Summary of Extensive-Form Games

We have considered extensive-form games (i.e., games on trees)

- with perfect information
- with imperfect information
- with chance nodes (and both perfect and imperfect information)

We have considered pure, mixed and behavioral strategies.

We have considered Nash equilibria (NE) and subgame perfect equilibria (SPE) in pure and behavioral strategies.

Summary of Extensive-Form Games (Cont.)

For perfect information we have shown that

- mixed and behavioral strategies are equivalent
- there is a pure strategy SPE in both pure as well as behavioral strategies
- SPE can be computed using backward induction in polynomial time

For imperfect information we have shown that

- mixed and behavioral strategies are not equivalent in general (but they are equivalent for games with perfect recall)
- backward induction can be used to propagate values through "perfect information nodes", but "imperfect information parts" have to be solved by different means
- solving imperfect information games is at least as hard as solving games in strategic-form; however, even in the zero-sum case, most decision problems are NP-hard (for details see the lecture).

Chance nodes do not interfere with any of the above results.

Summary of Extensive-Form Games (Cont.)

Finally, we discussed repeated games. We considered both, finitely as well as infinitely repeated games.

For finitely repeated games we considered the average payoff and discussed existence of pure strategy NE and SPE with respect to existence of NE in the original strategic-form game.

For infinitely repeated games we considered both

- discounted payoff: We have proved that
  - one-shot deviation property is equivalent to SPE
  - "grim trigger" strategy profiles can be used to implement any vector of payoffs strictly dominating payoffs for a Nash equilibrium in the original strategic-form game (Simple Folk Theorem)
- long-run average payoff: We have proved that all feasible and individually rational vectors of payoffs can be achieved by Nash equilibria (a variant of grim trigger)

Games of INcomplete Information

Bayesian Games

Auctions

Auctions

The (General) problem: How to allocate (discrete) resources among selfish agents in a multi-agent system?

Auctions provide a general solution to this problem.

As such, auctions have been heavily used in real life, in consumer, corporate, as well as government settings:

- eBay, art auctions, wine auctions, etc.
- advertising (Google adWords)
- governments selling public resources: electromagnetic spectrum, oil leases, etc.
- ...

Auctions also provide a theoretical framework for understanding resource allocation problems among self-interested agents: Formally, an auction is any protocol that allows agents to indicate their interest in one or more resources and that uses these indications to determine both the resource allocation and payments of the agents.

Auctions: Taxonomy

Auctions may be used in various settings depending on the complexity of the resource allocation problem:

- *Single-item auctions*: Here n bidders (players) compete for a single indivisible item that can be allocated to just one of them. Each bidder has his own private value of the item in case he wins (gets zero if he loses). Typically (but not always) the highest bid wins. How much should he pay?
- *Multiunit auctions*: Here a fixed number of identical units of a homogeneous commodity are sold. Each bidder submits both a number of units he demands and a unit price he is willing to pay. Here also the highest bidders typically win, but it is unclear how much they should pay (pay-as-bid vs uniform pricing)
- *Combinatorial auctions*: Here bidders compete for a set of distinct goods. Each player has a valuation function which assigns values to *subsets* of the set (some goods are useful only in groups etc.) Who wins and what he pays?

(We mostly concentrate on the single-item auctions.)

Single Unit Auctions

There are many single-item auctions, we consider the following well-known versions:

- *open auctions*:
  - *The English Auction*: Often occurs in movies, bidders are sitting in a room (by computer or a phone) and the price of the item goes up as long as someone is willing to bid it higher. Once the last increase is no longer challenged, the last bidder to increase the price wins the auction and pays the price for the item.
  - *The Dutch Auction*: Opposite of the English auction, the price starts at a prohibitively high value and the auctioneer gradually drops the price. Once a bidder shouts "buy", the auction ends and the bidder gets the item at the price.
- *sealed-bid-auction*:
  - *k-th price Sealed-Bid Auction*: Each bidder writes down his bid and places it in an envelope; the envelopes are opened simultaneously. The highest bidder wins and then pays the *k-th maximum bid*. (In a reverse auction it is the k-the minimum.) The most prominent special cases are *The First-Price Auction* and *The Second-Price Auction*.

Single Unit Auctions (Cont.)

