0368.4170: Cryptography and Game Theory	Ran Canetti and Alon Rosen
Lecture 6: Basics of Game Theory	
25 November 2009	Fall 2009

Scribes: D. Teshler

Lecture Overview

1. What is a Game?
2. Solution Concepts:
   - 2.1. Dominant Strategy
   - 2.2. Pure Nash Equilibrium
   - 2.3. Mixed Strategies

1. What is a Game?

A game is a model of interaction between 2 or more players (decision makers). The model includes the constraints on actions that players can take and the players' interests, but does not specify the actions that the players do take.
We start by defining strategic games in which each decision-maker chooses his plan of action once and for all (the game is played once), and these choices are made simultaneously (not turn based or sequential).

Definition 1 (Strategic Game)**.** *A* strategic game *is a triplet $G = (N, (A_i)_{i\in N}, (u_i)_{i\in N})$ such that:*

1. *$N = 1, ..., n$ is the set of players*
2. *$\forall i \in N, A_i$ is the set of* actions *of player i (legal actions)*
3. *$\forall i \in N, u_i : A_1 \times ... \times A_N \to \mathbb{R}$ is a* utility function *for player i. The outcome of G is the vector $a = (a_1, ..., a_n)$ of the actions actually taken by the players.*

Given the outcome of a game, the payoffs of the players, are completely determined using the u_i functions. In general, a player needs to know all the actions of all of the other players in order to calculate his payoff. But, in certain cases, the players may only care about is their own actions or the actions of a strict subset of players.

The model of a game is quite general. For example, it allows for players to be: Humans, Animals/Flowers, Government/Board of directors, Computer Programs (These are the players that we will talk about).

The models we study assume that decision-makers are "rational" in the sense that they are aware of their alternatives, form expectations about any unknowns, have clear preferences, and choose their action deliberately after some process of optimization. In our case, the preference of all "rational" players is to maximize their own payoff.

Actions: The set of available actions essentially defines the rules of the game, what the players are allowed to do (all the rest is not allowed and thus not explicitly stated). In our model the set of available actions is unrestricted, meaning that any actions which we decide upon are allowed (without size restrictions on the set). In our context, if we are talking about computer programs, we shall consider actions to be finite strings (messages between the players). This will allow us to model protocols via games. In theory we can have an infinite length game (infinite amount of actions) but we will talk only about *finite* games, and consider is an *infinite sequences* of games where each game is final.

Utilities: The utilities define the payoffs of each player given an outcome of the game. In our model these are concrete numerical payoffs (real numbers), along with an induced preference relation for the payoffs (larger is better typically). The preference relations are typically (but not necessarily) total, reflexive and transitive. This reflects an Ordinal (qualitative) point of view where we look at the payoff as a measure to what outcome we prefer more or less. In a Cardinal (quantitative) point of view there is also a measure of how much more or less do we prefer a certain outcome.

Representing a game: In a standard (strategic) form we explicitly state the payoff for each possible outcome $(a_1, ..., a_n) \in A_1 \times ... \times A_n$. A finite strategic game in which there are two players can be described conveniently in a table like the one below:

	L	R
T	(w_1, w_2)	(x_1, x_2)
B	(y_1, y_2)	(z_1, z_2)

One player's actions are identified with the rows and the other player's with the columns. The two numbers in the box formed by row r and column c are the players' payoffs when the row player chooses r and the column player chooses c, the first component being the payoff of the row player, and the second number being the payoff of the column player. Thus, in the game in the figure above, the set of actions of the row player is {T,B} and that of the column player is {L,R}. For example, the row player's payoff from the outcome (T,L) is y_1 ($y_1 = u_1(T, L)$) and the column player's payoff is y_2. If the players' names are 1 and 2 then the convention is that the row player is player 1 and the column player is player 2.
For multiple players, such a representation may be prohibitively large. Other, more manageable (in terms of size), forms of representing games are implicit representation using circuits, extensive form games (we will talk about them) and Graphical Games (a special case of implicit representation, we will not talk about that).

2. Solution Concepts

A game is just a model of interaction, it doesn't specify how to play, what actions should the players take, in order to achieve their objectives (maximizing payoff in our case). The central question that game theory is trying to answer is how players choose their strategies of play. There are two main points of view:

1. Prescriptive - Analyze and specify how players *should* play (Recommendation).
2. Descriptive - Analyze and predict how players *will* play (explanations, prediction). This is the leading point of view in social sciences, economics...

These two approaches are related, we want to design games using the prescriptive approach and analyze them using the descriptive approach. In the end, we want the descriptive analysis to predict that players would play in the way we intended when we designed the game. We focus on the prescriptive approach. This is because our main goal is to design games which have a good equilibrium(s), a good outcome for all players. A solution concept offers recommendations on what actions to take.

