

Market Driven Systems - Game Theory

Bo Bernhardsson

Department of Automatic Control LTH,
Lund University

Learning Goals

After the course you should be able to

- Formulate different control and decision situations as games and apply game theoretical ideas
- Calculate mixed strategies for small zero-sum games
- Find Nash and Stackelberg equilibria for multiagent games
- Describe Braess' paradox, Cournot's model of Duopoly, and basic auction theory
- Read literature containing game theoretical concepts



Game Theory - Why?

Useful on many levels:

- Analysis of production - customers - markets - competition
- Distributed control on plant level
 - Distributed control of communication networks (examples: TCP rate control, WCDMA power control)
 - Distributed scheduling of multicore computer systems
 - Robust Control
 - ...

Game Theory

Distributed decision-making leads to situations where new intuition and techniques are needed.

This lecture is intended as a short introduction to some of the central concepts and results.

The hope is that it gives an interest for further study of this fascinating subject.




Contents

- Motivation - Some Game Theory Situations
- Theory of Two Player Zero-Sum Games
- Multiplayer Nash and Stackelberg Equilibria
- Mechanism Design
- Auctions

Game Theory - Where?

The control hierarchy



Game Theory, or “Multiagent decision making”

Optimization theory:

Situation with ONE unit with CENTRALIZED information

Many courses at LTH.

Common theme in many of our other control courses.

What if “decision-making” is done by decentralized units, not having the same optimization criteria, not having the same information?

⇒ **Game theory**

What is Game Theory?

Multi-person decision making

Underlying basic assumptions:


Decision makers pursue well-defined objectives (they are *rational*), and take into account their knowledge or expectations of *other* decision-maker's behavior (they reason *strategically*)

(These assumptions are relaxed in *evolutionary game theory*)

Dynamic: the order in which decisions are made is important (otherwise *static*)

Cooperative: Binding agreements can be made (otherwise *noncooperative*)

Motivating Example - Traffic Optimization



Assume $A + B = 4000$ cars, each choosing quickest route time via A: $A/100 + 45$ min.
time via B: $B/100 + 45$ min.
At equilibrium $A = B = 2000$ and travel time is **65 minutes**

Braess' Paradox

Introduction of the new road has increased travel time by 15 minutes for everyone !

Is this against your intuition?


This is not only a mathematical curiosity. There are several indications that the phenomenon occurs in real world, see references.

Called Braess' paradox or "Cost of Anarchy"

Similar phenomenon can occur in other fields of applications, such as electrical networks and mechanical constructions.

Traffic Planning

A new fast road (0 min) between A and B is then built to shorten travel time.



Noone is now using the 45-min roads (since faster alternatives always exist, taking at most 40 min)

Total travel time is now **40+40=80 minutes**

Flow Control in Communication Networks

m distributed sources in a network want to send information through a common bottleneck node.

Performance degrades when many send simultaneously, either because of increased error probability or increased delay

The players have more or less accurate information about the state of the network

All nodes are greedy, want the highest possible throughput

Is there a strategy, which all can agree on and which does not encourage cheating, that maximizes total throughput in the network?

Air-Traffic Control



Want to give flight references to n airplanes

Airplanes can not change heading or height momentarily; high safety requirements

When giving references to one plane and calculating safety margins, consider all other objects as hostile players trying to do their best to collide with this plane

C. Tomlin, G. J. Pappas, and S. S. Sastry, "Conflict Resolution for Air Traffic Management: A Case Study in Multi-Agent Hybrid Systems," IEEE Trans. Automatic Control

Robust Control (H_∞ theory)

Assume u, v and y are related via

$$\begin{aligned}\dot{x} &= Ax + B_1 u + B_2 v \\ z &= C_1 x \\ y &= C_2 x + Dv\end{aligned}$$

Then the question can be answered by solving a differential game with two players

u the controller
 v nature


and a performance criterium

$$\min_u \max_v \int_0^\infty (z^2 + u^2 - v^2) dt$$

Robust Control

The part G of the system is known and the part Δ is unknown.