Observe that

- the English auction is essentially equivalent to the second price auction if the increments in every round are very small.

  There exists a "continuous" version, called Japanese auction, where the price continuously increases. Each bidder may drop out at any time. The last one who stays gets the item for the current price (which is the dropping price of the "second highest bid").

- similarly, the Dutch auction is equivalent to the first price auction. Note that the bidder with the highest bid stops the decrement of the price and buys at the current price which corresponds to his bid.

Now the question is, which type of auction is better?

Objectives

The goal of the bidders is clear: To get the item at as low price as possible (i.e., they maximize the difference between their private value and the price they pay)

We consider self-interested non-communicating bidders that are rational and intelligent.

There are at least two goals that may be pursued by the auctioneer (in various settings):

- Revenue maximization

  This may lead to auctions that do not always sell the item to the highest bid
- Incentive compatibility: We want the bidders to spontaneously bid their true value of the item

  This means, that such an auction cannot be strategically manipulated by lying.

Auctions vs Games

Consider *single-item sealed-bid auctions* as strategic form games: $G = (N, (B_i)_{i \in N}, (u_i)_{i \in N})$ where

- The set of players N is the set of bidders
- $B_i = [0, \infty)$ where each $b_i \in B_i$ corresponds to the bid b_i

  (We follow the standard notation and use b_i to denote pure strategies (bids))
- To define u_i, we assume that each bidder has his own private value v_i of the item, then given bids $b = (b_1, \ldots, b_n)$:

First Price: $$u_i(b) = \begin{cases} v_i - b_i & \text{if } b_i > \max_{j \neq i} b_j \\ 0 & \text{otherwise} \end{cases}$$

Second Price: $$u_i(b) = \begin{cases} v_i - \max_{j \neq i} b_j & \text{if } b_i > \max_{j \neq i} b_j \\ 0 & \text{otherwise} \end{cases}$$

Is this model realistic? Not really, usually, the bidders are not perfectly informed about the private values of the other bidders.

Can we use (possibly imperfect information) extensive-form games?

Incomplete Information Games

A *(strict) incomplete information game* is a tuple $G = (N, (A_i)_{i \in N}, (T_i)_{i \in N}, (u_i)_{i \in N})$ where

- $N = \{1, \ldots, n\}$ is a set of players,
- Each A_i is a set of *actions* available to player i,
  We denote by $A = \prod_{i=1}^{n} A_i$ the set of all *action profiles* $a = (a_1, \ldots, a_n)$.
- Each T_i is a set of *possible types* of player i,
  Denote by $T = \prod_{i=1}^{n} T_i$ the set of all *type profiles* $t = (t_1, \ldots, t_n)$.
- u_i is a type-dependent payoff function

$$u_i : A_1 \times \cdots \times A_n \times T_i \to \mathbb{R}$$

  Given a profile of actions $(a_1, \ldots, a_n) \in A$ and a type $t_i \in T_i$, we write $u_i(a_1, \ldots, a_n; t_i)$ to denote the corresponding payoff.

A *pure strategy* of player i is a function $s_i : T_i \to A_i$. As before, we denote by S_i the set of all pure strategies of player i, and by S the set of all pure strategy profiles $\prod_{i=1}^{n} S_i$.

Dominant Strategies

- A pure strategy s_i *very weakly dominates* s_i' if for every $t_i \in T_i$ the following holds: For all $a_{-i} \in A_{-i}$ we have

$$u_i(s_i(t_i), a_{-i}; t_i) \geq u_i(s_i'(t_i), a_{-i}; t_i)$$

  A pure strategy s_i *weakly dominates* s_i' if for every $t_i \in T_i$ the following holds: For all $a_{-i} \in A_{-i}$ we have

$$u_i(s_i(t_i), a_{-i}; t_i) \geq u_i(s_i'(t_i), a_{-i}; t_i)$$

  and the inequality is strict for at least one a_{-i}

  (Such a_{-i} may be different for different t_i.)

- A pure strategy s_i *strictly dominates* s_i' if for every $t_i \in T_i$ the following holds: For all $a_{-i} \in A_{-i}$ we have

$$u_i(s_i(t_i), a_{-i}; t_i) > u_i(s_i'(t_i), a_{-i}; t_i)$$

Definition 88

s_i is *(very weakly, weakly, strictly) dominant* if it (very weakly, weakly, strictly, resp.) dominates all other pure strategies.

