University of California at Berkeley
Department of Statistics
Fall Semester, 2019

STAT155 — Game Theory
Lecture Notes

Rodrigo Henriquez-Auba

Contents

1 Combinatorial Games

For this class, we define a Progressively Bounded Impartial Combinatorial Game (PBICG), as a game with the following characteristics:

- 2 players take turns to play.
- There is perfect information for both players.
- There is no randomness in the game.
- The game must end in finite moves, and no draws are allowed. That is, every game must have a terminal position on which no moves can be done.
- It is an impartial game, that is, every player has access to the same moves. On the contrary, in a partisan game, players have access to different moves (e.g. chess).

With that, in these games we can a define a position in the game, denoted as x, and a set of possible positions, X, on which there are rules on which moves a player can take. For that, we define a set of rules $F(x) \subseteq X$, as the set of a position that can be reached from x. We write $y \in F(x)$ if there exists a move from x to y.

Thus, for any PBICG, we can define the corresponding directed graph $G = (V, E)$. Each vertex $x \in V = X$ represent a state or position of the game. Edges between states $(x, y) \in E$ denotes if the player can go to the y position (or vertex) from the position x. A terminal position is a position on which no edges go out that vertex.

Definition 1.1. A combinatorial game with set X is **progressively bounded** if $\forall x \in X$, the game will terminate in finite number of moves. That is:

$$\forall x \in X,\ \exists n < \infty \ \text{ s.t. } \ B(x) \leq n$$

where $B(x)$ is the maximum number of moves from x to a terminal position.

Definition 1.2. For a game we define a **normal play** when the last person to move wins, or equivalently the person who can't move (or faces the terminal position) loses.

Definition 1.3. We define a **misère play** if the person who can't move (or face the terminal position) wins, or equivalently, the last person to move loses.

On the following sections we will study PBICG with normal play.

1.1 Take-away games

On these games there is a pile of chips (or several pile of chips), and players must take turns to remove chips from the piles. An example is the game Nim.

1.2 Subtraction games

These are similar to the take-away games but there are rules on how many chips a player can remove from the board. We define the **Subtraction Set** S, that specify the allowed number of chips that can be removed per turn. For example $S = \{1, 2, 3\}$ means that a player can remove only 1,2 or 3 chips.

Definition 1.4. We say that classify a game it is equivalent to classify all the positions, that is, for each position x of the game, classify which player will win (or can ensure a win using a winning strategy).

Definition 1.5. We say that x is a P-position if the **Previous** player (the one that already played) win (or will win) the game.

Definition 1.6. We say that x is an N-position if the **Next** player (the player whose turn is next) win (or will win) the game.

Example 1.1. Let us consider a subtraction game with a single pile of chips, with a subtraction set $S = \{1, 2, 3\}$. We want to classify the positions x (remaining number of chips) depending on which player will win the game. For that, we can construct the following table:

Position x	$F(x)$	Who wins	Comments
0	$\emptyset$	P	Previous player won, since next player face a terminal position (no chips).
1	$\{0\}$	N	Next player wins, since it can leave 0 chips for the previous player.
2	$\{0, 1\}$	N	Next player wins, since it can leave 0 chips for the previous player.
3	$\{0, 1, 2\}$	N	Next player wins, since it can leave 0 chips for the previous player.
4	$\{1, 2, 3\}$	P	Previous player wins, since next player must leave 1, 2 or 3 chips that will go to zero in the following turn.
5	$\{2, 3, 4\}$	N	
6	$\{3, 4, 5\}$	N	
7	$\{4, 5, 6\}$	N	
8	$\{5, 6, 7\}$	P	
9	$\{6, 7, 8\}$	N	

Table 1.1: Classification of positions for Example 1.1.

With that is clear that the classification of this game is given by:

$$x \in P \Leftrightarrow x \in 4\mathbb{Z} \quad \text{or equivalently} \quad x \in P \Leftrightarrow x = 4k, k = 0, 1, 2, \ldots$$

Example 1.2. Let us classify the same game, but with a subtraction set given by $S = \{1, 3, 4\}$. Constructing the table we have:

x	0	1	2	3	4	5	6	7	8	9	10
$F(x)$	$\emptyset$	$\{0\}$	$\{1\}$	$\{0,2\}$	$\{0,1,3\}$	$\{1,2,4\}$	$\{2,3,5\}$	$\{3,4,6\}$	$\{4,5,7\}$	$\{5,6,8\}$	$\{6,7,9\}$
P/N	P	N	P	N	N	N	N	P	N	P	N

Table 1.2: Classification of positions for Example 1.2.

With that we can classify the game as:

$$x \in P \Leftrightarrow x = 7k \text{ or } x = 7k+2, k = 0,1,2,\dots \quad \text{or equivalently} \quad x \in P \Leftrightarrow 0,2 \equiv x \pmod 7$$

We observe that if in a particular position x, $F(x)$ can only go to N positions, then we classify that position x as P (see position $x = 7$ for example), while on the other hand, if there exists at least one P position in $F(x)$, we classify that position x as N (see for example position $x = 8$ that can go to $F(x) = 7$). Using this, we will show the classification of the game using mathematical induction:

For $k = 0$, we have that $x = 0$ and $x = 2$ are indeed P positions. We assume that is true for $k = n$, that is the positions $x = 7n$ and $x = 7n + 2$ are P, and so the positions $x = 7n+1, 7n+3, 7n+4, 7n+5, 7n+6$ are N. We now want to prove that is true for $k = n+1$. We will show for each position that this is true:

Position x	$F(x) = P/N$	P/N
$7(n+1)$	$7n+6 = N, 7n+4 = N, 7n+3 = N$	P
$7(n+1)+1$	$7(n+1) = P, 7n+5 = N, 7n+4 = N$	N
$7(n+1)+2$	$7(n+1)+1 = N, 7n+6 = N, 7n+5 = N$	P
$7(n+1)+3$	$7(n+1)+2 = P, 7(n+1) = P, 7n+6 = N$	N
$7(n+1)+4$	$7(n+1)+3 = N, 7(n+1)+1 = N, 7(n+1) = P$	N
$7(n+1)+5$	$7(n+1)+4 = N, 7(n+1)+2 = P, 7(n+1)+1 = N$	N
$7(n+1)+6$	$7(n+1)+5 = N, 7(n+1)+3 = N, 7(n+1)+2 = P$	N

Table 1.3: Classification of positions for the mathematical induction proof.

As expected $\square$. Implicitly in the proof, we used a theorem, that if $x \notin P$, then it must be in N.

Theorem 1.1. Consider a PBICG with a set of positions X. Let $\mathcal{P} \subseteq X$ be the set of all P-positions and $\mathcal{N} \subseteq X$ be the set of all N-positions. For every $x \in X$, then $x \in \mathcal{N} \cup \mathcal{P}$.

Proof: [From Theorem 1.1.5 in [1]]
Recall that $B(x)$ is the maximum number of moves from x to a terminal position.
Let $\mathcal{N}_i$ (respectively $\mathcal{P}_i$) the set of positions from the next player (respectively, the previous player) can guarantee a win with at most i moves. It is clear that $\mathcal{N} = \bigcup_{i \geq 1} \mathcal{N}_i$ and $\mathcal{P} = \bigcup_{i \geq 0} \mathcal{P}_0$.
We will prove by induction on n, that all positions x with $B(x) \leq n$ are in $\mathcal{N}_n \cup \mathcal{P}_n$:

Certainly, for all x such that $B(x) = 0$, we have that $x \in \mathcal{P}_0 \subseteq \mathcal{P}$. Assume that is true for $B(x) \leq n$, that is for every position x on which $B(x) \leq n$, then $x \in \mathcal{N}_n \cup \mathcal{P}_n$.

