Evolutionary Game Theory

Luc Rey-Bellet

Classnotes for Math 697EG

Umass Amherst, Fall 2009

1 Introduction

We introduce some ideas and examples of game theory in an informal way. Let us consider a 2-players game where the player are called α and β. Each player has a certain number of strategies at his disposal.

- Player α has n strategies $s_1, \cdots, s_n$.
- Player β has m strategies $t_1, \cdots, t_m$.

A game (normal form game) is defined by specifying the payoffs received by the players if they choose a certain strategy. You can think of the payoff as, for example, a monetary payoff paid to the player after they choose their strategy and played the game but many interesting other interpretations will arise.

- $\pi_\alpha(i,j)$ is the payoff for player α if α chooses strategy s_i and β chooses strategy t_j.
- $\pi_\beta(i,j)$ is the payoff for player β if α chooses strategy s_i and β chooses strategy t_j.

Let $\boldsymbol{\pi}_\alpha$ and $\boldsymbol{\pi}_\beta$ be $n \times m$ matrices with entries $\pi_\alpha(i,j)$ and $\pi_\beta(i,j)$. The game is now conveniently represented in compact notation by using the bi-matrix notation. The bi-matrix has n rows and m columns, and each entry of the matrix gives the payoff for player α in the lower left corner and the payoff for β in the upper-right corner.

	t_1	$\cdots$	t_m
s_1	$\pi_\alpha(1,1)$ | $\pi_\beta(1,1)$	$\cdots$	$\pi_\alpha(1,m)$ | $\pi_\beta(1,m)$
$\vdots$	$\vdots$		$\vdots$
s_n	$\pi_\alpha(n,1)$ | $\pi_\beta(n,1)$	$\cdots$	$\pi_\alpha(n,m)$ | $\pi_\beta(n,m)$

Bi-matrix for normal form games

We assume here that each player knows all the strategies and all the payoffs for himself and the other player, i.e., *complete information* and moreover we assume that the game is *non-cooperative*, i.e., the players do not communicate with each other and cannot prearrange some choice of strategies. Let us illustrate theses concepts with a few examples.

Example 1.1 Two oil producing countries SA and IR can each produce either 2 millions or 4 millions barrels per day. For each of the countries the price of producing one barrels is \$5. The total production level will be either 4, 6, or 8 millions barrels per day. Due to the market demand the corresponding price per barrel will be \$25, \$17, or \$12. If we think the of payoff as the revenue for each country the bi-matrix payoff is given by

	$t_1 = 2$	$t_2 = 4$
$s_1 = 2$	40 | 40	24 | 48
$s_2 = 4$	48 | 24	28 | 28

One can see easily that for player SA it is more advantageous to produce 4 million barrels irrespective of what player IR does. If player IR chooses to produce a production level 2 millions, then SA is better off by choosing a production level of 4 millions and increase his payoff from 40 to 48. If player IR chooses to produce a production level 4 millions, then SA is again better off by choosing a production level of 4 millions and increase his payoff from 24 to 28. Here $s_1 = 2$ is an example of *a strictly dominated strategy* for the player SA. One can read dominated strategies directly from the payoff matrix

$$\boldsymbol{\pi}_{SA} = \begin{pmatrix} 40 & 24 \\ 48 & 28 \end{pmatrix} \tag{1.1}$$

we have $\pi_{SA}(1,j) < \pi_{SA}(2,j)$ for all j.

The payoff situation is the same for SA and IR and so both player rational choice is to choose a production level of 4 millions with a payoff of 28 millions. It would be of course much more advantageous for both to choose the lower level of 2 millions and earn a payoff of 40 millions. But this is not the rational choice in non-cooperative game theory. This game is known (with a slightly different narrative) as the *Prisonner's dilemna*. It shows that the rational behavior is sometimes to "defect" rather than to "cooperate" even though cooperation would result in mutual benefit. There is no real "dilemna" here, but rather just a lack of mechanism to ensure cooperation in non-cooperative games.

Example 1.2 The second example is the child game known as *Rock-Paper-Scissors*. In this game each player choose among the strategies s_1=Rock, s_2=Paper, or s_3=Scissor, and Paper wins over Rock, Rock wins over Scissor, and Scissor wins over Paper. The winner gets \$1 from the looser and no money is exchanged in case of a tie. The bi-matrix of the game is given by

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

This game is an example of a *zero-sum game*, i.e., the payoff for one player is the negative of the payoff for the other player. This is characteristic of games understood in their usual sense, say chess, or tic-tac-toe. In this case one player wins and the other looses or there is tie.

Every child who plays this game will learn very quickly that the best way to play this game is to play at random and choose any of the three strategies with equal probability 1/3. This is an example of a *mixed strategy* which can represented using a probability vector

$$\mathbf{p} = (p_1, \cdots, p_n), \quad p_i \geq 0, \quad \sum_{i=1}^{n} p_i = 1 .$$

Example 1.3 In traditional game theory which was applied primarily to economics the players are thought of being rational beings who act in such a way as to maximize their payoffs. In the 1970's the biologist John Maynard's Smith was the first to point out that concepts of game theory can be applied animal behavior to explain evolution. In this case the players are not necessarily rational but simple have a strategy dictated by a certain genotype. The payoff is interpreted as the fitness (or fitness increase) associated to a certain behavior, and fitness is, by definition, a measure of reproductionary success.

Imagine a species of territorial animals who engage in in ritual fights over territories. The behavior comes in 2 variants.

- The "Hawk" behavior is an aggressive behavior where the animal fights until either victory or injury ensues.
- The "Dove" behavior is to display hostility at first but to retreat at the first sign of attack from the opponent.

Let us assign a value v to the territory won after a fight, imagine for example that the winner will have access to more food and thus increase is fitness by v. We assign the value w to the cost of injury.

If two hawks meets they will escalate and fight, we will assume each one will win with probability 1/2, we take the the payoff to be the expected value of the increase of fitness, i.e.

$$v \times \ Pr\{ \text{ Win the fight} \} - w \times Pr\{ \text{ Loose the fight} \} = \frac{v - w}{2} .$$

Similarly if two doves meets, we assume that each of the dove might win, by default, with probability 1/2 and the payoff is $v/2$. Thus the payoff bi-matrix is then given by

	H	D
H	$\frac{v-w}{2}$, $\frac{v-w}{2}$	v, 0
D	0, v	$\frac{v}{2}$, $\frac{v}{2}$

If $v > w$, i.e., the cost of victory v exceeds the cost of injury, then the Dove strategy is dominated by the Hawk strategy. But if $w > v$ no strategy is dominated and a mixed strategy seems to be more appropriate. If we think here of (non-rational) animals we need to explain the meaning of mixed strategy. It does not really make sense to think of a genotype as giving such or such behavior with certain probabilities. It makes more sense to interpret a mixed strategy in tems of a population state. Let us imagine a very large population of N, if there are N_1 Hawks and N_2 Doves then we set $p_1 = N_1/N$ and $p_2 = N_2/N$. We can then think of playing against the mixed strategy $\mathbf{p}$ as the same as playing against a randomly choosen member in this large population.

If the population consists mostly of Hawks then they will spend a lot of time fighting each other and, as a result, they will loose fitness, (since $v < w$). It is then advantageous to be a Dove in such a population. On the contrary in a population which consists mostly of Doves, it is more advantageous to be a Hawk since fights will be won easily. So one would expect expect that there should be an "equilibrium" proportion of Hawks and Doves. To compute this equilibrium let us assume that the population state is $\mathbf{p} = (p_1, p_2)$ then the fitness increase for Hawks

$$\text{Fitness increase of Hawks} = p_1 \frac{v-w}{2} + p_2 v \,,$$

while for Doves it is

$$\text{Fitness increase of Doves} = p_2 \frac{v}{2} \,.$$

The fitnesses will equilibrate exactly when $p_1 = v/w$ and one would expect that a population would evolve over time to exhibit such proportion of Hawks and Doves. To make this conclusion more concrete one needs to specify a suitable evolution dynamics, which we will do later.

2 Payoffs and utilities

In this section we discuss several models of decision:

- Decision under certainty.
- Decision under risk.
- Decision under noise.

Decision under certainty means that each action leads to a specific determined outcome. In this case it is not difficult to assign a payoff function to represent the preferences in a consistent manner. In applications the payoff is the monetary outcome of having a given strategy (economic applications), or the reproductive fitness associated with a genotype (biological applications), or (minus) the energy associated to a state (physical applications). Then decision under certainty simply assumes that a *best response* is chosen, i.e., one maximizes the payoff function. If there are more than one best response, a suitable tie-breaker rule is chosen.

Decision under risk is discussed in subsection 2.2. In this case the strategies may involve chance and bets, but all probabilities are supposed to be known. For example one may think of a strategy which consist in buying a lottery ticket where you may win \$100 with probability p. The question we discuss here is to determine under which conditions we can assign preferences consistently between risky bets using a payoff function. We will show that under reasonnable conditions we can define a *utility function* which assigns a certain value to the various alternatives involved in the bets and that the payoff can be taken as the *expected value of the utility function* (Von Neumann–Morgenstern utility theory).

Decision under noise involves agents which may make error. There are many possible ways to model such behavior. One such models which connect payoff with information theory is discussed in subsection 2.3

We do not discuss here the case of *decision under uncertainty* in which the strategies involve chance but the probability are *unknown*. This leads to a Bayesian approach to payoff and to the so-called Bayesian games.

2.1 Strategies and payoffs, decisions under certainty

Let

$$S = \{s_1, \cdots, s_n\}$$

Figure 1: The simplex in $\mathbb{R}^2$ and $\mathbb{R}^3$

be a finite set containing n elements. Think of S as the set of possible actions or strategies being taken by a certain agent in a certain situation. We call s_i a *pure strategy* or *pure state*. The adjective "pure" alludes to the fact we also consider randomized strategies or *mixed strategies* or *mixed states*. In a mixed strategy an agent chooses the pure strategy s_i with some probability p_i. A mixed strategy is represented using a *probability vector*

$$\mathbf{p} = (p_1, \cdots, p_n)\,, \quad p_i \geq 0\,, \quad \sum_{i=1}^{n} p_i = 1\,.$$

The set of all strategies (mixed or pure) is denoted by $\Delta = \Delta(S)$. We can think of Δ as a subset of the n-dimensional euclidean space $\mathbb{R}^n$ and Δ is called the unit n-dimensional *simplex*. We can identify a pure strategy s_i with the standard basis vector $\mathbf{e}_i$ and we can write

$$\mathbf{p} = \sum_{i=1}^{n} p_i \mathbf{e_i}\,.$$

The set $\Delta(S)$ is a *convex set*, i.e., given two mixed strategies $\mathbf{p}$ and $\mathbf{q}$ and $0 \leq \alpha \leq 1$, the convex combination $\alpha\mathbf{p} + (1-\alpha)\mathbf{q}$ is also a mixed strategy. Geometrically it means for any two elements $\mathbf{p}$ and $\mathbf{q}$ in Δ , the line segment between $\mathbf{p}$ and $\mathbf{q}$ belongs to Δ. The pure strategies $\mathbf{e}_i$ are the *extremal points* of the convex set Δ, i.e., they are the points in Δ which cannot be written as a nontrivial convex combination of two other points.

For a state $\mathbf{p}$ we let

$$\Sigma_{\mathbf{p}} \;=\; \{i\,;\, p_i > 0\}\,.$$

and we call $\Sigma_{\mathbf{p}}$ the *support* of $\mathbf{p}$. If $i \in \Sigma_{\mathbf{p}}$ then the pure strategy s_i has positive probability in the mixed strategy $\mathbf{p}$. We call $\mathbf{p}$ an *interior point* of the simplex if every i belongs $\Sigma_{\mathbf{p}}$, i.e., all strategies have positive probability in the mixed strategy $\mathbf{p}$.

A *payoff function* $\boldsymbol{\pi} \,:\, S \to \mathbb{R}$ is a function which assigns a numerical value π_i to each pure strategy s_i. We identify the function $\boldsymbol{\pi}$ with the vector

$$\boldsymbol{\pi} = (\pi_1, \cdots, \pi_n) \in \mathbb{R}^n\,.$$

If one chooses a mixed strategy $\mathbf{p}$ then we can think of the payoff for $\mathbf{p}$ as the random variable which takes the value π_i with probability p_i and we define the payoff for the mixed strategy p to be the expected payoff

$$E_{\mathbf{p}}[\boldsymbol{\pi}] \equiv \sum_i \pi_i p_i = \langle \boldsymbol{\pi}, \mathbf{p} \rangle ,$$

where $\langle \cdot, \cdot \rangle$ denote the usual scalar product in $\mathbb{R}^n$.

Decision under certainty simply assumes that given a choice among strategies one chooses the one with the largest payoff, i.e., one maximizes the payoff over all strategies. We say then that s_i is a *pure best response* for the payoff $\boldsymbol{\pi}$ if

$$\pi_i = \max_k \pi_k .$$

There could be more that one pure best response strategy, in that case any mixed strategy $\mathbf{p}$ whose support contains only pure best responses is also a best response.

Formally let us call a strategy (pure or mixed) $\mathbf{p}^*$ a *best response* for the payoff $\boldsymbol{\pi}$ if $\mathbf{p}^*$ maximizes the payoff $\langle \boldsymbol{\pi}, \mathbf{p} \rangle$ and we will use the notation $\mathrm{BR}(\boldsymbol{\pi})$ for the set of best responses. We have then

$$\mathbf{p}^* \in \mathrm{BR}(\boldsymbol{\pi}) \quad \text{iff} \quad \langle \boldsymbol{\pi}, \mathbf{p}^* \rangle = \max_{\mathbf{p} \in \Delta} \langle \boldsymbol{\pi}, \mathbf{p} \rangle .$$

The following important fact is very easy to prove but is very important when we compute Nash equilibria.

Lemma 2.1 *For any payoff $\boldsymbol{\pi} \in \mathbb{R}^n$, the best response $\mathrm{BR}(\boldsymbol{\pi})$ is a face of the simplex Δ. In particular if $\mathbf{p}^*$ is a mixed strategy best response then all pure strategies in $\Sigma_{\mathbf{p}^*}$ are also best responses.*

Proof: If $\max_k \boldsymbol{\pi}_k$ is attained for a single i then the extremal point is $\mathbf{e}_i$ is the unique element in $\mathrm{BR}(\boldsymbol{\pi})$. If $\max_k \boldsymbol{\pi}_k$ is attained for several pure strategies $\mathbf{e}_{i_1}, \cdots, \mathbf{e}_{i_l}$, then any mixed strategy $\mathbf{p}$ with support $\Sigma_{\mathbf{p}} = \{i_1, \cdots, i_l\}$ maximizes the payoff $\langle \boldsymbol{\pi}, \mathbf{p} \rangle$.

Conversely suppose that $\mathbf{p}^*$ is a mixed strategy best response, then for all $i, j \in \Sigma_{\mathbf{p}^*}$ we must have $\boldsymbol{\pi}_i = \boldsymbol{\pi}_j$. If it were not the case and that for some i we have $\boldsymbol{\pi}_i \geq \boldsymbol{\pi}_l$ for all $l \in \Sigma_{\mathbf{p}^*}$ and $\boldsymbol{\pi}_i > \boldsymbol{\pi}_j$ for at least one $j \in \Sigma_{\mathbf{p}^*}$. Then we have

$$\pi_i = \pi_i \sum_{l \in \Sigma_{\mathbf{p}^*}} p_l^* > \sum_{l \in \Sigma_{\mathbf{p}^*}} \boldsymbol{\pi}_l p_l^* = \langle \boldsymbol{\pi}, \mathbf{p}^* \rangle$$

and this contradicts the assumption that p^* is a best response. ∎

2.2 Decision under risk; the Von Neumann-Morgenstern theorem

In many situations of interest the decisions are rather *decision under risk.* Each action leads to a set of possible outcomes, each outcome occurs with a certain probability. For examplethe action or strategy would be to play a bet which gives say \$10 with probability 1/2 and \$5 with probability 1/4 and \$0 with probability 1/4. In such a situation it is fairly natural to think of the payoff as the *expected value* of the bet, i.e., \$6.25 in the previous example.

The following famous example (the *St Petersbourg Paradox* due to D. Bernoulli) should serve as a cautionary tale about using monetary value too liberally. The best consists of tossing a fair coin repeteadly until a head appears. If head appears first on the n-th toss you shall be rewarded with 2^n dollars, with probability 1/2 you will win \$1, with probability 3/4 \$2 or less, etc... How much would you be willing to bet and lay this game? Well the expected money to be won in this game is $\sum_{n=1}^{\infty} \frac{1}{2^n} 2^n = 1 + 1 + 1 + \cdots = \infty$ therefore you should bet your car, your home, and your future earnings on that game, should'nt you.

The situation is however not entirely hopeless. We describe here an axiomatic treatement of decision under risk which is due (modulo some variants) to Von Neumann and Morgenstern. We follow here the presentation in [**?**] and [**?**]. The result is that under reasonnable assumptions which you could expect from a rational player then consistent preferences for decisions under risk can be achieved by computing the expectation of "something". This something is known has a utility function and the results is known as the *expected utility theorem*

If risk (and probabilities) are involved the strategies are usually called "lotteries" in this context to emphasize their random nature. We assume that all lotteries are built up from a finite set of basic alternatives or prizes

$$A_1, A_2, \cdots, A_r \, .$$

A *lottery* is a chance mechanism which yields the prize A_i with probability q_i. We will denote a lottery by

$$L = (A_1, q_1; \cdots, A_r, q_r) \, .$$

Note that we view a lottery as being conducted exactly once and not something repeated many times. A *utility function* is simply a function

$$u_i = u(A_i)$$

which assigns a utility value to the prize A_i. For a given lottery L the expected value will be

$$E_L[u] = \sum_i q_i u_i \,.$$

The first axiom is say that all basic prizes can be given consistent preferences, in mathematical terms there exists a total order.

Axiom A1. There exists a total order $\succeq$ on the set $(A_1, \cdots, A_r)$, i.e., for any A_i, A_j either $A_i \succeq A_j$ or $A_j \succeq A_j$ and if $A_i \succeq A_j$ and $A_j \succeq A_k$ then $A_i \succeq A_k$.

We will say that A_i is preferred to A_j if $A_i \succeq A_j$ and we say that A_i and A_j are equivalent if $A_i \succeq A_j$ and $A_j \succeq A_i$ and we write then $A_i \sim A_j$. Without loss of generality we can relabel the prizes so that

$$A_1 \succeq \cdots \succeq A_r \,,$$

i.e., the prizes are ordered.

Our second axiom involves lotteries with have exactly two prizes and and asserts that a player will choose a lottery which favors his preferred outcome.

Axiom A2. A player will prefer the lottery $(A_1, p; A_2(1-p))$ over the lottery $(A_1, q; A_2(1-q))$ if $p > q$.

This is a pretty reasonnable assumption and this allows to assign a utility to the set of lotteries involving only two prizes. We define a utility function $u_i = u(A_i)$ with $u_1 = a > u_2 = b$. Then the expected utility for the lottery $L = (A_1, p; A_2(1-p))$ is

$$E_L[u] \;=\; pa + (1-p)b = b + p(a-b) \,,$$

and since $a > b$, $E_L[u]$ is *increasing* in p and it thus represents the preferences of the player between all lotteries involving only A_1 and A_2. Note that the choice of a and b is irrelevant here, as long as $a > b$ and we see that if u is a utility then so is $v = \alpha u + \beta$ for any $\alpha > 0$ and $\beta \in \mathbb{R}$.

Imagine next that the lottery involve three prizes.

Axiom A3. If $A_1 \succeq A_2 \succeq A_3$ then the prize A_2 is equivalent to a lottery involving A_1 and A_3, i.e.

$$A_2 \sim (A_1, q; A_3(1-q)) \,.$$

for some probability q.

This is a *continuity assumption.* Since $A_1 \succeq A_2 \succeq A_3$ the assumption (A3) tells us that we can measure the relative preference between A_1 and A_i and A_i and A_r.

Suppose for example that $A_1 = \$100$ and $A_3 = \$0$ and A_2= a bag of apples. To asses the preference of the player ask her if she is willing to exchange her bag of apples against a lottery ticket paying \$100 with probability, say p. The smallest p^* for which she would accept the exchange with determine the value that the player attach the her bag of apples.

In this way we can assign a "value" or utility to any each prize A_i, with $A_1 \succeq A_i \succeq A_r$. By (A3) and (A2) there exists number $1 \geq q_2 \geq q_3 \cdots \geq 0$ such that $A_i \sim (A_1, q_j; A_3(1-q_j))$. In this way we can naturally define a utility function u_i for each prize A_i.

We need to verify that such utility function indeed represent the preferences of the player. For this we need two more assumptions.

Axiom A4. Player are indifferent if a prize in a lottery is replaced by another prize they see as equivalent.

Using the previous example it means that the bag of apples can be replaced by a lottery ticket with probability to win \$100 with some probability p^* and that these two lotteries are indifferent. As this example shows there is nothing which prevents us to choose a lottery as a prize in a lottery (compound lotteries). Think for example of buying a lottery ticket whose prize is the possibility to play at a poker tournament.

Finally our last axiom allows to reduce every lottery to a lottery involving only two prizes only and for such lotteries we have already defined a utility function.

Axiom A5. Suppose $L_j = (A_1, p_1^{(j)}; \cdots A_r, p_r^{(j)})$, $j = 1, \ldots, s$ are s lotteries each involving the same prizes $A_1, \ldots A_r$. Then the compound lotteries $(L_1, q_1, \cdots, L_s, q_s)$ is equivalent to $(A_1, p_1; \cdots A_r, q_r)$ where

$$p_i = \sum_{j=1}^{s} q_j p_i^{(j)} .$$

The assumption is seemingly innocuous and reasonnable and tells us that experiments can be combined a more complex lottery. There is however an assumption of *statistical independence* hidden behind (A5).

We can now conclude. Given a lottery $L = (A_1, p_1; \cdots; A_r, p_r)$ we use

Axiom (A3) and replace by $A_2, \dots, A_{r-1}$ by equivalent prizes (lotteries)

$$\tilde{A}_j = (q_j A_1; (1-q_j)A_r) \quad j = 2, \cdots r-1\,.$$

By axiom (A4) the lottery L is equivalent to a compound lottery $\tilde{L}$

$$L \sim \tilde{L} = (A_1, p_1; \tilde{A}_2, p_2, \cdots; \tilde{A}_{r-1}, p_{r-1} A_r, p_r)\,.$$

and $\tilde{L}$ only involve the prize A_1 and A_r.

By Axiom (A5) $\tilde{L}$ is equivalent to the lottery

$$\tilde{L} \sim \hat{L} = (A_1 s; A_r(1-s))$$

where

$$s = p_1 \times 1 + p_2 \times q_2 + \cdots p_{r-1} q_{r-1} + p_r \times 0\,.$$

This procedure defines a total among all lotteries involving r prizes A_1 and A_r: let us define the utility

$$u_1 = 1, u_2 = q_2, \cdots, u_r = 0$$

where the q_j are found using (A3). Then $E_L[u] = p_1 \times 1 + p_2 \times q_2 + \cdots p_{r-1} q_{r-1} + p_r \times 0$.

The previous discussion is summarized in

Theorem 2.2 *(Von Neuman–Morgenstern utility "theorem") Assume Axioms (A1)-(A5). Then there exists a total order among all lotteries involving a finite number of prizes. The order is given by computing the expectation $E_L[u]$ of a utility function u with $u_i = u(A_i)$. The utility u is uniquely defined up to a linear transformation $u' = \alpha u + \beta$ with $\alpha > 0$ and $\beta \in \mathbb{R}$.*

Remark 2.3 The economists use the shape of the utility function u as predicting if some agent is *risk-averse* or *risk-loving*. Imagine you have lottery where you win \$1 with probability $p = 3/4$ and win \$9 with probability $(1-p) = 1/4$, so the expected gain is \$3. You will be risk averse if you exchange your lottery ticket for an amount less than \$3 or risk loving if you are willing to exchange your lottery for more than \$3. You might repeat this experiment for various p and thus infer a utility function $u : [1,9] \to \mathbb{R}$. If u is linear we will say that the agent is risk neutral, if u is concave he is risk averse, while if u is convex he is risk loving loving.