Nash Equilibrium

In order to generalize Nash equilibria to incomplete information games, we use the following notation: Given a pure strategy profile $(s_1, \dots, s_n) \in S$ and a type profile $(t_1, \dots, t_n) \in T$, for every player i write

$$s_{-i}(t_{-i}) = (s_1(t_1), \dots, s_{i-1}(t_{i-1}), s_{i+1}(t_{i+1}), \dots, s_n(t_n))$$

Definition 89
A strategy profile $s = (s_1, \dots, s_n) \in S$ is an *ex-post-Nash equilibrium* if for *every* $t_1, \dots, t_n$ we have that $(s_1(t_1), \dots, s_n(t_n))$ is a Nash equilibrium in the strategic-form game defined by the t_i's.

Formally, $s = (s_1, \dots, s_n) \in S$ is an *ex-post-Nash equilibrium* if for all $i \in N$ and all $t_1, \dots, t_n$ and all $a_i \in A_i$:

$$u_i(s_1(t_1), \dots, s_n(t_n); t_i) \geq u_i(a_i, s_{-i}(t_{-i}); t_i)$$

Example: Single-Item Sealed-Bid Auctions

Consider *single-item sealed-bid auctions* as strict incomplete information games: $G = (N, (B_i)_{i \in N}, (V_i)_{i \in N}, (u_i)_{i \in N})$ where

- The set of players N is the set of bidders
- $B_i = [0, \infty)$ where each action $b_i \in B_i$ corresponds to the bid b_i
- $V_i = [0, \infty)$ where each type $v_i \in V_i$ corresponds to the private value v_i
- Let $v_i \in V_i$ be the type of player i (i.e. his private value), then given an action profile $b = (b_1, \ldots, b_n)$ (i.e. bids) we define

First Price: $$u_i(b; v_i) = \begin{cases} v_i - b_i & \text{if } b_i > \max_{j \neq i} b_j \\ 0 & \text{otherwise.} \end{cases}$$

Second Price: $$u_i(b; v_i) = \begin{cases} v_i - \max_{j \neq i} b_j & \text{if } b_i > \max_{j \neq i} b_j \\ 0 & \text{otherwise.} \end{cases}$$

Note that if there is a tie (i.e., there are $k \neq \ell$ such that $b_k = b_\ell = \max_j b_j$), then all players get 0.

Are there dominant strategies? Are there ex-post-Nash equilibria?

Second-Price Auction

For every i, we denote by v_i the pure strategy s_i for player i defined by $s_i(v_i) = v_i$.

Intuitively, such a strategy is *truth telling*, which means that the player bids his own private value truthfully.

Theorem 90

Assume the Second-Price Auction. Then for every player i we have that v_i is a weakly dominant strategy. Also, v is the unique ex-post-Nash equilibrium.

Proof. Let us fix a private value v_i and a bid $b_i \in B_i$ such that $b_i \neq v_i$. We show that for all bids of opponents $b_{-i} \in B_{-i}$:

$$u_i(v_i, b_{-i}; v_i) \geq u_i(b_i, b_{-i}; v_i)$$

with the strict inequality for at least one b_{-i}.

Intuitively, assume that player i bids b_i against b_{-i} and compare his payoff with the payoff he obtains by playing v_i against b_{-i}.

There are two cases to consider: $b_i < v_i$ and $b_i > v_i$.

Second-Price Auction (Cont.)

Case $b_i < v_i$: We distinguish three sub-cases depending on b_{-i}.

A. If $b_i > \max_{j \neq i} b_j$, then

$$u_i(b_i, b_{-i}; v_i) = v_i - \max_{j \neq i} b_j = u_i(v_i, b_{-i}; v_i)$$

Intuitively, player i wins and pays the price $\max_{j \neq i} b_j < b_i$. However, then bidding v_i, player i wins and pays $\max_{j \neq i} b_j$ as well.

