EVOLUTION AND THE THEORY OF GAMES

STEFAN GERITZ, HELSINKI, 2011

Transcript of my notes of 03-11-2011

7. Previous lecture we introduced the notion of a 'game' as a model of a situation of conflict where the payoff to one player depends not only on his own strategy but on the strategies of the other players as well. In particular, a 2-person game was defined by the following three things:

(1) Two strategy sets X and Y.

(2) A payoff function $\pi : X \times Y \to \mathbb{R}^2$.

(3) A solution concept.

Modeling a situation of conflict of interests thus amounts to choosing (1) the strategy sets, (2) a payoff function and (3) a solution concept in a way that best suits the situation. It should be emphasized that these are modeling choices; they do not follow in any way from the mathematics.

So far we have seen one solution concept: the dominant strategy solution. A strategy $x \in X$ is said to be dominated by $x' \in X$ if $\pi_1(x, y) \leq \pi_1(x', y)$ for all $y \in Y$ with a strict inequality for at least one value of y. If the inequality is strict for all y we speak of 'strong' domination, otherwise of 'weak' domination. We have seen how removal of dominated strategies can lead to a single strategy pair, which we then call a dominated strategy solution.

The example of Big Joe and Little Joe under the banana tree showed that once we have eliminated all dominated strategies for each player, it may turn out that a strategy that was not dominated at the onset now is dominated among the strategies that remain. And so we may be able to remove dominated strategies in several rounds. This is called the method of iterated removal of dominated strategies.

It should be kept in mind, however, that the game is *not played* in rounds; the elimination of dominated strategies in one or more sequential steps happens only in the mind of the players before they choose their strategy and therefore before the game is played. The dominant strategy solution thus presumes a rather sophisticated level of rationality for all players. The following game requires the players to think four rounds ahead:

	y_1	y_2	y_3	y_4
x_1	4,5	5,3	5,6	4,4
x_2	5,3	2,1	3,5	5,2
x_3	2,6	6,3	4,2	5,5

The actual solving of this game is left as an exercise.

8. Sometimes there is no dominant strategy solution. Take the following example:

	y_1	y_2	y_3
x_1	1,4	2,1	$4,1\frac{1}{2}$
x_2	2,1	4,4	3,2

There is no dominated strategy, and so the elimination method stalls already in the first round. To enforce a solution we could change the solution concept, or we could extend the strategy space. Let's do the latter by adding the mixed strategy $y^* = \frac{1}{2}y_1 + \frac{1}{2}y_2$ (i.e., play y_1 or y_1, each with a probability of one-half) to the strategy set of the column-player:

	y_1	y_2	y_3	y^*
x_1	1,4	2,1	$4,1\frac{1}{2}$	$1\frac{1}{2}, 2\frac{1}{2}$
x_2	2,1	4,4	3,2	$3, 2\frac{1}{2}$

Now y_3 is dominated by y^* and therefore can be eliminated, which gives

	y_1	y_2	y^*
x_1	1,4	2,1	$1\frac{1}{2}, 2\frac{1}{2}$
x_2	2,1	4,4	$3, 2\frac{1}{2}$

Now x_1 is dominated by x_2 and can be removed:

	y_1	y_2	y^*
x_2	2,1	4,4	$3, 2\frac{1}{2}$

Now both y_1 and y^* are dominated by y_2 and can be eliminated, and so we end up with the dominant strategy solution (x_2,y_2).

Note that y^* does not occur in the solution and yet was needed to complete the elimination procedure. Actually, this is quite absurd, and maybe we should have included mixed strategies from the beginning. Taking mixed strategies aboard, however, presumes something about the rationality of the players, namely, that they are actually able to generate random numbers, an assumption that may or may not be warranted depending on the context.

