Mechanism Design and Auctions

Kevin Leyton-Brown & Yoav Shoham

Chapter 7 of *Multiagent Systems* (MIT Press, 2012)

Drawing on material that first appeared in our own book,
Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations
(Cambridge University Press, 2009)

Mechanism Design

- Goal: pick a way of mapping from agents' actions to social choices in a way that will cause rational agents to behave in a desired way, specifically maximizing the mechanism designer's own "utility" or objective function
  - each agent holds private information, in the Bayesian game sense
  - often, we're interested in settings where agents' action space is identical to their type space, and an action can be interpreted as a declaration of the agent's type
- Various equivalent ways of looking at this setting
  - perform an optimization problem, given that the values of (some of) the inputs are unknown
  - choose the Bayesian game out of a set of possible Bayesian games that maximizes some performance measure
  - design a game that *implements* a particular social choice function in equilibrium, given that the designer no longer knows agents' preferences and the agents might lie

Overview

1. Mechanism Design with Unrestricted Preferences
   - Implementation
   - Revelation Principle
   - Impossibility of general, dominant-strategy implementation
2. Quasilinear Preferences
3. Efficient Mechanisms
4. Single-Good Auctions
5. Position Auctions
6. Combinatorial Auctions

Bayesian Game Setting

- Social choice in a setting where agents can't be relied upon to disclose their preferences honestly.
- Start with a set of agents in a Bayesian game setting (but no actions).

Definition (Bayesian game setting)

A Bayesian game setting is a tuple (N, O, Θ, p, u), where

- N is a finite set of n agents;
- O is a set of outcomes;
- $\Theta = \Theta_1 \times \cdots \times \Theta_n$ is a set of possible joint type vectors;
- p is a (common prior) probability distribution on Θ; and
- $u = (u_1, \ldots, u_n)$, where $u_i : O \times \Theta \mapsto \mathbb{R}$ is the utility function for each player i.

Mechanism Design

Definition (Mechanism)

A mechanism (for a Bayesian game setting (N, O, Θ, p, u)) is a pair (A, M), where

- $A = A_1 \times \cdots \times A_n$, where A_i is the set of actions available to agent $i \in N$; and
- $M : A \mapsto \Pi(O)$ maps each action profile to a distribution over outcomes.

Thus, the designer gets to specify

- the action sets for the agents (though they may be constrained by the environment)
- the mapping to outcomes, over which agents have utility
- can't change outcomes; agents' preferences or type spaces

Overview

1. Mechanism Design with Unrestricted Preferences
   - Implementation
   - Revelation Principle
   - Impossibility of general, dominant-strategy implementation
2. Quasilinear Preferences
3. Efficient Mechanisms
4. Single-Good Auctions
5. Position Auctions
6. Combinatorial Auctions

Implementation in Dominant Strategies

Definition (Implementation in dominant strategies)

Given a Bayesian game setting (N, O, Θ, p, u), a mechanism (A, M) is an implementation in dominant strategies of a social choice function C (over N and O) if for any vector of utility functions u, the game has an equilibrium in dominant strategies, and in any such equilibrium a^* we have $M(a^*) = C(u)$.

Implementation in Bayes-Nash equilibrium

Definition (Bayes–Nash implementation)

Given a Bayesian game setting (N, O, Θ, p, u), a mechanism (A, M) is an implementation in Bayes–Nash equilibrium of a social choice function C (over N and O) if there exists a Bayes–Nash equilibrium of the game of incomplete information (N, A, Θ, p, u) such that for every $\theta \in \Theta$ and every action profile $a \in A$ that can arise given type profile θ in this equilibrium, we have that $M(a) = C(u(\cdot, \theta))$.

Bayes-Nash Implementation Comments

Bayes-Nash Equilibrium Problems:

- there could be more than one equilibrium
  - which one should I expect agents to play?
- agents could miscoordinate and play none of the equilibria
- asymmetric equilibria are implausible

Refinements:

- Symmetric Bayes-Nash implementation
- *Ex-post* implementation

Implementation Comments

We can require that the desired outcome arises

- in the only equilibrium
- in every equilibrium
- in at least one equilibrium

Forms of implementation:

- Direct Implementation: agents each simultaneously send a single message to the center
- Indirect Implementation: agents may send a sequence of messages; in between, information may be (partially) revealed about the messages that were sent previously like extensive form

Overview

Revelation Principle

- It turns out that any social choice function that can be implemented by any mechanism can be implemented by a truthful, direct mechanism!
- Consider an arbitrary, non-truthful mechanism (e.g., may be indirect)

Revelation Principle

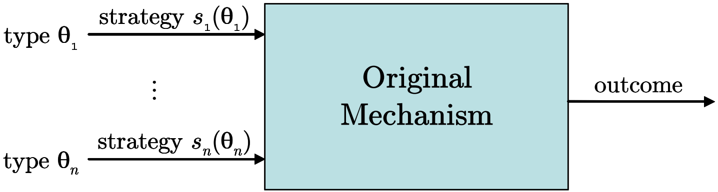

- It turns out that any social choice function that can be implemented by any mechanism can be implemented by a truthful, direct mechanism!
- Consider an arbitrary, non-truthful mechanism (e.g., may be indirect)

Revelation Principle

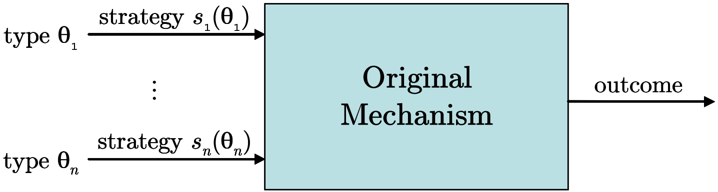

- It turns out that any social choice function that can be implemented by any mechanism can be implemented by a truthful, direct mechanism!
- Consider an arbitrary, non-truthful mechanism (e.g., may be indirect)
- Recall that a mechanism defines a game, and consider an equilibrium $s = (s_1, \ldots, s_n)$

Revelation Principle

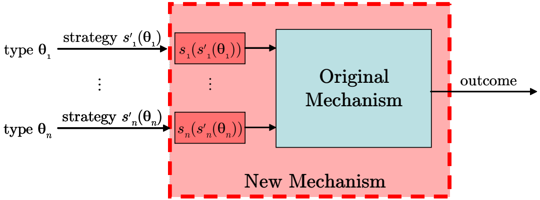

- We can construct a new direct mechanism, as shown above
- This mechanism is truthful by exactly the same argument that s was an equilibrium in the original mechanism
- "The agents don't have to lie, because the mechanism already lies for them."

