EVOLUTION AND THE THEORY OF GAMES

STEFAN GERITZ, HELSINKI, 2011

Transcript of my notes of 01-11-2011

1. The course "Evolution and the theory of games" is part of the biomathematics curriculum at the Department of Mathematics and Statistics of the University of Helsinki. The focus is on animal behavior: can we understand animal behavior from the theory of games?

Parts of this course are taken from the books *"Game theory evolving: a problem-centered introduction to modeling strategic interaction'* by Herbert Gintis (2009, 2nd edition, Princeton University Press) and *"Evolution and the theory of games'* by John Maynard Smith (1982, Cambridge University Press).

A game is a model of a situation of conflicts of interests where the pay-off to one player depends not only on his own strategy but also on the strategies of the other players. The phrase "*but also on the strategies of the other players*" makes that a game is not a mere optimization problem.

For example, the Traveling Salesman Problem (of finding the shortest route through a finite number of towns) is an optimization problem. However, if there are several salesmen and a bonus is given for arriving in a town first (presumably because who comes first strikes the best deal), then the problem is no longer a problem of mere optimization because the best route is not only short but also takes into account the choices of the others.

What games do animals play? What counts as a *solution* of a game, and how do we find it? These are questions we address first by means of examples.

2. The *prisoner's dilemma* (PD): Two men have been sentenced to serve two years in prison for the possession of illegal recreational substances. It is suspected, however, that they are not just users but also dealers. If proven true, this will give them two extra years in prison. Each prisoner is offered the following deal: if one gives testimony against the other, he gets one year reduction (for helping the police) and the other gets two extra years (because now there is proof that he is a dealer). The prisoner's are not allowed to talk to one another and so have to make their choices independently.

What would the prisoners chose: testify against the other or not? The payoffs depend on their combined choices as given in the following matrix:

	not testify	testify
not testify	-2, -2	-4, -1
testify	-1, -4	-3, -3

The payoffs are given as negative the number of years to serve in prison. The first of each pair is the payoff to the row-player (on the left) and the second of each pair is the payoff to the column-player (at the top).

It can be seen that if the column-player chooses not to testify, the row-player better does testify as this gives him only one year in prison rather than the two years he serves if he does not testify. If instead the column-player choses to testify, then the row-player also better testifies, because in that way he serves only three years instead of the four he serves if he does not testify.

So, whatever the choice of the column-player, the row-player better testifies. By symmetry the same holds for the column player, and both players end up testifying against the other with the result that each serves three years in prison. The strategy 'testify' is said to dominate the strategy 'not testify', and the outcome (both players testify) is called a dominant strategy solution. And this is the paradox: by attempting to minimize the time in prison, both players end up with a sub-optimal outcome: if both refuse to testify against the other they get only two years in prison.

The prisoner's dilemma is used as a basic model for the evolution of cooperation among humans or animals. Note that the prisoner's dilemma does not explain cooperation: it predicts non-cooperation (i.e., between the players, not with the police), which raises the question why there exists cooperation at all. Later we shall see that if the prisoner's dilemma is played repeatedly between the same players, under some circumstances cooperation will be the outcome.

3. The canonical form of the prisoner's dilemma is given by the payoff matrix

	C	D
C	R, R	S, T
D	T, S	P, P

where C stands for 'cooperation' (with the other player) and D for 'defection' (i.e., non-cooperation). The payoffs are denoted by T (*temptation*), R (*reward*), P (*punishment*) and S (*sucker's payoff*) and are ordered as $S < P < R < T$. One readily checks that D is a dominant strategy, i.e. gives a higher payoff than the alternative irrespective of the choice of the opponent.

Do animals play the prisoner's dilemma? Animal grooming (i.e., cleaning the fur or skin of one individual by removing parasites and dirt by another individual) may be an example of the PD in nature. Let c denote the cost of grooming someone else, and let b denote the benefit of being groomed yourself, and assume that the benefit exceeds the cost, i.e., $b > c > 0$. The payoff matrix then is

	C	D
C	$b - c,\ b - c$	$-c,\ b$
D	$b,\ -c$	$0,\ 0$

where C now stands for *grooming* and D for *not grooming*. It can be seen immediately that the payoff matrix has the ordering of a prisoner's dilemma game. Other examples have been given in the lecture.

4. The *hawk-dove game* (HD) describes a contest between two individuals for a resource like a food item or a territory and is used as a basic model for the evolution of aggression. There are two strategies: *hawk* (H) and *dove* (D). Someone playing H will fight for the resource until someone gives up or gets hurt. Someone playing D will not fight and either give up or share the resource evenly if the other player also has D. Here is the payoff matrix:

	H	D
H	$\frac{1}{2}R - \frac{1}{2}C,\ \frac{1}{2}R - \frac{1}{2}C$	$R,\ 0$
D	$0,\ R$	$\frac{1}{2}R,\ \frac{1}{2}R$

where R is the value of the resource and C is the cost of injury. The factor $\frac{1}{2}$ appears because each player in a H$\times$H-contest wins or looses with a probability one-halve, and likewise, in a D$\times$D-contest, each player gets the resource with a probability one-halve.

If $R > C$ (i.e., the value of the resource outweighs the potential cost of injury), then H is a dominant strategy which is preferred above D no matter how the opponent choses to play. The outcome of the game therefore is that both players chose H. This is the dominant strategy solution of the game. And here we see the same paradox as with the prisoner's dilemma: if both players chose H, their payoff is less than if both chose D.

