Stat 155 Game theory, Yuval Peres **Fall 2004**

Lectures 11,12,13

A betting game. Suppose that there are two players, a hider and a chooser. The hider has two coins. At the beginning of any given turn, he decides either to place one coin in his left hand, or two coins in his right. He does so, unseen by the chooser. The chooser then selects one of his hands, and wins the coins hidden there. That means he may get nothing if the hand is empty, or one or two coins. How should each of the agents play if he wants to maximise his gain, or minimize his loss? Calling the chooser player I and the hider player II, we record the outcomes in a *normal* or *strategic* form:

II / I	L	R
L	2	0
R	0	1

If the players choose non-random strategies, and each seeks to minimise his worst loss or to assure his gain, what are these amounts? In general, consider a pay-off matrix $(a_{i,j})_{m,n}$ (so that player I may play one of m possible plays, and player II one of n possibilities). The meaning of the entries is that a_{ij} is the amount that II pays I in the event that I plays i and II plays j. Let's calculate the assured payment for player I if pure strategies are used. If he announces to player II that he will play i, then II will play that j for which $\min_j a_{ij}$ is attained. If he were announcing his choice beforehand, player I would therefore play that i attaining $\max_i \min_j a_{ij}$. On the other hand, if player II has to annonce his intention for the coming round to player I, then a similar argument shows that he plays j, where j attains $\min_i \max_j a_{ij}$.

In the example, the assured value for II is 1, and the assured value for I is zero. In plain words, the hider can assure only losing one unit by placing one coin in his left hand, whereas the chooser knows that he will never lose anything by playing.

It's always true that the assured values satisfy the inequality: $\min_j a_{ij} \geq \max_i \min_j a_{ij}$. Intuitively, this is because player II cannot be assured of winning an amount greater than that amount more than which player I can be assured of not paying. Mathematically, let j^* denote the value of j that attains the minimum of $max_i a_{ij}$, and let $\hat{i}$ denote the value of i that attains the maximum of $\min_j a_{ij}$. Then

$$\min_j \max_i a_{ij} = \max_i a_{ij^*} \geq a_{\hat{i}j} \geq \min_j a_{\hat{i}j^*} = \max_i \min_j a_{ij}.$$

The fact that the assured values are not equal means that it makes sense to consider random strategies for the players. Suppose then that in the example I player L with probability x and R the rest of the time, whereas II plays L with probability t, and R with probability $1 - t$. Suppose that I announces to II his choice for x. How would II react? If he plays L, his expected payment is $2x$, if R, then $1 - x$. He minimizes the payout and achieves $\min\{2x, 1 - x\}$. Knowing that II will react in this way to hearing the value of x, I will seek to maximize his payoff by choosing x to maximize $\min\{2x, 1 - x\}$. He is choosing the value of x at which the two lines in this graph cross:

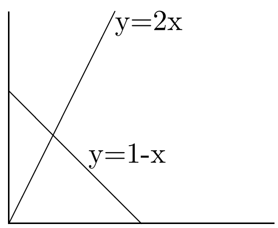

So his choice is $x = 1/3$. Looking at things the other way round, suppose that player II has to announce t first. The payoff for player I becomes $2t$ if he picks left and $1 - t$ if he picks right. Player should choose $t = 1/3$ to minimize his expected payment. This assures him of not paying more than 2/3 on the average.

Let's look at another example. Suppose we are dealing with a game that has the following payoff matrix:

II I	L	R
T	0	2
B	5	1

Suppose that player I plays T with probability x and B with probability $1 - x$, and that player II plays L with probability t and R with probability $1 - t$. If player II has declared the value of y, then Player I has expected payoff of $2(1 - y)$ if he plays T, and $4y + 1$ if he plays B. The maximum of these quantities is the expected payoff for player I under his optimal strategy, given that he knows y. Player II minimizes and so chooses $t = 1/6$ to obtain an expected payoff of 5/3.