We finish this section by discussing an example, called the *Allais paradox*, which shows that finding a consistent order of preferences between lotteries might be tricky. Consider the prizes

$$A_1 = \$0\,, \quad A_2 = \$1'000'000\,, \quad A_3 = 5'000'000$$

and the lotteries

$$L_1 = (A_1, 0; A_2, 1; A_3, 0) \quad L_2 = (A_1, 0.01; A_2, 0.89; A_3, 0.1)$$

Many people, not necessarily particularly risk averse will choose $L_1 \succeq L_2$, attracted by the sure payoff of \$1 million. Consider now the two lotteries

$$L_3 = (A_1, 0.89; A_2, 0.11; A_3, 0) \quad L_4 = (A_1, 0.9; A_2, 0; A_3, 0.1)$$

Then it is fairly natural to choose $L_4 \succeq L_3$. Let us consider a utility function which we assume, wlog, to have the form $u_1 = 0$, $u_2 = x$ and $u_3 = 1$ and x. Then we have

$$E_{L_1}[u] = x \tag{2.1}$$
$$E_{L_2}[u] = .89x + .1 \tag{2.2}$$
$$E_{L_3}[u] = .11x \tag{2.3}$$
$$E_{L_4}[u] = .01 \tag{2.4}$$

Since $L_1 \succeq L_2$ we should have $x > .89x + .1$ or $x > 10/11$. But since $L_4 \succeq L_3$ we have $.11x < .10$ or $x < 10/11$. This is a contradiction and so there is no utility which represents your choice of preferences.

2.3 Decision under noise; entropy

A perfectly rational agent would maximize his payoff function $\boldsymbol{\pi}$ and find the best response pure strategie(s). It is however useful to model agents who are not completely rational and make mistakes or maybe their decisions are submitted to some noise. There are of course many ways to model such thing. For example we could imagine that the agent choose *a perturbed best response*. Let $\epsilon > 0$ describe a noise level, then with probability ϵ the agent chooses a random strategy and with probability $1 - \epsilon$ it chooses the best response strategy. Instead of this we consider here a model based on information theoretic considerations which ensure that in the distribution of strategies, the ones with high payoff are more likely.

We introduce the concept of *relative entropy*. Given a probability vector $\mathbf{q} = (q_1, \cdots, q_n)$ we will say that another probability vector $\mathbf{p}$ is absolutely continuous with respect to $\mathbf{q}$ if $\Sigma_{\mathbf{p}} \subset \Sigma_{\mathbf{q}}$ (i.e., if $p_i = 0$ then $q_i = 0$) and we will use the notation $\mathbf{p} \ll \mathbf{q}$.

The function $f(x) = x\ln(x)$ define on $(0, \infty)$ is strictly convex since $f'(x) = 1/x$. Since $\lim_{x\to 0+} x\ln(x) = 0$ from now on we will extend f as a function on $[0, \infty)$ and use the convention$0\ln 0 = 0$.

Definition 2.4 *The* relative entropy of $\mathbf{p}$ with respect to $\mathbf{q}$, $H(\mathbf{p}\,|\,\mathbf{q})$, *is defined by*

$$H(\mathbf{p}\,|\,\mathbf{q}) = \begin{cases} \sum_{i\in\Sigma_{\mathbf{q}}} p_i \ln\left(\frac{p_i}{q_i}\right) = \sum_{i\in\Sigma_{\mathbf{q}}} \frac{p_i}{q_i} \ln\left(\frac{p_i}{q_i}\right) q_i & \text{if } \mathbf{p} \ll \mathbf{q} \\ +\infty & \text{otherwise} \end{cases} . \tag{2.5}$$

We have

Proposition 2.5 *The function* $\mathbf{p} \mapsto H(\mathbf{p}\,|\,\mathbf{q})$ *is a convex and continuous function on the set* $\{\mathbf{p} \in \Delta\,,\, \mathbf{p} \ll \mathbf{q}\}$. *Moreover* $H(\mathbf{p}\,|\,\mathbf{q}) \ge 0$ *and* $H(\mathbf{p}\,|\,\mathbf{q}) = 0$ *if and only if* $\mathbf{p} = \mathbf{q}$.

Proof: Since the function $x\ln x$ is a convex continuous function of x $H(\mathbf{p}\,|\,\mathbf{q})$ is continuous and convex in $\mathbf{p}$, whenever it is finite.

To see that $H(\mathbf{p}\,|\,\mathbf{q}) \ge 0$ we use the strict concavity of $\ln x$ which implies that $\ln x \le x - 1$ with equality iff $x = 1$. Note that if i is such that $p_i = 0$ the corresponding term in $H(\mathbf{p}\,|\,\mathbf{q})$ vanishes so that

$$\begin{aligned} H(\mathbf{p}\,|\,\mathbf{q}) &= \sum_{i\in\Sigma_p} p_i \ln\left(\frac{p_i}{q_i}\right) = -\sum_{i\in\Sigma_p} p_i \ln\left(\frac{q_i}{p_i}\right) \ge \sum_{i\in\Sigma_p} p_i\left(1 - \frac{q_i}{p_i}\right) \\ &= 1 - \sum_{i\in\Sigma_p} q_i \ \ge\ 0\,. \end{aligned} \tag{2.6}$$

Furthermore we equalities iff and only iff $p_i = q_i$ for all i. ∎

$H(\mathbf{p}\,|\,\mathbf{q})$ is a *lower semi-continuous function* of p, that is if $p^{(n)}$ is

Let $\mathbf{p}_u = (\frac{1}{n}, \cdots, \frac{1}{n})$ be the uniform probability distribution on S. We call

$$H(\mathbf{p}) \equiv H(\mathbf{p}\,|\,\mathbf{p}_u)$$

the *entropy of the probability vector* p. We have

Proposition 2.6 *The entropy $H(\mathbf{p})$ satisfies*

$$0 \leq H(\mathbf{p}) \leq \ln(n)\,.$$

and $H(\mathbf{p}) = \ln(n)$ iff p is a pure strategy and $H(\mathbf{p}) = 0$ iff $\mathbf{p} = \mathbf{p}_u = (\frac{1}{n}, \cdots, \frac{1}{n})$.

Proof: We have

$$H(\mathbf{p}) = H(\mathbf{p}\,|\,\mathbf{p}_u) \;=\; \sum_i p_i \ln(p_i) + \ln(n)\,.$$

Note that $x\ln(x) \leq 0$ if $x \in [0,1]$ and $x\ln(x) = 0$ iff $x = 0$ or 1. This implies that $H(\mathbf{p}) \leq \ln(n)$ and $H(\mathbf{p}) = \ln(n)$ iff $\mathbf{p} = \mathbf{e}_i$ is a pure strategy. On the other hand by Proposition 2.5 $H(\mathbf{p}) = 0$ iff $\mathbf{p} = \mathbf{p}_u$. ■

The entropy of the state $\mathbf{p}$ is a measure of the amount of order or information of the state $\mathbf{p}$. The larger the entropy the more order is associated to the state. The pure states corresponds to the maximal amount of order while the uniform distribution corresponds to the most disordered state. More generally you may think as $H(\mathbf{q}|\mathbf{q})$ as measuring a sort of "distance" between the probability distribution $\mathbf{q}$ and $\mathbf{q}$. But it is not a distance in the mathematical sense (it is not symmetric in $\mathbf{p}$ and $\mathbf{q}$.)

Remark: Different authors used different conventions for the entropy and relative entropy. In particular, in the physics literature, the entropy is usually defined as $-\sum_i p_i \ln(p_i)$, that the entropy is concave instead of convex and it is maximal for the uniform distribution: in this a larger entropy means more disorder.

It is not clear, a-priori, why the entropy should be a good choice of as the measure of the information of a probability vector. There are certainly other functions of p which are minimal for the uniform distribution and maximal for pure strategies. We sketch here a brief "axiomatic" justification of the choice of entropy as a measure of information.

We introduce three axioms that an information function $E(\mathbf{p})$ should satisfy

(A1) $E(\mathbf{p})$ has the form

$$E(\mathbf{p}) \;=\; \frac{1}{n}\sum_{i=1}^{n} G\left(\frac{p_i}{1/n}\right)\,.$$

with G continuous, $G(x) \geq 0$, and $G(0) = 0$.

(A2) Let $q = (q_1, \cdots, q_n) \in \mathbb{R}^n$ and $r = (r_1, \cdots, r_m) \in \mathbb{R}^m$ are two probability vectors let $p = q \otimes r$ be the product state given by the probability vector $p \in \mathbb{R}^{nm}$ with $p_{ij} = q_i r_j$, $1 \leq i \leq n$, $1 \leq j \leq m$. We have then

$$E(\mathbf{p}) = E(\mathbf{q}) + E(\mathbf{r}) .$$

(A3) For any states $\mathbf{p}_1$ and $\mathbf{p}_2$ and $0 \leq \theta \leq 1$ we have

$$E(\theta \mathbf{p}_1 + (1-\theta)\mathbf{p}_2) \leq \theta E(\mathbf{p}_1) + (1-\theta) E(\mathbf{p}_2) .$$

Note that (**A1**) is a mild assumption on the form of $E(p)$. The assumption that $G(0) = 0$ is simply a normalization since we can replace $G(x)$ by $G(x) + \alpha(x-1)$ without changing E. The axiom **(A2)** is crucial but very reasonable: if the state p is obtained by taking two *independent* decisions given by q and r then the information is the sum of the respective informations. The third axiom **(A3)** says that if two states are mixed then the information should decrease. This axiom ensures that the pure states maximize the entropy. But as the following theorem shows **(A3)** is used only to determine the overall sign of E!

Theorem 2.7 *If the function* $E(\mathbf{p})$ *satisfies* **(A1)** *and* **(A2)** *then* $G(x) = cx\ln(x)$ *and thus* $E(\mathbf{p}) = cH(\mathbf{p})$ *for some* $c \in \mathbb{R}$*. If* $E(\mathbf{p})$ *satisfies in addition* **(A3)** *then* $c \geq 0$.

Proof: Let $\mathbf{q} \in \mathbb{R}^{n+1}$ be a state with $q_j = \frac{1}{n}$ for $1 \leq j \leq n$ and $q_{n+1} = 0$. We consider the state $\mathbf{p} \in \mathbb{R}^{m(n+1)}$ which is the product of m independent copies of $\mathbf{q}$. Among the $(n+1)^m$ components of $\mathbf{p}$, n^m are equal to $\frac{1}{n^m}$ and the remaining $(n+1)^m - n^m$ are 0. Using the axioms **(A1)**

$$E(\mathbf{p}) \;=\; \frac{1}{(n+1)^m} n^m G\left(\frac{1/n^m}{1/(n+1)^m}\right) \;=\; (1+1/n)^{-m} G\left((1+1/n)^m\right) . \tag{2.7}$$

We take the limits $n \to \infty$ and $m \to \infty$ with $m/n \to \alpha$, then (2.7) tends to the limit $e^{-\alpha} G(e^{\alpha})$. Using (**A2**) we have

$$E(\mathbf{p}) \;=\; mE(\mathbf{q}) \;=\; m\frac{n}{n+1} G\left(1 + \frac{1}{n}\right) . \tag{2.8}$$

For (2.8) to have limit we need the limit $\lim_{n\to\infty} nG(1+1/n) \equiv c$ to exist.

$$G(e^\alpha) \;=\; c\alpha e^\alpha \tag{2.9}$$

and so we have proved that $G(x) = cx\ln(x)$ for $x \geq 1$.

To prove the formula for $x \leq 1$ we take $k \leq n$ and let $\mathbf{q} \in \mathbb{R}^{n+1}$ be the state with with $q_j = \frac{k}{n(n+1)}$ for $j = 1, \cdots, n$ and $q_{n+1} = 1 - \frac{k}{n+1}$. Let $\mathbf{r} \in \mathbb{R}^n$ be the state $r_j = 0$ with $1 \leq j \leq n-1$ and $r_n = 1$. Let $\mathbf{p} = \mathbf{q} \otimes \mathbf{r}$ be the product state, then by (**A2**)

$$\begin{aligned} E(\mathbf{p}) \;&=\; \frac{1}{n(n+1)}\left[nG(k) + G(n(n+1-k))\right] \\ = \; E(\mathbf{q}) + E(\mathbf{r}) \;&=\; \frac{1}{n+1}\left[nG(k/n) + G(n+1-k)\right] + \frac{1}{n}G(n) \end{aligned} \tag{2.10}$$

Using that $G(x) = cx\ln(x)$ for $x \geq 1$ one obtains that

$$G\left(\frac{k}{n}\right) \;=\; c\frac{k}{n}\ln\left(\frac{k}{n}\right)\,. \tag{2.11}$$

By the continuity of g we have $G(x) = x\ln(x)$ for all $x \geq 0$.

Finally we use **(A3)** to get $c \geq 0$. ■

Imagine that an agent has no payoff function at hand, then it is a fairly rational decision to choose a strategy at random, i.e., to follow the mixed strategy $\mathbf{p}_u$. As we have argued above it is reasonable to assume that the agent is actually minimizing the entropy function. A way to model an agent prone to some error and uncertainty is make him maximize a combination of his payoff function $\boldsymbol{\pi}$ and the entropy functional $H(\mathbf{p})$. Let β be a nonnegative number which measures, for example, the level of trust of an agent in his payoff function. The choice $\beta = 0$ corresponds to not believing in one's payoff function at all while $\beta = \infty$ corresponds to complete trust in one's payoff. This leads us to the maximization problem

$$\sup_{\mathbf{p}\in\Delta}\left[\beta\langle\boldsymbol{\pi},\mathbf{p}\rangle - H(\mathbf{p})\right] \tag{2.12}$$

and let us denote by $BR_\beta(\boldsymbol{\pi}, \mathbf{p}_u)$ the set of maximizers. Since the function $\beta\langle\boldsymbol{\pi},\mathbf{p}\rangle - H(\mathbf{p})$ is a continuous function and Δ is compact the set $BR_\beta(\boldsymbol{\pi}, \mathbf{p}_u)$ is not empty.

In the definition of the entropy $H(\mathbf{p}) = H(\mathbf{p}|\mathbf{p}_u)$ as a measure of order we have implicitly assumed that the most disordered state is the uniform state $\mathbf{p}_u$. There are situations where other choices of reference states $\mathbf{q}$ are possible. For example some partial knowledge of the problem or some inherent bias might make us prefer such or such strategy. For this purpose it is natural to consider the more general maximization problem problem

$$\sup_{\mathbf{p}\in\Delta} \left[\beta\langle\boldsymbol{\pi}, \mathbf{p}\rangle - H(\mathbf{p}\,|\,\mathbf{q})\right] . \tag{2.13}$$

We denote by $BR_\beta(\boldsymbol{\pi}, \mathbf{q})$ the set of maximizer. The function $\beta\langle\boldsymbol{\pi}, \mathbf{p}\rangle - H(\mathbf{p}\,|\,\mathbf{q})$ is upper-semicontinuous, it is equal to $-\infty$ if $\mathbf{p} \not\ll \mathbf{q}$ and is continuous on the face of Δ given by $\{\mathbf{q}\,,\,\mathbf{q} \ll \mathbf{p}\}$.

It turns out that the set $BR_\beta(\boldsymbol{\pi}, \mathbf{q})$ always contain a unique element and it is easy to compute it explicitly.

Theorem 2.8 *Let $\mathbf{p}(\beta, \mathbf{q})$ be the state given by*

$$p_i(\beta, \mathbf{q}) \;=\; \begin{cases} Z(\beta, q)^{-1} q_i e^{\beta\pi_i} & i \in \Sigma_{\mathbf{q}} \\ 0 & \text{otherwise} \end{cases} \tag{2.14}$$

where

$$Z(\beta, \mathbf{q}) \;=\; \sum_{k\in\Sigma_q} q_k e^{\beta\pi_k} \tag{2.15}$$

a the normalization constant. Let

$$F(\beta, \mathbf{q}) \;:=\; \ln Z(\beta, \mathbf{q})\,. \tag{2.16}$$

Then

$$F(\beta, \mathbf{q}) \;=\; \sup_{\mathbf{p}\in\Delta} \left[\beta\langle\boldsymbol{\pi}, \mathbf{p}\rangle - H(\mathbf{p}\,|\,\mathbf{q})\right] . \tag{2.17}$$

and the maximizer in (2.17) is unique and is given by $p(\beta, \mathbf{q})$

Proof: We can assume that $\mathbf{p} \ll \mathbf{q}$ and note that $\mathbf{p} \ll \mathbf{q}$ if and only if $\mathbf{p} \ll \mathbf{p}(\beta, \mathbf{q})$. We have

$$\begin{aligned} H(\mathbf{p}\,|\,\mathbf{p}(\beta, \mathbf{q})) \;&=\; \sum_{i\in\Sigma_{p(\beta,\mathbf{q})}} p_i \ln\left(\frac{p_i}{p_i(\beta, \mathbf{q})}\right) \\ &=\; \sum_{i\in\Sigma_{\mathbf{q}}} p_i \left[\ln\left(\frac{p_i}{q_i}\right) - \beta\pi_i + \ln\left(\sum_{k\in\Sigma_{\mathbf{q}}} q_k e^{\beta\pi_k}\right)\right] \\ &=\; H(\mathbf{p}\,|\,\mathbf{q}) - \beta\langle\pi, p\rangle + F(\beta, q) \end{aligned} \tag{2.18}$$

By theorem 2.5 the relative entropy is nonnegative and thus

$$F(\beta, \mathbf{q}) \geq \sup_{p \in \Delta} [\beta \langle \boldsymbol{\pi}, \mathbf{p} \rangle - H(\mathbf{p} \,|\, \mathbf{q})] \,. \tag{2.19}$$

Since $H(\mathbf{p}|\mathbf{p}(\beta, \mathbf{q})) = 0$ iff $\mathbf{p} = \mathbf{p}(\beta, \mathbf{q})$ this concludes the proof of Theorem 2.8. ■

Remark: In physics, the probability distribution $\mathbf{p}(\beta, \mathbf{q})$ is called a *Gibbs distribution* and $\beta = T^{-1}$ is the inverse of the temperature T. Low temperature T means very small noise or large β. Conversely high temperature T means large noise or small β.

More precisely let us consider in details the case $q = q_u$ in which case we have

$$p_i(\beta) = \frac{e^{\beta \pi_i}}{\sum_k e^{\beta \pi_k}} \,, .$$

As $\beta \to 0$ we have

$$\lim_{\beta \to 0} p_i(\beta) = \lim_{\beta \to 0} \frac{e^{\beta \pi_i}}{\sum_k e^{\beta \pi_k}} = \frac{1}{n}$$

For $\beta \to \infty$ let us suppose that we have a unique best reponse i^*, i.e., π_i is maximized for a unique i^*. Then we have

$$\lim_{\beta \to 0} p_j(\beta) = \lim_{\beta \to \infty} \frac{1}{\sum_k e^{\beta(\pi_k - \pi_j)}} = \left\{ \begin{array}{ll} 1 & j = i^* \\ 0 & j \neq i^* \end{array} \right. ,$$

i.e., $\mathbf{p}(\beta)$ converges to the best response pure strategy. More generally if there are L pure best responses then

$$\lim_{\beta \to 0} p_j(\beta) = \left\{ \begin{array}{ll} \frac{1}{L} & i \in BR(\boldsymbol{\pi}) \\ 0 & \text{otherwise} \end{array} \right. ,$$

i.e., $\mathbf{p}(\beta)$ converges to a mixed strategy for which all best response pure strategies occur with equal probability.

Remark: The variational principle is a Legendre transform. Let $f : \mathbb{R}^n \to \mathbb{R} \cup \{+\infty\}$ be a lower-semicontinuous function, then the Legendre transform of f, $f^*(\lambda)$ is defined as

$$f^*(\lambda) = \sup_{x \in \mathbb{R}^n} [\langle \lambda, x \rangle - f(x)] \tag{2.20}$$

Then $f^*(\lambda)$ is a lower-semicontinuous convex function. Moreover if f is itself convex then one has the duality formula

$$f(x) \;=\; \sup_{\lambda \in \mathbb{R}^n} \left[\langle \lambda, x \rangle - f^*(\lambda) \right] \tag{2.21}$$

In our case we can extend the function $H(\mathbf{p}\,|\,\mathbf{q})$ to $\mathbb{R}^n$ by setting it equal to $+\infty$ if $p \notin \Delta$. Then we have the formula

$$H(\mathbf{p}\,|\,\mathbf{q}) \;=\; \sup_{\boldsymbol{\pi} \in \mathbb{R}^n} \left[\langle \boldsymbol{\pi}, \mathbf{p} \rangle - F(1, \mathbf{q}) \right] . \tag{2.22}$$

3 Nash equilibria of normal form games

3.1 Nash equilibrium

We consider in this section *finite games in normal form* (also called *strategic form)*. Such a game is formally defined by the following data

- A finite set
$$\Gamma = \{\gamma_1, \cdots, \gamma_N\}$$
containing N elements. We call an element $\gamma \in \Gamma$ a *player* of type γ.

- For each $\gamma \in \Gamma$ a finite set S_γ is given. It is the set of pure strategies available to available to a player of type γ. We denote by $n_\gamma = \mathrm{card} S_\gamma$ the number of strategies available to player of type γ. The set
$$S \;=\; \times_{\gamma \in \Gamma} S_\gamma \,,$$
is the set of *strategy profiles* and an element
$$\mathbf{s} = (s_\gamma)_{\gamma \in \Gamma}$$
assigns a pure strategy to each player. There are $\prod_{\gamma \in \Gamma} n_\gamma$ different strategy profiles.

- For each $\gamma \in \Gamma$ a payoff function
$$\boldsymbol{\pi}_\gamma \;:\; S \to \mathbb{R} \tag{3.1}$$
is given which assigns a payoff to player γ given a strategy profile $\mathbf{s}$. We write
$$\pi_\gamma(i_1, \cdots, i_N)$$
for the payoff for player γ if the player γ_1 has strategy s_{i_1,γ_1}, the player γ_2 has strategy s_{i_2,γ_2}, and son on...

We assume throughout that the game is *of complete information* which means that all strategies and payoffs for every player are known to all players. We also assume that the game is *non-cooperative* which means that a player γ does not know what the other players are up to and has no way to communicate with, cooperate with, or manipulate the other players. The players will be able to choose pure or mixed strategies and the non-cooperative assumption means that each player will choose his strategy *independently* of the other players.

We denote by $\mathbf{p}_\gamma$ a (possibly mixed) strategy for the player γ, i.e., $\mathbf{p}_\gamma$ is an element of the simplex $\Delta(S_\gamma) \subset \mathbb{R}^{n_\gamma}$. The set of *strategy profiles* is given by the product space

$$\Delta \,=\, \times_{\gamma\in\Gamma}\Delta(S_\gamma) \tag{3.2}$$

We can view Δ as a subset of $\mathbb{R}^n$ where $n = \sum_\gamma n_\gamma$ and it is a product of unit simplices. We will denote the element of Δ by $\mathbf{p}$ and called them *mixed strategy profiles*. We write

$$\mathbf{p} = (\mathbf{p}_{\gamma_1}, \cdots, \mathbf{p}_{\gamma_N}),$$

with

$$\mathbf{p}_\gamma = (p_{1,\gamma}, \cdots, p_{n_\gamma,\gamma}) \qquad \sum_{i=1}^{n_\gamma} p_{i,\gamma} \,=\, 1$$

and we associate represent a pure strategy $s_{j,\gamma}$ for player γ with standard basis element $\mathbf{e}_{j\gamma} \in \mathbb{R}^{n_\gamma}$.