Is it possible to find a controller K that stabilizes the system under all perturbations Δ that satisfy $\|\Delta(r)\|^2 \leq \|r\|^2$?



Optimal Laziness

A worker, hired for a fixed wage normalized to 0, can either shirk or work. The cost for him to work is w . The gain for his boss if he works is g .

The boss can choose to inspect the worker at a cost of i . If the worker is caught shirking he has to pay the fine f to the boss.

		boss	
		not inspect	inspect
worker	work	-w,g	-w,g-i
	shirk	0,0	-f,f-i

How often should the worker shirk; how often should the boss inspect?

Market-driven Systems - Game Theory

Contents

- Motivation - Some Game Theory Situations
- [Theory of Two Player Zero-Sum Games](#)
- Multiplayer Nash and Stackelberg Equilibria
- Mechanism Design
- Auctions

Rock-Paper-scissor



		Player Y		
		rock	paper	scissor
Player X	rock	0	1	-1
	paper	-1	0	1
	scissor	1	-1	0

$$\min_x (\max_y A_{xy}) = 1$$

$$\max_y (\min_x A_{xy}) = -1$$

The Lunch Problem

Two polite academics independently choose between two nearby restaurants:

- QUICKY BAR where lunches take 20 minutes
- SLOWFOOD INN where lunch takes 50 minutes.

In case having chosen the same restaurant they spend lunch together.

Academic Y likes the company of X and would like to spend the maximum amount of lunch time together, whereas the opposite applies to X, who would like to minimize lunch time with Y; however being too polite to openly say so.

Find the optimal strategy for X, minimizing average lunch time with Y !

Solution

What if X chooses QUICKY with probability q such that average time spent at QUICKY and SLOWFOOD are equal, i.e.

$$20q = 50(1 - q) \implies q = 5/7$$

This gives in average $100/7 = 14.3$ minutes spent at each restaurant.

Common lunch time will now be in average 14.3 minutes whatever strategy Y chooses.

This is the best strategy for X (assuming Y is rational).

Similarly, by using the same strategy, Y can always achieve 14.3 min whatever X does.

Our first zero-sum game

Player X should pay Player Y an amount selected below

		Player Y			
		col 1	col 2	col3	col4
Player X	row 1	1	1	-5	2
	row 2	2	4	1	5
	row 3	3	9	2	3

Version 1: X chooses row first, then Y chooses column. **What choices should rational X and Y make?**

Version 2: Y chooses column first, then X chooses row. **What choices should rational X and Y make?**

If A_{xy} gives the element on row x and col y , result is either

$$\min_x (\max_y A_{xy}) \quad \text{or} \quad \max_y (\min_x A_{xy})$$

Which formula solves Version 1 and which Version 2?

Two player zero sum games

One can easily prove (exercise) that

$$\min_x (\max_y A_{xy}) \geq \max_y (\min_x A_{xy})$$

This means that it is an advantage to know the other players action before choosing own action.

What if neither player knows the other's action?

Lets solve an example before we describe the theory

Solution

Two player zero-sum game

		Player Y	
		QUICKY	SLOWFOOD
Player X	QUICKY	20	0
	SLOWFOOD	0	50

By going to QUICKY all the time, 20 min can be guaranteed.

Any predictable deviations can be learned by Y and will then increase common lunch time.

But can better than 20 min be achieved by X ?

Two player zero sum games

Let the random strategies be described by the probability vectors

$$x = (x_1 \dots x_n)^T \quad \text{and} \quad y = (y_1 \dots y_n)^T$$

where x and y are probability vectors, i.e. $x_i \geq 0$ and $\sum_i x_i = 1$.

Then the average outcome of the game is given by

$$x^T A y$$

The fundamental result of zero sum games (von Neumann) says that for any matrix A one has

$$\min_x (\max_y x^T A y) = \max_y (\min_x x^T A y)$$

This is defined as the value of the game described by the matrix A

The value of the game can be calculated using Linear Programming software, such as the command `lp` in matlab.

