

14.12 Game Theory Lecture Notes*

Lectures 3-6

Muhamet Yildiz[†]

We will formally define the games and some solution concepts, such as Nash Equilibrium, and discuss the assumptions behind these solution concepts.

In order to analyze a game, we need to know

- who the players are,
- which actions are available to them,
- how much each player values each outcome,
- what each player knows.

Notice that we need to specify not only what each player knows about external parameters, such as the payoffs, but also about what they know about the other players' knowledge and beliefs about these parameters, etc. In the first half of this course, we will confine ourselves to the games of complete information, where everything that is known by a player is common knowledge.¹ (We say that X is common knowledge if

*These notes are somewhat incomplete — they do not include some of the topics covered in the class.

[†]Some parts of these notes are based on the notes by Professor Daron Acemoglu, who taught this course before.

¹Knowledge is defined as an operator on the propositions satisfying the following properties:

1. if I know X, X must be true;
2. if I know X, I know that I know X;
3. if I don't know X, I know that I don't know X;
4. if I know something, I know all its logical implications.

everyone knows X, and everyone knows that everyone knows X, and everyone knows that everyone knows that everyone knows X, ad infinitum.) In the second half, we will relax this assumption and allow players to have asymmetric information, focusing on informational issues.

1 Representations of games

The games can be represented in two forms:

1. The normal (strategic) form,
2. The extensive form.

1.1 Normal form

Definition 1 (*Normal form*) An n -player game is any list $G = (S_1, \dots, S_n; u_1, \dots, u_n)$, where, for each $i \in N = \{1, \dots, n\}$, S_i is the set of all strategies that are available to player i , and $u_i : S_1 \times \dots \times S_n \rightarrow \mathbb{R}$ is player i 's von Neumann-Morgenstern utility function.

Notice that a player's utility depends not only on his own strategy but also on the strategies played by other players. Moreover, each player i tries to maximize the expected value of u_i (where the expected values are computed with respect to his own beliefs); in other words, u_i is a von Neumann-Morgenstern utility function. We will say that player i is rational iff he tries to maximize the expected value of u_i (given his beliefs).²

It is also assumed that it is common knowledge that the players are $N = \{1, \dots, n\}$, that the set of strategies available to each player i is S_i , and that each i tries to maximize expected value of u_i given his beliefs.

When there are only two players, we can represent the (normal form) game by a bimatrix (i.e., by two matrices):

| | | |
|------|------|-------|
| 1\2 | left | right |
| up | 0,2 | 1,1 |
| down | 4,1 | 3,2 |

²We have also made another very strong "rationality" assumption in defining knowledge, by assuming that, if I know something, then I know all its logical consequences.

Here, Player 1 has strategies up and down, and 2 has the strategies left and right. In each box the first number is 1's payoff and the second one is 2's (e.g., $u_1(\text{up},\text{left}) = 0$, $u_2(\text{up},\text{left}) = 2$.)


1.2 Extensive form

The extensive form contains all the information about a game, by defining who moves when, what each player knows when he moves, what moves are available to him, and where each move leads to, etc., (whereas the normal form is more of a 'summary' representation). We first introduce some formalisms.


Definition 2 A tree is a set of nodes and directed edges connecting these nodes such that

1. there is an initial node, for which there is no incoming edge;
2. for every other node, there is one incoming edge;
3. for any two nodes, there is a unique path that connect these two nodes.


Imagine the branches of a tree arising from the trunk. For example,



is a tree. On the other hand,



is not a tree because there are two alternative paths through which point A can be reached (via B and via C).



is not a tree either since A and B are not connected to C and D.

Definition 3 (*Extensive form*) A Game consists of a set of players, a tree, an allocation of each node of the tree (except the end nodes) to a player, an informational partition, and payoffs for each player at each end node.

The set of players will include the agents taking part in the game. However, in many games there is room for chance, e.g. the throw of dice in backgammon or the card draws in poker. More broadly, we need to consider “chance” whenever there is uncertainty about some relevant fact. To represent these possibilities we introduce a fictional player: Nature. There is no payoff for Nature at end nodes, and every time a node is allocated to Nature, a probability distribution over the branches that follow needs to be specified, e.g., Tail with probability of 1/2 and Head with probability of 1/2.

