EVOLUTION AND THE THEORY OF GAMES

STEFAN GERITZ, HELSINKI, 2011

Transcript of my notes of 01-12-2011

41. A *multi-stage game* Γ consists over several rounds of the same or different games $\Gamma_1, \dots, \Gamma_N$ called the stages of the game. In each stage the payoff is a real payoff plus a "ticket" to the same or another stage of the game played against the same or another opponent. The strategy sets $X_1, \dots, X_N$ and $Y_1, \dots, Y_N$ of the row- and the column-player for different stages may be different or the same. The payoff functions $\pi_i : X_i \times Y_i \to \mathbb{R}^2$, too, have to be specified for each stage of the game separately. We thus have

$$\Gamma = (\Gamma_1, \Gamma_2, \dots, \Gamma_N)$$

with strategy sets

$$X = X_1 \times X_2 \times \cdots \times X_N \quad \text{and} \quad Y = Y_1 \times Y_2 \times \cdots \times Y_N$$

and overall payoff function

$$\pi : X \times Y \to \mathbb{R}^2$$

and a solution concept, which we shall take to be the ESS. The rub is how to calculate the overall payoff function π from the payoff functions π_i of the separate stages.

42. An *iterated game* is a special form of a multi-stage game $\Gamma = (\Gamma_1, \Gamma_2, \dots)$ with infinitely many stages, all with the same strategy sets $X_1 = X_2 = \dots$ and $Y_1 = Y_2 = \dots$. For example, in the iterated Prisoner's Dilemma we have $X_i = Y_i = \{C, D\}$ for all i, and the payoff functions π_i are defined by the payoff matrix

Γ_i	C	D
C	$R + \delta\, \Gamma_{i+1},\ R + \delta\, \Gamma_{i+1}$	$S + \delta\, \Gamma_{i+1},\ T + \delta\, \Gamma_{i+1}$
D	$T + \delta\, \Gamma_{i+1},\ S + \delta\, \Gamma_{i+1}$	$P + \delta\, \Gamma_{i+1},\ P + \delta\, \Gamma_{i+1}$

where $\delta \in (0, 1)$ is the discounting factor (i.e., the probability of that there is a next round) and Γ_{i+1} represents the "ticket" to the next stage of the game.

43. Consider the Hawk-Dove game where the looser of a H×H contest must skip one round, or more, to recover (R) from its injuries. The winner, on the other hand, continues with another round of the Hawk-Dove game against another opponent.

The strategy of the new opponent is the same as that of the previous one, because we want to calculate the payoff to a rare invader strategy against the resident strategy. This situation can be modeled as a two-stage game $\Gamma = (\Gamma_1, \Gamma_2)$ where Γ_1 is the Hawk-Dove game with payoff matrix

Γ_1	H	D
H	$\frac{1}{2}(R+\delta\Gamma_1)+\frac{1}{2}\delta\Gamma_2$	$R+\delta\Gamma_1$
D	$\delta\Gamma_1$	$\frac{1}{2}R+\delta\Gamma_1$

(symmetric game: payoffs to the row player)

and Γ_2 is a recovery round with payoff matrix

Γ_2	
R	$\varepsilon\delta\Gamma_1+(1-\varepsilon)\delta\Gamma_2$

(one-player game)

where ε is the probability of recovery and δ is the probability of playing another round. Notice that Γ_2 has only one player: the loser of the H×H contest. Also notice that the cost of injury (previously denoted by C) in Γ_1 is now replaced by not being able to play and gather resources for one or more rounds.

If only pure strategies are allowed, then strategy sets of the overall game $\Gamma = (\Gamma_1, \Gamma_2)$ are $X = Y = \{(\text{H,R}),(\text{D,R})\}$. The calculation of the overall payoffs is not much different from how we calculated payoffs for iterated games.

