CSC2556

Algorithms for Collective Decision Making

Nisarg Shah

Introduction

- People
  - Instructor: Nisarg Shah (/~nisarg, nisarg@cs)
  - TA: Evi Micha (emicha@cs), Soroush Ebadian (soroush@cs)
- Info
  - Course Page: `cs.toronto.edu/~nisarg/teaching/2556s22/`
  - Discussion Board: `piazza.com/utoronto.ca/winter2022/csc2556`
- Meeting
  - Lectures: Online (Zoom) until at least Jan 31
  - Questions? Schedule 1-1 meeting by emailing me

What?

- Collective decision making by groups of agents
  - Literature of *computational social choice* at the intersection of computer science and economics
- Single-agent problems
  - E.g., the traveling salesman problem
  - A single agent wants to find the optimal route
- Multi-agent problems
  - What if multiple traveling salesmen want to share a bus?
  - Each agent has a different optimal route
  - Tradeoff → fairness, efficiency, strategic manipulations, …

What?

- Models will differ in various considerations, e.g.:
  - Can agents form binding contracts?
    - Entering in contracts allows agents to hedge uncertainties
  - Can agents exchange/pay/receive money?
    - Maybe we make a decision that is less preferable to an agent but pay the agent to compensate
  - What is the structure of the outcome space?
    - Is there a common decision that affects everyone (e.g., voting) or does each agent receive something (e.g., resource allocation)?

Logistics

Optional Reference Textbooks

- Handbook of Computational Social Choice
  - By Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia
- Algorithmic Game Theory
  - By Noam Nisan, Tom Roughgarden, Eva Tardos and Vijay Vazirani
- Networks, Crowds and Markets
  - By David Easley and Jon Kleinberg
- Online versions available on the course web page

Grading Policy

- 2 assignments: 40%
- Final project: 50%
- Embedded Ethics Module: 5%
  - Pre-module survey: 1%
  - Post-module survey: 1%
  - Post-module assignment: 3%
- Class participation: 5%

Assignments

- Theoretical
  - They will require deriving intricate proofs

- We will assume…
  - Strong familiarity with abstract reasoning and proof techniques
  - Adequate familiarity of CS concepts (e.g., algorithm design, worst-case approximation, NP-hardness)
  - Adequate familiarity of math concepts (e.g., probability, statistics, linear algebra, calculus)
  - No prior background in economics

Assignments

- Individual assignments
  - Free to discuss with classmates or read online material
  - Must write solutions in your own words
    - Easier if you do not take any pictures/notes from the discussions
  - Plagiarism will be dealt with strictly!

- Citation
  - For each question, you must cite the peer (write the name) or the online sources (provide links) referred, if any
  - Failing to do this is also plagiarism!

Other Policies

- “No Garbage” Policy
  - Borrowed from: Prof. Allan Borodin (citation!)

1. Partial marks for viable approaches
2. Zero marks if the answer makes no sense
3. 20% marks if you admit to not knowing how to solve

- 20% > 0% !!

Course Timeline

- (Approximate dates)
- $\approx$ Feb 1: HW1 posted
- $\approx$ Feb 15: HW1 due
- $\approx$ Mar 1: HW2 posted, project proposal due
- $\approx$ Mar 15: HW2 due
- $\approx$ Week of Mar 15: Mid-project check-in
- Last 1.5-2 lectures: Project presentations
- $\approx$ April 10: Project reports due

Course Project

- How?
  - Groups of 1-3
    - Larger groups are better
    - Find partners early, but maybe after the enrollment stabilizes

- What?
  - **Empirical:** Quantitative analysis of algorithms presented in class (or your own) using simulations or real data
  - **Theoretical:** Prove new observations about the algorithms or design new algorithms for a problem
  - **Ideal:** A bit of both

Project Topic

- From your own research area of interest
  - We'll introduce broad concepts that you may be able to apply to your own research area in order to find a project topic
  - E.g., fairness, allocation efficiency, preference elicitation, …
- From the course
  - I'll mention some open problems as we go along
  - Later, I'll also post sample projects from previous years as well as sample project ideas for this year
  - You can also study realistic variants of problems that we see in class

