

©Martin Osborne and Ariel Rubinstein. These **lecture notes** are distributed for the exclusive use of students in Econ 501, Princeton University Fall 2001. They are basically inserts from our book *A Course of Game Theory*, MIT Press 1994 combined with some additional comments.

Lecture 14: Mixed Strategy Equilibrium
Readings: Osborne and Rubinstein Ch 3.1-2

The notion of mixed strategy NE is designed to model a steady state of a game in which the participants' choices are regulated by probabilistic rules or that the players hold none-point beliefs about their opponents. We need to enrich the model with a specification of the players preference relations over *lotteries* on the set of outcomes. We adopt the vNM assumptions. Assume that player i 's preferences are represented by the expected value of $u_i: A \rightarrow \mathfrak{R}$. The basic game model is now $\langle N, (A_i), (u_i) \rangle$.

Let $\Delta(A_i)$ be the set of probability distributions over A_i . A member of $\Delta(A_i)$ is called a **mixed strategy** of player i ; we think about the players' mixed strategies as independent randomizations. We refer to a member of A_i as a **pure strategy**.

A comment on zero sum games:

A strictly competitive game is a two player game with diagrammatically opposing preferences. Once we extend our discussion to a world with lotteries we should notice that the fact that the players's preferences over three outcomes a^1, a^2 , and a^3 are $a^1 \succ_1 a^2 \succ_1 a^3$ and $a^3 \succ_2 a^2 \succ_2 a^1$, does not imply that they have opposing preferences on the set of lotteries over these three outcomes. (both may prefer a_2 to $1/2a_1 \oplus 1/2a_3$). When we talk on zero sum games and allow mixed strategies, we will request that whenever player 1 prefer the lottery L_1 to L_2 , player 2 prefers L_2 to L_1 . This means that we can assume without loss of generality that $u_2 = -u_1$.

Once again we construct another game through which we define the equilibrium for the model.

Definition The **mixed extension** of the strategic game $\langle N, (A_i), (u_i) \rangle$ is the strategic game $\langle N, (\Delta(A_i)), (U_i) \rangle$ in which $U_i: \times_{j \in N} \Delta(A_j) \rightarrow \mathfrak{R}$ assigns to each $\alpha \in \times_{j \in N} \Delta(A_j)$ the expected value under u_i of the lottery over A that is induced by α (so that $U_i(\alpha) = \sum_{a \in A} (\prod_{j \in N} \alpha_j(a_j)) u_i(a)$ if A is finite).

Note that U_i is multilinear (for $\forall \alpha, \beta_i, \gamma_i, i$, and $\lambda \in [0, 1]$)

$$U_i(\alpha_{-i}, \lambda \beta_i + (1 - \lambda) \gamma_i) = \lambda U_i(\alpha_{-i}, \beta_i) + (1 - \lambda) U_i(\alpha_{-i}, \gamma_i).$$

When each A_i is finite $U_i(\alpha) = \sum_{a_i \in A_i} \alpha_i(a_i) U_i(\alpha_{-i}, e(a_i))$

where $e(a_i)$ is the generate lottery which assign probability 1 to a_i .

Definition A **mixed strategy NE** of a strategic game is a NE of its mixed extension.

Every NE of the game with the ordinal preferences induced from (u_i) is a Nash equilibrium in this game.

Lemma Let $G = \langle N, (A_i), (u_i) \rangle$ be a finite strategic game. Then $\alpha^* \in \times_{i \in N} \Delta(A_i)$ is a mixed strategy NE of G if and only if $\forall i \in N$ every pure strategy in the support of α_i^* is a best response to α_{-i}^* .

Proof. First, if a_i in the support of α_i^* is not a best response to α_{-i}^* then by linearity of U_i in α_i player i can increase his payoff by transferring probability from a_i to an action that is a best response; hence α_i^* is not a best response to α_{-i}^* .

