Web-Mining Agents
Game Theory and Social Choice

Özgür L. Özcep
Universität zu Lübeck
Institut für Informationssysteme

Literature

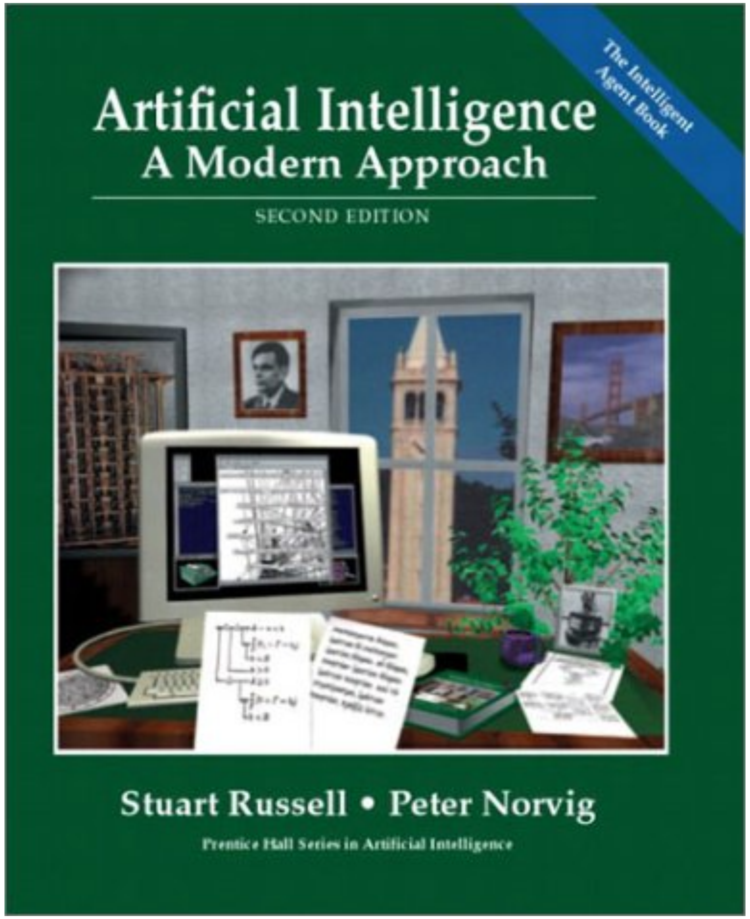

Chapter 17

Presentations from CS 886
Advanced Topics in
AI Electronic Market Design
Kate Larson
Waterloo Univ.

Full vs bounded rationality

Full rationality

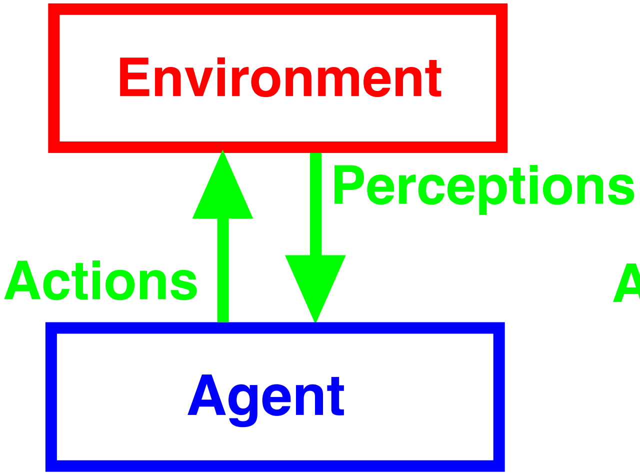

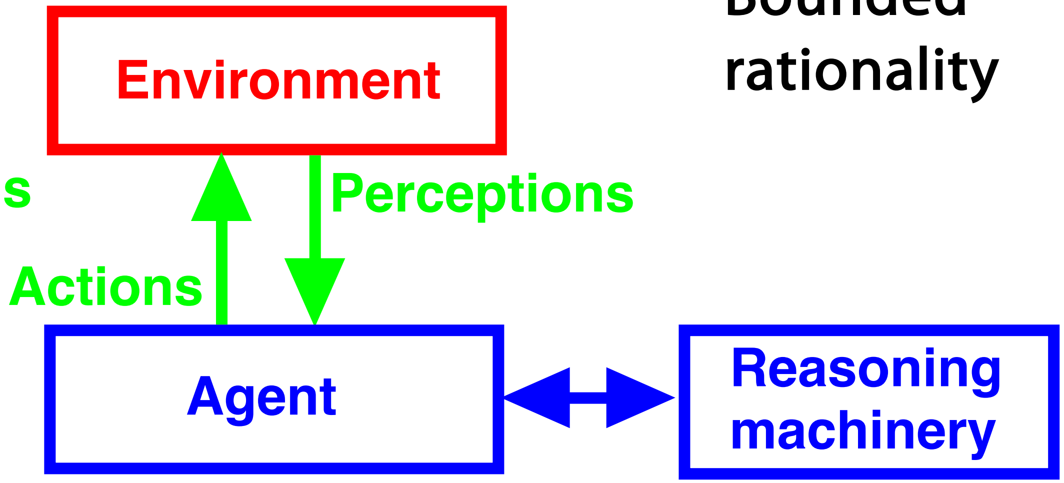

Bounded rationality

Descriptive vs. prescriptive theories of bounded rationality

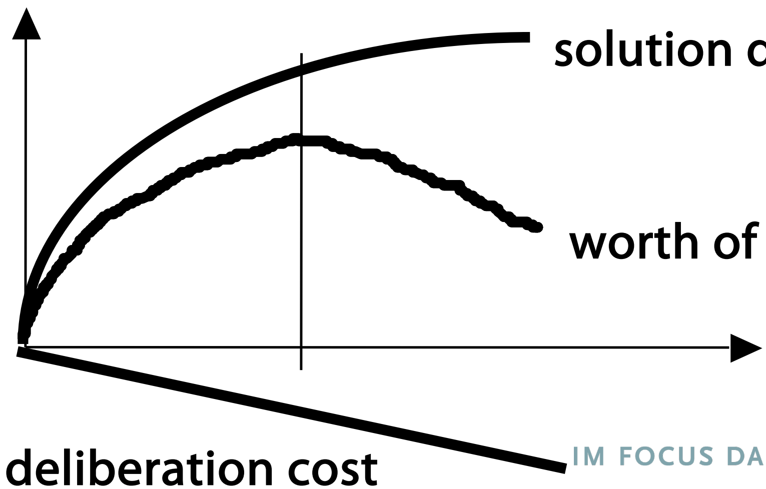

Multiagent Systems: Criteria

- Social welfare: $\max_{outcome} \sum_i u_i(outcome)$
- Surplus: social welfare of outcome – social welfare of status quo
  - Constant sum games have 0 surplus.
  - Markets are not constant sum
- Pareto efficiency: An outcome o is Pareto efficient if there exists no other outcome o′ s.t. some agent has higher utility in o′ than in o and no agent has lower
  - Implied by social welfare maximization
- Individual rationality: Participating in the negotiation (or individual deal) is no worse than not participating
- Stability: No agents can increase their utility by changing their strategies (aka policies)
- Symmetry: No agent should be inherently preferred, e.g. dictator

Game Theory: The Basics

- **A game:** Formal representation of a situation of strategic interdependence
  - Set of agents, I (|I|=n)
    - Aka players
  - Each agent, j, has a set of actions, A_j
    - Aka moves
  - Actions define outcomes
    - For each possible action there is an outcome.
  - Outcomes define payoffs
    - Agents' derive utility from different outcomes

Normal form game* (matching pennies)