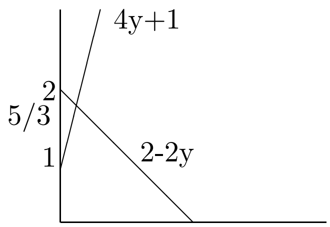

If player I has declared the value of x, then player II has expected payment of $5(1-x)$ if he plays L and $1+x$ if he plays R. He minimizes this, and then player II chooses x to maximize the resulting quantity. He therefore picks $x = 2/3$, with expected outcome of $5/3$.

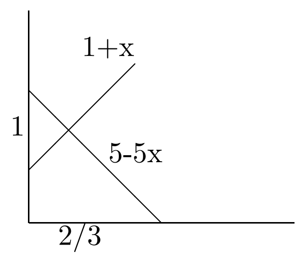

In general, player I can choose a probability vector

$$(x_1, \ldots, x_m), \ \sum_{i=1}^{m} x_i = 1,$$

where x_i is the probability that he plays i. Player II similarly chooses a strategy $(y_1, \ldots, y_n)^T$. The resulting expected payoff is given by $\sum x_i a_{ij} y_j = x^T Ay$. We will prove Von Neumann's minimax theorem, which states that $\min_y \max_x x^T Ay = \max_x \min_y x^T Ay$.

Example: Two players choose numbers in $\{1, 2, \ldots, n\}$. The player whose number is higher than that of her opponent by one wins a dollar, but if it exceeds the other number by two or more, she loses 2 dollars. In the event of a tie, no money changes hands. We write the payoff matrix for the game:

II	1	2	3	4	$\cdots$		n
I							
1	0	-1	2	2	$\cdots$		2
2	1	0	-1	2	$\cdots$		2
3	-2	1	0	-1	2	$\cdots$	2
.							
.							
n-1	-2	-2	$\cdots$		1	0	-1
n	-2	-2	$\cdots$			1	0

This apparently daunting example can be reduced by a new technique:

Domination: if row i has each of its elements at least the corresponding element in row $\hat{i}$: that is, is $a_{ij} \geq a_{\hat{i}j}$ for each j, then, for the purpose of determining the value of the game, we may erase row $\hat{i}$. (The value of the game is defined as the value arising from Von Neumann's minimax theorem). Similarly, there is a notion of domination for player II. If $a_{ij} \leq a_{ij^*}$ for each i, then we can eliminate column j^* without affecting the value of the game. More precisely, assuming that $a_{ij} \leq a_{ij^*}$, if player II changes a mixed strategy y to another z by letting $z_j = y_j + y_{j^*}$, $z_{j^*} = 0$ and $z_l = y_l$ for all $l \neq j, j^*$, then

$$\sum_{i,j} x_i a_{i,l} y_l = x^T A y \geq \sum_{i,j} x_i a_{i,l} z_l = x^T A z,$$

because $\sum_{i,j} x_i (a_{i,j} y_j + a_{i,j^*} y_{j^*}) \geq \sum_{i,j} x_i a_{i,j} (y_j + y_{j^*})$.

In the example in question, we may eliminate each row and column indexed by four or greater. We obtain the reduced game:

II / I	1	2	3
1	0	-1	2
2	1	0	-1
3	-2	1	0

Consider (x_1, x_2, x_3) as a strategy for player I. The expected payments made by player II under his pure strategies 1,2 and 3 are

$$(x_2 - 2x_3, -x_1 + x_3, 2x_1 - x_3). \tag{1}$$

Player II seeks to minimize his expected payment. Player I is choosing (x_1, x_2, x_3): for the time being, suppose that he fizes x_3, and optimises his choice for x_1. Eliminating x_2, (1) becomes

$$(1 - x_1 - 3x_3, -x_1 + x_3, 3x_1 + x_3 - 1).$$

Computing the choice of x_1 for which the maximum of the minimum of these quantities is attained, and then maximises this over x_3 yields an optimal strategy for each player of $(1/4, 1/2, 1/4)$ and a value for the game of zero.

