Computational Game Theory **Spring Semester, 2011/2012**

Lecture 1: March 7

Lecturer: Amos Fiat *Scribe: Eran Nir and Yael Amsterdamer*

1.1 Computational Game Theory

1.1.1 Agenda

The course will cover the following material:

- Introduction to Game Theory
- Examples
- Matrix form Games
- Utility
- Solution concepts
- Dominant Strategies
- Nash Equilibria
- Complexity
- Mechanism Design: reverse game theory

1.1.2 CGT in Computer Science

The study of Game Theory in the context of Computer Science exists only in the last 15 years. The purpose is to understand problems from the perspective of computability and algorithm design. Computing involves many different selfish entities, and therefore involves game theory. An example to this is the internet:

- Many players (end-users, ISVs, Infrastructure Providers).

- Players wish to maximize their own benefit and act accordingly.
- The trick is to design a system where its beneficial for the player to follow the rules.

The interest in game theory is divided into theory studying and industry use:

Theory:

- Algorithm design
- Complexity
- Quality of game states (Equilibrium states in particular)
- Study of dynamics

Industry:

- Sponsored search
- Other auctions

1.1.3 Game Theory

Some explication in a nutshell:

- Rational Player: Prioritizes possible actions according to utility or cost, and strives to maximize utility or to minimize cost.
- Competitive Environment: More than one player at the same time.

Game Theory analyzes how rational players behave in competitive environments.

1.2 The Prisoner's Dilemma

The prisoner's dilemma is the most known example of a game in which two individuals might not cooperate, even if it appears that it is in their best interest to do so.

The dilemma is as follows: two men are arrested upon committing a crime, but the police don't have enough evidence for conviction. The two are held in separate rooms for interrogation, and offered the same deal by the police: If one of them confesses and testifies against his partner, the betrayer receives 2 years in jail, while the one who kept silent receives a 6 years sentence. if both keep silent, both are sentenced to 3 years in jail (for other charges). if both cooperate with the police, both receive 5 years. The table representation of the game is as follows, where for each pair (x, y), x is the utility of the row player, and y is the utility of the column player:

	Keep silent	Defect
Keep silent	(3,3)	(6,2)
Defect	(2,6)	(5,5)

Since the game is played only once (the situation in the police station is assumed not to return on itself), each player is only concerned with getting less jail time. Regardless of what the other prisoner chooses to do - keep silent or betray - one will benefit more by betraying. The symmetry of the game means both prisoners will betray, therefore receiving the highest possible total jail time of all options (10 total years), while cooperating and keeping silent would've ended with the lowest possible jail time (6).

In such case, where no matter what the other players in the game do, one prefers a specific action, we call this action a "dominant strategy". In this game, betraying is the dominant strategy.

1.3 ISP Routing

Another variation of the Prisoner's Dilemma is the ISP (Internet Service Providers) Routing Game. ISPs often share their physical networks for free, and so in some cases an ISP can either choose to route traffic in its own network or via a partner network. the routing choice made by the originating ISP also affects the load at the destination ISP.

Consider two ISPs, as seen in Figure 1.4, each having its own separate network. The two networks can exchange traffic via two points, A and B. In each ISP There are origin and destination points (s_i, t_i), but the traffic has to cross between the domains. The cost of routing along each edge is 1.

Suppose that ISP_1 needs to send traffic from point s_1 in his own domain to point t_i in ISP_2's domain. ISP_1 has two choices for sending its traffic, either via A or via B. Assume that the ISPs behave selfishly and try to minimize their own costs, by sending traffic to the closest transit point. The ISP with the destination node must route the traffic, no matter where it enters its domain. Using A as a transit point (see 1.5), ISP_1 incurs a cost of 1 unit. ISP_2 has to transmit the traffic via 4 edges, causing a cost of 4. The total cost for both ISPs is 5.

Suppose ISP_2 is not as selfish, and it chooses to use point B as a transit point (see 1.5): In this case the cost for ISP_2 is 2, while the cost caused to ISP_1 is 1. The total cost for both ISPs is 3. By not playing selfishly, one can reduce the total cost of both players.