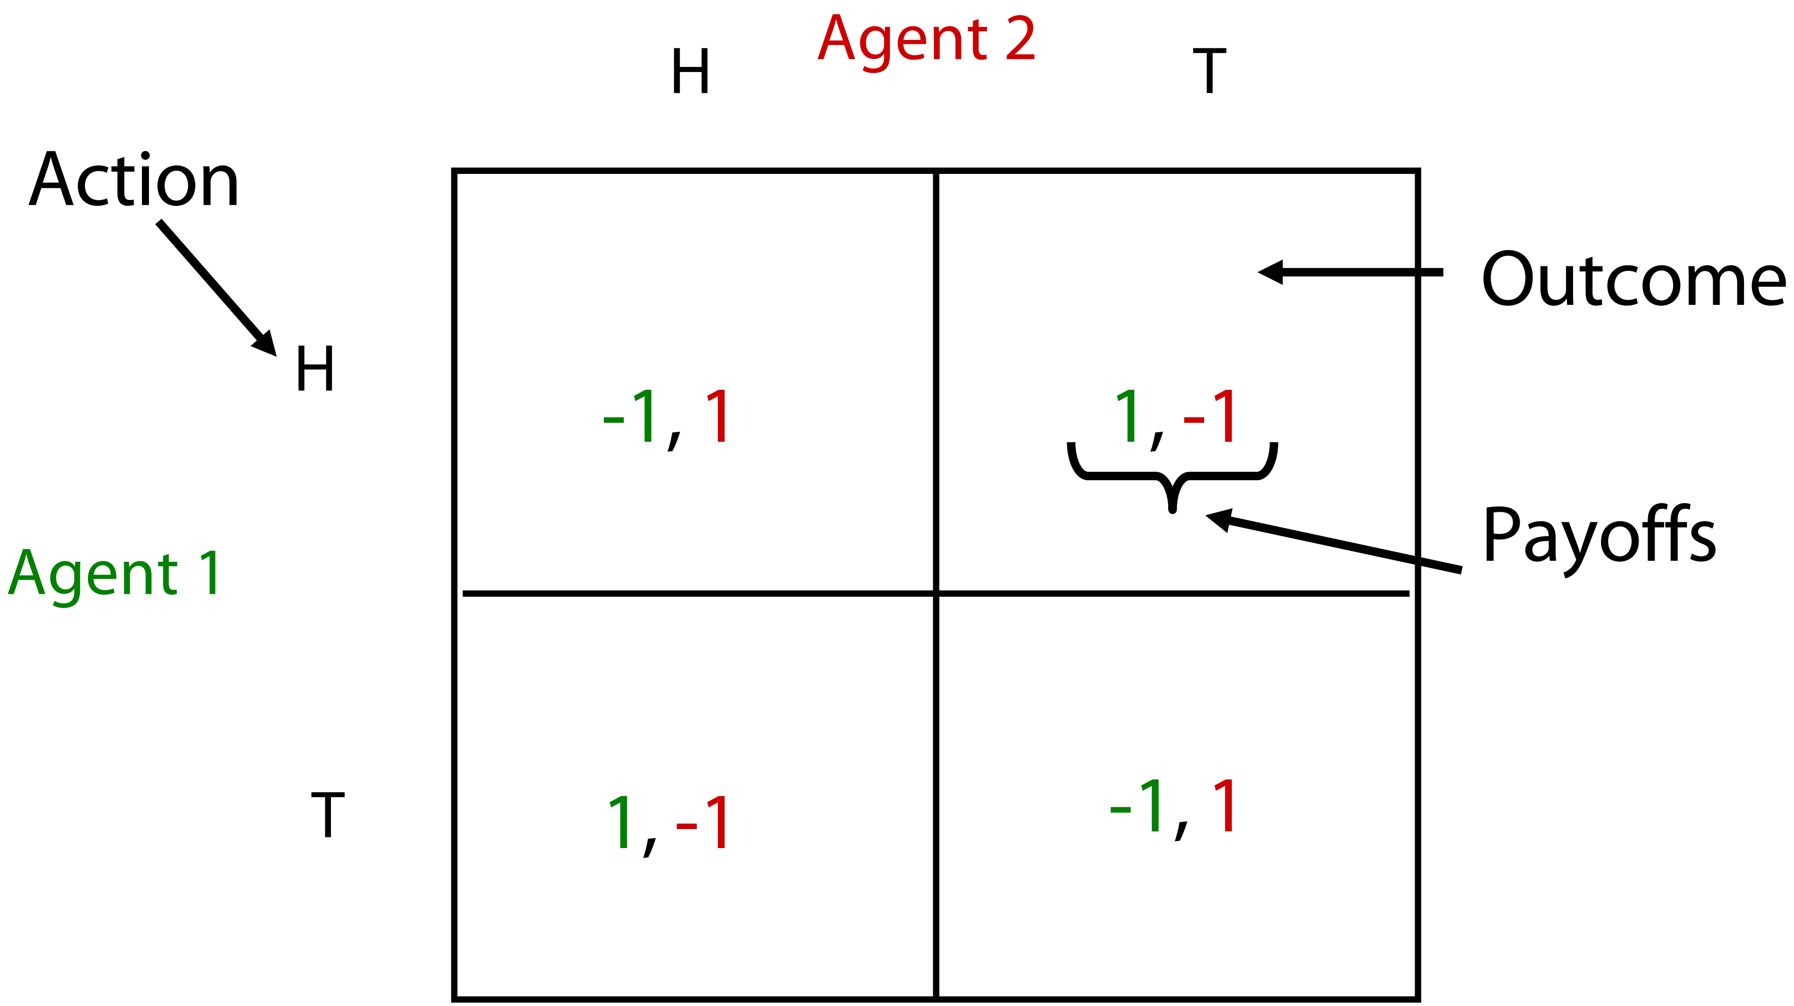

*aka strategic form, matrix form

Extensive form game (matching pennies)