Course Project: Timeline

- Find partners and think about a project idea
- Submission 1: Project proposal
  - Ideally 1 page but up to 2 pages excluding references
  - Outline of the idea, prior work, reasonable goals
- Mid-project meetings
  - Optional, 1-1 with me, 30-minute
- Class presentations
- Submission 2: Final project report
  - Up to 5 pages excluding references and appendix
  - Focus on quality academic writing

Introductions

Brief Introductions

- What to say?
  - Which program?
  - Which year?
  - Who are you working with (if any)?
  - What is your area of interest (if any)?
  - Anything else you'd like to share

Overview of the Course

Social Choice, Mechanism Design

- Social choice
  - Given the preferences of the agents, which collective decision is the most desirable?
  - Fairness, welfare, ethics, resource utilization, …

- Mechanism design
  - Agents have private information, which they may lie about
  - How to design the "rules of the game" such that selfish agent behavior results in desirable outcomes
  - We call this "implementing" the social choice rule

Mechanism Design

- With money
  - Principal can "charge" the agents (require payments)
  - Helps significantly
  - **Example:** auctions

- Without money
  - Monetary transfers are not allowed
  - Incentives must be balanced otherwise
  - Often impossible without sacrificing the objective a little
  - **Example:** elections, kidney exchange

Example: Auction

Objective: The one who really needs it more should have it.

Rule 1: Each would tell me his/her value. I'll give it to the one with the higher value.

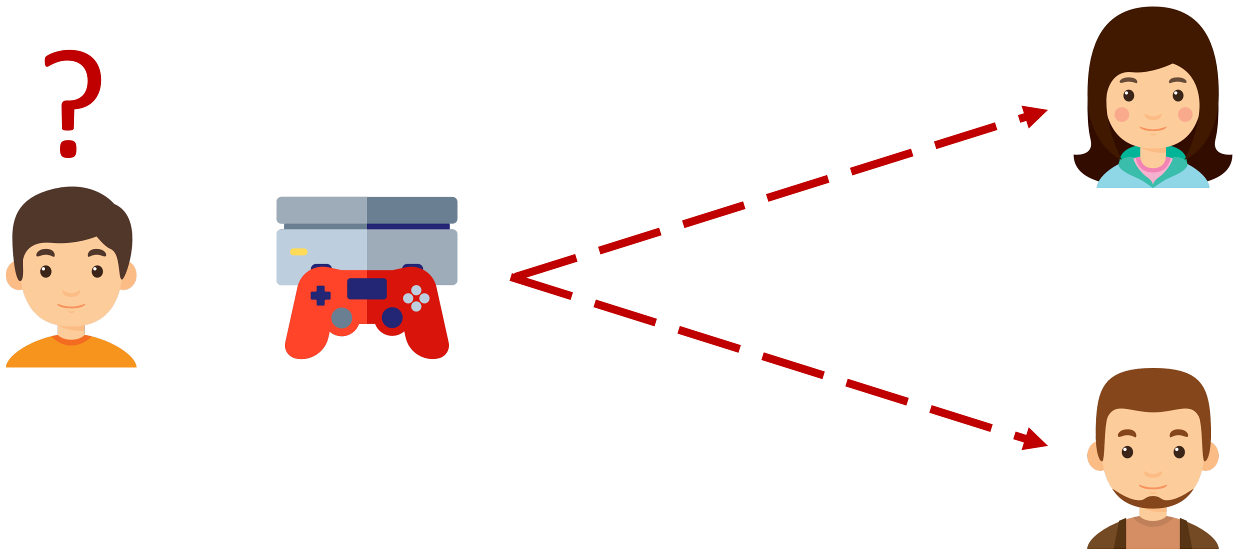

Image Courtesy: Freepik

Example: Auction

Objective: The one who really needs it more should have it.

Rule 2: Each would tell me his/her value. I'll give it to the one with the higher value, but they must pay me that value.