Assumptions:

1. The game is given (sometimes it is designed by us). Everybody knows the rules (all available actions). It is well defined and players can't change it.
2. Awareness: All utilities of all of the players are known, and a player knows all actions available to him.
3. Perfect Information: A player knows all the actions of other players. For *strategic form games* this is the same as 2. For *extensive form games* this is not necessary.
4. No computational restrictions (here cryptography will come into play). Is finding a strategy feasible? Is playing a strategy feasible?).
5. Players are rational: they try to maximize payoff.
6. All actions of all players are taken simultaneously. For *strategic form games* this is the same as 2.

When referring to the actions of the players in a strategic game as "simultaneous" we do not necessarily mean that these actions are taken at the same point in time. One situation that can be modeled as a strategic game is the following: The players are at different locations, in front of terminals. First the players' possible actions and payoffs are described publicly (so that they are common knowledge among the players). Then each player chooses an action by sending a message to a central computer; the players are informed of their payoffs when all the messages have been received. This example requires that there be a trusted party in the form of a central computer (may want to use cryptography for that). However, the model of a strategic game is much more widely applicable than this example suggests. For a situation to be modeled as a strategic game it is important only that the players make decisions independently, no player being informed of the choice of any other player prior to making his own decision.

2.1. Dominant Strategy

A player's *strategy* in a game is a complete plan of action, for whatever situation might arise; this fully determines the player's behaviour. A player's strategy will determine the action(s) the player will take at any stage of the game, for every possible history of play up to that stage.

Notation:

- $a_{-i} = (a_1, ..., a_{i-1}, a_{i+1}, ..., a_n)$ - Drop i^{th} element.
- $(a, a_{-i}) = (a_1, ..., a_{i-1}, a, a_{i+1}, ..., a_n)$ - Replace i^{th} element with a.

We start with a basic (and robust) solution concept.

Definition 2 (Dominant Strategy)**.** *A strategy* $a_i \in A_i$ *is* dominant *if* $\forall a_{-i}$ *and* $a'_i \neq a_i :\ u_i(a_i, a_{-i}) \geq u_i(a'_i, a_{-i})$

No matter what other players do, player i maximizes payoff by playing a_i. Thus, the action of every player in a dominant strategy equilibrium is a best response to every collection of actions for the other players, a dominant strategy is optimal no matter what the other players do. The outcome of games which have a dominant strategy, given that all the players are rational, is fairly easy to predict. This is because there is only one "sensible" outcome. However, a dominant strategy doesn't always exist:

	L	R
T	$(1,1)$	$(0,0)$
B	$(0,0)$	$(1,1)$

Here the actions of both players defines the payoff. If the column player decides to play left, the row should play top to maximize his payoff. If the column player decides to play right, the row player should play bottom. Thus, each player's strategy depends on the other player's actions and there is no strategy which is best for all possible actions of the other player.

Example 1. *Prisoner's Dilemma*

Two suspects in a crime are put into separate cells. If they both confess, each will be sentenced to three years in prison. If only one of them confesses, he will be freed and used as a witness against the other, who will receive a sentence of four years. If neither confesses, they will both be convicted of a minor offense and spend one year in prison. Choosing a convenient payoff representation for the preferences, we have the game in the table below:

	Don't Confess	Confess
Don't Confess	$(3,3)$	$(0,4)$
Confess	$(4,0)$	$(1,1)$

The payoff here represents the amount of years in prison. Thus, a higher payoff is worse. The dominant strategy here is "Don't Confess". However, if they confess they will sit less time in jail. Thus, if they play the dominant strategy, they will both get worse results in the end (but they will maximize the winning chance of the other by doing so). Thus, the equilibrium is not always the best strategy in the real world (with our payoff and preference relation). The best strategy here is cooperation. This demonstrates the shortcomings of rationality (may lead to undesirable outcomes).

One example where dominant strategies are useful is mechanism design. It is often desirable to design the game in such a way that the best strategy (for us as the designers) is dominant.

Example 2. *Sealed Bid Auction*

- An object is to be assigned to a player in $\{1, .., n\}$.
- One Auctioneer who holds the object.
- Each player has his own valuation of the object. Player i's valuation of the object is denoted v_i. We further assume that $v_1 > v_2 > ... > 0$.

Game:

1. The players simultaneously submit their bids $b_1, ..., b_n$.
2. Auctioneer opens all bids. The object is given to the player with the highest bid (or to a random player among the ones bidding the highest value).
3. The Utility function for each of the players is as follows:
   - The winner receives his valuation of the object minus the price he pays.
   - Everyone else receives 0.
4. Determining winning player's payment can be done in several ways. For example:
   - Winner pays his bid (first price auction).