9. A *mixed strategy* is a probability distribution over a given set of pure strategies. Given the pure strategies $x_1, \ldots, x_k$, we can represent a mixed strategy by a vector of probabilities $p_1, \ldots, p_k$ where $p_i \geq 0\ \forall i$ and $\sum p_i = 1$ and where p_i is the probability of playing strategy x_i. The set $\{x_i : \ p_i > 0,\ i = 1, \ldots, k\}$ is called the support of the mixed strategy. The notion of mixed strategy as a probability distribution is readily generalized to countably many pure strategies or even a continuum of pure strategies.

If we allow for mixed strategies, the existence of a dominant strategy solution is still not guaranteed. Take the Hawk-Dove game:

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

If $R > C$, then (H,H) is the dominant strategy solution, but with $R < C$ there is no dominant strategy solution, neither pure (which is obvious) nor mixed. To see that there is no mixed solution, let $x = (p, 1-p)$ and $y = (q, 1-q)$ be mixed strategies where p and q are the probabilities of playing Hawk. Then

$$\begin{aligned} \pi_1(x,y) &= \tfrac{1}{2}(R-C)pq + Rp(1-q) + 0(1-p)q + \tfrac{1}{2}R(1-p)(1-q) \\ &= \tfrac{1}{2}R(1-q) + \tfrac{1}{2}(R-qC)p \end{aligned}$$

So, π_1 is an increasing function of p if $q < R/C$ but a decreasing function of p if $q > R/C$. Hence, there do not exist any two values p and p' such that one gives a higher payoff *for all* q, because the ordering of the respective payoffs is reversed if we change from $q < R/C$ to $q > R/C$.

The dominant strategy solution is a very credible solution concept, but it need not exist. Moreover, the presumed level of rationality of the players may limit its usefulness in the context of animal behavior. There exist many other solution, but each has its pros and contras. Some of these alternatives we have a look at now.

10. The *Hicks optimum* maximizes total payoff as a function of the strategies of all player. For example, the Prisoner's Dilemma with $T > R > P > S$ and $2R > T + S$ has a unique Hicks optimum, namely (C,C).

In the context of the evolution of animal behavior, accepting the Hicks optimum would amount to accepting solutions *for the good of the species* or, on a smaller scale, for the good of the population. Such a solution, however, is not stable against cheaters: in the Prisoner's Dilemma, for example, a defector would enjoy a higher payoff than the cooperators, and so the temptation for a player to defect destroys the coherence of the group and makes the Hicks optimum (C,C) an unlikely outcome.

The *Pareto optimum* is defined such that any change of strategies (for one or more players) is disadvantageous for at least one of the players. For example, the Prisoner's Dilemma with $T > R > P > S$ has a three Pareto optima: (C,C), (D,C) and (C,D), but again none of these is free of the temptation for one or both players to change their strategy from cooperation to defection. With the Pareto optimum the emphasis shifts to pre-game negotiation and contract, which in the context of animal behavior may not be very realistic.

The *minimax solution* is such that each player chooses a strategy that minimizes his maximum loss, which is equivalent to maximizing one's minimum payoff. In the Hawk-Dove game, the minimum payoff to Hawk is $\frac{1}{2}(R-C)$, and the minimum payoff to Dove is 0:

	H	D	$\min\{\pi_1\}$
H	$\frac{1}{2}(R-C)$, $\frac{1}{2}(R-C)$	R, 0	$\frac{1}{2}(R-C)$
D	0, R	$\frac{1}{2}R$, $\frac{1}{2}R$	0
$\min\{\pi_2\}$	$\frac{1}{2}(R-C)$	0	

If $R > C$, then for each player the minimum payoff is maximized by choosing Hawk, and so the minimax solution is (H,H). If $R < C$, then for each player the minimum payoff is maximized by choosing Dove, and so the minimax solution is (D,D). In the latter case, however, both player are tempted to change unilaterally from D to H, because Hawk against Dove gives a higher payoff. In a zero-sum game (i.e., a game where one player's loss is the other player's gain) such temptation does not occur, but the Hawk-Dove game is not a zero-sum game, and moreover, zero-sum games are rather special.