<http://www.hofstra.edu/~matscw/gametheory/games.html>

has a Java-solver for zero-sum matrix games up to size 5×5 . (Note however that the row player is the maximizer there)

A Matlab-solver for zero-sum two player games

For Matlab ver 6.1, uses function
`X=lp(f,A,b,VLB,VUB,X0,N)`
 which solves the linear programming problem:
 $\min f'x$ subject to: $Ax \leq b$, $VLB \leq x \leq VUB$
 where the first N constraints defined
 by A and b are equality constraints

```
function [value,x,y]=game(A)
[nx,ny]=size(A);
first we solve the primal for the minimizer x
bigA = [ones(1,nx) 0; A' -ones(ny,1)];
bigb = [1;zeros(ny,1)];
f = [zeros(nx,1);1];
VLB = [zeros(nx,1);-inf];
VUB = [inf*ones(nx,1);inf];
sol = lp(f,bigA,bigb,VLB,VUB,[],1);
x=sol(1:nx,1);
valuex=sol(nx+1,1);
```

A Matlab-solver for zero-sum two player games

```
then we solve the dual for the maximizer y
bigA = [ones(1,ny) 0; -A ones(nx,1)];
bigb = [1;zeros(nx,1)];
f = [zeros(ny,1);-1];
VLB = [zeros(ny,1);-inf];
VUB = [inf*ones(ny,1);inf];
sol = lp(f,bigA,bigb,VLB,VUB,[],1);
y=sol(1:ny,1);
value=sol(ny+1,1);
if abs(value-valuex)>1e-3 error('bad LP solution'); end
```

Market-driven Systems - Game Theory

Contents

- Motivation - Some Game Theory Situations
- Theory of Two Player Zero-Sum Games
- [Multiplayer Nash and Stackelberg Equilibria](#)
- Mechanism Design
- Auctions

Static N -player Games of Complete Information

The **Normal form** of such a game consists of

- A finite set N of **players**
- for each $i \in N$ a set S_i of **strategies** available for player i
- for each player $i \in N$ a **payoff** function u_i on $S = \times_{i \in N} S_i$, which the players try to **maximize**

Will only discuss **Noncooperative** version, i.e. no binding agreements can be made between players (such as exchanging payoff after the game).

Example Prisoner's Dilemma

Provides insight into the difficulty in maintaining cooperation.

	Don't Confess	Confess
Don't Confess	-1, -1	-10, 0
Confess	0, -10	-7, -7

The first number is the payoff for the row player, the second number the column player. Both players now tries to maximize.

Dominating strategy is "Confess"

Pure Nash Equilibrium

Game $(S_1, \dots, S_N, u_1, \dots, u_N)$

Notation: If $s = (s_1, s_2, \dots, s_N)$ is a vector of (pure) strategies then (s_{-i}, a) denotes the (pure) strategy obtained from s by replacing s_i with a

The (pure) strategies $s^* = (s_1^*, \dots, s_N^*)$ constitute a **Pure Nash Equilibrium** if s_i is a best-response for s_{-i} for all i , i.e.

$$u_i(s_{-i}^*, s_i) \leq u_i(s_{-i}^*, s_i^*) \quad \text{for all } s_i \in S_i$$

"There should be no incitement for one-player deviations"

The only Nash equilibrium in the Prisoner's dilemma is (Confess, Confess)

No pure Nash equilibria in the rock-paper-scissor game

Exercise

What are the Pure Nash equilibria of

	L	C	R
T	0,4	4,0	5,3
M	4,0	0,4	5,3
B	3,5	3,5	6,6

Mixed Strategies and Expected Outcome

Game $\{S_1, \dots, S_N; u_1, \dots, u_N\}$

Suppose $S_i = \{s_{i1}, \dots, s_{iK}\}$. Then a **mixed strategy** for player i is a probability distribution $p_i = (p_{i1}, \dots, p_{iK})$, where $0 \leq p_{ik}$, $k = 1, \dots, K$, and $p_{i1} + \dots + p_{iK} = 1$.