An *information set* is a collection of points (nodes) $\{n_1, \dots, n_k\}$ such that

1. the same player i is to move at each of these nodes;
2. the same moves are available at each of these nodes.

Here the player i , who is to move at the information set, is assumed to be unable to distinguish between the points in the information set, but able to distinguish between the points outside the information set from those in it. For instance, consider the game in Figure 1. Here, Player 2 knows that Player 1 has taken action T or B and not action X; but Player 2 cannot know for sure whether 1 has taken T or B. The same game is depicted in Figure 2 slightly differently.




Figure 1:




Figure 2:


An *information partition* is an allocation of each node of the tree (except the starting and end-nodes) to an information set.

To sum up: at any node, we know: which player is to move, which moves are available to the player, and which information set contains the node, summarizing the player's information at the node. Of course, if two nodes are in the same information set, the available moves in these nodes must be the same, for otherwise the player could distinguish the nodes by the available choices. Again, all these are assumed to be common knowledge. For instance, in the game in Figure 1, player 1 knows that, if player 1 takes X, player 2 will know this, but if he takes T or B, player 2 will not know which of these two actions has been taken. (She will know that either T or B will have been taken.)

Definition 4 A strategy of a player is a complete contingent-plan determining which action he will take at each information set he is to move (including the information sets that will not be reached according to this strategy).

For certain purposes it might suffice to look at the reduced-form strategies. A reduced form strategy is defined as an incomplete contingent plan that determines which action the agent will take at each information set he is to move and that has not been precluded by this plan. But for many other purposes we need to look at all the strategies. Let us now consider some examples:

Game 1: Matching Pennies with Perfect Information



The tree consists of 7 nodes. The first one is allocated to player 1, and the next two to player 2. The four end-nodes have payoffs attached to them. Since there are

two players, payoff vectors have two elements. The first number is the payoff of player 1 and the second is the payoff of player 2. These payoffs are von Neumann-Morgenstern utilities so that we can take expectations over them and calculate expected utilities.

The informational partition is very simple; all nodes are in their own information set. In other words, all information sets are singletons (have only 1 element). This implies that there is no uncertainty regarding the previous play (history) in the game. At this point recall that in a tree, each node is reached through a unique path. Therefore, if all information sets are singletons, a player can construct the history of the game perfectly. For instance in this game, player 2 knows whether player 1 chose Head or Tail. And player 1 knows that when he plays Head or Tail, Player 2 will know what player 1 has played. (Games in which all information sets are singletons are called *games of perfect information*.)

In this game, the set of strategies for player 1 is {Head, Tail}. A strategy of player 2 determines what to do depending on what player 1 does. So, his strategies are:

HH = Head if 1 plays Head, and Head if 1 plays Tail;

HT = Head if 1 plays Head, and Tail if 1 plays Tail;

TH = Tail if 1 plays Head, and Head if 1 plays Tail;

TT = Tail if 1 plays Head, and Tail if 1 plays Tail.


What are the payoffs generated by each strategy pair? If player 1 plays Head and 2 plays HH, then the outcome is [1 chooses Head and 2 chooses Head] and thus the payoffs are (-1,1). If player 1 plays Head and 2 plays HT, the outcome is the same, hence the payoffs are (-1,1). If 1 plays Tail and 2 plays HT, then the outcome is [1 chooses Tail and 2 chooses Tail] and thus the payoffs are once again (-1,1). However, if 1 plays Tail and 2 plays HH, then the outcome is [1 chooses Tail and 2 chooses Head] and thus the payoffs are (1,-1). One can compute the payoffs for the other strategy pairs similarly.

Therefore, the normal or the strategic form game corresponding to this game is

| | HH | HT | TH | TT |
|------|------|------|------|------|
| Head | -1,1 | -1,1 | 1,-1 | 1,-1 |
| Tail | 1,-1 | -1,1 | 1,-1 | -1,1 |

Information sets are very important! To see this, consider the following game.

Game 2: Matching Pennies with Imperfect Information




Games 1 and 2 appear very similar but in fact they correspond to two very different situations. In Game 2, when she moves, player 2 does not know whether 1 chose Head or Tail. This is a game of imperfect information (That is, some of the information sets contain more than one node.)

