CS344M
Autonomous Multiagent Systems

Prof: Peter Stone

Department of Computer Science
The University of Texas at Austin

Good Afternoon, Colleagues

Are there any questions?

Good Afternoon, Colleagues

Are there any questions?

- Why is the sequential auction difficult?
- Was there negative social utility in the Clarke Tax Algorithm?

Logistics

- Peer reviews due next Thursday

Logistics

- Peer reviews due next Thursday
- Final projects due sooner than you think!

Logistics

- Peer reviews due next Thursday
- Final projects due sooner than you think!
  - Code due Tuesday, November 30th.
  - Written reports due Thursday, December 2nd.

Logistics

- Peer reviews due next Thursday
- Final projects due sooner than you think!
  - Code due Tuesday, November 30th.
  - Written reports due Thursday, December 2nd.
- FAI talk on Friday at 11 - poker: PAI 3.14

Distributed Rational Decision Making

Self-interested, rational agent

Distributed Rational Decision Making

Self-interested, rational agent

- Self-interested: maximize own goals
  - No concern for global good
- Rational:

Distributed Rational Decision Making

Self-interested, rational agent

- Self-interested: maximize own goals
  - No concern for global good
- Rational: agents are smart
  - Ideally, will act *optimally*

The protocol is key

Auctions vs. voting

- Auctions: maximize profit
  - result affects buyer and seller
- Voting: maximize social good
  - result affects all

Gibbard-Satterthwaite

- Example: Bush, Gore, or Nader?

Gibbard-Satterthwaite

- Example: Bush, Gore, or Nader?
  - Assume your preference is Nader $>$ Gore $>$ Bush
  - For whom should you vote?

Gibbard-Satterthwaite

- Example: Bush, Gore, or Nader?
  - Assume your preference is Nader $>$ Gore $>$ Bush
  - For whom should you vote?
  - What if we change the system?

Gibbard-Satterthwaite

- Example: Bush, Gore, or Nader?
  - Assume your preference is Nader $>$ Gore $>$ Bush
  - For whom should you vote?
  - What if we change the system?
  - Plurality, Binary, Borda?

Gibbard-Satterthwaite

- Example: Bush, Gore, or Nader?
  - Assume your preference is Nader $>$ Gore $>$ Bush
  - For whom should you vote?
  - What if we change the system?
  - Plurality, Binary, Borda?
- 3+ candidates $\Longrightarrow$ only dictatorial system eliminates need for tactical voting
  - One person appointed

Gibbard-Satterthwaite

- Example: Bush, Gore, or Nader?
  - Assume your preference is Nader $>$ Gore $>$ Bush
  - For whom should you vote?
  - What if we change the system?
  - Plurality, Binary, Borda?
- 3+ candidates $\Longrightarrow$ only dictatorial system eliminates need for tactical voting
  - One person appointed
- No point thinking of a "better" voting system
- Assumption: no restrictions on preferences

Gibbard-Satterthwaite

- Example: Bush, Gore, or Nader?
  - Assume your preference is Nader $>$ Gore $>$ Bush
  - For whom should you vote?
  - What if we change the system?
  - Plurality, Binary, Borda?
- 3+ candidates $\Longrightarrow$ only dictatorial system eliminates need for tactical voting
  - One person appointed
- No point thinking of a "better" voting system
- Assumption: no restrictions on preferences

$\boxed{\text{What about Clarke tax algorithm?}}$

Types of Tactical Voting

- Compromising: Rank someone higher to get him/her elected
  - e.g. Gore instead of Nader

Types of Tactical Voting

- Compromising: Rank someone higher to get him/her elected
  - e.g. Gore instead of Nader
- Burying: Rank someone lower to get him/her defeated
  - e.g. in Borda protocol

Types of Tactical Voting

- Compromising: Rank someone higher to get him/her elected
  - e.g. Gore instead of Nader
- Burying: Rank someone lower to get him/her defeated
  - e.g. in Borda protocol
- Push-over: Rank someone higher to get someone else elected
  - e.g. in a protocol with multiple rounds

Arrow's Theorem

Universality.

Arrow's Theorem

Universality. The voting method should provide a complete ranking of all alternatives from any set of individual preference ballots.

Arrow's Theorem

Universality. The voting method should provide a complete ranking of all alternatives from any set of individual preference ballots.