Computational Criticism of the Revelation Principle

- computation is pushed onto the center
  - often, agents' strategies will be computationally expensive
    - e.g., in the shortest path problem, agents may need to compute shortest paths, cutsets in the graph, etc.
  - since the center plays equilibrium strategies for the agents, the center now incurs this cost
- if computation is intractable, so that it cannot be performed by agents, then in a sense the revelation principle doesn't hold
  - agents can't play the equilibrium strategy in the original mechanism
  - however, in this case it's unclear what the agents will do

Discussion of the Revelation Principle

- The set of equilibria is not always the same in the original mechanism and revelation mechanism
  - of course, we've shown that the revelation mechanism does have the original equilibrium of interest
  - however, in the case of indirect mechanisms, even if the indirect mechanism had a unique equilibrium, the revelation mechanism can also have new, bad equilibria
- So what is the revelation principle good for?
  - recognition that truthfulness is not a restrictive assumption
  - for analysis purposes, we can consider only truthful mechanisms, and be assured that such a mechanism exists
  - recognition that indirect mechanisms can't do (inherently) better than direct mechanisms

Overview

1. Mechanism Design with Unrestricted Preferences
   - Implementation
   - Revelation Principle
   - Impossibility of general, dominant-strategy implementation
2. Quasilinear Preferences
3. Efficient Mechanisms
4. Single-Good Auctions
5. Position Auctions
6. Combinatorial Auctions

Impossibility Result

Theorem (Gibbard-Satterthwaite)

Consider any social choice function C of N and O. If:

1. $|O| \geq 3$ *(there are at least three outcomes);*
2. C *is onto; that is, for every* $o \in O$ *there is a preference profile* $[\succ]$ *such that* $C([\succ]) = o$ *(this property is sometimes also called citizen sovereignty); and*
3. C *is dominant-strategy truthful,*

then C is dictatorial.

What does this mean?

- We should be discouraged about the possibility of implementing arbitrary social-choice functions in mechanisms.
- However, in practice we can circumvent the Gibbard-Satterthwaite theorem in two ways:
  - use a weaker form of implementation
    - note: the result only holds for dominant strategy implementation, not e.g., Bayes-Nash implementation
  - relax the onto condition and the (implicit) assumption that agents are allowed to hold arbitrary preferences

Overview

Quasilinear Utility

Definition (Quasilinear preferences)

Agents have quasilinear preferences in an n-player Bayesian game when the set of outcomes is

$$O = X \times \mathbb{R}^n$$

for a finite set X, and the utility of an agent i given joint type θ is given by

$$u_i(o, \theta) = u_i(x, \theta) - p_i,$$

where $o = (x, p)$ is an element of O, $u_i : X \times \Theta \mapsto \mathbb{R}$ is an arbitrary function.

Quasilinear utility

- $u_i(o, \theta) = u_i(x, \theta) - p_i$
- We split the mechanism into a choice rule and a payment rule:
  - $x \in X$ is a discrete, non-monetary outcome
  - $p_i \in \mathbb{R}$ is a monetary payment (possibly negative) that agent i must make to the mechanism
- Implications:

Quasilinear utility

- $u_i(o, \theta) = u_i(x, \theta) - p_i$
- We split the mechanism into a choice rule and a payment rule:
  - $x \in X$ is a discrete, non-monetary outcome
  - $p_i \in \mathbb{R}$ is a monetary payment (possibly negative) that agent i must make to the mechanism
- Implications:
  - $u_i(x, \theta)$ is not influenced by the amount of money an agent has
  - agents don't care how much others are made to pay (though they *can* care about how the choice affects others.)

Overview

1. Mechanism Design with Unrestricted Preferences

2. Quasilinear Preferences
    - Mechanism design in the quasilinear setting

3. Efficient Mechanisms

4. Single-Good Auctions

5. Position Auctions

6. Combinatorial Auctions

Quasilinear Mechanism

Definition (Quasilinear mechanism)

A mechanism in the quasilinear setting (for a Bayesian game setting $(N, O = X \times \mathbb{R}^n, \Theta, p, u)$) is a triple $(A, \chi, \mathcal{p})$, where

- $A = A_1 \times \cdots \times A_n$, where A_i is the set of actions available to agent $i \in N$,
- $\chi : A \mapsto \Pi(X)$ maps each action profile to a distribution over choices, and
- $\mathcal{p} : A \mapsto \mathbb{R}^n$ maps each action profile to a payment for each agent.

Direct Quasilinear Mechanism

Definition (Direct quasilinear mechanism)

A direct quasilinear mechanism (for a Bayesian game setting $(N, O = X \times \mathbb{R}^n, \Theta, p, u)$) is a pair (χ, p). It defines a standard mechanism in the quasilinear setting, where for each i, $A_i = \Theta_i$.

Definition (Conditional utility independence)

A Bayesian game exhibits conditional utility independence if for all agents $i \in N$, for all outcomes $o \in O$ and for all pairs of joint types θ and $\theta' \in \Theta$ for which $\theta_i = \theta'_i$, it holds that $u_i(o, \theta) = u_i(o, \theta')$.

Quasilinear Mechanisms with Conditional Utility Independence

- Given conditional utility independence, we can write i's utility function as $u_i(o, \theta_i)$
  - it does not depend on the other agents' types
- An agent's valuation for choice $x \in X$: $v_i(x) = u_i(x, \theta_i)$
  - the maximum amount i would be willing to pay to get x
  - in fact, i would be indifferent between keeping the money and getting x
- Alternate definition of direct mechanism:
  - ask agents i to declare $v_i(x)$ for each $x \in X$
- Define $\hat{v}_i$ as the valuation that agent i declares to such a direct mechanism
  - may be different from his true valuation v_i
- Also define the tuples $\hat{v}$, $\hat{v}_{-i}$

Truthfulness

Definition (Truthfulness)

A quasilinear mechanism is truthful if it is direct and $\forall i \forall v_i$, agent i's equilibrium strategy is to adopt the strategy $\hat{v}_i = v_i$.

- Our definition before, adapted for the quasilinear setting

Efficiency

Definition (Efficiency)

A quasilinear mechanism is strictly Pareto efficient, or just efficient, if in equilibrium it selects a choice x such that

$$\forall v \forall x', \sum_i v_i(x) \geq \sum_i v_i(x').$$

- An efficient mechanism selects the choice that maximizes the sum of agents' utilities, disregarding monetary payments.

Efficiency

Definition (Efficiency)

A quasilinear mechanism is strictly Pareto efficient, or just efficient, if in equilibrium it selects a choice x such that

$$\forall v \forall x', \sum_i v_i(x) \geq \sum_i v_i(x').$$

- Called economic efficiency to distinguish from other (e.g., computational) notions
- Also called social-welfare maximization
- Note: defined in terms of true (not declared) valuations.

Budget Balance

Definition (Budget balance)

A quasilinear mechanism is budget balanced when

$$\forall v, \sum_i p_i(s(v)) = 0,$$

where s is the equilibrium strategy profile.

- regardless of the agents' types, the mechanism collects and disburses the same amount of money from and to the agents

Budget Balance

Definition (Budget balance)

A quasilinear mechanism is budget balanced when

$$\forall v, \sum_i p_i(s(v)) = 0,$$

where s is the equilibrium strategy profile.