B. If there is $k \neq i$ such that $b_k > \max_{j \neq k} b_j$, then

$$u_i(b_i, b_{-i}; v_i) = 0 \leq u_i(v_i, b_{-i}; v_i)$$

Moreover, if $b_i < b_k < v_i$, then we get the strict inequality

$$u_i(b_i, b_{-i}; v_i) = 0 < v_i - b_k = u_i(v_i, b_{-i}; v_i)$$

Intuitively, if another player k wins, then player i gets 0 and increasing b_i to v_i does not hurt. Moreover, if $b_i < b_k < v_i$, then increasing b_i to v_i strictly increases the payoff of player i.

C. If there are $k \neq \ell$ such that $b_k = b_\ell = \max_j b_j$, then

$$u_i(b_i, b_{-i}; v_i) = 0 \leq u_i(v_i, b_{-i}; v_i)$$

Intuitively, there is a tie in (b_i, b_{-i}) and hence all players get 0.

Second-Price Auction (Cont.)

Case $b_i > v_i$: We distinguish four sub-cases depending on b_{-i}.

A. If $b_i > \max_{j \neq i} b_j > v_i$, then

$$u_i(b_i, b_{-i}; v_i) = v_i - \max_{j \neq i} b_j < 0 = u_i(v_i, b_{-i}; v_i)$$

So in this case the inequality is strict.

B. If $b_i > v_i \geq \max_{j \neq i} b_j$, then

$$u_i(b_i, b_{-i}; v_i) = v_i - \max_{j \neq i} b_j = u_i(v_i, b_{-i}; v_i)$$

Note that this case also covers $v_i = \max_{j \neq i} b_j$ where decreasing b_i to v_i causes a tie with zero payoff for player i.

C. If there is $k \neq i$ such that $b_k > \max_{j \neq k} b_j > v_i$, then

$$u_i(b_i, b_{-i}; v_i) = 0 = u_i(v_i, b_{-i}; v_i)$$

D. If there are $k \neq k'$ such that $b_k = b_{k'} = \max_j b_j > v_i$, then

$$u_i(b_i, b_{-i}; v_i) = 0 = u_i(v_i, b_{-i}; v_i)$$

First-Price Auction

Consider the First-Price Auction.

Here the highest bidder wins and pays his bid.

Let us impose a (reasonable) assumption that no player bids more than his private.

Question: Are there any dominant strategies?
Answer: No, to obtain a contradiction, assume that s_i is a very weakly dominant strategy.

Intuitively, if player i wins against some bids of his opponents, then his bid is strictly higher than bids of all his opponents. Thus he may slightly decrement his bid and still win with a better payoff.

Formally, assume that all opponents bid 0, i.e., $b_j = 0$ for all $j \neq i$, and consider $v_i > 0$.
If $s_i(v_i) > 0$, then

$$u_i(s_i(v_i), b_{-i}; v_i) = v_i - s_i(v_i) < v_i - s_i(v_i)/2 = u_i(s_i(v_i)/2, b_{-i}; v_i)$$

If $s_i(v_i) = 0$, then

$$u_i(s_i(v_i), b_{-i}; v_i) = 0 < v_i/2 = u_i(v_i/2, b_{-i}; v_i)$$

Hence, s_i cannot be weakly dominant.

First-Price Auction (Cont.)

Question: Is there a pure strategy Nash equilibrium?
Answer: No, assume that $(s_1, \dots, s_n)$ is a Nash equilibrium.

If there are $v_1, \dots, v_n$ such that some player i wins, i.e., his bid $s_i(v_i)$ satisfies $s_i(v_i) > \max_{j \neq i} s_j(v_j)$, then

$$u_i(s_i(v_i), s_{-i}(v_{-i}); v_i) = v_i - s_i(v_i) < v_i - (s_i(v_i) - \varepsilon) = u_i(s_i(v_i) - \varepsilon, s_{-i}(v_{-i}); v_i)$$

for $\varepsilon > 0$ small enough to satisfy $s_i(v_i) - \varepsilon > \max_{j \neq i} s_j(v_j)$
(i.e., player i may help himself by decreasing the bid a bit)