Pareto optimality.

Arrow's Theorem

Universality. The voting method should provide a complete ranking of all alternatives from any set of individual preference ballots.

Pareto optimality. If everyone prefers X to Y, then the outcome should rank X above Y.

Arrow's Theorem

Universality. The voting method should provide a complete ranking of all alternatives from any set of individual preference ballots.

Pareto optimality. If everyone prefers X to Y, then the outcome should rank X above Y.

Criterion of independence of irrelevant alternatives.

Arrow's Theorem

Universality. The voting method should provide a complete ranking of all alternatives from any set of individual preference ballots.

Pareto optimality. If everyone prefers X to Y, then the outcome should rank X above Y.

Criterion of independence of irrelevant alternatives. If one set of preference ballots would lead to an an overall ranking of alternative X above alternative Y and if some preference ballots are changed without changing the relative rank of X and Y, then the method should still rank X above Y.

Citizen Sovereignty.

Citizen Sovereignty. Every possible ranking of alternatives can be achieved from some set of individual preference ballots.

Citizen Sovereignty. Every possible ranking of alternatives can be achieved from some set of individual preference ballots.

Non-dictatorship.

Citizen Sovereignty. Every possible ranking of alternatives can be achieved from some set of individual preference ballots.

Non-dictatorship. There should not be one specific voter whose preference ballot is always adopted.

Arrow's Theorem

Universality.

Arrow's Theorem

Universality. Complete rankings

Arrow's Theorem

Universality. Complete rankings

Pareto optimality.

Arrow's Theorem

Universality. Complete rankings

Pareto optimality. X > Y if all agree

Arrow's Theorem

Universality. Complete rankings

Pareto optimality. X > Y if all agree

Citizen Sovereignty.

Arrow's Theorem

Universality. Complete rankings

Pareto optimality. X > Y if all agree

Citizen Sovereignty. Any ranking possible

Arrow's Theorem

Universality. Complete rankings

Pareto optimality. X > Y if all agree

Citizen Sovereignty. Any ranking possible

Non-dictatorship.

Arrow's Theorem

Universality. Complete rankings

Pareto optimality. X > Y if all agree

Citizen Sovereignty. Any ranking possible

Non-dictatorship. No one voter decides

Arrow's Theorem

Universality. Complete rankings

Pareto optimality. X > Y if all agree

Citizen Sovereignty. Any ranking possible

Non-dictatorship. No one voter decides

Independence of irrelevant alternatives.

Arrow's Theorem

Universality. Complete rankings

Pareto optimality. X > Y if all agree

Citizen Sovereignty. Any ranking possible

Non-dictatorship. No one voter decides

Independence of irrelevant alternatives. Removing or adding a non-winner doesn't change winner

Arrow's Theorem

Universality. Complete rankings

Pareto optimality. X > Y if all agree

Citizen Sovereignty. Any ranking possible

Non-dictatorship. No one voter decides

Independence of irrelevant alternatives. Removing or adding a non-winner doesn't change winner

Not all possible!

Condorcet Voting

- Strategy proof under weaker irrelevant alternatives criterion

Condorcet Voting

- Strategy proof under weaker irrelevant alternatives criterion
- A pairwise method

Condorcet Voting

- Strategy proof under weaker irrelevant alternatives criterion
- A pairwise method
- Smith set: smallest set of candidates such that each candidate in the set preferred over each candidate not in the set

Condorcet Voting

- Strategy proof under weaker irrelevant alternatives criterion
- A pairwise method
- Smith set: smallest set of candidates such that each candidate in the set preferred over each candidate not in the set
- Every candidate in the Smith set is relevant

Condorcet Example

- 48: $A > B > C$
- 40: $B > C > A$
- 12: $C > B > A$

Condorcet Example

- 48: $A > B > C$
- 40: $B > C > A$
- 12: $C > B > A$

- A vs. B :

Condorcet Example

- 48: $A > B > C$
- 40: $B > C > A$
- 12: $C > B > A$

- A vs. B : 48 – 52 $\Longrightarrow B > A$

Condorcet Example

- 48: A $>$ B $>$ C
- 40: B $>$ C $>$ A
- 12: C $>$ B $>$ A

- A vs. B : 48 – 52 $\Longrightarrow$ B $>$ A
- A vs. C : 48 – 52 $\Longrightarrow$ C $>$ A
- B vs. C : 88 – 12 $\Longrightarrow$ B $>$ C