- regardless of the agents' types, the mechanism collects and disburses the same amount of money from and to the agents
- relaxed version: weak budget balance:

$$\forall v, \sum_i p_i(s(v)) \geq 0$$

  - the mechanism never takes a loss, but it may make a profit

Budget Balance

Definition (Budget balance)

A quasilinear mechanism is budget balanced when

$$\forall v, \sum_i p_i(s(v)) = 0,$$

where s is the equilibrium strategy profile.

- regardless of the agents' types, the mechanism collects and disburses the same amount of money from and to the agents
- Budget balance can be required to hold *ex ante*:

$$\mathbb{E}_v \sum_i p_i(s(v)) = 0$$

  - the mechanism must break even or make a profit only on expectation

Individual Rationality

Definition (*Ex interim* individual rationality)

A mechanism is ex interim individual rational when
$\forall i \forall v_i,\ \mathbb{E}_{v_{-i}|v_i} v_i(\chi(s_i(v_i), s_{-i}(v_{-i}))) - p_i(s_i(v_i), s_{-i}(v_{-i})) \geq 0,$
where s is the equilibrium strategy profile.

- no agent loses by participating in the mechanism.
- *ex interim* because it holds for *every* possible valuation for agent i, but averages over the possible valuations of the other agents.

Individual Rationality

Definition (*Ex interim* individual rationality)

A mechanism is ex interim individual rational when $\forall i \forall v_i,\ \mathbb{E}_{v_{-i}|v_i} v_i(\chi(s_i(v_i), s_{-i}(v_{-i}))) - p_i(s_i(v_i), s_{-i}(v_{-i})) \geq 0$, where s is the equilibrium strategy profile.

- no agent loses by participating in the mechanism.
- *ex interim* because it holds for *every* possible valuation for agent i, but averages over the possible valuations of the other agents.

Definition (*Ex post* individual rationality)

A mechanism is ex post individual rational when $\forall i \forall v,\ v_i(\chi(s(v))) - p_i(s(v)) \geq 0$, where s is the equilibrium strategy profile.

Tractability

Definition (Tractability)

A mechanism is tractable when $\forall \hat{v}$, $\chi(\hat{v})$ and $p(\hat{v})$ can be computed in polynomial time.

- The mechanism is computationally feasible.

Revenue Maximization

We can also add an objective function to our mechanism. One example: revenue maximization.

Definition (Revenue maximization)

A mechanism is revenue maximizing when, among the set of functions χ and p that satisfy the other constraints, the mechanism selects the χ and p that maximize $\mathbb{E}_\theta \sum_i p_i(s(\theta))$, where $s(\theta)$ denotes the agents' equilibrium strategy profile.

- The mechanism designer can choose among mechanisms that satisfy the desired constraints by adding an objective function such as revenue maximization.

Revenue Minimization

- The mechanism may not be intended to make money.
- Budget balance may be impossible to satisfy.
- Set weak budget balance as a constraint and add the following objective.

Definition (Revenue minimization)

A quasilinear mechanism is revenue minimizing when, among the set of functions χ and p that satisfy the other constraints, the mechanism selects the χ and p that minimize $\max_v \sum_i p_i(s(v))$ in equilibrium, where $s(v)$ denotes the agents' equilibrium strategy profile.

- Note: this considers the worst case over valuations; we could consider average case instead.

Fairness

- Fairness is hard to define. What is fairer:
  - an outcome that fines all agents $100 and makes a choice that all agents hate equally?
  - an outcome that charges all agents $0 and makes a choice that some agents hate and some agents like?

Fairness

- Fairness is hard to define. What is fairer:
  - an outcome that fines all agents \$100 and makes a choice that all agents hate equally?
  - an outcome that charges all agents \$0 and makes a choice that some agents hate and some agents like?
- Maxmin fairness: make the least-happy agent the happiest.

Definition (Maxmin fairness)

A quasilinear mechanism is maxmin fair when, among the set of functions χ and p that satisfy the other constraints, the mechanism selects the χ and p that maximize

$$\mathbb{E}_v\left[\min_{i\in N} v_i(\chi(s(v))) - p_i(s(v))\right],$$

where $s(v)$ denotes the agents' equilibrium strategy profile.

Price of Anarchy Minimization

- When an efficient mechanism is impossible, we may want to get as close as possible
- Minimize the worst-case ratio between optimal social welfare and the social welfare achieved by the given mechanism.

Definition (Price-of-anarchy minimization)

A quasilinear mechanism minimizes the price of anarchy when, among the set of functions χ and p that satisfy the other constraints, the mechanism selects the χ and p that minimize

$$\max_{v \in V} \frac{\max_{x \in X} \sum_{i \in N} v_i(x)}{\sum_{i \in N} v_i\left(\chi(s(v))\right)},$$

where $s(v)$ denotes the agents' equilibrium strategy profile in the *worst* equilibrium of the mechanism—i.e., the one in which $\sum_{i \in N} v_i(\chi(s(v)))$ is the smallest.

Overview

Overview

A positive result

- Recall that in the quasilinear utility setting, a mechanism can be defined as a choice rule and a payment rule.
- The Groves mechanism is a mechanism that satisfies:
  - dominant strategy (truthfulness)
  - efficiency
- In general it's not:
  - budget balanced
  - individual-rational

  ...though we'll see later that there's some hope for recovering these properties.

The Groves Mechanism

Definition (Groves mechanism)

The Groves mechanism is a direct quasilinear mechanism (χ, p), where

$$\chi(\hat{v}) = \arg\max_{x} \sum_{i} \hat{v}_i(x)$$

$$p_i(\hat{v}) = h_i(\hat{v}_{-i}) - \sum_{j \neq i} \hat{v}_j(\chi(\hat{v}))$$

The Groves Mechanism

$$\chi(\hat{v}) = \arg\max_x \sum_i \hat{v}_i(x)$$

$$p_i(\hat{v}) = h_i(\hat{v}_{-i}) - \sum_{j \neq i} \hat{v}_j(\chi(\hat{v}))$$

- The choice rule should not come as a surprise (why not?)

The Groves Mechanism

$$\chi(\hat{v}) = \arg\max_x \sum_i \hat{v}_i(x)$$

$$p_i(\hat{v}) = h_i(\hat{v}_{-i}) - \sum_{j \neq i} \hat{v}_j(\chi(\hat{v}))$$

- The choice rule should not come as a surprise (why not?) because the mechanism is both truthful and efficient: these properties entail the given choice rule.