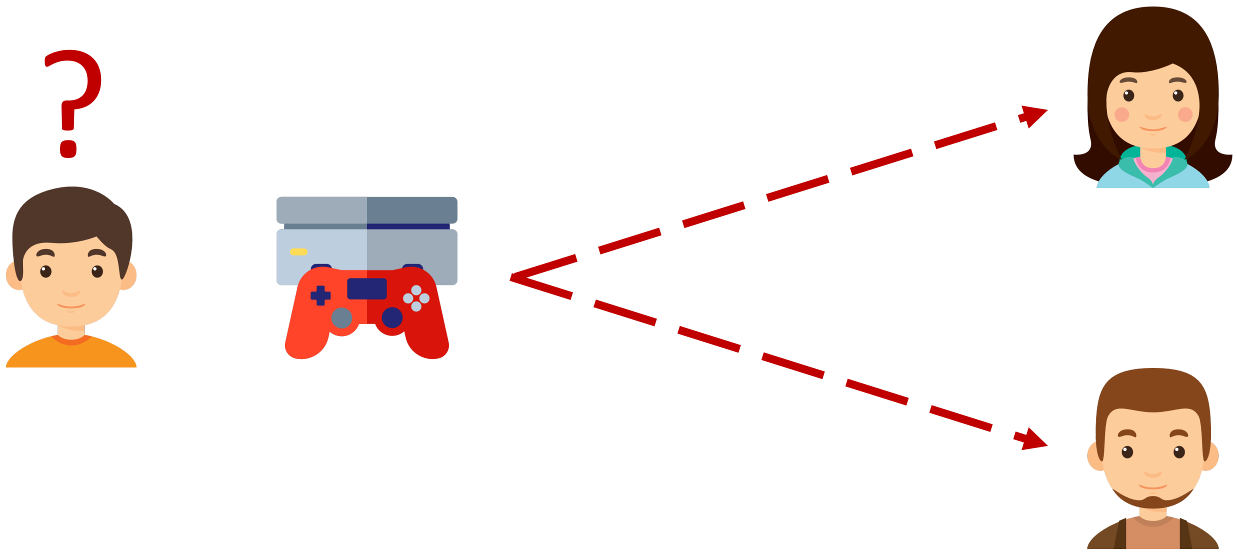

Image Courtesy: Freepik

Example: Auction

Objective: The one who really needs it more should have it.

Question: Can I make it easier so that each can just truthfully tell me how much they value it?

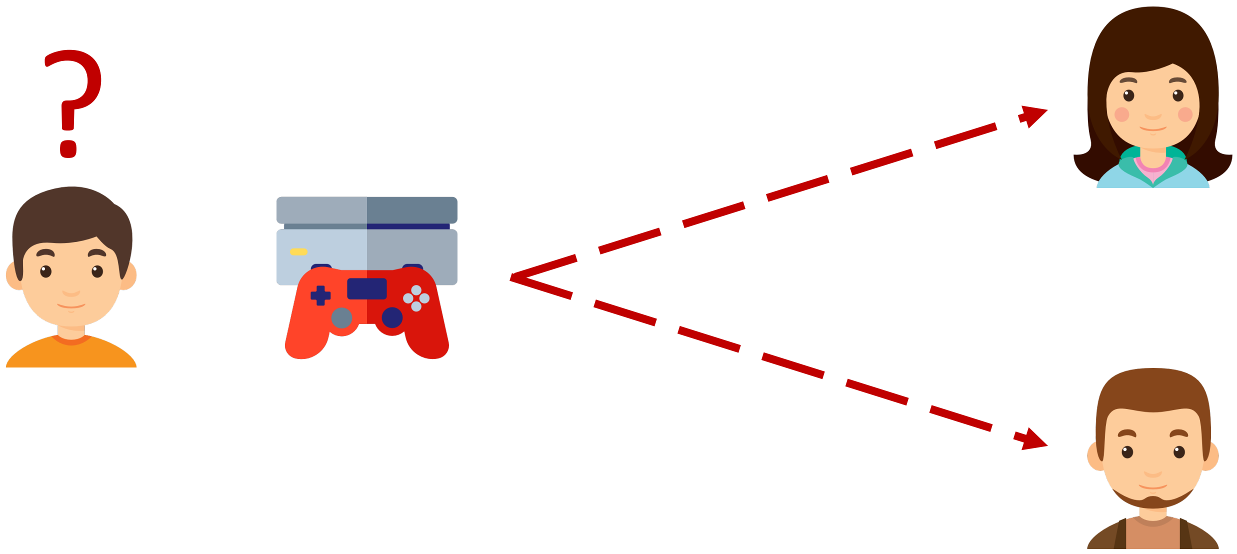

Image Courtesy: Freepik

Real-World Applications

- Auctions form a significant part of mechanism design with money
- Auctions are ubiquitous in the real world!
  - A significant source of revenue for many large organizations (including Facebook and Google)
  - Often run billions of tiny auctions everyday
  - Need the algorithms to be fast