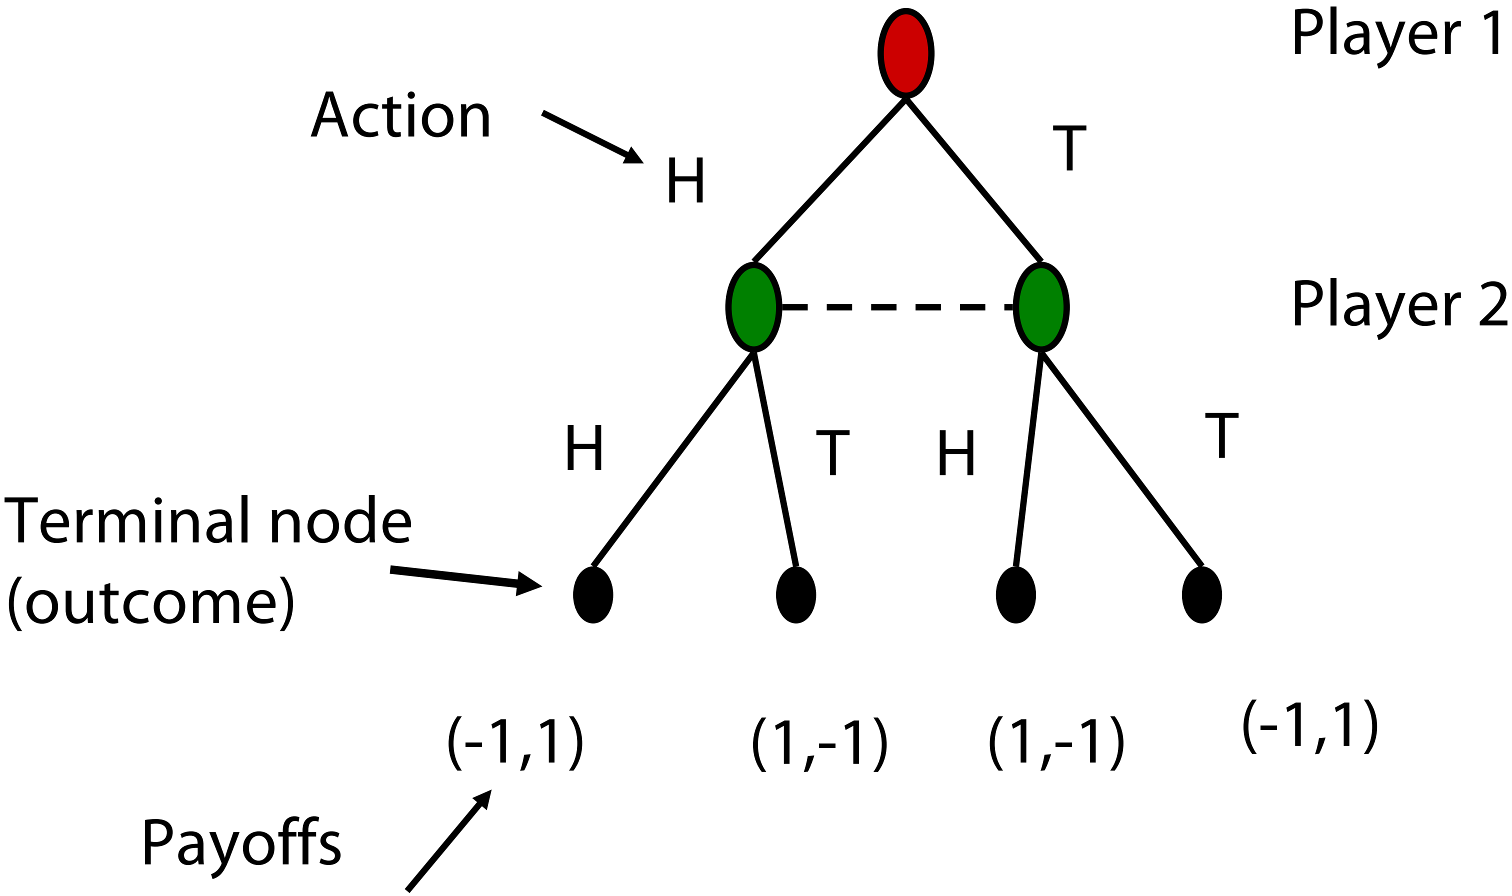

Strategies (aka Policies)

- Strategy:
  - A strategy, s_j, is a **complete contingency plan**; defines actions agent j should take for all possible states of the world
- Strategy profile: $s=(s_1,\ldots,s_n)$
  - $s_{-i} = (s_1,\ldots,s_{i-1},s_{i+1},\ldots,s_n)$
- Utility function: $u_i(s)$
  - Note that the utility of an agent depends on the strategy profile, not just its own strategy
  - We assume agents are **expected utility maximizers**

Normal form game*
(matching pennies)

Agent 2

	H	T
Agent 1: H	-1, 1	1, -1
Agent 1: T	1, -1	-1, 1

Strategy for agent 1: H

Strategy profile (H,T)

U1((H,T))=1
U2((H,T))=-1

*aka strategic form, matrix form

Extensive form game
(matching pennies)

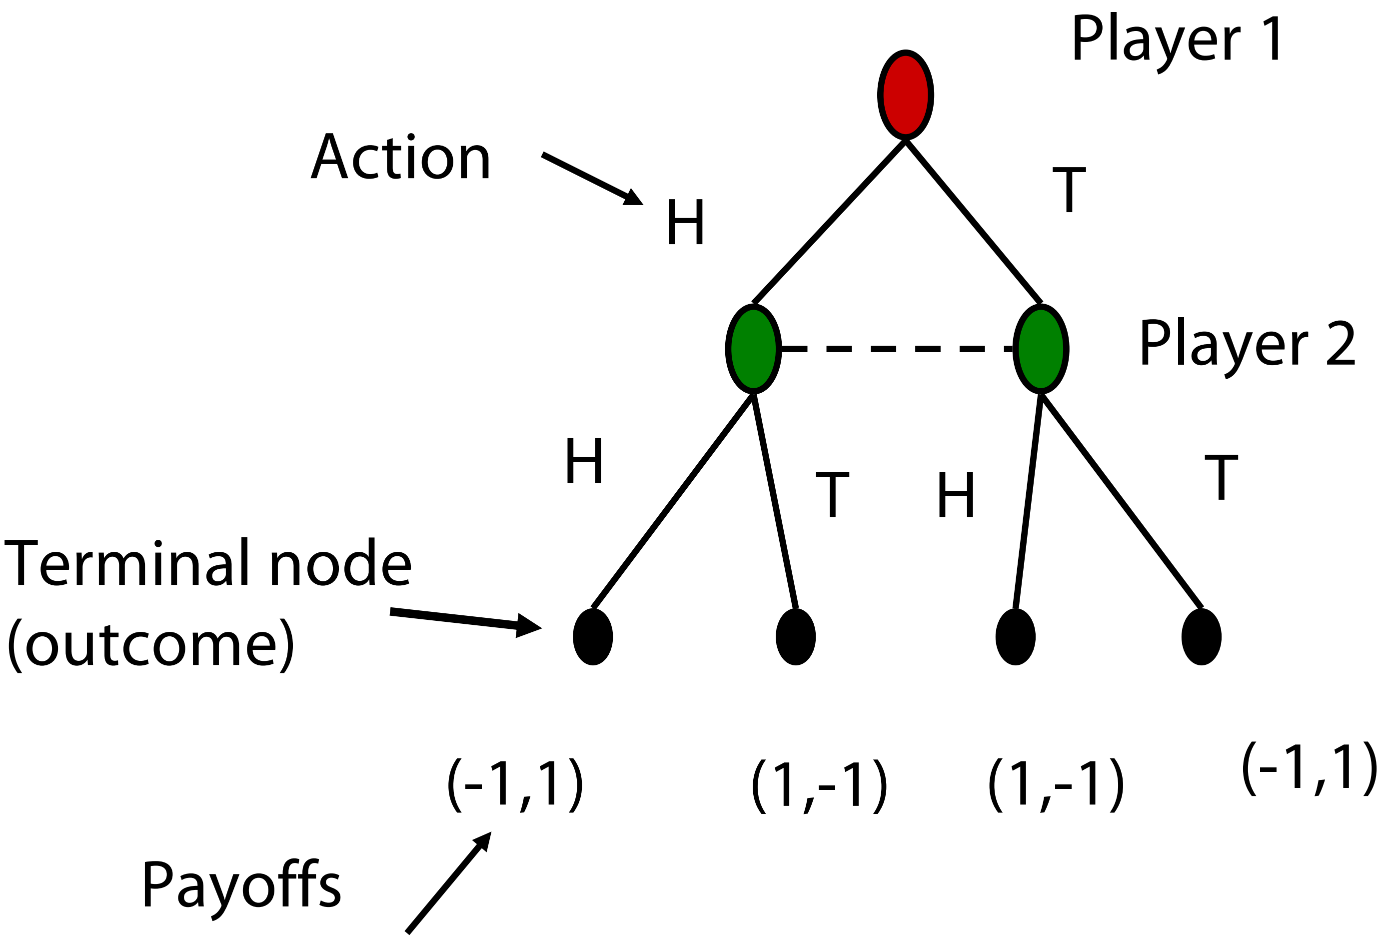

Strategy for agent 1: T

Strategy profile: (T,T)

U1((T,T))=-1

U2((T,T))=1

Extensive form game
(matching pennies, seq moves)

