Auctions and Mechanism Design **Winter Semester, 2012/2013**

Lecture 2: October 30

Lecturer: Amos Fiat *Scribe: Nir Hemed*

2.1 Dominant Strategy

Let us consider a case where we are interested in **selling a product**.
There are various mechanisms that we can use to do so:

- Sealed bid first price auction *(highest bidder wins)*
- Sealed bid second price auction *(highest bidder wins, but the price paid is the second-highest bid)*
- Lottery
- English auction *("going once, going twice, sold!")*
- Dutch auction *(similar to English auction, only here the price decreases)*

It is interesting to differentiate between such mechanisms from the agents perspective - can an agent chooses a bidding strategy that will be better than all others?

Definition A Dominant Strategy (DS) *is a strategy profile which is at least as good as all other strategies regardless of the strategies of all other agents.*

Some examples:

- In a sealed bid first price auction there is no dominant strategy: every agent will bid his value and the utility will be zero.
- In a sealed bid second price auction there exists a dominant strategy: every agent should bid his value. This is in fact a special case of the VCG mechanism studied in lecture 1.
- In an English Auction the dominant strategy will be to stay in the game for as long as the bidding price is lower than the value.

2.2 Complete Information Games

In games of complete information all players are assumed to know precisely the pay-off structure of all other players for all possible outcomes of the game. Formally, the information can be represented as a *pay-off matrix*: a multi-dimensional matrix of vectors, where the dimension is the number of agents (players), and the vectors represent the pay-off for the chosen strategy. So in complete information games all the agents know this matrix.
A classic example of such a game is the *prisoners' dilemma*, the story for which is as follows:
Two prisoners who have jointly committed a crime, are being interrogated in separate quarters. Unfortunately, the interrogators are unable to prosecute either prisoner without a confession. Each prisoner is offered the following deal: If he confesses and their accomplice does not, he will be released and his accomplice will serve the full sentence of ten years in prison. If they both confess, they will share the sentence and serve five years each. If neither confesses, they will both be prosecuted for a minimal offence and each receive a year of prison.
This story can be expressed as the following bimatrix game where entry (a, b) represents row player's pay-off a and column player's pay-off b.

	Silent	Confess
Silent	(-1,-1)	(-5,0)
Confess	(0,-5)	(-3,-3)

A simple thought experiment enables prediction of what will happen in the prisoners' dilemma. Suppose the row player is silent. What should the column player do? Remaining silent as well results in one year of prison while confessing results in immediate release. Clearly confessing is better. Now suppose that the row player confesses. Now what should the column player do? Remaining silent results in ten years of prison while confessing as well results in only five. Clearly confessing is better. In other words, no matter what the row player does, the column player is better of by confessing. We thus see that the strategy profile (confess, confess) is a dominant strategy. It is in fact also a *Dominant Strategy Equilibrium* as we will now define.

2.2.1 Dominant Strategy Equilibrium

- $S = (s_1, ..., s_n)$ is a vector of *strategies*.

- A_i - set of possible *actions* for agent i. Note that $s_i \in A_i$.

- s_i is a *Dominant Strategy* for agent i if for all $S_{-i} = (s_1, ..., s_{i-1}, s_{i+1}, ..., s_n)$ the utility (pay-off) of agent i is maximized by choosing s_i.

- A *Dominant Strategy Equilibrium* (DSE) is a vector $S = (s_1, ..., s_n)$ s.t s_i is a Dominant Strategy for *every* agent i.

As it turns out, DSE is quite a rare phenomenon. A weaker notion of equilibrium was defined by John Forbes Nash, who added the concept of *response to a strategy.*
Let us observe a game called "Battle of the Sexes":
Imagine a couple that agreed to meet this evening, but cannot recall if they will be attending the opera or a football match (and the fact that they forgot is common knowledge). The husband would most of all like to go to the football game. The wife would like to go to the opera. Both would prefer to go to the same place rather than different ones.
A possible pay-off matrix is:

	Opera	Football
Opera	(3,2)	(0,0)
Football	(0,0)	(2,3)

If the couple cannot communicate, where should they go? If the husband knows his wife went to the opera, it is clear that he should go to the opera - this will maximizes his pay-off. Similarly, if he knows the wife went to the football game he should also join her. Note that there is no DSE here. What we have shown is that both strategy profiles (opera, opera) and (football, football) are mutual best responses, a.k.a., *Nash Equilibriums.*

2.2.2 Nash Equilibrium

- s_i - some distribution over A_i.

- S_i - *all* distributions over A_i.

- A *Nash Equilibrium* (NE) is a vector of strategies $S = (s_1, ..., s_n)$ s.t $\forall i, s_i$ is a *best response* for S_{-i}.

Note that $DSE \subsetneq NE$. Also, in the example above, the strategies of the players correspond directly to actions in the game, a.k.a., *pure strategies*. In general, Nash equilibrium strategies can be randomizations over actions in the game, a.k.a., *mixed strategies*.

