

8 Mixed Strategies

	Rock	Scissors	Paper
Rock	0,0	1,-1	-1,1
Scissors	1,-1	0,0	1,-1
Paper	1,-1	-1,1	0,0

Figure 1: A Game with no Pure Strategy Nash Equilibrium

Clearly, there is no pure strategy equilibrium in “Rock, Scissors, and Paper”

Definition 1 Let $G = (n, S, u)$ be a finite normal form game. A mixed strategy for player $i \in I$, denoted σ_i , is probability distribution over the set of pure strategies S_i .

We let the set of mixtures of S_i be denoted $\Delta(S_i)$.

Remark 1 Pure strategies are corners of the mixed strategy simplex

We extend the pure strategy utility functions $u : S \rightarrow R^n$ by assuming that all agents evaluate lotteries according to a von Neumann-Morgenstern utility function, that is $u_i : \times_{j=1}^n \Delta(S_j) \rightarrow R$ is given by

$$u_i(\sigma_1, \dots, \sigma_n) = \sum_{s_1 \in S_1} \dots \sum_{s_n \in S_n} u_i(s_1, \dots, s_n) \sigma_1(s_1) \times \dots \times \sigma_n(s_n)$$

or in short

$$u_i(\sigma) = \sum_{s \in S} u_i(s) \prod_{j=1}^n \sigma_j(s_j).$$

We note that playing pure strategy $s_i \in S_i$ against σ_{-i} gives payoff

$$u_i(s_i, \sigma_{-i}) = \sum_{s_{-i} \in S_{-i}} u_i(s_i, s_{-i}) \prod_{j \neq i} \sigma_j(s_j)$$

So

$$\begin{aligned} u_i(\sigma_i, \sigma_{-i}) &= \sum_{s_i \in S_i} \sum_{s_{-i} \in S_{-i}} u_i(s_i, s_{-i}) [\prod_{j \neq i} \sigma_j(s_j)] \sigma_i(s_i) \\ &= \sum_{s_i \in S_i} u_i(s_i, \sigma_{-i}) \sigma_i(s_i) \end{aligned}$$

PAYOFF IS LINEAR IN OWN STRATEGY, which has some very important consequences.

Definition 2 A Nash equilibrium in G is a profile $\sigma^* = (\sigma_1^*, \dots, \sigma_n^*) \in \times_{j=1}^n \Delta(S_j)$ such that

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(\sigma_i, \sigma_{-i}^*) \text{ for every } i \in I \text{ and } \sigma_i \in \Delta(S_i)$$

Remark 2 Using linearity we have that σ^* is a Nash equilibrium if and only if

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(s_i, \sigma_{-i}^*) \text{ for every } i \in I \text{ and } s_i \in S_i.$$

Proof. (Only if) Trivial as $S_i \subset \Delta(S_i)$. (If) Suppose that $u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(s_i, \sigma_{-i}^*)$ for every $i \in I$ and $s_i \in S_i$ and assume that σ^* is not Nash. Then, there exists some player $i \in I$ and a strategy $\sigma'_i \in \Delta(S_i)$ such that

$$u_i(\sigma'_i, \sigma_{-i}^*) > u_i(\sigma_i^*, \sigma_{-i}^*)$$

But

$$\begin{aligned} u_i(\sigma'_i, \sigma_{-i}^*) &= \sum_{s_i \in S_i} u_i(s_i, \sigma_{-i}^*) \sigma'_i(s_i) \\ &\leq \max_{s_i \in S_i} u_i(s_i, \sigma_{-i}^*) \equiv u_i(\bar{s}_i, \sigma_{-i}^*) \end{aligned}$$

Hence, there exists some $\bar{s}_i \in S_i$ such that

$$u_i(\bar{s}_i, \sigma_{-i}^*) \geq u_i(\sigma'_i, \sigma_{-i}^*) > u_i(\sigma_i^*, \sigma_{-i}^*),$$

contradicting that $u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(s_i, \sigma_{-i}^*)$ for every $i \in I$ and $s_i \in S_i$. ■

Remark 3 The notion of a Nash equilibrium is often taken as a minimal condition for “reasonable” play. However, when randomizations are introduced one may sometimes argue

that it may make sense to assume that players can correlate the randomizations (for example by conditioning on sunspots). Such an assumption gives us an alternative equilibrium concept which is called correlated equilibrium, which is a probability distribution $\sigma \in \Delta(S)$ such that

$$\sum_{s \in S} \sigma(s) u_i(s_i, s_{-i}) \geq \sum_{s \in S} \sigma(s) u_i(f(s_i), s_{-i}),$$

for every $f : S_i \rightarrow S_i$. The set of Nash equilibria are contained in the set of correlated equilibria. Many results in the repeated game literature relies on correlated randomizations, which are handy because this convexifies the payoffs.