Given a payoff $\boldsymbol{\pi}_\gamma(\mathbf{s})$ for the player γ given a profile a pure strategies we extend the payoff to mixed strategies by applying the rules of probability. If $\mathbf{p}$ is the strategy profiles of the players then $p_{i_1,\gamma_1} \times \cdots \times p_{i_N,\gamma_N}$ is the probability that they choose the strategy profiles $\mathbf{s} = (s_{i_1,\gamma_1}, \cdots, s_{i_N,\gamma_N})$ and thus the (expected) payoff for the profile $\mathbf{p}$ is

$$E_{\mathbf{p}}[\boldsymbol{\pi}] \,\equiv\, \boldsymbol{\pi}_\gamma(\mathbf{p}) \,=\, \sum_{i_1=1}^{n_{\gamma_1}} \cdots \sum_{i_N=1}^{n_{\gamma_N}} \pi_\gamma(i_1, \cdots, i_N) p_{i_1,\gamma_1} \cdots p_{i_N,\gamma_N} \tag{3.3}$$

By our identification of pure strategies $s_{i,\gamma}$ with basis elements $\mathbf{e}_{i,\gamma}$ we have the identity

$$\pi_\gamma(i_1, \cdots, i_N) \,=\, \boldsymbol{\pi}_\gamma(\mathbf{e}_{i_1,\gamma_1}, \cdots, \mathbf{e}_{i_N,\gamma_N}).$$

It will be useful to introduce the notations $\mathbf{s}_{-\alpha}$ to denote the pure strategy profile for all players except α, i.e.,

$$\mathbf{s}_{-\alpha} \,=\, (s_\gamma)_{\substack{\gamma\in\Gamma \\ \gamma\neq\alpha}} \in \times_{\substack{\gamma\in\Gamma, \\ \gamma\neq\alpha}} S_\gamma$$

Similarly we denote by $\mathbf{p}_{-\alpha}$ the profile of mixed startegies for all player except player α, i.e.,

$$\mathbf{p}_{-\alpha} = (\mathbf{p}_\gamma)_{\substack{\gamma\in\Gamma,\\ \gamma\neq\alpha}} \in \times_{\substack{\gamma\in\Gamma,\\ \gamma\neq\alpha}} \Delta(S_\gamma) .$$

Using this we can write, with a slight abuse of notation, $\mathbf{p} = (\mathbf{p}_\alpha, \mathbf{p}_{-\alpha})$ and we use the notation $(\mathbf{q}_\alpha, \mathbf{p}_{-\gamma})$ for to suggest a strategy profile where player α changes his strategy from $\mathbf{p}_\alpha$ to $\mathbf{q}_\alpha$ while all other players keep their strategies unchanged to $\mathbf{p}_\gamma$. For example we have $(\mathbf{e}_{i\alpha}, \mathbf{p}_{-\alpha})$ is the strategy profile where α has the pure strategy $s_{i,\gamma}$ while the other players have strategies $\mathbf{p}_{-\alpha}$. It is easy to see from (3.3) that we have the identity

$$\boldsymbol{\pi}_\alpha(\mathbf{p}) = \sum_{j=1}^{n_\alpha} \boldsymbol{\pi}_\alpha(\mathbf{e}_{j\alpha}, \mathbf{p}_{-\alpha}) p_{j\alpha} \tag{3.4}$$

In other words the payoff for player α can be computed as the weighted payoff of his own pure strategies (all other's player mixed strategies being fixed). It is useful to think

$$(\boldsymbol{\pi}_\alpha(\mathbf{e}_{1\alpha}, \mathbf{p}_{-\alpha}) , \cdots , \boldsymbol{\pi}_\alpha(\mathbf{e}_{n_\alpha\alpha}, \mathbf{p}_{-\alpha}))$$

has the payoff vector for α given the strategies $\mathbf{p}_{-\alpha}$ for all other players.

Note that the function $\boldsymbol{\pi}$ given in equation (3.3) can be extended to all of $\mathbb{R}^{n_{\gamma_1}} \times \cdots \times \mathbb{R}^{n_{\gamma_N}}$ and is a *multilinear function*, i.e., linear in each $\mathbf{p}_\gamma$.

If there are $N = 2$ players called α and β with respectively n and m strategies then we can use a matrix notation to express our payoffs. We can think of $\boldsymbol{\pi}_\alpha$ and $\boldsymbol{\pi}_\beta$ as $n \times m$ matrices with entries $(\pi_\alpha(i,j))$ and $(\pi_\beta(i,j))$. We write then

$$\boldsymbol{\pi}_\alpha \mathbf{p}_\beta = \begin{pmatrix} \pi_\alpha(1,1) & \cdots \pi_\alpha(1,n) \\ \vdots & \vdots \\ \pi_\alpha(m,1) & \cdots \pi_\alpha(m,n) \end{pmatrix} \begin{pmatrix} p_{1,\beta} \\ \vdots \\ p_{n,\beta} \end{pmatrix}$$

and this the payoff vector for player α given strategy $\mathbf{p}_\beta$ for player β. Furthermore

$$\langle \mathbf{p}_\alpha \, , \, \boldsymbol{\pi}_\alpha \mathbf{p}_\beta \rangle$$

is the payoff for α given the strategy profile $\mathbf{p} = (\mathbf{p}_\alpha, \mathbf{p}_\beta)$. Similarly, if A^T denotes the transpose of the matrix A, $\boldsymbol{\pi}_\beta^T \mathbf{p}_\alpha$ is the payoff vector for β if α's strategy is p_α. Therefore

$$\langle \boldsymbol{\pi}_\beta^T \mathbf{p}_\alpha \, , \, \mathbf{p}_\beta \rangle = \langle \mathbf{p}_\alpha \, , \, \boldsymbol{\pi}_\beta \mathbf{p}_\beta \rangle$$

is the payoff for α given the strategy profile $\mathbf{p} = (\mathbf{p}_\alpha, \mathbf{p}_\beta)$.

Example 3.1 In the example 1.1 we have $N = 2$ players and the payoff matrices are

$$\boldsymbol{\pi}_\alpha = \begin{pmatrix} 40 & 24 \\ 48 & 28 \end{pmatrix} \qquad \boldsymbol{\pi}_\beta = \begin{pmatrix} 40 & 48 \\ 24 & 28 \end{pmatrix}$$

If $\mathbf{p}_{SA} = (p_1, p_2)$ and $\mathbf{p}_{IR} = (q_1, q_2)$ we have the payoffs

$$\begin{aligned} \boldsymbol{\pi}_{SA}(\mathbf{p}) &= p_1(40q_1 + 24q_2) + p_2(48q_1 + 28q_2)\,, \\ \boldsymbol{\pi}_{IR}(\mathbf{p}) &= q_1(40p_1 + 24p_2) + q_2(48p_1 + 28p_2)\,. \end{aligned} \tag{3.5}$$

Imagine that each of the N player is trying to maximize his own payoff $\boldsymbol{\pi}_\gamma$ with an optimal strategy $\mathbf{p}_\gamma$. In a non-trivial game the payoff $\boldsymbol{\pi}_\gamma$ will of course depend on the choice of strategies of other players. A *Nash equilibrium* is a strategy profile $\mathbf{p}$ in which no player can improve his payoff by changing his strategy if the other players leave their own strategy unchanged. Formally

Definition 3.2 *A* Nash equilibrium *(NE) for the game*

$$(\Gamma\,, S\,, \{\boldsymbol{\pi}_\gamma\}_{\gamma\in\Gamma})$$

is a strategy profile $\mathbf{p}^* \in \Delta = \times_{\gamma\in\Gamma}\Delta(S_\gamma)$ *such that, for every* γ, $\mathbf{p}^*_\gamma$ *is a best response for the player* γ *given the startegy profile* $\mathbf{p}_{-\gamma}$ *of the other players. That is,*

$$\boldsymbol{\pi}_\gamma(\mathbf{p}^*) = \max_{\mathbf{q}_\gamma} \boldsymbol{\pi}_\gamma(\mathbf{q}_\gamma, \mathbf{p}^*_{-\gamma})\,, \qquad \text{for all } \gamma \in \Gamma\,.$$

Consider for example a $N = 2$ player games with payoff matrices $\boldsymbol{\pi}_\alpha$ and $\boldsymbol{\pi}_\beta$. The strategies $\mathbf{p}_\alpha$ and $\mathbf{p}_\beta$ are a Nash equilibrium if and only if

$$\mathbf{p}_\alpha \text{ is a best response for the payoff } \boldsymbol{\pi}_\alpha\mathbf{p}_\beta$$

and

$$\mathbf{p}_\beta \text{ is a best response for the payoff } \boldsymbol{\pi}_\beta^T\mathbf{p}_\alpha$$

Using (3.4) we can rephrase the equilibrium condition as

Lemma 3.3 *The strategy* $\mathbf{p}^*$ *is a NE if and only if*

$$\boldsymbol{\pi}_\alpha(\mathbf{e}_{j\alpha}, \mathbf{p}^*_{-\alpha}) \le \boldsymbol{\pi}_\alpha(\mathbf{p}^*)\,, \tag{3.6}$$

for $\alpha \in \Gamma$ *and for all* $j \in S_\alpha$.

Furthermore we have

$$\boldsymbol{\pi}_\alpha(\mathbf{e}_{j\alpha}, \mathbf{p}^*_{-\alpha}) < \boldsymbol{\pi}_\alpha(\mathbf{p}^*) \text{ iff } j \notin \Sigma_{\mathbf{p}^*_\alpha}$$

and

$$\boldsymbol{\pi}_\alpha(\mathbf{e}_{j\alpha}, \mathbf{p}^*_{-\alpha}) = \boldsymbol{\pi}_\alpha(\mathbf{p}^*) \text{ iff } j \in \Sigma_{\mathbf{p}^*_\alpha}$$

This means in Nash equilibrium, for every player α, every pure strategy in the support of $\mathbf{p}_\alpha$ is a best response.

Proof: Since

$$\boldsymbol{\pi}_\alpha(\mathbf{q}_\alpha, \mathbf{p}_{-\alpha}) \;=\; \sum_i q_{i\alpha} \boldsymbol{\pi}_\alpha(\mathbf{e}_{i\alpha}, \mathbf{p}_{-\alpha})$$

we have

$$\sup_{\mathbf{q}_\alpha} \boldsymbol{\pi}_\alpha(\mathbf{q}_\alpha, \mathbf{p}_{-\alpha}) \;=\; \max_i \boldsymbol{\pi}_\alpha(\mathbf{e}_{i\alpha}, \mathbf{p}_{-\alpha})$$

To conclude use Lemma 2.1. ∎

Remark 3.4 Lemma 3.3 is a very important tool to actually compute Nash equilibria. It shows that it is enough to consider the pure strategies to determine if one is in a NE or not. But it does not mean that Nash equilibria consists only pure strategies as the following example demonstrates.

Example 3.5 (Matching pennies game). Two children, holding a penny apiece, independently choose which side of their coin to show. Child 1 wins if both coins show the same side and child 2 wins otherwise. In bi-matrix form the payoffs are

$$\begin{array}{c|c|c} & H & T \\ \hline H & \begin{matrix} & -1 \\ 1 & \end{matrix} & \begin{matrix} & 1 \\ -1 & \end{matrix} \\ \hline T & \begin{matrix} & 1 \\ -1 & \end{matrix} & \begin{matrix} & -1 \\ 1 & \end{matrix} \end{array} \tag{3.7}$$

This game has no pure strategy Nash equilibrium. For example if child 2 chooses the pure strategy H, then the payoff for child 1 is $\boldsymbol{\pi}_{C1} = (1, -1)$ and so the best response strategy for child 1 is to choose the pure strategy H. However if child 1 choose strategy H, the payoff for child 2 is $\boldsymbol{\pi}_{C1} = (-1, 1)$

and thus the best response to for child 2 is the pure strategy T. So H is a NE for child 2. By similar arguments one sees easily that no pure strategy is a Nash equilibrium.

Let us assume next that child 2 chooses the mixed strategy $(q, 1-q)$ then the payoff for child 1 is $(2q-1, 1-2q)$. If $q > 1/2$ then child 1 best response is the pure strategy H but the best response for child 2 to H is T, so this is does not lead to a NE. Similarly $q > 1/2$ can be eliminated. This leaves us with $q = 1/2$ in which case any strategy for child 1 is a best response. By reversing the argument we conclude that the mixed strategies profiles $((1/2, 1/2), (1/2, 1/2))$ is the unique Nash equilibrium of the game.

Note furthermore that for Nash equilibria we let open the possibility that the payoff stays unchanged for some change of strategies, i.e., the inequality in (3.6) needs not be strict. This motivates the definition

Definition 3.6 *A* strict Nash equilibrium *for the game* Γ, S, $(\boldsymbol{\pi}_\gamma)_{\gamma\in\Gamma}$ *is a strategy* $\mathbf{p}^*$ *such that, for any player* γ, $\mathbf{p}^*_\gamma$ *is the* unique *best response for* γ *given the strategy profile* $\mathbf{p}_{-\gamma}$ *of the other players.*

As we can see from Lemma 3.3 the best response is unique only if and only if it is a pure strategy and so a strict Nash equilibrium always is a profile of pure strategies (note that the opposite need not to hold).

Example 3.7 Consider again example 3.1. Given any choice $\mathbf{p}_{IR} = (q_1, q_2)$ for player IR the payoff $\boldsymbol{\pi}_{SA}$ is a decreasing function of p_1 and so the best response is the pure strategy 2. Conversely for any $\mathbf{p}_{SA} = (p_1, p_2)$ the best response for IR is the pure startegy 2 again and so the unique (strict) Nash equilibrium consist of the pair of pure strategies $(0, 1), (0, 1))$.

It is not obvious, a-priori, that a strategy profile which achieves a Nash equilibrium, is actually always attainable for any game. But indeed we have

Theorem 3.8 *(Nash) A finite game has at least one Nash equilibrium.*

Proof: The proof is an application of Brouwer's fixed point which asserts that if $f : K \to K$ is a continuous maps on the compact convex set K then f has at least one fixed point.

To apply this theorem we consider the product of mixed strategies simplexes Δ and we construct a continuous map $f : \Delta \to \Delta$ whose fixed points are exactly the Nash equilibria. Let us define first the continuous maps $f_{j\gamma} : \Delta \to \mathbb{R}$ given by

$$f_{j\gamma}(\mathbf{p}) \;=\; \max\left(0, \boldsymbol{\pi}_\gamma(\mathbf{e}_{j\gamma}, \mathbf{p}_{-\gamma}) - \boldsymbol{\pi}_\gamma(\mathbf{p})\right) \tag{3.8}$$

i.e., the maps $f_{j\gamma}$ measure the potential increase in payoff obtained by player γ if he changes from his current strategy $\mathbf{p}_\gamma$ to the pure strategy $\mathbf{e}_{j\gamma}$. Let us now define a map $T : \Delta \to \Delta$ by $T(\mathbf{p}) = \mathbf{p}'$ where

$$\mathbf{p}'_\gamma \;=\; \frac{\mathbf{p}_\gamma + \sum_j f_{j\gamma}(\mathbf{p})\mathbf{e}_{j\gamma}}{1 + \sum_j f_{j\gamma}(\mathbf{p})} \;=\; \sum_j \frac{p_{j\gamma} + f_{j\gamma}}{1 + \sum_l f_{l\gamma}} \mathbf{e}_{j\gamma}\,. \tag{3.9}$$

The maps T increases the probability of the the pure states if they are more favorable and it does so continuously in $\mathbf{p}$.

Suppose that $\mathbf{p}$ is a fixed point for the map T, then for all γ and j, $f_{j\gamma}(\mathbf{p}) = 0$ and thus no payoff is increased by switching to a new (pure) strategy. By Lemma 3.3 this implies that $\mathbf{p}$ is a NE. Conversely if $\mathbf{P}$ is a NE then clearly $f_{j\gamma}(\mathbf{p}) = 0$ and thus $T(\mathbf{p}) = \mathbf{p}$. ∎

Remark 3.9 The hard part in this theorem is to prove Brouwer's fixed point theorem. We shall not do this here.

From the examples we have considered so far you might start to think that games have unique Nash equilibria but this far from being the case.

Example 3.10 *(Battle of the sexes).* Let us consider the following game. Lola likes movies but she like theater much more. Bob, on the contrary, likes movies (M) more than theater (T). They however both equally dislike to be apart. They both know they are having a date tonight but they can't remember where. The payoff bi-matrix can be taken to have the form

$$\begin{array}{c|cc} & M & T \\ \hline M & \begin{array}{|cc|}\hline & 1 \\ 2 & \\ \hline\end{array} & \begin{array}{|cc|}\hline & 0 \\ 0 & \\ \hline\end{array} \\ T & \begin{array}{|cc|}\hline & 0 \\ 0 & \\ \hline\end{array} & \begin{array}{|cc|}\hline & 3 \\ 1 & \\ \hline\end{array} \end{array} \tag{3.10}$$

This game is an example of *coordination game*, it is better for both players to coordinate and choose the same strategy.

Let us look first at the pure strategies, for example Bob and Lola both chooses movies. Then given Bob's choices Lola has no reason to switch because she would be worse off and it would decrease her payoff from 1 to 0 and similarly for Bob. Therefore the strategy $(1, 0), (1, 0)$ is a NE. By the same argument $(0, 1), (0, 1)$ is also a NE. So Bob and Lola reach an equilibrium if they choose the same pure strategies. But how are they supposed to choose it without communicating? This suggest that there may be another equilibrium of mixed strategies. Suppose Lola's strategy is $(q, (1-q))$, then Bob's payoff vector is $(2q, 1-q)$. If q is such that $2q = 1-q$, i.e. $q = 1/3$ then Bob's strategies will be equivalent. For any $q \neq 1/3$ Bob will prefer a pure strategy in which case Lola best response should be a pure best response too. Arguing similarly, if Bob's chooses $(p, 1-p)$ then for $p = 3/4$ all of Lola's strategies will be equivalent. This leads to exactly three Nash equilibria, two pure, one mixed.

$$((1,0),(1,0))\ , \quad ((0,1),(0,1))\ , \quad ((3/4,1/4),(1/3,2/3))$$

with payoffs

$$2 \text{ and } 1\,, \qquad 2/3 \text{ and } 3/4\,, \qquad 1 \text{ and } 3$$

Which one would you choose?

3.2 Dominated strategies and their elimination

A useful tool to find Nash equilibria is to eliminate strategies from the game, a-priori, because they will never be in the support of a Nash equilibrium strategy.

Definition 3.11 *(Weak domination) A strategy* $\mathbf{q}_\gamma$ *for player* γ weakly dominates *the strategy* $\mathbf{p}_\gamma$ *if*

$$\boldsymbol{\pi}_\gamma(\mathbf{q}_\gamma, \mathbf{p}_{-\gamma}) \geq \boldsymbol{\pi}_\gamma(\mathbf{p}_\gamma, \mathbf{p}_{-\gamma})$$

for every $\mathbf{p}_{-\gamma}$ *and*

$$\boldsymbol{\pi}_\gamma(\mathbf{q}_\gamma, \mathbf{p}_{-\gamma}) > \boldsymbol{\pi}_\gamma(\mathbf{p}_\gamma, \mathbf{p}_{-\gamma})$$

for at least one $\mathbf{p}_{-\gamma}$.

Definition 3.12 *(Strict domination) A strategy* $\mathbf{q}_\gamma$ *for player* γ strictly dominates *the strategy* $\mathbf{p}_\gamma$ *if*

$$\boldsymbol{\pi}_\gamma(\mathbf{q}_\gamma, \mathbf{p}_{-\gamma}) > \boldsymbol{\pi}_\gamma(\mathbf{p}_\gamma, \mathbf{p}_{-\gamma})$$

for every $\mathbf{p}_{-\gamma}$.

Example 3.13 Consider the Prisonner's dilemna example with payoff matrices

$$\boldsymbol{\pi}_\alpha = \begin{pmatrix} 40 & 24 \\ 48 & 28 \end{pmatrix} \quad \boldsymbol{\pi}_\beta = \begin{pmatrix} 40 & 48 \\ 24 & 28 \end{pmatrix}.$$

For α, the pure strategy $\mathbf{e}_{2,\alpha}$ strictly dominates $\mathbf{e}_{1,\alpha}$ since $\pi_\alpha(2,j) > \pi_\alpha(1,j)$ for all j. Similarly for β the pure strategy $\mathbf{e}_{2,\beta}$ strictly dominates $\mathbf{e}_{1,\beta}$ since $\pi_\beta(j,2) > \pi_\beta(j,1)$ for all j.

Example 3.14 Consider a two player games with players α (3 strategies) and β (2 strategies) and payoff matrix

$$\boldsymbol{\pi}_\alpha = \begin{pmatrix} 3 & 0 \\ 0 & 3 \\ 1 & 1 \end{pmatrix} \quad \boldsymbol{\pi}_\beta = \begin{pmatrix} 0 & 1 \\ -1 & 0 \\ 2 & -1 \end{pmatrix}$$

For α no pure strategy strategy is dominated by any other. However the pure strategy $\mathbf{e}_{3,\alpha}$ is strictly dominated by the mixed strategy $\mathbf{q}_\alpha = (\frac{1}{2}, \frac{1}{2}, 0)$: For any pure strategy $\mathbf{e}_{j,\beta}$ for β we have

$$\boldsymbol{\pi}(\mathbf{e}_{3,\alpha}, \mathbf{e}_{j,\beta}) = 1 < \boldsymbol{\pi}(\mathbf{q}_\alpha, \mathbf{e}_{j,\beta}) = 3/2$$

There are no dominated strategy for β.

A rational player should never use a strictly dominated strategy. Also an rational opponent to such player will also know that such strategy will not be used and so it make sense to simply eliminate the strictly dominated pure strategies from the game.

Indeed we have

Proposition 3.15 *Suppose* $e_{j,\gamma}$ *is a strictly dominated strategy for player* γ *in a game* $(\Gamma, S, (\boldsymbol{\pi}_\gamma)_{\gamma\in\Gamma})$*. Then for any Nash equilibrium* $\mathbf{p}^*$

$$j \notin \Sigma_{\mathbf{p}^*_\gamma}$$

Proof: This is an immediate consequence of the characterization of NE in Lemma 3.3. ∎

This allows to eliminate strictly dominated strategies in an iterative manner. Suppose $\mathbf{e}_{j,\alpha}$ is strictly dominated by some other (pure or mixed) strategy $\mathbf{p}_\alpha$ for some player α, Then consider a new game $(\Gamma, S', \boldsymbol{\pi}'_\gamma)$ where $S'_\gamma = S_\gamma$ if $\gamma \neq \alpha$ and

$$S'_\alpha = S_\alpha \setminus \{s_{j,\alpha}\}\,.$$

The new payoffs $\boldsymbol{\pi}'_\gamma$ are simply the old payoffs $\boldsymbol{\pi}_\gamma$ restricted to S'.

Example 3.16 For the example 3.14 we can eliminate $s_{3,\alpha}$ and obtain new payoff matrices

$$\boldsymbol{\pi}'_\alpha = \begin{pmatrix} 3 & 0 \\ 0 & 3 \end{pmatrix}, \quad \boldsymbol{\pi}'_\beta = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}$$

and the NE are easier to compute.

We can of course repeat this elimination procedure as many times as possible.