The **expected outcome** U_i of a game given a certain set of mixed strategies p_i , $i = 1, \dots, N$ is given by the expected value of u_i : (here illustrated for the case $N = 2$)

$$U_i := \sum_j \sum_k p_{1j} p_{2k} u_i(s_{1j}, s_{2k}), \quad i = 1, 2$$

Existence Result



Theorem [Nash] Every static N -player game with complete information has **at least one mixed Nash equilibrium**

Proof: uses fixed point theorem in mathematics

J.F. Nash. Equilibrium Points in n -Person Games. Proc. National Academy of Sciences of the USA, 36:48-49, 1950

There can be several Nash equilibria. Each giving different outcome vectors. Think of them as corresponding to "local minima" in optimization.

Oligopoly

Model assumptions

- ▶ Few producers, so actions of individual competitors must be considered strategically
- ▶ Very many consumers (continuum); will model their common behavior with a demand vs price curve
- ▶ High barriers for entrants, so one needs not consider new firms arising the market
- ▶ every actor has full information of other players production cost, the price-demand curve, etc

Example: Q4 2008, Verizon, AT&T, Sprint Nextel, and T-Mobile together control 89% of the US cellular phone market.

Example – Bach or Stravinsky?

		Player 2	
		Bach	Stravinsky
Player 1	Bach	2,1	0,0
	Stravinsky	0,0	1,2

What are the pure Nash equilibria?

Mixed Nash Equilibrium

The probability vectors $p_1^*, \dots, p_N^*$ constitute a **mixed Nash Equilibrium** if for all players i we have

$$U_i(p_{-i}^*, p_i) \leq U_i(p_{-i}^*, p_i^*)$$

In words: the mixed strategy p_i^* should be a best response to the other players' mixed strategies.

(This is true iff all pure strategies s_i in "the support" of p_i^* yield the same, optimal, value $u_i(p_{-i}^*, s_i)$).

Monopoly/Oligopoly/Perfect Market

Let's look on a market situation with "producers" and "consumers".

Monopoly - from Greek (mono) "alone/single" + (polein) "to sell"

Oligopoly - from Greek (oligoi) "few" + (polein) "to sell"



Cournot Model of Duopoly

Let q_1 and q_2 denote the quantities produced by firms 1 and 2.

Will for simplicity assume linear price vs volume curve

$$\text{Market clearing price} \quad P(q_1, q_2) = a - (q_1 + q_2)$$

$$\text{Production cost} \quad C_i(q_i) = c q_i, \quad i = 1, 2$$

$$\text{Profit} \quad u_i(q_1, q_2) = q_i(a - q_1 - q_2) - c q_i$$

Strategic choice: What quantities should be produced?

Note: There is a similar (Bertrand) model of Duopoly where the strategic choices instead are the prices asked by the two firms.

Monopoly/Duopoly/Perfect Market

Monopoly situation

$$\begin{aligned} \text{Production} \quad q_m^* &= (a - c)/2 \\ \text{Profit} \quad &(a - c)^2/4 \end{aligned}$$

Duopoly situation, unique Nash equilibrium

$$\begin{aligned} \text{Production} \quad q_1^* &= q_2^* = (a - c)/3 \\ \text{Total Profit} \quad &2(a - c)^2/9 \end{aligned}$$

In a "perfect market" (for the buyer), production increases until price equals production cost for the least efficient producer (for which profit can be very small)

$$\begin{aligned} \text{Production} \quad \sum q_i^* &= (a - c) \\ \text{Total Profit} \quad &0 \end{aligned}$$

Stackelberg Games

Two actions should be taken: first one by the leader, then one by the follower. The game is then over.