Example: Facility Location

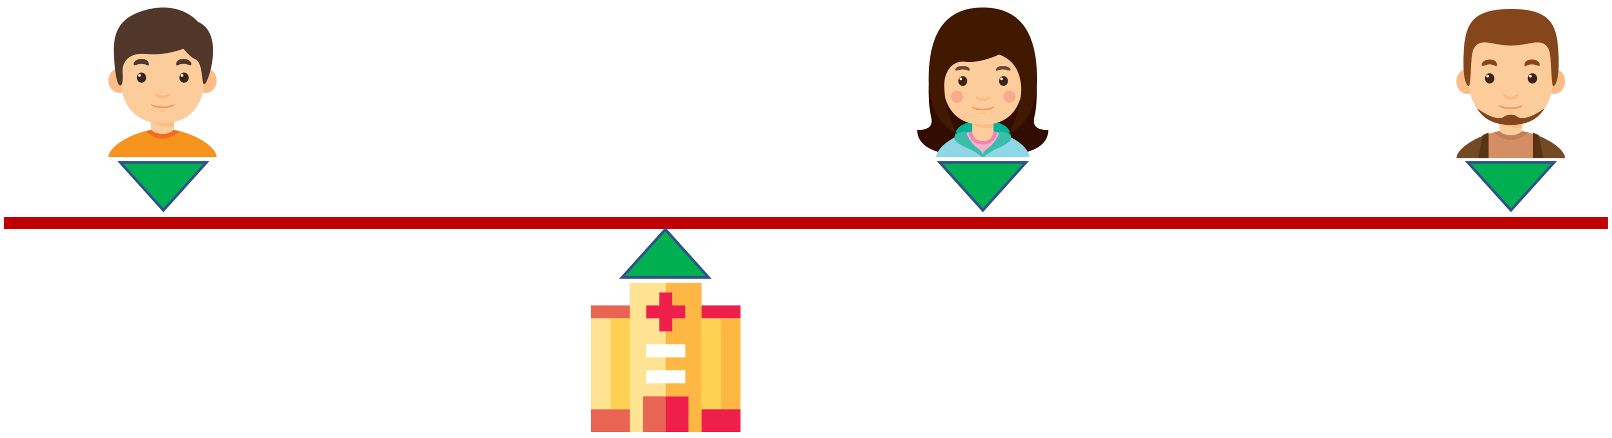

Cost to each agent: Distance from the hospital

Objective: Minimize the sum of costs

Constraint: No money

Image Courtesy: Freepik

Example: Facility Location

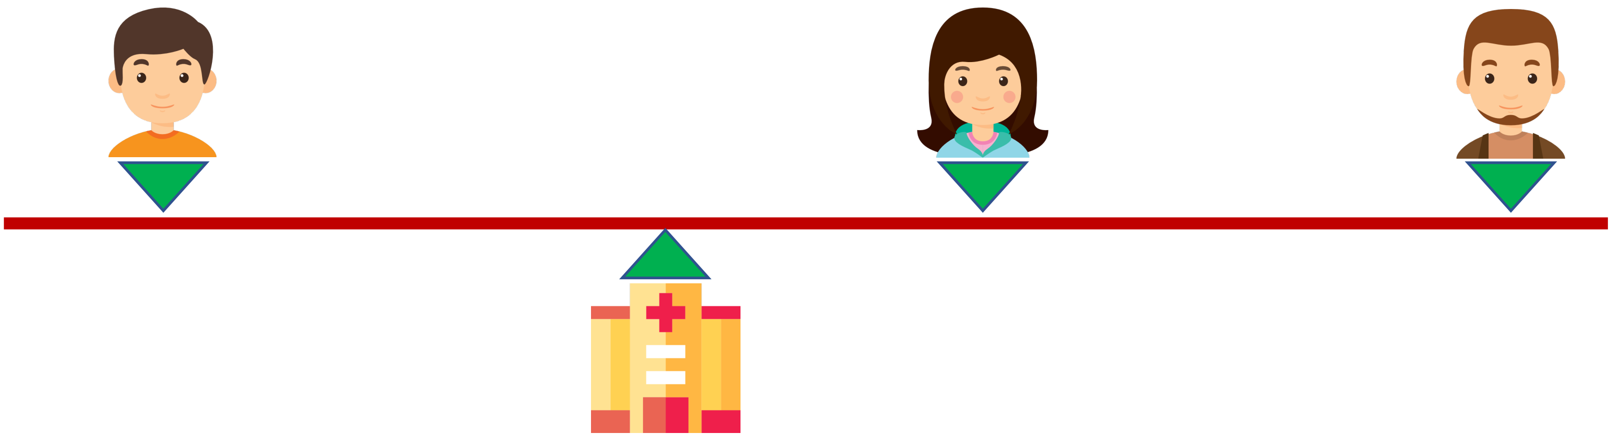

Q: What is the optimal hospital location?

Q: If we decide to choose the optimal location, will the agents really tell us where they live?