Second, each action a_i^* in the support of α^* is a best response to α_{-i}^* and thus for all actions in the support $U_i(\alpha_{-i}^*, a_i^*)$ is identical and by the linearity of U_i we have $U_i(\alpha_{-i}^*, \alpha_i^*) = U_i(\alpha_{-i}^*, a_i^*) \geq U_i(\alpha_{-i}^*, a_i)$ for all a_i . By the linearity of U_i we get $U_i(\alpha_{-i}^*, \alpha_i^*) \geq U_i(\alpha_{-i}^*, \alpha_i)$ for all α_i

The lemma leads to an alternative definition of NE:

A *mixed strategy NE* of a finite strategic game is $(\alpha_i^*)_{i \in N}$ with the property that for every i every action in the support of α_i^* is a best response to α_{-i}^* .

Comment:Note that the assumption that the players' preferences can be represented by expected payoff functions plays a key role in this characterization of mixed strategy equilibrium.

Existence

Proposition Every finite strategic game $G = \langle N, (A_i), (u_i) \rangle$ has a mixed strategy NE.

Proof Denote $m_i = \#A_i$. We can identify $\Delta(A_i)$ with $\{(p_1, \dots, p_{m_i}) \mid p_k \geq 0 \forall k \text{ and } \sum_{k=1}^{m_i} p_k = 1\}$. This set is non-empty, convex, and compact in \mathfrak{R}^{m_i} . Since expected payoff is linear in the probabilities, each player's payoff function in the mixed extension of G is both quasi-concave in his own strategy and continuous. Thus the mixed extension of G satisfies all the requirements of the existence theorem in Lecture 12.

Examples

	B	S
(BoS) B	2,1	0,0
S	0,0	1,2

Suppose that (α_1, α_2) is a mixed strategy NE. If $\alpha_1(B)$ is zero or one, we obtain the two pure NE (B, B) and (S, S) . If $0 < \alpha_1(B) < 1$ then, given α_2 , player 1's actions B and S must yield the same payoff, so that we must have $2\alpha_2(B) = \alpha_2(S)$ and thus $\alpha_2(B) = \frac{1}{3}$. Similarly $\alpha_1(B) = 2\alpha_1(S)$, or $\alpha_1(B) = \frac{2}{3}$.

Interpretations of Mixed Strategy Nash Equilibrium

▲Mixed Strategies as Objects of Choice

By the naive interpretation, a mixed strategy is thought of as a deliberate decision to use a random device to select actions.

Criticism: There certainly are cases in which players introduce randomness into their behavior. For example, players randomly "bluff" in poker, governments randomly audit taxpayers, and some stores randomly offer discounts.

However, the notion of a mixed strategy equilibrium in a strategic game does not capture, in my opinion, the players' motivation to introduce randomness into their behavior. Usually a player deliberately randomizes in order to influence the other players' behavior. Consider, for example, the children's version of *Matching Pennies* in which the players choose to display number of fingers. Does a child choose randomly? It is more a response to a guess about the

other player's choice or attempt to mislead the other player. Or, tax authorities audit taxpayers randomly. They would like the taxpayer to know their strategy and are not indifferent between the possible mixed strategies.

▲ *Mixed Strategy Nash Equilibrium as a Steady State*

An n -player game as a model of the interaction of n large populations. Each occurrence of the game takes place after n players are randomly drawn, one from each population. The probabilities in player i 's equilibrium mixed strategy are interpreted as the steady state frequencies with which the actions of A_i are used by the i th population. This interpretation could fit the case that the players are animals or even flowers... However, the use of such concepts requires to specify the evolutionary forces which prevent other steady situations.

Evolutionary Equilibrium

In the following we will refer to situation where in each "match" between two players each of the players have to choose from a set B of "modes of behavior" and there is a function $u(a, b)$ with the interpretation of the payoff of a player who takes the action a against a player who takes the action b .