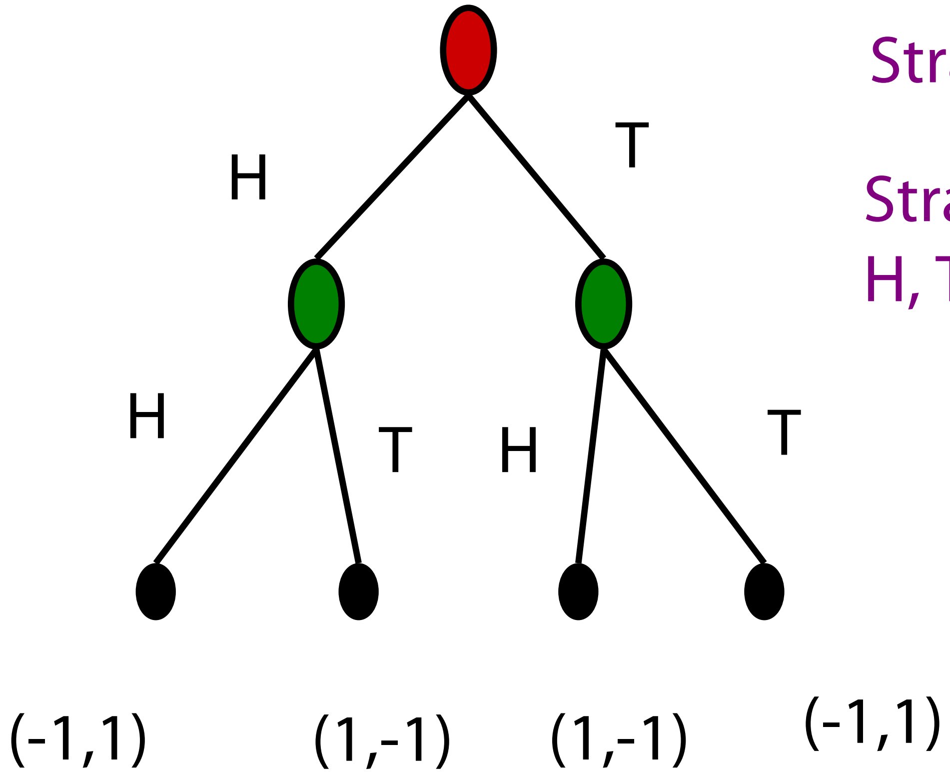

Recall: A strategy is a contingency plan for all states of the game

Strategy for agent 1: T

Strategy for agent 2: H if 1 plays H, T if 1 plays T (H,T)

Strategy profile: (T,(H,T))

U1((T,(H,T)))=-1

U2((T,(H,T)))=1

Game Representation

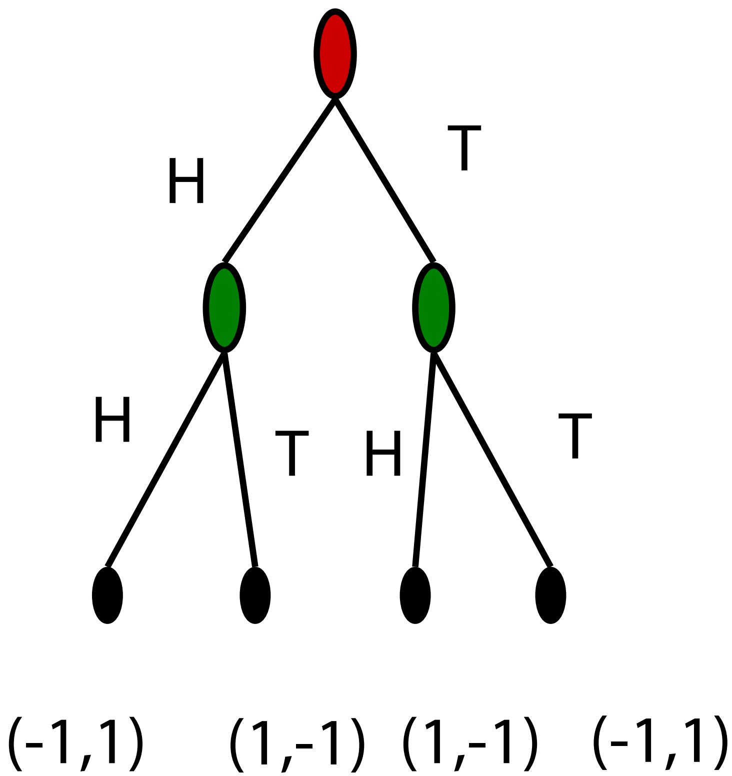

	H,H	H,T	T,H	T,T
H	-1,1	-1,1	1,-1	1,-1
T	1,-1	-1,1	1,-1	-1,1

Potential combinatorial explosion

Example: Ascending Auction

- State of the world is defined by (x,p)
  - $x \in \{0,1\}$ indicates if the agent has the object
  - p is the current next price
- Strategy $s_i((x,p))$

$$s_i((x,p)) = \begin{cases} p, \text{ if } v_i \geqq p \text{ and } x=0 \\ \text{No bid otherwise} \end{cases}$$

(v_i is the value agent i ascribes to the object

Dominant Strategies

- Recall that
  - Agents' utilities depend on what strategies other agents are playing
  - Agents are expected utility maximizers
- Agents will play best-response strategies

s_i^* is a best response if $u_i(s_i^*, s_{-i}) \geq u_i(s_i', s_{-i})$ for all s_i'

- A dominant strategy is a best-response for all s_{-i}
  - They do not always exist
  - Inferior strategies are called dominated

Dominant Strategy Equilibrium

- A dominant strategy equilibrium is a strategy profile where the strategy for each player is dominant
  - $s^*=(s_1^*,\dots,s_n^*)$
  - $u_i(s_i^*,s_{-i})\geq u_i(s_i',s_{-i})$ for all i, for all s_i', for all s_{-i}
- **GOOD**: Agents do not need to counterspeculate!

Example: Prisoner's Dilemma

Two people are arrested for a crime.

- If neither suspect confesses, both are released (ÖÖ: but sentenced semi-heavy).
- If both confess then they get sent to jail.
- If one confesses and the other does not, then the confessor gets a light sentence and the other gets a heavy sentence.

	A: Confess	A: Don't Confess
B:Confess	B=-5, A=-5	B=-1, A=-10
B:Don't Confess	B=-10, A=-1	B=-2, A=-2

Dom. Str. Eq

Pareto Optimal Outcome

Dominant strategy exists but is not Pareto efficient

Example: Split or Steal

Does communication help?

Only if agents do not lie

		A: Steal	A: Split	
Dom. Str. Eq	B:Steal	B=0, A=0	B=100, A=-10	
	B:Split	B=-10, A=100	B=50, A=50	Pareto Optimal Outcome

ÖÖ: Example from British Game Show „Golden Balls“
See http://blogs.cornell.edu/info2040/2012/09/21/split-or-steal-an-analysis-using-game-theory/
And may be...
https://www.youtube.com/watch?v=p3Uos2fzIJ0

Vickrey *) Auctions

- Vickrey auctions are:
  - *second-price*
  - *sealed-bid*
- Good is awarded to the agent that made the highest bid; at the price of the *second highest* bid
- *Bidding to your true valuation is dominant strategy in Vickrey auctions*
- Vickrey auctions susceptible to *antisocial* behavior

*) Russel/Norvig add in a FN:
Named after William Vickrey (1914–1996), who won the 1996 Nobel Prize in economics for this work and died of a heart attack three days later

Example: Vickrey Auction (2nd price sealed bid)

- Each agent i has value v_i
- Strategy $b_i(v_i) \in [0,\infty)$
- $b^* :=$ 2nd best bid.

$$u_i(b_i, b_{-i}) = \begin{cases} v_i - b^* & \text{if } b_i > b^* \\ 0 & \text{otherwise} \end{cases}$$

Given value v_i, $b_i(v_i)=v_i$ is dominant.