Let $B_2(a_1)$ be the follower's best response to action a_1 . Assume B_2 always is a singleton. The **stackelberg solution** is the one obtained when the leader solves

$$\max_{a_1 \in A_1} u_1(a_1, B_2(a_1))$$

Stackelberg Model of Cournot Duopoly

First Firm 1 chooses quantity q_1 , then Firm 2 solves

$$\max_{q_2} q_2(a - q_1 - q_2 - c) \Rightarrow B_2(q_1) = (a - q_1 - c)/2$$

So Firm 1 solves

$$\max_{q_1} q_1(a - q_1 - B_2(q_1) - c) \Rightarrow q_1^* = \frac{a - c}{2}, \quad q_2^* = \frac{a - c}{4}$$

Result: Firm 1 is better off, Firm 2 is worse off, total profit is lower.

Note that in nonzero-sum games a player can be worse off having more information (or, more precisely, having it known to the other players that he has more information)

Incomplete Information

Uncertainty about the other players payoff (their "type")

Players' types are stochastic variables, ("drawn by nature")

Player i knows her own type and assigns probabilities

$$p(t_{-i} | t_i)$$

to all the others.

The probabilities $p(t)$ are "common knowledge".

Cournot Model of Duopoly

Note that both firms would prefer a situation where $q_1 = q_2 = q_m/2$, i.e. shared production at the monopoly rate.

This is however an unstable situation, since both firms then would profit from deviating.

Would need binding agreements to be able to sustain the monopoly situation.

Countries often has law against such agreements (cartels).



Stackelberg

If this optimization problem has a unique solution, a_1^* then

$$(a_1^*, B_2(a_1^*))$$


is called the **backward-induction outcome** or the **Stackelberg solution** of the game.

No "noncredible threats" are taken into account. When the second stage arrives, player 2 will respond in a way that is purely in his self-interest.

In zero-sum games it was better to be follower

Depending on the game, both roles can be favorable (exercise)

Graphical Solution



Where are the Stackelberg solutions? The Nash solution?

Example Incomplete Information

Example: Cournot Duopoly Game, where production cost c for firm 2 is either high $c = c_H$ or low $c = c_L$. Assume $c_2 = c_H$ with probability θ

$$\begin{aligned} q_2^*(c_H) &= \arg \max_{q_2} (a - q_1^* - q_2 - c_H)q_2 \\ q_2^*(c_L) &= \arg \max_{q_2} (a - q_1^* - q_2 - c_L)q_2 \end{aligned}$$

Firm 1 anticipates $q_2^*(c_H)$ with prob. θ and solves

$$q_1^* = \arg \max_{q_1} \theta(a - q_1 - q_2^*(c_H) - c)q_1 + (1 - \theta)(a - q_1 - q_2^*(c_L) - c)q_1$$

Cournot Duopoly under Asymmetric Information

Solution is

$$\begin{aligned} q_2^*(c_H) &= \frac{a - 2c_H + c}{3} + \frac{1 - \theta}{6}(c_H - c_L) \\ q_2^*(c_L) &= \frac{a - 2c_L + c}{3} - \frac{\theta}{6}(c_H - c_L) \\ q_1^* &= \frac{a - 2c + \theta c_H + (1 - \theta)c_L}{3} \end{aligned}$$

Compare with solution in perfect information

$$q_i^* = \frac{a - 2c_i + c_j}{3}$$

Market-driven Systems - Game Theory

Contents

- ▶ Motivation - Some Game Theory Situations
- ▶ Theory of Two Player Zero-Sum Games
- ▶ Multiplayer Nash and Stackelberg Equilibria
- ▶ Mechanism Design
- ▶ Auctions

Example

You are an almighty judge who wants to assign a valuable object to either player 1 or 2.

Valuations: The valuations $v_L < v_H$ are known, but unknown which player has the higher valuation.

Goal: Assign the object to the player with the highest valuation.

You have the right to administer fines, but you prefer an outcome where no fines are used.

You must make your decision rules public in advance and then follow them

Can this problem be solved ?