Example 3.17 Consider a two-player game, each player has three strategies and the payoff matrices are

$$\boldsymbol{\pi}_\alpha = \begin{pmatrix} 3 & 1 & 6 \\ 0 & 0 & 4 \\ -1 & 2 & 5 \end{pmatrix}, \quad \boldsymbol{\pi}_\beta = \begin{pmatrix} 3 & 0 & -1 \\ 1 & 0 & 2 \\ 6 & 4 & 5 \end{pmatrix}$$

The pure strategies $\mathbf{e}_{2,\alpha}$ and $\mathbf{e}_{2,\beta}$ are strictly dominated by $\mathbf{e}_{1,\alpha}$ and $\mathbf{e}_{1,\beta}$ respectively. The reduced game has payoff matrices

$$\boldsymbol{\pi}'_\alpha = \begin{pmatrix} 3 & 6 \\ -1 & 5 \end{pmatrix}, \quad \boldsymbol{\pi}'_\beta = \begin{pmatrix} 3 & -1 \\ 6 & 5 \end{pmatrix}$$

and for the reduced game the strategies $\mathbf{e}'_{3,\alpha}$ and $\mathbf{e}'_{3,\beta}$ are strictly dominated by $\mathbf{e}'_{3,\alpha}$ and $\mathbf{e}'_{1,\beta}$ respectively. So we can now reduce the game to a trivial game where α and β only have one strategy each with payoffs $\boldsymbol{\pi}_\alpha(1,1) = 3$ and $\boldsymbol{\pi}_\beta(1,1) = 3$. We conclude from this that the only Nash equilibrium for the original game is the pure Nash equilibrium $((1,0,0),(1,0,0))$.

We conclude this section by considering an example which shows that if we eliminate weakly dominated strategies the resulting games may depend upon the order with which the strategies are eliminated.

Example 3.18 Consider a 2-person game where Player α has three strategies and player β has 2 strategies. The payoff matrices are

$$\boldsymbol{\pi}_\alpha = \begin{pmatrix} 3 & 2 \\ 1 & 0 \\ 0 & 1 \end{pmatrix} \qquad \boldsymbol{\pi}_\beta = \begin{pmatrix} 2 & 2 \\ 1 & 0 \\ 0 & 1 \end{pmatrix}$$

For player α both pure strategies $\mathbf{e}_{2,\alpha}$ and $\mathbf{e}_{3,\alpha}$ are strictly dominated by $\mathbf{e}_{1,\alpha}$. Let us consider several

1. Let us eliminate both $\mathbf{e}_{2,\alpha}$ and $\mathbf{e}_{3,\alpha}$ for player α. Then the reduced game has payoff matrices

$$\boldsymbol{\pi}'_\alpha = \begin{pmatrix} 3 & 2 \end{pmatrix}, \quad \boldsymbol{\pi}'_\beta = \begin{pmatrix} 2 & 2 \end{pmatrix}$$

   The NE are $\mathbf{p}_\alpha = (1,0,0)$ and $\mathbf{q}_\beta = (q, 1-q)$ for any $0 \le q \le 1$.

2. Let us eliminate only $\mathbf{e}_{3,\alpha}$ and so the payoff matrices become

$$\boldsymbol{\pi}'_\alpha = \begin{pmatrix} 3 & 2 \\ 1 & 0 \end{pmatrix} \qquad \boldsymbol{\pi}'_\beta = \begin{pmatrix} 2 & 2 \\ 1 & 0 \end{pmatrix}$$

   For player β, the pure strategy $\mathbf{e}_{2,\beta}$ is weakly dominated dominated by $\mathbf{e}_{1,\beta}$ and if we eliminate $\mathbf{e}_{2,\beta}$ we obtain the payoff matrices

$$\boldsymbol{\pi}''_\alpha = \begin{pmatrix} 3 \\ 1 \end{pmatrix} \qquad \boldsymbol{\pi}''_\beta = \begin{pmatrix} 2 \\ 1 \end{pmatrix}$$

   This lead to a NE $(1,0,0), (1,0)$.

3. If we eliminate only $\mathbf{e}_{2,\alpha}$ we obtain

$$\boldsymbol{\pi}'_\alpha = \begin{pmatrix} 3 & 2 \\ 0 & 1 \end{pmatrix} \qquad \boldsymbol{\pi}'_\beta = \begin{pmatrix} 2 & 2 \\ 0 & 1 \end{pmatrix}$$

   in which case for player β, the pure strategy $\mathbf{e}_{1,\beta}$ is weakly dominated dominated by $\mathbf{e}_{2,\beta}$. If we eliminate $\mathbf{e}_{1,\beta}$ we have the new payoff matrices

$$\boldsymbol{\pi}''_\alpha = \begin{pmatrix} 2 \\ 1 \end{pmatrix} \qquad \boldsymbol{\pi}''_\beta = \begin{pmatrix} 2 \\ 1 \end{pmatrix}$$

   and this lead to a NE $(1,0,0)$ and $(1,0,0)$.

The previous example shows that eliminating weakly dominated strategies can eliminate some NE from the game. It can be shown (see exercises) that the NE equilibria for the reduced game are always NE from the original game.

3.3 An algorithm to compute Nash equilibria

In this section we give a fairly general algorithm to compute the NE for a general N-persons game. The algorithm is general and it also shows that computing the NE for 2-players games reduce to solve a number of *linear* equations. it will also show that computing the NE for even a small number (let us say 4) strategies can become quickly quite tedious. For games with N player solving the NE is equivalent to solve a system of polynomial equations in degree $N-1$ which is very difficult except in special case.

The algorithm relies on the following fact that we have already noted and used before

Fact: A sufficient condition for $\mathbf{p}_\gamma$ to be a NE strategy for player γ is that all the payoffs

$$\boldsymbol{\pi}_\gamma(\mathbf{e}_{j,\gamma}, \mathbf{p}_{-\gamma})$$

for all the pure strategies $\mathbf{e}_{j,\gamma}$ with j in the support of $\mathbf{p}_\gamma$ are equal to each other.

This leads to the following

Algorithm 3.19 (Algorithm to compute the NE)

1. *Select for each player γ a support $\Sigma_{\mathbf{p}_\gamma}$ of his equilibrium strategy $\mathbf{p}_\gamma$. This equivalent to choosing a subset $\tilde{S}_\gamma \subset S_\gamma$ for each γ.*

2. *For each player γ write down the set of equations*

$$u_\gamma = \boldsymbol{\pi}_\gamma(\mathbf{e}_{j,\gamma}, \mathbf{p}_{-\gamma}), \quad j \in \Sigma_{\mathbf{p}_\gamma},$$

with the constraints

$$p_{j,\gamma} = 0, \quad \text{if } j \notin \Sigma_{\mathbf{p}_\gamma},$$

and

$$\sum_{j \in \Sigma_{\mathbf{p}_\gamma}} p_{j,\gamma} = 1.$$

3. *Solve the set of equations in Step 2, if possible. It might happen that no solution exists or that a solution exists but some of $p_{j,\gamma}$ are negative. If a solution exists with positive $p_{j,\gamma}$ then it is a* candidate *for a NE but it might happen that the payoff for some pure strategy $\mathbf{e}_{l,\gamma}$ with $l \notin \Sigma_{p_\gamma}$ is bigger than the payoff for playing p_γ. This happen for example if you have excluded from the support a pure strategy who strictly dominate some other strategy. Therefore one needs to check that, for every γ,*

$$p_{j,\gamma} \geq 0\,, \quad j \in \Sigma_{\mathbf{p}_\gamma}\,,$$

and

$$\boldsymbol{\pi}_\gamma(\mathbf{p}_\gamma, \mathbf{p}_{-\gamma}) \geq \boldsymbol{\pi}_\gamma(\mathbf{e}_{l,\gamma}, \mathbf{p}_{-\gamma})\,, \quad l \notin \Sigma_{\mathbf{p}_\gamma}\,.$$

4. *Go back to Step 1 and choose another support for the equilibrium strategy.*

Let us illustrate our algorithm with

Example 3.20 Consider a 2-players game with payoff bi-matrix

	t_1	t_2	t_3
s_1	7, 2	2, 7	3, 6
s_2	2, 7	7, 2	4, 5

(3.11)

Case 1: Let us assume that the NE are pure strategies. If α chooses s_1 then β's best response is t_2. But α's best response to t_2 is s_2. If α chooses s_2, then β's best response to t_1 and α's best response to t_1 is s_1. This rules out equilibrium pure strategies for α. By arguing similarly we can exclude pure strategies for β (details left to the reader).
Case 2: Let us assume $\Sigma_{\mathbf{p}_\alpha} = \{1,2\}$ and $\mathbf{p}_\alpha = (p_1, p_2)$ and $\Sigma_{\mathbf{p}_\beta} = \{1,2,3\}$ and $\mathbf{p}_\beta = (q_1, q_2, q_3)$. Then we have the equations

$$\begin{aligned} u_\alpha &= 7q_1 + 2q_2 + 3q_3 \\ &= 2q_1 + 7q_2 + 4q_3 \end{aligned} \tag{3.12}$$

as well as

$$
\begin{aligned}
u_\beta &= 2p_1 + 7p_2 \\
&= 7p_1 + 2p_2 \\
&= 6p_1 + 5p_2 \qquad (3.13)
\end{aligned}
$$

The first two equations in (3.13) implies $p = 1/2$ which is incompatible with the the third equation in (3.13). This rules out NE.
Case 3: Let us assume $\Sigma_{\mathbf{p}_\alpha} = \{1,2\}$ and $\mathbf{p}_\alpha = (p_1, p_2)$ and $\Sigma_{\mathbf{p}_\beta} = \{2,3\}$ and $\mathbf{p}_\beta = (0, q_2, q_3)$. Then we must have

$$
\begin{aligned}
u_\alpha &= 2q_2 + 3q_3 \\
&= 7q_2 + 4q_3 \qquad (3.14)
\end{aligned}
$$

and

$$
\begin{aligned}
u_\beta &= 7p_1 + 2p_2 \\
&= 6p_1 + 5p_2 \qquad (3.15)
\end{aligned}
$$

Equation (3.14) yields $5q_2 = -q_3$ which is not a probability. This rules out NE in this case.
Case 4: Let us assume $\Sigma_{\mathbf{p}_\alpha} = \{1,2\}$ and $\mathbf{p}_\alpha = (p_1, p_2)$ and $\Sigma_{\mathbf{p}_\beta} = \{1,2\}$ and $\mathbf{p}_\beta = (q_1, q_2, 0)$. Then we must have

$$
\begin{aligned}
u_\alpha &= 7q_1 + 2q_2 \\
&= 2q_1 + 7q_2 \qquad (3.16)
\end{aligned}
$$

and

$$
\begin{aligned}
u_\beta &= 7p_1 + 2p_2 \\
&= 2p_1 + 7p_2 \qquad (3.17)
\end{aligned}
$$

These equation have solutions $\mathbf{p}_\alpha = (1/2, 1/2)$ and $\mathbf{p}_\beta = (1/2, 1/2, 0)$. To make sure that we have a NE we need to compare $\mathbf{p}_\beta$ with the pure strategy $\mathbf{e}_{3,\beta}$. Since we have

$$
\boldsymbol{\pi}_\beta(\mathbf{p}_\alpha, \mathbf{e}_{3,\beta}) = \frac{1}{2} \times 6 + \frac{1}{2} \times 5 = 5.5 > \boldsymbol{\pi}_\beta(\mathbf{p}_\alpha, \mathbf{q}_\beta) = u_\beta = 7 \times \frac{1}{2} + 2\frac{1}{2} = 4.5\,,
$$

and so this is not a NE.

Case 5: The last case is $\Sigma_{\mathbf{p}_\alpha} = \{1,2\}$ and $\mathbf{p}_\alpha = (p_1, p_2)$ and $\Sigma_{\mathbf{p}_\beta} = \{1,3\}$ and $\mathbf{p}_\beta = (q_1, 0, q_3)$ with payoff equations

$$\begin{aligned} u_\alpha &= 7q_1 + 3q_3 \\ &= 2q_1 + 4q_3 \end{aligned} \tag{3.18}$$

and

$$\begin{aligned} u_\beta &= 2p_1 + 7p_2 \\ &= 6p_1 + 5p_2 \end{aligned} \tag{3.19}$$

with solutions and $\mathbf{p}_\alpha = (1/3, 2/3)$ and $\mathbf{p}_\beta = (1/5, 0, 5/6)$ with payoffs $u_\alpha = 11/3$ and $u_\beta = 16/3$. As a final check one verifies that

$$\boldsymbol{\pi}_\beta(\mathbf{p}_\alpha, \mathbf{e}_{2,\beta}) = 7 \times \frac{1}{3} + 2 \times \frac{2}{3} = \frac{11}{3} < \frac{16}{3}\,,$$

and so

$$\mathbf{p}_\alpha = (1/3, 2/3)\,, \quad \mathbf{p}_\beta = (1/5, 0, 5/6)\,,$$

is the unique NE.

3.4 Invariance of Nash equilibria under payoff transformation

In this section we discuss another fact which is very useful to classify the games and very often can be used to simplify the computation of the equilibria. The basic idea is that given a payoff vector $\boldsymbol{\pi} = (\pi(1), \cdots, \pi(n))$ then the new payoff

$$\boldsymbol{\pi}' = a\boldsymbol{\pi} + b\,, \quad a > 0, b \in \mathbb{R}$$

is also a payoff which represents the same preferences. The choice of a and b should be thought as a choice of physical units to measure the payoff. The constant a means that we can measure a payoff in US dollars or in Euros and the constant b means that we can choose, say, an arbitrary 0 value for the payoffs.

We consider the following payoff transformations

(a) Global linear payoff rescaling. For player γ choose constant $a_\gamma > 0$ and $b_\gamma \in \mathbb{R}$. Set

$$\boldsymbol{\pi}'_\gamma = a_\gamma \boldsymbol{\pi}_\gamma + b_\gamma\,.$$

(b) Strategy dependent linear payoff rescaling. For player γ let $\tilde{s}_{-\gamma}$ be a *fixed* strategy profile of the other players and choose a constant $d = d(\tilde{s}_{-\gamma})$. Set

$$\boldsymbol{\pi}'_\gamma(s) \;=\; \begin{cases} \boldsymbol{\pi}_\gamma(s) + d(\tilde{s}_{-\gamma}) & \text{if} \quad s = (s_\gamma, \tilde{s}_{-\gamma}) \\ \boldsymbol{\pi}_\gamma(s) & \text{if} \quad s = (s_\gamma, s_{-\gamma})\,, \quad s_{-\gamma} \neq \tilde{s}_{-\gamma} \end{cases}$$

Definition 3.21 *(Nash equivalence)*

- *Two payoffs* $\boldsymbol{\pi}_{gamma}$ *and* $\boldsymbol{\pi}'_\gamma$ *are* Nash equivalent *if there exists a collection of linear rescaling (either global of payoff dependent) which transform* $\boldsymbol{\pi}_\gamma$ *into* $\boldsymbol{\pi}_{\gamma'}$*. We write*

  $$\boldsymbol{\pi}_\gamma \overset{NE}{\sim} \boldsymbol{\pi}_{\gamma'}$$

  if $\boldsymbol{\pi}_\gamma$ *and* $\boldsymbol{\pi}_{\gamma'}$ *are Nash equivalent.*

- *Two games* $(\Gamma, S, \boldsymbol{\pi}_\gamma)$ *and* $(\Gamma, S, \boldsymbol{\pi}'_\gamma)$ *are* Nash equivalent *if* $\boldsymbol{\pi}_\gamma \overset{NE}{\sim} \boldsymbol{\pi}'_\gamma$ *for all* $\gamma \in \Gamma$*.*

In the case of 2 players this means that we can add a constant to each row of the payoff matrix $\boldsymbol{\pi}_\alpha$ and add a constant to each column of the payoff matrix $\boldsymbol{\pi}_\beta$. If we have

$$\pi'_\alpha(i,j) \;=\; a\pi_\alpha(i,j) + b(j)\,, \quad \pi'_\beta(i,j) \;=\; c\pi_\alpha(i,j) + d(i)$$

for $a > 0$ and $c > 0$ and arbitrary $b(j)$ and $d(i)$ then the two games are Nash equivalent.

The motivation for the definition is

Proposition 3.22 *Nash equivalent games have the same Nash equilibria.*

Proof: If $\boldsymbol{\pi}_\gamma$ and $\boldsymbol{\pi}'_\gamma$ are related by linear rescalings then for each γ there is a constant a_γ such that

$$\boldsymbol{\pi}_\gamma(\mathbf{p}_\gamma, \mathbf{p}_{-\gamma}) - \boldsymbol{\pi}_\gamma(\mathbf{q}_\gamma, \mathbf{p}_{-\gamma}) \;=\; a\left(\boldsymbol{\pi}'_\gamma(\mathbf{p}_\gamma, \mathbf{p}_{-\gamma}) - \boldsymbol{\pi}'_\gamma(\mathbf{q}_\gamma, \mathbf{p}_{-\gamma})\right)$$

This implies that the NE are the same for both games.

3.5 Symmetric games

In this section we discuss a special class of games called *symmetric games*, in which one should think that all players γ have the same strategies and same payoffs. In this case one may think of γ as merely a label (say to distinguish between different teams). If we consider two players $\gamma_k, \gamma_l \in \Gamma$ and fix any strategy profile for the other players then the payoff for γ_k if he plays a pure strategy $\mathbf{e}_{i,\gamma_k}$ against $\mathbf{e}_{j,\gamma_l}$ should be the same as the payoff for γ_l if he plays $\mathbf{e}_{i,\gamma_l}$ against $\mathbf{e}_{j,\gamma_k}$.

Definition 3.23 *A game $(\Gamma, S, (\boldsymbol{\pi}_\gamma)_{\gamma\in\Gamma})$ is called a* symmetric game *if S_γ is independent of γ and the payoff $\boldsymbol{\pi}_\gamma$ have the the following symmetry: for any players γ_l, γ_k and any i_l, i_k*

$$\begin{aligned} & \pi_{\gamma_l}(i_1, \cdots, i_l, \cdots, i_k, \cdots, i_N) \\ = \quad & \pi_{\gamma_k}(i_1, \cdots, i_k, \cdots, i_l, \cdots, i_N) \end{aligned} \tag{3.20}$$

Example 3.24 A 2 player symmetric game must satisfy $\pi_\alpha(i,j) = \pi_\beta(j,i)$, that is, we have

$$\boldsymbol{\pi}_\beta = \boldsymbol{\pi}_\alpha^T .$$

In symmetric games it is natural to ask whether there exists NE which are symmetric, that is the strategies $\mathbf{p}_\gamma$ are all equal. Let

$$\Delta_{sym} = \{\mathbf{p} \in \Delta\,;\, \mathbf{p} = (\mathbf{q}, \cdots, \mathbf{q})\}$$

denote the set of symmetric strategy profile. It is easy to verify that Δ_{sym} is a convex and compact set.

Theorem 3.25 *A symmetric game has at least one symmetric NE $\mathbf{p}^* = (\mathbf{q}, \cdots, \mathbf{q}) \in \Delta_{sym}$.*

Proof: The proof is the same as for the existence of Nash equilibrium, Theorem 3.8. We use the same map T (see Equations (3.8) and (3.9)) and we need only to check that T maps Δ_{sym} into itself. From the definition of symmetric games we have for any γ_l , γ_k and j

$$\boldsymbol{\pi}_{\gamma_l}(\mathbf{q}, \cdots, \overbrace{\mathbf{e}_j}^{\gamma_l}, \cdots, \overbrace{\mathbf{q}}^{\gamma_k}, \cdots, \mathbf{q}) = \boldsymbol{\pi}_{\gamma_k}(\mathbf{q}, \cdots \overbrace{\mathbf{q}}^{\gamma_l}, \cdots, \overbrace{\mathbf{e}_j}^{\gamma_k}, \cdots, \mathbf{q})$$

and thus all the function $f_{j\gamma}$ coincide on Δ_{sym}. Therefore the map T defined in (3.9) maps Δ_{sym} into itself. The Brouwer's fixed point theorem implies that there exists a symmetric equilibrium. ∎

Symmetric games are very important in evolutionary game theory as they will be used to describe one single population of individuals who has each the same set of strategies. In this context of playing against the pure strategy $\mathbf{p}$ as equivalent to say that the proportion of strategies in the population is $\mathbf{p}$ and you choose an opponent at random in this population. In such a situation the relevant Nash equilibrium are the symmetric ones and they describe the state of a population of agents. For this reason if we have a symmetric games and we restrict ourselves to symmetric one states and Nash equilibria, we call such a game a *one population game*

The theorem of course does not mean that all NE have to be symmetric but just that at least one NE is symmetric.

Example 3.26 *(Hawk and Dove revisited).*Let us consider a 2-player hawk and dove games with payoff matrix

$$\boldsymbol{\pi}_\alpha = \boldsymbol{\pi}_\beta^T = \begin{pmatrix} -2 & 1 \\ 0 & 1/2 \end{pmatrix}$$

Let us look at strict NE. If player α chooses Hawk, i.e., $\mathbf{p}_\alpha = (1,0)$ then the payoff for player β is $(-2,0)$ and the best response for β is Dove i.e. $\mathbf{p}_\beta = (0,1)$. If β chooses Dove then the payoff for α is $(1,1/2)$ and so the best response to Dove is Hawk. By symmetry we find two pure, strict, and non-symmetric NE

$$((1,0)(0,1)) \qquad \text{and } ((0,1)(1,0))$$

If player α chooses the mixed strategy $(p,(1-p)$, the payoff for β is $(1-3p, 1/2-p/2)$. Unless the two payoffs are equal this will lead to a pure strategy best response and no NE. So we take $p = 1/5$ and obtain a strict symmetric NE

$$((1/5,4/5)(1/5,4/5))\ .$$

To conclude we classify the generic symmetric games with two player and two strategies. Using the rescaling of payoffs of section 3.4 we have

$$\boldsymbol{\pi}_\alpha = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \overset{NE}{\sim} \begin{pmatrix} a-c & 0 \\ 0 & d-b \end{pmatrix},$$

$$\boldsymbol{\pi}_\beta = \begin{pmatrix} a & c \\ b & d \end{pmatrix} \overset{NE}{\sim} \begin{pmatrix} a-c & 0 \\ 0 & d-b \end{pmatrix}.$$

We consider here only the generic cases where $a \neq c$ and $b \neq d$. There are generically three distinct cases

Generic symmetric two strategies games

- **(Symmetric coordination game)**: $a > c$ and $d > b$. There are 2 pure strict Nash equilibria with $p = (1,0)$ and $p = (0,1)$ and one mixed Nash equilibrium $\left(\frac{d-b}{(a-c)+(d-b)}, \frac{a-c}{(a-c)+(d-b)}\right)$
- **(Hawk and Dove)**: $a < c$ and $d < b$. There is a single mixed NE $\left(\frac{d-b}{(a-c)+(d-b)}, \frac{a-c}{(a-c)+(d-b)}\right)$
- **(Prisonner's dilemna)** $a > c$ and $d < b$ or $a < c$ and $d > b$. There is a single (pure) NE $(1,0)$ or $(0,1)$.

You may think of non-generic games as bifurcations between generic games. If $a > c$ and $d = b$ (or $a = c$ and $d > b$) then there are 2 pure NE (only one of which is strict) and this is a bifurcation from coordination game to a prisonner's dilemna game. If $a < c$ and $d = b$ or $a = c$ and $d < b$ then this a bifurcation between Hawk and prisonner's dilemna and there is one NE which is pure but not strict.

3.6 Potential Games

We discuss in this section a special class of games, called *potential games*, for which the computation of the Nash equilibria can be reduced to a standard optimization problem.