(H,R)×(H,R) – Let E_1 and E_2 denote the payoff to the row-player if starting with, respectively, Γ_1 or Γ_2. Then

$$\begin{cases} E_1 = \frac{1}{2}(R+\delta E_1)+\frac{1}{2}\delta E_2 \\ E_2 = \varepsilon\delta E_1+(1-\varepsilon)\delta E_2 \end{cases}$$

from which we solve

$$E_1 = \frac{R\big(1-\delta(1-\varepsilon)\big)}{(1-\delta)\big(2-\delta(1-2\varepsilon)\big)}$$

Calculation of the payoffs for the other strategy combinations is trivial (see lecture on iterated games) because they do not involve switching between the different stages of the game. The overall payoff matrix then becomes

Γ	(H,R)	(D,R)
(H,R)	$\frac{R(1-\delta(1-\varepsilon))}{(1-\delta)(2-\delta(1-2\varepsilon))}$	$\frac{R}{1-\delta}$
(D,R)	0	$\frac{R}{2(1-\delta)}$

(symmetric game: payoffs to the row player)

It follows that (H,R) is evolutionarily stable. Since the support of an ESS cannot be a subset of the support of another ESS, (H,R) is the only ESS of the game, and so we do not have to look for a mixed ESS.

44. As a variation on the previous game, assume that the winner of a Hawk×Hawk contest is not paired with another opponent, but simply takes (T) the resource each round until his opponent has recovered (R) and can play again. We then have a two-stage game $\Gamma = (\Gamma_1, \Gamma_2)$ where Γ_1 is a symmetric Hawk-Dove game:

Γ_1	H	D
H	$\frac{1}{2}(R + \delta\Gamma_2^{\text{col}}) + \frac{1}{2}\delta\Gamma_2^{\text{row}}$	$R + \delta\Gamma_1$
D	$\delta\Gamma_1$	$\frac{1}{2}R + \delta\Gamma_1$

(symmetric game: payoffs to the row player)

Notice that the winner of the H×H contest in Γ_1 gets a "ticket" to play the column player in Γ_2 while the looser gets a "ticket" to play the row player. So, Γ_2 is an asymmetric game with two roles: injured (row player) and uninjured (column player):

Γ_2	T
R	$\varepsilon\delta\Gamma_1 + (1-\varepsilon)\delta\Gamma_2^{\text{row}}$, $R + \varepsilon\delta\Gamma_1 + (1-\varepsilon)\delta\Gamma_2^{\text{col}}$

(asymmetric game: injured (row player), not injured (column player))

The strategy sets are

$$X = Y = \left\{ \left(\text{H}, \begin{pmatrix}\text{T}\\ \text{R}\end{pmatrix}\right), \left(\text{D}, \begin{pmatrix}\text{T}\\ \text{R}\end{pmatrix}\right) \right\}$$

Here is how we calculate the payoffs for the overall game Γ:

(H,(T,R))×(H,(T,R)) – Let E_1 denote the payoff to the row-player if starting with Γ_1, and let E_2^{row} and E_2^{col} be the payoff to players if starting with Γ_2 in the role of row player and column player, respectively. Then

$$\begin{cases} E_1 & = & \frac{1}{2}(R + \frac{1}{2}\delta E_2^{\text{col}}) + \frac{1}{2}\delta E_2^{\text{row}} \\ E_2^{\text{row}} & = & \varepsilon\delta E_1 + (1-\varepsilon)\delta E_2^{\text{row}} \\ E_2^{\text{col}} & = & R + \varepsilon\delta E_1 + (1-\varepsilon)\delta E_2^{\text{col}} \end{cases}$$

If we write $E_2 \overset{def}{=} \frac{1}{2}(E_2^{\text{row}} + E_2^{\text{col}})$, then the above simplifies to

$$\begin{cases} E_1 &= \frac{1}{2}R + \delta E_2 \\ E_2 &= \frac{1}{2}R + \varepsilon\delta E_1 + (1-\varepsilon)\delta E_2 \end{cases}$$

from which we solve

$$E_1 = \frac{R}{2(1-\delta)}$$

The payoffs for the other strategy combinations stays the same as in the previous example. The overall payoff matrix thus becomes

Γ	(H,(T,R))	(D,(T,R))
(H,(T,R))	$\frac{R}{2(1-\delta)}$	$\frac{R}{1-\delta}$
(D,(T,R))	0	$\frac{R}{2(1-\delta)}$

(symmetric game: payoffs to the row player)

It follows that (H,(T,R)) is evolutionarily stable. Since the support of an ESS cannot be a subset of the support of another ESS, this is also the only ESS of the game.