Assume that for no $v_1, \dots, v_n$ there is a winner (this itself is a bit weird). Consider $0 < v_1 < \dots < v_n$. Since there is no winner, there are two players i, j such that $i < j$ satisfying

$$s_j(v_j) = s_i(v_i) \geq \max_{\ell} s_\ell(v_\ell)$$

But then, due to our assumption, $s_j(v_j) = s_i(v_i) \leq v_i < v_j$ and thus

$$u_j(s_j(v_j), s_{-j}(v_{-j}); v_j) = 0 < v_j - (s_j(v_j) + \varepsilon) = u_j(s_j(v_j) + \varepsilon, s_{-j}(v_{-j}); v_j)$$

for $\varepsilon > 0$ small enough to satisfy $s_j(v_j) + \varepsilon < v_j$.
(i.e., player j can help himself by increasing his bid a bit)

Summary

Second Price Auction:

- There is an ex-post Nash equilibrium in weakly dominant strategies
- It is incentive compatible (players are self-motivated to bid their private values)

First Price Auction:

- There are neither dominant strategies, nor ex-post Nash equilibria

Question: Can we modify the model in such a way that First Price Auction has a solution?

Answer: Yes, give the players at least some information about private values of other players.

Bayesian Games

A *Bayesian Game* $G = (N, (A_i)_{i \in N}, (T_i)_{i \in N}, (u_i)_{i \in N}, P)$ where $(N, (A_i)_{i \in N}, (T_i)_{i \in N}, (u_i)_{i \in N})$ is a strict incomplete information game and P is a distribution on types, i.e.,

- $N = \{1, \dots, n\}$ is a set of players,
- A_i is a set of *actions* available to player i,
- T_i is a set of *possible types* of player i,

  Recall that $T = \prod_{i=1}^{n} T_i$ is the set of type profiles, and that $A = \prod_{i=1}^{n} A_i$ is the set of action profiles.
- u_i is a type-dependent payoff function

$$u_i : A_1 \times \cdots \times A_n \times T_i \to \mathbb{R}$$

- P is a *(joint) probability distribution over T* called *common prior*.

  Formally, P is a probability measure over an appropriate measurable space on T. However, I will not go into measure theory and consider only two special cases: finite T (in which case $P : T \to [0, 1]$ so that $\sum_{t \in T} P(t) = 1$) and $T_i = \mathbb{R}$ for all i (in which case I assume that P is determined by a (joint) density function p on $\mathbb{R}^n$).

Bayesian Games: Strategies & Payoffs

A play proceeds as follows:

- First, a type profile $(t_1, \dots, t_n) \in T$ is randomly chosen according to P.
- Then each player i learns his type t_i.
  (It is a common knowledge that every player knows his own type but not the types of other players.)
- Each player i chooses his action based on t_i.
- Each player receives his payoff $u_i(a_1, \dots, a_n; t_i)$.

A *pure strategy* for player i is a function $s_i : T_i \to A_i$.
As before, we use S to denote the set of pure strategy profiles.

Properties

- We assume that u_i depends only on t_i and not on t_{-i}. This is called **private values** model and can be used to model auctions. This model can be extended to **common values** by using $u_i(a_1, \ldots, a_n; t_1, \ldots, t_n)$.
- We assume the *common prior* P. This means that all players have *the same* beliefs about the type profile. This assumption is rather strong. More general models allow each player to have
  - his own individual beliefs about types
  - ... his own beliefs about beliefs about types
  - .... beliefs about beliefs about beliefs about types
  - .....
  - (we get an infinite hierarchy)

  There is a generic result of Harsanyi saying that the hierarchy is not necessary: It is possible to extend the type space in such a way that each player's "extended type" describes his original type as well as all his beliefs.

  (This does not mean that common prior suffices.)

Example: Battle of Sexes

Assume that player 1 may suspect that player 2 is angry with him/her (the choice is yours) but cannot be sure.

In other words, there are two types of player 2 giving two different games.

