EVOLUTION AND THE THEORY OF GAMES

STEFAN GERITZ, HELSINKI, 2011

Transcript of my notes of 08-11-2011

12. A Nash equilibrium in an N-person game with strategy sets $X_1, \dots, X_N$ is a strategy combination $(\hat{x}_1, \dots, \hat{x}_N) \in X_1 \times \cdots \times X_N$ such that

$$\pi_i(\hat{x}_1, \dots, \hat{x}_{i-1}, x_i, \hat{x}_{i+1} \dots, \hat{x}_N) \le \pi_i(\hat{x}_1, \dots, \hat{x}_N) \quad \forall x_i \in X_i$$

for all $i \in \{1, \dots, N\}$.

Theorem (*Nash existence theorem*). Every N-person with finitely many pure strategies game has at least one Nash equilibrium if mixed strategies are allowed. □

For the proof we shall need the following result:

Theorem (*Brouwer's fixed point theorem*). A continuous function that maps a closed, bounded and convex subset of $\mathbb{R}^n$ into itself has at least one fixed point. □

We don't prove Bouwer's fixed point theorem, but make sure you understand what it says: the meaning of a set being closed and bounded is clear. A set $C \in \mathbb{R}^n$ is convex if for every $c_1, c_2 \in C$ and $0 \le p \le 1$ the weighed average $pc_1 + (1-p)c_2 \in C$. The set of all mixed strategies over a finite set of pure strategies is convex, because any mix of mixed strategies is also a mixed strategy. Furthermore, x is a fixed point of a function f if $f(x) = x$.

Proof (*Nash existence theorem*). The proof is for the case $N = 2$ but is readily generalized to any number of players. First some notation and definitions:

Let $x_1, \dots, x_m$ and $y_1, \dots, y_n$ be pure strategies for, respectively, the row- and column-player, and let $x = (p_1, \dots, p_m) \in X$ and $y = (q_1, \dots, q_n) \in Y$ be mixed strategies. Define the functions

$$\begin{aligned} c_i(x, y) &= \max\{0, \pi_1(x_i, y) - \pi_1(x, y)\} \\ d_i(x, y) &= \max\{0, \pi_2(x, y_i) - \pi_2(x, y)\} \end{aligned}$$

and

$$f(x,y) = \begin{pmatrix} f_X(x,y) \\ \\ f_Y(x,y) \end{pmatrix} = \begin{pmatrix} \frac{p_1+c_1(x,y)\ ,\ \dots\ ,\ p_m+c_m(x,y)}{1+\sum_i c_i(x,y)} \\ \\ \frac{q_1+d_1(x,y)\ ,\ \dots\ ,\ q_m+d_n(x,y)}{1+\sum_i d_i(x,y)} \end{pmatrix}$$

We claim that f has a fixed point $(\hat{x},\hat{y}) \in X \times Y$ and this fixed point is a Nash equilibrium.

To see that f has a fixed point, note that X and Y are convex, and so is $X \times Y$. Moreover, $f_X(x,y)$ and $f_Y(x,y)$ are probability distributions over $x_1,\dots,x_m$ and $y_1,\dots,y_n$ (check this!), and so f takes values in $X \times Y$. Finally, f is continuous, and so we can apply Brouwer's fixed point theorem.

To see that a fixed point of f is a Nash equilibrium, we prove that $\pi_1(x,\hat{y}) \leq \pi_1(\hat{x},\hat{y})\ \forall x \in X$. (The proof of $\pi_2(\hat{x},y) \leq \pi_2(\hat{x},\hat{y})\ \forall y \in Y$ is similar and therefore will not be given.)

Write $\hat{x} = (\hat{p}_1,\dots,\hat{p}_m)$, and take $i_0 \in \{1,\dots,m\}$ such that $c_{i_0}(\hat{x},\hat{y}) = 0$. Note that such i_0 exists, because if it did not, then $c_i(\hat{x},\hat{y}) > 0\ \forall i$, i.e., $\pi_1(x_i,\hat{y}) > \pi_1(\hat{x},\hat{y})\ \forall i$, and so $\pi_1(\hat{x},\hat{y}) = \sum_i \hat{p}_i \pi_1(x_i,\hat{y}) > \sum_i \hat{p}_i \pi_1(\hat{x},\hat{y}) = \pi_1(\hat{x},\hat{y})$, which is a contradiction.