The strategies for player 1 are again Head and Tail. This time player 2 has also only two strategies: Head and Tail (as he does not know what 1 has played). The normal form representation for this game will be:


| 1\2 | Head | Tail |
|------|------|------|
| Head | -1,1 | 1,-1 |
| Tail | 1,-1 | -1,1 |

Game 3: A Game with Nature:



Here, we toss a fair coin, where the probability of Head is $1/2$. If Head comes up, Player 1 chooses between Left and Right; if Tail comes up, Player 2 chooses between Left and Right.

Exercise 5 *What is the normal-form representation for the following game:*



Can you find another extensive-form game that has the same normal-form representation?

[Hint: For each extensive-form game, there is only one normal-form representation (up to a renaming of the strategies), but a normal-form game typically has more than one extensive-form representation.]

In many cases a player may not be able to guess exactly which strategies the other players play. In order to cover these situations we introduce the mixed strategies:

Definition 6 *A mixed strategy of a player is a probability distribution over the set of his strategies.*

If player i has strategies $S_i = \{s_{i1}, s_{i2}, \dots, s_{ik}\}$, then a mixed strategy σ_i for player i is a function on S_i such that $0 \leq \sigma_i(s_{ij}) \leq 1$ and $\sigma_i(s_{i1}) + \sigma_i(s_{i2}) + \dots + \sigma_i(s_{ik}) = 1$. Here σ_i represents other players' beliefs about which strategy i would play.

2 How to play?

We will now describe the most common “solution concepts” for normal-form games. We will first describe the concept of “dominant strategy equilibrium,” which is implied by the rationality of the players. We then discuss “rationalizability” which corresponds to the common knowledge of rationality, and finally we discuss the Nash Equilibrium, which is related to the mutual knowledge of players' conjectures about the other players' actions.

2.1 Dominant-strategy equilibrium

Let us use the notation s_{-i} to mean the list of strategies s_j played by all the players j other than i , i.e.,

$$s_{-i} = (s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_n).$$

Definition 7 A strategy s_i^* strictly dominates s_i if and only if

$$u_i(s_i^*, s_{-i}) > u_i(s_i, s_{-i}), \forall s_{-i} \in S_{-i}.$$

That is, no matter what the other players play, playing s_i^* is strictly better than playing s_i for player i . In that case, if i is rational, he would never play the strictly dominated strategy s_i .³

A mixed strategy σ_i dominates a strategy s_i in a similar way: σ_i strictly dominates s_i if and only if

$$\sigma_i(s_{i1})u_i(s_{i1}, s_{-i}) + \sigma_i(s_{i2})u_i(s_{i2}, s_{-i}) + \dots + \sigma_i(s_{ik})u_i(s_{ik}, s_{-i}) > u_i(s_i, s_{-i}), \forall s_{-i} \in S_{-i}.$$

A rational player i will never play a strategy s_i iff s_i is dominated by a (mixed or pure) strategy.

Similarly, we can define weak dominance.

Definition 8 A strategy s_i^* weakly dominates s_i if and only if

$$u_i(s_i^*, s_{-i}) \geq u_i(s_i, s_{-i}), \forall s_{-i} \in S_{-i}$$

and

$$u_i(s_i^*, s_{-i}) > u_i(s_i, s_{-i})$$

for some $s_{-i} \in S_{-i}$.

That is, no matter what the other players play, playing s_i^* is at least as good as playing s_i , and there are some contingencies in which playing s_i^* is strictly better than s_i . In that case, if rational, i would play s_i only if he believes that these contingencies will never occur. If he is cautious in the sense that he assigns some positive probability for each contingency, he will not play s_i .

³That is, there is no belief under which he would play s_i . Can you prove this?

Definition 9 A strategy s_i^d is a (weakly) dominant strategy for player i if and only if s_i^d weakly dominates all the other strategies of player i . A strategy s_i^d is a strictly dominant strategy for player i if and only if s_i^d strictly dominates all the other strategies of player i .

If i is rational, and has a strictly dominant strategy s_i^d , then he will not play any other strategy. If he has a weakly dominant strategy and cautious, then he will not play other strategies.