Now consider any position z for which $B(z) = n + 1$. Then we will show that every move from z leads to a position w such that $B(w) \leq n$. There are two cases:

- Case 1: Each move from z leads to a position in $\mathcal{N}_n$. Then, since all moves leads to an N position, it must be that $z \in \mathcal{P}_{n+1}$.
- Case 2: There is a move that leads to a position w such that $w \notin \mathcal{N}_n$. Since we already did a move, it is clear $B(w) \leq n$, and hence by induction hypothesis, $w \in \mathcal{N}_n \cup \mathcal{P}_n$. Since z can move to a non-N position, it is clear that $z \in \mathcal{N}_n$.

Hence, all positions lie in $\mathcal{N} \cup \mathcal{P}$. If the starting position is in $\mathcal{N}$, the first player (next player) has a winning strategy, otherwise is in $\mathcal{P}$ and the second player has a winning strategy.

1.3 Algorithm to classify combinatorial games

Based on the previous examples we can construct an algorithm on how to classify combinatorial games:

1. Label every terminal position as P.
2. Label a position N if there exists a move from this position to a P-position.
3. Label a position P if all moves from this position are to N-positions.
4. Check if all positions are labeled, if not, go to 2.

1.4 Chomp

Chomp is a game created by David Gale. It is based on an $m \times n$ rectangular chocolate bar, on which players take turns chomping the bar. The main rule is that you must choose a chunk and eat the chunk and everything that is on the up and on the right of that chunk. The player that eat the bottom left chunk (poisoned one) loses the game.

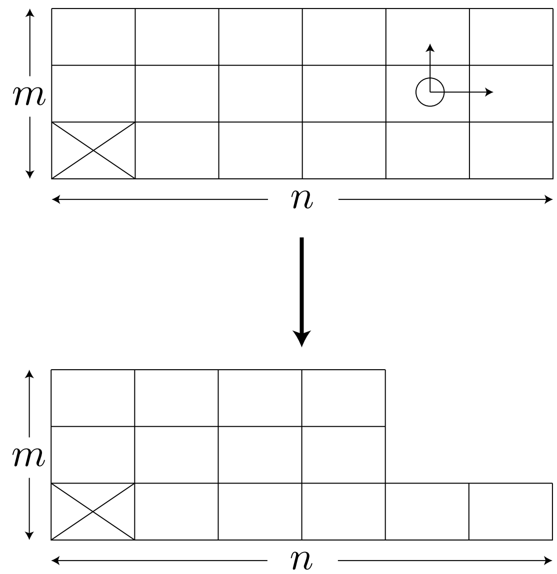

Figure 1.1: Illustration of a move in the game Chomp.

Now we are interested in classifying the positions based on the board:

i. It is clear that a terminal position is just the bottom left chunk. That means the next player has to eat it and lose. Thus, we say that just the bottom left chunk is a terminal position, and hence a P position.

ii. A position with just one extra chunk (on top or right) is an N position, since the next player can eat the remaining chunk and leave the second player (previous player) with just the poisoned chunk.

iii. A position with one chunk at top and one at right of the poisoned one is a P position, since it can only move to an N position (the cases with only one remaining chunk besides the poisoned one).

iv. A position with one chunk on top, and two on the right of the poisoned one is an N position, since you can move to a P position (the one mentioned at iii.).

v. We continue classifying the positions recursively based on the algorithm.

Some particular boards are easy to classify, like square $n \times n$ board or $2 \times n$ (or $m \times 2$) board.

1. For the case of square boards is quite easy to show that is an N position. Start by chunking the chunk in the upper right diagonal next to the poisoned chunk. That will leave the board with 1 column (at the left) and one row (at the bottom). With that, mirror the moves from the second player, and eventually will leave to the position [i.] or [iii.] for the second player, that you know are P positions, and hence you win.

2. For $2 \times n$ boards, is also an N position. The strategy is quite simply, we just want to leave the board like a rectangular grid with the top right chunk removed. Eventually, we will leave the

board as the position [iii.] for the other player (that is a 2×2 board with the top right chunk removed), that we know is a P position, and hence we win. To do this, start by chomping the top right chunk. Then depending on what the other player do:

- If the other player eat one chunk of the bottom row, in your next move you simply eat the top right one in order to leave the board again as a rectangular grid with the top right chunk removed.
- If the other player eat one chunk of the top row, in your next move you simply eat the chunk in the bottom row, but not exactly below the one the other player ate, but the one right next to it. This condition will leave the board again for the other player as a rectangular grid with top right chunk removed.

1.4.1 Strategy Stealing Argument

Consider in Chomp an (n, m) bar, and the first move to remove the top right chunk. We say that x is the initial position for the rectangular grid, and y as the same grid without the top right corner. We seek to classify the initial position x. Since Chomp is a PBICG, then y is either in P or N. If $y \in P$, then $x \in N$ and we are done. Now, consider the $y \in N$, that implies that $\exists z \in P$ such that there is a move from y to z. However, by the rules of Chomp, that move is also available at the beginning from $x \to z$. Then $x \in N$. With that, we have shown that any non-terminal rectangular grid is N.

Note that this strategy allows us to prove that such strategy exists, but we actually do not know what the strategy is.

1.5 Nim

Nim is a PBICG that is played with a set of pile of chips. At each turn, a player must choose a pile of chips and remove as many chips as the player wants. The game is over when there are no remaining chips. The last player to move wins the game (in a normal play), or equivalently the player who cannot remove chips, loses the game.

Note that, for a 1-pile game, the game is always N, since you can remove all the chips at the first move. For a 2-piles game, denoted as (n, m), where n is the remaining number of chips on pile 1, and m is the reimaining on pile 2, it is quite straightforward to classify it:

Position x	$F(x)$	Who wins	Comments
(0,0)	$\emptyset$	P	
(0,1) or (1,0)	(0,0)	N	
(1,1)	(0,1) or (1,0)	P	Only goes to N.
(n, n)		P	Simply mirror the moves until (1,1).
(n, m)	$(\min(n, m), \min(n, m))$	N	You can move to two equals piles, that is N.

Table 1.4: Classification of positions for 2-piles Nim.

For more piles, the trick is partition the piles of powers of 2, for example:

$$(5, 3, 9) \to (2^2 + 2^0, 2^1 + 2^0, 2^3 + 2^0)$$

and we say that P-positions are those in which every power of 2 occurs evenly many times. In this example, that does not happens, and hence $(5, 3, 9)$ is N.

To properly do that, we will write the numbers in binary, and use the XOR sum without carrying.

Example 1.3. Classify the Nim game $(23, 27, 22, 15)$ as N or P. If it is N, find a move to a P-position.

We first write the numbers in binary and do the XOR sum (or simply nim-sum):

2^4	2^3	2^2	2^1	2^0	
1	0	1	1	1	(23)
1	1	0	1	1	(27)
1	0	1	1	0	(22)
0	1	1	1	1	(15)
1	0	1	0	1	

since the nim-sum is not zero, then this is an N-position.

To win, we get to a P-position by modifying a pile to convert the nim-sum to zero. Note, that in Nim you can only remove chips, so since you need to flip the 2^4 position, the pile with 15 chips do not have a winning move. The winning moves are the following:

- (23): (1,0,1,1,1) $\rightarrow$ (0,0,0,1,0) (2)
- (27): (1,1,0,1,1) $\rightarrow$ (0,1,1,1,0) (14)
- (22): (1,0,1,1,0) $\rightarrow$ (0,0,0,1,1) (3)

Theorem 1.2. **[Bouton's Theorem]** A position $x = (x_1, \ldots, x_k)$ in Nim is in P if and only if $x_1 \oplus x_2 \oplus \ldots \oplus x_k = 0$

Proof: Let $\hat{P}$ be the set of all positions with nim-sum equals to zero, and let $\hat{N}$ be everything else, $\hat{N} = X - \hat{P}$.