Since $(\hat{x},\hat{y}) = f(\hat{x},\hat{y})$, we have in particular

$$\hat{p}_{i_0} = \frac{\hat{p}_{i_0} + c_{i_0}(\hat{x},\hat{y})}{1+\sum_i c_i(\hat{x},\hat{y})} = \frac{\hat{p}_{i_0}}{1+\sum_i c_i(\hat{x},\hat{y})}$$

and so $\sum_i c_i(\hat{x},\hat{y}) = 0$. Since by definition the c_i are non-negative, it follows that $c_i(\hat{x},\hat{y}) = 0$, i.e., $\pi_1(x_i,\hat{y}) \leq \pi_1(\hat{x},\hat{y})\ \forall i$. Hence, for arbitrary $x = (p_1,\dots,p_m) \in X$ we have

$$\pi_1(x,\hat{y}) = \sum_i p_i \pi_1(x_i,\hat{y}) \leq \sum_i p_i \pi_1(\hat{x},\hat{y}) = \pi_1(\hat{x},\hat{y})$$

which is what we set out to prove. □

13. In the previous lecture we have seen that the game with payoff matrix

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

does not have a Nash equilibrium of pure strategies. However, the Nash existence theorem tells us that there is a Nash equilibrium if we allow for mixed strategies. We used the Swastika method to find out what that equilibrium is. The Swastika method, however, is practical for 2×2-payoff matrices only. Here is a result that can be used more generally:

Theorem (*Bishop-Cannigns theorem*). If $(\hat{X}, \hat{y})$ is a Nash equilibrium, then $\pi_1(x, \hat{y}) = \pi_1(\hat{x}, \hat{y})$ for every pure strategy x in the support of $\hat{x}$, and $\pi_2(\hat{x}, y) = \pi_2(\hat{x}, \hat{y})$ for every pure strategy y in the support of $\hat{y}$. $\square$

The proof is given as an exercise. To see how the Bishop-Cannings theorem is used, consider the popular Rock-Paper-Scissors game with payoff matrix:

	R	P	S
R	0 , 0	-1 , 1	1 , -1
P	1 , -1	0 , 0	-1 , 1
S	-1 , 1	1 , -1	0 , 0

Obviously there is no *pure strategy* Nash equilibrium: the best reply to Rock is Paper, and the best reply to Paper is Scissors, and the best reply to Scissors is Rock, and so there is no pair of strategies that are best replies to each other. Still, from the Nash existence theorem we know that there does exist at least one *mixed strategy* Nash equilibrium.

We first look for a solution $(\hat{x}, \hat{y})$ with full support for both strategies. From the Bishop-Cannings theorem we have

$$\pi_1(R, \hat{y}) = \pi_1(P, \hat{y}) = \pi_1(S, \hat{y}) = \pi_1(\hat{x}, \hat{y})$$

Written out for $\hat{y} = (q_1, q_2, q_3)$ and the given payoff matrix this becomes

$$-q_2 + q_3 = +q_1 - q_3 = -q_1 + q_2 = \pi_1(\hat{x}, \hat{y})$$

which, together with $q_1 + q_2 + q_3 = 1$, form a system of four independent linear equations in q_1, q_2, q_3 and $\pi_1(\hat{x}, \hat{y})$. Solving these equations gives

$$q_1 = q_2 = q_3 = \frac{1}{3} \quad \& \quad \pi_1(\hat{x}, \hat{y}) = 0$$

Similarly, for $\hat{x} = (p_1, p_2, p_3)$ we find

$$p_1 = p_2 = p_3 = \frac{1}{3} \quad \& \quad \pi_2(\hat{x}, \hat{y}) = 0$$

The conditions in the definition of the Nash equilibrium are readily verified to show that $(\hat{x}, \hat{y})$ is a Nash equilibrium, indeed.

In fact, the above equilibrium is also the only equilibrium of the game. The proof of this is left as an exercise.

There are two known examples of the Rock-Paper-Scissors game in nature. The first is found in a species of Californian lizards. The males come in three types: (1) the ultra-dominant polygynous orange-throated males,(2) the mate-guarding monogamous blue-throated males, and (3) the female-mimicking yellow-throated males. Orange wins against Blue but looses against Yellow, because the latter looks like a female and can sneak into the harem of Orange to mate with the females. Yellow, however, looses against the Blue, because the latter guards a single female against any intruder. So, the best reply to Orange is Yellow, and the best reply to Yellow is Blue, and the best reply to Blue is Orange. (*For further details see the paper by Sinervo and Livley in* Nature*(1996) 340: 240-243.*)

The second example is found in *E. coli*, the common gut bacteria. There are three types: (1) type C produces a toxin called colicin but is itself not affected, (2) type R is resistant to colicin but cannot make it, and (3) type S is susceptible and gets killed when exposed to colicin. Type R grows faster than type C and outcompetes it in an initially mixed culture, because it does not have the cost of producing the toxin. Type S grows even faster than R, because it does not have the cost of developing and maintaining a detoxification mechanism. Type C, however, invades and outcompetes a S, because the latter has no defense against the toxin. So R beats C, C beats S, and S beats R.