Example:

| | | |
|------------|-----------|-------|
| 1\2 | work hard | shirk |
| hire | 2,2 | 1,3 |
| don't hire | 0,0 | 0,0 |

In this game, player 1 (firm) has a strictly dominant strategy which is to "hire." Player 2 has only a weakly dominated strategy. If players are rational, and in addition player 2 is cautious, then we expect player 1 to "hire", and player 2 to "shirk".⁴

| | | |
|------------|----------------|----------------|
| 1\2 | work hard | shirk |
| hire | 2,2 \implies | 1,3 |
| don't hire | 0,0 \uparrow | 0,0 \uparrow |

Definition 10 A strategy profile $s^d = (s_1^d, s_2^d, \dots, s_N^d)$ is a dominant strategy equilibrium, if and only if s_i^d is a dominant strategy for each player i .

As an example consider the Prisoner's Dilemma.

| | | |
|---------------|---------|---------------|
| 1\2 | confess | don't confess |
| confess | -5,-5 | 0,-6 |
| don't confess | -6,0 | -1,-1 |

"Confess" is a strictly dominant strategy for both players, therefore ("confess", "confess") is a dominant strategy equilibrium.

| | | |
|---------------|-----------------|-------------------------------|
| 1\2 | confess | don't confess |
| confess | -5,-5 | \Leftarrow 0,-6 |
| don't confess | -6,0 \uparrow | \Leftarrow -1,-1 \uparrow |

⁴This is the only outcome, provided that each player is rational and player 2 knows that player 1 is rational. Can you show this?

Example: (second-price auction) We have an object to be sold through an auction. There are two buyers. The value of the object for any buyer i is v_i , which is known by the buyer i . Each buyer i submits a bid b_i in a sealed envelope, simultaneously. Then, we open the envelopes;

the agent i^* who submits the highest bid

$$b_{i^*} = \max \{b_1, b_2\}$$

gets the object and pays the second highest bid (which is b_j with $j \neq i^*$). (If two or more buyers submit the highest bid, we select one of them by a coin toss.)

Formally the game is defined by the player set $N = \{1, 2\}$, the strategies b_i , and the payoffs

$$u_i(b_1, b_2) = \begin{cases} v_i - b_j & \text{if } b_i > b_j \\ (v_i - b_j) / 2 & \text{if } b_i = b_j \\ 0 & \text{if } b_i < b_j \end{cases}$$

where $i \neq j$.

In this game, bidding his true valuation v_i is a dominant strategy for each player i . To see this, consider the strategy of bidding some other value $b'_i \neq v_i$ for any i . We want to show that b'_i is weakly dominated by bidding v_i . Consider the case $b'_i < v_i$. If the other player bids some $b_j < b'_i$, player i would get $v_i - b_j$ under both strategies b'_i and v_i . If the other player bids some $b_j \geq v_i$, player i would get 0 under both strategies b'_i and v_i . But if $b_j = b'_i$, bidding v_i yields $v_i - b_j > 0$, while b'_i yields only $(v_i - b_j) / 2$. Likewise, if $b'_i < b_j < v_i$, bidding v_i yields $v_i - b_j > 0$, while b'_i yields only 0. Therefore, bidding v_i dominates b'_i . The case $b'_i > v_i$ is similar, except for when $b'_i > b_j > v_i$, bidding v_i yields 0, while b'_i yields negative payoff $v_i - b_j < 0$. Therefore, bidding v_i is dominant strategy for each player i .

Exercise 11 *Extend this to the n -buyer case.*

When it exists, the dominant strategy equilibrium has an obvious attraction. In that case, the rationality of players implies that the dominant strategy equilibrium will be played. However, it does not exist in general. The following game, the Battle of the Sexes, is supposed to represent a timid first date (though there are other games from animal behavior that deserve this title much more). Both the man and the woman

want to be together rather than go alone. However, being timid, they do not make a firm date. Each is hoping to find the other either at the opera or the ballet. While the woman prefers the ballet, the man prefers the opera.

| | | |
|-----------|-------|--------|
| Man\Woman | opera | ballet |
| opera | 1,4 | 0,0 |
| ballet | 0,0 | 4,1 |

Clearly, no player has a dominant strategy:

| | | |
|-----------|----------------------------|-----------------------------|
| Man\Woman | opera | ballet |
| opera | 1,4 | $\leftarrow \Downarrow$ 0,0 |
| ballet | $0,0 \Uparrow \Rightarrow$ | 4,1 |

2.2 Rationalizability or Iterative elimination of strictly dominated strategies

Consider the following Extended Prisoner’s Dilemma game:

| | | | |
|---------------|---------|---------------|----------|
| 1\2 | confess | don’t confess | run away |
| confess | -5,-5 | 0,-6 | -5,-10 |
| don’t confess | -6,0 | -1,-1 | 0,-10 |
| run away | -10,-6 | -10,0 | -10,-10 |