The Groves Mechanism

$$\chi(\hat{v}) = \arg\max_x \sum_i \hat{v}_i(x)$$

$$p_i(\hat{v}) = h_i(\hat{v}_{-i}) - \sum_{j \neq i} \hat{v}_j(\chi(\hat{v}))$$

- The choice rule should not come as a surprise (why not?) because the mechanism is both truthful and efficient: these properties entail the given choice rule.
- So what's going on with the payment rule?
  - the agent i must pay some amount $h_i(\hat{v}_{-i})$ that doesn't depend on his own declared valuation
  - the agent i is paid $\sum_{j \neq i} \hat{v}_j(\chi(\hat{v}))$, the sum of the others' valuations for the chosen outcome

Groves Truthfulness

Theorem

Truth telling is a dominant strategy under the Groves mechanism.

Consider a situation where every agent j other than i follows some arbitrary strategy $\hat{v}_j$. Consider agent i's problem of choosing the best strategy $\hat{v}_i$. As a shorthand, we will write $\hat{v} = (\hat{v}_{-i}, \hat{v}_i)$. The best strategy for i is one that solves

$$\max_{\hat{v}_i} \left(v_i(\chi(\hat{v})) - p(\hat{v}) \right)$$

Substituting in the payment function from the Groves mechanism, we have

$$\max_{\hat{v}_i} \left(v_i(\chi(\hat{v})) - h_i(\hat{v}_{-i}) + \sum_{j \neq i} \hat{v}_j(\chi(\hat{v})) \right)$$

Since $h_i(\hat{v}_{-i})$ does not depend on $\hat{v}_i$, it is sufficient to solve

$$\max_{\hat{v}_i} \left(v_i(\chi(\hat{v})) + \sum_{j \neq i} \hat{v}_j(\chi(\hat{v})) \right).$$

Groves Truthfulness

$$\max_{\hat{v}_i}\left(v_i(\chi(\hat{v})) + \sum_{j\neq i}\hat{v}_j(\chi(\hat{v}))\right).$$

The only way the declaration $\hat{v}_i$ influences this maximization is through the choice of x. If possible, i would like to pick a declaration $\hat{v}_i$ that will lead the mechanism to pick an $x \in X$ which solves

$$\max_{x}\left(v_i(x) + \sum_{j\neq i}\hat{v}_j(x)\right). \qquad (1)$$

Under the Groves mechanism,

$$\chi(\hat{v}) = \arg\max_{x}\left(\sum_{i}\hat{v}_i(x)\right) = \arg\max_{x}\left(\hat{v}_i(x) + \sum_{j\neq i}\hat{v}_j(x)\right).$$

The Groves mechanism will choose x in a way that solves the maximization problem in Equation (1) when i declares $\hat{v}_i = v_i$. Because this argument does not depend in any way on the declarations of the other agents, truth-telling is a dominant strategy for agent i.

Proof intuition

- externalities are internalized
  - agents may be able to change the outcome to another one that they prefer, by changing their declaration
  - however, their utility doesn't just depend on the outcome—it also depends on their payment
  - since they get paid the (reported) utility of all the other agents under the chosen allocation, they now have an interest in maximizing everyone's utility rather than just their own
- in general, DS truthful mechanisms have the property that an agent's payment doesn't depend on the amount of his declaration, but only on the other agents' declarations
  - the agent's declaration is used only to choose the outcome, and to set other agents' payments

Groves Uniqueness

Theorem (Green–Laffont)

An efficient social choice function $C : \mathbb{R}^{Xn} \rightarrow X \times \mathbb{R}^n$ can be implemented in dominant strategies for agents with unrestricted quasilinear utilities only if $p_i(v) = h(v_{-i}) - \sum_{j \neq i} v_j(\chi(v))$.

- it turns out that the same result also holds for the broader class of Bayes–Nash incentive-compatible efficient mechanisms.

Overview

VCG

Definition (Clarke tax)

The Clarke tax sets the h_i term in a Groves mechanism as

$$h_i(\hat{v}_{-i}) = \sum_{j \neq i} \hat{v}_j \left(\chi(\hat{v}_{-i})\right).$$

Definition (Vickrey-Clarke-Groves (VCG) mechanism)

The Vickrey-Clarke-Groves mechanism is a direct quasilinear mechanism (χ, p), where

$$\chi(\hat{v}) = \arg\max_x \sum_i \hat{v}_i(x)$$

$$p_i(\hat{v}) = \sum_{j \neq i} \hat{v}_j \left(\chi(\hat{v}_{-i})\right) - \sum_{j \neq i} \hat{v}_j(\chi(\hat{v}))$$

VCG discussion

$$\chi(\hat{v}) = \arg\max_x \sum_i \hat{v}_i(x)$$

$$p_i(\hat{v}) = \sum_{j \neq i} \hat{v}_j\left(\chi(\hat{v}_{-i})\right) - \sum_{j \neq i} \hat{v}_j(\chi(\hat{v}))$$

- You get paid everyone's utility under the allocation that is actually chosen
  - except your own, but you get that directly as utility
- Then you get charged everyone's utility in the world where you don't participate
- Thus you pay your social cost

VCG discussion

$$\chi(\hat{v}) = \arg\max_x \sum_i \hat{v}_i(x)$$

$$p_i(\hat{v}) = \sum_{j \neq i} \hat{v}_j\left(\chi(\hat{v}_{-i})\right) - \sum_{j \neq i} \hat{v}_j(\chi(\hat{v}))$$

Questions:

- who pays 0?

VCG discussion

$$\chi(\hat{v}) = \arg\max_x \sum_i \hat{v}_i(x)$$

$$p_i(\hat{v}) = \sum_{j \neq i} \hat{v}_j\left(\chi(\hat{v}_{-i})\right) - \sum_{j \neq i} \hat{v}_j(\chi(\hat{v}))$$

Questions:

- who pays 0?
  - agents who don't affect the outcome

VCG discussion

$$\chi(\hat{v}) = \arg\max_{x} \sum_{i} \hat{v}_i(x)$$

$$p_i(\hat{v}) = \sum_{j \neq i} \hat{v}_j\left(\chi(\hat{v}_{-i})\right) - \sum_{j \neq i} \hat{v}_j(\chi(\hat{v}))$$

Questions:

- who pays 0?
  - agents who don't affect the outcome
- who pays more than 0?

VCG discussion

$$\chi(\hat{v}) = \arg\max_{x} \sum_{i} \hat{v}_i(x)$$

$$p_i(\hat{v}) = \sum_{j \neq i} \hat{v}_j\left(\chi(\hat{v}_{-i})\right) - \sum_{j \neq i} \hat{v}_j(\chi(\hat{v}))$$

Questions:

- who pays 0?
  - agents who don't affect the outcome
- who pays more than 0?
  - (pivotal) agents who make things worse for others by existing

VCG discussion

$$\chi(\hat{v}) = \arg\max_x \sum_i \hat{v}_i(x)$$

$$p_i(\hat{v}) = \sum_{j \neq i} \hat{v}_j\left(\chi(\hat{v}_{-i})\right) - \sum_{j \neq i} \hat{v}_j(\chi(\hat{v}))$$

Questions:

- who pays 0?
  - agents who don't affect the outcome
- who pays more than 0?
  - (pivotal) agents who make things worse for others by existing
- who gets paid?