Definition 3.27 *Let us define*

1. *A game* $(\Gamma, S, (\boldsymbol{\pi}_\gamma)_{\gamma\in\Gamma})$ *is a* potential game, in the strict sense, *if there exists a function* $V : S \to \mathbb{R}$ *such that*

$$\boldsymbol{\pi}_\gamma = V \text{ for all } \gamma \in \Gamma\,, \tag{3.21}$$

   i.e., for any strategy profile $\mathbf{s}$ *the payoff for all players is the same.*

2. *A game* $(\Gamma, S, (\boldsymbol{\pi}_\gamma)_{\gamma\in\Gamma})$ *is a* potential game *if there exists a function* $\boldsymbol{\pi} : S \to \mathbb{R}$ *such that*

$$\boldsymbol{\pi}_\gamma \stackrel{NE}{\sim} V \text{ for all } \gamma \in \Gamma\,, \tag{3.22}$$

   i.e., the game is Nash equivalent to a potential game in the strict sense.

The function V is called the potential function *of the game* $(\Gamma, S, (\boldsymbol{\pi}_\gamma)_{\gamma\in\Gamma})$.

We can extend the function $\boldsymbol{\pi}$ as a multilinear function on the simplex Δ and to $\mathbb{R}^n$ by the formula

$$V(\mathbf{x}) \;=\; \sum_{i_1=1}^{n_{\gamma_1}} \cdots \sum_{i_N=1}^{n_{\gamma_N}} V(i_i, \cdots, i_N) x_{i_1}, \cdots x_{i_n} \,.$$

Using this we have that for any player γ and any strategy $i \in S_\gamma$

$$\frac{\partial V}{\partial x_{i,\gamma}}(\mathbf{p}) \;=\; \boldsymbol{\pi}_\gamma(\mathbf{e}_{i,\gamma}, \mathbf{p}_{-\gamma}) \,, \tag{3.23}$$

i.e., the gradient of the the potential function $\boldsymbol{\pi}$ gives the payoff for all players pure strategies.

Example 3.28 *Potential symmetric games.* A two-player games with payoff matrix $\boldsymbol{\pi}_\alpha$ and $\boldsymbol{\pi}_\beta$ is symmetric (in the strict sense) and a potential game (in the strict sense) if we have

$$\boldsymbol{\pi}_\alpha = \boldsymbol{\pi}_\beta^T \text{ (symmetric)} \quad \boldsymbol{\pi}_\alpha = \boldsymbol{\pi}_\beta \text{ (potential)}$$

and thus we must have

$$\boldsymbol{\pi}_\alpha = \boldsymbol{\pi}_\alpha^T$$

i.e., the payoff matrix is selfadjoint. More generally the game is a symmetric potential game if both $\boldsymbol{\pi}_\alpha$ and $\boldsymbol{\pi}_\beta$ are Nash equivalent to a self-adjoint matrix $\boldsymbol{\pi}$. In particular any two players, two strategies game is a symmetric game.

Definition 3.29 *The potential function for a one-population game is given by*

$$V(\mathbf{p}) \equiv \frac{1}{2}\boldsymbol{\pi}(\mathbf{p}, \mathbf{p}) \;=\; \frac{1}{2}\langle \mathbf{p} \,, \boldsymbol{\pi}\mathbf{p} \rangle$$

We can extend the function on $\mathbb{R}^n$ by multilinearity and we have

$$\frac{\partial V}{\partial x_i}(\mathbf{p}) \;=\; \frac{1}{2}\left((\boldsymbol{\pi}\mathbf{p})_i + (\boldsymbol{\pi}^T\mathbf{p})_i \right) \;=\; (\boldsymbol{\pi}\mathbf{p})_i = \boldsymbol{\pi}(\mathbf{e}_i, \mathbf{p}) \,.$$

and thus the gradient of the potential f gives the payoff of the pure strategies against the population state $\mathbf{p}$.

Let us consider the optimization problem which consist of finding

$$\max V(\mathbf{p})\,, \quad \mathbf{p} \in \Delta$$

i.e., we want to maximize the potential $V(p)$ on the simplex Δ, in other words under the constraints

$$\sum_i p_{i,\gamma} \;=\; 1\,,\;\; \gamma \in \Gamma\,,$$

$$p_{i,\gamma} \geq 0\,,\;\; i \in S_\gamma, \gamma \in \Gamma\,.$$

It is well-known that if f and h_j, $j = 1, \cdots, J$, are continuously differentiable functions of n variables $x = (x_1, \cdots, x_n)$ then the maximization problem

$$\max f(x)$$

subject to

$$h_j(x) \;=\; 0\,, \quad j = 1, \cdots, J$$

can be analyzed using Lagrange multipliers $\mu = (\mu_1, \cdots, \mu_J) \in \mathbb{R}^J$ and the Lagrange function

$$L(\mathbf{x}, \lambda) \;=\; F(x) - \sum_{j=1}^{J} \mu_j g_j(x)\,.$$

A necessary condition for x^* to be a maximizer is that $\mathbf{x}^*$ and μ^* satisfy the *Lagrange conditions*

$$\begin{aligned}
\frac{\partial L}{\partial x_i}(\mathbf{x}^*, \mu^*) \quad &= \quad \frac{\partial f}{\partial x_i}(\mathbf{x}^*) + \sum_j \mu_j^* \frac{\partial g_j}{\partial x_i}(\mathbf{x}^*) \;=\; 0 \\
\frac{\partial L}{\partial \mu_j}(\mathbf{x}^*, \mu^*) \quad &= \quad g_j(\mathbf{x}^*) \;=\; 0
\end{aligned}$$

Our optimization problem has both *equality and inequality constraints* and there is a generalization of Lagrange Theorem which gives a necessary condition to solve the optimization problem.

Let $f(x)$, $h_j(x)$, $j = 1, \cdots, J$, $g_k(x)$, $k = 1, \cdots, K$ be continuously differentiable functions of $\mathbf{x} = (x_1, \cdots, x_n)$. Consider

Optimization problem: Determine

$$\max f(x) \tag{3.24}$$

subject to

$$h_j(x) \;=\; 0\,,\;\; j=1,\cdots,J\,, \tag{3.25}$$
$$g_k(x) \;\geq\; 0\,,\;\; k=1,\cdots,K\,. \tag{3.26}$$

and let us introduce the Lagrange function

$$L(\mathbf{x},\mu,\lambda) \;=\; f(x) + \sum_{j=1}^{J}\mu_j h_j(x) + \sum_{k=1}^{K}\lambda_k g_k(x)\,. \tag{3.27}$$

We will use, but not prove, the following theorem.

Theorem 3.30 *(***Karush-Kuhn-Tucker Theorem)** *A necessary condition for $\mathbf{x}^*$ to solve the maximization problem (3.24)–(3.25)–(3.26) is that $\mathbf{x}^*$, λ^*, μ^* satisfy the Karush-Kuhn-Tucker conditions*

$$\begin{aligned}
\frac{\partial L}{\partial x_i}(\mathbf{x}^*,\mu^*,\lambda^*) &= \frac{\partial f}{\partial x_i}(\mathbf{x}^*) + \sum_j \mu_j \frac{\partial h_j}{\partial x_i}(\mathbf{x}^*) + \sum_k \lambda_k \frac{\partial g_k}{\partial x_i}(\mathbf{x}^*) \;=\; 0 \\
\frac{\partial L}{\partial \mu_j}(\mathbf{x}^*,\mu^*,\lambda^*) &= h_j(\mathbf{x}^*) \;=\; 0 \\
\frac{\partial L}{\partial \lambda_k}(\mathbf{x}^*,\mu^*,\lambda^*) &= g_k(\mathbf{x}^*) \;\geq\; 0 \\
\lambda_j &\geq 0 \\
\lambda_j g_j(\mathbf{x}^*) &\geq 0
\end{aligned}$$

For a potential game with potential function V, if we maximize V over the simplex, the Lagrange function is

$$L(\mathbf{x},\mu,\lambda) \;=\; V(x) + \sum_{\gamma\in\Gamma}\mu_\gamma\left(1-\sum_{j=1}^{n_\gamma}x_{j,\gamma}\right) + \sum_{\gamma\in\Gamma}\sum_{j=1}^{n_\gamma}\lambda_{j,\gamma}x_{j,\gamma}\,, \tag{3.28}$$

and the Karush-Kuhn-Tucker conditions are

$$\frac{\partial V}{\partial x_{i,\gamma}}(\mathbf{p}^*) \;=\; \boldsymbol{\pi}_\gamma(\mathbf{e}_{i,\gamma},\mathbf{p}^*_{-\gamma}) \;=\; \mu_\gamma - \lambda_{i,\gamma}\,, \tag{3.29}$$
$$\sum_i p^*_{i,\gamma} \;=\; 1\,, \tag{3.30}$$
$$p^*_{i,\gamma} \;\geq\; 0\,, \tag{3.31}$$
$$\lambda_{j,\gamma}p^*_{j,\gamma} \;=\; 0\,, \tag{3.32}$$
$$\lambda_{j,\gamma} \;\geq\; 0\,. \tag{3.33}$$

It turns out that the Karush-Kuhn-Tucker conditions are equivalent to the Nash equilibrium conditions.

Theorem 3.31 *Let* $(\Gamma, S, (\boldsymbol{\pi}_\gamma)_\gamma)$ *be a potential game with potential function* $V(p)$ *then* $\mathbf{p}^*$ *is a Nash equilibrium if and only if* $\mathbf{p}^*$ *satisfies the Kuhn-Tucker conditions.*

Proof: (i) Let us assume that $\mathbf{p}^*$ is a NE. For a potential game we have $\frac{\partial V}{\partial x_{i,\gamma}}(\mathbf{p}^*) = \boldsymbol{\pi}_\gamma(\mathbf{e}_{i,\gamma}, \mathbf{p}^*_{-\gamma})$. Let us choose

$$\begin{aligned} \mu_\gamma &= \max_j \boldsymbol{\pi}_\gamma(\mathbf{e}_{j,\gamma}, \mathbf{p}^*_{-\gamma}) \,, \\ \lambda_{i,\gamma} &= \mu_\gamma - \boldsymbol{\pi}_\gamma(\mathbf{e}_{i,\gamma}, \mathbf{p}^*_{-\gamma}) \,. \end{aligned}$$

Then the conditions (3.29), (3.30), and (3.31) are trivially satisfied. Since $\mathbf{p}^*$ is a NE, we have $\lambda_{i,\gamma} \geq 0$ and so (3.33) is satisfied. Furthermore if $i \in \Sigma_{\mathbf{p}^*}$ (i.e., $p_{i,\gamma} > 0$) then $\lambda_{i,\gamma} = 0$ and so (3.32) is also satisfied.

(ii) Conversely let us assume that the KTT conditions (3.29)–(3.33) are satisfied. If $p^*_{i,\gamma} > 0$ then $\lambda_{i,\gamma} = 0$ and thus by (3.29)

$$\frac{\partial V}{\partial x_{i,\gamma}}(\mathbf{p}^*) \,=\, \boldsymbol{\pi}_\gamma(\mathbf{e}_{i,\gamma}, \mathbf{p}^*_{-\gamma}) \,=\, \mu_\gamma \,,$$

is independent of i. On the other hand if $p^*_{i,\gamma} = 0$ then using (3.33) we have

$$\frac{\partial V}{\partial x_{i,\gamma}}(\mathbf{p}^*) \,=\, \boldsymbol{\pi}_\gamma(\mathbf{e}_{i,\gamma}, \mathbf{p}^*_{-\gamma}) \,=\, \mu_\gamma - \lambda_{i,\gamma} \,\leq\, \mu_\gamma \,.$$

This implies that $\mathbf{p}^*$ is a NE. ∎

Example 3.32 Let us compute the NE for the one-population game with payoff matrix

$$\boldsymbol{\pi} \,=\, \begin{pmatrix} 1 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 3 \end{pmatrix} \,.$$

The potential function $V(\mathbf{p})$ is given by

$$V(\mathbf{x}) \,=\, \frac{1}{2}\langle \mathbf{x}, \boldsymbol{\pi}\mathbf{x} \rangle \,=\, \frac{1}{2}x_1^2 + x_2^2 + \frac{3}{2}x_3^2 \,.$$

The KTT conditions are

$$
\begin{aligned}
p_1 &= \mu - \lambda_1 \\
2p_2 &= \mu - \lambda_2 \\
3p_3 &= \mu - \lambda_2 \\
\lambda_i p_i &= 0 \\
\lambda_i &\geq 0
\end{aligned}
$$

We need to distinguish several cases:

(i) $\lambda_1 = \lambda_2 = \lambda_3 = 0$, i.e. $p_i > 0$ for all i. This gives $p_1 = 2p_2 = 3p_3$ and so $\mathbf{p} = (6/11, 3/11, 2/11)$.

(ii) Let us assume that $\lambda_1 > 0$ but $\lambda_2 = \lambda_3 = 0$. This gives $p_1 = 0$ and so $\lambda_1 = \mu$ and $2p_2 = 3p_3$ and so $\mathbf{p} = (0, 3/5, 2/5)$. Similarly we also find the NE $(2/3, 1/3, 0)$ and $(3/4, 0, 1/4)$,

(iii) We also have the pure NE $(1,0,0)$, $(0,1,0)$ and $(0,0,1)$ this corresponds to choosing one single $\lambda_j > 0$ in the KTT conditions.

If you are given game a particular it is not obvious to determine a priori if it is Nash equivalent to a potential game. To this effect the following criterion is useful since it reduces the potential condition to a simple finite algorithm. We will need the following.

Definition 3.33 *A* path of length L, $\mathcal{P} = \{\mathbf{s}^{(0)}, \mathbf{s}^{(1)}, \cdots, \mathbf{s}^{(L)}\}$ *for a game* $(\Gamma, S, \boldsymbol{\pi})$ *is a collection of* $L+1$ *pure strategy profile such that* $\mathbf{s}^{(l-1)}$ *and* $\mathbf{s}^{(l)}$ *differ for exactly one player, i.e., for every* $1 \leq l \leq L$ *there exists a player* $\gamma(l)$ *and a strategy profile* $\mathbf{s}^{(l)}_{-\gamma(l)}$ *such that* $\mathbf{s}^{(l-1)} = (s_{\gamma(l)}, \mathbf{s}^{(l)}_{-\gamma(l)})$ *and* $\mathbf{s}^{(l)} = (s'_{\gamma(l)}, \mathbf{s}^{(l)}_{-\gamma(l)})$.

A path is said to be a closed path *if* $\mathbf{s}_0 = \mathbf{s}_L$ *and it is said to be a* simple path *if* $\mathbf{s}_k \neq \mathbf{s}_l$ *for* $k \neq l$ *with* $0 \leq k, l \leq L-1$.

Definition 3.34 *The* weight *of a path* $\mathcal{P}$*, denoted by* $I(\mathcal{P})$*, is given by*

$$
I(\mathcal{P}) = \sum_{l=1}^{L} \boldsymbol{\pi}_{\gamma(l)}(\mathbf{s}^{(l)}) - \boldsymbol{\pi}_{\gamma(l)}(\mathbf{s}^{(l-1)}) \,, \tag{3.34}
$$

where $\gamma(l)$ *is the unique player whose strategy is different in* $\mathbf{s}_{l-1}$ *and* $\mathbf{s}_l$.

Also we need the following

Lemma 3.35 *A game* $(\Gamma, S, \boldsymbol{\pi})$ *is Nash equivalent to a potential game if and only if of there exist* $V : S \to \mathbb{R}$ *and constants* a_γ *such that*

$$\boldsymbol{\pi}_\gamma(s_\gamma, \mathbf{s}_{-\gamma}) - \boldsymbol{\pi}_\gamma(s'_\gamma, \mathbf{s}_{-\gamma}) \;=\; a_\gamma \left[V(s_\gamma, \mathbf{s}_{-\gamma}) - V(s'_\gamma, \mathbf{s}_{-\gamma})\right] \,, \tag{3.35}$$

for all s_γ, s'_γ *and* $\mathbf{s}_{-\gamma}$.

Proof: If the game is Nash equivalent to a potential game then (**??**) follows immdiately. Conversely if (3.35) holds then we have with $\mathbf{s} = (s_\gamma, \mathbf{s}_{-\gamma})$ and for all s'_γ

$$\boldsymbol{\pi}(\mathbf{s}) \;=\; a_\gamma V(\mathbf{s}) + \big(\boldsymbol{\pi}_\gamma(s'_\gamma, \mathbf{s}_{-\gamma}) + a_\gamma V(s'_\gamma, \mathbf{s}_{-\gamma})\big) \tag{3.36}$$

$$\equiv\; a_\gamma V(s) + b_\gamma(\mathbf{s}_{-\gamma}) \tag{3.37}$$

The fact that equation (3.36) holds for all s'_γ implies that the quantity $b_\gamma(\mathbf{s}_{-\gamma})$ is independent of s_γ. ■

Theorem 3.36 *For a game* $(\Gamma, S, \boldsymbol{\pi})$ *the following conditions are equivalent*

1. *There exists* $V(\mathbf{s})$ *and* $b_\gamma(\mathbf{s}_{-\gamma})$ *such that for all* $\gamma \in \Gamma$

$$\boldsymbol{\pi}_\gamma(\mathbf{s}) \;=\; V(\mathbf{s}) + b_\gamma(\mathbf{s}_{-\gamma}) \,.$$

2. $I(\mathcal{P}) = 0$ *for every closed paths* $\mathcal{P}$.

3. $I(\mathcal{P}) = 0$ *for every simple closed paths* $\mathcal{P}$.

4. $I(\mathcal{P}) = 0$ *for every simple closed paths* $\mathcal{P}$ *of length* 4.

Proof: Before we start the proof we note that for simple closed paths of length 2 and 3, we note that $\gamma(l)$ is the same for every l and thus $I(\mathcal{P}) = 0$ for every game. Therefore the length 4 is the smallest nontrivial length for which a vanishing weight is a non-trivial condition

Trivially we have $2 \Rightarrow 3 \Rightarrow 4$.

We have $1 \Rightarrow 2$ since 1 implies that for any path $\mathcal{P}$

$$I(\mathcal{P}) \;=\; \sum_{l=1}^{L} V(\mathbf{s}^{(l)}) - V(\mathbf{s}^{(l-1)}) \;=\; V(\mathbf{s}^{(L)}) - V(\mathbf{s}^{(0)}) \,.$$

and this is 0 is the path is closed.

Conversely let us assume that $I(\mathcal{P}) = 0$ for all closed path. Then we claim for any 2 paths $\mathcal{P}_1 = \{\mathbf{s}', \mathbf{s}_1, \cdots, \mathbf{s}_{L-1}, \mathbf{s}\}$ and $\mathcal{P}_2 = \{\mathbf{s}', \mathbf{t}_1, \cdots, \mathbf{t}_{K-1}, \mathbf{s}\}$ starting at $\mathbf{s}'$ and ending at $\mathbf{s}'$ we have $I(\mathcal{P}_1) = I(\mathcal{P}_2)$. Suppose this not true then consider the closed path $\mathcal{P}$ which consists of going first from $\mathbf{s}'$ to $\mathbf{s}$ along $\mathcal{P}_1$ and then going from $\mathbf{s}$ to $\mathbf{s}'$ along the reversed path $\overline{\mathcal{P}}_2 = \{\mathbf{s}, \mathbf{t}_{K-1}, \cdots, \mathbf{t_1}, \mathbf{s}'\}$. Then we have $I(\mathcal{P}) = I(\mathcal{P}_1) - I(\mathcal{P}_2) \, = \, 0$ and this is a contradiction. This proves that $2 \Rightarrow 1$.

Finally we show that $4 \Rightarrow 2$. Let $L \geq 5$ be the smallest length for which there is a path $\mathcal{P}$ of length L with $I(\mathcal{P}) \neq 0$. Let $\gamma(1)$ be the player whose strategy changes between $\mathbf{s}^{(0)}$ and $\mathbf{s}^{(1)}$, then there exists at least one $1 < j \leq L$ such that $\gamma(j) = \gamma(1)$ since the path is closed. Suppose that $j = 2$, then the same player changes his strategy twice in a row along the path, let us consider the path of length $L-1$, $\tilde{\mathcal{P}} = \{\mathbf{s}^{(0)}, \mathbf{s}^{(2)}, \cdots, \mathbf{s}^{(L))}\}$. We have $I(\mathcal{P}) = I(\tilde{\mathcal{P}}) \neq 0$ and this contradicts our assumption that L was minimal. Since the path is closed, by the same argument we can exclude $j = L$. So we must have $2 < j < L$. We claim that we can construct a path

$$\overline{\mathcal{P}} \, = \, \{\mathbf{s}^0, \cdots, \mathbf{s}^{(j-2)}, \mathbf{t}, \mathbf{s}^{(j)}, \cdots, \mathbf{s}^{(L)}\}\,,$$

with $I(\mathcal{P}) = I(\overline{\mathcal{P}})$ and $\gamma(j-1) = \gamma(1)$. This is a contradiction since by repeating this argument we have $\gamma(2) = \gamma(1)$ which is impossible. To construct $\overline{\mathcal{P}}$ let us write

$$\begin{aligned}
\mathbf{s}^{(j-2)} &= (s_{\gamma(j-1)}, s_{\gamma(j)}, \mathbf{s}^{(j-2)}_{-(\gamma(j-1),\gamma(j))}) \\
\mathbf{s}^{(j-1)} &= (\tilde{s}_{\gamma(j-1)}, s_{\gamma(j)}, \mathbf{s}^{(j-2)}_{-(\gamma(j-1),\gamma(j))}) \\
\mathbf{s}^{(j)} &= (\tilde{s}_{\gamma(j-1)}, \tilde{s}_{\gamma(j)}, \mathbf{s}^{(j-2)}_{-(\gamma(j-1),\gamma(j))})
\end{aligned}$$

and let us define

$$\mathbf{t} = (s_{\gamma(j-1)}, \tilde{s}_{\gamma(j)}, \mathbf{s}^{(j-2)}_{-\gamma(j-1),\gamma(j))})$$

By construction we have

$$I(\overline{\mathcal{P}}) = I(\mathcal{P}) + I(\mathcal{Q})\,,$$

where $\mathcal{Q} = \{\mathbf{s}^{(j-2)}, \mathbf{s}^{(j-1)}, \mathbf{s}^{(j)}, \mathbf{t}, \mathbf{s}^{(j-2)}\}$ is a simple closed path of length 4. By assumption $I(\mathcal{Q}) = 0$. ∎

3.7 Two person zero-sum games

Zero-sum games is a class of games which has been widely studied, especially in the early work in game theory.

Definition 3.37 *A two player game* $\Gamma = (\{\alpha, \beta\}, S_\alpha \times S_\beta, \boldsymbol{\pi}_\alpha, \boldsymbol{\pi}_\beta)$ *is called a* zero-sum game *if*

$$\boldsymbol{\pi}_\alpha = -\boldsymbol{\pi}_\beta \tag{3.38}$$

Example 3.38 The matching pennies game of Example 3.5 is a zero-sum game.

Example 3.39 The Rock-Paper-Scissor game of Example 1.2 is zero-sum game. Since $\boldsymbol{\pi}_\alpha$ is anti-selfadjoint, $\boldsymbol{\pi}_\alpha = -\boldsymbol{\pi}_\alpha^T$ we have

$$\boldsymbol{\pi}_\alpha = -\boldsymbol{\pi}_\alpha^T = -(-\boldsymbol{\pi}_\beta)^T = \boldsymbol{\pi}_\beta^T \,,$$

and thus RPS is a zero-sum and symmetric.

We have already found the NE for these two games. Before we prove a general theorem about zero-sum game let us consider the following example which we can analyze using elementary considerations.