Example: Bilateral Trade

Individual S owns an indivisible object and considers to sell this to prospective buyer B.

The object is worth s to S and b to B.

The valuations are private information. But known (after normalization) to lie in $[0, 1]$.


If the object is sold at price p then the utilities are changed with

- ▶ for S: $p - s$
- ▶ for B: $b - p$

What kind of mechanism could they use to trade?

Cournot Duopoly

Best responses and Nash equilibrium when $c_1 = 1$, $c_H = c_L = c_2 = 0.75$



Compare with previous figure where $c_1 = c_2 = 1$.

What has happened with the production rates? The price ?

Implementation Theory – Mechanism Design

Rather than fix a game and look for the set of outcomes given by some solution concept, one fixes a set of outcomes and look for a game that yields that set of outcomes as equilibria.

Favorable outcomes could e.g. mean maximizing

- ▶ profit for seller or buyer
- ▶ society benefit
- ▶ market efficiency
- ▶ fairness

Also called "inverse game theory"

Solution

Ask player 1 of his evaluation:

- ▶ If v_L assign object to player 2.

Otherwise ask player 2 of her evaluation.

- ▶ If v_L assign to player 1,
- ▶ otherwise assign to player 2, but administer fines $(\epsilon, (v_H + v_L)/2)$.

This will force truth-telling of both players (assuming of course they are rational).




Example: Bilateral Trade

Suppose the individuals drawn from population with valuations randomly located with uniform density on the unit square

Both players have an incentive to lie about their type when the price is negotiated.


Mechanism to optimize total utility: Is there a way to achieve that the object is sold if $b > s$ and not if $b < s$?




Example: Bilateral Trade

One possibility is that S makes a take-it-or-leave-it offer $p_s(s)$ to B. Trade then occurs if $b > p_s(s)$

Another that B makes such an offer $p_b(b)$ to S. Trade occurs if $s < p_b(b)$



Example: Bilateral trade



Trade occurs when $b > s + 1/4$ (dashed)

Valuations

Agents have different valuations

Private versus common values

Values often not perfectly known: secret / random / interdependent

Risk-neutral vs risk-averse

Example of model: Buyer i gains $v_i - p$ if he wins object, otherwise 0, where v_i is his valuation and p the payment.

A risk-neutral buyer maximises $\text{Prob}(\text{win}) \cdot (v_i - E(p|\text{win}))$

A risk-neutral seller maximizes $E(p)$

The Revelation Principle

Gibbard, Green and Laffont, Dasgupta, Myerson

Theorem Any (Bayesian) Nash Equilibrium of any (Bayesian) type game can be represented by an incentive-compatible mechanism.

Proof: Not very difficult

Makes it easier to construct game mechanisms, or prove that they do not exist, since restriction can be made to mechanisms where truth-telling is optimal.

Example: Bilateral trade

A third possibility would be **Double Auction**: Both parties announce a price and if $p_B > p_S$ trade occurs (for example at $(p_B + p_S)/2$).

The double auction has an affine Nash Equilibrium where

$$p_S = \max(2s/3 + 1/4, s)$$

$$p_B = \min(2b/3 + 1/12, b)$$

If $s < 3/4$ seller asks for more than true valuation s

If $b > 1/4$, buyer bids below true valuation b

Example: Bilateral trade

None of the mechanisms suggested leads to trade for the entire region $s < b$, hence they are not optimally socially efficient.

Does a more efficient mechanism exist?

Answer: No !

The impossibility was established by Laffont and Maskin 1979.

In fact, the double auction mechanism described above maximises the potential total gains from trade. No better design exists !

Revelation Principle

Theory shows that type-game design (under very general assumptions) can be restricted to direct mechanisms:

1. The players simultaneously make (possibly dishonest) claims about their types. Player i can claim to be any type τ_i , no matter what true type t_i he is
2. Given claims $\tau_1, \dots, \tau_N$ a result, a pdf over the set of outcomes, is chosen

A direct mechanism in which truth-telling, $\tau_i(t_i) = t_i$, is a Nash equilibrium is called **incentive-compatible**.