1. Terminal position: $(0, \ldots, 0) \rightarrow 0 \oplus \ldots \oplus 0 = 0 \rightarrow (0, \ldots, 0) \in \hat{P}$.

2. Consider a position $x \in \hat{N}$ so $x_1 \oplus \ldots \oplus x_k > 0$. Pick the largest pile, and consider the left most column with the numbers in binary notation with odd number of 1's and flip that to zero. Continue with the next columns flipping odd numbers of 1's to zero. This is a legal move because only remove chips and a move from $\hat{N} \rightarrow \hat{P}$.

3. If $\hat{x} \in \hat{P} \Rightarrow x_1 \oplus \ldots \oplus x_k = 0$. Suppose w.l.o.g we can change x_1 to $x_1' < x_1$. Can $x_1' \oplus x_2 \oplus \ldots \oplus x_k = 0$? Suppose by the sake of contradiction that yes. Then:

$$x_1' \oplus x_2 \oplus \ldots x_k = x_1 \oplus \ldots \oplus x_k \;\Rightarrow\; x_1' = x_1 \quad \rightarrow\leftarrow$$

   since $x_i \oplus x_i = 0$, and we said that $x_1' < x_1$. Then $x' = (x_1', \ldots, x_k) \in \hat{N}$. Then, every $\hat{P}$ position go to $\hat{N}$, the same as the algorithm, and thus, $\hat{P} = P$ and $\hat{N} = N$ $\square$.

1.6 Sprague-Grundy Theorem

1.6.1 Minimal Excluded Integer

Definition 1.7. Define mex as Minimal Excluded Integer, defined on sets of non-negative integers. Let $A \subseteq \mathbb{Z}_0^+$:

$$\text{mex}(A) = \min\{n, n \geq 0, n \notin A\}$$

Example 1.4. Consider the following examples:

- $\text{mex}(\text{primes}) = 0$.
- $\text{mex}(\{0,1,2\}) = 3$.
- $\text{mex}(\{0,1,3,4\}) = 2$.

1.6.2 Sprague-Grundy function

Definition 1.8. The Sprague-Grundy function $g(x)$, is defined on directed graph, like any PBICG:

$$g(x) = \text{mex}\{g(y) : y \in F(x)\}$$

Note that we take the mex of the Sprague-Grundy values (SG values) of the followers of x, not the value of the followers.

Example 1.5. Let x be a terminal position, then $F(x) = \emptyset$, and so $g(x) = 0$. Now, let y be a position that only goes to a terminal position. Then it is clear that $g(y) = 1$, since all terminal positions have SG values of 0, and so the mex is 1.

1.6.3 Classification of Sprague-Grundy values

Consider a game G and let $Q \subset X$ be the set of all $x \in X$ such that $g(x) = 0$:

$$Q = \{x \in X : \ g(x)\}$$

and define $R = X - Q$. Note the following:

1. All terminal positions are in Q.
2. Let $x \in Q \ \Rightarrow \ g(x) = 0$, or equivalently:

   $$0 = g(x) = \text{mex}\{g(y) : \ y \in F(x)$$

   that is, $\exists$ a move from x to y in the game (if x is not a terminal position).

   This implies in $y \in F(x)$, there are none $g(y) = 0$. That is, all the moves from x that goes to y, satisfies that $g(y) > 0$, i.e. every move from x goes to $R = Q^c$.
3. Let $y \in Q$. That is:

   $$0 < g(y) = \text{mex}\{g(z) : \ z \in F(y)\}$$

   Since $g(y) > 0$, then that means there is some $z \in F(z)$ such that $g(z) = 0$, by the properties of the mex function. Thus, $\exists$ a move from y into Q.

4. The previous considerations imply that $Q = P$ and $R = N$ from our characterization of the sets P and N and the algorithm that classifies P and N positions. That is:

$$x \in P \Leftrightarrow g(x) = 0$$

Example 1.6. Consider the subtraction game $S = \{1, 2, 3\}$. The Sprague-Grundy values are as follow:

x	0	1	2	3	4	5	6	7	8	9	10
$g(x)$	0	1	2	3	0	1	2	3	0	1	2
P/N	P	N	N	N	P	N	N	N	P	N	N

Table 1.5: Classification of positions for Example 1.6.

The conjecture is that $g(x) \equiv x \bmod 4$, or equivalently:

$$g(x) = \begin{cases} 0, & \text{if } x = 4k \\ 1, & \text{if } x = 4k+1 \\ 2, & \text{if } x = 4k+2 \\ 3, & \text{if } x = 4k+3 \end{cases}$$

for $k = 0, 1, 2, \ldots$. We will show that this is true by induction. The base is clearly true by looking Table 1.5. For the induction hypothesis, we assume that is true for $0 \leq k \leq n$. Then, the induction step is simply to show for $k = n + 1$:

- For $x = 4(n+1)$ we can go to $F(x) = \{4n+3, 4n+2, 4n+1\}$, and so $g\{F(x)\} = \{3, 2, 1\}$ by induction hypothesis. Thus, $g(x) = 0$.
- For $x = 4(n+1)+1$ we can go to $F(x) = \{4n+4, 4n+3, 4n+2\}$, and so $g\{F(x)\} = \{0, 3, 2\}$. Thus, $g(x) = 1$.
- For $x = 4(n+1)+2$ we can go to $F(x) = \{4n+5, 4n+4, 4n+3\}$, and so $g\{F(x)\} = \{1, 0, 3\}$. Thus, $g(x) = 2$.
- For $x = 4(n+1)+3$ we can go to $F(x) = \{4n+6, 4n+5, 4n+4\}$, and so $g\{F(x)\} = \{2, 1, 0\}$. Thus, $g(x) = 3$ □.

1.6.4 Sum game of several PBCIG

Let $G_1, G_2, \ldots, G_n$ be PBCIG with the game G_i having the position set X_i and the function F_i describing the moves of G_i. Then $G = G_1 + G_2 + \ldots + G_n = (X, F)$, where $X = X_1 \times \ldots \times X_n$ and if $x \in X$, then $x = (x_1, x_2, \ldots, x_n)$, where x_i is a position in G_i. We say that $y \in F(x)$ if we can obtain $y \in X$ by moving exactly one of the games G_i, that is: $y = (x_1, x_2, \ldots, y_i, \ldots, x_n)$, where $y_i \in F_i(x_i)$ and so:

$$\begin{aligned} F(x) = & F_1(x_1) \times \{x_2\} \times \ldots \times \{x_n\} \\ & \cup \{x_1\} \times F(x_2) \times \ldots \times \{x_n\} \\ & \vdots \\ & \cup \{x_1\} \times \ldots \times \{x_{n-1}\} \times F(x_n) \end{aligned}$$

Example 1.7. Think a nim game with 2 piles as a sum game of two 1-pile nim games.

As showed before, we know that for a game if $g(x) = 0$, then $x \in P$, and $x \in N$ otherwise. The main question is how to compute the Sprague-Grundy value of the position for a sum game. To do this we use a generalization of Bouton's theorem:

Theorem 1.3. [**Sprague-Grundy Theorem**] If g_i is the Sprague-Grundy function of G_i for $1 \leq i \leq n$, then the sum game G has Sprague-Grundy function g given by:

$$g(x_1, x_2, \ldots, x_n) = g_1(x_1) \oplus g_2(x_2) \oplus \ldots \oplus g_n(x_n)$$

That is, the Sprague-Grundy value of a position in the sum game is the nim-sum of the SG values of the component positions.

Proof [From Ferguson [2]]: Let $x = (x_1, \ldots, x_n)$ be a position in X. Let $b = g_1(x_1) \oplus \ldots \oplus g_n(x_n)$. We have to show two things for the function $g(x)$:

(1) For every non-negative integer $a < b$, there is a follower of $(x_1, \ldots, x_n)$ that has a SG value a.

(2) No follower of x has SG value b.