VCG discussion

$$\chi(\hat{v}) = \arg\max_x \sum_i \hat{v}_i(x)$$

$$p_i(\hat{v}) = \sum_{j \neq i} \hat{v}_j\left(\chi(\hat{v}_{-i})\right) - \sum_{j \neq i} \hat{v}_j(\chi(\hat{v}))$$

Questions:

- who pays 0?
  - agents who don't affect the outcome
- who pays more than 0?
  - (pivotal) agents who make things worse for others by existing
- who gets paid?
  - (pivotal) agents who make things better for others by existing

VCG properties

$$\chi(\hat{v}) = \arg\max_x \sum_i \hat{v}_i(x)$$

$$p_i(\hat{v}) = \sum_{j \neq i} \hat{v}_j\left(\chi(\hat{v}_{-i})\right) - \sum_{j \neq i} \hat{v}_j(\chi(\hat{v}))$$

- Because only pivotal agents have to pay, VCG is also called the pivot mechanism
- It's dominant-strategy truthful, because it's a Groves mechanism

Selfish routing example

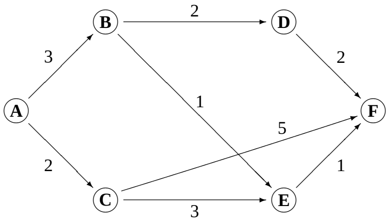

- What outcome will be selected by χ?

Selfish routing example

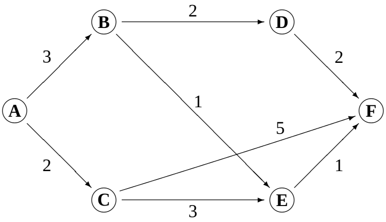

- What outcome will be selected by χ? path $ABEF$.

Selfish routing example

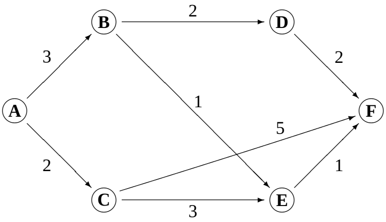

- What outcome will be selected by χ? path $ABEF$.
- How much will AC have to pay?

Selfish routing example

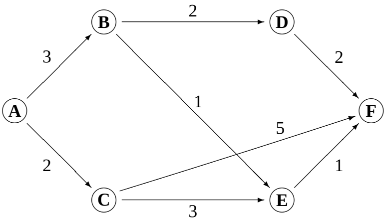

- What outcome will be selected by χ? path $ABEF$.
- How much will AC have to pay?
  - The shortest path taking his declaration into account has a length of 5, and imposes a cost of -5 on agents other than him (because it does not involve him). Likewise, the shortest path without AC's declaration also has a length of 5. Thus, his payment $p_{AC} = (-5) - (-5) = 0$.
  - This is what we expect, since AC is not pivotal.
  - Likewise, BD, CE, CF and DF will all pay zero.

Selfish routing example

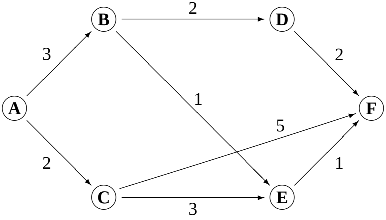

- How much will AB pay?

Selfish routing example

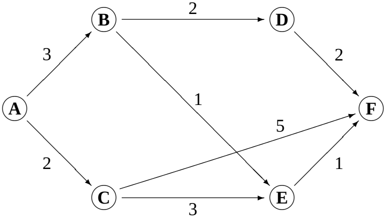

- How much will AB pay?
  - The shortest path taking AB's declaration into account has a length of 5, and imposes a cost of 2 on other agents.
  - The shortest path without AB is $ACEF$, which has a cost of 6.
  - Thus $p_{AB} = (-6) - (-2) = -4$.

Selfish routing example

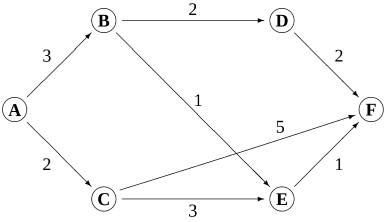

- How much will BE pay?

Selfish routing example

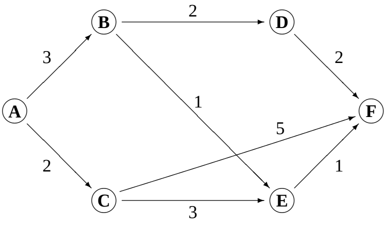

- How much will BE pay? $p_{BE} = (-6) - (-4) = -2.$

Selfish routing example

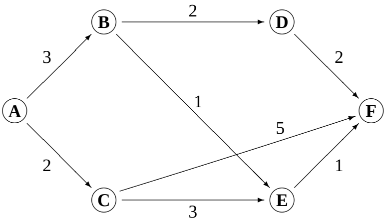

- How much will BE pay? $p_{BE} = (-6) - (-4) = -2$.
- How much will EF pay?

Selfish routing example

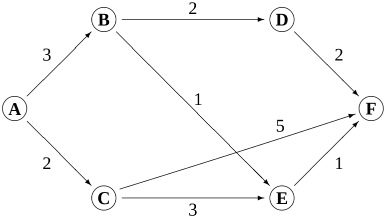

- How much will BE pay? $p_{BE} = (-6) - (-4) = -2.$
- How much will EF pay? $p_{EF} = (-7) - (-4) = -3.$

Selfish routing example

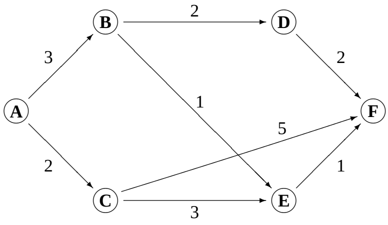

- How much will BE pay? $p_{BE} = (-6) - (-4) = -2$.
- How much will EF pay? $p_{EF} = (-7) - (-4) = -3$.
  - EF and BE have the same costs but are paid different amounts. Why?

Selfish routing example

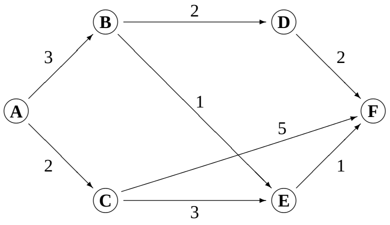

- How much will BE pay? $p_{BE} = (-6) - (-4) = -2$.
- How much will EF pay? $p_{EF} = (-7) - (-4) = -3$.
  - EF and BE have the same costs but are paid different amounts. Why?
  - EF has more *market power*: for the other agents, the situation without EF is worse than the situation without BE.

Overview

VCG and Individual Rationality

Definition (Choice-set monotonicity)

An environment exhibits choice-set monotonicity if $\forall i$, $X_{-i} \subseteq X$.