Solution of chomp: The game of chomp is progressively bounded, so that each position is N or P. We will show that each rectangular position is in N. Otherwise, it is in P. Consider the move by Player I of chomping the upper-right hand corner. The resulting position is in N. This means that player II has a move to P. However, player I can play this move to start with, because each move after the upper-right square of chocolate is gone is available when it was still there. So player I can move to P, a contradiction. Note that it may not be that chomping the upper-right hand corner is a winning move. This argument, called *strategy stealing*, proves that player I has a winning strategy, without identifying it.

We can be precise about what is meant by a winning strategy: a strategy is a function assigning to each possible position, a legal move. It is winning if the player who has the first move will win by using it, whatever strategy his opponent uses.

Example: Wythoff nim A position consists of a pair of (n, m) of natural numbers, $n, m \geq 0$. A legal move is one of the following: to reduce n to some value between 0 and $n-1$ without changing m, to reduce m to some value between 0 and $m-1$ without changing n, or to reduce each of n and m by the same amount, so that the outcome is a pair of natural numbers. Consider the following recursive definition of a sequence of pairs of natural numbers: $(a_0, b_0) = (0,0), (a_1, b_1) = (1,2)$, and, for each $k \geq 1$,

$$a_k = mex\{a_1, \ldots, a_{k-1}, b_0, \ldots, b_{k-1}\}$$

and $b_k = a_k + k$. Each natural number greater than zero is equal to precisely one of the a_i or the b_i. To see this, note that a_j cannot be equal to any of $a_1, \ldots, a_{j-1}$ or $b_1, \ldots, b_{j-1}$. We have that $a_k > a_j$ because otherwise a_j would have taken the slot that a_k did, so that $b_k = a_k + k > a_j + j = b_j$.

Set $\alpha_k(\theta) = \lfloor k/\theta \rfloor$ and $\beta_k(\theta) = \lfloor k/(1-\theta) \rfloor$. Firstly, we will suppose that there exists $\theta \in (0,1)$ for which

$$\alpha_k(\theta) = a_k \text{ and } \beta_k(\theta) = b_k, \tag{2}$$

and find that there is only one number in $(0,1)$ for which this might be true. Since $b_k = a_k + k$, (2) implies that $\lfloor k/\theta \rfloor + k = \lfloor k/(1-\theta) \rfloor$. Dividing by k and noting that

$$0 \leq k/\theta - \lfloor k/\theta \rfloor < 1,$$

so that

$$0 \leq 1/\theta - (1/k)\lfloor k/\theta \rfloor < 1/k,$$

we find that

$$1/\theta + 1 = 1/(1-\theta). \tag{3}$$

Thus, $\theta^2 + \theta - 1 = 0$, so that θ or $1/\theta$ equal $2/(1+\sqrt{5})$. Thus, if there is a solution in $(0,1)$, it must be this value.

We now define $\theta = 2/(1+\sqrt{5})$. Note that (3) implies that

$$1/\theta + 1 = 1/(1-\theta),$$

so that

$$\lfloor k/(1-\theta) \rfloor = \lfloor k/\theta \rfloor + k.$$

This means that $\beta_k = \alpha_k + k$. We need to verify that

$$\alpha_k = mex\{\alpha_0, \ldots, \alpha_{k-1}, \beta_0, \ldots, \beta_{k-1}\}.$$

We checked earlier that α_k is not one of these values. Why is it equal to their mex? Define z to be this mex. If $z \neq \alpha_k$, then $Z < \alpha_k \leq \alpha_l \leq \beta_l$ for all $l \geq k$. Since z is defined as a mex, $z \neq \alpha_i, \beta_i$ for $i \in \{0, \ldots, k-1\}$.