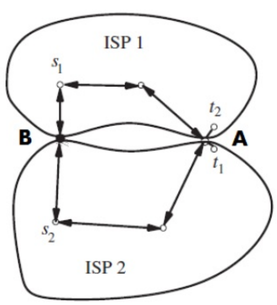

Figure 1.1: Networks of two ISPs, with possible links for traffic

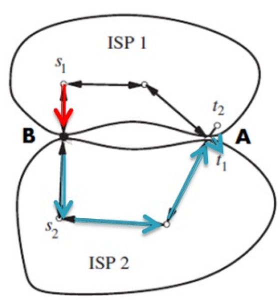

Figure 1.2: Traffic from s_1 to t_1 transferred through ISP_2

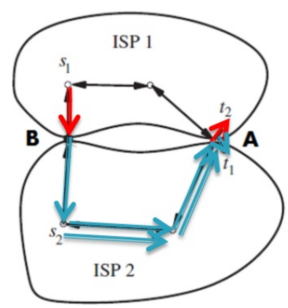

Figure 1.3: Traffic from s_2 to t_2 transferred through ISP_2

If we look at the table representation of the game, we'll see a familiar one:

	Using A	Using B (playing selfishly)
Using A	(3,3)	(6,2)
Using B (playing selfishly)	(2,6)	(5,5)

While keeping the traffic in one's own network as much as possible would've been the best option if both ISPs chose it, the dominant strategy is clearly getting rid of the traffic as soon as possible. Therefore the outcome of this 'dilemma' causes both ISPs pay a cost of 5, instead a cost of 3. The ratio between the selfish solution, to the globally optimal one, is called "the price of anarchy". Later in the course we'll look for ways to keep this price as low as possible.

1.4 Strategic Games

We will now give a formal model to a strategic game under the following assumptions:

- The game consists of only one 'turn' - we'll not refer to games with multiple turns, where considering future dynamics may influence actions taken in current time.

- All the players play simultaneously and are unaware of what the other players do.

- Players are selfish, seek to maximize their own benefit.

1.4.1 Formal Model

The Model is defined as follows:

- There are $N = \{1, ..., n\}$ players.

- Player i has actions/strategies $A_i = \{a_{i1}, a_{i2}, \ldots, a_{im}\}$.

- The space of all possible action vectors is $A = A_1 \times A_2 \times \cdot \times A_n$.

- A joint action is the vector $a \in A$.

- Player i has a utility function $u_i : A \to \mathcal{R}$. If the utility is negative we may call it cost.

In total the game is defined by $\langle N, \langle A_i \rangle_{i=1}^n, \langle u_i \rangle_{i=1}^n \rangle$.

1.4.2 Dominant Strategies

Definition Action a_i of player i is a *weakly* dominant strategy if

$$\forall b_{-i} \in A_{-i}, \forall b_i \in A_i : u_i(b_{-i}, b_i) \leq u_i(b_{-i}, a_i)$$

Definition Action a_i of player i is a *strongly* dominant strategy if

$$\forall b_{-i} \in A_{-i}, \forall b_i \in A_i : u_i(b_{-i}, b_i) < u_i(b_{-i}, a_i)$$

Definition An outcome a of a game is Pareto optimal if for every other outcome b, some player will lose by changing to b. More formally: $\forall b \in A \exists i \in N$ such that $u_i(b) < u_i(a)$.

1.5 Rationality Axioms

1.5.1 St. Petersburg Paradox

The St. Petersburg paradox is a paradox which comes to describe a situation where a decision based only on the expected value of some variable, supposedly a rational decision, is a decision that no rational person would be willing to take. The paradox was published by Bernoulli in 1738, and is named after the magazine it was published in - the Commentaries of the Imperial Academy of Science of Saint Petersburg.

The paradox goes as follows: A game of chance is offered for a single player, in which a fair coin is tossed many times. The prize in the game depends on the number of times n the coin comes up heads, until it first comes up tails. The prize starts at 1 dollar and is doubled every time a head appears. The first time a tail appears, the game ends and the player wins whatever accumulated so far. Thus the player wins 1 dollar if a tail appears on the first toss, 2 dollars if a head appears on the first toss and a tail on the second, 4 dollars if a head appears on the first two tosses and a tail on the third, and so on. In general, the player wins 2^n dollars if the coin is tossed n times until the first tail appears.