Let $b' = \max_{j \neq i} b_j$. If $b' < v_i$ then any bid $b_i(v_i) \geq b'$ is optimal. If $b' \geq v_i$, then any bid $b_i(v_i) \leq v_i$ is optimal. Bid $b_i(v_i)=v_i$ satisfies both constraints.

Dominant strategy is Pareto efficient

Phone Call Competition Example

- Customer wishes to place long-distance call
- Carriers simultaneously bid, sending proposed prices
- Phone automatically chooses the carrier (dynamically)

Best Bid Wins

- Phone chooses carrier with lowest bid
- Carrier gets amount that it bid

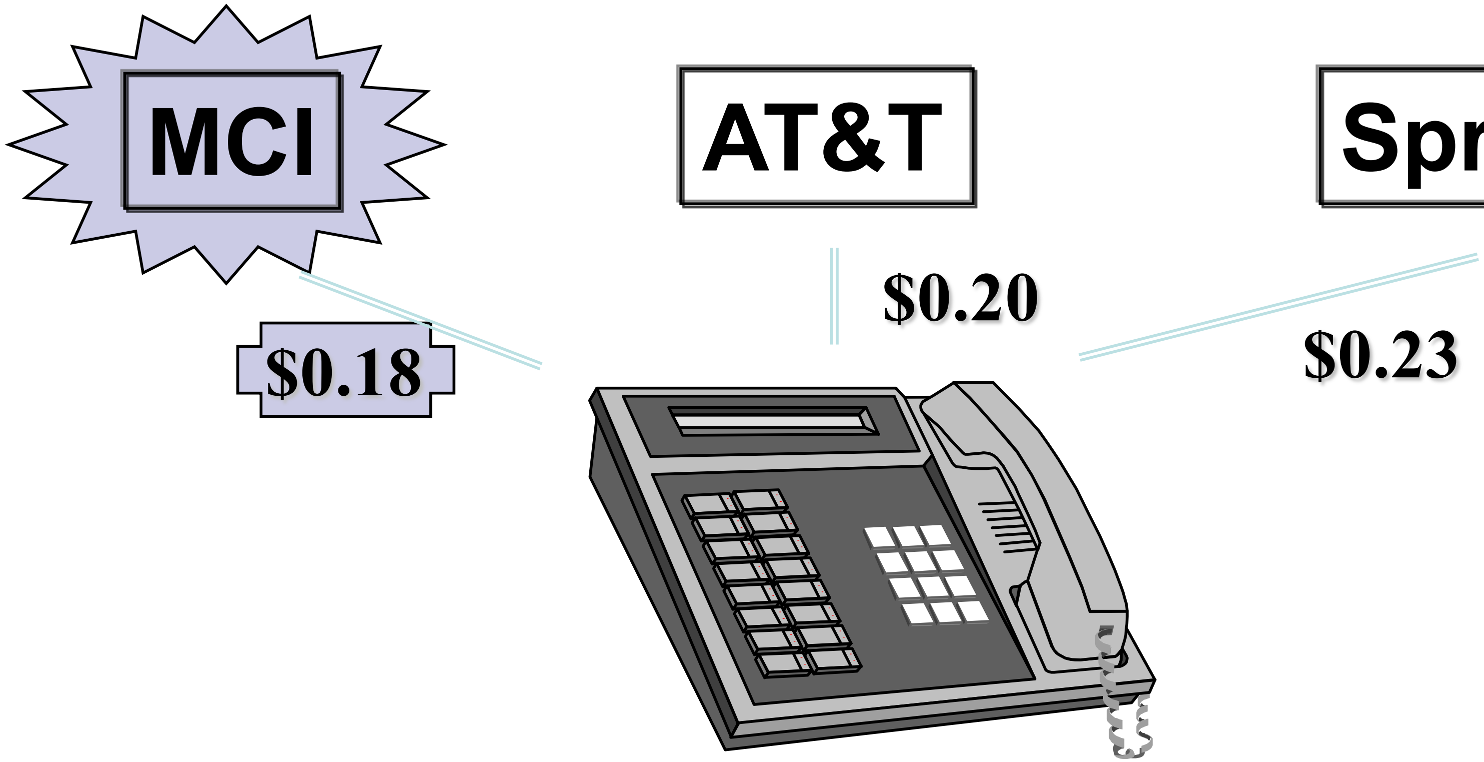

Attributes of the Mechanism

- ✓ *Distributed*
- ✓ *Symmetric*
- ✗ *Stable*
- ✗ *Simple*
- ✗ *Efficient*

Carriers have an incentive to invest effort in strategic behavior

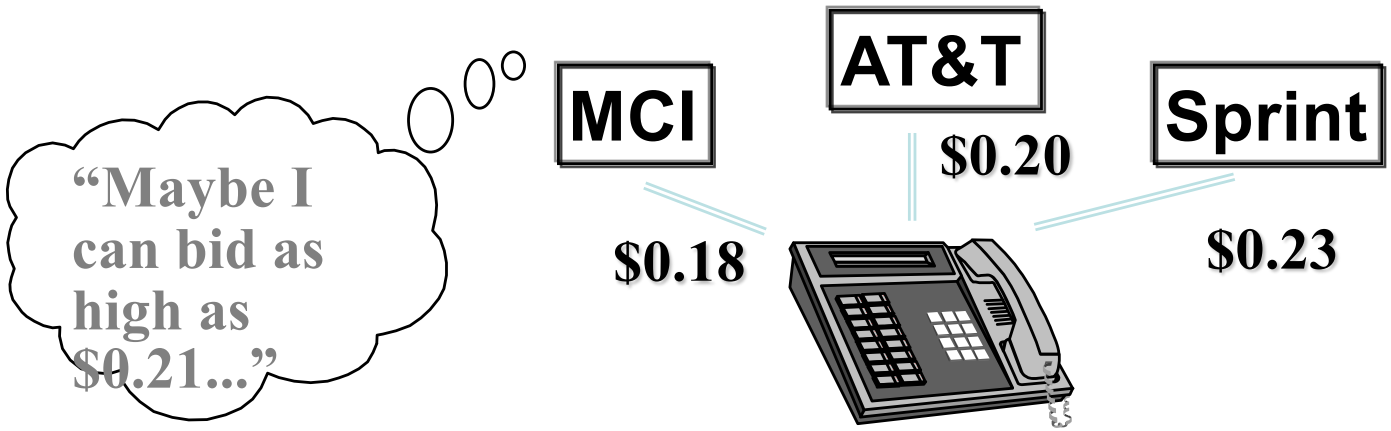

Best Bid Wins, Gets Second Price (Vickrey Auction)

- Phone chooses carrier with lowest bid
- Carrier gets amount of second-best price