11. A *Nash equilibrium* in a two-person game with strategy sets X and Y is a strategy point $(\hat{x}, \hat{y}) \in X \times Y$ such that

$$\begin{cases} \pi_1(x, \hat{y}) \leq \pi_1(\hat{x}, \hat{y}) & \forall x \in X \\ \pi_2(\hat{x}, y) \leq \pi_2(\hat{x}, \hat{y}) & \forall y \in Y \end{cases}$$

In other words: $\hat{x}$ and $\hat{y}$ are best responses to one another, and hence unilateral change of strategy does not increase a player's payoff. So, the absence of temptation to change one's strategy is build into the definition of the Nash equilibrium.

The definition of a Nash equilibrium readily generalizes to an N-person game.

For example, take the Hawk-Dove game: if $R > C$, then (H,H) is a Nash equilibrium, and if $R < C$, then (D,C) and (C,D) are Nash equilibria. As a next example, consider

	y_1	y_2
x_1	0,2	2,1
x_2	3,1	1,3

which does not have a Nash equilibrium of pure strategies. The question is whether there is an equilibrium with mixed strategies.

Theorem. *Every game with finitely many pure strategies for each player has at least one Nash equilibrium if mixed strategies are allowed.*

The proof is given in a later lecture.

The theorem implies that the last example must have a Nash equilibrium which necessarily involves mixed strategies. To find this solution, let $x = (p, 1-p)$ and $y = (q, 1-q)$ be mixed strategies. The payoff to the row-player is $\pi_1(x, y) = 1 + 2q + (1-4q)p$, and so the best response to $q < \frac{1}{4}$ is $p = 1$, and the best response to $q > \frac{1}{4}$ is $p = 0$, while for $q = \frac{1}{4}$ any $p \in [0, 1]$ is a best response, as indicated by the thick line in the next figure:

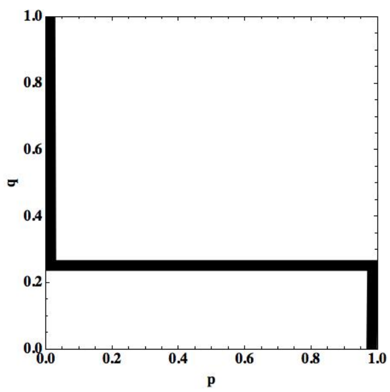

The payoff to the column-player is $\pi_1(x, y) = 3 - 2p + (3p - 2)q$, and so the best response to $p < \frac{2}{3}$ is $q = 0$, and the best response to $p > \frac{2}{3}$ is $q = 1$, and for $p = \frac{2}{3}$ any $q \in [0, 1]$ is a best response:

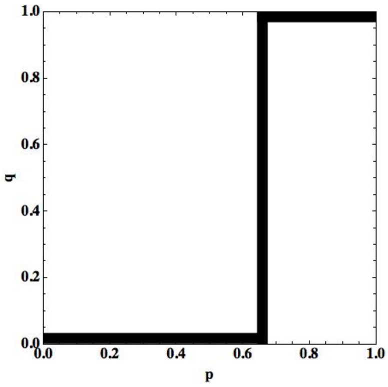

Since the first figure gives the set of best responses to q and the second figure the best responses to p, their superposition gives at the intersection of the thick lines the pair of strategies $(\hat{p}, \hat{q}) = (\frac{2}{3}, \frac{1}{4})$ that are best responses to one another and hence is a Nash equilibrium:

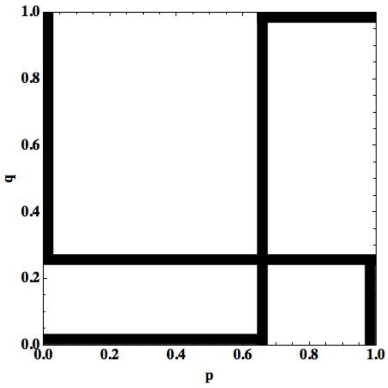

For obvious reasons this graphical method is called the *Swastika method*. Unfortunately, this method is only practical for two-person games with two pure strategies for each player. Next lecture we introduce a more general method.