Mixed Strategies: Minimax/Maximin and Nash Equilibrium

In the preceding lecture we analyzed maximin strategies. A maximin strategy is an **assurance** strategy: it achieves the **best** expected payoff a player can possibly **assure** himself, i.e., it's the mixture that yields a player his *best* worst-case expectation.

Maximin value or payoff: the best expected payoff a player can assure himself.

In our pursuit-evasion example: 2/9 for the Row player, and 7/9 for the Column player.

Maximin strategy or mixture: the mixture (mixed strategy) that assures the maximin value.

In our pursuit-evasion example: 2/3 Left and 1/3 Right for each player.

Instead of asking the question we've been asking,

(1) "How high a payoff can I assure myself?" ← My **maximin** value

we could instead ask the questions

(2) "How ***low*** a payoff can I limit my opponent to?" ← Opponent's **minimax** value

(3) "How ***low*** a payoff can my opponent limit me to?" ← My **minimax** value

These three questions are related to one another, but they're different questions. And they're also related to the Nash equilibria of the game.

Maximin and Minimax

Definintions:

Maximin value: The highest value a player can assure himself.

Maximin strategy or mixture: A strategy that assures a player of his maximin value.

Minimax value: The lowest value a player's opponent can limit him to.

Minimax strategy or mixture: A strategy that limits a player's opponent to his minimax value.

From now on, we'll use $\mathbf{m_i}$ to denote player i's maximin value, and $\mathbf{M_i}$ to denote player i's minimax value.

It seems pretty obvious that in our pursuit-and-evasion game,

$$\mathbf{M_{Col}} = 7/9, \text{ because } \mathbf{m_{Row}} = 2/9; \quad \text{and} \quad \mathbf{M_{Row}} = 2/9, \text{ because } \mathbf{m_{Col}} = 7/9.$$

If I'm the Row player, for example, I can assure myself a 2/9 probability of winning a point, but not more than that: $\mathbf{m_{Row}} = 2/9$. So it's clear that I can limit my opponent to a 7/9 probability of winning, but not less than that: $\mathbf{M_{Col}} = 7/9$.

This is because in each of the game matrix's four cells, the two players' payoffs are probabilities, and one player or the other has to win the point – the payoffs sum to 1 in each cell. It's a **constant-sum,** or **strictly competitive** game.

Constant-Sum (Strictly Competitive) Games

Constant-sum Game: A game is **constant-sum** (also called *strictly competitive*) if the sum of all the players' payoffs is the same at every profile of strategies. For constant-sum games, we'll use K to denote the sum of payoffs in each cell. A **zero-sum** game is one for which K = 0.

Several facts about two-player constant-sum games are obvious:

(1) A player's maximin value and his opponent's minimax value must sum to K:

$$\mathbf{m_{Row} + M_{Col} = K} \quad \text{and} \quad \mathbf{m_{Col} + M_{Row} = K} \,.$$

(2) A maximin mixture for a player is also a minimax mixture for him, and vice versa.

(3) A player's maximin value cannot exceed his minimax value: $\mathbf{m_i \leq M_i}$.

(4) The two players' maximin values cannot sum to more than K, nor can their minimax values:

$$\mathbf{m_{Row} + m_{Col} \leq K} \quad \text{and} \quad \mathbf{M_{Col} + M_{Row} \leq K} \,.$$

And here's a fact that's *not* obvious:

(5) A player's minimax and maximin values are the same: $\mathbf{m_i = M_i}$.

In other words, whatever expected payoff a player can be held to by his opponent is also the expected payoff the player can assure himself. This is von Neumann's famous **Minimax Theorem**, which was the beginning of game theory. We will not give a proof of the theorem here.

Nash Equilibrium in Mixed Strategies: The Strategy Sets

In our 2x2 “point game” you may not want to play as conservatively as maximin/minimax prescribes. Of course, if your opponent is playing maximin/minimax, then it won’t matter what you do. But if he is not playing that way, then his play is “exploitable” – since he’s not playing in a way that assures himself of his maximin/minimax value, there is some way for you to play that will keep him *below* that value, thereby doing *better* for you than your maximin value.

In fact, we might want to determine your best response to any strategy by your opponent, and we might want to find out the Nash equilibrium of the game – a profile of strategies that are mutual best responses.

Since we’re now allowing for mixed strategies in our 2x2 game, each player’s set of available strategies has expanded from just two (pure) strategies to all mixtures over his two pure strategies – which we can represent (for the Row player) by the mixtures **r** on Left, since the mixture on Right will be **1-r**. In other words, the Row player’s set of strategies is now all numbers **r** between 0 and 1, i.e., all numbers in the “unit interval,” and similarly for the Column player:

|-------------------------| |-------------------------|

0   ← **r** →   1   0   ← **c** →   1

Payoff Functions and Best Response Functions

So now a player's payoff function has two *continuous* variables: $\pi_{Row}(\mathbf{r}, \mathbf{c})$ and $\pi_{Col}(\mathbf{r}, \mathbf{c})$.