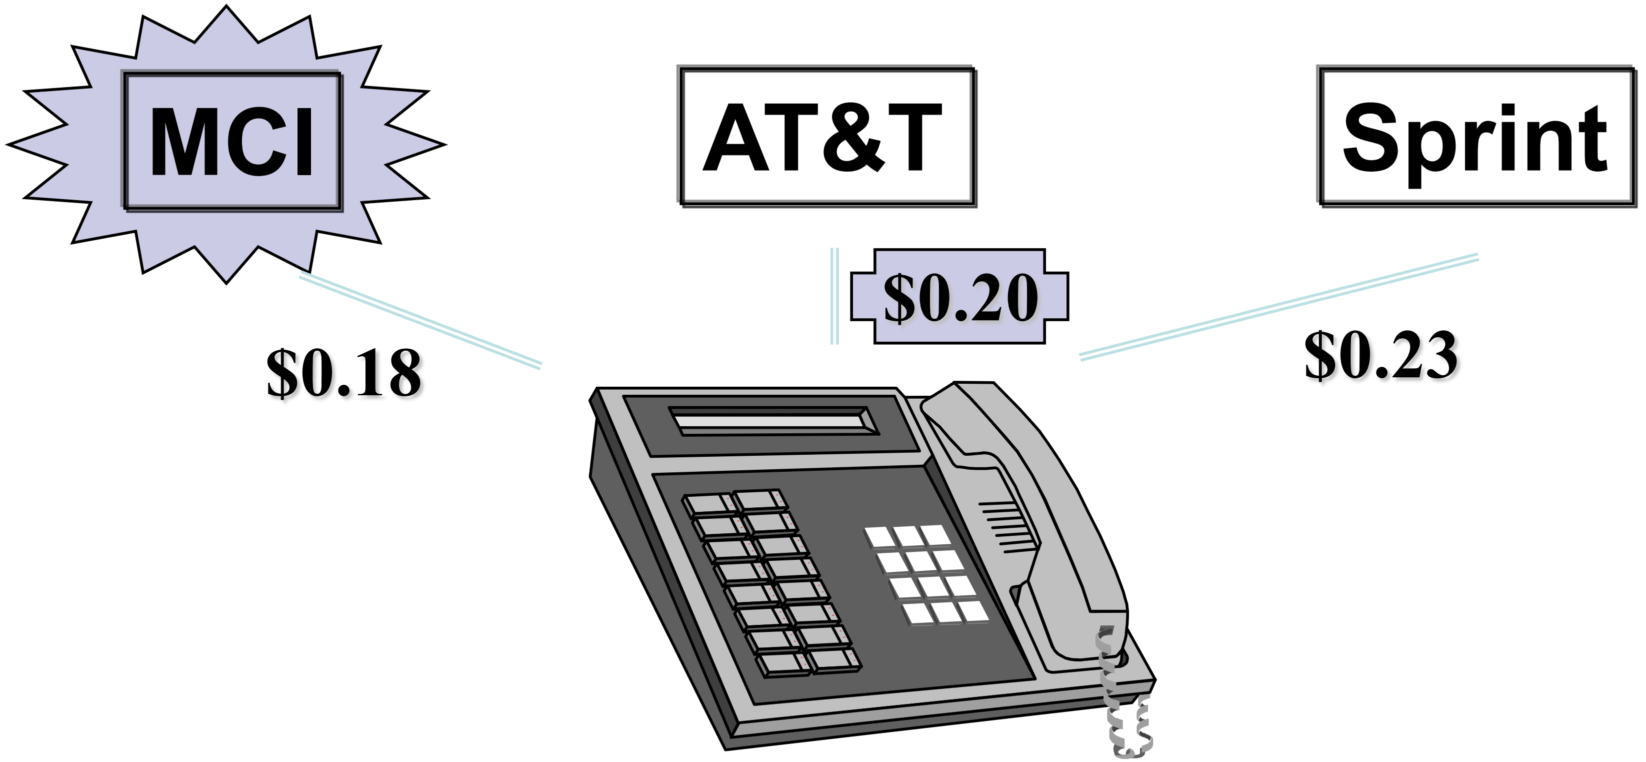

Attributes of the Vickrey Mechanism

- ✓ *Distributed*
- ✓ *Symmetric*
- ✓ *Stable*
- ✓ *Simple*
- ✓ *Efficient*

Carriers have *no* incentive to invest effort in strategic behavior

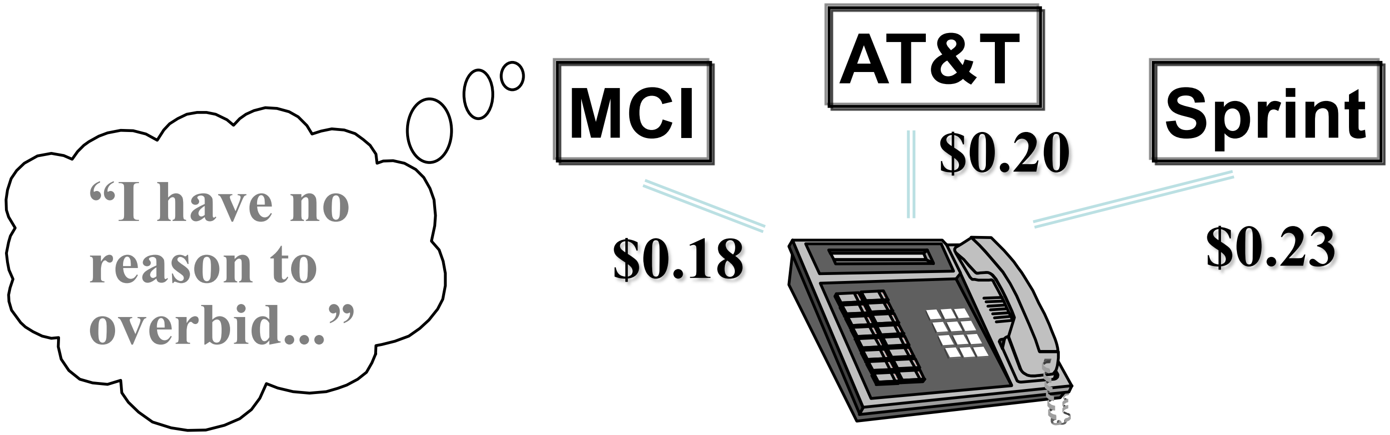

Example: Bach or Stravinsky

- A couple likes going to concerts together. One loves Bach but not Stravinsky. The other loves Stravinsky but not Bach. However, they prefer being together than being apart.

	B	S
B	2,1	0,0
S	0,0	1,2

No dom. str. equil.

Nash Equilibrium

- Sometimes an agent's best-response depends on the strategies other agents are playing
  - No dominant strategy equilibria
- A strategy profile is a **Nash equilibrium** if no player has incentive to deviate from his strategy given that others do not deviate:
  - for every agent i, $u_i(s_i^*, s_{-i}^*) \geq u_i(s_i', s_{-i}^*)$ for all s_i'

	B	S
B	2,1	0,0
S	0,0	1,2

Nash Equilibrium

- Interpretations:
  - Focal points, self-enforcing agreements, stable social convention, consequence of rational inference..
- Criticisms
  - They may not be unique (Bach or Stravinsky)
    - Ways of overcoming this
      - Refinements of equilibrium concept, Mediation, Learning
  - Do not exist in all games (in the form defined above)
  - They may be hard to find
  - People don't always behave based on what equilibria would predict (ultimatum games and notions of fairness,…)

Example: Matching Pennies

	H	T
H	-1, 1	1,-1
T	1,-1	-1, 1

So far we have talked only about **pure** (deterministic) strategy equilibria.

Not all games have pure strategy equilibria. Some equilibria are **mixed** (randomzied) strategy equilibria.