Example 3.40 Let us consider a zero-sum game with $S_\alpha = \{s_1, \cdots, s_5\}$ and $S_\beta = \{t_1, \cdots, t_4\}$ and

$$\boldsymbol{\pi}_\alpha = \begin{pmatrix} 18 & 3 & 0 & 2 \\ 0 & 3 & 8 & 20 \\ 5 & 4 & 5 & 5 \\ 16 & 4 & 2 & 25 \\ 9 & 3 & 0 & 20 \end{pmatrix}, \quad \boldsymbol{\pi}_\beta = -\boldsymbol{\pi}_\alpha$$

From the point view of α we have the following best responses

β choice	t_1	t_2	t_3	t_4
α's best response	s_1	s_3 or s_4	s_2	s_4
α's payoff	18	4	8	25

Similarly from the point view of β we have (use $\boldsymbol{\pi}_\alpha = -\boldsymbol{\pi}_\beta$)

α choice	s_1	s_2	s_3	s_4	s_5
β's best response	t_3	t_1	t_2	t_3	t_3
α's payoff	0	0	4	2	0

These two tables suggest the following solutions. From the second tables we see that if α chooses s_3 he will maximize the minimal payoff he can get, in this case 4. That is β minimizes $\pi_\alpha(i,j)$ over j and then α maximizes over i. From the point of view of β it makes sense to choose t_2 since it will get a payoff which at least -4 compared to -18, -8, or -25 for the other choices. In other words α minimizes $\pi(i,j)$ over i for fixed j and then β minimizes over i. Since we have that s_3 is a best response to t_2 and t_2 is a best response to s_3 then s_3 and t_2 is a NE. Furthermore we have

$$\pi_\alpha(3,2) \;=\; \max_i \min_j \pi_\alpha(i,j) \;=\; \min_j \max_i \pi_\alpha(i,j)\,.$$

This procedure, which is specific to zero-sum games, is called a *maximin solution.* Player α assumes that β will actually figure out which pure strategy α chooses and therefore α chooses the strategy which guarantees him the maximum minimal level of payoff. This leads to

- Find the minimum in each row and the maximum of all these numbers
- Find the maximum in each column and the minimum of all these numbers.
- If these two quantities coincide for a pair (i_0, j_0) then s_{i_0} and t_{j_0} are a (pure strategy) NE for the game.

This procedure does not always lead to a NE. For example for Rock-Paper-Scissor we find

$$\max_i \min_j \pi_\alpha(i,j) \;=\; -1\,, \quad \min_j \max_i \pi_\alpha(i,j) \;=\; 1\,.$$

This is consistent with the fact that RPS has no pure Nash equilibrium.

It also easy to see that one can find several minimax solutions, for example

$$\boldsymbol{\pi}_\alpha \;=\; \begin{pmatrix} 4 & 5 & 4 \\ 3 & 0 & 1 \end{pmatrix}$$

leads to pure NE s_1, t_1 and s_1, t_3.

The following theorem shows that the minimax strategy actually always work provided we also allow mixed strategies and that this actually leads to all NE of two person zero-sum games.

Theorem 3.41 *(***Minimax Theorem***)*

1. *If* $(\mathbf{p}_\alpha^*, \mathbf{p}_\beta^*)$ *is a NE of a zero-sum game then we have the formula*

$$\boldsymbol{\pi}_\alpha(\mathbf{p}_\alpha^*, \mathbf{p}_\beta^*) = \max_{q_\alpha} \min_{q_\beta} \boldsymbol{\pi}_\alpha(\mathbf{q}_\alpha, \mathbf{q}_\beta) = \min_{q_\beta} \max_{q_\alpha} \boldsymbol{\pi}_\alpha(\mathbf{q}_\alpha, \mathbf{q}_\beta), \tag{3.39}$$

   in particular all NE yields the same payoff.

2. $(\mathbf{p}_\alpha^*, \mathbf{p}_\beta^*)$ *is a NE of a zero-sum game if and only if*

$$\mathbf{p}_\alpha^* \text{ is a maximum of } \quad \mathbf{q}_\alpha \mapsto \min_{q_\beta} \boldsymbol{\pi}_\alpha(\mathbf{q}_\alpha, \mathbf{q}_\beta) \tag{3.40}$$

$$\mathbf{p}_\beta^* \text{ is a minimum of } \quad \mathbf{q}_\beta \mapsto \max_{q_\alpha} \boldsymbol{\pi}_\alpha(\mathbf{q}_\alpha, \mathbf{q}_\beta) \tag{3.41}$$

Proof: If $(\mathbf{p}_\alpha^*, \mathbf{p}_\beta^*)$ is a NE then we have

$$\boldsymbol{\pi}_\alpha(\mathbf{p}_\alpha^*, \mathbf{p}_\beta^*) = \max_{\mathbf{q}_\alpha} \boldsymbol{\pi}_\alpha(\mathbf{q}_\alpha, \mathbf{p}_\beta^*) \geq \max_{\mathbf{q}_\alpha} \min_{\mathbf{q}_\beta} \boldsymbol{\pi}_\alpha(\mathbf{q}_\alpha, \mathbf{q}_\beta), \tag{3.42}$$

$$\boldsymbol{\pi}_\alpha(\mathbf{p}_\alpha^*, \mathbf{p}_\beta^*) = \min_{\mathbf{q}_\beta} \boldsymbol{\pi}_\alpha(\mathbf{p}_\alpha^*, \mathbf{q}_\beta) \leq \min_{\mathbf{q}_\beta} \max_{\mathbf{q}_\alpha} \boldsymbol{\pi}_\alpha(\mathbf{q}_\alpha, \mathbf{q}_\beta). \tag{3.43}$$

The equality in (3.42) follows from the definition of NE while the inequality is simply the fact that $f(x) \leq g(x)$ implies $\max_x f(x) \leq \max_x g(x)$. The inequality (3.43) is similar.

On the other hand we have the inequalities

$$\max_{\mathbf{q}_\alpha} \min_{\mathbf{q}_\beta} \boldsymbol{\pi}_\alpha(\mathbf{q}_\alpha, \mathbf{q}_\beta) \geq \min_{\mathbf{q}_\beta} \boldsymbol{\pi}_\alpha(\mathbf{p}_\alpha^*, \mathbf{q}_\beta) \tag{3.44}$$

$$\min_{\mathbf{q}_\beta} \max_{\mathbf{q}_\alpha} \boldsymbol{\pi}_\alpha(\mathbf{q}_\alpha, \mathbf{q}_\beta) \leq \max_{\mathbf{q}_\alpha} \boldsymbol{\pi}_\alpha(\mathbf{q}_\alpha, \mathbf{p}_\beta^*) \tag{3.45}$$

The only possibility for (3.42), (3.43), (3.44), and (3.45) to be true is that all the inequalities are actually equalities and thus we have

$$\boldsymbol{\pi}_\alpha(\mathbf{p}_\alpha^*, \mathbf{p}_\beta^*) = \max_{\mathbf{q}_\alpha} \boldsymbol{\pi}_\alpha(\mathbf{q}_\alpha, \mathbf{p}_\beta^*) = \max_{\mathbf{q}_\alpha} \min_{\mathbf{q}_\beta} \boldsymbol{\pi}_\alpha(\mathbf{q}_\alpha, \mathbf{q}_\beta), \tag{3.46}$$

$$\boldsymbol{\pi}_\alpha(\mathbf{p}_\alpha^*, \mathbf{p}_\beta^*) = \min_{\mathbf{q}_\beta} \boldsymbol{\pi}_\alpha(\mathbf{p}_\alpha^*, \mathbf{q}_\beta) = \min_{\mathbf{q}_\beta} \max_{\mathbf{q}_\alpha} \boldsymbol{\pi}_\alpha(\mathbf{q}_\alpha, \mathbf{q}_\beta). \tag{3.47}$$

This implies that item 1. and the only if part of item 2.

Conversely let us assume that $\mathbf{p}_\alpha^*$ and $\mathbf{p}_\beta^*$ satisfy (3.40) and (3.41). Since there exists at least one Nash equilibrium by Theorem 3.8, we have equality between min-max and max-min. We have

$$\begin{aligned}
\boldsymbol{\pi}_\alpha(\mathbf{p}_\alpha^*, \mathbf{p}_\beta^*) &\geq \min_{\mathbf{q}_\beta} \boldsymbol{\pi}_\alpha(\mathbf{p}_\alpha^*, \mathbf{q}_\beta) \\
&= \max_{\mathbf{q}_\alpha} \min_{\mathbf{q}_\beta} \boldsymbol{\pi}_\alpha(\mathbf{q}_\alpha, \mathbf{q}_\beta) \quad (\text{ by } (3.40)) \\
&= \min_{\mathbf{q}_\beta} \max_{\mathbf{q}_\alpha} \boldsymbol{\pi}_\alpha(\mathbf{q}_\alpha, \mathbf{q}_\beta) \quad (\text{ by } (3.39)) \\
&= \max_{\mathbf{q}_\alpha} \boldsymbol{\pi}_\alpha(\mathbf{q}_\alpha, \mathbf{p}_\beta^*) \quad (\text{ by } (3.41)) \\
&\geq \boldsymbol{\pi}_\alpha(\mathbf{p}_\alpha^*, \mathbf{p}_\beta^*) \,. \qquad (3.48)
\end{aligned}$$

Therefore we must have equalities and this shows that $(\mathbf{p}_\alpha^*, \mathbf{p}_\beta^*)$ is a Nash equilibrium. ∎

Since all NE for zero-sum games lead to the same payoff we have

Definition 3.42 *The* value v *of a 2 players zero-sum game is the payoff for the corresponding Nash equilibria given in (3.39).*

Example 3.43 Let us solve the matching pennies game using the minimax Theorem. Let $\mathbf{q}_\alpha = (p, 1-p)$ and $\mathbf{q}_\beta = (q, 1-q)$ we have

$$\boldsymbol{\pi}_\alpha(\mathbf{q}_\alpha, \mathbf{q}_\beta) = \langle \mathbf{q}_\alpha, \boldsymbol{\pi}_\alpha \mathbf{q}_\beta \rangle = \left\langle \begin{pmatrix} p \\ 1-p \end{pmatrix}, \begin{pmatrix} 2q-1 \\ 1-2q \end{pmatrix} \right\rangle = (2p-1)(2q-1)$$

and thus

$$\min_q (2p-1)(2q-1) = \begin{cases} 1-2p & \text{if} \quad p \geq 1/2 \\ 2p-1 & \text{if} \quad p \leq 1/2 \end{cases} .$$

The maximum over p of this function is attained at $p = 1/2$ and is equal to 0 and so $\mathbf{p}_\alpha^* = (1/2, 1/2)$. By symmetry we find $\mathbf{q}_\alpha^* = (1/2, 1/2)$ and the equilibrium payoff for both players is 0.

As we have done for symmetric games let us classify the generic two strategy zero-sum games. Let us assume that the payoff has the form

$$\begin{array}{cc} & \begin{array}{cc} t_1 & t_2 \end{array} \\ \begin{array}{c} s_1 \\ s_2 \end{array} & \begin{pmatrix} a & b \\ c & d \end{pmatrix} \end{array}$$

and we assume that the game is generic in the sense no pair of entries in the matrix are equal.

Note first that if the pair of pure strategies s_1 and t_1 is a Nash equilibrium iff

$$s_1, t_1 \text{ N.E.} \quad \text{iff} \quad b > a > c\,.$$

For example $b > a$ ensures that a is the minimum in its rows while $a > c$ ensures that it is the maximin. Similarly $a > c$ ensures that a is the max of its colum and $b > a$ ensures that the minimax is a. Note also that this corresponds to a strictly dominated strategy for one the player. If $b > a > c > d$ or $b > a > d > c$ then s_1 strictly dominates s_2 for player α while if $b > d > a > c$ or $d > b > a > c$ then t_1 strictly dominates t_2 for player β.

Using the same reasonning we have

$$\begin{array}{lll} s_1, t_2 \text{ N.E.} & \text{iff} & a > b > d\,, \\ s_2, t_1 \text{ N.E.} & \text{iff} & d > c > a\,, \\ s_2, t_2 \text{ N.E.} & \text{iff} & c > d > b\,. \end{array}$$

This leaves only the games to consider where the payoff one diagonal of the matrix are bigger than the payoff on the other diagonal, e.g.

$$\min\{a, d\} > \max\{b, c\} \text{ or } \min\{b, c\} > \max\{a, d\}$$

All these conditions are the same up to relabelling of the payoffs so wlog let us asume that

$$a > d > b > c$$

To compute the NE we use the minimax theorem with $\mathbf{q}_\alpha = (p, 1-p)$ and $\mathbf{q}_\beta = (q, 1-q)$. Then we have

$$\boldsymbol{\pi}_\alpha(\mathbf{q}_\alpha, \mathbf{q}_\beta) \;=\; pq\left[(a-b) + (d-c)\right] - p(d-b) - q(d-c) + d$$

and so

$$\min_{\mathbf{q}_\beta} \boldsymbol{\pi}_\alpha(\mathbf{q}_\alpha, \mathbf{q}_\beta) \;=\; \begin{cases} d - p(d-b) & \text{if} \quad p \geq \frac{d-c}{(a-b)+(d-c)} \\ p(a-c) + c & \text{if} \quad p \leq \frac{d-c}{(a-b)+(d-c)} \end{cases}$$

and the maximum over p is attained iff $p = \frac{d-c}{(a-b)+(d-c)}$. On the other hand we have

$$\max_{\mathbf{p}_\alpha} \boldsymbol{\pi}_\alpha(\mathbf{q}_\alpha, \mathbf{q}_\beta) \;=\; \begin{cases} d - p(d-c) & \text{if} \quad q \geq \frac{d-b}{(a-c)+(d-b)} \\ q(a-b) + b & \text{if} \quad q \leq \frac{d-b}{(a-c)+(d-b)} \end{cases}$$

and the maximum over q is attained iff $q = \frac{d-b}{(a-c)+(d-b)}$. The value of the game is

$$v = \frac{da - bc}{(a+d)-(b+c)}\,.$$

Generic two strategies zero-sum games: Generically there are 2 zero-sum 2 strategy games.

- **One mixed NE** This occurs if $\min\{a,d\} > \max\{b,c\}$ or $\min\{b,c\} > \max\{a,d\}$.
- **One pure NE** This occurs if there is one strictly dominating strategy for α or β.

Next we explore the deep connection between the minimax theorem of game theory and linear programming. From Theorem 3.39 we know that , on one hand, if α chooses the NE strategy $\mathbf{p}_\alpha^*$ then α's payoff will be at least v no matter what β's strategy is. On the other hand if β chooses the NE $\mathbf{q}_\beta^*$ then $\alpha's$ payoff will be no more than v, no matter what α's strategy is. This shows that the following two problems have a solution with the same optimal value v^* for v.

Player α problem: Find the largest v and $\mathbf{x} = (x_1, \cdots, x_n) \in \mathbb{R}^n$ with $x_i \ge 0$, $\sum_{i=1}^n x_i = 1$, and

$$\sum_{i=1}^{n} x_i \pi_\alpha(i,j) \;=\; (\boldsymbol{\pi}_\alpha^T x)_j \;\ge\; v\,, \quad j = 1, \cdots, m\,. \tag{3.49}$$

Player β problem: Find the smallest v and $\mathbf{y} = (y_1, \cdots, y_m) \in \mathbb{R}^m$ with $y_j \ge 0$, $\sum_{j=1}^m y_j = 1$, and

$$\sum_{i=1}^{m} \pi_\alpha(i,j) y_j \;=\; (\boldsymbol{\pi}_\alpha y)_i \;\le\; v\,, \quad i = 1, \cdots, n\,. \tag{3.50}$$

Let us see how this formulation can be useful to solve a game

Example 3.44 Let us consider the zero-sum game with payoff matrix

$$\boldsymbol{\pi}_\alpha = \begin{pmatrix} 4 & 0 & 3 \\ -3 & 2 & 0 \end{pmatrix}$$

It is easy to check that there are no pure NE. The optimization problem for the payoff are, with $\mathbf{p}_\alpha = (p, 1-p)$ and $\mathbf{q}_\beta = (q, r, 1-q-r)$

$$
\begin{aligned}
4p - 3(1-p) &\geq v & (3.51)\\
2(1-p) &\geq v & (3.52)\\
3p - (1-p) &\geq v & (3.53)\\
4q + 3(1-q-r) &\leq v & (3.54)\\
-3q + 2r - (1-q-r) &\leq v & (3.55)
\end{aligned}
$$

First note that we could find a solution if we assume that all $\geq$ are actually equality. Since we have 4 unknowns but 5 equations this system is overdetermined and might not have a solution. Let use (3.51)-(3.52) with equalities to solve for p and v, we find $p = 5/9$ and $v = 8/9$. Using the value of v to solve (3.54)–(3.55) we find $q = 2/9$ and $r = 4/9$. Finally for (3.53) we have the (strict) inequality $11/9 > 8/9$. This yields a NE $((5/9, 4/9), (2/9, 7/9, 0))$. We could have solved for p by using equalities with (3.51)–(3.53) and would have found $p = 2/3$ and $v = 5/3$ but this is *not consistent* with (3.52). If we solve for p and v using (3.52)–(3.53) we find $p = 1/2$ and $v = 1$ but this is inconsistent with (3.51). We see that the strategy t_3 for β is not played even though it is not strictly dominated.

Without loss of generality we can assume that $v > 0$. Indeed by rescaling the payoff as

$$\boldsymbol{\pi}'_\alpha = \boldsymbol{\pi}_\alpha + A\,, \quad \boldsymbol{\pi}'_\beta = \boldsymbol{\pi}_\beta - A$$

for a sufficiently large A we do not change the structure of the game (it is still a zero-sum game with the same NE) but it ensures that all matrix elements and thus also the value of the game v are positive.

Using this rescaling we can reformulate the two optimization problem by dividing by v both sides of (3.49) and (3.50). We then set $u_i = x_i/v$ and $w_i = y_i/v$ Note that $\sum_i u_i = \sum w_i = 1/v$ and thus maximizing (resp. minimizing) v is equivalent to minimizing the linear form $\sum_i u_i$ (resp. maximizing the linear form $\sum_i w_i$). Using this we can reformulate the optimization problems for the players as

Player α problem: Find $\mathbf{u} = (u_1, \cdots, u_n) \in \mathbb{R}^n$ such as to

$$\text{miminize } \sum_{i=1}^{n} u_i$$

with the constraints

$$u_i \geq 0\,, \qquad i = 1, \cdots, n\,, \tag{3.56}$$
$$(\boldsymbol{\pi}_\alpha^T \mathbf{u})_j \geq 1\,, \qquad j = 1, \cdots, m\,. \tag{3.57}$$

Player β problem: Find $\mathbf{v} = (v_1, \cdots, v_m) \in \mathbb{R}^m$ such as to

$$\text{maximize} \sum_{j=1}^{m} w_j\,,$$

with the constraints

$$w_j \geq 0\,, \qquad j = 1, \cdots, m\,, \tag{3.58}$$
$$(\boldsymbol{\pi} \mathbf{w})_j \leq 1\,, \qquad i = 1, \cdots, n\,. \tag{3.59}$$

The problems for α and β are called *dual linear programming problems.* Note that player α can get at least $1/\sum_i u_i$ and at most $1/\sum_j w_j$ and so we must have

$$\frac{1}{\sum_i u_i} \leq \frac{1}{\sum_j w_j}$$

But from the minimax theorem we know that they are actually equal and so we have a symmetric version of these dual programs.

Symmetric problem: Find $\mathbf{u} \in \mathbb{R}^n$ which satisfies (3.56)–(3.57) and $\mathbf{w} \in \mathbb{R}^m$ which satisfies (3.58)–(3.59) so that

$$\sum_i u_i = \sum_j w_j\,.$$

These two problems are special cases of linear programing problems. To formulate the general problems let $\mathbf{b} = (b_1, \cdots, b_m) \in \mathbb{R}^m$ and $\mathbf{c} = (c_1, \cdots, c_n) \in \mathbb{R}^n$ be two given vector and let $\boldsymbol{\pi}$ be $n \times m$ matrix. No positivity assumption is required for the entries of the vectors and matrix. The linear programming problems are

Minimization problem: Find $\mathbf{u} = (u_1, \cdots, u_n)$ such as to

$$\text{miminize}\ \langle \mathbf{c}\,,\, \mathbf{u} \rangle$$

with the constraints

$$u_i \geq 0\,, \qquad i = 1, \cdots, n\,. \tag{3.60}$$

$$(\boldsymbol{\pi}^T \mathbf{u})_j \;\geq\; b_j\,, \qquad j = 1, \cdots, m\,. \tag{3.61}$$

Maximization problem: Find $\mathbf{w} = (w_1, \cdots, w_m)$ such as to

$$\text{maximize } \langle \mathbf{b}\,,\, \mathbf{w} \rangle$$

with the constraints

$$w_j \geq 0\,, \qquad j = 1, \cdots, m\,, \tag{3.62}$$

$$(\boldsymbol{\pi}\mathbf{w})_i \leq c_i\,, \qquad i = 1, \cdots, n\,. \tag{3.63}$$

Symmetric problem: Find $\mathbf{u}$ which satisfies the constraints (3.60), (3.61) and $\mathbf{w}$ which satisfies the constraints (3.62), and (3.63) such that

$$\langle \mathbf{c}\,,\, \mathbf{u} \rangle \;=\; \langle \mathbf{b}\,,\, \mathbf{w} \rangle\,. \tag{3.64}$$

We have

Theorem 3.45 (Main theorem of linear programing)

1. *If* $\mathbf{u}$ *satisifies the constraints (3.60)–(3.61) and* $\mathbf{w}$ *satisifies the constraints (3.62)–(3.63) then*

$$\langle \mathbf{c}\,,\, \mathbf{u} \rangle \;\geq\; \langle \mathbf{b}\,,\, \mathbf{w} \rangle\,.$$

2. *If* $\mathbf{u}^*$ *and* $\mathbf{w}^*$ *are a solution of the symmetric problem, then* $\mathbf{u}^*$ *is a solution of the minimization problem and* $\mathbf{w}^*$ *is a solution of the maximization problem.*

3. *If* $\mathbf{u}^*$ *is a solution of the minimization problem and* $\mathbf{w}^*$ *is a solution of the maximization problem then*

$$\langle \mathbf{c}\,,\, \mathbf{u} \rangle \;=\; \langle \mathbf{b}\,,\, \mathbf{w} \rangle\,.$$

4. *If a solution exists to one problem, then the solution exists to all three problems.*

5. *If there exists at least one* $\mathbf{u}$ *which satisifies the constraints (3.60)–(3.61) and at least one* $\mathbf{w}$ *satisifies the constraints (3.62)–(3.63) then all three problems have a solution.*

Proof: We give a proof which is based on the minimax theorem.

1. This follows from (3.61)–(3.63):

$$\langle \mathbf{b}\,,\,\mathbf{w}\rangle = \sum_j b_j w_j \le \sum_j \sum_i \pi(i,j) u_i w_j \le \sum_i c_i u_i \;=\; \langle \mathbf{c}\,,\,\mathbf{u}\rangle\,.$$

2. If u^* and w^* solve the symmetric problem then they each satisfy their respective constraints and from 1. we see that u^* minimizes $\langle \mathbf{c}\,,\,\mathbf{u}\rangle$ and w^* maximizes $\langle \mathbf{b}\,,\,\mathbf{w}\rangle$.