Thus, the SG value of x, being the smallest SG value not assumed by the one if its followers must be b.

To show (1), let $d = a \oplus b$, and k be the number of digits in the binary expansion of d, so that $2^{k-1} \leq d < 2^k$ and d has a 1 in the k-th position (from the right). Since $a < b$, b has a 1 in the k-th position and a has a 0 there. Since $b = g_1(x_1) \oplus \ldots \oplus g_n(x_n)$, there is at least one x_i such that the binary expansion of $g_i(x_i)$ has a 1 in the k-th position. Suppose w.l.o.g that $i = 1$. Then $d \oplus g(x_1) < g(x_1)$ so that there is a move from x_1 to some x'_i with $g_1(x'_1) = d \oplus g_1(x_1)$. Then the move from $(x_1, x_2, \ldots, x_n)$ to $(x'_1, x_2, \ldots, x_n)$ is a legal move in the sum game G and:

$$g_1(x'_1) \oplus g_2(x_2) \oplus \ldots g_n(x_n) = d \oplus g_1(x_1) \oplus g_2(x_2) \oplus \ldots \oplus g_n(x_n) = d \oplus b = a$$

Finally, to show (2), suppose for the sake of contradiction that $(x_1, \ldots, x_n)$ has a follower with the same SG value, and suppose w.l.o.g. that this involves a move in the first game. That is, we suppose that $(x'_1, x_2, \ldots, x_n)$ is a follower of $(x_1, x_2, \ldots x_n)$ and that $g_1(x'_1) \oplus g_2(x_2) \oplus \ldots \oplus g_n(x_n) = g_1(x_1) \oplus g_2(x_2) \oplus \ldots \oplus g_n(x_n)$. By the cancellation law, $g_1(x'_1) = g_1(x_1)$. But this is a contradiction, since no position can have a follower of the same SG value. □

Example 1.8. Let $S_m = \{1, 2, \ldots, m\}$. We claim that:

$$g(x) := g_m(x) \equiv x \bmod m + 1 \quad 0 \leq g_m(x) \leq m$$

The proof is an extension of Example 1.6, by simply using induction on all the cases.

Example 1.9. Let G_m be the game defined by the subtraction game S_m defined in example 1.8. Let $G = G_3 + G_5 + G_7$, and let $x = (9, 10, 14)$. Determine if x is in P or N.

It is clear that $g_3(9) = 1$, $g_5(10) = 4$, $g_7(14) = 6$, and so:

2^2	2^1	2^0	
0	0	1	$g_3(9) = 1$
1	0	0	$g_5(10) = 4$
1	1	0	$g_7(14) = 6$
0	1	1	

since the nim-sum is different than zero, then this is an N-position by the Sprague-Grundy Theorem. The winning move, is removing one chip from the third pile (14 chips to 13 chips, that is fliping the 2^1 one to zero, and the 2^0 zero to one).

1.6.5 Lasker's Nim

Lasker's Nim is similar to nim, but with one additional move, on which each player can split into two non-empty piles in a move. Clearly, for the one-pile game it satisfies that $g(0) = 0$ and $g(1) = 1$. The followers of 2, are 0,1 and (1,1) with respective SG values of 0,1 and $1 \oplus 1 = 0$. Hence $g(2) = 2$. The followers of 3 are $0, 1, 2$ and $(1, 2)$ with SG values of $0, 1, 2$ and $1 \oplus 2 = 3$. Hence, $g(3) = 4$. Continuing with this analysis we got:

x	0	1	2	3	4	5	6	7	8	9	10	11	12	...
$g(x)$	0	1	2	4	3	5	6	8	7	9	10	12	11	...

Table 1.6: Classification of positions for Lasker's Nim.

We see that a possible conjecture is such that:

$$g(x) = \begin{cases} 4k+1 & \text{if } x = 4k+1 \\ 4k+2 & \text{if } x = 4k+2 \\ 4k+4 & \text{if } x = 4k+3 \\ 4k+3 & \text{if } x = 4k+4 \end{cases}$$

for all $k \geq 0$. The result can be proved by induction as follow. For $k = 0$ the table shows that it is satisfied. Now assuming that is true for all $x \leq 4k$ we have:

- The followers of $4k + 1$ that consist of a single pile have SG value from 0 to $4k$. Those that consist of two piles, $(4k + 1), (4k - 1, 2), \ldots, (2k + 1, 2k)$, have even SG values, and therefore $g(4k + 1) = 4k + 1$.

- The followers of $4k + 2$ that consists on a single pile has SG values from 0 to $4k + 1$. Those that consist of two piles, $(4k + 1, 1), (4k, 2), \ldots, (2k + 1, 2k + 1)$, have SG values alternately divisible by 4 and odd, so that $g(4k + 2) = 4k + 2$.

- The followers of $4k + 3$ that consist of a single pile have values from 0 to $4k + 2$. Those that consist of two piles, $(4k + 2, 1), (4k + 1, 2), \ldots, (2k + 2, 2k + 1)$, have odd SG values and in particular $(4k + 2, 1) = 4k + 3$. Hence $g(4k + 3) = 4k + 4$.

- Finally, the followers of $4k + 4$ that consist of a single pile have SG values from 0 to $4k + 2$, and $4k + 4$. Those that consist of two piles, $(4k + 3, 1), (4k + 2, 2), \ldots, (2k + 2, 2k + 2)$ have SG values alternately equal to $1 \mod 4$ and even. Hence, $g(4k + 4) = 4k + 3$.

1.6.6 Wythoff's Game

The Wythoff's game are played in a board given by a chip and a chessboard. Players, take turns moving the chips. At each turn, a player must move the chip like a queen in a chess game, but only vertically down, horizontally left or diagonally down to the left. When the chip reaches the lower left corner, the game is over and the player to move last wins. The game is usually played on a 8×8 board. It seems that there is no close form for positions. The solution for the 8×8 board game is given by:

7	8	6	9	0	1	4	5
6	7	8	1	9	10	3	4
5	3	4	0	6	8	10	1
4	5	3	2	7	6	9	0
3	4	5	6	2	0	1	9
2	0	1	5	3	4	8	6
1	2	0	4	5	3	7	8
0	1	2	3	4	5	6	7

Table 1.7: Classification of positions for Wythoff's Game.

2 Two-person zero-sum games

A description of a zero sum game can be done using a payoff matrix. Player I (or PI) has m possible actions and Player II (or PII) has n possible actions available. For each pair (i, j) there is a payoff to PI: A_{ij}. This gives the payoff matrix $A = (A_{ij})$. If $A_{ij} > 0$, PII pays PI, and if $A_{ij} < 0$ then PI pays PII.

In this sense, PI seeks to maximize his payoff while PII seeks to minimize his payments (based on the defined payoff matrix).

Example 2.1. The inspection game is a classical zero-sum game. There is one firm and one environmental entity that wants the following: the firm wants to dispose waste in the river, that can be done today or tomorrow, while the environmental entity has only the resources to send an employee only day, today or tomorrow.

We say that is a win for the entity if they visit the plant the same day the waste is being disposed, and assign a value of 1, while we say that is a loss for the entity if they visit a different day, assigning a value of -1.

		PII: Firm		
		Today		Tomorrow
PI:	Today	$+1$	$\rightarrow$	-1
Environmental		$\uparrow$		$\downarrow$
Entity	Tomorrow	-1	$\leftarrow$	$+1$

Table 2.1: Inspection game description.

In this case, PI prefers positive values, while player II prefers negative values. It is assumed that there is no information and both players know all the strategies. Strategies along the rows are for PI, while strategies along the columns are for PII.

- Horizontal arrows represent the preference of the column players (i.e. PII).
- Vertical arrows represent the preference of the row player (PI).

