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

Lecture 2

We continue the overview of topics. Another example of a mathematical tool that we will encounter is

Theorem 1 (Brouwer's fixed point theorem) *: If $K \subseteq \mathbb{R}^d$ is closed, bounded and convex, and $T : K \to K$ is continuous, then T has a fixed point. That is, there exists $x \in K$ for which $T(x) = x$.*

The assumption of convexity can be weakened, but not discarded entirely. To see this, consider the example of the annulus $C = \{x \in \mathbb{R}^2 : 1 \leq |x| \leq 2\}$, and the mapping $T : C \to C$ that sends each point to its rotation by 90 degrees anticlockise about the origin. Then T is *isometric*, that is, $|T(x) - T(y)| = |x - y|$ for each pair of points $x, y \in C$. Certainly then, T is continuous, but it has no fixed point.

Although from its statement, it is not evident what the connection of Brouwer's theorem to game theory might be, we will see that the theorem is of central importance in proving the existence of Nash equilibria.

Example: Pie cutting. As another example, consider the problem of a pie, different parts of whose interior are composed of different ingredients. The game has two or more players, who each have their own preferences regarding which parts of the pie they would most like to have. If there are just two players, there is a well-known method for dividing the pie: one splits it into two halves, and the other chooses which he would like. Each obtains at least one-half of the pie, as measured according to each own preferences. But what if there are three or more players? We will study this question, and that where we ask that the pie be cut in such a way that each player judges that he gets at least as much as anyone else, according to his own criterion.

Example: Secret sharing. Suppose that we plan to give a secret to two people. We do not trust either of them entirely, but want the secret to be known to each of them provided that they co-operate. If we look for a physical solution to this problem, we might just put the secret in a room, put two locks on the door, and give each of the players the key to one of the doors. In a computing context, we might take a password and split it in two, giving the each half to one of the players. However, this would force the length of the password to be high, if one or other half is not to be guessed by repeated tries. A more ambitious goal is to split the secret in two in such a way that neither person has any useful information on his own. And here is how to do it: suppose that the secret s is an integer that lies between 0 and some large value M, for example, $M = 10^6$. We who hold the secret at the start produce a random integer x, whose distribution is uniform on the interval $\{0, \ldots, M-1\}$ (*uniform* means that each of the M possible outcomes is equally likely, having probability $1/M$). We tell the number x to the first person, and the number $y = (s-x) mod M$ to the second person (mod M means adding the right multiple of M so that the value lies on the interval $\{0, \ldots, M-1\}$). The first person has no useful information. What about the second? Note that

$$\mathbb{P}(y = j) = \mathbb{P}((s-x)\, mod\, M = j) = 1/M,$$

the last equality because $(s-x) mod M$ equals y if and only if the uniform random variable x happens to hit one particular value on $\{0, \ldots, M_1\}$. So the second person himself only has a uniform random variable, and, thus, no useful information. Together, however, the players can add the values they have been given, reduce the answer mod M, and get the secret s back. A variant of this scheme can work with any number of players. We might have ten of them, in such a way that any nine of them have no useful information if they pool their resources, and with the ten together being able to unlock the secret.

Example: Cooperative games. These games deal with the formation of coalitions, and their mathematical solution involves the notion of *Shapley value*. As an example, suppose that three people, I,II and III, sit in a store, the first two bearing a left-handed glove, while the third has a right-handed one. A wealthy tourist, ignorant of the bitter local climactic conditions, enters the store in dire need of a pair of gloves. She refuses to deal with the glove-bearers individually, so that it becomes their job to form coalitions to make a sale of a left and right-handed glove to her. The third player has an advantage, because his commodity is in scarcer supply. This means that he should be able to obtain a higher fraction of the payment that the tourist makes than either of the other players. However, if he holds out for too high a fraction of the earnings, the other players may agree between them to refuse to deal with him at all, blocking any sale, and thereby risking his earnings. We will prove results in terms of the concept of the Shapley value that provide a solution to this type of problem.