Image Courtesy: Freepik

Example: Facility Location

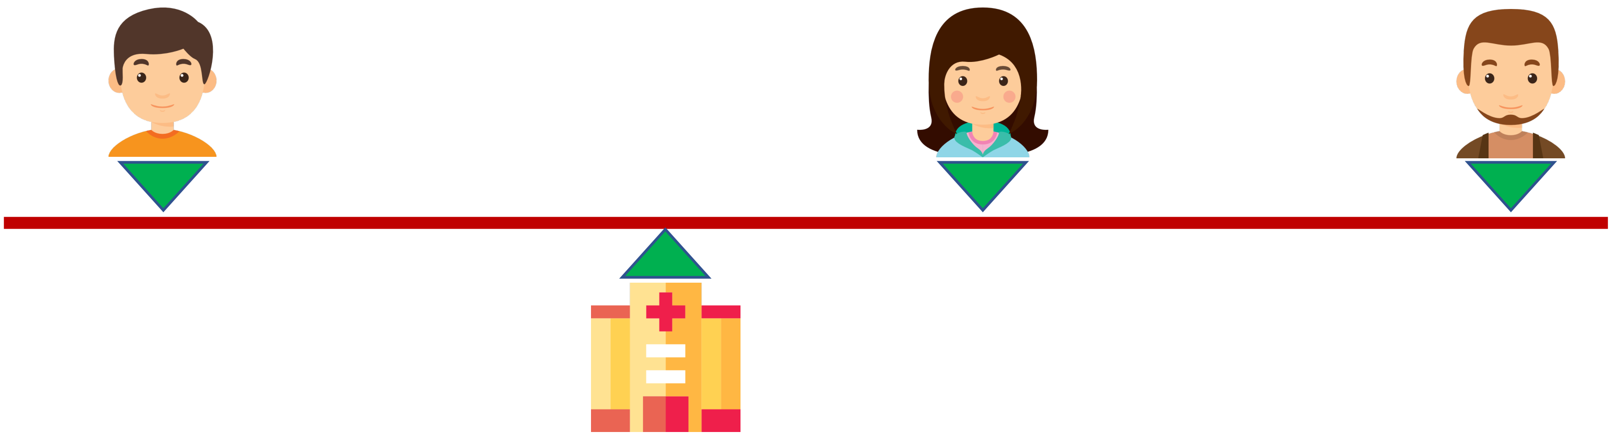

Cost to each agent: Distance from the hospital

Objective: Minimize the maximum cost

Constraint: No money

Image Courtesy: Freepik

Example: Facility Location

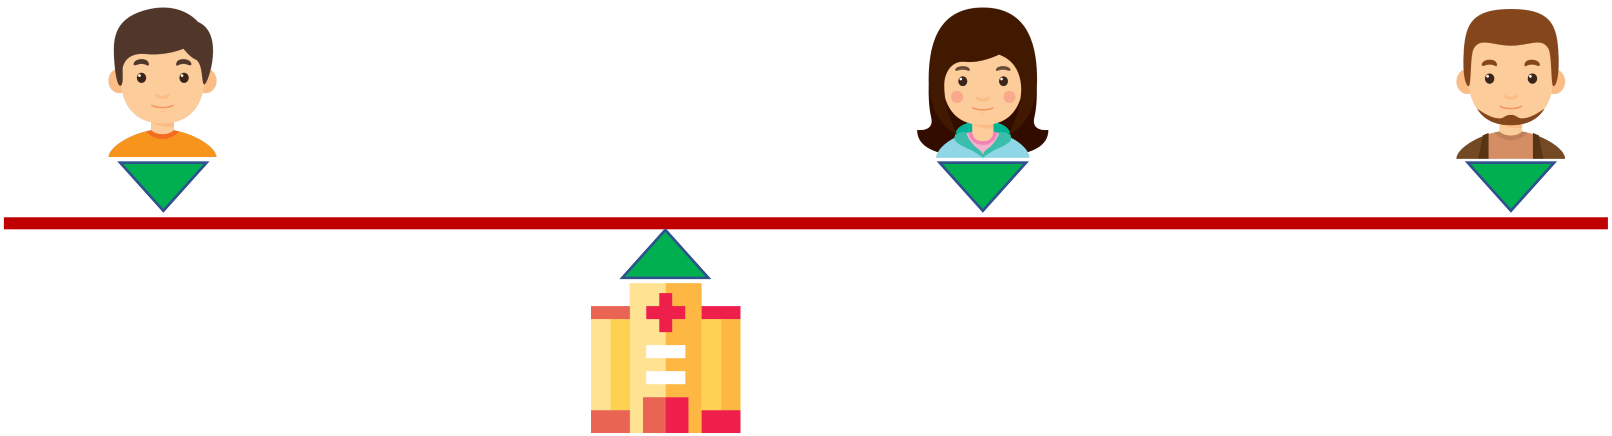

Q: What is the optimal hospital location?

Q: If we decide to choose the optimal location, will the agents really tell us where they live?

Image Courtesy: Freepik

Real-World Applications

Roth

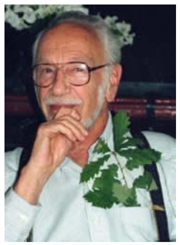

Gale

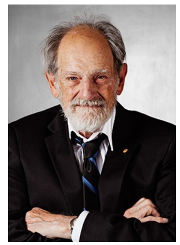

Shapley

Matching

- National Resident Matching Program (NRMP)
- School Choice (New York, Boston)

Fair Division

splíddít

Voting

ROBOVOTE

Voting Theory

Social Choice Theory

- Mathematical theory for aggregating individual preferences into collective decisions

Voting Theory

- Originated in ancient Greece
- Formal foundations
  - 13th Century (Ramon Llull)
  - 18th Century (Marquis de Condorcet and Jean-Charles de Borda)
  - 19th Century: Charles Dodgson (a.k.a. Lewis Carroll)
  - 20th Century: Nobel prizes to Kenneth Arrow and Amartya Sen




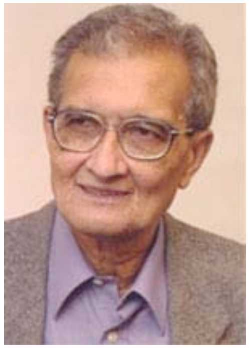

Voting Theory

- We want to select a collective decision based on (possibly different) individual preferences
  - Presidential election, restaurant/movie selection for group activity, committee selection, facility location, …
- Resource allocation is a special case
  - You can think of all possible allocations as the different "outcomes"
    - A very restricted case due to lots of ties
    - An agent is indifferent among all allocations in which the resources *she* gets are the same
  - We want to study the general case

Voting Framework

- Set of voters $N = \{1, \dots, n\}$
- Set of alternatives A, $|A| = m$
- Voter i has a preference ranking $\succ_i$ over the alternatives
- Preference profile $\overrightarrow{\succ}$ is the collection of all voters' rankings

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

Voting Framework

- Social choice function f
  - Takes as input a preference profile $\vec{\succ}$
  - Returns an alternative $a \in A$

- Social welfare function f
  - Takes as input a preference profile $\vec{\succ}$
  - Returns a societal preference $\succ^*$

- For now, voting rule = social choice function

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

Voting Rules

- Plurality
  - Each voter awards one point to her top alternative
  - Alternative with the most point wins
  - Most frequently used voting rule
  - Almost all political elections use plurality

- Problem?