The organisms do not consciously choose actions; they either inherit modes of behavior from their forebears or are assigned them by mutation. The function u has a *new* meaning: it measures each organism's ability to survive: if an organism takes the action a and faces the distribution β of opponents' actions, then its ability to survive is assumed to be the expectation of $u(a, b)$ under β .

A candidate for an evolutionary equilibrium is an action in B . An equilibrium is a steady state in which all organisms take this action and no mutant can invade the population. More precisely, for $\forall b \in B$ we view the evolutionary process occasionally transforming a small fraction of the population into mutants who follow b . In an equilibrium any such mutant must obtain an expected payoff strictly lower than that of the equilibrium action, so that it dies out. Now, if the fraction $\epsilon > 0$ of the population consists of mutants taking the action b while all other organisms take the action b^* , then the average payoff of a mutant is $(1 - \epsilon)u(b, b^*) + \epsilon u(b, b)$ while the average payoff of a non-mutant is $(1 - \epsilon)u(b^*, b^*) + \epsilon u(b^*, b)$. Therefore for b^* to be an evolutionary equilibrium we require $(1 - \epsilon)u(b, b^*) + \epsilon u(b, b) < (1 - \epsilon)u(b^*, b^*) + \epsilon u(b^*, b)$ for $\forall \epsilon$ sufficiently small. This inequality is satisfied iff $\forall b \neq b^*$ either

$$u(b, b^*) < u(b^*, b^*), \text{ or}$$

$$u(b, b^*) = u(b^*, b^*) \text{ and } u(b, b) < u(b^*, b).$$

Definition Let $G = \langle \{1, 2\}, (B, B), (u_i) \rangle$ be a symmetric strategic game, where $u_1(a, b) = u_2(b, a) = u(a, b)$ for some function u . An *evolutionarily stable strategy* (ESS) of G is an action $b^* \in B$ for which (b^*, b^*) is a NE of G and $u(b, b) < u(b^*, b)$ for every best response $b \in B$ to b^* with $b \neq b^*$.

It is immediate from Definition that if (b^*, b^*) is a symmetric strict NE of a symmetric game (where $u_1(a, b) = u_2(b, a)$ for all a and b) then b^* is an ESS.

Example:

γ	1	-1
-1	γ	1
1	-1	γ

 $\gamma \geq 0$

This game has a unique NE in which each player's mixed strategy is $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. A mutant who uses any of the three pure strategies obtains an expected payoff of 0 when it encounters a non-mutant, but the payoff $\gamma \geq 0$ when it encounters another mutant. Hence the game has no ESS.

▲Mixed Strategies as Pure Strategies in an Extended Game

Under the next interpretation, before a player selects his action he consciously or not receives random private information, inconsequential from the point of view of the players, on which he depends his action. A mixed strategy NE captures the dependence of behavior on factors that the players perceive as irrelevant. A mixed strategy NE, viewed in this way, is a description of a steady state of the system reflecting elements missing from the model. In Bos for example, the strategy might reflect a mood of the player (insisting or giving).

Criticisms:

First, it is hard to accept that the deliberate behavior of a player depends on factors that are irrelevant. We usually give reasons for choices.

Second, the behavior predicted by an equilibrium under this interpretation is very fragile. If a manager's behavior is determined by the type of breakfast he eats, then factors outside the model, such as a change in his diet or the price of eggs, may change the frequency with which he chooses his actions, thus inducing changes in the beliefs of the other players and causing instability.

Finally, in order to interpret an equilibrium of a real life problem in this way one needs to indicate the exogenous variables on which the players base their behavior. For example, to interpret a mixed strategy NE in a model of price competition one should both specify the unmodeled factors that serve as the basis for the firms' pricing policies and show that the information structure is rich enough to span the set of all mixed strategy NE. Those who apply the notion of mixed strategy equilibrium rarely do so.

Correlated Equilibrium

If a mixed strategy NE is a steady state in which each player's action depends on a signal that he receives from "nature" one might wonder why to exclude equilibria with signals which are not independent.

