on the order revealing scheme

- The property of revealing the ground truth for certain:

$$\forall x, y \in \mathcal{A},\ \exists B \in \mathcal{B} \text{ such that } x, y \in B.$$

The limitation on the order revealing scheme

- The property of revealing the ground truth for certain:

$$\forall x, y \in \mathcal{A},\ \exists B \in \mathcal{B} \text{ such that } x, y \in B.$$

  - Suppose NO bundle contains both $x, y \in \mathcal{A}$.
  - Let $\prec, \prec'$ be two complete rankings.
    - x, y are in the first two positions in $\prec, \prec'$;
    - $\prec$ and $\prec'$ differs only in the order of x and y.
  - Clearly, partial rankings within the bundles are identical in both cases.
  - No way to identify whether $\prec$ or $\prec'$ is the ground truth.

The limitation on the order revealing scheme

- The property of revealing the ground truth for certain:

  $$\forall x, y \in \mathcal{A},\ \exists B \in \mathcal{B} \text{ such that } x, y \in B.$$

  - Suppose NO bundle contains both $x, y \in \mathcal{A}$.
  - Let $\prec, \prec'$ be two complete rankings.
    - x, y are in the first two positions in $\prec, \prec'$;
    - $\prec$ and $\prec'$ differs only in the order of x and y.
  - Clearly, partial rankings within the bundles are identical in both cases.
  - No way to identify whether $\prec$ or $\prec'$ is the ground truth.
- To reveal the ground truth with certainty: $k = \Omega(\sqrt{n})$.
  - $n \cdot \binom{k}{2} \geq \binom{n}{2}$.

Seeking for approximate order-revealing grading schemes

- Use a bundle graph with a very low degree k (independent of n).
- Randomly permute the elements by $\pi : U \mapsto \mathcal{A}$ before associating them to the nodes of U of the bundle graph.
- Aiming at $\frac{\#\text{correctly recovered pairwise relations}}{\binom{n}{2}}$.

The main result

Theorem (Caragiannis, Krimpas, Voudouris@AAMAS'15)

When

- Borda is applied as the aggregation rule, and
- all the partial rankings are consistent to the ground truth,

then the expected fraction of correctly recovered pairwise relations is $1 - O(1/\sqrt{k})$.

Question

- What will happen if we assign for each student only two assignments and each assignment is graded by exactly two students?