45. In the previous two example, the players could not choose what strategy to play in the second stage of the game: there was only one (conditional) strategy in Γ_2. Here is another variation on the multi-stage Hawk-Dove game that is a bit more interesting:

Suppose that an injured player still actively participates in the game but now only with a probability $p < \frac{1}{2}$ of winning the H×H contest against an uninjured opponent. If a player gets injured twice in a row, then he dies. His opponent, however, continues against another (uninjured) opponent. If an injured player plays Dove, or plays Hawk against Dove, an escalated fight is avoided, and the player recovers from his injury. An injured player that plays Hawk against Hawk does not recover, also not if he wins the contest.

The above situation can be modeled as a three-stage game $\Gamma = (\Gamma_1, \Gamma_2, \Gamma_3)$ where Γ_1 is the ordinary symmetric Hawk-Dove game between two uninjured players:

Γ_1	H	D
H	$\frac{1}{2}(R + \delta\Gamma_2^{\text{col}}) + \frac{1}{2}\delta\Gamma_2^{\text{row}})$	$R + \delta\Gamma_1$
D	$\delta\Gamma_1$	$\frac{1}{2}R + \delta\Gamma_1$

(symmetric game: payoffs to the row player)

The second stage, Γ_2, is an asymmetric Hawk-Dove game with two roles: Injured (row player) and Uninjured (column player):

Γ_2	H	D
H	$p(R+\delta\Gamma_3)$, $p\delta\Gamma_3+(1-p)(R+\delta\Gamma_1)$	$R+\delta\Gamma_1$, $\delta\Gamma_1$
D	$\delta\Gamma_1$, $R+\delta\Gamma_1$	$\frac{1}{2}R+\delta\Gamma_1$, $\frac{1}{2}R+\delta\Gamma_1$

(asymmetric game: row player is injured; column player is uninjured)

Notice that the column player is already injured and therefore in a H×H contest has only a probability $p < \frac{1}{2}$ of winning against the column player who is not injured. If the row player wins, he gets the resource R but he does not recover, and so both players are now injured and go over to stage Γ_3. If the row player looses, however, which happens with probability $1-p$, then the game for him is over (because he got injured twice in a row), and the column player gets the resource and (since he is still uninjured) starts afresh in stage Γ_1 with a new (uninjured) opponent.

In the third stage Γ_3 both players are injured, but otherwise the game is like an ordinary symmetric Hawk-Dove game:

Γ_3	H	D
H	$\frac{1}{2}(R+\delta\Gamma_2^{\text{row}})$	$R+\delta\Gamma_1$
D	$\delta\Gamma_1$	$\frac{1}{2}R+\delta\Gamma_1$

(symmetric game: payoffs to the row player)

Notice that whichever player looses the H×H contest got injured twice in a row, and so for him the game is over. The winner, however, gets the resource R and continues to play in stage Γ_2 as row player (because he is still injured) against a fresh (uninjured) opponent.

There are two pure strategies for Γ_1, four pure (conditional) strategies for Γ_2 and two pure strategies for Γ_3. For the game $\Gamma = (\Gamma_1, \Gamma_2, \Gamma_3)$ there are thus sixteen different pure strategies. To illustrate how to calculate payoffs for the overall game Γ, we compare here only two strategies, i.e.,

$$\left(H, \binom{H}{H}, H\right) \quad \times \quad \left(H, \binom{H}{H}, D\right)$$

The first strategy always plays Hawk; the second strategy plays Hawk except when in Γ_3, where both he and his opponent are injured. I do not suggest that these are particular good strategies; as a matter of fact, I do not know which of the sixteen possible pure strategies are evolutionarily stable. As an exercise it might be interesting to propose your own champion and test that one against other strategies.

Let E_1 denote the payoff to the row-player if starting with Γ_1, and let $E_2^{\rm row}$ and $E_2^{\rm col}$ be the payoff to players if starting with Γ_2 in the role of, respectively, the row player and the column player, and let E_3 denote the payoff if starting in Γ_3.