And the best response functions, or reaction curves, will look a lot like the continuous ones we obtained in our Cournot oligopoly analysis:

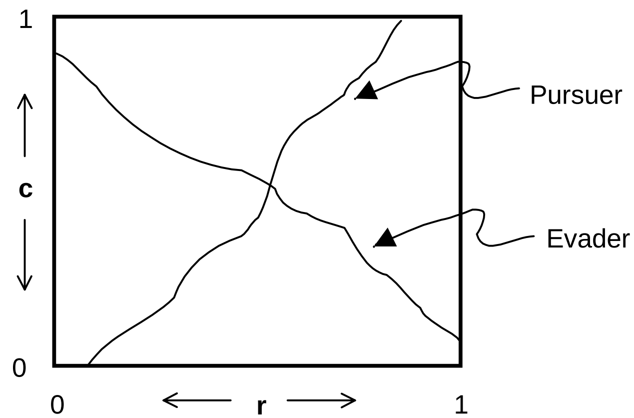

Note that the Pursuer's reaction curve has been drawn upward-sloping and the Evader's downward-sloping. You should be able to see why (for our 2x2 point game).

We can actually determine the exact payoff functions, and therefore the exact best response functions – which we do on the following page.

Payoff Functions and Best Response Functions

What happens when the players use the profile of *mixed* strategies **(r,c)**? We can determine the frequencies with which they'll end up with any particular profile of *pure* strategies – i.e., the frequencies with which they'll end up in any of the four cells of the payoff matrix:

	L	R
L	**rc**	**r(1-c)**
R	**(1-r)c**	**(1-r)(1-c)**

The Row player's payoff function:

$$\pi_{Row}(r,c) = rc\pi_{Row}(L,L) + r(1-c)\pi_{Row}(L,R) + (1-r)c\pi_{Row}(R,L) + (1-r)(1-c)\pi_{Row}(r,R)$$

$$= \frac{1}{3}rc + \frac{2}{3}(1-r)(1-c)$$

$$= \frac{1}{3}rc + \frac{2}{3}(1-r-c+rc)$$

$$= \frac{2}{3}(1-c) + \frac{1}{3}cr - \frac{2}{3}r + \frac{2}{3}cr$$

$$= (\frac{2}{3} - \frac{2}{3}c) + (c - \frac{2}{3})r.$$

Similar algebra also gives us the Column player's payoff function:

$$\pi_{Col}(r,c) = rc\pi_{Col}(L,L) + r(1-c)\pi_{Col}(L,R) + (1-r)c\pi_{Col}(R,L) + (1-r)(1-c)\pi_{Col}(r,R)$$
$$= \frac{2}{3}rc + 1r(1-c) + 1(1-r)c + \frac{1}{3}(1-r)(1-c)$$
$$= (\frac{1}{3} + \frac{2}{3}r) + (\frac{2}{3} - r)c.$$

In order to determine each player's best response function (i.e., reaction curve), notice that …

- if r < 2/3, then π_{Col} (**r, c**) is largest when **c** = 1
- if r > 2/3, then π_{Col} (**r, c**) is largest when **c** = 0
- if r = 2/3, then π_{Col} (**r, c**) is the same for every choice of **c**

and

- if c < 2/3, then π_{Row} (**r, c**) is largest when **r** = 0
- if c > 2/3, then π_{Row} (**r, c**) is largest when **r** = 1
- if c = 2/3, then π_{Row} (**r, c**) is the same for every choice of **r**.

Consequently, the players' reaction functions are as shown on the following page.

The Players' Reaction Functions

The Column player's (the Evader's) reaction function:

If **r** < 2/3, then **c** = 1; if **r** > 2/3, then **c** = 0; if **r** = 2/3, then any **c** is best.

The Row player's (the Pursuer's) reaction function:

If **c** < 2/3, then **r** = 0; if **c** > 2/3, then **r** = 1; if **c** = 2/3, then any **r** is best.

Diagrammatically:

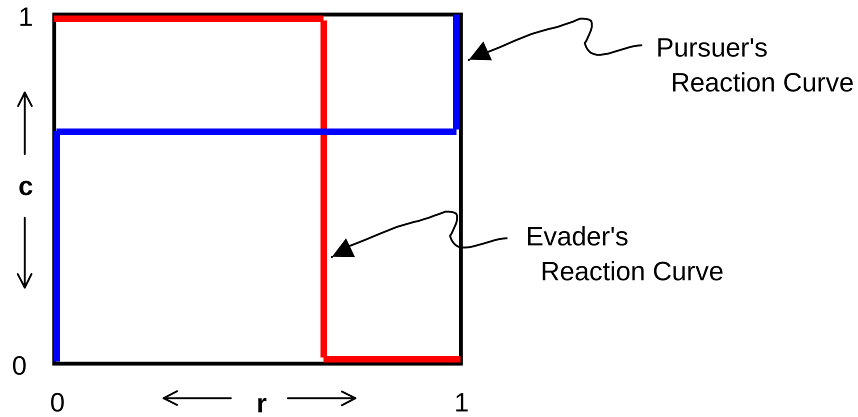

The unique Nash equilibrium is **r = 2/3, c = 2/3 --** i.e., it is for each player to use his minimax mixture!

General Results for Games

(if each player has a finite set of pure strategies)

What we've just discovered about our 2x2 point game – that the Nash equilibrium is for each player to play his minimax (= maximin) mixed strategy – is true for all two-player constant-sum games:

- **In any two-player constant-sum game:**

  It is a Nash equilibrium for each player to play a minimax strategy.

In games that have more players, and in games that are not constant-sum, minimax and maximin are not so useful, but we still always have a Nash equilibrium:

- **In *any* game (any number of players; any payoffs – they needn't be constant-sum):**

  The game will always have at least one Nash equilibrium (but may have no NE in pure strategies).

This second result is Nash's Theorem, formulated and proved by John Nash as a graduate student at Princeton in 1950.