8.0.1 Example

Let's consider an asymmetric version of the battle of the sexes

$$\begin{array}{cc} & O & F \\ o & 1, 2 & 0, 0 \\ f & 0, \frac{1}{2} & 2, 1 \end{array}$$

where the column player actually gets something out of opera even without the other. Let $p = \Pr[o]$ and $q = \Pr[O]$. Hence

$$\begin{aligned} u_1(o, q) &= q \\ u_1(f, q) &= 2(1 - q) \end{aligned}$$

Hence

$$\beta_1(q) = \begin{cases} \{1\} & \text{if } q > \frac{2}{3} \\ [0, 1] & \text{if } q = \frac{2}{3} \\ \{0\} & \text{if } q < \frac{2}{3} \end{cases}$$

and

$$\begin{aligned} u_2(p, O) &= 2p + \frac{1}{2}(1 - p) \\ u_2(p, F) &= (1 - p) \end{aligned}$$

so

$$\beta_2(p) = \begin{cases} \{1\} & \text{if } p > \frac{1}{5} \\ [0, 1] & \text{if } p = \frac{1}{5} \\ \{0\} & \text{if } p < \frac{1}{5} \end{cases}$$

DRAW!

8.1 Domination and Rationalizability

DRAW

	L	R
T	3, ·	0, ·
M	2, ·	2, ·
B	0, ·	3, ·

It cannot be an equilibrium to put positive probability on both T and B as mixing requires that the player is indifferent between T and B ,

$$\begin{aligned} 3p + 0(1-p) &= 0 + 3(1-p) \\ &\Leftrightarrow \\ p &= \frac{1}{2} \end{aligned}$$

in which case the payoff from either T or B is $1.5 < 2$. We can think of this as saying that there are no beliefs about other players strategy that makes non-degenerate mixtures of T and B optimal. We'll say that mixtures of T and B are not *rationalizable*.

DRAW

	L	R
T	3, ·	0, ·
M	1, ·	1, ·
B	0, ·	3, ·

Here M is never a best response.

Formally:

Definition 3 Let $\tilde{\Delta}^k(S_i) = \Delta(S_i)$ and define recursively

$$\tilde{\Delta}^k(S_i) = \left\{ \begin{array}{l} \sigma_i \in \tilde{\Delta}^{k-1}(S_i) \text{ s.t. there is } \sigma_{-i} \in \times_{j \neq i} CO \left[\tilde{\Delta}^{k-1}(S_j) \right] \\ \text{s.t. } u_i(\sigma_i, \sigma_{-i}) \geq u_i(\sigma'_i, \sigma_{-i}) \forall \sigma'_i \in \tilde{\Delta}^{k-1}(S_i) \end{array} \right\}.$$

σ_i is rationalizable if $\sigma_i \in \cap_{k=0}^{\infty} \tilde{\Delta}^k(S_i)$.

- Iteration 1-get rid of strategies that are never a best response.
- Iteration 2-get rid of strategies that are never a best response against strategies surviving first iteration etc.
- Convex hull used because a mixture of σ_j and σ'_j need not be in $\tilde{\Delta}^{k-1}(S_j)$ even if both mixtures are (illustrated in example above).

Theorem 1 In two player games, rationalizability and iterated elimination of strictly dominated strategies are equivalent.

To get idea, we'll prove:

Proposition 1 In a two player game, $s'_1 \in S_1$ is not strictly dominated by any (pure or mixed) strategy $\sigma_1 \in \Delta(S_1)$ if and only if there exists some $\sigma_2 \in \Delta(S_2)$ such that s'_1 is a best reply to σ_2 .