Example 2.2. Odd-even: Players call out 1 or 2. If sum is even, PII wins the sum, if odd, PI wins the amount. The payoff matrix is:

PI/PII	1	2
1	-2	3
2	3	-4

Table 2.2: Odd-even game: payoff matrix

As can be seen, if PI know PII's strategy, then it will play a specific strategy, and in that case the PII will realize and modify his strategy, but then PI will realize and will modify its strategy. If this cyclic behavior occurs then we realize that randomize the strategy is a good idea to on average obtain a good result.

2.1 Optimization of worst case scenario

In the following games, we assume that both players play optimizing of their respective worst-case scenario.

In PI's case, what we want to guarantee is some minimum gain, i.e. a floor gain. In the case of PII we want to guarantee a maximum loss, i.e. a ceil loss.

In the following section we will show that in zero-sum games, the floor gain for PI and ceil loss PII are the same under optimal mixed strategies (see Definition 2.2) due to Minimax theorem.

2.2 Randomization of plays

Consider the Example 2.2, and let's say that its player consecutively. As PI, we are interested in finding a strategy of playing 1 or 2 in order to maximize our expected payoff. Let p_1 be the probability of playing 1, and hence $1 - p_1$ the probability of playing 2. Now, if PII plays 1, PI expected payoff is given by $-2p_1 + 3(1 - p_1) = 3 - 5p_1$, while if PII plays 2, PI expected payoff is given by $3p_1 - 4(1 - p_1) = 7p_1 - 4$. Thus the two cases are:

- PII plays 1: $\mathbb{E}[\text{payoff}] = 3 - 5p_1$
- PII plays 2: $\mathbb{E}[\text{payoff}] = 7p_1 - 4$

Since PI wants to defend himself against the worst-case expected gain, in the worst-case his profit will be $\min\{3 - 5p_1, 7p_1 - 4\}$. Since these two are simply straight lines in p_1, then the worst-case is maximized when both lines intersect. Thus, we simply solve:

$$3 - 5p_1 = 7p_1 - 4 \to p_1 = \frac{7}{12}$$

We say the $p = (7/12, 5/12)$ is the optimal **mixed strategy** for PI.

Definition 2.1. We say that **pure strategies** are such that a player strategy is playing the same action, i.e. $p = (0, \ldots, 1, \ldots, 0)$.

Definition 2.2. We say that **mixed strategies** are such that a player strategy is playing under a randomization mechanism to choose which action to take at each round, i.e. $p = (p_1, \ldots, p_i, \ldots, p_m)$.

Definition 2.3. We define the value of a zero-sum game (in the sense of mixed strategies) as the floor gain V_- that PI can guarantee or the ceiling loss that PII can guarantee V_+. Due to the Minimax theorem we will have that: $V_- = V_+ = V$.

2.3 Dominating Strategies

Consider the payoff matrix:

$$A = \begin{bmatrix} 12 & -1 & 1 & 0 \\ 5 & 1 & 7 & -20 \\ 3 & 2 & 4 & 3 \\ -16 & 0 & 0 & 16 \end{bmatrix}$$

As Player II you never play column 3, since its payoff are always worse than column 2:

$$\begin{array}{ccc} -1 & \leq & 1 \\ 1 & \leq & 7 \\ 2 & \leq & 4 \\ 0 & \leq & 0 \end{array}$$

In this case we say that strategy S **dominates** strategy T (in this case $S = 2$ and $T = 3$ for PII), since every outcome in S is at least as good as every outcome in T, and at least one outcome in S is strictly better than the corresponding outcome in T. This leads to the **dominance principle**, on which *a rational player should never play a dominated strategy.*

In this case, this implies that we can reduce the game to:

$$A = \begin{bmatrix} 12 & -1 & 0 \\ 5 & 1 & -20 \\ 3 & 2 & 3 \\ -16 & 0 & 16 \end{bmatrix}$$

In addition, mixed strategies can be used to dominates rows or columns.

2.4 Saddle points

A saddle point (in the sense of pure strategies) is a particular entry a_{ij} in a payoff matrix that simultaneously satisfies that:

- It is the largest value in its column. That is, if PII is playing is playing that column j, then i is the best action for PI:

$$\max_{\alpha} a_{\alpha j} = a_{ij}$$

- It is the smallest value in its row. That is, if PI is playing the row i, then j is the best action for PII:

$$\min_{\beta} a_{i\beta} = a_{ij}$$

Thus, a saddle point is the best option for PI and PII to guarantee their worst-case scenario. In a zero-sum game with saddle points (that has the same value) then the value of the game is exactly the value of the entry of the saddle point.

2.5 Solving zero-sum games

We are interested in finding the value and optimal strategies for zero-sum games. The algorithm is as follows:

- Find saddle points in the payoff matrix. If there exists one, then the value of the game is given by the saddle point a_{ij}, and the optimal strategies are pure strategies that involves playing i for PI and j for PII.

- If no saddle points can be found, use dominanting strategies to reduce the payoff matrix.

- Use a method to find the optimal mixed strategies $p = (p_1, \ldots, p_m)$ and $q = (q_1, \ldots, q_m)$. Methods will be described in the following subsections.

First we will describe the optimization problem for Player I. Consider a zero-sum game with payoff matrix $A \in \mathbb{R}^{m\times n}$. As PI we are interested in guaranteeing a floor gain by choosing a mixed strategy p, independently of what PII chooses as a strategy q. This can be written in the following optimization problem:

$$p^* = \underset{p\in\Delta_m}{\operatorname{argmax}} \ \min_{q\in\Delta_n} p^\top A q$$

on which Δ_m and Δ_n are the m, n respective simplex:

$$\Delta_m = \left\{ p = (p_1, \ldots, p_m) : \ \sum_{i=1}^{m} p_i = 1, \ p_i \geq 0, i = \{1, \ldots, m\} \right\}$$

$$\Delta_n = \left\{ q = (q_1, \ldots, q_n) : \ \sum_{j=1}^{n} q_j = 1, \ q_j \geq 0, j = \{1, \ldots, n\} \right\}$$

The key property that we will be using to solve the problems is that due to the convexity of this problem, we can reduce it to simply consider pure strategies for PII.

That is we can simplify the simplex Δ_n as the collection of coordinate vectors $e_j^\top = (0, \ldots, 1, \ldots 0)$ for $j = \{1, \ldots, n\}$. Thus, we need to find p only looking in a finite number of linear functions over p:

$$p^* = \underset{p\in\Delta_m}{\operatorname{argmax}} \ \min\{p^\top A e_1, \ p^\top A e_2, \ \ldots \ , \ p^\top A e_n\}$$

A similar analysis can be done for PII:

$$q^* = \underset{q\in\Delta_n}{\operatorname{argmin}} \ \max\{e_1^\top A q, \ e_2^\top A q, \ \ldots \ , \ e_m^\top A q\}.$$

2.5.1 Solving using linear programming

The following subsection was not covered in class and goes into how to find optimal strategies using linear programming. This is adapted from [3]. The optimization technique requires to have a payoff matrix with only positive values, so we use the following theorem that is stated as Lemma 2.2 in [3].

Theorem 2.1 (Equivalent Payoff Matrices)**.** Let A and B be two (m, n) dimensional matrices related as follow:

$$A = B + c\mathbb{1}_m \mathbb{1}_n^\top$$

where c is a positive value and $\mathbb{1}_m$ is a vector of ones of dimension m. That is, $a_{ij} = b_{ij} + c$, for every (i, j). Then:

i. Every pair of optimal mixed strategies (p^*, q^*) of the zero-sum game of A also constitutes a pair of optimal mixed strategies of the zero-sum game of B, and viceversa.

ii. $V(A) = V(B) + c$. That is, the value of the game A is equal to the value of game B plus c.