Mixed strategy equilibria

- Let Σ_i be the set of probability distributions over S_i
- All possible pure strategy profiles: $S = S_1 \times \cdots \times S_n$
- σ_i in Σ_i
- Strategy profile: $\sigma = (\sigma_1, \ldots, \sigma_n)$
- Expected utility for pure strategy $s_i \in \sigma_i$ for agent i

$$u_i(s_i, \sigma_{-i}) = \sum_{s \in S_{-i}} \left(\prod_{1 \le j \le n, j \ne i} \sigma_j(s_j) \right) u_i(s_i, s)$$

- Expected utility for strategy profile σ:

$$u_i(\sigma) = \sum_{s \in S} \left(\prod_{1 \le j \le n} \sigma_j(s_j) \right) u_i(s)$$

Mixed strategy equilibria

- Nash Equilibrium:
  - σ^* is a (mixed) Nash equilibrium iff

    $u_i(\sigma^*_i, \sigma^*_{-i}) \geq u_i(\sigma_i, \sigma^*_{-i})$ for all $\sigma_i \in \sum_i$, for all i

Example: Matching Pennies

		q H	1-q T	: σ_2
σ_1:	p H	−1, 1	1, −1	
	1-p T	1, −1	−1, 1	

Want to play each strategy with a certain probability so that the competitor is indifferent between its own strategies.

$$u_2(H, \sigma_1) = u_2(T, \sigma_1)$$

1p+(-1)(1-p)=(-1)p+1(1-p) ⇒ p=1/2

q-(1-q)=-q+(1-q) ⇒ q=1/2

Mixed Nash Equilibrium

- **Theorem (Nash 50):**
  - Every game in which the strategy sets, $S_1,\ldots,S_n$ have a finite number of elements has a mixed strategy equilibrium.
- Complexity of finding Nash Equilibria
  - "Together with prime factoring, the complexity of finding a Nash Eq is, in my opinion, the most important concrete open question on the boundary of P today" (Papadimitriou)
  - (Daskalakis, Goldberg/Papadimitriou, 2005): Finding Nash equilibrium is very hard (though not NP complete): PPAD complete (Polynomial Parity Arguments on Directed graphs)

Imperfect Information about Strategies and Payoffs

- So far we have assumed that agents have complete information about each other (including payoffs)
  - Very strong assumption!
- Assume agent i has type $\theta_i \in \Theta_i$, which defines the payoff $u_i(s, \theta_i)$
- Agents have common prior over distribution of types $p(\theta)$
  - Conditional probability $p(\theta_{-i} | \theta_i)$ (obtained by Bayes Rule when possible)

Bayesian–Nash Equilibrium

- Strategy: $\sigma_i(\theta_i)$ is the (mixed) strategy agent i plays if its type is θ_i
- Strategy profile: $\sigma = (\sigma_1, \ldots, \sigma_n)$
- Expected utility:

$$EU_i(\sigma_i(\theta_i), \sigma_{-i}(\quad), \theta_i) = \sum_{\theta_i} p(\theta_{-i}|\theta_i) u_i(\sigma_i(\theta_i), \sigma_{-i}(\theta_{-i}), \theta_i)$$

- Bayesian Nash Eq: Strategy profile σ^* is a Bayesian-Nash Eq iff for all i, for all θ_i,

$$EU_i(\sigma^*_i(\theta_i),\sigma^*_{-i}(),\theta_i) \geq EU_i(\sigma_i(\theta_i),\sigma^*_{-i}(),\theta_i)$$

(best responding w.r.t. its beliefs about the types of the other agents, assuming they are also playing a best response)

Harsanyi, John C., "Games with Incomplete Information Played by Bayesian Players, I-III." Management Science 14 (3): 159-183 (Part I), 14 (5): 320-334 (Part II), 14 (7): 486-502 (Part III) (1967/68)

John Harsanyi was a co-recipient along with John Nash and Reinhard Selten of the 1994 Nobel Memorial Prize in Economics

Social Choice Theory

Assume a group of agents make a decision

1. Agents have preferences over alternatives
   - Agents can rank order the outcomes
     - a>b>c=d is read as "a is preferred to b which is preferred to c which is equivalent to d"
2. Voters are sincere
   - They truthfully tell the center their preferences
3. Outcome is enforced on all agents

The problem

- Majority decision:
  - If more agents prefer a to b, then a should be chosen
- Two outcome setting is easy
  - Choose outcome with more votes!

- What happens if you have 3 or more possible outcomes?

Case 1: Agents specify their top preference

Ballot

Election System

- Plurality Voting
  - One name is ticked on a ballot
  - One round of voting
  - One candidate is chosen

Is this a "good" system?

What do we mean by good?

Example: Plurality

- 3 candidates
  - Lib, NDP, C
- 21 voters with the preferences
  - 10 Lib>NDP>C
  - 6 NDP>C>Lib
  - 5 C>NDP>Lib
- Result: Lib 10, NDP 6, C 5
  - But a majority of voters (11) prefer all other parties more than the Libs!

What can we do?

- Majority system
  - Works well when there are 2 alternatives
  - Not great when there are more than 2 choices
- Proposal:
  - Organize a series of votes between 2 alternatives at a time
  - How this is organized is called an agenda
    - Or a cup (often in sports)

Agendas

- 3 candidates {a,b,c}
- Agenda a,b,c

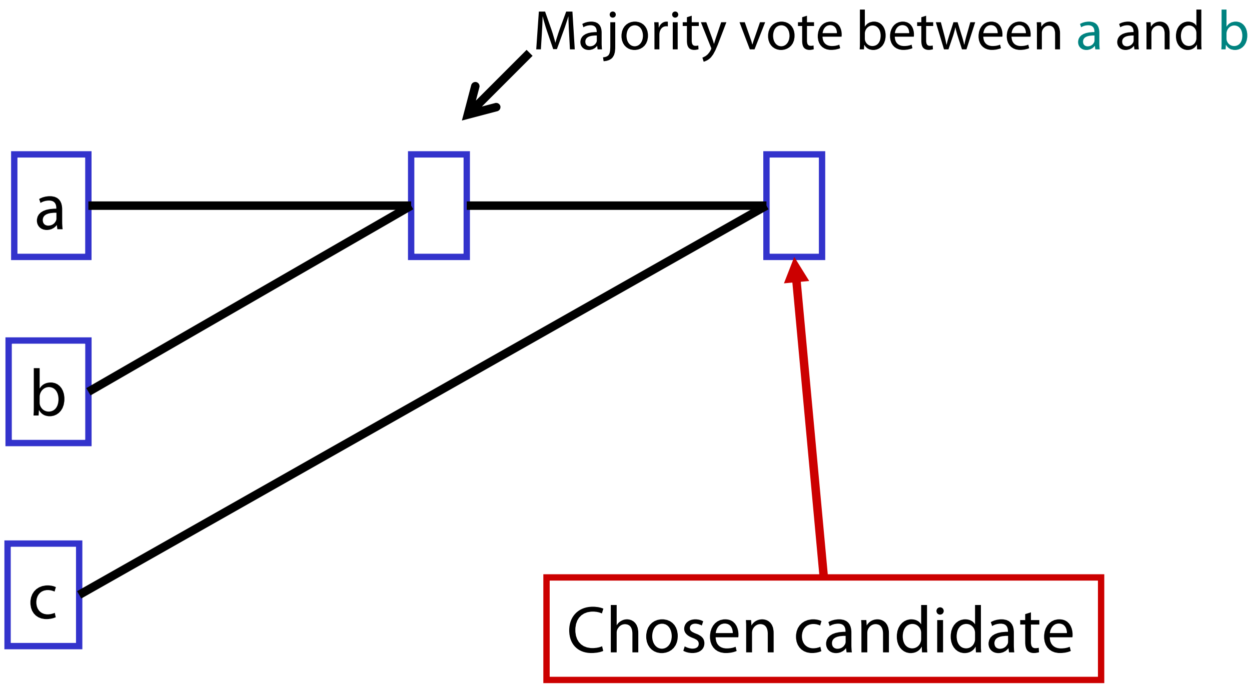

Agenda paradox

- *Binary protocol (majority rule) = cup*
- Three types of agents:
  1. x > z > y (35%)
  2. y > x > z (33%)
  3. z > y > x (32%)