3.-4.-5. The idea is to consider the zero-sum symmetric game with $n + m + 1$ strategies and payoff matrix

$$\begin{pmatrix} \mathbf{0} & -\boldsymbol{\pi}^{\mathbf{T}} & \mathbf{b} \\ \boldsymbol{\pi} & \mathbf{0} & -\mathbf{c} \\ -\mathbf{b}^T & \mathbf{c}^T & 0 \end{pmatrix}$$

Since the game is symmetric its value is equal to 0, let us write the minimax strategy for player β as

$$\begin{pmatrix} \mathbf{y} \\ \mathbf{x} \\ z \end{pmatrix}$$

with $\mathbf{y} \in \mathbb{R}^m$, $\mathbf{x} \in \mathbb{R}^n$, and $z \in \mathbb{R}$. By the mimimax theorem this strategy gives a payoff for α no bigger than 0 and so we have

$$-(\boldsymbol{\pi}^T\mathbf{x})_j + z b_j \;\le\; 0 \tag{3.65}$$
$$(\boldsymbol{\pi}\mathbf{y})_i - z c_i \;\le\; 0 \tag{3.66}$$
$$-\langle \mathbf{b}\,,\,\mathbf{y}\rangle + \langle \mathbf{c}\,,\,\mathbf{x}\rangle \;\le\; 0\,. \tag{3.67}$$

There are now 2 cases to distinguish.

Case 1. *There exists a minimax strategy for player* 2 *with* $z > 0$*.* Then from (3.65)–(3.66) the vectors

$$\mathbf{u} = \frac{1}{z}\mathbf{x}\,, \quad \mathbf{w} = \frac{1}{z}\mathbf{y}$$

do satisfy the constraints of the dual problems while from (3.67) we have

$$\langle \mathbf{b}\,,\,\mathbf{w}\rangle \geq \langle \mathbf{c}\,,\,\mathbf{u}\rangle\,.$$

Since, from 1. the opposite inequality always holds we actually have equality and $\mathbf{u}$, $\mathbf{w}$ are solution of the symmetric problem. By 2. $\mathbf{u}$ and $\mathbf{w}$ are solution of the maximizing and minimizing problem respectively.

Case 2. *There exists no minimax strategy for player* 2 *with* $z > 0$.

The crucial remark is that there exists a minimax strategy for player β such that the payoff for α if he plays strategy s_{n+m+1} against this strategy is strictly less than 0. Indeed, by contradiction, assume that all minimax strategies for β gives a payoff 0 for α if he plays s_{n+m+1}, then player α has a maximin strategy which gives positive probability to s_{n+m+1}. But by symmetry of the problem this would imply that β has a minimax strategy which assigns positive probability to t_{n+m+1} and this is a contradiction. As a consequence we have, for that strategy,

$$-\langle \mathbf{b}\,,\,\mathbf{y}\rangle + \langle \mathbf{c}\,,\,\mathbf{x}\rangle < 0\,.$$

Let us assume next that there exists $\mathbf{w}$ which satisfies the constraints (3.62)–(3.63) then we have using (3.65) (with the minimax strategy with $z = 0$)

$$\langle \mathbf{b}\,,\,\mathbf{y}\rangle > \langle \mathbf{c}\,,\,\mathbf{x}\rangle \geq \langle \boldsymbol{\pi}\mathbf{w}\,,\,\mathbf{x}\rangle \;=\; \langle \mathbf{w}\,,\,\boldsymbol{\pi}^T\mathbf{x}\rangle \geq 0 \tag{3.68}$$

If we assume that there exists $\mathbf{u}$ which satisfies the constraints (3.60)–(3.61) then with the minimax strategy with $z = 0$ we have using (3.66)

$$\langle \mathbf{b}\,,\,\mathbf{y}\rangle \leq \langle \boldsymbol{\pi}^T\mathbf{u}\,,\,\mathbf{y}\rangle = \langle \mathbf{u}\,,\,\boldsymbol{\pi}\mathbf{y}\rangle \leq 0 \tag{3.69}$$

As a consequence there cannot exist both an $\mathbf{u}$ which satisfies the constraints (3.60)–(3.61) and an $\mathbf{w}$ which satisfies the constraints (3.62)–(3.63), since (3.68) and (3.69) contradict each other.

Let us assume then that there exists $\mathbf{w}$ which satisfies (3.62)–(3.63) and then for y given by the minimax strategy with $z = 0$, $\mathbf{w} + \lambda\mathbf{y}$, for any $\lambda > 0$, also satisfies the constraints (3.62)–(3.63). Indeed we have

$$(\boldsymbol{\pi}(\mathbf{w} + \lambda\mathbf{y}))_i \;=\; (\boldsymbol{\pi}\mathbf{w})_i + \lambda(\boldsymbol{\pi}\mathbf{y})_i \leq (\boldsymbol{\pi}\mathbf{w})_i \leq c_i$$

But we have

$$\langle \mathbf{b}\,,\,\mathbf{w} + \lambda\mathbf{y}\rangle \;=\; \langle \mathbf{b}\,,\,\mathbf{w}\rangle + \lambda\langle \mathbf{b}\,,\mathbf{y}\rangle$$

Since by (3.68) $\langle \mathbf{b}, \mathbf{y} \rangle > 0$, $\max_{\mathbf{w}} \langle \mathbf{b}, \mathbf{w} \rangle = +\infty$.

One can prove in a similar way the dual assertion that if exists $\mathbf{u}$ which satisfies (3.60)–(3.61) then there $\langle \mathbf{c}, \mathbf{u} \rangle$ has no mimimum.

We can now conclude the proof. If we are in case 1. then the minimax theorem implies that there exists a solution for all three problems. Case 2 shows that 4. and 5. hold. ∎

4 Population game dynamics

4.1 Evolutionary stability

In this section we restrict ourselves to *symmetric games* with strategy set $S = \{1, \cdots n\}$ and we will always take the population point of view. We think of one (large) population of players where all players are indistinguishable from each other. We interpret a mixed strategy $p \in \Delta$ as representing the state of the population, i.e., p_i is the proportion of players playing strategy i. In this context only the symmetric NE are of interest.

A key concept in evolutionary game theory is the that of an *evolutionary stable* strategy which was introduced by Maynard Smith. Such a strategy is robust to evolutionary pressures in the following sense. Suppose that all individuals are genetically programmed to play a given strategy $p \in \Delta$ and that now a small population share of individuals who are likewise programmed to play another strategy $\mathbf{q}$ are injected into this population. The incumbent strategy is said to be evolutionary stable if, for each such mutant strategy, there exists a positive invasion barrier such that if the mutant population share falls below this barrier then the incumbent strategy earns a higher payoff than the mutant strategy. This criterion refers implicitly to a close connection between payoffs and the spreading of strategies in the population. How to model the spreading of strategies over time will be discussed in the next sections. The payoff in the game are supposed to represent here the gain in biological fitness (that is, reproductive value) obtained from a choice of strategy. In this biological interpretation the evolutionary stability criterion can be thought as a generalization of Darwin's survival of the fittest. How such an equilibrium is reached is not explained by the evolutionary stability property but it asks whether, once reached, it is robust with respect evolutionary changes.

Since the game is symmetric we only need to specify the $n \times n$ payoff matrix $\boldsymbol{\pi} = \boldsymbol{\pi}_\alpha$ for player α and we recall that the payoff for player strategy $\mathbf{p}$ against strategy $\mathbf{q}$ is

$$\boldsymbol{\pi}(\mathbf{p}, \mathbf{q}) = \langle \mathbf{p}\, , \boldsymbol{\pi}\mathbf{q} \rangle\, .$$

Note that the r.h.s of this formula allows to extend the function $\boldsymbol{\pi} : \Delta \times \Delta \to \mathbb{R}$ as a function $\boldsymbol{\pi} : \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}$.

Suppose the state of the population is the *incumbent strategy* $\mathbf{p}$ and that a small group appear within the population which all use the *mutant strategy* $\mathbf{q}$. Let us assume that the proportion of mutant is ϵ. If we pick an individual in this new population he will play the strategy $\mathbf{p}$ with probability $1-\epsilon$ and the strategy $\mathbf{q}$ with probability ϵ so the payoff against a randomly chosen individual is the same as the payoff against an individual playing the strategy $\epsilon\mathbf{q} + (1-\epsilon)\mathbf{p}$. The strategy $\mathbf{p}$ is evolutionary stable if it more advantageous to play $\mathbf{p}$ than $\mathbf{q}$ against this population state.

Definition 4.1 *A strategy* $\mathbf{p}$ *is an* evolutionary stable strategy *(ESS) if for every strategy* $\mathbf{q} \neq \mathbf{p}$ *there exists* $\bar{\epsilon}_{\mathbf{q}}$ *such that*

$$\boldsymbol{\pi}(\mathbf{p}, \epsilon\mathbf{q} + (1-\epsilon)\mathbf{p}) > \boldsymbol{\pi}(\mathbf{q}, \epsilon\mathbf{q} + (1-\epsilon)\mathbf{p}) \quad \text{for} \quad 0 < \epsilon < \bar{\epsilon}_{\mathbf{q}}\, . \tag{4.1}$$

The following proposition shows that every ESS is a NE and but that the ESS condition is, in general, stronger that the NE condition.

Proposition 4.2 *Let* $\mathbf{p}$ *be an ESS.*

1. *The strategy* $\mathbf{p}$ *is a NE.*
2. *Let* $\mathbf{q}$ *be any best response to* $\mathbf{p}$*, with* $\mathbf{q} \neq \mathbf{p}$ *then*

$$\boldsymbol{\pi}(\mathbf{q}, \mathbf{q}) < \boldsymbol{\pi}(\mathbf{p}, \mathbf{q})\, ,$$

that is any alternative best response $\mathbf{q}$ *to* $\mathbf{p}$ *other than* $\mathbf{p}$ *earns against itself a smaller payoff than what* $\mathbf{p}$ *would earn against this alternative best response.*

Proof: 1. Suppose that $\mathbf{p}$ is an ESS then $\mathbf{p}$ is a NE. It it were not a NE there would exists a strategy $\mathbf{q}$ with $\boldsymbol{\pi}(\mathbf{q}, \mathbf{p}) > \boldsymbol{\pi}(\mathbf{p}, \mathbf{p})$. Then by continuity of the payoff function for ϵ small enough we have $\boldsymbol{\pi}(\mathbf{q}, \epsilon\mathbf{q} + (1-\epsilon)\mathbf{p}) > \boldsymbol{\pi}(\mathbf{p}, \epsilon\mathbf{q} + (1-\epsilon)\mathbf{p})$ and this contradict the ESS property.

2. Let $\mathbf{q} \neq \mathbf{p}$ be a best response to $\mathbf{p}$ (i.e. $\boldsymbol{\pi}(\mathbf{q}, \boldsymbol{\pi}) = \boldsymbol{\pi}(\mathbf{p}, \mathbf{p})$) and let us assume that $\mathbf{q}$ earns against itself a payoff which is at least as much as what $\mathbf{p}$ earns against $\mathbf{q}$ (i.e., $\boldsymbol{\pi}(\mathbf{q}, \mathbf{q}) \geq \boldsymbol{\pi}(\mathbf{p}, \mathbf{q})$). Then by the multilinearity of the payoff function we have

$$\boldsymbol{\pi}(\mathbf{q}, \epsilon\mathbf{q} + (1-\epsilon)\mathbf{p}) \geq \boldsymbol{\pi}(\mathbf{p}, \epsilon\mathbf{q} + (1-\epsilon)\mathbf{p}) .$$

and so $\mathbf{p}$ is not an ESS. ∎

We can reformulate the ESS condition as

Corollary 4.3 *A strategy* $\mathbf{p}$ *is an ESS if and only if it satisfies*

1. *First order condition:* $\boldsymbol{\pi}(\mathbf{q}, \mathbf{p}) \leq \boldsymbol{\pi}(\mathbf{p}, \mathbf{p})$ *for all* $q \in \Delta$.
2. *Second order condition:* $\boldsymbol{\pi}(\mathbf{q}, \mathbf{p}) = \boldsymbol{\pi}(\mathbf{p}, \mathbf{p})$ *then* $\boldsymbol{\pi}(\mathbf{q}, \mathbf{q}) < \boldsymbol{\pi}(\mathbf{p}, \mathbf{q})$.

This is actually the original defintion of ESS due to Maynard Smith and the equivalent definition we have given is due to Taylor and Jonker.

We can also illustrate the two equivalent definition of the ESS as follows: for any $\mathbf{q} \in \Delta$ and $\epsilon > 0$ let us define

$$f_{\mathbf{p}}(\epsilon, \mathbf{q}) \;=\; \boldsymbol{\pi}(\mathbf{p} - \mathbf{q}, \epsilon\mathbf{q} + (1-\epsilon)\mathbf{p}) \;=\; \boldsymbol{\pi}(\mathbf{p} - \mathbf{q}, \mathbf{p}) + \epsilon\boldsymbol{\pi}(\mathbf{p} - \mathbf{q}, \mathbf{q} - \mathbf{p}) .$$

and so $f_{\mathbf{p}}(\epsilon, \mathbf{q})$ is a linear function with vertical intercept $\boldsymbol{\pi}(\mathbf{p} - \mathbf{q}, \mathbf{p})$ and slope $\boldsymbol{\pi}(\mathbf{p} - \mathbf{q}, \mathbf{q} - \mathbf{p})$. The strategy $\mathbf{p}$ is an ESS if and only if the function $f_{\mathbf{p}}(\epsilon, \mathbf{q})$ is nonnegative for any $\mathbf{q}$ and for any sufficiently small neighborhood of 0. If $\mathbf{q}$ is not a best response to $\mathbf{p}$ then the vertical intercept is positive and thus $f_{\mathbf{p}}$ is a positive for ϵ small enough. If $\mathbf{q}$ is a best response to $\mathbf{p}$ then $f_{\mathbf{p}}$ is positive iff $\boldsymbol{\pi}(\mathbf{p} - \mathbf{q}, \mathbf{p} - \mathbf{q}) < 0$.

The following lemma gives some information on the structure of the set of ESS's.

Lemma 4.4 *We have*

1. *If* $\mathbf{p}$ *is a strict NE then* $\mathbf{p}$ *is an ESS.*
2. *If* $\mathbf{p}$ *is an ESS then the support of* $\mathbf{p}$ *does not contain the support of another NE. In particular if* $\mathbf{p}$ *is an ESS in the interior of* Δ *then* $\mathbf{p}$ *is the unique NE.*

Proof: 1. If $\mathbf{p}$ is a strict NE then $\mathbf{p}$ is a pure strategy and the unique best response to $\mathbf{p}$. This implies that $\mathbf{p}$ is an ESS.

2. Suppose $\mathbf{q}$ is a NE and $\Sigma_{\mathbf{q}} \subset \Sigma_{\mathbf{p}}$. Since any pure strategy in Σ_p is a best response to $\mathbf{p}$ we have $\boldsymbol{\pi}(\mathbf{q},\mathbf{p}) = 0$. The second order condition implies then that $\boldsymbol{\pi}(\mathbf{p},\mathbf{q}) > \boldsymbol{\pi}(\mathbf{q},\mathbf{q})$ but this contradicts the fact that $\mathbf{q}$ is a NE. ■

Example 4.5 Two strategies games Note that the definition of ESS only involve difference of payoff and thus an ESS is left invariant under the linear rescaling of payoffs given in Section 3.4. So without loss of generality we can assume that the payoff matrix $\boldsymbol{\pi}$ has the form

$$\boldsymbol{\pi} = \begin{pmatrix} a_1 & 0 \\ 0 & a_2 \end{pmatrix} .$$

Case (i): $a_1 > 0$**,** $a_2 > 0$**.** There are 3 NE, $\mathbf{e}_1$, $\mathbf{e}_2$ and the mixed strategy NE $\left(\frac{a_2}{a_1+a_2}, \frac{a_1}{a_1+a_2}\right)$. The two pure NE are strict and thus also ESS by Lemma 4.4, 1. By Lemma 4.4, 2., the mixed strategy is not a ESS.

Case (ii): $a_1 < 0$**,** $a_2 < 0$**.** In this case there is a unique NE $\mathbf{p} = \left(\frac{a_2}{a_1+a_2}, \frac{a_1}{a_1+a_2}\right)$. It is an ESS since the matrix $\boldsymbol{\pi}$ is negative definite and so $\boldsymbol{\pi}(\mathbf{p}-\mathbf{q},\mathbf{p}-\mathbf{q}) < 0$ for any $\mathbf{q}$.

Case (iii): $a_1 < 0$**,** $a_2 > 0$**, or** $a_1 > 0$**,** $a_2 < 0$**.** In this case there is a unique (pure and strict) NE which is also an ESS

Example 4.6 Rock-Paper-Scissor Let us consider the symmetric game with three strategies and payoff matrix

$$\boldsymbol{\pi} = \begin{pmatrix} 0 & 1+a & -1 \\ -1 & 0 & 1+a \\ 1+a & -1 & 0 \end{pmatrix}$$

We need to distinguish several cases

(i) If $a < -1$ then $\mathbf{e}_1$, $\mathbf{e}_2$, $\mathbf{e}_3$ are strict NE and thus also ESS.

(ii) If $a > -1$ then $\mathbf{p} = (1/3, 1/3, 1/3)$ is the unique NE. Against $\mathbf{p}$ the payoff is $(a/3, a/3, a/3)$ and so $\boldsymbol{\pi}(\mathbf{q},\mathbf{p}) = \boldsymbol{\pi}(\mathbf{p},\mathbf{p})$ for any $\mathbf{q} \in \Delta$. To check the second order equation we first note that for any $\mathbf{x} \in \mathbb{R}^3$

$$\langle \mathbf{x} \,, \boldsymbol{\pi}\mathbf{x} \rangle = a(x_1x_2 + x_2x_3 + x_3x_1) \,. \tag{4.2}$$

If $\mathbf{x} = \mathbf{p} - \mathbf{q}$ with $\mathbf{p}, \mathbf{q} \in \Delta$ then $\sum_i x_i = 0$. So we have

$$0 = (x_1 + x_2 + x_3)^2 \; = \; \langle \mathbf{x} \, , \, \mathbf{x} \rangle + 2(x_1 x_2 + x_2 x_3 + x_3 x_1) \tag{4.3}$$

and thus for such x we have

$$\langle \mathbf{x} \, , \boldsymbol{\pi} \mathbf{x} \rangle \; = \; -\frac{a}{2} \langle \mathbf{x} \, , \mathbf{x} \rangle$$

So we have

(ii)(a) If $-1 \le a \le 0$ there is no ESS.

(ii)(b) If $0 < a$ the interior NE $(1/3, 1/3, 1/3)$ is an ESS.

In the definition of ESS one needs, in principle to check the conditions for all $\mathbf{q} \in \Delta$. The following proposition shows that the ESS property is actually a local one.

Proposition 4.7 *A strategy* $\mathbf{p}$ *is an ESS if and only if*

$$\boldsymbol{\pi}(\mathbf{p}, \mathbf{q}) \ge \boldsymbol{\pi}(\mathbf{q}, \mathbf{q}) \, , \tag{4.4}$$

for all $\mathbf{q}$ *in a neighborhood of* $\mathbf{p}$ *in* Δ *with strict inequality if* $\mathbf{q} \neq \mathbf{p}$.

Proof: Let us assume that $\mathbf{p}$ is an ESS then we can write any $\mathbf{q}$ close to $\mathbf{p}$ as $\mathbf{q} = \epsilon \mathbf{r} + (1 - \epsilon)\mathbf{p}$. Actually one can choose $\mathbf{r}$ in the *compact* set

$$C = \{ \mathbf{r} \in \Delta \, ; \, r_i = 0 \text{ for some } i \in \Sigma_{\mathbf{p}} \} \, .$$

which consist of the faces of the simplex Δ which do not contain $\mathbf{p}$.

For every $\mathbf{r} \in C$ there exists $\epsilon_{\mathbf{r}}$ such that

$$\boldsymbol{\pi}(\mathbf{p}, \epsilon \mathbf{r} + (1 - \epsilon)\mathbf{p}) > \boldsymbol{\pi}(\mathbf{r}, \epsilon \mathbf{r} + (1 - \epsilon)\mathbf{p}) \, , \tag{4.5}$$

for all $0 < \epsilon < \epsilon_{\mathbf{r}}$. it is easy to see that we can choose $\epsilon_{\mathbf{r}}$ as a continuous function of $\mathbf{r}$ and thus, since C is compact $\delta \equiv \inf_{\mathbf{r} \in C} \epsilon_{\mathbf{r}}$ exist and is positive. Then for $0 < \epsilon < \delta$ we have

$$\begin{aligned} &(1 - \epsilon)\boldsymbol{\pi}(\mathbf{p}, \epsilon \mathbf{r} + (1 - \epsilon)\mathbf{p}) + \epsilon \boldsymbol{\pi}(\mathbf{p}, \epsilon \mathbf{r} + (1 - \epsilon)\mathbf{p}) \\ &> (1 - \epsilon)\boldsymbol{\pi}(\mathbf{p}, \epsilon \mathbf{r} + (1 - \epsilon)\mathbf{p}) + \epsilon \boldsymbol{\pi}(\mathbf{r}, \epsilon \mathbf{r} + (1 - \epsilon)\mathbf{p}) \, , \end{aligned} \tag{4.6}$$

i.e., we have (4.4) with $\mathbf{q} = \epsilon \mathbf{r} + (1 - \epsilon)\mathbf{p}$.

Conversely if (4.4) holds then we can rewrite it as (4.6) and this shows that $\mathbf{p}$ is an ESS. ∎

In special the case of potential games, the ESS concept has a nice and simple interpretation. For a symmetric potential game we can assume, by rescaling the payoff if necessary, that $\boldsymbol{\pi}$ is a self-adjoint matrix, i.e., $\boldsymbol{\pi} = \boldsymbol{\pi}^T$.

Proposition 4.8 *For symmetric potential games with payoff matrix $\boldsymbol{\pi} = \boldsymbol{\pi}^T$ a strategy $\mathbf{p}$ is an ESS if and only if $\mathbf{p}$ is a local maximum of the potential function $V(p) = \frac{1}{2}\boldsymbol{\pi}(\mathbf{p}, \mathbf{p})$.*

Proof: Let $\mathbf{r} = \frac{1}{2}\mathbf{p} + \frac{1}{2}\mathbf{q}$ then we have by multilinearity of the payoff

$$\boldsymbol{\pi}(\mathbf{q}, \mathbf{q}) \;=\; \boldsymbol{\pi}(\mathbf{p}, \mathbf{p}) - 2\boldsymbol{\pi}(\mathbf{p}, \mathbf{r}) - 2\boldsymbol{\pi}(\mathbf{r}, \mathbf{p}) + 4\boldsymbol{\pi}(\mathbf{r}, \mathbf{r})$$

For potential games we have $\boldsymbol{\pi}(\mathbf{p}, \mathbf{q}) = \boldsymbol{\pi}(\mathbf{p}, \mathbf{q})$ for all $\mathbf{p}, \mathbf{q}$ and so we find

$$\boldsymbol{\pi}(\mathbf{p}, \mathbf{p}) - \boldsymbol{\pi}(\mathbf{q}, \mathbf{q}) \;=\; 4\left[\boldsymbol{\pi}(\mathbf{p}, \mathbf{r}) - \boldsymbol{\pi}(\mathbf{r}, \mathbf{r})\right] .$$

Note next that if $\mathbf{p}$ and $\mathbf{q}$ are at distance ϵ then $\mathbf{p}$ and $\mathbf{r}$ are at distance $\epsilon/2$. We can now conclude the proof by invoking Proposition 4.7. ∎

4.2 The replicator equation for one population games

In this section we derive and analyze a set of differential equations which describe the selection mechanism over time in a population which has different types.

Let us assume that the population comes in n different types and we let $p_i(t)$ to be the proportion of type i at time t and we write $\mathbf{p} = (p_1, \cdots, p_n) \in \Delta$ for the state of the population.