Proof: It is easy to see that:

$$p^*(A) = \underset{p\in\Delta_m}{\text{argmax}} \min_{q\in\Delta_n} p^\top Aq = \underset{p\in\Delta_m}{\text{argmax}} \min_{q\in\Delta_n} p^\top Bq + cp^\top \mathbb{1}_m \mathbb{1}_n^\top q = \underset{p\in\Delta_m}{\text{argmax}} \min_{q\in\Delta_n} p^\top Bq + c$$

Since c is a constant, then both optimization problems yield the same solution. The relationship of ii. follows directly from the previous equation. □

The previous theorem showcases that any zero-sum game can be converted to a game with only positive entries in the payoff matrix.

Let $A \in \mathbb{R}^{m\times n}$ be a payoff matrix with only positive entries of a zero-sum game with player I using strategy p and player II using strategy q. If p^*, q^* are optimal strategies of the zero-sum game with payoff matrix A and game value V, then the following linear problems can be used to find both p^*, q^* and the value of the game V.

To find p^* we use the following linear program $\mathcal{P}_1$:

$$\begin{aligned} M = \min_{z\in\mathbb{R}^m} \quad & z^\top \mathbb{1}_m \\ \text{s.t.} \quad & A^\top z \geq \mathbb{1}_n \\ & z \geq 0 \end{aligned}$$

Then, if M is the optimal value of $\mathcal{P}_1$ and z^* the optimal solution of $\mathcal{P}_1$, we have:

$$V = \frac{1}{M}, \qquad p^* = \frac{z^*}{M} = Vz^*$$

Similarly, to find q^* we use the following linear program $\mathcal{P}_2$:

$$\begin{aligned} N = \max_{y\in\mathbb{R}^n} \quad & y^\top \mathbb{1}_n \\ \text{s.t.} \quad & Ay \leq \mathbb{1}_m \\ & y \geq 0 \end{aligned}$$

Then, if N is the optimal value of $\mathcal{P}_2$ and y^* the optimal solution of $\mathcal{P}_2$, we have:

$$V = \frac{1}{N} = \frac{1}{M}, \qquad q^* = \frac{y^*}{N} = Vy^*$$

Example 2.3. Consider the following payoff matrix of a zero-sum game:

$$A = \begin{pmatrix} 8 & 3 & 4 & 1 \\ 4 & 7 & 1 & 7 \\ 0 & 3 & 8 & 5 \end{pmatrix}$$

The following Julia code can be used to find p, q and V:

```
using LinearAlgebra
using JuMP
using GLPK
A = [8 3 4 1; 4 7 1 7; 0 3 8 5]
```

```
#Find p
m = Model(with_optimizer(GLPK.Optimizer))
@variable(m, z[1:3] >=0)
@constraint(m, A'*z .>= ones(4))
@objective(m, Min,  z'*ones(3))
optimize!(m)
obj = objective_value(m)
p_sol = value.(z) ./ obj
V = 1/obj

#Find q
m = Model(with_optimizer(GLPK.Optimizer))
@variable(m, y[1:4] >=0)
@constraint(m, A*y .<= ones(3))
@objective(m, Max,  y'*ones(4))
optimize!(m)
obj = objective_value(m)
q_sol = value.(y) ./ obj
V = 1/obj
```

that yields:

$$V = 0.36111 \qquad p^* = \begin{bmatrix} 0.3611 \\ 0.3333 \\ 0.3055 \end{bmatrix} \qquad q^* = \begin{bmatrix} 0.3333 \\ 0 \\ 0.29629 \\ 0.37037 \end{bmatrix}$$

2.5.2 2×2 games

If we are in the case of a 2×2 matrix:

$$A = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix}$$

we can directly solve for $p = (p_1, 1 - p_1)$ as:

$$p^* = \underset{p \in \Delta_m}{\operatorname{argmax}} \ \min\{a_{11}p_1 + a_{21}(1 - p_1),\ a_{12}p_1 + a_{22}(1 - p_1)\}$$

Assuming that the there are not optimal pure strategies, the minimum of those lines over p will simply be the intersection of the lines. Then, to find p we simply equate both lines:

$$a_{11}p_1 + a_{21}(1 - p_1) = a_{12}p_1 + a_{22}(1 - p_1)$$

A similar analysis can be made for $q = (q_1, 1 - q_1)$:

$$q^* = \underset{q \in \Delta_n}{\operatorname{argmin}} \ \max\{a_{11}q_1 + a_{12}(1 - q_1),\ a_{21}q_1 + a_{22}(1 - q_1)\}$$

or equivalently:

$$a_{11}q_1 + a_{12}(1 - q_1) = a_{21}q_1 + a_{22}(1 - q_1)$$

With the optimal values of p^* or q^* the value of the game V is simply:

$$\begin{aligned} V = \mathbb{E}[\text{payoff}] &= a_{11}p_1^* + a_{21}(1 - p_1^*) \\ &= a_{12}p_1^* + a_{22}(1 - p_1^*) \\ &= a_{11}q_1^* + a_{12}(1 - q_1^*) \\ &= a_{21}q_1^* + a_{22}(1 - q_1^*) \end{aligned}$$

Example 2.4. Consider the game with payoff matrix given by:

$$A = \begin{bmatrix} +1 & -2 \\ -7 & +8 \end{bmatrix}$$

To obtain the game probabilities we repeat the same process. For PI, by equating the expected values (when PII plays strategy one S_1 or two S_2), we have:

$$p - 7(1 - p_1) = -2p_1 + 8(1 - p) \;\rightarrow\; p = \frac{5}{6} \;\rightarrow\; \mathbb{E}[\text{value}] = \frac{5}{6} - 7 \cdot \frac{1}{6} = -\frac{1}{3}$$

Equivalently for PII, by equating the expected values (when PI plays S_1 and S_2), we have:

$$q - 2(1 - q) = -7q + 8(1 - q) \;\rightarrow\; q = \frac{5}{9} \;\rightarrow\; \mathbb{E}[\text{value}] = \frac{5}{9} - 2\frac{4}{9} = -\frac{1}{3}$$

same expected value, as expected from minimax theorem.

Note that, for PII, the worst-case is the maximum of both lines $\max\{3q - 2, 8 - 15q\}$, since PII wants to minimize the maximum payment of both cases. While for PI, the worst-case is the minimum of both lines $\min\{8p - 7, 8 - 10p\}$, since PI wants to maximize the minimum payoff of both cases.

Example 2.5. Consider and solve the game:

$$A = \begin{pmatrix} 10 & 0 & 7 & 1 \\ 2 & 6 & 4 & 7 \\ 6 & 3 & 3 & 5 \end{pmatrix}$$

Observe that column 2 is always smaller than column 4, so C_2 dominates C_4. And we got:

$$A = \begin{pmatrix} 10 & 0 & 7 \\ 2 & 6 & 4 \\ 6 & 3 & 3 \end{pmatrix}$$

Now observe that a mixed strategy of $(1/2, 1/2, 0)$ yields a expected payoff of:

$$\begin{pmatrix} 6 & 3 & 5.5 \end{pmatrix}$$

that is always greater or equal than row 3. Thus, this mixed strategy dominates row 3. Then:

$$A = \begin{pmatrix} 10 & 0 & 7 \\ 2 & 6 & 4 \end{pmatrix}$$

And now for columns 1 and 2 with mixed strategy $(1/2, 1/2)$ we have: $(5, 4)^\top$ that is always less or equal than column 3, thus:

$$A = \begin{pmatrix} 10 & 0 \\ 2 & 6 \end{pmatrix}$$

Then this is a 2×2 system that can be solved by equating payoff with optimal solution $p = (2/7, 2/7, 0)$ and $q = (3/7, 4/7, 0, 0)$ with value $V = 30/7$.

Example 2.6. Consider the game with payoff matrix:

$$A = \begin{pmatrix} 0 & 2 \\ t & 1 \end{pmatrix}$$

For all ranges of the parameter t, find the optimal mixed strategies for both players.

