Mixed Strategies and Expected Payoffs

An important concept is a player's *beliefs* about what strategies the other players are choosing.

Sometimes we may be certain that a player will play a particular strategy, but sometimes we may think that a player is "likely" to play a particular strategy, or that the strategy is "plausible."

We formalize this strategic uncertainty as a probability distribution over the opponent's strategies.

For example, if player 1 is sure that player 2 will defect in the Prisoner's Dilemma, he assigns a probability of zero to player 2 cooperating and a probability of one to player 2 defecting.

In the Matching Pennies Game, if player 1 believes that player 2 is equally likely to choose heads and tails, he assigns a probability of $\frac{1}{2}$ to player 2 choosing heads and a probability of $\frac{1}{2}$ to player 2 choosing tails.

Definition: A *belief* of player i is a probability distribution over the strategies of the other players, which we denote as $\theta_{-i} \in \triangle S_{-i}$.

To be a valid probability distribution, the probability of each strategy profile of the other players must be non-negative, and the sum of the probabilities across all profiles must equal one:

$$\begin{aligned} \theta_{-i}(s_{-i}) &\geq 0 \quad \text{for all } s_{-i} \in S_{-i} \\ \sum_{s_{-i} \in S_{-i}} \theta_{-i}(s_{-i}) &= 1. \end{aligned}$$

For the Matching Pennies Game, if player 1 believes that player 2 is equally likely to choose heads and tails, his beliefs according to our notation would be:

$$\begin{aligned} \theta_2(H) &= \frac{1}{2} \\ \theta_2(T) &= \frac{1}{2} \end{aligned}$$

Example with 3 Players:

$S_1 = \{a, b\}$

$S_2 = \{c, d\}$

$S_3 = \{e, f\}$

Then $\theta_{-2}(a, f)$ or $\theta_{1,3}(a, f)$ is the probability that player 2 assigns to **both** player 1 choosing a and player 3 choosing f.

Back to matching pennies, one interpretation of beliefs is that player 2 is consciously trying to randomize between H and T. (Maybe he is flipping his coin to determine his action.)

Another interpretation of beliefs is that player 2 is not randomizing at all, but that player 1 is uncertain of player 2's strategy choice.

–Here is a subtle point: In games with 3 or more players, is it sensible for player 1 to believe that players 2 and 3 randomize in a correlated (coordinated) fashion? We allow, but do not require, a player to believe that other players' actions are correlated with each other.

Related to beliefs is the concept of a mixed strategy. A *mixed strategy* for a player arises when the player selects his strategy according to a probability distribution. For a mixed strategy by player i, we use the notation $\sigma_i \in \triangle S_i$. If all players are choosing mixed strategies, their random choices are independent of each other.

For the special case in which a player assigns probability one to a particular strategy, this is called a *pure strategy*. Thus, a pure strategy is a special case of a mixed strategy.

In the text, Watson (Chapter 4) distinguishes between (i) beliefs that player 1 has about player 2's strategy choice and (ii) the mixed strategy chosen by player 2. Other game theorists believe that player 1's beliefs about player 2's strategy is another interpretation of a mixed strategy for player 2.

–Wrestling example

Beliefs and Payoffs

When player i has beliefs about other players' strategy choices, she faces an uncertain outcome. How does she evaluate her payoff?

We use the concept of expected payoff, which is the mathematical expectation of her payoff. Basically, we take an average, but "weight" the outcomes with the probability of that outcome occurring.

$$u_i(s_i, \theta_{-i}) = \sum_{s_{-i} \in S_{-i}} \theta_{-i}(s_{-i}) u_i(s_i, s_{-i})$$

When player i chooses a mixed strategy:

$$u_i(\sigma_i, \theta_{-i}) = \sum_{s \in S} \sigma_i(s_i) \theta_{-i}(s_{-i}) u_i(s_i, s_{-i})$$

When all players choose a mixed strategy:

$$u_i(\sigma) = \sum_{s \in S} \sigma_1(s_1) \cdots \sigma_i(s_i) \cdots \sigma_n(s_n) u_i(s_i, s_{-i})$$

Here is an example from the text.

		player 2		
		L	M	R
	U	$8, 1$	$0, 2$	$4, 0$
player 1	C	$3, 3$	$1, 2$	$0, 0$
	D	$5, 0$	$2, 3$	$8, 1$

Suppose player 1 believes that player 2 will play strategy L with probability $\frac{1}{2}$, strategy M with probability $\frac{1}{4}$, and strategy R with probability $\frac{1}{4}$.