- removing any agent weakly decreases—that is, never increases—the mechanism's set of possible choices X

Definition (No negative externalities)

An environment exhibits no negative externalities if $\forall i \forall x \in X_{-i},\ v_i(x) \geq 0$.

- every agent has zero or positive utility for any choice that can be made without his participation

Theorem

The VCG mechanism is ex-post individual rational when the choice set monotonicity and no negative externalities properties hold.

Example: road referendum

Example

Consider the problem of holding a referendum to decide whether or not to build a road.

- The set of choices is independent of the number of agents, satisfying choice-set monotonicity.
- No agent negatively values the project, though some might value the situation in which the project is not undertaken more highly than the situation in which it is.

Example: simple exchange

Example

Consider a market setting consisting of agents interested in buying a single unit of a good such as a share of stock, and another set of agents interested in selling a single unit of this good. The choices in this environment are sets of buyer-seller pairings (prices are imposed through the payment function).

- If a new agent is introduced into the market, no previously-existing pairings become infeasible, but new ones become possible; thus choice-set monotonicity is satisfied.
- Because agents have zero utility both for choices that involve trades between other agents and no trades at all, there are no negative externalities.

VCG and weak budget balance

Definition (No single-agent effect)

An environment exhibits no single-agent effect if $\forall i$, $\forall v_{-i}$, $\forall x \in \arg\max_y \sum_j v_j(y)$ there exists a choice x' that is feasible without i and that has $\sum_{j \neq i} v_j(x') \geq \sum_{j \neq i} v_j(x)$.

Example

Consider a single-sided auction. Dropping an agent just reduces the amount of competition, making the others better off.

Theorem

The VCG mechanism is weakly budget-balanced when the no single-agent effect property holds.

Drawbacks of VCG

1. Agents must fully disclose private information.
2. VCG is susceptible to collusion.
3. VCG is not "frugal": prices can be many times higher than the true value of the best allocation involving no winning agents.
4. Excluding bidders can (unboundedly) increase revenue.
5. It is impossible to return all of VCG's revenue to the agents without distorting incentives.
6. The problem of identifying the argmax can be computationally intractable.

Budget Balance and Efficiency

Theorem (Green–Laffont; Hurwicz)

No dominant-strategy incentive-compatible mechanism is always both efficient and weakly budget balanced, even if agents are restricted to the simple exchange setting.

Theorem (Myerson–Satterthwaite)

No Bayes-Nash incentive-compatible mechanism is always simultaneously efficient, weakly budget balanced and ex-interim individual rational, even if agents are restricted to quasilinear utility functions.

Overview

Motivation

- Auctions are any mechanisms for allocating resources among self-interested agents
- Very widely used
  - government sale of resources
  - privatization
  - stock market
  - request for quote
  - FCC spectrum
  - real estate sales
  - eBay

CS Motivation

- resource allocation is a fundamental problem in CS
- increasing importance of studying distributed systems with heterogeneous agents
- markets for:
  - computational resources (JINI, etc.)
  - P2P systems
  - network bandwidth
- currency needn't be real money, just something scarce
  - that said, real money trading agents are also an important motivation

Overview

Some Canonical Auctions

- English
- Japanese
- Dutch
- Sealed Bid

English Auction

English Auction

- auctioneer starts the bidding at some "reservation price"
- bidders then shout out ascending prices
- once bidders stop shouting, the high bidder gets the good at that price

Japanese Auction

Japanese Auction

- Same as an English auction except that the auctioneer calls out the prices
- all bidders start out standing
- when the price reaches a level that a bidder is not willing to pay, that bidder sits down
  - once a bidder sits down, they can't get back up
- the last person standing gets the good

- analytically more tractable than English because jump bidding can't occur
  - consider the branching factor of the extensive form game...

Dutch Auction

Dutch Auction

- the auctioneer starts a clock at some high value; it descends
- at some point, a bidder shouts "mine!" and gets the good at the price shown on the clock

Sealed-Bid Auctions

First-Price Auction

- bidders write down bids on pieces of paper
- auctioneer awards the good to the bidder with the highest bid
- that bidder pays the amount of his bid

Second-Price Auction

- bidders write down bids on pieces of paper
- auctioneer awards the good to the bidder with the highest bid
- that bidder pays the amount bid by the second-highest bidder

Overview

1. Mechanism Design with Unrestricted Preferences
2. Quasilinear Preferences
3. Efficient Mechanisms
4. Single-Good Auctions
   - Canonical auction families
   - Auctions as Bayesian mechanisms
   - Second-price auctions
   - First-price auctions
   - Revenue equivalence
5. Position Auctions
6. Combinatorial Auctions

Modeling an auction as a Bayesian mechanism

- The possible outcomes O consist of all possible ways of allocating the good—the set of choices X—and of charging the agents.
- The agents' action sets vary in different auction types.
  - In a sealed-bid auction, each set A_i is an interval from $\mathbb{R}$: the declaration of a bid amount between some minimum and maximum value.
  - A Japanese auction is an imperfect-information extensive-form game with chance nodes, and so A_i is the space of all policies i could follow.
- χ and p depend on the objective of the auction, such as achieving an efficient allocation or maximizing revenue.
- common prior: agent's valuations are drawn independently from a known distribution ("independent private values" model)

Overview

Second-Price

Theorem

Truth-telling is a dominant strategy in a second-price auction.

- In fact, we know this already (do you see why?)
- However, we'll look at a simpler, direct proof.

Second-Price proof

Theorem

Truth-telling is a dominant strategy in a second-price auction.

Proof.

Assume that the other bidders bid in some arbitrary way. We must show that i's best response is always to bid truthfully. We'll break the proof into two cases:

1. Bidding honestly, i would win the auction
2. Bidding honestly, i would lose the auction

Second-Price proof (2)

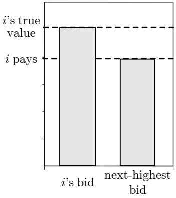

- Bidding honestly, i is the winner

Second-Price proof (2)

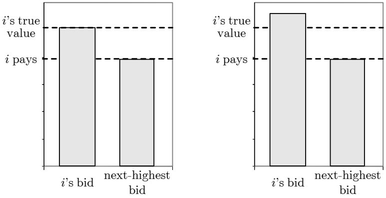

- Bidding honestly, i is the winner
- If i bids higher, he will still win and still pay the same amount

Second-Price proof (2)

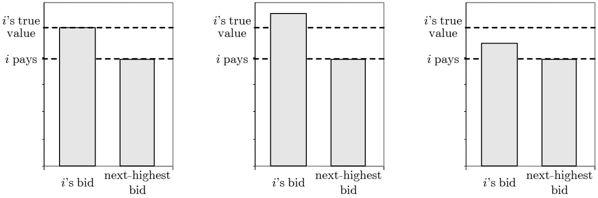

- Bidding honestly, i is the winner
- If i bids higher, he will still win and still pay the same amount
- If i bids lower, he will either still win and still pay the same amount. . .