Proof. If there exists σ_2 so that $s'_1 \in S_1$ best response it is obvious that s'_1 is not strictly dominated. For the other direction, suppose that s'_1 is not strictly dominated. We want to show that there exists $\sigma_2 \in \Delta(S_2)$ such that s'_1 is a best response to σ_2 . Define

$$X = \left\{ x \in R^{|S_2|} \mid \text{exists } \sigma_1 \in \Delta(S_1) \text{ s.t. } x = u_1(\sigma_1, \vec{s}_2) = \begin{pmatrix} u_1(\sigma_1, s_2^1) \\ \dots \\ u_1(\sigma_1, s_2^k) \\ \dots \\ u_1(\sigma_1, s_2^{|S_2|}) \end{pmatrix} \right\},$$

which is a convex set. To see this, take any $x', x'' \in X$ and let σ'_1, σ''_1 be randomizations that $x' = u_1(\sigma'_1, \vec{s}_2)$ and $x'' = u_1(\sigma''_1, \vec{s}_2)$ consider a convex combination $x^\lambda = \lambda x' + (1 - \lambda)x''$. Then, the randomization $\sigma_1^\lambda = \lambda \sigma'_1 + (1 - \lambda)\sigma''_1$ will be such that $x^\lambda = u_1(\sigma_1^\lambda, \vec{s}_2)$.

Consider the set

$$C = X - \{u_1(s'_1, \vec{s}_2)\}$$

which is convex since linear combinations of convex sets are convex. We note that $\vec{0} \in C$, since $\tilde{\sigma}_1$ with $\tilde{\sigma}_1(s'_1) = 1$ and $\tilde{\sigma}_1(s_1) = 0$ for all $s_1 \neq s'_1$. We also note that there exists no $x \in C$ s.t. $x_k > 0$ for every coordinate k since that would make s'_1 strictly dominated. The set $R_{++}^{|S_2|}$ is obviously convex and since

$$C \cap R_{++}^{|S_2|}$$

there exists a hyperplane separating C and $R_{++}^{|S_2|}$, that is there exists some $p \in R^{|S_2|}$ such that

$$\begin{aligned} px &\geq p \vec{0} = 0 \text{ for all } x \in R_{++}^{|S_2|} \\ px &\leq 0 \text{ for all } x \in C. \end{aligned}$$

The only possibility for $px = 0$ to hold for all $x \in R_{++}^{|S_2|}$ is if $p \geq 0$. Hence, we may use the normalized vector

$$\sigma_2 = \left(\frac{p^1}{\sum p^k}, \dots, \frac{p^{|S_2|}}{\sum p^k} \right) \in \Delta(S_2),$$

for which we conclude that

$$\begin{aligned} \sigma_2 c &= \sigma_2 [x - u_1(s'_1, \vec{s}_2)] \\ &= \sigma_2 [u_1(\sigma_1, \vec{s}_2) - u_1(s'_1, \vec{s}_2)] \leq 0 \end{aligned}$$

for every $\sigma_1 \in \Delta(S_1)$. We conclude that s'_1 is a best reply to σ_2 . ■

Remark 4 *To prove that Rationalizability and iterated elimination of strictly dominated strategies coincide in two player games. Let $S^n = (S_1^n, S_2^n)$ denote the strategies that remain after n rounds of elimination. Let $\Delta(S_j^n)$ denote the corresponding mixed strategies. Also,*

let $\tilde{\Delta}^n(S_j)$ denote the strategies that survive n rounds in the definition of rationalizability.

Obviously

$$\Delta(S_j^0) = \tilde{\Delta}^0(S_j).$$

Now suppose that $\Delta(S_j^n) = \tilde{\Delta}^n(S_j)$. Change the proof as follows:

1. Redefine X by using only $s_2 \in S_2^n$ and $\sigma_1 \in \tilde{\Delta}^n(S_1)$

2. Pick $s'_1 \in S_1^{n+1}$

Proceed by induction.

8.2 Existence of Nash equilibria

Proposition 2 $\sigma^* \in NE(G) \Leftrightarrow \sigma^* \in \beta(\sigma^*)$

I.e., need existence of a fixed point. For finite games the typical fixed point theorem used.

Theorem 2 Let $F : X \rightrightarrows X$ be a correspondence and $X \subset R^n$, then there exists $x^* \in F(x^*)$ if:

1. X is non-empty, convex and compact

2. F is convex valued, non-empty valued and upper-hemi continuous .

Theorem 3 Every Finite Normal form game has at least one Nash equilibrium.