- If $t < 0$, then 0 at $(1, 1)$ is a saddle point, since 0 is the minimum of the first row and the maximum of the first column.
- If $t \in [0, 1]$ then t at $(2, 1)$ is a saddle point, since t is the minimum of the second row, and the maximum of the first column.
- If $t > 1$, then there are no saddle points, so we use equalizing payoffs:

  $$t(1-p) = tp + 1 - p \to p = \frac{t-1}{t+1}$$

  and so $\left(\frac{t-1}{t+1}, \frac{2}{t+1}\right)$ is an optimal strategy for PI with value $\frac{2t}{t+1}$. Similarly for PII:

  $$2(1-q) = tq + 1 - q \to q = \frac{1}{t+1}$$

  and hence the optimal strategy for PII is $\left(\frac{1}{t+1}, \frac{t}{t+1}\right)$.

2.5.3 $2 \times n$ or $m \times 2$ games

In this section we explore how to solve games with payoff matrix of $2 \times n$ or $m \times 2$ size. We should try to reduce the matrix using dominating strategies. After reducing as much as possible we can use a graphical to solve the problem. This is illustrated in the following example.

Example 2.7. Consider the game with the payoff matrix:

$$A = \begin{pmatrix} 0 & 1 & 2 \\ 4 & 3 & 1 \end{pmatrix}$$

It seems a priori that dominating strategies cannot be used used. This is a 2 size game for PI, so we will work solving for PI (in $m \times 2$ games we work for PII). Let $p \equiv (p, 1-p)$ the probabilities of a mixed strategy for PI. The expected payoff for PI under pure strategies of PII are:

$$\mathbb{E}[\text{payoff}] = \min\{4(1-p),\ p + 3(1-p),\ 2p + (1-p)\} = \min\{4 - 4p,\ 3 - 2p,\ 1 + p\}$$

We can plot these lines and find the minimum visually:

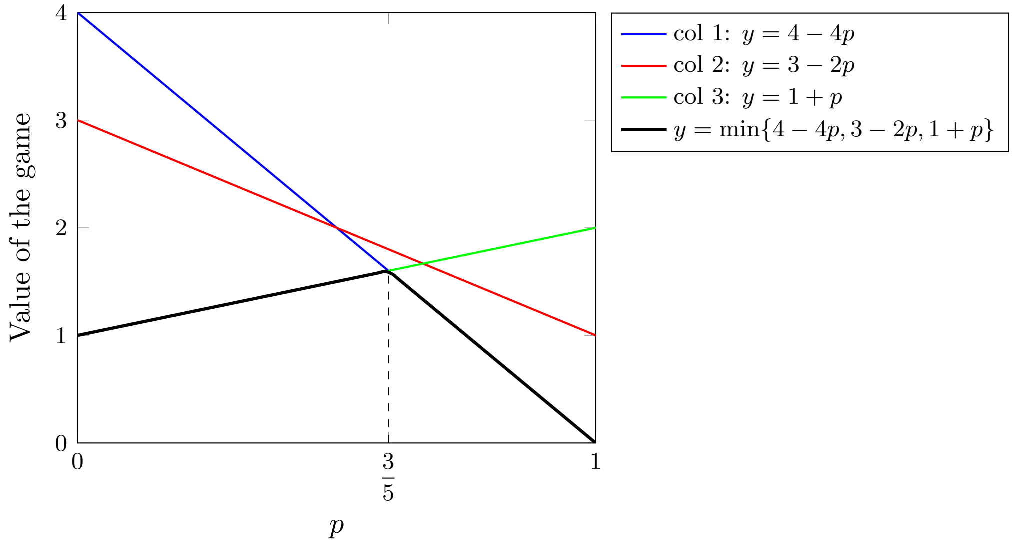

Figure 2.1: Graphical solution of $\max_p \min\{4 - 4p, 3 - 2p, 1 + p\}$.

The minimum is found in the intersection between column 1 and column 3:

$$4 - 4p = 1 + p \;\rightarrow\; p = \frac{3}{5}$$

with expected payoff: $\mathbb{E}[\text{payoff}] = 1+(3/5) = 8/5$, and so the optimal PI strategy is $p^* = (3/5,\ 2/5)$.

Now, to solve for PII, we note that t_2 (the second column) is never used in the optimal strategy for PI. Thus, as PII, we will never play this strategy, and we can discard column 2 of our payoff matrix. With that we got a 2×2 system that we can solve using equating strategies:

$$A = \begin{pmatrix} 0 & 2 \\ 4 & 1 \end{pmatrix}$$

and hence

$$q^* = \arg\min_q \;\; \max\{0 \cdot 2(1-q),\ 4q + 1(1-q)\} = \arg\min_q \;\; \max\{2 - 2q, 1 + 3q\}$$

that is solved: $2 - 2q = 3q + 1 \;\rightarrow\; q = \frac{1}{5}$. Then, the optimal PII strategy is:

$$q = \begin{pmatrix} 1/5 \\ 0 \\ 4/5 \end{pmatrix}$$

with value $\mathbb{E}[\text{loss}] = 2 - 2 \cdot \frac{1}{5} = \frac{8}{5}$ as expected.

2.5.4 Skew-Symmetric Payoff Matrices

Consider the game "Plus-One". In such game, each player chooses from $(1, 2, 3, \ldots, n)$. Then they compare choices. Let r be PI's choice and s be PII choice. If $|r - s| = 1$, then the player with the higher number gets \$1. If $|r - s| > 1$, the player with the lower number gets \$2. In case of a tie, no cash is exchanges. We can set up the payoff matrix as:

$$A = \begin{bmatrix} 0 & -1 & 2 & 2 & \ldots & 2 & 2 \\ 1 & 0 & -1 & 2 & \ldots & 2 & 2 \\ -2 & 1 & 0 & -1 & \ldots & 2 & 2 \\ -2 & -2 & 1 & 0 & \ldots & 2 & 2 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ -2 & -2 & -2 & -2 & \ldots & 0 & -1 \\ -2 & -2 & -2 & -2 & \ldots & 1 & 0 \end{bmatrix}$$

that leads to a skew-symmetric payoff matrix. We observe that row 1 is always better than row 4 for PI, but not only that, any row i, for $i \geq 4$, row 1 has a better payoff for PI. So all rows from 4 to n can be discared since they are dominated by row 1. Similarly, column 1 dominates all columns from 4 to n for PII, that allow us to discard all columns from 4 to n. This leads to the reduced game:

$$A = \begin{bmatrix} 0 & -1 & 2 \\ 1 & 0 & -1 \\ -2 & 1 & 0 \end{bmatrix}$$

Now, we will show the main property of skew-symmetric games:

Theorem 2.2. Let A be an antisymmetric payoff matrix in $\mathbb{R}^{n \times n}$ (i.e. $A^\top = -A$) of a zero-sum game. Let V be the value of the game. Then $V = 0$.

Proof: There exists a vector $p \in \Delta_n$ such that for every $q \in \Delta_n$, it holds that:

$$p^\top A q \geq V$$

and in particular for $p = q$ we have: $p^\top A p \geq V$. By taking the transpose we have:

$$p^\top A^\top p = -p^\top A p \geq V$$

Thus by adding both inequalities we have: $0 \geq 2V \rightarrow V \leq 0$. Now, similarly there exists a vector $\hat{q} \in \Delta_n$ such that for every $\hat{p} \in \Delta_n$, it holds that:

$$\hat{p} A^\top \hat{q} \leq V$$

By taking $\hat{p} = \hat{q}$ we have:

$$\hat{q}^\top A \hat{q} \leq V \rightarrow -\hat{q}^\top A \hat{q} \leq V$$

Adding both inequalities yields that $V \geq 0$ and hence $V = 0$. □

That is, skew-symmetric games have a value of the game equal to zero. In the particular

"Plus-One" game we have the following for PI with a mixed strategy of $p = (p_1, p_2, p_3)$:

$$\begin{aligned} p_2 - 2p_3 &\geq 0 \\ -p_1 + p_3 &\geq 0 \\ 2p_1 - p_2 &\geq 0 \end{aligned}$$

Suppose that one of the inequalities is actually strict. Observe that if we add the first and third equation, and then add two times the second one we got $0 > 0$, that is a contradiction. Thus, no inequality can be strict and hence all of them are equalities. Using two equalities and that p has to be in the simplex we have:

$$\begin{aligned} p_1 + p_2 + p_3 &= 1 \\ -p_1 + p_3 &= 0 \\ 2p_1 - p_2 &= 0 \end{aligned}$$

that yields $p_1 + 2p_1 + p_1 = 1 \to p_1 = 1/4$, and hence $p_2 = 1/2$ and $p_3 = 1/4$.