Second-Price proof (2)

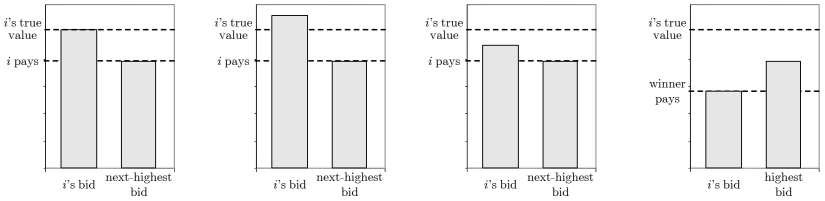

- Bidding honestly, i is the winner
- If i bids higher, he will still win and still pay the same amount
- If i bids lower, he will either still win and still pay the same amount. . . or lose and get utility of zero.

Second-Price proof (3)

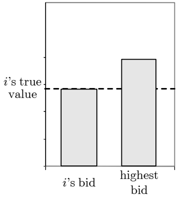

- Bidding honestly, i is not the winner

Second-Price proof (3)

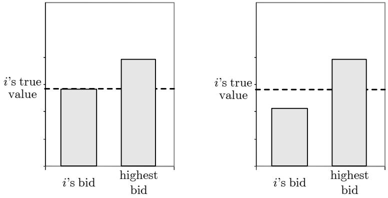

- Bidding honestly, i is not the winner
- If i bids lower, he will still lose and still pay nothing

Second-Price proof (3)

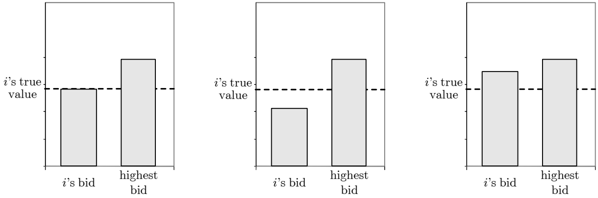

- Bidding honestly, i is not the winner
- If i bids lower, he will still lose and still pay nothing
- If i bids higher, he will either still lose and still pay nothing. . .

Second-Price proof (3)

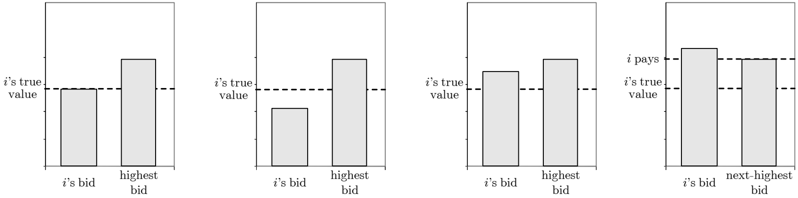

- Bidding honestly, i is not the winner
- If i bids lower, he will still lose and still pay nothing
- If i bids higher, he will either still lose and still pay nothing. . . or win and pay more than his valuation.

English and Japanese auctions

- A much more complicated strategy space
  - extensive form game
  - bidders are able to condition their bids on information revealed by others
  - in the case of English auctions, the ability to place jump bids
- intuitively, though, the revealed information doesn't make any difference in the IPV setting.

English and Japanese auctions

- A much more complicated strategy space
  - extensive form game
  - bidders are able to condition their bids on information revealed by others
  - in the case of English auctions, the ability to place jump bids
- intuitively, though, the revealed information doesn't make any difference in the IPV setting.

Theorem

Under the independent private values model (IPV), it is a dominant strategy for bidders to bid up to (and not beyond) their valuations in both Japanese and English auctions.

Overview

First-Price and Dutch

Theorem

First-Price and Dutch auctions are strategically equivalent.

- In both first-price and Dutch, a bidder must decide on the amount he's willing to pay, conditional on having placed the highest bid.
  - despite the fact that Dutch auctions are extensive-form games, the only thing a winning bidder knows about the others is that all of them have decided on lower bids
    - e.g., he does not know *what* these bids are
    - this is exactly the thing that a bidder in a first-price auction assumes when placing his bid anyway.
- Note that this is a stronger result than the connection between second-price and English.

Discussion

- So, why are both auction types held in practice?
  - First-price auctions can be held asynchronously
  - Dutch auctions are fast, and require minimal communication: only one bit needs to be transmitted from the bidders to the auctioneer.
- How should bidders bid in these auctions?

Discussion

- So, why are both auction types held in practice?
  - First-price auctions can be held asynchronously
  - Dutch auctions are fast, and require minimal communication: only one bit needs to be transmitted from the bidders to the auctioneer.
- How should bidders bid in these auctions?
  - They should clearly bid less than their valuations.
  - There's a tradeoff between:
    - probability of winning
    - amount paid upon winning
  - Bidders don't have a dominant strategy anymore.

Equilibrium

- First-price auctions are not incentive compatible
  - hence, unsurprisingly, not equivalent to second-price auctions

Theorem

In a first-price sealed bid auction with n risk-neutral agents whose valuations are independently drawn from a uniform distribution on the same bounded interval of the real numbers, the unique symmetric equilibrium is given by the strategy profile $(\frac{n-1}{n}v_1, \ldots, \frac{n-1}{n}v_n)$.

- This equilibrium can be verified using straightforward but somewhat involved calculus
- But, how do we identify such an equilibrium in the first place?

Overview

Revenue Equivalence

- Which auction should an auctioneer choose? To some extent, it doesn't matter...

Theorem (Revenue Equivalence Theorem)

Assume that each of n risk-neutral agents has an independent private valuation for a single good at auction, drawn from a common cumulative distribution $F(v)$ that is strictly increasing and atomless on $[\underline{v}, \bar{v}]$. Then any auction mechanism in which

- *the good will be allocated to the agent with the highest valuation; and*
- *any agent with valuation $\underline{v}$ has an expected utility of zero;*

yields the same expected revenue, and hence results in any bidder with valuation v making the same expected payment.