Example: Keeping the metereologist honest. The employer of a weatherman is determined that he should provide a good prediction of the weather for the following day. The weatherman's instruments are good, and he can, with sufficient effort, tune them to obtain the correct value for the probability of rain on the next day. There are many days, and, on the i-th of them, this correct probability is called p_i. On the evening of the $i-1$-th day, the weatherman submits his estimate $\hat{p}_i$ for the probability of rain on the following day, the i-th one. Which scheme should we adopt to reward or penalize the weatherman for his predictions, so that he is motivated to correctly determine p_i (that is, to declare $\hat{p}_i = p_i$)? The employer does not know what p_i is, because he has no access to technical equipment, but he does know the $\hat{p}_i$ values that the waetherman provides, and he knows whether or not it is raining on each day.

One suggestion is to pay the weatherman, on the i-th day, the amount $\hat{p}_i$ (or some dollar multiple of that amount) if it rains, and $1-\hat{p}_i$ if it shines. If $\hat{p}_i = p_i = 0.6$, then the payoff is

$$\begin{aligned} & \hat{p}_i \mathbb{P}(\text{it rains}) + (1-\hat{p}_i)\mathbb{P}(\text{it shines}) = \hat{p}_i p_i + (1-\hat{p}_i)(1-p_i) \\ = \; & 0.6 \times 0.6 + 0.4 \times 0.4 = 0.52. \end{aligned}$$

But in this case, even if the weatherman does correctly compute that $p_i = 0.6$, he is tempted to report the $\hat{p}_i$ value of 1, because, by the same formula, in this case, his earnings are 0.6.

Another idea is to wait for a long time, one year, say, and reward the weatherman according to how accurate his predictions have been on the average. More concretely, suppose for the sake of simplicity that both the weatherman is only able to report $\hat{p}_i$ values on a scale of 0.1: so that he has eleven choices, namely $\{k/10 : k \in \{0, \ldots, 10\}\}$. When a year has gone by, the days of that year may be divided into eleven types, according to the $\hat{p}_i$-value that the weatherman declared. We would compare the fraction of days it rained in the class of days on which $\hat{p}_i$ was declared to the value $\hat{p}_i$, and reward the weatherman if this value is low, for each of the eleven values of i.

A scheme like this is quite reasonable and would certainly ensure that the weatherman tuned his instruments at the beginning of the year. However, it is not perfect. For example, ordinary random fluctuation mean that the weatherman will be probably be slightly wrong as the end of the year approaches. It might be that it rained on 95 percent of the days on which the weatherman declared $\hat{p} = 0.9$, while, on those on which he declared at 0.6, the average in reality has been 55 percent. Suppose that on the evening of one of the later days he sees that his instruments accurately predict 0.9. He knows that it very likely to rain on the next day. He is tempted to declare the next day at 0.6 - that is, to set $\hat{p}_i = 0.6$ for the i in question, because doing so will boost his 0.6 category and allow his 0.9 category to catch up with the downpour.

Infact, it is possible to design a scheme whereby we decide day-by-day how to reward the weatherman only on the basis of his declaration from the previous evening, without encountering the kind of problem that the last scheme had. Suppose that we pay $f(\hat{p}_i)$ to the weatherman if it rains, and $f(1-\hat{p}_i)$ if it shines on day i. If $p_i = p$ and $\hat{p}_i = x$, then the expected payment made on day i is equal to

$$pf(x) + (1-p)f(1-x) = g_p(x).$$

(We are defining $g_p(x)$ to be this left-hand-side, because we are interested in how the payout is expected to depend on the prediction x of the weatherman on a given day where the probability of rain is p). Our aim is to reward the weatherman if his $\hat{p}_i$ equals p_i, in other words, to ensure that the expected payout is maximised when $x = p$. This means that the function $g_p : [0, 1] \to \mathbb{R}$ should satisfy $g_p(p) > g_p(x)$ for all $x \in [0, 1] - \{p\}$.

Homework. Using this approach, determine a good choice of f. Or find another good scheme for rewarding the weatherman.