Condorcet Example

- 48: A $>$ B $>$ C
- 40: B $>$ C $>$ A
- 12: C $>$ B $>$ A

- A vs. B : 48 – 52 $\Longrightarrow$ B $>$ A
- A vs. C : 48 – 52 $\Longrightarrow$ C $>$ A
- B vs. C : 88 – 12 $\Longrightarrow$ B $>$ C

$\boxed{\text{Overall: B} > \text{C} > \text{A}}$

Condorcet Example

- 48: $A > B > C$
- 40: $B > C > A$
- 12: $C > B > A$

- A vs. B : 48 – 52 $\Longrightarrow B > A$
- A vs. C : 48 – 52 $\Longrightarrow C > A$
- B vs. C : 88 – 12 $\Longrightarrow B > C$

$$\boxed{\text{Overall: } B > C > A}$$

- Does that solve everything?

Condorcet Example

- 48: A $>$ B $>$ C
- 40: B $>$ C $>$ A
- 12: C $>$ B $>$ A

- A vs. B : 48 – 52 $\Longrightarrow$ B $>$ A
- A vs. C : 48 – 52 $\Longrightarrow$ C $>$ A
- B vs. C : 88 – 12 $\Longrightarrow$ B $>$ C

$$\boxed{\text{Overall: B} > \text{C} > \text{A}}$$

- Does that solve everything? What about cycles?

Bargaining

$$\boxed{\text{small market, both can come out favorably}}$$

Bargaining

small market, both can come out favorably

- Two people bargaining, each with a preference over outcomes O
- Let o^* be the selected outcome

Bargaining

small market, both can come out favorably

- Two people bargaining, each with a preference over outcomes O
- Let o^* be the selected outcome
- Example: "split the dollar"

Bargaining

small market, both can come out favorably

- Two people bargaining, each with a preference over outcomes O
- Let o^* be the selected outcome
- Example: "split the dollar"
  - One person makes offer o
  - Other rejects with probaility $p(o)$ — based on offer
  - If rejects, both get nothing

Bargaining

small market, both can come out favorably

- Two people bargaining, each with a preference over outcomes O
- Let o^* be the selected outcome
- Example: "split the dollar"
  - One person makes offer o
  - Other rejects with probaility $p(o)$ — based on offer
  - If rejects, both get nothing
- Another version
  - One person makes an offer
  - Other accepts, rejects, or counters
  - If counters, \$.05 lost
  - Game ends with an accept or reject

Nash Bargaining Solution

Unique solution that satisfies:

Nash Bargaining Solution

Unique solution that satisfies:

Invariance: only preference *orders* matter

Anonymity: no discrimination

Pareto efficiency: if one does better, other does worse

Independence of irrelevant alternatives: removing outcomes doesn't change things

Nash Bargaining Solution

Unique solution that satisfies:

Invariance: only preference *orders* matter

Anonymity: no discrimination

Pareto efficiency: if one does better, other does worse

Independence of irrelevant alternatives: removing outcomes doesn't change things

$$\boxed{\textbf{Maximize } u_1(o) * u_2(o)}$$

General Equilibrium

Consumers: utilities, endowments

Producers: production possibility sets

Variables: prices on goods

General Equilibrium

Consumers: utilities, endowments

Producers: production possibility sets

Variables: prices on goods

Equilibrium: allocation (prices) such that consumers maximize preferences, producers maximize profits

General Equilibrium

Consumers: utilities, endowments

Producers: production possibility sets

Variables: prices on goods

Equilibrium: allocation (prices) such that consumers maximize preferences, producers maximize profits

- Assumption: agent doesn't affect prices

General Equilibrium

Consumers: utilities, endowments

Producers: production possibility sets

Variables: prices on goods

Equilibrium: allocation (prices) such that consumers maximize preferences, producers maximize profits

- Assumption: agent doesn't affect prices
  - Only true if market is infinitely large
  - Else, strategic bidding (like bargaining) possible

General Equilibrium

Consumers: utilities, endowments

Producers: production possibility sets

Variables: prices on goods

Equilibrium: allocation (prices) such that consumers maximize preferences, producers maximize profits

- Assumption: agent doesn't affect prices
  - Only true if market is infinitely large
  - Else, strategic bidding (like bargaining) possible
- Assumption: no externalities