The question arouse is what would be a fair price to pay in order to play the game? To answer the question we first want to find what will be the average prize a participant will win. We can see that as the prize is exponentially growing, the probability to win it is exponentially growing smaller: With probability 1/2, the player wins 1 dollar. with probability 1/4 he wins 2 dollars; with probability 1/8 he wins 4 dollars and so on. Thus the expected value of the game is given by:

$$
\begin{aligned}
E[Prize] = \frac{1}{2} * 1 + \frac{1}{4} * 2 + \frac{1}{8} * 4 + \frac{1}{16} * 8 + ... \\
&= \frac{1}{2} + \frac{1}{2} + \frac{1}{2} + \frac{1}{2} + ... \\
&= \sum_{i=1}^{\infty} \frac{1}{2} \\
&= \infty
\end{aligned}
$$

The outcome is since the expected value of the prize is infty, any amount of money would be a fair enough "entering fee" to the game. The paradox os that in reality (obviously) very few people will agree paying a high amount of money to enter such a game. This can be explained by either the low probability (regardless the high expected value) to actually win a big prize, or the fact that the utility of the money in reality is not proportional to the amount. One way to solve the paradox is to define a non-linear utility function over the prize, which give a finite expected value.

1.5.2 Von NeumannMorgenstern Rationality Axioms (1944)

In 1944, (John von) Neumann and (Oskar) Morgenstern introduced four axioms of "rationality" that refer to preferences over options called "lotteries", such that every person that is satisfying the axioms has a utility function. This means that a person is "rational" (by satisfying the axioms) if and only if there exists a real-valued function u defined on possible outcomes such that every preference of the person is given by maximizing the expected value of u. Given lotteries L, M, N, the four rationality axioms are:

- Completeness:

$$L \prec M, \; M \prec L, \text{ or}, \; M = L.$$

- Transitivity:

$$\text{if } L \preceq M \preceq N, \text{ then } L \preceq N.$$

- Continuity:

$$\text{if } L \preceq M \preceq N \text{ then there exists } p \in [0,1] \text{ s.t. } pL + (1-p)N = M.$$

- Independence:

$$\text{if } L \prec M, \text{ then for any } N \text{ and } p \in (0,1] : pL+(1-p)N \prec pM+(1-p)N.$$

As written above, Von Neumann and Morgenstern proved that given those axioms, we have a real-valued utility function over lotteries, and holds: given two lotteries, $u(\alpha L_1 + (1-\alpha)L_2) = \alpha u(L_1) + (1-\alpha)u(L_2)$.

1.5.3 Allais Paradox (1953)

The Allais paradox arises when comparing participants' choices in two different experiments, where each experiment consists of a choice between two gambles. The winning probabilities and prizes are as follows:

	Chance	Winnings
Gamble A	100%	1 Million $
Gamble B	89%	1 Million $
	10%	5 Million $
	1%	Nothing

Figure 1.4: Experiment 1

	Chance	Winnings
Gamble C	11%	1 Million $
	89%	Nothing
Gamble D	10%	5 Million $
	90%	Nothing

Figure 1.5: Experiment 2

Several studies (including one done interactively in class) had shown that given the choice between A and B, many people would choose A, while given the choice between C and D, many people would choose D. However, for the same person to choose both A and D is inconsistent with expected utility theory. A short proof to this claim goes as follows:

$$\begin{aligned} & E(A) = 1M, E(B) = 1.39M \\ u(A) &> u(B) \\ u(1M) &> 0.1 \cdot u(5M) + 0.89u(1M) + 0.01u(0) \\ 0.11 \cdot u(1M) + 0.89 \cdot u(0) &> 0.1 \cdot u(5M) + 0.9 \cdot u(0), \\ &\rightarrow C \succ D. \end{aligned}$$

1.5.4 Expected Utility Theory

Later on in this lesson we will learn about different kings of Nash Equilibrium. Here we'll just mention that Von NeumannMorgenstern Rationality Axioms $\Rightarrow$ Expected Utility Maximization $\Rightarrow$ Mixed Nash Equilibrium exists.

1.6 The Tragedy of the Commons

The following example shows a situation where

- All the players are equivalent, in the sense that the utility functions are computed in the same manner.
- The domain of user strategies is *continuous*.
- Like in the prisoner's dilemma, there is a stable solution which is not Pareto optimal.
- No dominant strategies exist.

We model a situation where increased use of a common resource reduces its utility. In such a case, if there is one leadership, it can force all the players to play according to the strategies that lead to an optimal global solution. However, when each player decides on a strategy, and wants to maximize his own utility, the obtained equilibrium may be far from optimal.