That is, $\theta_2(L) = \frac{1}{2}, \theta_2(M) = \frac{1}{4}, \theta_2(R) = \frac{1}{4}$.

Then player 1's expected payoff from playing U is

$$u_1(U, \theta_2) = \frac{1}{2}(8) + \frac{1}{4}(0) + \frac{1}{4}(4) = 5.$$

		player 2		
		L	M	R
player 1	U	$8,1$	$0,2$	$4,0$
	C	$3,3$	$1,2$	$0,0$
	D	$5,0$	$2,3$	$8,1$

Suppose $\theta_1(U) = \frac{1}{10}, \theta_1(C) = \frac{2}{10}, \theta_1(D) = \frac{7}{10}$. Then $u_2(M, \theta_1) = \frac{1}{10} \cdot 2 + \frac{2}{10} \cdot 2 + \frac{7}{10} \cdot 3 = \frac{27}{10}$.

Suppose $\sigma_1 = (\frac{1}{3}, \frac{2}{3}, 0)$ and $\sigma_2 = (\frac{1}{4}, \frac{3}{4}, 0)$. Then $u_1(\sigma_1, \sigma_2) = \frac{1}{12} \cdot 8 + \frac{3}{12} \cdot 0 + \frac{2}{12} \cdot 3 + \frac{6}{12} \cdot 1 = \frac{5}{3}$.

Rationality and Common Knowledge

Game theorists usually assume that players are rational. The definition of *rationality* is that a player selects the strategy that he most prefers. In other words, players seek to maximize their expected payoff, given their beliefs about the strategies of the other players.

Notice that once we know the players' beliefs, rationality tells us how to solve the game. The difficult part of solving games is figuring out which beliefs make sense.

Assuming that players are rational does not mean that they are selfish or seek to maximize their own monetary gains. A "payoff" is not necessarily the same as a monetary gain. Our framework is consistent with both altruism and risk aversion:

1. Altruism can be modeled as making a player's payoff increase when the monetary gains of other players increase. Consider the Dictator Game, where player 1 decides how to split \$100 between herself and player 2, and player 2's only strategy is to accept the money.

Thus, $S_1 = \{0, 1, 2, ..., 100\}$ and $S_2 = \{accept\}$.

If player 1 is selfish, his payoffs are

$$u_1(s_1, s_2) = s_1.$$

If player 1 is very altruistic, and cares only about the monetary payoff of the most disadvantaged player, his payoffs are

$$\begin{aligned} u_1(s_1, s_2) &= s_1 \quad \text{if } s_1 \leq 50 \\ u_1(s_1, s_2) &= 100 - s_1 \quad \text{if } s_1 \geq 50. \end{aligned}$$

If player 1 is Mother Teresa, and cares only about the monetary payoff of player 2, her payoffs are $u_1(s_1, s_2) = 100 - s_1$.

More complicated sorts of interdependent preferences are possible.

2. Risk aversion can be modeled as making a player's payoff a concave function of his monetary payoff. Maximizing expected payoff does not mean that you are willing to take a fair bet.