Example

	<i>L</i>	<i>R</i>
<i>T</i>	6, 6	2, 7
<i>B</i>	7, 2	0, 0

The NE payoff profiles are (2, 7) and (7, 2) (pure) and $(4\frac{2}{3}, 4\frac{2}{3})$ (mixed). Consider the case that $\Omega = \{x, y, z\}$ and $\pi(x) = \pi(y) = \pi(z) = \frac{1}{3}$; Player 1's information partition is $\{\{x\}, \{y, z\}\}$ and player 2's $\{\{x, y\}, \{z\}\}$. Define the strategies as follows: $\sigma_1(\{x\}) = B$ and $\sigma_1(\{y, z\}) = T$; $\sigma_2(\{x, y\}) = L$ and $\sigma_2(\{z\}) = R$.

Player 1's behavior is optimal given player 2's: Once he is informed $\{x\}$, player 1 knows

that player 2 plays L and thus it is optimal for him to play B ; in $\{y, z\}$ he assigns equal probabilities to player 2 using L and R , so that it is optimal for him to play T . Symmetrically, for player 2. The payoff profile is $(5, 5)$.

Definition A **correlated equilibrium** of a strategic game $\langle N, (A_i), (u_i) \rangle$ is a BNE of a Baysean game $\langle N, (A_i), (u_i), \Omega, \pi, (\tau_i) \rangle$ for some $(\Omega, \pi, (\tau_i))$

Note, the probability space and information partition are not exogenous but are part of the equilibrium description.

Proposition For every mixed strategy NE α of a finite strategic game $\langle N, (A_i), (u_i) \rangle$ there is a correlated equilibrium in which for each player $i \in N$ the distribution on A_i induced by σ_i is α_i .

Proof Let $\Omega = A (= \times_{j \in N} A_j)$. Define π by $\pi(a) = \prod_{j \in N} \alpha_j(a_j)$. For each $i \in N$ and $a \in A$ let $\tau_i(a) = a_i$. Define $\sigma_i(a) = a_i$ for each $a \in A$. The vector (σ_i) is a correlated equilibrium of the Baysean game $\langle N, (A_i), (u_i), \Omega, \pi, (\tau_i) \rangle$.

Proposition Let $G = \langle N, (A_i), (u_i) \rangle$ be a strategic game. Any convex combination of correlated equilibrium payoff profiles of G is a correlated equilibrium payoff profile of G .

Proof. We will construct the correlated equilibrium by construction interpreted as: first a public random device determines which of the K correlated equilibria is to be played, and then the random variable corresponding to the k th correlated equilibrium is realized.

Let u^1, \dots, u^K be correlated equilibrium payoff profiles and let $(\lambda^1, \dots, \lambda^K) \in \mathfrak{R}^K$ with $\lambda^k \geq 0$ for all k and $\sum_{k=1}^K \lambda^k = 1$. For each value of k let (σ_i^k) be a Baysean equilibrium of $\langle N, (A_i), (u_i), \Omega^k, \pi^k, (\tau_i^k) \rangle$ that generates the payoff profile u^k . Without loss of generality assume that the sets Ω^k are disjoint and the signal sets T_i^k be disjoint. Let $\Omega = \cup_k \Omega^k$. For any $\omega \in \Omega^k$ define $\tau_i(\omega) = \tau_i^k(\omega)$ and $\pi(\omega) = \lambda^k \pi^k(\omega)$. For each $i \in N$ define σ_i by $\sigma_i(t_i) = \sigma_i^k(t_i)$ where k is such that $t_i \in T_i^k$. We constructed a correlated equilibrium for which the payoff profile is $\sum_{k=1}^K \lambda^k u^k$.

Proposition Let $G = \langle N, (A_i), (u_i) \rangle$ be a finite strategic game. Every probability distribution over outcomes that can be obtained in a correlated equilibrium of G can be obtained in a correlated equilibrium in which the set of states is A and for each $i \in N$ $\tau_i(a) = a_i$.