1.6.1 Problem Setting

There exists a shared resource (e.g., a pasture field or network bandwidth). There are N users, and each user i decides how much load x_i he wants to put on the resource.

In a case when the sum of user loads $\sum_{i=1}^{N} x_i$ is ≤ 1, the resource is overloaded and for every user i the utility $u_i = 0$. In any other case, the utility of a user i is computed by the formula $u_i = x_i \left(1 - \sum_{j=1}^{N} x_j\right)$. I.e., increasing the user load x_i will not always increase the utility u_i (because of the second factor, that decreases as the overall load increases). The optimal choice of x_i thus depends on the choices of the other uses.

1.6.2 The Rational Solution

If the choices of the other users are fixed (as unknown parameters) we can compute the maximum point of the utility function for a single user, as follows:

$u_i = x_i \left(1 - \sum_{j=1}^{N} x_j\right) = x_i \left(1 - x_i - \sum_{j \neq i} x_j\right)$ – *developing the utility formula.*

$\frac{du_i}{dx_i} = -2x_i + \left(1 - \sum_{j \neq i} x_j\right) = 0$ – *deriving according to x_i and setting the derivative to 0.*

$\Leftrightarrow 2x_i = 1 - \sum_{j \neq i} x_j$

$\Leftrightarrow x_i = 1 - x_i - \sum_{j \neq i} x_j = 1 - \sum_{j=1}^{N} x_j$ – *moving one x_i to the right-hand size.*

$\Rightarrow \sum_{j=1}^{N} x_j = N - N \sum_{j=1}^{N} x_j$ – *since all the players are equivalent.*

$\Leftrightarrow (N+1) \sum_{j=1}^{N} x_j = N$

$\Leftrightarrow \sum_{j=1}^{N} x_j = \frac{N}{N+1}$ – *we got an expression for the sum of loads.*

$\Rightarrow x_i = 1 - \sum_{j=1}^{N} x_j = 1 - \frac{N}{N+1} = \frac{N+1-N}{N+1} = \frac{1}{N+1}$ – *using the expression in the formula for x_i.*

$\Rightarrow u_i = x_i \left(1 - \sum_{j=1}^{N} x_j\right) = \frac{1}{N+1} \cdot \frac{1}{N+1} = \frac{1}{(N+1)^2}$ – *the resulting utility.*

Assuming that the players are rational, *each of them must* choose $x_i =$

$\frac{1}{N+1}$ and consequently get utility $u_i = \frac{1}{(N+1)^2}$. Is this state Pareto optimal?

If every user i chose $x_i = \frac{1}{2N}$ instead, the utility of each user would be

$$u_i = x_i \left(1 - \sum_{j=1}^{N} x_j\right) = \frac{1}{2N}\left(1 - N \cdot \frac{1}{2N}\right) = \frac{1}{2N}\left(1 - \frac{1}{2}\right) = \frac{1}{2N} \cdot \frac{1}{2} = \frac{1}{4N}$$

. This utility is asymptotically larger than $\frac{1}{(N+1)^2}$. Every player would increase his utility by moving from the rational solution $\forall i\ x_i = \frac{1}{(N+1)^2}$ to the new solution, $\forall i\ x_i = \frac{1}{2N}$. Thus, the rational solution is not Pareto optimal. Still, given the new solution, as our previous computation shows, every user will have the motivation to change his choice, and increase his utility at the expense of the others.

1.7 Nash Equilibrium

he obtained solution in the “tragedy of the commons” was not globally optimal due to the assumption that players are not collaborating with each other. Each player tries to maximize his own utility given the choices of the others. When no player can increase his utility alone (by changing his choice), we get a *Nash equilibrium* (named after John Nash, who first defined this situation in game theory). The formal definition:

Definition If x_i, x_{-i} are certain strategies of user i and the other users respectively, $\{x_i, x_{-i}\}$ is called a *Nash equilibrium* if

$$\forall 1 \le i \le N, \forall y_i \in A_i : u_i\left(x_i, x_{-i}\right) \ge u_i\left(y_i, x_{-i}\right)$$

According to this definition, we actually proved in the tragedy of the commons that $\forall 1 \le i \le N : x_i = \frac{1}{N+1}$ is a Nash Equilibrium.