For **(H,(H,H),H)×(H,(H,H),H)** we get

$$\begin{cases} E_1 &= \frac{1}{2}(R+\delta E_2^{\rm col}) + \frac{1}{2}\delta E_2^{\rm row} \\ E_2^{\rm col} &= p\delta E_3 + (1-p)(R+\delta E_1) \\ E_2^{\rm row} &= p(R+\delta E_3) \\ E_3 &= \frac{1}{2}(R+\delta E_2^{\rm row}) \end{cases}$$

from which we solve

$$E_1 = R\cdot\frac{2+2\delta+p\delta^2-p(1-2p)\delta^3}{4-2\delta^2+p(1-p)\delta^4} \tag{1}$$

For **(H,(H,H),H)×(H,(H,H),D)** we have

$$\begin{cases} E_1 &= \frac{1}{2}(R+\delta E_2^{\rm col}) + \frac{1}{2}\delta E_2^{\rm row} \\ E_2^{\rm col} &= p\delta E_3 + (1-p)(R+\delta E_1) \\ E_2^{\rm row} &= p(R+\delta E_3) \\ E_3 &= R+\delta E_1 \end{cases}$$

from which we solve

$$E_1 = R\cdot\frac{1+\delta+2p\delta^2}{2-(1-p)\delta^2-2p\delta^3} \tag{2}$$

For **(H,(H,H),D)×(H,(H,H),H)** we have

$$\begin{cases} E_1 &= \frac{1}{2}(R+\delta E_2^{\rm col}) + \frac{1}{2}\delta E_2^{\rm row} \\ E_2^{\rm col} &= p\delta E_3 + (1-p)(R+\delta E_1) \\ E_2^{\rm row} &= p(R+\delta E_3) \\ E_3 &= \delta E_1 \end{cases}$$

which gives

$$E_1 = R \cdot \frac{1+\delta}{2-(1-p)\delta^2 - 2p\delta^3} \tag{3}$$

And for **(H,(H,H),D)**$\times$**(H,(H,H),D)** we have

$$\begin{cases} E_1 &= \frac{1}{2}(R+\delta E_2^{\text{col}}) + \frac{1}{2}\delta E_2^{\text{row}} \\ E_2^{\text{col}} &= p\delta E_3 + (1-p)(R+\delta E_1) \\ E_2^{\text{row}} &= p(R+\delta E_3) \\ E_3 &= \frac{1}{2}R + \delta E_1 \end{cases}$$

from which we get

$$E_1 = R \cdot \frac{1+\delta+p\delta^2}{2-(1-p)\delta^2 - 2p\delta^3} \tag{4}$$

It can be seen from equations (1)-(4) that the actual value of R does not matter for the ordering of the payoffs. So, without loss of generality we can put $R = 1$. Then, for the payoff matrix of the overall game Γ we have

Γ	(H,(H,H),H)	(H,(H,H),D)
(H,(H,H),H)	$\frac{2+2\delta+p\delta^2-p(1-2p)\delta^3}{4-2\delta^2+p(1-p)\delta^4}$	$\frac{1+\delta+2p\delta^2}{2-(1-p)\delta^2-2p\delta^3}$
(H,(H,H),D)	$\frac{1+\delta}{2-(1-p)\delta^2-2p\delta^3}$	$\frac{1+\delta+p\delta^2}{2-(1-p)\delta^2-2p\delta^3}$

(symmetric game: payoffs to the row player)

From the matrix, one readily shows that (H,(H,H),D) can always be invaded by (H,(H,H),H). However (H,(H,H),H) cannot be invaded by (H,(H,H),D) whenever

$$\frac{2+2\delta+p\delta^2-p(1-2p)\delta^3}{4-2\delta^2+p(1-p)\delta^4} > \frac{1+\delta}{2-(1-p)\delta^2-2p\delta^3}$$

which corresponds to points (δ, p) in the region left of the thick line in the following figure:

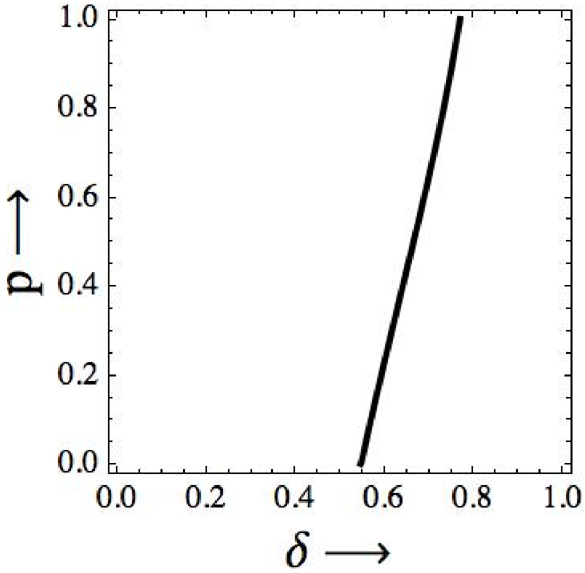

So, for (δ, p) on the left side of the thick line, (H,(H,H),H) is evolutionarily stable; for (δ, p) on the right side of the line, (H,(H,H),H) and (H,(H,H),D) can mutual invade, which means that they will coexist (and also that there may exist some other strategy that cannot be invaded by either), but we shall not pursue this issue further.

Since p is the probability that an injured player wins a H×H contest against a player that is not injured, only values of p less than one-half (and probably much less) are reasonable.

46. What is the expected length of the game Γ (measured in number of rounds) for given p and δ and a given strategy combination?

Let L_1 denote the expected length for a given strategy combination starting in Γ_1, and let L_2^{col} and L_2^{row} denote the expected length if starting in Γ_2 in the roles of, respectively, the column player and the row player, and let L_3 denote the expected length if starting in Γ_3.

Then for **(H,(H,H),H)×(H,(H,H),H)** we find

$$\begin{cases} L_1 & = & 1 + \frac{1}{2}\delta L_2^{\mathrm{col}} + \frac{1}{2}\delta L_2^{\mathrm{row}} \\ L_2^{\mathrm{col}} & = & 1 + p\delta L_3 + (1-p)\delta L_1 \\ L_2^{\mathrm{row}} & = & 1 + p\delta L_3 \\ L_3 & = & 1 + \frac{1}{2}\delta L_2^{\mathrm{row}} \end{cases}$$

from which we solve

$$L_1 = \frac{4 + 4\delta + 2p\delta^2}{4 - 2\delta^2 - p(1-p)\delta^4} \tag{5}$$

For all the other contests we find

$$\begin{cases} L_1 &= 1 + \frac{1}{2}\delta L_2^{\text{col}} + \frac{1}{2}\delta L_2^{\text{row}} \\ L_2^{\text{col}} &= 1 + p\delta L_3 + (1-p)\delta L_1 \\ L_2^{\text{row}} &= 1 + p\delta L_3 \\ L_3 &= 1 + \delta L_1 \end{cases}$$

which gives

$$L_1 = \frac{2 + 2\delta + 2p\delta^2}{2 - (1-p)\delta^2 - 2p\delta^3} \tag{6}$$

The following figure gives the expected number of rounds for the **(H,(H,H),H)** $\times$**(H,(H,H),H)** contest (solid line) and the other contests (dashed line) for fixed $\delta = 0.9$:

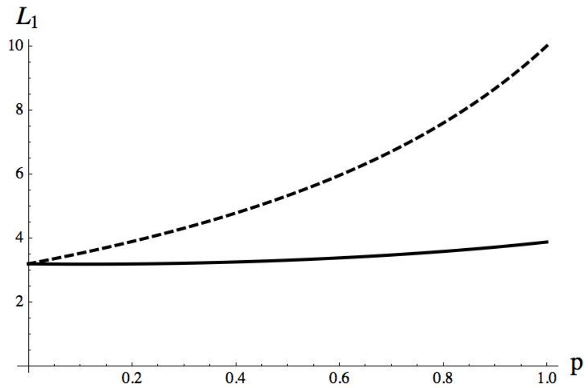

It can be seen that the first kind of contest on average lasts significantly shorter than the other kinds of contest, especially for larger values of p.