First and Second-Price Auctions

- The k^{th} order statistic of a distribution: the expected value of the k^{th}-largest of n draws.
- For n IID draws from $[0, v_{max}]$, the k^{th} order statistic is

$$\frac{n+1-k}{n+1} v_{max}.$$

First and Second-Price Auctions

- The k^{th} order statistic of a distribution: the expected value of the k^{th}-largest of n draws.
- For n IID draws from $[0, v_{max}]$, the k^{th} order statistic is

$$\frac{n+1-k}{n+1} v_{max}.$$

- Thus in a second-price auction, the seller's expected revenue is

$$\frac{n-1}{n+1} v_{max}.$$

First and Second-Price Auctions

- The k^{th} order statistic of a distribution: the expected value of the k^{th}-largest of n draws.
- For n IID draws from $[0, v_{max}]$, the k^{th} order statistic is

$$\frac{n+1-k}{n+1} v_{max}.$$

- Thus in a second-price auction, the seller's expected revenue is

$$\frac{n-1}{n+1} v_{max}.$$

- First and second-price auctions satisfy the requirements of the revenue equivalence theorem
  - every symmetric game has a symmetric equilibrium
  - in a symmetric equilibrium of this auction game, higher bid $\Leftrightarrow$ higher valuation

Applying Revenue Equivalence

- Thus, a bidder in a FPA must bid his expected payment conditional on being the winner of a second-price auction
  - this conditioning will be correct if he does win the FPA; otherwise, his bid doesn't matter anyway
  - if v_i is the high value, there are then $n-1$ other values drawn from the uniform distribution on $[0, v_i]$
  - thus, the expected value of the second-highest bid is the first-order statistic of $n-1$ draws from $[0, v_i]$:

$$\frac{n+1-k}{n+1}v_{max} = \frac{(n-1)+1-(1)}{(n-1)+1}(v_i) = \frac{n-1}{n}v_i$$

- This provides a basis for our earlier claim about n-bidder first-price auctions.
  - However, we'd still have to check that this is an equilibrium
  - The revenue equivalence theorem doesn't say that every revenue-equivalent strategy profile is an equilibrium!

Overview

1. Mechanism Design with Unrestricted Preferences
2. Quasilinear Preferences
3. Efficient Mechanisms
4. Single-Good Auctions
5. Position Auctions
6. Combinatorial Auctions

Position Auctions

- Search engines make most of their money—billions of dollars—by selling advertisements through position auctions.
  - Keyword-specific "slots" in a list on the right-hand side of a page of search results are simultaneously offered to advertisers.
  - Slots are more valuable the closer they are to the top: more likely to be clicked.
  - Every time a user searches for a keyword, an auction is held.
  - Advertisers pay only if a user clicks on their ad.

Formal Model

- Define the setting:
  - N: the set of bidders (advertisers)
  - v_i: i's (commonly known) valuation for a click
  - $b_i \in \mathbb{R}_+$: i's bid
  - $b_{(j)}$: the jth-highest bid, or 0 if there are fewer than j bids
  - $G = \{1, \ldots, m\}$: the set of goods (slots)
  - α_j: the expected number of clicks (the click-through rate) that an ad will receive if it is listed in the ith slot
- Observe:
  - α does not depend on a bidder's identity
  - the auction is modeled as unrepeated
  - we assume that agents know each other's valuations

Generalized First-Price Auctions

- The generalized first-price auction was the first position auction to be used by search engines.

Definition (Generalized first-price auction)

The generalized first-price auction (GFP) awards the bidder with the jth-highest bid the jth slot. If bidder i's ad receives a click, he pays the auctioneer b_i.

- These auctions do not always have pure-strategy equilibria, even in the unrepeated, perfect-information case.
  - if bidders bid by best responding to each other, their bid amounts can cycle: a low bidder increases bids to try to get a slot; he is outbid by a high bidder; eventually the low bidder drops out; the high bidder reduces his bid; . . .

Generalized Second-Price Auctions

- The instability of bidding under the GFP led to the introduction of the generalized second-price auction.
- It is now the dominant mechanism in practice.

Definition (Generalized second-price auction)

The generalized second-price auction (GSP) awards the bidder with the jth-highest bid the jth slot. If bidder i's ad is ranked in slot j and receives a click, he pays the auctioneer $b_{(j+1)}$.

VCG in the position auction setting

- GSP seems very similar to the VCG mechanism. However, these two mechanisms are actually quite different, as becomes clear when we apply the VCG formula to the position auction setting.

Definition (VCG)

In the position auction setting, the VCG mechanism awards the bidder with the jth-highest bid the jth slot. If bidder i's ad is ranked in slot j and receives a click, he pays the auctioneer $\frac{1}{\alpha_j}\sum_{k=j+1}^{m+1} b_{(k)}(\alpha_{k-1}-\alpha_k)$.

- the key difference: GSP does not charge an agent his social cost, which depends on the differences between click-through rates that other agents would receive with and without his presence.

Equilibria of GSP

- Truthful bidding is not an equilibrium of the GSP.
- Perfect-information setting: the GSP has many equilibria.
  - the most stable configurations will be locally envy free: no bidder will wish that he could switch places with the bidder who won the slot directly above his.
  - There exists a locally envy-free equilibrium of the GSP that achieves exactly the VCG allocations and payments.
  - All other locally envy-free equilibria lead to higher revenues for the seller, and hence are worse for the bidders.
- Beyond perfect information: one can construct a generalized English auction that corresponds to the GSP, and to show that this English auction has a unique equilibrium in which payoffs are again the same as the VCG payoffs, and the equilibrium is ex post, meaning that it is independent of the underlying valuation distribution.

Overview

Valuations for heterogeneous goods

- now consider a case where multiple, heterogeneous goods are being sold.
- consider the sorts of valuations that agents could have in this case:
  - complementarity: for sets S and T, $v(S \cup T) > v(S) + v(T)$
    - e.g., a left shoe and a right shoe
  - substitutability: $v(S \cup T) < v(S) + v(T)$
    - e.g., two tickets to different movies playing at the same time
- substitutability is relatively easy to deal with
  - e.g., just sell the goods sequentially, or allow bid withdrawal
- complementarity is trickier...

Combinatorial auctions

- running a simultaneous ascending auction is inefficient
  - exposure problem
  - inefficiency due to fear of exposure
- if we want an efficient outcome, why not just run VCG?
  - unfortunately, it again requires solving an NP-complete problem
  - let there be n goods, m bids, sets C_j of XOR bids
  - weighted set packing problem:

$$
\begin{aligned}
\max \sum_{i=1}^{m} x_i p_i & & \\
\text{subject to} \sum_{i \mid g \in S_i} x_i \leq 1 & & \forall g \\
x_i \in \{0,1\} & & \forall i \\
\sum_{k \in C_j} x_k \leq 1 & & \forall j
\end{aligned}
$$

Combinatorial auctions

$$
\begin{aligned}
\max \sum_{i=1}^{m} x_i p_i & & \\
\text{subject to} \sum_{i \mid g \in S_i} x_i \le 1 & & \forall g \\
x_i \in \{0,1\} & & \forall i \\
\sum_{k \in C_j} x_k \le 1 & & \forall j
\end{aligned}
$$

- we don't need the XOR constraints
  - instead, we can introduce "dummy goods" that don't correspond to goods in the auction, but that enforce XOR constraints.
  - amounts to exactly the same thing: the first constraint has the same form as the third

Winner determination problem

How do we deal with the computational complexity of the winner determination problem?

- Require bids to come from a restricted set, guaranteeing that the WDP can be solved in polynomial time
  - problem: these restricted sets are *very* restricted...
- Use heuristic methods to solve the problem
  - this works pretty well in practice, making it possible to solve WDPs with many hundreds of goods and thousands of bids.