Definition Given the strategies of other players x_{-i}, the *best response* of a player i is the set of all strategies that together with x_{-i} maximize the utility for player i. It is denoted by

$$BR_i(x_{-i}) = \arg\max_{x_i} u_i\left(x_i, x_{-i}\right)$$

In a Nash equilibrium, every player's strategy must be a best response, according to the definition above. We get an equivalent definition for Nash equilibrium.

Definition (An equivalent definition) For users $1, \ldots, N$, a set of choices $x_1 \ldots x_N$ is a *Nash equilibrium* if

$$\forall 1 \leq i \leq N : \; x_i \in BR_i(x_{-i})$$

1.7.1 Battle of the Sexes

We next describe some examples for games where Nash equilibria exist. In the first example, "the battle of the sexes", the situation is as follows: a couple tries to decide whether to go to the opera or to a sports game. One partner favors the sports game, and the other favors the opera, but in any case they prefer to stay together. The following table expresses their utility, where the row player is the one who favors the sports game. Recall that if a cell table contains a pari (x, y), x is the utility of the row player for the combination of strategies, and y is the utility of the column player.

	Sports	Opera
Sports	(4,3)	(2,2)
Opera	(1,1)	(3,4)

In this situation we have two Nash equilibria. For the strategy combination $\{\text{Sports}, \text{Sports}\} \in A$, each of the players will lose by changing their strategy: if the row player switches to the Opera strategy, the gain of this player will be reduced from 4 to 1; and if the column player switches to the opera, the utility of this player will be reduced to from 3 to 2. The strategy combination $\{\text{Opera}, \text{Opera}\} \in A$ is also a Nash equilibrium from similar reasons.

Unlike the tragedy of the commons, here there is no clear strategy that a rational player should choose. We do not know *how* to achieve one of the possible Nash equilibria; we only know that if we initialize the players' strategies according to one of the equilibria, and the players are selfish and not collaborating with each other, then none of them will have an incentive to switch to a different strategy.

1.7.2 Routing Game

The second example for a Nash equilibrium is similar, only that here the players have an incentive to choose different strategies rather than the same strategy. Consider two common routes that connect the users to the internet. Route A consists of only 1 link, on which the cost of a single packet is 1, and the cost of 2 packets together is 3 per packet. Route B consists of 2 links, where the price of a single packet per link is 1, and the price of 2 packets is 2 per packet per link. This is summarized in the following table (the row and column players are symmetric).

	A	B
A	(3,3)	(1,2)
B	(2,1)	(4,4)

Here, (1,2) and (2,1) are Nash equilibrium states, since if either the row player or the column player change their strategy alone, they will only increase their cost.

1.8 Mixed Strategies

1.8.1 Matching Pennies

Consider a situation, where two players have to choose heads or tails each. The row player wins a point if they make the same choice, and loses one if they choose differently. The column player, however, wins a point if they choose differently, and loses a point if they choose the same. We get the following utility table:

	Heads	Tails
Heads	(1,-1)	(-1,1)
Tails	(-1,1)	(1,-1)

In each of the four states, one player wins and one loses. The loser always has an incentive to change his strategy, and become a winner. Thus, in this game, no Nash equilibrium exists.

What can the players do in such a situation? Assume that the players are allowed not to choose a single strategy, but a distribution over the strategies. The row player will choose heads with probability p (and tails with probability $1-p$, and similarly the column player will choose heads with probability q. We can now compute the *expected* utility of, e.g., the row player, for each choice he makes: if he chooses heads, he will win one point with probability q (the column player also chooses heads; we assume that the players are independent). With probability $1-q$, the column player will choose tails and the row player will lose one point. Thus the expected utility of the row player for choosing heads is $u_{\text{heads}} = 1 \cdot q + (-1) \cdot (1-q) = 2q - 1$. Choosing tails is symmetric, thus the expected utility for the row player for choosing tails is $u_{\text{tails}} = (-1) \cdot q + 1 \cdot (1-q) = 1 - 2q$.

The row player will choose a mixed strategy only if one or more strategies have the maximum utility. In this case, this means that $u_{\text{heads}} = u_{\text{tails}} \Leftrightarrow 2q - 1 = 1 - 2q \Leftrightarrow q = \frac{1}{2}$.

1.8.2 Strategy Distributions

If A_i is the domain of possible strategies for the player i, then $\Delta(A_i)$ is the domain of all possible *probabilistic distributions* over A_i. It can be viewed as the simplex of strategies, where each strategy has a probability between 0 and 1, and the sum of probabilities is 1.

