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

Lecture 1

In this course on game theory, we will be studying a range of mathematical models of conflict and cooperation between two or more agents. The course will attempt an overview of a broad range of the models that are studied in game theory, and that have found application in, for example, economics and evolutionary biology. This first lecture gives an inkling of what is to come in the form of numerous examples.

One class of games that we begin by studying are combinatorial games. An example of a combinatorial game is that of *hex*, which is played on an hexagonal grid shaped as a square: think of a large square-shaped region that is tiled by a grid of small hexagons. Two players R and G alternately colour in hexagons of their choice either red or green, the red player aiming to produce a red crossing from left to right in the square and the green player aiming to form a green one from top to bottom. Finding the optimal strategy for either player remains an unsolved problem, except in a few cases, where the number of hexagons is small. An interesting variant of the game is that in which, instead of taking turns to play, a coin is tossed at each turn, so that each player plays the next turn with probability one half.

A second example which is simpler to analyse is the game of *nim*. There are two players, and several piles of sticks at the start of the game. The players take turns, and at each turn, most remove at least one stick from one pile. The player can remove any number of sticks that he pleases, but these must be drawn from a given pile. The aim of the game is to force the opponent to take the last stick remaining in the game. We will find the solution to *nim*: it is not one of the harder examples.

Thirdly, there are *congestion games*. Imagine two drivers, I and II who aim respectively to travel from cities B to D, or A to C:

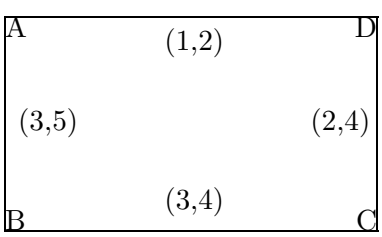

The costs incurred to the drivers depend on whether they travel the roads alone or along with the other driver. The vectors (a, b) attached to each road mean that the cost paid by wither driver for the use of the road is a if he travels the road alone, and b if he

shares its use with the other driver. For example, if I and II each use the road AB - which means that I chooses the route via A and II that via B - then each pays 5 units for doing so, whereas if only one of them uses that road, the cost is 3 units to that driver. We write a cost matrix to describe the game:

II / I	B	D
A	(6,8)	(4,3)
C	(5,6)	(7,5)

The vector notation $(\cdot, \cdot)$ denotes the costs to players I and II of their joint choice.

A fourth example is that of *penalty kicks*, in which there are two participants, the penalty-taker and the goalkeeper. The notion of left and right will be determined by the goalkeeper, not the penalty-taker. The penalty-taker chooses to hit the ball either to the left or the right, and the goalkeeper dives in one of these directions. We display the probabilities that the penalty is scored in the following table:

GK / PT	L	R
L	0.8	1
R	1	0.5

That is, if the goalie makes the wrong choice, he has no chance. The penalty-taker has a strong 'left' foot, and has a better chance if he plays left. The goalkeeper aims to minimise the probability of the penalty being scored, and the penalty-taker aims to maximise it. We could write a payoff matrix for the game, but, since it is zero-sum, with the interests of the players being diametrically opposed, one value sums up the outcome of each joint strategy for the players, so that doing so is redundant. We will determine the optimal strategy for the players for a class of games that include this one, and, often, it will turn out to be a random one.

Two person zero-sum games have been applied in a lot of contexts: in sprots, like this example, in military contexts, in a few economic applications, in evolutionary biology. These games have a very complete theory, so that it has been tempting to try to apply them. Real life is often more complicated, however, with the possibility of cooperation between players to realize a mutual advantage. The theory of games that model such an effect is much less complete.

The mathematics associated to *zero-sum games* is that of convex geometry. A convex set is one where, for any two points in the set, the straight line segment whose endpoints are these two itself lies in the set.

The relevant geometric fact for this aspect of game theory is that, given any closed convex set in the plane and a point lying outside of it, we can find a find that separates the set from the point. There is an analogous statement in higher dimensions. Von Neumann exploited this fact to solve zero sum games using a *minimax* variational principle. We will prove this result.

A related concept is that of Nash equilibrium: is there a 'rational' choice for the two players, and if so, what could it be? The meaning of 'rational' here and in many contexts is a valid subject for discussion. There are anyway often many Nash equilibria and further criteria are required to pick out relevant ones.

A development of the last twenty that we will discuss is the application of game theory to evolutionary biology. In economic applications, it is often assumed that the agents are acting 'rationally', and a neat theorem should not distract us from remembering that this can be a hazardous assumption. In some biological applications, we can however see Nash equilibria arising as stable points of evolutionary systems composed of agents who are 'just doing their own thing', without needing to be 'rational'.

Another interesting topic is that of signalling. If one player has some information that another does not, that may be to his advantage. But if he plays differently, might he give away what he knows, thereby removing this advantage?

A quick mention of other topics: that of voting. Arrow's impossibility theorem states roughly that if there is an election with more than two candidates, then no matter which system one chooses to use for voting, there is trouble ahead: at least one desirable property that we might wish for the election will be violated. A recent topic is that of eliciting truth. In an ordinary auction, there is a temptation to underbid. For example, if a bidder values an item at 100 dollars, then he has no motive to bid any more or even that much, because by exchanging 100 dollars for the object at stake, he has gained an item only of the same value to him as his money. The second-price auction is an attempt to overcome this flaw: in this scheme, the lot goes to the highest bidder, but at the price offered by the second-highest bidder. This problem and its solutions are very relevant to bandwidth auctions made by governments to cellular phone companies.

A final question for this first lecture: suppose that we seek to reward a metereologist who is asked for the probability of rain tomorrow. She offers a probability between 0 and 1, and we want to devise a strategy for paying here so that she is motivated to get at least roughly the right answer. We will later see a good example of a scheme for doing this.