If $R < C$, there is no dominant strategy. To find a 'solution' of the game (in some sense or another) we have to develop another solution concept than the dominant strategy solution. We do that later; first we develop a more formal language to describe games in general.

5. A *two-person game* is fully characterized by the following three things:

(1) The strategy sets X and Y, one for each player but not necessarily the same.

(2) The payoff function $\pi : X \times Y \to \mathbb{R}^2$ where $\pi(x, y)$ is the pair of payoffs to the first and second players if they choose the strategies $x \in X$ and $y \in Y$.

(3) The solution concept, like the 'dominant strategy solution' as in the previous examples, but there are other solution concepts too.

(An N-person game has N players with strategy sets $X_1, \ldots, X_N$ and payoff function $\pi : X_1 \times \cdots \times X_N \to \mathbb{R}^N$.)

In the example of the prisoner's dilemma we have $X = Y = \{\mathrm{C}, \mathrm{D}\}$. The payoff function is given by the entries of the payoff matrix, and so

$$\begin{cases} \pi(\mathrm{C},\mathrm{C}) & = & (R, R) \\ \pi(\mathrm{C},\mathrm{D}) & = & (S, T) \\ \pi(\mathrm{D},\mathrm{C}) & = & (T, S) \\ \pi(\mathrm{D},\mathrm{D}) & = & (P, P) \end{cases}$$

The solution concept we used was the 'dominant strategy solution', which gives us the strategy pair (D,D).

6. Before we go into alternative solution concepts, have a look at the following example taken from *"Game theory evolving"* by Herbert Gintis (but with bananas instead of coconuts):

Big Joe and Little Joe are two monkeys sitting under a banana tree. The idea is that at least one of them has to climbs up and knocks the bananas down so that they can be eaten. The dilemma is this: if BJ climbs up while LJ waits below, then LJ can already start eating before BJ is back on the ground, and *vice versa*.

The costs and benefits are the following: The total energetic value of the bananas is 10 kcal. The cost of climbing is 2 kcal for BJ and 0 kcal for LJ, who is the better climber. If both climb the tree, then after they're back on the ground, BJ gets a share worth of 7 kcal while LJ gets only 3 kcal. If BJ climbs the tree while LJ waits down below, this will be 6 kcal for BJ versus 4 kcal for LJ. If BJ waits and LJ climbs, this becomes 9 kcal versus 1 kcal. If both wait, they don't get any bananas, and so both get 0 kcal.

There are various possible scenarios, depending on which player makes the first move, leading to different games with different strategy sets and different payoff functions. Here we consider the case where BJ moves first; the other possibilities are left as an exercise.

We may represent the game as a so-called game tree (see figure): at the top is the 'root' indicated by BJ, who decides either to wait (w) or to climb (c) as shown by the first two branches. Once BJ has made his move, LJ will have to make his move as shown by the second-level branches. At the bottom of the figure are the payoffs to BJ (first of each pair) and to LJ (second of each pair) depending on their combined moves (and taken into account that BJ is a lousy climber).

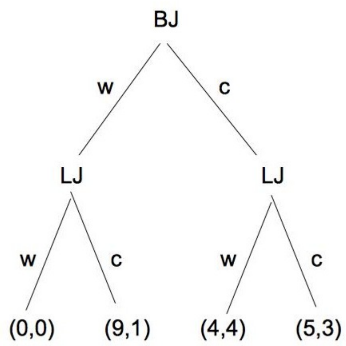

From the tree it can be seen that if BJ chooses to wait, then LJ better climbs, and if BJ chooses to climb then LJ better waits. Assuming that LJ indeed chooses what is best for him, BJ should choose to wait in order to maximize his payoff.

How to model this game in the way presented in the previous section? If BJ moves first, his strategy set is simply $X_{\mathrm{BJ}} = \{\mathrm{w,c}\}$: wait or climb. However, LJ can use the information about BJ's first move, i.e., he can use conditional strategies: 'if BJ waits I do this and if BJ climbs I do that'. His strategy set therefore is $X_{\mathrm{LJ}} = \{\mathrm{ww,wc,cw,cc}\}$, where the first letter of each pair tells what LJ does if BJ waits while the second letter tells what LJ does if BJ climbs. For example, 'cw' means 'if BJ waits, I climb, and if BJ climbs, I wait'. The payoff function is given by the payoff matrix

	cc	cw	wc	ww
w	9,1	9,1	0,0	0,0
c	5,3	4,4	5,3	4,4

where BJ is the row-player and LJ the column player. It can be seen that the pay-off to strategy cw is always better (or at least not worse) than the other strategies of LJ irrespective BJ's: in other words cw is a dominant strategy. Eliminating the other strategies of LJ, we obtain a reduced game with the payoff matrix

	cw
w	9,1
c	4,4

In this reduced game, the strategy w is a dominant strategy for BJ, and so we end up with the dominant strategy solution (w,cw), which is also what we found by inspection of the game tree.

Some remarks: (1) The payoffs for (w,cw) are the same as the payoffs for (w,cc), but cw and cc are essentially different strategies, which becomes immediately clear if you compare the payoffs for (c,cw) and (c,cc). (2) The strategy cw is strongly dominating wc but only weakly dominating cc and ww, because some of the payoffs are the same. The solution (w,cw) is therefore also called a weakly dominant strategy solution. (3) Note that w is not a dominant strategy in the first payoff matrix, i.e., before we reduced the size of the game by eliminating the (weakly or strongly) dominated strategies for LJ.