Definition For the player i, a *mixed strategy* is a choice of $p_i \in \Delta(A_i)$.

A single possible strategy (as opposed to a mixed strategy) is also called a *pure strategy.*

Definition For players $1, \ldots, N$, a *joint mixed strategy* is a vector of mixed strategies $\vec{P} = \{p_1, \ldots, p_N\}$. The *outcome* of the game is a joint mixed strategy.

Definition A *mixed Nash equilibrium* is a joint mixed strategy where for every player i, given the (mixed) strategies of the other players p_{-i}, i cannot increase his expected utility by changing his strategy. More formally,

$$\forall i \in 1 \ldots N, \forall q_i \in \Delta(A_i) : \; E_{x_i \sim p_i, x_{-i} \sim p_{-i}} \left[u_i(x_1, \ldots, x_N) \right] \geq E_{x_i \sim q_i, x_{-i} \sim p_{-i}} \left[u_i(x_1, \ldots, x_N) \right]$$

(p_{-i} is used to denote the joint distribution of all the players but i).

We will define two useful notations.

Definition The *support* of a mixed strategy p_i is the set of all strategies with non-zero probability in p_i, denoted by

$$\textit{support}(p_i) = \{a_i \mid p_i(a_i) > 0\}$$

Definition (In the mixed case) Given the mixed strategies of other players p_{-i}, the *best response* of a player i is the set of all *pure* strategies that together with p_{-i} give the maximal expected utility for player i. The best response set is denoted by

$$BR_i(p_{-i}) = \arg\max_{a_i} E_{x_{-i} \sim p_{-i}} \left[u_i\left(a_i, x_{-i}\right)\right]$$

For each mixed strategy, there exists a pure strategy with a greater or equal expected utility. To see that, recall that an expectation is always smaller than the maximal case. Take a mixed strategy $p_i \in \Delta\left(A_i\right)$, and the joint mixed strategy of the other players p_{-i}. We can choose a pure strategy $a_i \in \textit{support}(p_i)$, such that the expected utility of a_i, p_{-i} is maximal. I.e., this is the "maximal case", and the expected utility of p_i, p_{-i} is the "expectation". Thus, the expected utility of a_i, p_{-i} has to be greater.

In a Nash equilibrium, we choose a positive probability only for strategies a_i that maximize the utility with respect to the mixed strategies of the others. This actually means that

$$\forall i: \ \textit{support}(p_i) \subseteq BR_i(p_{-i}) \text{(an important property of mixed Nash equilibrium).}$$

For such a mixed strategy there exist only pure strategies with equal expected utility, and not strictly higher. We can use this property to formulate an equivalent definition of mixed Nash equilibrium.

Definition (Equivalent definition) A *mixed Nash equilibrium* is a joint mixed strategy where for every player i, given the (mixed) strategies of the other players p_{-i}, i cannot increase his expected utility by changing his strategy to a (different) pure strategy. Formally,

$$\forall i \in 1 \ldots N, \forall a_i \in A_i: \ E_{x_i \sim p_i, x_{-i} \sim p_{-i}}\left[u_i\left(x_1, \ldots, x_N\right)\right] \geq E_{x_{-i} \sim p_{-i}}\left[u_i\left(a_i, x_{-i}\right)\right]$$

1.8.3 Rock Paper Scissors

The following utility table describes the results of the rock-paper-scissors game.

	Rock	Paper	Scissors
Rock	(0,0)	(-1,1)	(1,-1)
Paper	(1,-1)	(0,0)	(-1,1)
Scissors	(-1,1)	(1,-1)	(0,0)

Similarly to the matching Pennies game, here there exists no pure Nash equilibrium. However, there exists a mixed one: $\frac{1}{3}$ probability for each strategy for each of the players. If we fix this strategy for the column player, than the expected utility for each choice of the row player is 0. This means that any of the choices is a best response and can be in the support of the mixed strategy of the row player. To achieve an equilibrium (i.e., to allow all strategies to be best responses also for the column player), we should give equal probability to each of the row player choices.

