Two-Person Games

Game theory was developed, starting in the 1940's, as a model of situations of conflict. Such situations and interactions will be called *games* and they have participants who are called *players*. We will focus on games with exactly two players. These two players compete for a *payoff* that one player pays to the other. These games are called *zero-sum games* because one player's loss is the other player's gain, and the payoff to both players for any given scenario adds to zero.

A coin-matching game

Roger and Colleen play a game. Each one has a coin. They will both show a side of their coin simultaneously. If both show heads, no money will be exchanged. If Roger shows heads and Colleen shows tails then Colleen will give Roger \$1. If Roger shows tails and Colleen shows heads, then Roger will pay Colleen \$1. If both show tails, then Colleen will give Roger \$2.

This is a *two-person game*, with players Roger and Colleen. It is also a *zero-sum game*. This means that Roger's gain is Colleen's loss (and vice-versa).

Like chess, this is a game where both players have full knowledge of what is happening. Unlike chess, it is a game where both players play *simultaneously*, as opposed to alternately.

A coin-matching game

We can use a 2×2 array to show all four situations that can arise in a single play of this game, and the results of each situation, as follows:

Colleen

Roger		Heads	Tails
	Heads	Roger pays/gets \$0 Colleen pays/gets \$0	Roger gets \$1 Colleen pays \$1
	Tails	Roger pays \$1 Colleen gets \$1	Roger gets \$2 Colleen pays \$2

The amount won by either player in any given situation is called the *pay-off* for that player.

A coin-matching game

Because this is a zero-sum game, we can deduce the pay-off for one player from that of the other, and so we can deduce all information about the game from the *pay-off matrix* shown below:

Colleen

R		H	T
o	H	0	1
g.	T	-1	2

The pay-off matrix for a game shows only the pay-off for the *row player* — this is a universal convention. A positive pay-off for the row player indicates that he gains that amount (and the column player losses that amount). A negative pay-off for the row player indicates that he loses that amount (and the column player gains that amount).

A coin-matching game

	Colleen		
R		H	T
o	H	0	1
g.	T	-1	2

A player's plan of action against the opponent is called a *strategy*.

In the present example, each player has two possible strategies; H and T. If, for example, Roger employs the strategy H and Colleen employs T, then the payoff is 1 — Roger receives \$1, Colleen loses \$1.

Our goal is to determine each player's best strategy, assuming both players want to *maximize* their pay-off. Our conclusions will make most sense when we consider players who are repeatedly playing the same game.

Summary of assumptions in general

- We will limit our attention to two-person zero-sum games.
- We assume that each player is striving to maximize their pay-off.
- Each player will have several options or strategies that (s)he can exercise.
- Each time the game is played, each player selects one option/strategy.
- The players decide on their strategies simultaneously and independently.
- Each player has full knowledge of the strategies available to himself and his opponent and the pay-offs associated to each possible scenario. (However neither player knows which strategy their opponent will choose.)

The language of the pay-off matrix

For a two-player, zero-sum game, we will call the two players R (for row) and C (for column).

For each such game, we can represent all of the information about the game in a matrix, called the *pay-off matrix*.

It is an array with a list of R's strategies as labels for the rows and a list of C's strategies as labels for the columns. The entries in the pay-off matrix are what R gains for each combination of strategies. If this is a negative number than it represents a loss for R.

N.B. — the pay-off matrix will **always** be presented from R's point of view. This is a completely arbitrary choice.

Two Finger Morra

Ruth and Charlie play a game. At each play, Ruth and Charlie simultaneously extend either one or two fingers and call out a number. The player whose call equals the total number of extended fingers wins that many pennies from the opponent. In the event that neither player's call matches the total, no money changes hands.

Here's the pay-off matrix for this game (here strategy $(1, 2)$ means that the player holds up one finger and shouts "2").

Charlie

		$(1,2)$	$(1,3)$	$(2,3)$	$(2,4)$
R	$(1,2)$	0	2	-3	0
u	$(1,3)$	-2	0	0	3
t	$(2,3)$	3	0	0	-4
h	$(2,4)$	0	-3	4	0

Two Finger Morra

Charlie

		$(1,2)$	$(1,3)$	$(2,3)$	$(2,4)$
R	$(1,2)$	0	2	-3	0
u	$(1,3)$	-2	0	0	3
t	$(2,3)$	3	0	0	-4
h	$(2,4)$	0	-3	4	0

(b) What is the payoff for Ruth if Ruth shows two fingers and calls out 4 and Charlie shows 1 finger and calls out 3? What is the payoff for Charlie in this situation?

$(2,4)$ for Ruth is row 4 and $(1,3)$ for Charlie is column 2. Hence -3 is Ruth's payoff so she gives Charlie 3 cents or Charlie's payoff is 3.