- Winner pays *2nd* highest bid (second price auction).

Claim 1. *In the first price auction there is no dominant strategy.*

Proof. To see that there is no dominant strategy, first we argue that there is no dominant strategy in which player 1 does not obtain the object. Suppose that player $i \neq 1$ submits the highest bid b_i and $b_1 < b_i$. If $b_i > v_2$ then player i's payoff is negative, so that he can increase his payoff by bidding 0. If $b_i \leq v_2$ then player 1 can deviate to the bid b_i and win, increasing his payoff. Now let the winning bid be b^*. We have $b^* \geq v_2$, otherwise player 2 can change his bid to some value in (v_2, b^*) and increase his payoff. Also $b^* \leq v_1$, otherwise player 1 can reduce her bid and increase her payoff. Finally, $b_j = b^*$ for some $j \neq 1$ otherwise player 1 can increase his payoff by decreasing his bid. □

Comment. An assumption in the exercise is that in the event of a tie for the highest bid the winner is the player with the lowest index.

Claim 2. *In the second price auction, there exists a dominant strategy.*

Proof. We shall show that the aforementioned strategy of bidding one's valuation (i.e., $b_i = v_i$), in fact, dominates over all other strategies. Consider the following picture proof where B^* represents the maximum of all bids excluding player i's bid, $B^* = max_{j \neq i} b_j$, and v^* is player i's valuation. The vertical axis is utility. The first graph shows the payoff for bidding one's valuation. In the second graph, which represents the case when a player bids lower than their valuation, notice that whenever $bi \leq B^* \leq v^*$,player i receives utility 0 because he loses the auction to whoever bid B^*.If he would have bid his valuation, he would have positive utility in this region (as depicted in the first graph). Similar analysis is made for the case when a player bids more than their valuation.

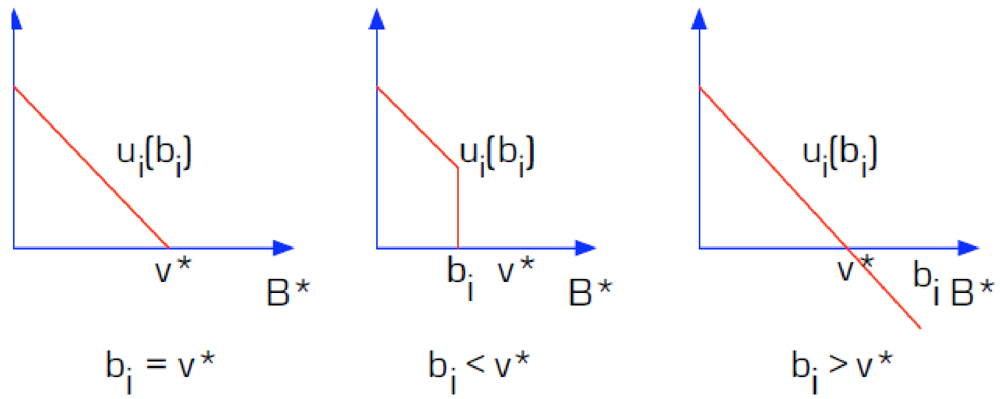

□

Note. In the graphs above, u_i denotes v_i in our notation (the player's value for the item).
Remark. Here the auctioneer knows all of the values (bids). This does not preserve the privacy of the players. We can use cryptography (protocols) to to preserve the players' privacy and hide the values from the auctioneer. A generalization of this game is a Combinatorial Auction (with several items).

2.2. Nash Equilibrium (NE)

The most commonly used solution concept in game theory is that of Nash equilibrium. This notion captures a steady state of the play of a strategic game in which each player holds the correct expectation about the other players' behavior and acts rationally. It does not attempt to examine the process by which a steady state is reached.

Definition 3 (Nash Equilibrium)**.** *A vector of strategies* $a = a(a_1, ..., a_n) \in A$ *is in* Nash Equilibrium *in* $G = (N, (A_i), (u_i))$ *if* $\forall i \in N$ *any strategy* $a'_i \neq a_i$*:* $u_i(a_i, a_{-i}) \geq u_i(a'_i, a_{-i})$

So, for a to be a Nash equilibrium it must be that no player i has an action yielding an outcome that he prefers to that generated when he chooses a_i, given that every other player chooses his equilibrium action a'_i. In other words, given that all players stick to a_{-i}, it doesn't payoff to deviate. That is, a_i is the *best response* to a_{-i}. Note that a dominant strategy is an even stronger form of equilibrium from of NE since it is not only best response for some a_{-i} but rather $\forall a_{-i}$.