1	2	3	4	5
a	a	a	b	b
b	b	b	c	c
c	c	c	d	d
d	d	d	e	e
e	e	e	a	a

Winner
a

Voting Rules

- Borda Count
  - Each voter awards $m - k$ points to alternative at rank k
  - Alternative with the most points wins
  - Proposed by Ramon Llull in the 13th Century but named after 18th Century work by Jean-Charles de Borda
  - Used for elections to the national assembly of Slovenia

1	2	3
a (2)	c (2)	b (2)
b (1)	a (1)	a (1)
c (0)	b (0)	c (0)

Total
a: 2+1+1 = 4
b: 1+0+2 = 3
c: 0+2+0 = 2

Winner
a

Current uses [edit]

Political uses [edit]

The Borda count is used for certain political elections in at least three countries, Slovenia and the tiny Micronesian nations of Kiribati and Nauru. In Slovenia, the Borda count is used to elect two of the ninety members of the National Assembly: one member represents a constituency of ethnic Italians, the other a constituency of the Hungarian minority. As noted above, members of the Parliament of Nauru are elected based on a variant of the Borda count that involves two departures from the normal practice: (1) multi-seat constituencies, of either two or four seats, and (2) a point-allocation formula that involves increasingly small fractions of points for each ranking, rather than whole points. In Kiribati, the president (or *Beretitenti*) is elected by the plurality system, but a variant of the Borda count is used to select either three or four candidates to stand in the election. The constituency consists of members of the legislature (*Maneaba*). Voters in the legislature rank only four candidates, with all other candidates receiving zero points. Since at least 1991, tactical voting has been an important feature of the nominating process.

The Republic of Nauru became independent from Australia in 1968. Before independence, and for three years afterwards, Nauru used instant-runoff voting, importing the system from Australia, but since 1971, a variant of the Borda count has been used.

The modified Borda count has been used by the Green Party of Ireland to elect its chairperson.[6][7]

The Borda count has been used for non-governmental purposes at certain peace conferences in Northern Ireland, where it has been used to help achieve consensus between participants including members of Sinn Féin, the Ulster Unionists, and the political wing of the UDA.

Other uses [edit]

The Borda count is used in elections by some educational institutions in the United States.

- University of Michigan
  - Central Student Government
  - Student Government of the College of Literature, Science and the Arts (LSASG)
- University of Missouri: officers of the Graduate-Professional Council
- University of California Los Angeles: officers of the Graduate Student Association
- Harvard University: officers of the Civil Liberties Union
- Southern Illinois University at Carbondale: officers of the Faculty Senate,
- Arizona State University: officers of the Department of Mathematics and Statistics assembly.
- Wheaton College, Massachusetts: faculty members of committees.
- College of William and Mary: members of the faculty personnel committee of the School of Business Administration (tie-breaker).

The Borda count is used in elections by some professional and technical societies.

- International Society for Cryobiology: Board of Governors.
- Tempo sustainable design network: management committee.
- U.S. Wheat and Barley Scab Initiative: members of Research Area Committees.
- X.Org Foundation: Board of Directors.

The OpenGL Architecture Review Board uses the Borda count as one of the feature-selection methods.

The Borda count is used to determine winners for Toastmasters International speech contests. Judges offer a ranking of their top three speakers, awarding them three points, two points, and one point, respectively. All unranked candidates receive zero points.

The modified Borda count is used to elect the President for the United States member committee of AIESEC.

The Borda count, and points-based systems similar to it, are often used to determine awards in competitions.

The Borda count is a popular method for granting sports awards in the United States. Uses include:

- MLB Most Valuable Player Award (baseball)
- Heisman Trophy (college football)[8]
- Ranking of NCAA college teams

The Eurovision Song Contest uses a positional voting method similar to the Borda count, with a different distribution of points: only the top ten entries are considered in each ballot, the favorite entry receiving 12 points, the second-placed entry receiving 10 points, and the other eight entries getting points from 8 to 1. Although designed to favor a clear winner, it has produced very close races and even a tie.

The People's Remix Competition uses a Borda variant where each voter ranks only the top three contestants.

The Borda count is used for wine trophy judging by the Australian Society of Viticulture and Oenology, and by the RoboCup autonomous robot soccer competition at the Center for Computing Technologies, in the University of Bremen in Germany.