Rock-Paper-Scissors

The players face each other and simultaneously display their hands in one of the three following shapes: a fist denoting a rock (R), the forefinger and middle finger extended and spread denoting scissors (S), or a downward facing palm denoting a sheet of paper (P). Rock beats (smashes) Scissors, Scissors beats (cuts) Paper, and Paper beats (covers) Rock. The winner collects a penny from the opponent and no money changes hands in the case of a tie.

Here's the pay-off matrix:

$$\begin{array}{c|rrr} & R & P & S \\ R & 0 & -1 & 1 \\ \hline P & 1 & 0 & -1 \\ S & -1 & 1 & 0 \end{array}$$

Football: Run or Pass?

In football, the offense selects a play and the defense lines up to defend. Here's a simple model, in which offense and defense simultaneously select their play. The offense may choose to run or to pass and the defense may choose a run or a pass defense. One can use data gathered from the past — the average yardage gained or lost in each situation — as payoffs, and so construct a payoff matrix for this two player zero-sum game.

Suppose that if the offense runs and the defense makes the right call, offense losses 5 yards on average; if offense runs and defense makes the wrong call, average gain is 5 yards; on a pass, the right defensive call usually results in an incomplete pass (0 gain/loss), and the wrong defensive call leads to a 10 yard gain.

Payoff matrix:

	DR	DP
OR	-5	5
OP	10	0

Constant-Sum Games

In some games, we have the same assumptions as above except that the pay-offs of both players add to a constant.

For example if both players are competing for a share of a market of fixed size, we can write pay-offs as percentage of the market for each player with the percentages adding to 100. All results and methods that we study for zero-sum games also work for constant sum games.

Example — percentages as payoffs

Rory and Corey own stores next to each other. Each day they announce a sale giving 10% or 20% off. If they both give 10% off, Rory gets 70% of the customers. If Rory goes 10% and Corey 20% sale, Rory gets 30% of the customers. If Rory announces a 20% sale and Corey a 10% sale, Rory gets 90% of the customers. If they both go 20%, Rory gets 50% of the customers. Between them Rory and Corey get all of the customers each day and each customer patronizes only one of the shops each day.

Denote the two choices by BS = 20% off and SS = 10% off (big sale verses small sale). The payoff matrix (Rory as row player) is

	BS	SS
BS	0.5	0.9
SS	0.3	0.7

Example — probabilities as payoffs

General Roadrunner and General Coyote are generals of opposing armies. Every day General Roadrunner sends out a bombing sortie consisting of a heavily armed bomber plane and a lighter support plane. The sortie's mission is to drop a single bomb on General Coyotes forces. However a fighter plane of General Coyote's army is waiting for them in ambush and it will dive down and attack one of the planes in the sortie.

The bomber has an 80% chance of surviving such an attack, and if it survives it is sure to drop the bomb right on the target. General Roadrunner also has the option of placing the bomb on the support plane. In this case, due to this plane's lighter armament and lack of proper equipment, the bomb will reach its target with a probability of only 50% or 90%, depending on whether or not it is attacked by General Coyote's fighter. Represent this information on a pay-off matrix for General Roadrunner.

Example — probabilities as payoffs

Coyote's strategies are attack the bomber B or attack the support S. Roadrunner's are to put the bomb on the bomber B or the support S.

If Coyote plays B and Roadrunner plays B, Coyote wins with probability 0.8. If Coyote plays B and Roadrunner plays S, Coyote wins with probability 1.

If Coyote plays S and Roadrunner plays B, Coyote wins with probability 0.9. If Coyote plays S and Roadrunner plays B, Coyote wins with probability .5.

Pay-off matrix:

	B	S
B	0.8	1.0
S	0.9	0.5

Endgame Basketball [Ruminski]

Often in late game situations, a team with the ball may find themselves down by two points with the shot clock turned off. In this situation, the offensive team must decide whether to shoot for two points, hoping to tie the game and win in overtime, or to try for a three pointer and win the game without overtime. The defending team must decide whether to defend the inside or outside shot. We assume that the probability of winning in overtime is 50% for both teams.

In this situation, the offensive team's coach will ask for a timeout in order to set up the play. Simultaneously, the defensive coach will decide how to set up the defense to ensure a win. Therefore we can consider this as a simultaneous move game with both coaches making their decisions without knowledge of the others strategy.

Endgame Basketball [Ruminski]

To calculate the probability of success for the offense, Ruminski (link to blog post) used NBA league-wide statistics on effective shooting percentages to determine probabilities of success for open and contested shots.