Market-driven Systems - Game Theory

Contents

- Motivation - Some Game Theory Situations
- Theory of Two Player Zero-Sum Games
- Multiplayer Nash and Stackelberg Equilibria
- Mechanism Design
- Auctions

Auction Theory

Simple protocols for one-to-many negotiations

Here: One seller, many buyers, single item to be sold

Many versions

- ▶ English Ascending
- ▶ Dutch Descending
- ▶ First-price sealed-bid
- ▶ Second-price sealed-bid (Vickrey)



Auction Design

Auction

- ▶ determines the buyer (highest valuer?)
- ▶ determines the price
- ▶ gets the item sold

Seller usually sets the rules, and can choose between different mechanisms, either to maximize purchase price, or optimize more complicated function of outcome, such as social effectiveness, fairness etc (government)

- ▶ open or closed bids?
- ▶ revise bids?
- ▶ share any information affecting bidders values?
- ▶ etc

Sealed first price

Requires more thought. Must balance

- ▶ risk of bidding much higher than 2nd highest
- ▶ risk of losing profitable opportunity by bidding below at least one bidder
- ▶ risk of bidding more than the item turns out to be worth (if valuation unknown)

With one round sealed bid first price, the bidders should bid lower than their true valuations ("shade their bids"), but how much?

Note: If other players valuations were known, one would never need considering bidding above 2nd highest valuation, so valuation information can now help.

The Winner's Curse for uncertain valuations



Example: Five people are invited to bid for a suitcase of money.

Can't look inside the suitcase, but each given a private estimate of X , the actual amount

It is publicly known that estimates are

$$X - 2, X - 1, X, X + 1, X + 2$$

What do you bid if you get an estimate of 37?

Valuations

Private: Inherent difference between bidders, such as people bidding for an item for personal use without thinking about reselling

Common value: Item has a single true value, winning it would be equally valuable for all, although how rewarding is partly uncertain to bidders at auction time, such as bidding for oil or spectrum rights.

During bidding, actors are trying to guess the true value with different pieces of incomplete information. Might gain information and revise estimate from

- ▶ how many remain in the bidding
- ▶ how aggressively others bid

Open English

Best strategy for bidder i in private value case: Remain in the bidding until the bid gets above your own, perfectly known, valuation v_i .

This is a dominant strategy, which disappears with sealed bids or with more complicated common value information structures

In general the winner pays less than its worth to him.

Not good for auctioneer !

Note: No value for a bidder to know other bidders valuations !



Shading in sealed-bid first price auctions

Example:

Assume a risk-neutral bidder BoB knows that everybody's evaluations lie in the interval $[0, 1]$, with uniform independent distributions. What should he bid in a first-price auction?

Answer:

$$\frac{n-1}{n} \cdot \text{BoB's own valuation}$$

where n is the total number of bidders (everybody bidding so is a Nash equilibrium)

Note: As the number of bidders becomes larger, the bids tend closer to the true valuations.

The Winner's Curse for uncertain valuations

If you knew all five estimates, you could infer the value. But you only know that X could be between 35 and 39.

You know that 37 is on average correct, so you might e.g. choose to bid 36 "to earn 1 in average if you win". This reasoning is however incorrect:

If all bid their estimates minus 1, the winner is the person with the highest estimate $X + 2$, who will bid $X + 1$, to make a loss of 1: the Winner's Curse

Expect that you overestimated the (common) value when you win the auction.

Don't use the mean $E(\text{value})$, use $E(\text{value} \mid \text{auction is won})$



BoB, an aged and wise parent, wishes to sell his used car to one of his (perfectly rational) many children.

His only concern is to make sure the child valuing it the most gets it, the price is secondary to him. But the children might be dishonest, each having an incentive to exaggerate the car's worth to him/her.

BoB would prefer that no information such as valuations or bids be exchanged between the children.