The Finnish Associations Act lists three different modifications of the Borda count for holding a proportional election. All the modifications use fractions, as in Nauru. A Finnish association may choose to use other methods of election, as well.[9]

Borda count in real life

Voting Rules

- Positional Scoring Rules
  - Defined by a score vector $\vec{s} = (s_1, \dots, s_m)$
  - Each voter gives s_k points to alternative at rank k

- A family containing many important rules
  - Plurality = $(1,0, \dots, 0)$
  - Borda = $(m-1, m-2, \dots, 0)$
  - k-approval = $(1, \dots, 1,0, \dots, 0)$ ← top k get 1 point each
  - Veto = $(0, \dots, 0, -1)$
  - …

Voting Rules

- Plurality with runoff
  - **First round:** two alternatives with the highest plurality scores survive
  - **Second round:** between these two alternatives, select the one that majority of voters prefer

- Similar to the French presidential election system
  - Problem: vote division
  - Happened in the 2002 French presidential election

Voting Rules

- Single Transferable Vote (STV)
  - $m - 1$ rounds
  - In each round, the alternative with the least plurality votes is eliminated
  - Alternative left standing is the winner
  - Used in Ireland, Malta, Australia, New Zealand, …

- STV has been strongly advocated for due to various reasons

STV Example

2 voters	2 voters	1 voter
a	b	c
b	a	d
c	d	b
d	c	a

→

2 voters	2 voters	1 voter
a	b	c
b	a	b
c	c	a

↓

2 voters	2 voters	1 voter
a	b	b
b	a	a

←

2 voters	2 voters	1 voter
b	b	b

Voting Rules

- Kemeny's Rule
  - Social welfare function (selects a ranking)
  - Let $n_{a \succ b}$ be the number of voters who prefer a to b
  - Select a ranking σ of alternatives = for every pair (a, b) where $a \succ_\sigma b$, we make $n_{b \succ a}$ voters unhappy
  - Total unhappiness $K(\sigma) = \sum_{(a,b):a \succ_\sigma b} n_{b \succ a}$
  - Select the ranking σ^* with minimum total unhappiness
- Social choice function
  - Choose the top alternative in the Kemeny ranking

Kemeny Example

2 voters	2 voters	1 voter
a	b	c
b	a	d
c	d	b
d	c	a

- $K(a \succ b \succ c \succ d)$
  - 0 x first 2 voters
  - 2 x next 2 voters
  - 5 x last voter
  - 9 in total

- $K(b \succ a \succ c \succ d)$
  - 1 x first 2 voters
  - 1 x next 2 voters
  - 4 x last voter
  - 8 in total

Condorcet Winner

- Definition
  - Alternative x defeats y in a pairwise election if a *strict* majority of voters prefer x to y
  - Alternative x is a Condorcet winner if it defeats every other alternative in a pairwise election
- Question
  - Can there be two Condorcet winners?
- Condorcet paradox
  - No Condorcet winner when the majority preference is cyclic

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

Majority Preference

$a \succ b$

$b \succ c$

$c \succ a$

Condorcet Consistency

- Condorcet consistency
  - A voting rule is Condorcet consistent if it selects the Condorcet winner whenever one exists
  - On preference profiles where there is no Condorcet winner, it is free to output any winner

- Among the rules we saw so far…
  - ***NOT* Condorcet consistent:** all positional scoring rules (plurality, Borda, …), plurality with runoff, STV
  - **Condorcet consistent:** Kemeny (Why?)

Majority Consistency

- Majority consistency
  - If a strict majority of voters rank alternative x first, then x must be the winner.

- Question: What is the relation between majority consistency and Condorcet consistency?
  1. Majority consistency $\Rightarrow$ Condorcet consistency
  2. Condorcet consistency $\Rightarrow$ Majority consistency
  3. Equivalent
  4. Incomparable

Condorcet Consistency

- Copeland
  - Score(x) = # alternatives x beats in pairwise elections
  - Select x^* with the maximum score
  - Condorcet consistent (Why?)

- Maximin
  - Score(x) = $\min_y n_{x \succ y}$
  - Select x^* with the maximum score
  - Also Condorcet consistent (Why?)