In this game, no agent has any dominant strategy, but there exists a dominated strategy: “run away” is strictly dominated by “confess” (both for 1 and 2). Now consider 2’s problem. She knows 1 is “rational,” therefore she can predict that 1 will not choose “run away,” thus she can eliminate “run away” and consider the smaller game

| | | | |
|---------------|---------|---------------|----------|
| 1\2 | confess | don’t confess | run away |
| confess | -5,-5 | 0,-6 | -5,-10 |
| don’t confess | -6,0 | -1,-1 | 0,-10 |

where we have eliminated “run away” because it was strictly dominated; the column player reasons that the row player would never choose it.

In this smaller game, 2 has a dominant strategy which is to “confess.” That is, if 2 is rational and knows that 1 is rational, she will play “confess.”

In the original game “don’t confess” did better against “run away,” thus “confess” was not a dominant strategy. However, player 1 playing “run away” cannot be *rationalized* because it is a dominated strategy. This leads to the *Elimination of Strictly Dominated Strategies*. What happens if we “Iteratively Eliminate Strictly Dominated” strategies? That is, we eliminate a strictly dominated strategy, and then look for another strictly dominated strategy in the reduced game. We stop when we can no longer find a strictly dominated strategy. Clearly, if it is common knowledge that players are rational, they will play only the strategies that survive this iteratively elimination of strictly dominated strategies. Therefore, we call such strategies *rationalizable*. **Caution:** we do eliminate the strategies that are dominated by some mixed strategies!

In the above example, the set of rationalizable strategies is once again “confess,” “confess.”

At this point you should stop and apply this method to the Cournot duopoly!! (See Gibbons.) Also, make sure that you can generate the rationality assumption at each elimination. For instance, in the game above, player 2 knows that player 1 is rational and hence he will not “run away;” and since she is also rational, she will play only “confess,” for the “confess” is the only best response for any belief of player 2 that assigns 0 probability to that player 1 “runs away.”

The problem is there may be too many rationalizable strategies. Consider the Matching Pannies game:

| | | |
|------|------|------|
| 1\2 | Head | Tail |
| Head | -1,1 | 1,-1 |
| Tail | 1,-1 | -1,1 |

Here, every strategy is rationalizable. For example, if player 1 believes that player 2 will play Head, then he will play Tail, and if player 2 believes that player 1 will play Tail, then she will play Tail. Thus, the strategy-pair (Head,Tail) is rationalizable. But note that the beliefs of 1 and 2 are not congruent.

The set of rationalizable strategies is in general very large. In contrast, the concept of dominant strategy equilibrium is too restrictive: usually it does not exist.

The reason for existence of too many rationalizable strategies is that we do not restrict players' conjectures to be 'consistent' with what the others are actually doing. For instance, in the rationalizable strategy (Head, Tail), player 2 plays Tail by conjecturing that Player 1 will play Tail, while Player 1 actually plays Head. We consider another concept — Nash Equilibrium (henceforth NE), which assumes mutual knowledge of conjectures, yielding consistency.

2.3 Nash Equilibrium

Consider the battle of the sexes

| | | |
|-----------|-------|--------|
| Man\Woman | opera | ballet |
| opera | 1,4 | 0,0 |
| ballet | 0,0 | 4,1 |

In this game, there is no dominant strategy. But suppose W is playing opera. Then, the best thing M can do is to play opera, too. Thus opera is a best-response for M against opera. Similarly, opera is a best-response for W against opera. Thus, at (opera, opera), neither party wants to take a different action. This is a Nash Equilibrium.

More formally:

Definition 12 For any player i , a strategy s_i^{BR} is a best response to s_{-i} if and only if

$$u_i(s_i^{BR}, s_{-i}) \geq u_i(s_i, s_{-i}), \forall s_i \in S_i$$

This definition is identical to that of a dominant strategy except that it is not for all $s_{-i} \in S_{-i}$ but for a specific strategy s_{-i} . If it were true for all s_{-i} , then S_i^{BR} would also be a dominant strategy, which is a stronger requirement than being a best response against some strategy s_{-i} .