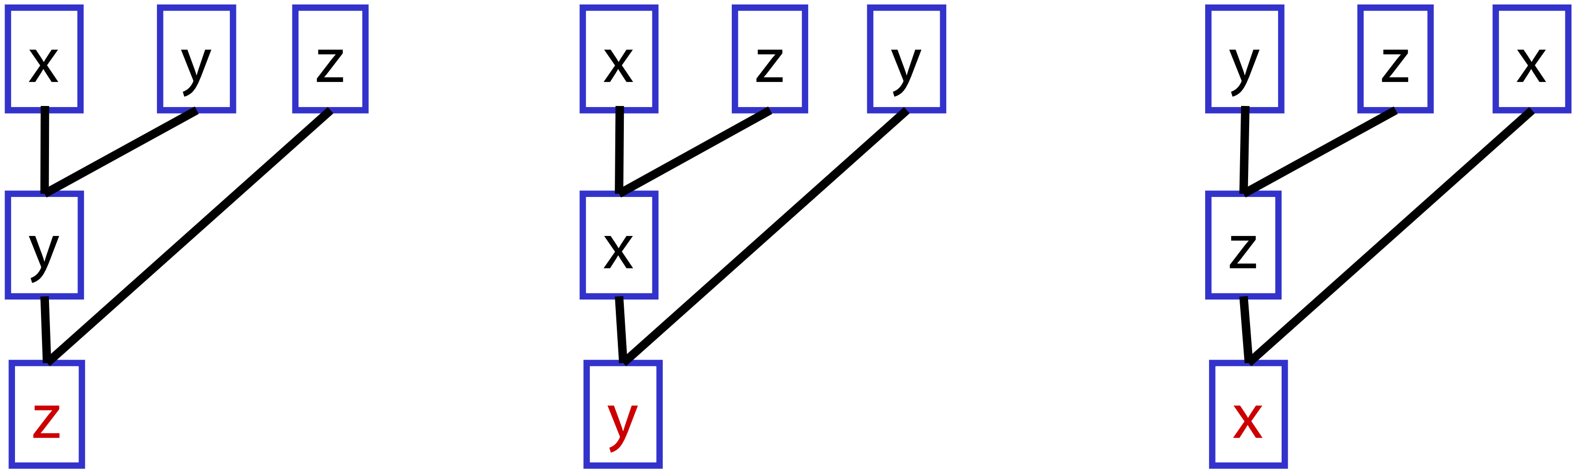

- Power of agenda setter (e.g. chairman)
- Vulnerable to irrelevant alternatives (z)
  - x vs. y only leads to winner y
  - But adding z may lead to x winning (last agenda)

Another problem: Pareto dominated winner paradox

Agents:

1. x > y > b > a
2. a > x > y > b
3. b > a > x > y

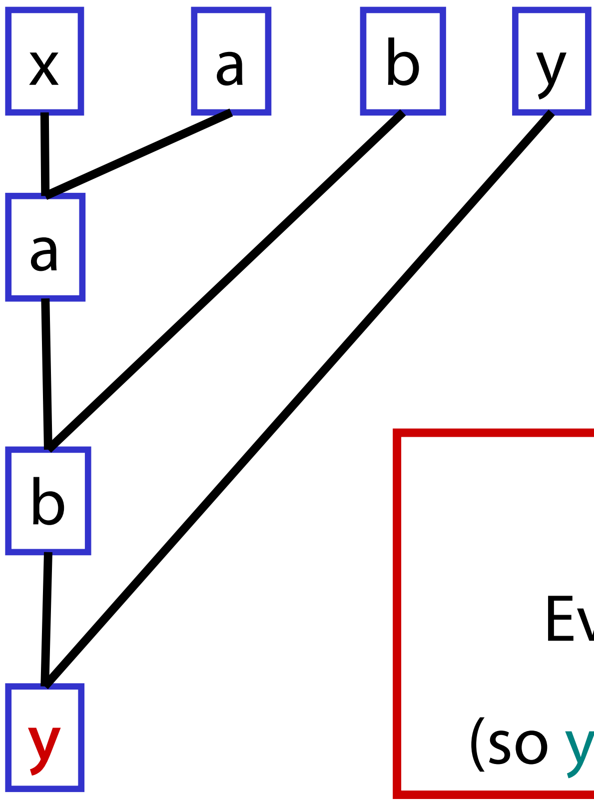

BUT

Everyone prefers x to y!

(so y pareto dominated by x)

Case 2: Agents specify their complete preferences

Maybe the problem was with the ballots!

Ballot

X>Y>Z

Now have more information

Condorcet

- Proposed the following
  - Compare each pair of alternatives
  - Declare "a" is socially preferred to "b" if more voters strictly prefer a to b

- Condorcet Principle: If one alternative is preferred to all other candidates then it should be selected

Example: Condorcet

- 3 candidates
  - Lib, NDP, C
- 21 voters with the preferences
  - 10 Lib>NDP>C
  - 6 NDP>C>Lib
  - 5 C>NDP>Lib
- Result:
  - NDP win! (11/21 prefer them to Lib, 16/21 prefer them to C)

A Problem

- 3 candidates
  - Lib, NDP, C
- 3 voters with the preferences
  - Lib>NDP>C
  - NDP>C>Lib
  - C>Lib>NDP
- Result:
  - No Condorcet Winner

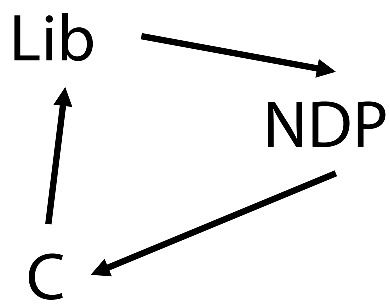

Borda Count

- Each ballot is a list of ordered alternatives
- On each ballot compute the rank of each alternative
- Rank order alternatives based on decreasing sum of their ranks

A>B>C

A>C>B

C>A>B

→

A: 4

B: 8

C: 6

Borda Count

- Simple
- Always a Borda Winner
- BUT does not always choose Condorcet winner!
- 3 voters
  - 2: b>a>c>d
  - 1: a>c>d>b

<u>Borda scores:</u>

a:5, b:6, c:8, d:11

Therefore a wins

BUT b is the Condorcet winner

Inverted-order paradox

- Borda rule with 4 alternatives
  - Each agent gives 1 point to best option, 2 to second best...
- Agents:
  1. x > c > b > a
  2. a > x > c > b
  3. b > a > x > c
  4. x > c > b > a
  5. a > x > c > b
  6. b > a > x > c
  7. x > c > b > a
- x=13, a=18, b=19, c=20
- Remove x: c=13, b=14, a=15

Borda rule vulnerable to irrelevant alternatives

- Three types of agents:

  1. x > z > y (35%)
  2. y > x > z (33%)
  3. z > y > x (32%)

- Borda winner is x
- Remove z: Borda winner is y

Desirable properties for a voting protocol

- No dictators
- Universality (unrestricted domain)
  - It should work with any set of preferences
- Non-imposition (citizen sovereignty)
  - Every possible societal preference order should be achievable
- Independence of irrelevant alternatives (IIA)
  - The comparison of two alternatives should depend only on their standings among agents' preferences, not on the ranking of other alternatives
- Monotonicity
  - An individual should not be able to hurt an option by ranking it higher.
- Paretian
  - If all all agents prefer x to y then in the outcome x should be preferred to y

Arrow's Theorem (1951)

If there are 3 or more alternatives and a finite number of agents then there is **no** protocol which satisfies all desired properties

Take-home Message

- Despair?
  - No ideal voting method
  - That would be boring!

- A group is more complex than an individual
- Weigh the pro's and con's of each system and understand the setting they will be used in

- Do not believe anyone who says they have the best voting system out there!