Formally we have a Bayesian Game $G = (N, (A_i)_{i \in N}, (T_i)_{i \in N}, (u_i)_{i \in N}, P)$ where

- $N = \{1, 2\}$
- $A_1 = A_2 = \{F, O\}$
- $T_1 = \{t_1\}$ and $T_2 = \{t_2^1, t_2^2\}$
- The payoffs are given by

t_1 :

t_2^1

	F	O
F	2, 1	0, 0
O	0, 0	1, 2

t_2^2

	F	O
F	2, 0	0, 2
O	0, 1	1, 0

- $P(t_2^1) = P(t_2^2) = \frac{1}{2}$

Example: Single-Item Sealed-Bid Auctions

Consider *single-item sealed-bid auctions* as Bayesian games: $G = (N, (B_i)_{i \in N}, (V_i)_{i \in N}, (u_i)_{i \in N}, P)$ where

- The set of players $N = \{1, \ldots, n\}$ is the set of bidders
- $B_i = [0, \infty)$ where each action $b_i \in B_i$ corresponds to the bid
- $V_i = \mathbb{R}$ where each type v_i corresponds to the private value
- Let $v_i \in V_i$ be the type of player i (i.e. his private value), then given an action profile $b = (b_1, \ldots, b_n)$ (i.e. bids) we define

  First Price: $$u_i(b; v_i) = \begin{cases} v_i - b_i & \text{if } b_i > \max_{j \neq i} b_j \\ 0 & \text{otherwise.} \end{cases}$$

  Second Price: $$u_i(b; v_i) = \begin{cases} v_i - \max_{j \neq i} b_j & \text{if } b_i > \max_{j \neq i} b_j \\ 0 & \text{otherwise.} \end{cases}$$

- P is a probability distribution of the private values such that $P(v \in [0, \infty)^n) = 1$. For example, we may (and will) assume that each v_i is chosen independently and uniformly from $[0, v_{\max}]$ where $v_{\max}$ is a given number. Then P is uniform on $[0, v_{\max}]^n$.

Finite-Type Bayesian Games: Payoffs

For now, let us assume that each player has only finitely many types, i.e., T is finite.

Given a type profile $t = (t_1, \ldots, t_n)$, we denote by $P(t_{-i} \mid t_i)$ the *conditional probability* that the opponents of player i have the type profile t_{-i} conditioned on player i having t_i, i.e.,

$$P(t_{-i} \mid t_i) := \frac{P(t_i, t_{-i})}{\sum_{t'_{-i}} P(t_i, t'_{-i})}$$

Intuitively, $P(t_{-i} \mid t_i)$ is the maximum information player i may squeeze out of P about possible types of other players once he learns his own type t_i.

Given a pure strategy profile $s = (s_1, \ldots, s_n)$ and a type $t_i \in T_i$ of player i the *expected payoff* for player i is

$$u_i(s; t_i) = \sum_{t_{-i} \in T_{-i}} P(t_{-i} \mid t_i) \cdot u_i(s_1(t_1), \ldots, s_n(t_n); t_i)$$

(this is the conditional expectation of u_i assuming the type t_i of player i)

Example: Battle of Sexes

t_2^1

t_1 :

	F	O
F	$2,1$	$0,0$
O	$0,0$	$1,2$

t_2^2

	F	O
F	$2,0$	$0,2$
O	$0,1$	$1,0$

$P(t_2^1) = P(t_2^2) = \frac{1}{2}$

Consider strategies s_1 of player 1 and s_2 of player 2 defined by

- $s_1(t_1) = F$
- $s_2(t_2^1) = F$ and $s_2(t_2^2) = O$

Then

- $u_1(s_1, s_2; t_1) = \frac{1}{2} \cdot 2 + \frac{1}{2} \cdot 0 = 1$
- $u_2(s_1, s_2; t_2^1) = 1$ and $u_2(s_1, s_2; t_2^2) = 2$

Infinite-Type Bayesian Games: Payoffs

Now assume that for each player i we have $T_i = \mathbb{R}$ and thus that $T = \mathbb{R}^n$. The concrete type is randomly chosen according to P, denote by $\mathbf{t} = (\mathbf{t}_1, \ldots, \mathbf{t}_n)$ the corresponding random vector with distribution P (each $\mathbf{t}_i$ is a random variable giving a type of player i).