Let us assume that n functions $f_i = f_i(\mathbf{p})$ are given which describe the fitness of the type i in a population state $\mathbf{p}$. We derive a differential equation for p_i following the basic tenet of Darwinism which is that the reproductive success of a given individual of a certain type should be proportional to the difference between the fitness $f_i(\mathbf{p})$ of type i and the average fitness $\bar{f}(\mathbf{p}) = \sum_{i=1}^n p_i f_i(\mathbf{p})$. So we obtain

$$\frac{dp_i/dt}{p_i} \;=\; \text{fitness of type } i - \text{ average fitness}$$

and this yields the *replicator equation*

$$\frac{dp_i}{dt} = p_i \left(f_i(\mathbf{p}) - \bar{f}(\mathbf{p}) \right) , \quad i = 1, \cdots , n \,. \tag{4.7}$$

Under the usual smoothness assumption on $f_i(\mathbf{p})$ (e.g., if f_i is locally Lipschitz continuous) the ordinary differential equation 4.7 with initial condition $\mathbf{p}^(0) = \mathbf{p}(0)$ has a unique solution (for sufficiently small time t) which we denote by $\mathbf{p}(t)$.

This equation has the following elementary properties

Lemma 4.9 *1. The simplex Δ is invariant under the dynamics (4.7). In particular any solution of (4.7) with $\mathbf{p}(0) \in \Delta$ exists for all $t \in \mathbb{R}$.*

2. The support of a strategy Σ_p is left invariant under the dynamics. In particular every face of the simplex Δ is invariant under the dynamics (4.7).

Proof: 1. Let $K = \sum_{i=1}^n p_i$, then we have

$$\frac{dK}{dt} = (1-K)\bar{f}(\mathbf{p}) \,.$$

and thus if $K(0) = \sum_i p_i(0) = 1$ then $\sum_i p_i(t) = 1$ for all t.

2. Note that if $p_i(t) = 0$ then $\frac{dp_i}{dt} = 0$ and thus $p_i(t) = 0$ for all t. This implies that if $p_i(t) > 0$ for some t then $p_i(t) > 0$ for all t. ∎

By Lemma 4.9, item 1. we can and will restrict the equation (4.7) to the simplex Δ. Furthermore item 2. means that if a certain type is not present in the population the dynamics will not introduce it. So the replicator dynamics does not include mutation but only selection mechanisms. If certain types are absent from the population then the population stays on the corresponding face of the simplex.

Of particular interest is the case of linear $f_i(\mathbf{p})$ which has a direct game theoretic interpretation. If $\boldsymbol{\pi}$ is the payoff matrix for a symmetric game then the fitness of the type i is identified as the the payoff for strategy i against a population $\mathbf{p}$, i.e., we have

$$f_i(\mathbf{p}) - \sum_i p_i f_i(\mathbf{p}) = \boldsymbol{\pi}(\mathbf{e}_i, \mathbf{p}) - \boldsymbol{\pi}(\mathbf{p}, \mathbf{p}) = (\boldsymbol{\pi}\mathbf{p})_i - \langle \mathbf{p}, \boldsymbol{\pi}\mathbf{p} \rangle \,.$$

and the replicator equation is

$$\frac{dp_i}{dt} = p_i \left[\boldsymbol{\pi}(\mathbf{e}_i, \mathbf{p}) - \boldsymbol{\pi}(\mathbf{p}, \mathbf{p}) \right] , \quad i = 1, \cdots, n \, . \tag{4.8}$$

If a certain strategy is absent at initial times, for example let us assume that $p_1(0) = 0$, then by Lemma 4.9 the trajectories of the dynamics do not leave the face of the simplex $\{\mathbf{p} \in \Delta, p_1 = 0\}$. Let us set $\tilde{\mathbf{p}} = (p_2, \cdots, p_n)$ and let $\tilde{\boldsymbol{\pi}}$ the $(n-1) \times (n-1)$ matrix obtained from $\boldsymbol{\pi}$ be deleting its first row and first column. Then we have

$$\frac{dp_i}{dt} = p_i \left[\tilde{\boldsymbol{\pi}}(\tilde{\mathbf{e}}_i, \tilde{\mathbf{p}}) - \tilde{\boldsymbol{\pi}}(\tilde{\mathbf{p}}, \tilde{\mathbf{p}}) \right] , \quad i = 2, \cdots, n \, . \tag{4.9}$$

and so the dynamics on a face of the simplex is simply the dynamics for the reduced game with strategies $\tilde{S} = \{2, \cdots, n\}$ and payoff matrix $\tilde{\boldsymbol{\pi}}$.

Under the linear rescaling of the payoff defined in section the replicator equation is left essentially unchanged, up to a rescaling of time.

Lemma 4.10 *Let* $\mathbf{p}$ *and* $\mathbf{p}'$ *be the solution of the replicator equation with payoff matrix* $\pi(i,j)$ *and* $\pi' = a\pi(i,j) + b(j)$ *and same initial conditions* $\mathbf{p}'(0) = \mathbf{p}(0)$*, then* $\mathbf{p}'(t) = \mathbf{p}(at)$*.*

Proof: Since $\boldsymbol{\pi}'(\mathbf{e}_i, \mathbf{p})_i - \boldsymbol{\pi}'(\mathbf{p}, \mathbf{p}) = a \left[\boldsymbol{\pi}(\mathbf{e}_i, \mathbf{p}) - \boldsymbol{\pi}(\mathbf{p}, \mathbf{p}) \right]$ and thus $dp_i'/dt = adp_i/dt$. This is equivalent to $\mathbf{p}'(t) = \mathbf{p}(at)$. ∎

Let us consider some examples of replicator dyanmics

Example 4.11 Two strategy games. By the invariance under linear rescaling of the payoff we can assume that

$$\boldsymbol{\pi} = \begin{pmatrix} a_1 & 0 \\ 0 & a_2 \end{pmatrix}$$

and the replicator equation is given by

$$\begin{aligned} \frac{dp_1}{dt} &= p_1(a_1 p_1 - a_1 p_1^2 - a_2 p_2^2) \\ \frac{dp_2}{dt} &= p_2(a_2 p_2 - a_1 p_1^2 - a_2 p_2^2) \end{aligned}$$

Since $p_2 = 1 - p_1$ we can rewrite the two equation in (4.10) as a single one equation for $p = p_1$ on the interval $[0, 1]$ and we obtain

$$\frac{dp}{dt} = p(1-p)[(a_1 + a_2)p - a_2] \tag{4.10}$$

(a) $a_1 > 0$, $a_2 > 0$. In the coordination game we have $dp/dt > 0$ if $p > a_2/(a_1 + a_2)$ and $dp/dt < 0$ if $p < a_2/(a_1 + a_2)$. Unless $\mathbf{p}(t)$ starts at the mixed NE, $\mathbf{p}(t)$ converges to one the pure NE (and ESS) depending on the initial $\mathbf{p}$.

(b) $a_1 < 0$, $a_2 < 0$. In a Hawk and dove game for every initial condition (except $p(0) = 0$ and $p(0) = 1$ which are stationary points) converge to the ESS.

(b) $a_1 < 0$, $a_2 > 0$ or $a_1 > 0$, $a_2 < 0$. Then $dp/dt < 0$ or > 0 on $[0, 1]$ and $p(t)$ converges to the unique NE (and ESS) unless $p(0) = 0$ or $p(0) = 1$.

Example 4.12 Let us consider the generalized Rock-Paper-Scissor with payoff matrix

$$\boldsymbol{\pi} = \begin{pmatrix} 0 & 1+a & -1 \\ -1 & 0 & 1+a \\ 1+a & -1 & 0 \end{pmatrix}.$$

The replicator equation is the system

$$\begin{aligned} \frac{dp_1}{dt} &= p_1\left[(p_2 - p_3) + ap_2 - \langle \mathbf{p}, \boldsymbol{\pi}\mathbf{p}\rangle\right] \\ \frac{dp_2}{dt} &= p_2\left[(p_3 - p_1) + ap_3 - \langle \mathbf{p}, \boldsymbol{\pi}\mathbf{p}\rangle\right] \\ \frac{dp_3}{dt} &= p_3\left[(p_1 - p_2) + ap_1 - \langle \mathbf{p}, \boldsymbol{\pi}\mathbf{p}\rangle\right] \end{aligned}$$

For any a, $\mathbf{p}^* = (1/3, 1/3, 1/3)$ is the unique NE in intΔ and the unique NE for $a > -1$. Let us consider the Liapunov function

$$V(\mathbf{p}) = \sum_{i=1}^{3} \log p_i \,,$$

(the choice of V will be expanded upon later on). Note that $\mathbf{p}$ is a concave function on intΔ with a unique maximum at $(1/3, 1/3, 1/3)$ and that $V(p)$ approaches $-\infty$ as $\mathbf{p}$ approaches the boundary of Δ.

Let us compute the derivative of V along the trajectories of the replicator equation We have

$$\begin{aligned} \frac{d}{dt}V(\mathbf{p}(t)) &= \sum_{i=1}^{3} \frac{1}{p_i}\frac{dp_i}{dt} \\ &= a(p_1+p_2+p_3) - 3a(p_1p_2+p_2p_3+p_3p_1)\,. \end{aligned} \tag{4.11}$$

Using the identity $1 = (p_1+p_2+p_3)^2 = \langle \mathbf{p}, \mathbf{p}\rangle + 2(p_1p_2+p_2p_3+p_3p_1)$ we have

$$\frac{d}{dt}V(\mathbf{p}) = \frac{a}{2}(3\langle \mathbf{p}, \mathbf{p}\rangle - 1)\,.$$

Note that $\langle \mathbf{p}, \mathbf{p}\rangle$ has its maximum (equal to 1) on the simplex at the pure strategies $\mathbf{p} = \mathbf{e}_i$ and its minimum equal to $1/2$ at the Nash equilibrium $\mathbf{p}^* = (1/3, 1/3, 1/3)$.

If $a > 0$ we have $dV/dt \geq 0$ with equality iff $\mathbf{p}(t) = \mathbf{p}^*$ and thus every trajectory starting in $\mathrm{int}(\Delta)$ converges to the ESS $(1/3, 1/3, 1/3)$. If $a = 0$ the trajectories are periodic orbits on the closed curves $\{V = const\}$ and if $a < 0$ then every trajectory moves away from $\mathbf{p}^*$. When $a < -1$ the dynamical behavior changes and reflects that there are 6 more NE than for $a > -1$.

Next we investigate how strictly dominated strategies behave under the replicator dynamics. One would expect that a "good" dynamics does eliminate the strictly dominated strategies. This holds true for the replicator dynamics if one starts in the interior of Δ.

Proposition 4.13 *Suppose that the strategy $i \in S$ is strictly dominated. If $\mathbf{p}(0) \in \mathrm{int}\Delta$ then $\lim_{t\to\infty} p_i(t) = 0$.*

Proof: : Suppose that the pure strategy $\mathbf{e}_i$ is dominated by the strategy $\mathbf{q}$ then $\pi(i,j) < \sum_i q_i \pi(i,j)$ for all j and so by compactness of Δ we have

$$\delta \equiv \inf_{p\in\Delta} [\boldsymbol{\pi}(\mathbf{q}, \mathbf{p}) - \boldsymbol{\pi}(\mathbf{e}_i, \mathbf{p})] > 0\,.$$

Let us consider the function

$$V_i(\mathbf{p}) = \ln(p_i) - \sum_{j=1}^{n} q_j \log(p_j).$$

Note that $V_i(p)$ approaches $-\infty$ as p_i approaches 0 and approaches $+\infty$ whenever p_j approaches 0 for some $j \in \Sigma_q$. Let us differentiate along the trajectories. We have

$$\frac{d}{dt}V_i(\mathbf{p}(t)) = \frac{1}{p_i}\frac{dp_i}{dt} - \sum_j q_j \frac{1}{p_j}\frac{dp_j}{dt} = [\boldsymbol{\pi}(\mathbf{e}_i, \mathbf{p}) - \boldsymbol{\pi}(\mathbf{q}, \mathbf{p})] \le -\delta < 0\,.$$

This implies that $V(\mathbf{p}(t)) \le V(\mathbf{p}(0)) - \delta t$ and thus decreases monotonically to $-\infty$ and this is possible only if $\lim_{t\to\infty} p_i(t) = 0$. ■

Note that the assumption that all strategies have positive probability initially is necessary. For example if the pure strategy $\mathbf{e}_i$ is strictly dominated by another pure strategy $\mathbf{e}_j$ but $p_j(0) = 0$ then since the replicator dynamics does not introduce absent strategies there is no reason why $p_i(t)$ should decrease in general.

We turn next to the relation between the Nash equilibrium of the game wih payoff matrix $\boldsymbol{\pi}$ and the stationary solutions of the replicator equation (4.8)

Proposition 4.14 *We have*

1. *Any NE is a stationary solution of (4.8).*

2. *The set of stationary solutions for (4.8) in* $\mathrm{int}\Delta$ *is a convex set and any stationary solution of (4.8) in* $\mathrm{int}\Delta$ *is a NE.*

3. *If* $\mathbf{p}^*$ *is a stationary solution for (4.8) if and only if* $\mathbf{p}^*$ *is a NE for some subgame* $(\tilde{S}, \tilde{\boldsymbol{\pi}})$ *of* $(S, \boldsymbol{\pi})$. *In particular all extreme point of the simplex are stationary solutions.*

Proof: : If $\mathbf{p}^*$ is a stationary solution for (4.8) then $p_i = 0$ or $\boldsymbol{\pi}(\mathbf{e}_i, \mathbf{p}) = \boldsymbol{\pi}(\mathbf{p}, \mathbf{p})$. Therefore for a stationary solution all pure strategies present in the population must earn the same payoff against this population. This is true of all NE of the game as well as for all the NE of all the subgames. This proves 1. and 3. If $\mathbf{p}^* \in \mathrm{int}\Delta$ then the condition for stationarity is equivalent to the NE condition. Finally let us suppose that $\mathbf{p}^*$ and $\mathbf{q}^*$ are stationary solutions in $\mathrm{int}\Delta$ then for $0 \le \alpha \le 1$ and $\mathbf{r} = \alpha\mathbf{p} + (1-\alpha)\mathbf{q}$ we have using the multilinearity of the payoff and the stationarity condition

$$\boldsymbol{\pi}(\mathbf{e}_i, \mathbf{r}) = \alpha\boldsymbol{\pi}(\mathbf{e}_i, \mathbf{p}) + (1-\alpha)\boldsymbol{\pi}(\mathbf{e}_i, \mathbf{q})$$

Since by stationarity all the pure strategies earn the same payoff against $\mathbf{p}$ and $\mathbf{q}$ the r.h.s is independent of i thus so is the l.h.s. and this means that all pure strategies earn the same payoff against $\mathbf{r}$. This proves 2. ■

We show next that even though there are stationary solutions which are not NE of the game they are not stable. Recall that for an ODE $dx/dt = f(x)$ with solution $x(t)$, a stationary point x^* is *stable* (or *Liapunov stable*) if for any (arbitrary small) neighborhood B of x^* there exists a neighborhood C of x^* such that $x(0) \in C$ implies that $x(t) \in B$ for all $t > 0$.

We have

Proposition 4.15 *Suppose the stationary solution* $\mathbf{p}^*$ *of (4.8) is stable then* $\mathbf{p}^*$ *is a NE.*

Proof: Suppose $\mathbf{p}^*$ is stationary but not a NE. Then for all $i \in \Sigma_{\mathbf{p}^*}$ all payoff $\boldsymbol{\pi}(\mathbf{e}_i, \mathbf{p}^*)$ are equal but there exists $j \notin \Sigma_{\mathbf{p}^*}$ such that $\boldsymbol{\pi}(\mathbf{e}_j, \mathbf{p}^*) > \boldsymbol{\pi}(\mathbf{p}^*, \mathbf{p}^*)$. By continuity there exists a neighborhood B of $\mathbf{p}^*$ such that $\boldsymbol{\pi}(\mathbf{e}_j, \mathbf{q}) - \boldsymbol{\pi}(\mathbf{q}, \mathbf{q}) > \delta$ for all $\mathbf{q} \in B \cap \Delta$. If $\mathbf{p}(0) \in B$ with $p_j(0) \neq 0$ then as long as $\mathbf{p}(t) \in B$ we have

$$\frac{dp_j}{dt} = p_j \left[\boldsymbol{\pi}(\mathbf{e}_j, \mathbf{p}) - \boldsymbol{\pi}(\mathbf{p}, \boldsymbol{\pi})\right] \geq \delta p_j \,,$$

and thus as long as $\mathbf{p}(t) \in B$ we have $p_j(t) > e^{\delta t} p_j(0)$. This shows that $\mathbf{p}^*$ is not stable. ■

Next we show that if $\mathbf{p}(t)$ converges as $t \to \infty$ then the limiting point must be a NE, provided we start with all strategy present.

Proposition 4.16 *Let* $\mathbf{p}(t)$ *be a solution of (4.8) with* $\mathbf{p}(0) \in \mathrm{int}\Delta$. *If* $\lim_{t\to\infty} \mathbf{p}(t) = \mathbf{p}^*$ *then* $\mathbf{p}^*$ *is a NE.*

Proof: Suppose $\mathbf{p}^*$ is not a NE, then there exists i such that $\boldsymbol{\pi}(\mathbf{e}_i, \mathbf{p}^*) - \boldsymbol{\pi}(\mathbf{p}^*, \mathbf{p}^*) = \epsilon > 0$. Since $\boldsymbol{\pi}$ is continuous and $\mathbf{p}(t) \to \mathbf{p}^*$ there exists T such that for $t > T$ we have

$$\frac{dp_i}{dt} = p_i \left[\boldsymbol{\pi}(\mathbf{e}_i, \mathbf{p}(t)) - \boldsymbol{\pi}(\mathbf{p}(t), \mathbf{p}(t))\right] \geq \frac{\epsilon}{2} p_i \,, \quad t \geq T \,,$$

and thus $p_i(t) \geq p_i(T) e^{(t-T)\epsilon/2}$ for all $t \geq T$. This is a contradiction. Hence $\mathbf{p}^*$ is a NE. ■

Next we discuss the properties of an evolutionary stable strategies $\mathbf{p}^*$ for the replicator dynamics. To this effect we consider the relative entropy defined by

$$H(\mathbf{p}^* \,|\, \mathbf{p}) \;=\; \sum_{i \in \Sigma_{\mathbf{p}}} p_i^* \ln \left(\frac{p_i^*}{p_i} \right) \,, \quad \text{if } \Sigma_{\mathbf{p}^*} \subset \Sigma_{\mathbf{p}} \,,$$

and we define

$$B_{\mathbf{p}^*} \;=\; \{ \mathbf{p} \in \Delta \,,\, \Sigma_{\mathbf{p}^*} \subset \Sigma_{\mathbf{p}} \}$$

Note that $B_{\mathbf{p}^*}$ is a neighborhood of $\mathbf{p}^*$ relative to $\mathbf{p}$ and is the domain of the relative entropy function. Moreover $H(\mathbf{p}^* \,|\, \mathbf{p})$ is convex in $\mathbf{p}$ and its minimum is attained only $\mathbf{p} = \mathbf{p}^*$.

Using this we prove

Theorem 4.17 *Let* $\mathbf{p}^*$ *be an ESS, then* $\mathbf{p}^*$ *is asymptotically stable.*

Proof: We use the Liapunov function $V(\mathbf{p}) = H(\mathbf{p}^* \,|\, \mathbf{p})$ and note that

$$\begin{aligned} \frac{dV}{dt}(\mathbf{p}(t)) &= -\sum_{i\Sigma_p} p_i^* \frac{1}{p_i} \frac{dp_i}{dt} \\ &= -\sum_{i\Sigma_p} p_i^* \left[\boldsymbol{\pi}(\mathbf{e}_i, \mathbf{p}) - \boldsymbol{\pi}(\mathbf{p}, \mathbf{p}) \right] \\ &= -\left[\boldsymbol{\pi}(\mathbf{p}^*, \mathbf{p}) - \boldsymbol{\pi}(\mathbf{p}, \mathbf{p}) \right] \end{aligned} \tag{4.12}$$

By the characterization of ESS in Proposition (4.7) we obtain that $dV/dt \leq 0$ in a neighborhood of $\mathbf{p}^*$ with equality iff $\mathbf{p}^* = \mathbf{p}$ and this implies asymptotic stability. ∎

If the ESS is in the interior of the simplex we can prove a more global result

Theorem 4.18 *Suppose* $\mathbf{p}^* \in \text{int}\Delta$ *is an ESS. Then for any* $\mathbf{p}(0) \in \text{int}\Delta$ *we have* $\lim_{t \to \infty} \mathbf{p}(t) = \mathbf{p}^*$.

Proof: Note that if $\mathbf{p}^* \in \text{int}\Delta$ then $B_{\mathbf{p}}^* = \text{int}\Delta$. Moreover since $\mathbf{p}^*$ is an ESS (and thus a NE) all strategies $\mathbf{q} \in \text{int}\Delta$ are best response to $\mathbf{p}^*$. So by the second-order condition for ESS (see Corollary 4.3) we have $\boldsymbol{\pi}(\mathbf{p}^*, \mathbf{p}) > \boldsymbol{\pi}(\mathbf{p}, \mathbf{p})$ for all $\mathbf{p} \in \Delta$. This implies however that $dV/dt \leq 0$ everywhere in intΔ with equality only at $\mathbf{p}^*$. This implies global asymptotic stability. ∎

For the special case of potential games, the potential function $V(\mathbf{p})$ naturally plays an important role.

Theorem 4.19 *For a potential game the potential* $V(\mathbf{p}) = \frac{1}{2}\boldsymbol{\pi}(\mathbf{p},\mathbf{p})$ *is a strict Liapunov function, i.e., if* $\mathbf{p}(t)$ *is a solution of the replicator equation then*

$$\frac{dV}{dt}(\mathbf{p}(t)) \geq 0$$

and

$$\frac{dV}{dt}(\mathbf{p}) = 0 \text{ iff } \mathbf{p} \text{ is a stationary solution}$$

Moreover the following are equivalent
(i) $\mathbf{p}$ *is asymptotically stable.*
(ii) $\mathbf{p}$ *is an ESS.*
(iii) $\mathbf{p}$ *is a local maximum of* V*.*

Proof: We have

$$\begin{aligned} \frac{dV}{dt}(\mathbf{p}(t)) &= \sum_i \frac{\partial V}{\partial p_i}\frac{dp_i}{dt} \\ &= \sum_i \boldsymbol{\pi}(\mathbf{e}_i,\mathbf{p}) p_i \left[\boldsymbol{\pi}(\mathbf{e}_i,\mathbf{p}) - \boldsymbol{\pi}(\mathbf{p},\mathbf{p})\right] . \end{aligned} \tag{4.13}$$

Note that we have the identity

$$0 = \sum_i \boldsymbol{\pi}(\mathbf{p},\mathbf{p}) p_i \left[\boldsymbol{\pi}(\mathbf{e}_i,\mathbf{p}) - \boldsymbol{\pi}(\mathbf{p},\mathbf{p})\right] .$$

and thus we have

$$\frac{dV}{dt}(\mathbf{p}(t)) = \sum_i p_i \left[\boldsymbol{\pi}(\mathbf{e}_i,\mathbf{p}) - \boldsymbol{\pi}(\mathbf{p},\mathbf{p})\right]^2 \leq 0\,. \tag{4.14}$$

We have $\frac{dV}{dt}(\mathbf{p}(t)) = 0$ if and only if, for every i either $p_i = 0$ or $\boldsymbol{\pi}(\mathbf{e}_i,\mathbf{p}) = \boldsymbol{\pi}(\mathbf{p},\mathbf{p})$. This is exactly the condition for $\mathbf{p}$ to be a stationary solution of the replicator equation. This implies that V increases along every solution $\mathbf{p}(t)$ and is constant only at the critical point of the replicator equation. Moreover if V has a local maxima at $\mathbf{p}^*$ if and only $\mathbf{p}^*$ is asymptotically stable and this proves the equivalence of (i) and (iii). The equivalence of (ii) and (iii) has been proved in Proposition 4.8.