1.8.4 Nash Theorem

Theorem 1.1 (Nash, 1951) *Any game with a finite set of players and a finite set of strategies has a mixed Nash equilibrium.*

We will show an algorithm for finding the mixed Nash equilibrium, in a game with two players. The following is the matrix of utilities:

$$\begin{pmatrix} (r_{11}, c_{11}) & \cdots & (r_{1n}, c_{1n}) \\ \vdots & \ddots & \vdots \\ (r_{m1}, c_{m1}) & \cdots & (r_{mn}, c_{mn}) \end{pmatrix}$$

We will also define two sets of variables, $p(1), \ldots, p(m)$ for the probabilities assigned to each of the strategies of the row player, and similarly $q(1), \ldots, q(n)$, for the column player. We showed that in a mixed Nash equilibrium, $support(p) \subseteq BR_{\text{row}}(q)$. This means that the expected utility of every strategy of the row player should be the same, and this works in a symmetrical way for the column player. Assume that the support of the mixed strategies of the players are fixed. If $I = \{i_1, \ldots, i_k\}$ are the indices of the strategies in *support*(p), and $J = \{j_1, \ldots, j_l\}$ are the indices for *support*(q), we can formulate this as two linear programs:

$q(j_1) \cdot r_{i_1,j_1} + \cdots + q(j_l) \cdot r_{i_1,j_l} = q(j_1) \cdot r_{i_2,j_1} + \cdots + q(j_l) \cdot r_{i_2,j_l}$ – *first equation: the expected utility of the first row player strategy is equal to that of the second strategy*

$q(j_1) \cdot r_{i_2,j_1} + \cdots + q(j_l) \cdot r_{i_2,j_l} = q(j_1) \cdot r_{i_3,j_1} + \cdots + q(j_l) \cdot r_{i_3,j_l}$

$\vdots$

$q(j_1) \cdot r_{i_{k-1},j_1} + \cdots + q(j_l) \cdot r_{i_{k-1},j_l} = q(j_1) \cdot r_{i_k,j_1} + \cdots + q(j_l) \cdot r_{i_k,j_l}$

$q(j_1) + \cdots + q(j_l) = 1 \quad \forall 1 \leq s \leq n : q(s) \geq 0$ – *enforce that q is a distribution with the support J*

... and the second system:

$p(i_1) \cdot c_{i_1,j_1} + \cdots + p(i_k) \cdot c_{i_k,j_1} = p(i_1) \cdot c_{i_1,j_2} + \cdots + p(i_k) \cdot c_{i_k,j_2}$ – *first equation for the column player*

$\vdots$

$p(i_1) \cdot c_{i_1,j_{l-1}} + \cdots + p(i_k) \cdot c_{i_k,j_{l-1}} = p(i_1) \cdot c_{i_1,j_l} + \cdots + p(i_k) \cdot c_{i_k,j_l}$

$p(i_1) + \cdots + p(i_k) = 1$ – *last equations for the row player, enforcing that p is a distribution with the support I*

$\forall s \in I : p(s) > 0 \quad \forall 1 \leq s \leq m : p(s) \geq 0$

The solution to each system, if exists, is unique, since there are more equations than variables. If there exists a solution to both systems, that is according to the selected support (i.e., $p(s) > 0$ for $s \in I$, $p(s) = 0$ for $s \notin I$, $q(s) > 0$ for $s \in J$ and $q(s) = 0$ for $s \notin J$) – then this solution is a mixed Nash equilibrium. In the worst case, to find an equilibrium we will have to try this for every possible support for p and q, a total of $(2^m - 1)(2^n - 1)$ combinations (the -1 is there since the supports cannot be empty).

1.9 Location / Lemonade Stand Game

In this example, computing the Nash equilibria is not simple. There are many veriations to this general problem, where N ice cream/lemonade vendors are choosing a location in a defined space. The utility of each vendor is determined by the distance between him, the neighbor vendors and the defined space boundaries.

First variation The vendors are spread on the segment $[0,1]$. The utility of each vendor is half of the distance between him and his neighbors, i.e., if the location of the vendor i is x_i, the utility of the vendor i is $u_i = \frac{x_{i+1} - x_{i-1}}{2}$. When $N = 2$, there is a pure Nash equilibrium with the two vendors as close as possible to the middle point $\frac{1}{2}$. When $N = 3$, no pure Nash equilibrium exists.

Second variation If the vendors are spread on a circle instead of a segment, there always exists a Nash equilibrium, with the vendors spread at even distances from each other.

Links for other variations of the game and related competitions:

- http://martin.zinkevich.org/lemonade/
- http://tech.groups.yahoo.com/group/lemonadegame/