Theorem 2.1 *(Nash, 1951) There always exists a mixed Nash Equilibrium*

Proof Will be covered in...

2.3 Incomplete Information Games

Now we turn to the case where the pay-off structure of the game is not completely known. We will assume that each agent has some private information and this information affects the pay-off of this agent in the game. We will refer to this information as the agent's "type", as opposed to "value".

2.3.1 Definitions

- T_i is the set of *all* possible types of agent i.
- $t_i \in T_i$ is some possible type of agent i.
- A_i is the set of actions.
- S_i is the set of strategies.
- $s_i \in S_i$ is a function:

$$s_i : T_i \longrightarrow \{\textit{Distributions over actions}\}$$

- In Incomplete Information Games a strategy $s_i \in S_i$ is a *Dominant Strategy* for agent i if $\forall t \in T_i, s_i(t)$ is a best response for *any* $b_{-i} = (a_1, ..., a_{i-1}, a_{i+1}, ..., a_n)$.
Note that this is equivalent of saying that $a_{k \neq i}$ has a positive probability.
- A *Dominant Strategy Equilibrium* (DSE) is a vector $S = (s_1, ..., s_n)$ s.t s_i is a Dominant Strategy (as defined herein) for *every* agent i.

The auctions described in section Dominant Strategy were games of incomplete information where an agent's private type was his value for receiving the item, i.e., $t_i = v_i$. As we described, strategies in the English auction were $s_i(v_i) =$ "drop out when the price exceeds v_i" and strategies in the second-price auction were $s_i(v_i) =$ "bid $b_i = v_i$". We refer to this latter strategy as *truth-telling*. Both of these strategy profiles are in dominant strategy equilibrium for their respective (incomplete information) games.
It is interesting to find an equivalent for Nash Equilibrium of incomplete information games as well.

2.3.2 Bayes-Nash Equilibrium

Recall that equilibrium is a property of a strategy profile. It is in equilibrium if each agent does not want to change his strategy given the other agents' strategies. Meaning, for an agent i, we want to the fix other agent strategies and let agent i optimize his strategy (meaning: calculate his best response for all possible types t_i he may have). This is an ill-specified optimization as just knowing the other agents' strategies is not enough to calculate a best response. Additionally, i's best response depends additionally on i's beliefs on the types of the other agents. The standard economic treatment addresses this by assuming a *common prior*. This is called **Bayesian Updating**.

Definition A Bayesian-Nash Equilibrium (BNE) *for an incomplete information game G and a Common Prior D is a **strategy profile** $S = (s_1, ..., s_n)$ s.t $\forall t \in T_i, s_i(t)$ is a best response when other agents play $s_{-i}(t_{-i})$, and $t_{-i} \sim D_{-i|t_i}$.*

The suffix ($t_{-i} \sim D_{-i|t_i}$) means that the agent utilizes an updated distribution of types given that he knows his own determined value.

In incomplete games we only have a prior assumption about how the actions are distributed, rather than specific knowledge about actions that will be taken by the other (rational) agents. This differentiates BNE from NE.
To illustrate BNE, consider using the **first-price auction** to sell a single item to one of two agents, each with valuation drawn independently and identically from $U[0, 1]$, i.e., $\mathbf{D} = D \times D$ with $D(z) = P_{u \sim D}[u < z] = z$. Here each agent's type is his valuation.
We will calculate the BNE of this game by the "guess and verify" technique.

First, we guess that there is a symmetric BNE with $s_i(z) = \frac{z}{2}, i \in \{1,2\}$. Second, we calculate agent 1's expected utility with value v_i and bid b_i under the standard assumption that the agent's utility u_i is his value less his payment(when he wins).

- $E[u_1] = (v_1 - u_1) \times P[1 \text{ wins}]$
- $P[1 \text{ wins}] = P[b_2 \leq b_1] = P[\frac{v_2}{2} \leq b_1] = P[v_2 \leq 2b_1] = P[D(2b_1)] = 2b_1$
- $\implies E[u_1] = (v_1 - u_1) \times 2b_1 = 2v_1b_1 - 2b_1^2$

Third, we optimize agent 1's bid. Agent 1 with value v_1 should maximize this quantity as a function of b_1, and to do so, can differentiate the function and set its derivative equal to zero. The result is

- $\frac{d}{db_1}(2v_1b_1 - 2b_1^2) = 2v_1 - 4b_1 = 0$
- $\implies$ the optimal bis is $b_1 = \frac{v_1}{2}$

This proves that agent 1 should bid as prescribed if agent 2 does; and vice versa. Thus, we conclude that the guessed strategy profile is in BNE.

Q: What will be the result for n players?

2.3.3 Bayesian Games

In Bayesian games it is useful to distinguish between stages of the game in terms of the knowledge sets of the agents. The three stages of a Bayesian game are:

1. *ex ante* - before agents know their types.
2. *interim* - immediately after the agents learn their types, but before playing in the game.
3. *ex post* - the game is played and the actions of all agents are known.