General Equilibrium

Consumers: utilities, endowments
Producers: production possibility sets
Variables: prices on goods
Equilibrium: allocation (prices) such that consumers maximize preferences, producers maximize profits

- Assumption: agent doesn't affect prices
  - Only true if market is infinitely large
  - Else, strategic bidding (like bargaining) possible
- Assumption: no externalities
  - Utilities or production sets don't depend on others'

General Equilibrium

Consumers: utilities, endowments
Producers: production possibility sets
Variables: prices on goods
Equilibrium: allocation (prices) such that consumers maximize preferences, producers maximize profits

- Assumption: agent doesn't affect prices
  - Only true if market is infinitely large
  - Else, strategic bidding (like bargaining) possible
- Assumption: no externalities
  - Utilities or production sets don't depend on others'
  - Braess' paradox

Other DRDM

- Contract nets: task allocation among agents

Other DRDM

- Contract nets: task allocation among agents
  - Contingencies
  - Leveled commitment (price)

Other DRDM

- Contract nets: task allocation among agents
  - Contingencies
  - Leveled commitment (price)
- Coalitions

Other DRDM

- Contract nets: task allocation among agents
  - Contingencies
  - Leveled commitment (price)
- Coalitions
  - Formation
  - Optimization within
  - Payoff division

Contract Nets

Task allocation among agents

Contract Nets

Task allocation among agents

- OCSM-contracts: original, cluster, swap, multiagent
  - Hill-climbing leads to optimum
  - Without any type, may be no sequence to optimum

Contract Nets

Task allocation among agents

- OCSM-contracts: original, cluster, swap, multiagent
  - Hill-climbing leads to optimum
  - Without any type, may be no sequence to optimum
- Backing out of contracts

Contract Nets

Task allocation among agents

- OCSM-contracts: original, cluster, swap, multiagent
  - Hill-climbing leads to optimum
  - Without any type, may be no sequence to optimum
- Backing out of contracts
  - Contingency (future events)

Contract Nets

Task allocation among agents

- OCSM-contracts: original, cluster, swap, multiagent
  - Hill-climbing leads to optimum
  - Without any type, may be no sequence to optimum
- Backing out of contracts
  - Contingency (future events)
  - Leveled commitment (price)

Contract Nets

Task allocation among agents

- OCSM-contracts: original, cluster, swap, multiagent
  - Hill-climbing leads to optimum
  - Without any type, may be no sequence to optimum
- Backing out of contracts
  - Contingency (future events)
  - Leveled commitment (price)
  - What are some of the tradeoffs?

Contingency vs. leveled commitment

Contingency problems:

Contingency vs. leveled commitment

Contingency problems:

1. Hard to track all contingencies

Contingency vs. leveled commitment

Contingency problems:

1. Hard to track all contingencies
2. Could be impossible to enumerate all possible contingencies

Contingency vs. leveled commitment

Contingency problems:

1. Hard to track all contingencies
2. Could be impossible to enumerate all possible contingencies
3. What if only one agent observes that relevant event happened?

Contingency vs. leveled commitment

Contingency problems:

1. Hard to track all contingencies
2. Could be impossible to enumerate all possible contingencies
3. What if only one agent observes that relevant event happened?

Leveled commitment problems:

Contingency vs. leveled commitment

Contingency problems:

1. Hard to track all contingencies
2. Could be impossible to enumerate all possible contingencies
3. What if only one agent observes that relevant event happened?

Leveled commitment problems:

1. Breacher's gain may be smaller than victim's loss

Contingency vs. leveled commitment

Contingency problems:

1. Hard to track all contingencies
2. Could be impossible to enumerate all possible contingencies
3. What if only one agent observes that relevant event happened?

Leveled commitment problems:

1. Breacher's gain may be smaller than victim's loss
2. May decommit insincerely (wait for other) - ineffiicent contracts executed.

Coalitions

- Formation
- Optimization within
- Payoff division

DRDM Summary

For many agents: voting, general equilibrium, auctions

For fewer agents: auctions, contract nets, bargaining

Possible in all: coalitions

DRDM Summary

For many agents: voting, general equilibrium, auctions

For fewer agents: auctions, contract nets, bargaining

Possible in all: coalitions

All self-interested, rational agents