Consider the game where players bet their \$10 lunch money on a matching pennies game, so their monetary payoffs are given by:

		player 2	
		heads	tails
player 1	heads	$20, 0$	$0, 20$
	tails	$0, 20$	$20, 0$

Winning means a fancy lunch, which is better than an ordinary lunch, but losing means going hungry, which is really bad. The actual utility payoffs (if the utility of an ordinary lunch is 0) might be something like:

		player 2	
		heads	tails
player 1	heads	$1, -2$	$-2, 1$
	tails	$-2, 1$	$1, -2$

If player 1 believes that player 2 is choosing heads and tails with equal probability, so $\theta_2(H) = \theta_2(T) = \frac{1}{2}$, then his expected payoff from choosing heads is:

$$u_1(H, \theta_2) = \frac{1}{2}(1) + \frac{1}{2}(-2) = -\frac{1}{2}.$$

Risk aversion makes the players better off not even playing the game and receiving a payoff of zero.

Common Knowledge

In modeling a strategic situation as a game, we assume that the players share a common understanding of the game.

If player 2 is not sure whether she is playing matching pennies or the prisoner's dilemma, then that uncertainty should have been incorporated into the payoff structure, so in fact a totally different game is being played.

If player 1 is not sure that player 2 knows player 1's payoffs, then again the true game must incorporate this uncertainty and is much more complicated.

A fact "F" is *common knowledge* if each player knows F, each player knows that the other player knows F, each player knows that the other player knows that each player knows F, and so on.

Thus, we assume that the game is common knowledge. One way to achieve common knowledge in an experimental setting is to have the instructions read to the players as they are sitting together in the same room.

Puzzles involving Common Knowledge

1. Suppose Jack has arranged to meet Jill at restaurant A, but finds out it is closed. Jack texts Jill to meet at restaurant B instead. Jill texts Jack that she received the message and will meet at restaurant B. When Jill arrives at restaurant B, Jack is not there at the scheduled time.

Is Jack just running late, or did he not receive the confirmation message and think that Jill would be at restaurant A?

The problem is that you can never achieve common knowledge by texting. Better to call instead.

2. One football official standing underneath each goal post when there is a field goal attempt. Why do the officials take longer to call a field goal good when it splits the uprights than call a field goal no good when it is wide left or wide right?

3. A group of 10 people in a room are blindfolded, given either a red hat or a black hat to wear, then told to take off their blindfold without looking at their hat or saying anything about other people's hat color. Thus, each person observes the color of everyone else's hat but not his/her own.

Then the host tells the 10 people that at least one of them has a black hat. The host then announces that she will count from 1 to 10, and that any person that knows his/her hat is black should shout "black hat." (You can assume that it is common knowledge that the 10 people are very smart, and that they like to show off and shout "black hat" as soon as they know it, but would never take a guess and risk the embarrassment of being wrong.)

Claim: If there are n black hats, then each of the people wearing a black hat will shout "black hat" when the host counts to n.

Solution to Hats Puzzle:

It is common knowledge that there is at least one black hat. If there is only one, the person with the black hat sees all red hats and shouts "black hat" when the host counts to 1.

When no one shouts "black hat" when the host counts to 1, it is common knowledge that there are at least two black hats. If there are exactly two black hats, those two people see only 1 black hat, and shout "black hat" when the host counts to 2.

When no one shouts "black hat" when the host counts to 2, it is common knowledge that there are at least three black hats, and so on.

To solve the Hats Puzzle, we needed to go beyond common knowledge of the game itself, and also assume common knowledge of rationality of the players.

We are now ready to move from the definition of the game to the solution of the game, starting with the implications of common knowledge of rationality.

How realistic is it to assume common knowledge of rationality?

In some situations, it is pretty reasonable.

In other situations, people are *boundedly rational*. Either they cannot perform required calculations or they face cognitive limitations preventing them from logically working out the rational strategy.

In still other situations, players are rational but they think that others may not be rational.

Here is a two-player game that illustrates bounded rationality:

The numbers 1-9 are in the center of a board. First player 1 selects a number and moves it from the center to her side of the board. Then player 2 selects one of the 8 remaining numbers and moves it from the center to his side of the board. Then player 1 selects one of the 7 remaining numbers, and so on.

The first player to have exactly three of her/his numbers add up to 15 wins the game, and the other loses. If neither player has exactly three of her/his numbers add up to 15, the game is a tie.