Example 3. *Bach or Stravinsky (Battle of the sexes) - BoS:*

Two people wish to go out together to a concert of music by either Bach or Stravinsky. But, one person prefers Bach and the other person prefers Stravinsky. Representing the individuals' preferences by payoff functions, we have the game in the table below:

	Bach	Stravinsky
Bach	$(2,1)$	$(0,0)$
Stravinsky	$(0,0)$	$(1,2)$

This game models a situation in which players wish to coordinate their behavior, but have conflicting interests. The row player prefers Bach and the column player prefers Stravinsky. There is no dominant strategy. The two NE are (B,B) (S,S), however a-priori it is still unclear how to play.

Example 4. *Coordination Game*

As in BoS, two people wish to go out together, but in this case they agree on the more desirable concert. A game that captures this situation is given in the table below:

	Mozart	Mahler
Mozart	$(2,2)$	$(0,0)$
Mahler	$(0,0)$	$(1,1)$

Like BoS, the game has two Nash equilibria: (Mozart; Mozart) and (Mahler; Mahler). In contrast to BoS, the players have a mutual interest in reaching one of these equilibria, namely (Mozart; Mozart); however, the notion of Nash equilibrium does not rule out a steady state in which the outcome is the inferior equilibrium (Mahler; Mahler).

Example 5. *Prisoner's Dilemma*

We shall look back at this game and try to find NEs (besides the existing dominant strategy which we have seen in *Example 1*).

	DC	C
DC	$(3,3)$	$(0,4)$
C	$(4,0)$	$(1,1)$

The dominant strategy is (DC,DC) which is a unique NE. Here the unique NE presents a paradox, it gives us more time in jail (as we have discussed before). Thus, rational reasoning is not always good.

Example 6. *Hawk Dove*

Two animals are fighting over some prey. Each can behave like a dove or like a hawk. The best outcome for each animal is that in which it acts like a hawk while the other acts like a dove; the worst outcome is that in which both animals act like hawks. Each animal prefers to be hawkish if its opponent is dovish and dovish if its opponent is hawkish. A game that captures this situation is shown in the table below:

	D	H
D	$(3,3)$	$(1,4)$
H	$(4,1)$	$(0,0)$

The game has two Nash equilibria, (D,H) and (H,D), corresponding to two different conventions about the player who yields.

Example 7. *Matching Pennies (Game theory variant of coin toss):*

Each of two people chooses either Head or Tail. If the choices differ, person 1 pays person 2 a dollar; if they are the same, person 2 pays person 1 a dollar. Each person cares only about the amount of money that he receives. A game that models this situation is shown in the table below:

	Head	Tail
Head	$(1,-1)$	$(-1,1)$
Tail	$(-1,1)$	$(1,-1)$

Here we have no NE. The game has also a zero sum property (the utilities always sum up to 0 or a constant), thus this is strictly a win-lose game. Such a game, in which the interests of the players are diametrically opposed, is called "strictly competitive". From this example we can see that we don't always have a stable state, or a strategy which leads to it. This is a problem for descriptive analysis.

2.3. Mixed Strategies

Motivated by the Matching Pennies example (Example 7), which has no NE, we consider "mixed" strategies, namely a probability distribution over strategies. For example, we will give each player a probability of $\frac{1}{2}$ to play head or tail. Thus, each player's chance of winning becomes $\frac{1}{2}$. The distributions are known to all players. However, what is actually played (the outcome) is not known in advance.

Notation:

- We let $\Delta(A_i)$ be a set of all of the distributions on A_i (the set of all available mixed strategies using player i's allowed actions A_i).
- Mixed strategy $s_i \in \Delta(A_i)$. A_i is referred to as pure strategy.
- $s_i(x) = Pr[a_i = x]\ x \in A_i$ - the probability of strategy x being played in mixed strategy s_i
- $support(s_i) = set\ of\ all\ x \in A\ s.t.\ s_i(x) > 0$ - The set of all strategies which can be played in mixed strategy s_i

Assumptions:

- Player's mixed strategies are *independently distributed*.
- Preference relation on utility function is total order, meaning that it can generate a payoff for every possible outcome of the mixed strategy.
- Players try to maximize *expected* payoff. In particular, we are not concerned with the distribution of the payoff, but only with its expected value. e.g. a player doesn't care if the distribution is $\begin{cases} 3 & \frac{1}{2} \\ 5 & \frac{1}{2} \end{cases}$ or $4\ w.p.\ 1$.

This is called *Risk Neutral*, where the players are looking to maximize the expected value. The literature sometimes discusses alternatives such as *Risk Seeking/Risk Averse*, where players try to maximize the actual value or maximize the worst case.
One possible critique of the notion is that it is unclear whether people take randomness into consideration since people are bad randomizers. This problem doesn't arise with computer programs which may be considered good randomizers.