14. In spite of its charms, the Nash equilibrium also has its problems: e.g., 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	$\frac{1}{2}R$, $\frac{1}{2}R$

with $R < C$ so that (H,D) and (D,H) are Nash equilibria. But how is such an outcome realized? The Hawk gets the resource and the Dove gets nothing. Why would one player voluntarily take the role of Dove? Without some form of coordination there is a high chance that the players end up with (H,H) or (D,D). In other words, accepting the Nash equilibria (H,D) and (D,H) as credible solutions of the game shifts the emphasis from the game to pre-game negotiations and contract. In the context of animal behavior, such negotiations may or may not be realistic, but this is the reason why we introduce one more (and last) solution concept.

15. An *evolutionarily stable strategy* (ESS) is a strategy such that, if adopted by a sufficiently large fraction of the population, then no other strategy can invade, i.e., increase in frequency (*Maynard Smith 1982*).

Note that "invasion of one strategy in a population of an other strategy" is changing population densities over time and hence about population dynamics. So, it should not come as a surprise that the *exact mathematical conditions* that make a strategy an ESS depend on the population dynamical embedding of the game. There are various ways how to do this, but the following is by far the most common:

Let $x', x \in X$ be two strategies that occur in the population with relative frequencies ε and $1 - \varepsilon$, and assume that opponents are assigned randomly so that the probability of being paired with an x'-player is ε, and the probability of being paired with a x-player is $1-\varepsilon$. Then the expected payoff to an x'-player is

$$\varepsilon\,\pi_1(x',x') + (1-\varepsilon)\pi_1(x',x)$$

and to an x-player is

$$\varepsilon\,\pi_1(x,x') + (1-\varepsilon)\pi_1(x,x)$$

(Note that the constancy of being paired with a player of a given type implies an infinitely large population or "sampling with replacement".)

Suppose further that x' cannot increase in relative frequency if and only if x'-players are "doing worse" than x-players in terms of their expected payoff, i.e., if

$$\varepsilon\,\pi_1(x',x') + (1-\varepsilon)\pi_1(x',x) < \varepsilon\,\pi_1(x,x') + (1-\varepsilon)\pi_1(x,x)$$

Collecting terms with the same order of ε, this can be rewritten as

$$\big[\pi_1(x',x) - \pi_1(x,x)\big] + \varepsilon\big[\pi_1(x',x') - \pi_1(x,x') - \pi_1(x',x) + \pi_1(x,x)\big] < 0$$

For sufficiently small ε the sign of the left hand side is dominated by the sign of $\pi_1(x',x) - \pi_1(x,x)$. Only when this term is zero, the sign of the remaining order-ε term $\pi_1(x',x') - \pi_1(x,x')$ matters.

Hence, x is evolutionarily stable (in the context of random pairing with replacement) if and only if for every $x' \neq x$ we have either

$$\pi_1(x',x) < \pi_1(x,x) \quad (\text{“}1^{\text{st}}\text{ ESS condition”})$$

or

$$\pi_1(x',x) = \pi_1(x,x)\ \&\ \pi_1(x',x') < \pi_1(x,x') \quad (\text{“}2^{\text{nd}}\text{ ESS condition”})$$

(Often these are given as the definition of an ESS, but in fact this is suitable only in the case of random pairing of opponents with replacement.)

16. Note that:

(1) If x is an ESS, then $\pi_1(x',x) \leq \pi_1(x,x)\ \forall x'$, and so the strategy pair (x,x) is a (symmetric) Nash equilibrium.

(2) Not every symmetric Nash equilibrium corresponds to an ESS, because the ESS conditions are more restrictive.

(3) Like the two strategies of a Nash equilibrium (x,y) are best responses to one another, so is an ESS a best response to itself.

(4) Since an ESS corresponds to a Nash equilibrium, the Bishop-Cannings theorem (see section 13) applies: if x is an ESS, then $\pi_1(x',x) = \pi_1(x,x)$ for every pure strategy x' in the support of x.

Next lecture we have a look at examples.