Definition 13 A strategy profile $(s_1^{NE}, \dots, s_N^{NE})$ is a Nash Equilibrium if and only if s_i^{NE} is a best-response to $s_{-i}^{NE} = (s_1^{NE}, \dots, s_{i-1}^{NE}, s_{i+1}^{NE}, \dots, s_N^{NE})$ for each i . That is, for all i , we have that

$$u_i(s_i^{NE}, s_{-i}^{NE}) \geq u_i(s_i, s_{-i}^{NE}) \quad \forall s_i \in S_i.$$

In other words, no player would have an incentive to deviate, if he knew which strategies the other players play.

If a strategy profile is a dominant strategy equilibrium, then it is also a NE, but the reverse is not true. For instance, in the Battle of the Sexes, both (O,O) and (B,B) are Nash equilibria, but neither are dominant strategy equilibria. Furthermore, a dominant strategy equilibrium is unique, but as the Battle of the Sexes shows, Nash equilibrium is not unique in general.

At this point you should stop, and compute the Nash equilibrium in Cournot Duopoly game!! Why does Nash equilibrium coincide with the rationalizable strategies. In general: Are all rationalizable strategies Nash equilibria? Are all Nash equilibria rationalizable? You should also compute the Nash equilibrium in Cournot oligopoly, Bertrand duopoly and in the commons problem.

The definition above covers only the pure strategies. We can define the Nash equilibrium for mixed strategies by changing the pure strategies with the mixed strategies. Again given the mixed strategy of the others, each agent maximizes his expected payoff over his own (mixed) strategies.⁵

Example Consider the Battle of the Sexes again where we located two pure strategy equilibria. In addition to the pure strategy equilibria, there is a mixed strategy equilibrium.

| | | |
|-----------|-------|--------|
| Man\Woman | opera | ballet |
| opera | 1,4 | 0,0 |
| ballet | 0,0 | 4,1 |

Let's write q for the probability that M goes to opera; with probability $1 - q$, he goes to ballet. If we write p for the probability that W goes to opera, we can compute her


⁵In terms of beliefs, this correspondes to the requirement that, if i assigns positive probability to the event that j may play a particular pure strategy s_j , then s_j must be a best response given j 's beliefs.

expected utility from this as

$$\begin{aligned}
 U_2(p; q) &= pq u_2(\text{opera, opera}) + p(1 - q) u_2(\text{ballet, opera}) \\
 &\quad + (1 - p) q u_2(\text{opera, ballet}) + (1 - p)(1 - q) u_2(\text{ballet, ballet}) \\
 &= p [q u_2(\text{opera, opera}) + (1 - q) u_2(\text{ballet, opera})] \\
 &\quad + (1 - p) [q u_2(\text{opera, ballet}) + (1 - q) u_2(\text{ballet, ballet})] \\
 &= p [q4 + (1 - q)0] + (1 - p) [0q + 1(1 - q)] \\
 &= p[4q] + (1 - p)[1 - q].
 \end{aligned}$$

Note that the term $[4q]$ multiplied with p is her expected utility from going to opera, and the term multiplied with $(1 - p)$ is her expected utility from going to ballet. $U_2(p; q)$ is strictly increasing with p if $4q > 1 - q$ (i.e., $q > 1/5$); it is strictly decreasing with p if $4q < 1 - q$, and is constant if $4q = 1 - q$. In that case, W's best response is $p = 1$ if $q > 1/5$, $p = 0$ if $q < 1/5$, and p is any number in $[0, 1]$ if $q = 1/5$. In other words, W would choose opera if her expected utility from opera is higher, ballet if her expected utility from ballet is higher, and can choose any of opera or ballet if she is indifferent between these two.

Similarly we compute that $q = 1$ is best response if $p > 4/5$; $q = 0$ is best response if $p < 4/5$; and any q can be best response if $p = 4/5$. We plot the best responses in the following graph.



The Nash equilibria are where these best responses intersect. There is one at $(0,0)$, when they both go to ballet, one at $(1,1)$, when they both go to opera, and there is one at $(4/5, 1/5)$, when W goes to opera with probability $4/5$, and M goes to opera with probability $1/5$.

Note how we compute the mixed strategy equilibrium (for 2x2 games). We choose 1's probabilities so that 2 is indifferent between his strategies, and we choose 2's probabilities so that 1 is indifferent.