Assume that the type $\mathbf{t}$ is absolutely continuous which means that there is a (joint) density function p such that for all rectangles $R = [a_1, b_1] \times \cdots \times [a_n, b_n]$

$$P[\mathbf{t} \in R] = \int_{a_1}^{b_1} \cdots \int_{a_n}^{b_n} p(t_1, \ldots, t_n) dt_n \cdots dt_1$$

Let p_i be the marginal density function of $\mathbf{t}_i$, i.e.,

$$p_i(t_i) = \int_{T_{-i}} p(t_i, t_{-i}) dt_{-i}$$

The conditional density of $\mathbf{t}_{-i} = (\mathbf{t}_1, \ldots, \mathbf{t}_{i-1}, \mathbf{t}_{i+1}, \ldots, \mathbf{t}_n)$ conditioned on $\mathbf{t}_i = t_i$ where $p_i(t_i) > 0$ is

$$p(t_{-i} \mid t_i) = p(t)/p_i(t_i)$$

(Here $t = (t_1, \ldots, t_n)$ is a type profile.)

Infinite-Type Bayesian Games: Payoffs

Given a pure strategy profile $s = (s_1, \dots, s_n)$ and a type $t_i \in T_i$ of player i, the *expected payoff* for player i is

$$u_i(s; t_i) = \int_{T_{-i}} u_i(s_1(t_1), \dots, s_n(t_n); t_i)\, p(t_{-i} \mid t_i)\, dt_{-i}$$

Example: First-Price Auction
Consider the first-price auction as a Bayesian game where the types of players are chosen uniformly and independently from $[0, v_{\max}]$.

Consider a pure strategy profile $v = (v_1/2, \dots, v_n/2)$ (i.e., each player i plays $v_i/2$). What is $u_i(v; v_i)$?

$$\begin{aligned}
u_i(v; v_i) &= P(\text{player } i \text{ wins}) \cdot v_i/2 + P(\text{player } i \text{ loses}) \cdot 0 \\
&= P(\text{all players except } i \text{ bid less than } v_i/2) \cdot v_i/2 \\
&= \left(\frac{v_i}{2v_{\max}}\right)^{n-1} \cdot v_i/2 \\
&= \frac{v_i^n}{2^n v_{\max}^{n-1}}
\end{aligned}$$

Risk Aversion

We assume that players *maximize* their expected payoff. Such players are called **risk neutral**.

In general, there are three kinds of players that can be described using the following experiment. A player can choose between two possibilities: Either get \$50 surely, or get \$100 with probability $\frac{1}{2}$ and 0 with probability $\frac{1}{2}$.

- risk neutral person has no preference
- risk averse person prefers the first alternative
- risk seeking person prefers the second one

Dominance and Nash Equilibria

A pure strategy s_i *weakly dominates* s_i' if for every $t_i \in T_i$ the following holds: For all $s_{-i} \in S_{-i}$ we have

$$u_i(s_i, s_{-i}; t_i) \geq u_i(s_i', s_{-i}; t_i)$$

and the inequality is strict for at least one s_{-i}.

The other modes of dominance are defined analogously. Dominant strategies are defined as usual.

Definition 91

A pure strategy profile $s = (s_1, \ldots, s_n) \in S$ in the Bayesian game is a *pure strategy Bayesian Nash equilibrium* if for each player i and each type $t_i \in T_i$ of player i and every strategy $s_i' \in S_i$ we have that

$$u_i(s_i, s_{-i}; t_i) \geq u_i(s_i', s_{-i}; t_i)$$

Example: Battle of Sexes

t_2^1

t_1 :	F	O
F	2, 1	0, 0
O	0, 0	1, 2

t_2^2

	F	O
F	2, 0	0, 2
O	0, 1	1, 0

$P(t_2^1) = P(t_2^2) = \frac{1}{2}$

Use the following notation: $(X, (Y, Z))$ means that player 1 plays $X \in \{F, O\}$, and player 2 plays $Y \in \{F, O\}$ if his/her type is t_2^1 and $Z \in \{F, O\}$ otherwise.

Are there pure strategy Bayesian Nash equilibria?

$(F, (F, O))$ is a Bayesian NE.

Even though O is preferred by player 2, the outcome (O, O) cannot occur with a positive probability in any BNE.

- To ever meet at the opera, player 1 needs to play O.
- The unique best response of player 2 to O is (O, F)
- But $(O, (O, F))$ is not a BNE:
  - The expected payoff of player 1 at $(O, (O, F))$ is $\frac{1}{2}$
  - The expected payoff of player 1 at $(F, (O, F))$ is 1

Second Price Auction

Consider the second-price sealed-bid auction as a Bayesian game where the types of players are chosen according to an arbitrary distribution.