Shot	Success rate
open 2pt.	62.5%
open 3pt.	50%
Contested 2pt.	35.7%
Contested 3pt.	22.8%

Using this and the 50% probability of winning in overtime for each team, we can figure out the probability of winning for each team in all four scenarios using the following tree diagram:

Endgame Basketball [Ruminski]

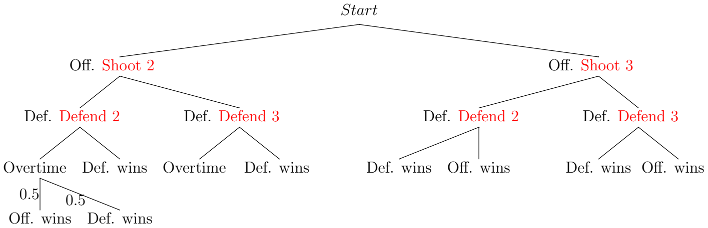

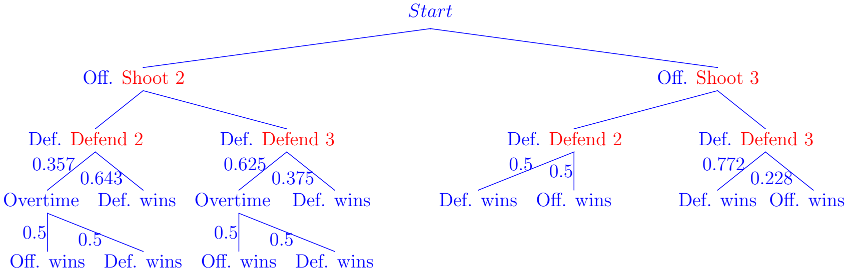

Endgame Basketball [Ruminski]

We can use those probabilities to fill in the probabilities of a win for the row player (offense) in the payoff matrix below. (Note the probability for a win for the defense team is 1 - prob. win for offense.)

		Defending	Team
		Defend 2	**Defend 3**
Offense	**Shoot 2**		
	Shoot 3		

Endgame Basketball [Ruminski]

What are the Offense's chances of winning if they try for 2 against the inside defense? This situation starts at the left node of row 2. The offense wins with probability $0.357 \cdot 0.5 = 0.1785$.

What are the Offense's chances of winning if they try for 2 against the outside defense? This situation starts at the second node of row 2. The offense wins with probability $0.625 \cdot 0.5 = 0.3125$.

What are the Offense's chances of winning if they try for 3 against the inside defense? This situation starts at the third node of row 2. The offense wins with probability 0.5.

What are the Offense's chances of winning if they try for 3 against the outside defense? This situation starts at the right hand node of row 2. The offense wins with probability 0.228.

Endgame Basketball [Ruminski]

The pay-off matrix is

	$D2$	$D3$
$S2$	0.1785	0.3125
$S3$	0.5	0.228

Old Exam Question

Rudolph (R) and Comet (C) play a game. They both choose a number between 1 and 4 simultaneously. Comet gives Rudolph a number of carrots equal to the sum of the two numbers chosen minus three. If this number is negative, Comet recieves carrots from Rudolph. Which of the following give the pay-off matrix for Rudolph?

Old Exam Question

<u>(a)</u>

		C			
	$No.$	1	2	3	4
	1	-1	0	1	2
R	2	0	1	2	3
	3	1	2	3	4
	4	2	3	4	5

(b)

		C			
	$No.$	1	2	3	4
	1	3	2	1	2
R	2	2	1	2	3
	3	1	2	3	4
	4	2	3	4	5

(c)

		C			
	$No.$	1	2	3	4
	1	2	3	4	5
R	2	1	0	-1	-2
	3	0	1	2	3
	4	-1	0	1	2

(d)

		C			
	$No.$	1	2	3	4
	1	-1	0	2	2
R	2	0	0	2	0
	3	1	2	3	3
	4	2	3	4	4

(e)

		C			
	$No.$	1	2	3	4
	1	3	4	5	6
R	2	4	5	6	7
	3	5	6	7	8
	4	6	7	8	9

Rock-Paper-Scissors in the real world

The biologists B. Sinervo and C. M. Lively wrote a report on a lizard species whose males are divided into three classes according to their mating behavior. Each male of the side-blotched lizards (Uta Stansburiana) exhibits one of three (genetically transmitted) behaviors:

a) *highly aggressive*, with a large territory that includes several females;

b) *monogamous*, with a smaller territory that includes one female;

c) nonaggressive *sneaker*, with no territory, who sneaks into the others' domains and tries his luck with the females there.

Here's a link to a wikipedia article.

Rock-Paper-Scissors in the real world

In a confrontation,

- the highly aggressive male has an advantage over the monogamous one (more females in his domain),
- the monogamous male has an advantage over the sneaker (again, more females in his domain), and
- the sneaker has some advantage over the highly aggressive male, because the highly aggressive males must split their time between their various consorts, and so they are vulnerable to sneakers.

The observed consequence of this is that the male populations cycle from a high frequency of monogamous to a high frequency of highly aggressives, then on to a high frequency of sneakers and back to a high frequency of monogamous.

Reference: Sinervo, B. and Lively, C. M., The Rock-Paper-Scissors Game and the evolution of alternative male strategies, *Nature*, **380** (1996), 240–243. link to paper