Proof.

This result allows us to confine attention, when calculating correlated equilibrium payoffs, to equilibria in which the set of states is the set of outcomes. Note however that such equilibria may have no natural interpretation.

▲ Mixed Strategies as beliefs

A mixed strategy NE is a profile β of *beliefs*, in which β_i is the *common* belief of all the *other* players about player i 's actions, with the property that $\forall i$ each action in the support of β_i is optimal given β_{-i} .

Each player chooses a single action rather than a mixed strategy. An equilibrium is a steady state of the players' **beliefs**, not their actions. These beliefs are required to satisfy two

properties: they are common among all players and are consistent with the assumption that every player is an expected utility maximizer.

Criticism: When we interpret mixed strategy equilibrium in this way the predictive content of an equilibrium is weak: equilibrium predicts only that each player uses an action that is a best response to the equilibrium beliefs. The set of such best responses includes any action in the support of a player's equilibrium mixed strategy and may even include actions outside the support of this strategy.

Problem set 14

1. **(Exercise)** (*Guess the average*) Each of n people announces a number in the set $\{1, \dots, K\}$. A prize of \$1 is split equally between all the people whose number is closest to $\frac{2}{3}$ of the average number. Show that the game has a unique mixed strategy NE, in which each player's strategy is pure. (see, <http://www.marietta.edu/~delemeeg/expnom/f99.html>)

2. **(Exercise)** (*An investment race*) Two investors are involved in a competition with a prize of \$1. Each investor can spend any amount in the interval $[0, 1]$. The winner is the investor who spends the most; in the event of a tie each investor receives \$0.50. Formulate this situation as a strategic game and find its mixed strategy Nash equilibria. (Note that the players' payoff functions are discontinuous, nevertheless the game has a mixed strategy NE.)

3. **(Easy Exercise)** (*Guessing right*) Players 1 and 2 each choose a member of the set $\{1, \dots, K\}$. If the players choose the same number then player 2 pays \$1 to player 1; otherwise no payment is made. Each player maximizes his expected monetary payoff. Find the mixed strategy Nash equilibria of this (strictly competitive) game.

4. **(Exercise)** (*Air strike*) Army A has a single plane with which it can strike one of three possible targets. Army B has one anti-aircraft gun that can be assigned to one of the targets. The value of target k is v_k , with $v_1 > v_2 > v_3 > 0$. Army A can destroy a target only if the target is undefended and A attacks it. Army A wishes to maximize the expected value of the damage and army B wishes to minimize it. Formulate the situation as a (strictly competitive) strategic game and find its mixed strategy Nash equilibria.

5. **(Exercise)** Consider the following symmetric game:

	L	R
L	0, 0	1, 2
R	2, 1	0, 0

▲Analyze the set of mixed strategy NE and the ESS of this game.

▲Compare with the analysis of the game which fit to the case that when two players interact they play the situation twice and a player can decide about what to do at the second period as a function of what has happened in the first period.

6. **(Exercise)** Consider the three-player game with the payoffs given in the table: (Player 1 chooses one of the two rows, player 2 chooses one of the two columns, and player 3 chooses one of the three tables, A , B or C .)

A	L	R	B	L	R	C	L	R
L	0, 0, 3	0, 0, 0	L	2, 2, 2	0, 0, 0	L	0, 0, 0	0, 0, 0
R	1, 0, 0	0, 0, 0	R	0, 0, 0	2, 2, 2	R	0, 1, 0	0, 0, 3

▲Show that the pure strategy equilibrium payoffs are $(1, 0, 0)$, $(0, 1, 0)$, and $(0, 0, 0)$.

▲Show that there is a correlated equilibrium in which player 3 chooses B and players 1 and 2 play (T, L) and (B, R) with equal probabilities.

▲ Explain the sense in which player 3 prefers not to have the information that players 1 and 2 use to coordinate their actions.