Proposition 7

In a second-price sealed-bid auction, with any probability distribution P, the truth revealing *profile of bids, i.e., $v = (v_1, \ldots, v_n)$, is a weakly dominant strategy profile.*

Proof.

The exact same proof as for the strict incomplete information games. Indeed, we do not need to assume that the players have a common prior for this! □

First Price Auction

Consider the first-price sealed-bid auction as a Bayesian game with some prior distribution P.

Note that bidding truthfully does *not* have to be a dominant strategy. For example, if player i knows that (with high probability) his value v_i is much larger than $\max_{j \neq i} v_j$, he will not *waste money* and bid less than v_i.

So is there a pure strategy Bayesian Nash equilibrium?

Proposition 8

Assume that for all players i the type of player i is chosen independently and uniformly from $[0, v_{\max}]$. Consider a pure strategy profile $s = (s_1, \ldots, s_n)$ where $s_i(v_i) = \frac{n-1}{n} v_i$ for every player i and every value v_i. Then s is a Bayesian Nash equilibrium.

Proof. We show that $s_i(v_i) = \frac{n-1}{n} v_i$ is the best response to s_{-i} for all i. Let us fix i and consider a pure strategy s_i' of player i.

Fix v_i and define $b_i = s_i'(v_i)$. We show (see the greenboard) that $b_i = \frac{n-1}{n} v_i$ maximizes $u_i(b_i, s_{-i}; v_i)$. This holds for all v_i, and thus $s_i' = s_i$ is the best response to s_{-i}. □

First Price Auction (Cont.)

More generally, assume only that the private values v_i are identically and independently distributed on $[v_{\min}, v_{\max}]$ (this is called **independent private values** model). Let $F(x)$ be the cumulative distribution function of the private value (for each player).

Let us restrict to *strictly increasing strategies*.
Note that this restriction is quite reasonable, intuitively it means, that the higher the private value, the higher is the bid.

Then one may show that there is a symmetric Bayesian Nash equilibrium $(s_1, \ldots, s_n)$ where each s_i is defined by

$$s_i(v_i) = v_i - \frac{\int_{v_{\min}}^{v_{\max}} [F(v_i)]^{n-1} dx}{[F(v_i)]^{n-1}}$$

That is, in particular, the bid is always smaller than the private value.

Expected Revenue

Consider the first and second price sealed-bid auctions. For simplicity, assume that the type of each player is chosen independently and uniformly from $[0, 1]$.

What is the expected revenue of the auctioneer from these two auctions when the players play the corresponding Bayesian NE?

- In the first-price auction, players bid $\frac{n-1}{n} v_i$. Thus the probability distribution of the revenue is

$$F(x) = P(\max_j \frac{n-1}{n} v_j \leq x) = P(\max_j v_j \leq \frac{nx}{n-1}) = \left(\frac{nx}{n-1}\right)^n$$

  It is straightforward to show that then the expected maximum bid in the first-price auction (i.e., the revenue) is $\frac{n-1}{n+1}$.

- In the second-price auction, players bid v_i. However, the revenue is the expected second largest value. Thus the distribution of the revenue is

$$F(x) = P(\max_j v_j \leq x) + \sum_{i=1}^{n} P(v_i > x \text{ and for all } j \neq i, v_j \leq x)$$

  Amazingly, this also gives the expectation $\frac{n-1}{n+1}$.

Revenue Equivalence (Cont.)

The result from the previous slide is a special case of a rather general **revenue equivalence theorem**, first proved by Vickrey (1961) and then generalized by Myerson (1981).

Both Vickrey and Myerson were awarded Nobel Prize in economics for their contribution to the auction theory.

Theorem 92 (Revenue Equivalence)

Assume that each of n risk-neutral players has independent private values drawn from a common cumulative distribution function $F(x)$ which is continuous and strictly increasing on an interval $[v_{\min}, v_{\max}]$ (the probability of $v_i \notin [v_{\min}, v_{\max}]$ is zero). Then any efficient auction mechanism in which any player with value $v_{\min}$ has an expected payoff zero yields the same expected revenue.

Here efficient means that the auction has a symmetric and increasing Bayesian Nash equilibrium and always allocates the item to the player with the highest bid.