Can this be solved?

Vickery

Answer: Yes, honesty, i.e. bidding one's true valuation, will be a weakly dominant strategy to all other strategies:

Overbidding: No improvement, the only times it changes the outcome, you lose out (think)

Underbidding: No improvement, the only time it changes things, you lose out (dito)

Outcome will be the same as in English auction. But achieved in one round, with sealed bids !

The Revenue Equivalence Theorem

No it wasn't! The most fundamental result in auction theory is (Myerson 1981, Riley and Samuelsson 1981):

Revenue Equivalence Theorem

Assume each of a given number of risk-neutral potential buyers have privately known valuations v in $[\underline{v}, \bar{v}]$ independently drawn from a common strictly increasing atom-less distribution (i.e. pdf has $0 < f(v) < \infty$).

Then any auction mechanism in which

- ▶ the object goes to the buyer with highest valuation
- ▶ a bidder with valuation $\underline{v}$ gets zero expected surplus.

yields **the same expected revenue for the auctioneer**

Proof Sketch

This gives

$$S_i(v) = S_i(\underline{v}) + \int_{\underline{v}}^v P_i(w) dw$$

Any two mechanisms with the same "not-interested" surpluses $S_i(\underline{v})$ and the same $P_i(\cdot)$, as will be the case when the object is given to the highest valuation, will hence lead to the same surplus versus valuation functions, $S_i(v)$.

Hence also, from

$$S_i(v) = vP_i(v) - E(\text{pay}(v)),$$

to the same expected pay(v). QED.

BoB devises the following scheme:

- ▶ asks each child to tell him confidentially (sealed-bid) their value v_i , (value is 0 if child doesn't get the car, don't care who gets it then).
- ▶ promises to give the car to the one with the highest value
- ▶ that person must pay the 2nd highest bidder's valuation

Will this scheme (Vickery), make honesty the best strategy?

Comparison

Natural question from seller's point of view: Which auction type gives the highest revenue?

Open, sealed-bid 1st price, sealed-bid 2nd price?

Example: As above with n bidders with evaluations randomly drawn uniform in $[0, 1]$. Auction outcome is expected value of

open: 2nd highest valuation

Sealed bid 1st price: $\frac{n-1}{n}$. highest valuation (shading)

Sealed bid 2nd price: 2nd highest valuation (truth-telling)

All of these have the same expected value: $\frac{n-1}{n+1}$ (exercise)

What a coincidence! Or ...

Proof Sketch

Let $S_i(v)$ be the expected surplus bidder i obtains in equilibrium from participating in the auction. Let $P_i(v)$ be her probability of receiving the object in equilibrium.

$$S_i(v) = vP_i(v) - E(\text{pay}(v))$$

Since for all $\bar{v}$

$$S_i(v) \geq S_i(\bar{v}) + (v - \bar{v})P_i(\bar{v})$$

(in equilibrium: using strategy for $\bar{v}$ instead of v does not improve outcome)

Using this with $\bar{v}$ slightly below and above v and using that $P_i(v)$ is continuous, one can deduce that

$$\frac{dS_i}{dv} = P_i(v)$$

Exceptions to Revenue-equivalence

Note that the Revenue-equivalence theorem applies to a rather idealised situation. There are many exceptions:

- ▶ For valuations that are not independent
- ▶ Collusion (group of bidders cooperate to keep price down): none of the schemes above are collusion proof
- ▶ Goods might be divisible, e.g. contain several parts where valuations on parts differ between bidders.
- ▶ ...

Control department home page (Education - Doctorate)
contains a PhD course in Game Theory

Wikipedia:

- ▶ Game Theory
- ▶ Braess' paradox:
- ▶ Oligopoly
- ▶ Cournot
- ▶ Auction Theory
- ▶ Mechanism Design

Mechanism Design Theory - Scientific background on
Economic Prize 2007, Royal Swedish Academy of Sciences
(Google it).

Osborne-Rubinstein, A Course in Game Theory