2.6 Nash Equilibrium

Let (p^*, q^*) be the optimal mixed strategies for players I and II respectively for a zero-sum game with payoff matrix $A \in \mathbb{R}^{m \times n}$. We say that (p^*, q^*) are:

- Optimal strategies: Since they attain the min and max respectively for each player.
- Safety strategies: Since they make the best of the worst case scenario for each player.
- Mixed Nash Equilibrium: Since if PI is playing p^*, for PII the best strategy is to play q^*. Similarly, if PII is playing q^*, then for PI the best strategy is to play p^*. That is, for a Nash equilibirum, there is no incentive to unilaterally change his strategy.

2.7 Von Neumann's Minimax Theorem

Define for player I the function that is the worst case scenario. That is, for any given mixed strategy p that PI chooses, PII will pick the counter strategy:

$$v_I(p) = \min_{q \in \Delta_n} p^\top A q$$

We say that is a safety strategy because p^* makes best of worst case scenario:

$$p^* = \arg\max_{p \in \Delta_m} v_I(p) = \arg\max_{p \in \Delta_m} \min_{q \in \Delta_n} p^\top A q$$

We say that the value of the game for PI is then:

$$V_I = \max_{p \in \Delta_m} \min_{q \in \Delta_n} p^\top A q$$

Similarly for PII:

$$v_{II}(q) = \max_{p \in \Delta_n} p^\top A q$$

we say that q^* makes the best of the worst case scenario:

$$q^* = \arg\min_{q\in\Delta_n} \; v_{II}(q) = \arg\min_{q\in\Delta_n}\max_{p\in\Delta_n} p^\top Aq$$

with the value of the game for PII as:

$$V_{II} = \min_{q\in\Delta_n}\max_{p\in\Delta_n} p^\top Aq$$

Von Neumann's Minimax Theorem states that $V_I = V_{II}$. To prove it, we will use the following lemmas.

Lemma 2.1. Separating Hyperplane Lemma: Let S_1 and S_2 be two disjoint ($S_1 \cap S_2 = \emptyset$) convex and close sets. Then there exists a hyperplane that can separate both sets:

$$\mathcal{H} = \{x \in \mathbb{R}^n : \; a^\top x = d\}$$

That is, for any $y \in S_1$ we have $a^\top y \leq d$, while for every $z \in S_2$ we have $a^\top z > d$. The proof can be found in [4].

Lemma 2.2. The ∞-norm is the dual norm of the 1-norm: For a fixed q, it is true that:

$$\max_{p\in\Delta_m} p^\top Aq = \max_{1\leq i\leq m}\{(Aq)_i\}$$

on which $(Aq)_i$ is the i-th entry of the vector Aq.

Proof: Let $Aq = x$. Then we want to solve:

$$\max_{p\in\Delta_m} p^\top x = \max_{p\in\Delta_n} p_1x_1 + p_2x_2 + \ldots + p_mx_m$$

Since p has to be in the simplex, we can sort x from the highest to the lowest. Let $x_{[i]}$ be the i-th position of the sorted x vector. Since $\sum_i p_i = 1$, the best we can do is to pick $p_{[1]} = 1$ and all the rest zero to maximize that function. Then, for sure the maximum is attained in one of the coordinate vectors, i.e.:

$$\max_{p\in\Delta_m} p^\top x = \max_{1\leq i\leq m}\{x_i\} = \max_{1\leq i\leq m}\{(Aq)_i\} \quad \square$$

Theorem 2.3. Von Neumann's Minimax Theorem: Let $A \in \mathbb{R}^{m\times n}$ be the payoff matrix of a zero-sum game. Let V_I be the value of the game for PI defined as:

$$V_I = \max_{p\in\Delta_m}\min_{q\in\Delta_n} p^\top Aq = \max_{p\in\Delta_n} v_I(p)$$

and V_{II} be the value of the game for PII defined as:

$$V_{II} = \min_{q\in\Delta_n}\max_{p\in\Delta_n} p^\top Aq = \min_{q\in\Delta_n} v_{II}(q)$$

then $V_I = V_{II}$.

Proof [adapted from David Blackwell's notes]: For each $q \in \Delta_n$ associate the point $w \equiv w(q) := Aq \in \mathbb{R}^m$. Define $\mathcal{W}$ as:

$$\mathcal{W} = \{w \in \mathbb{R}^m : \ w = Aq \text{ for } q \in \Delta_n\}$$

since Δ_n is closed bounded and convex and $w(q)$ is linear, then $\mathcal{W}$ is closed, bounded and convex. Now it is clear that:

$$v_I(p) = \min_{q \in \Delta_n} p^\top Aq \le \max_{p \in \Delta_m} p^\top Aq = v_{II}(q)$$

since $\min_q p^\top Aq \le p^\top Aq$ for any fixed p, in particular for p^* that maximize $p^\top Aq$ for any q. Then since $v_I(p) \le v_{II}(q)$ for any p and q, it holds for the maximum and minimum respectively, and so:

$$V_I \le V_{II}$$

To show the other direction note the following:

$$\begin{aligned} V_{II} &\le v_{II}(q) \\ &= \max_{p \in \Delta_n} p^\top Aq \\ &= \max_{1 \le i \le m} \{(Aq)_i\} \text{ by Lemma 2.2} \end{aligned}$$

Let $w = Aq \in \mathcal{W}$. Then:

$$V_{II} \le \max_{1 \le i \le m} \{w_i\}$$

Let z be some number such that $z < V_{II}$. Define the set $S \subseteq \mathbb{R}^m$ as:

$$s \in S \Leftrightarrow \max_{1 \le i \le m} \{s_i\} \le z$$

Thus, $\max_{1 \le i \le m} \{s_i\} < V_{II} \le \max_{1 \le i \le m} \{w_i\}$. Note also that $s \notin \mathcal{W}$, since z is strictly lower than V_{II}. Since S is also convex and closed, and S is disjoint with $\mathcal{W}$ we can use the separating hyperplane lemma. That is $\exists u \in \mathbb{R}^m$ and $d \in \mathbb{R}$ such that:

$$u^\top w > d, \ \forall w \in \mathcal{W} \qquad\qquad u^\top s < d, \ \forall s \in \mathcal{S}$$

Observe that note all u_i cannot be zero, since even if $d = 0$, we have that $0 > 0$ and $0 < 0$, that is not true. In addition, note that some $u_j < 0$ is also not possible, since we can take an s_j sufficiently negative, that no matter what d is chosen, $u_j s_j$ will always be larger than d (since S is not bounded from below).
Now, normalizing the hyperplane such that $\sum_i u_i = 1$, and considering that $u_i \ge 0$ and not all zero, implies that $u \in \Delta_m$, i.e.

$$u^\top w = u^\top Aq > d \ \ \forall q \in \Delta_n$$

and in particular for q^*:

$$v_I(u) = \min_{q \in \Delta_n} u^\top Aq = u^\top Aq^* \le V_I, \ \forall u \in \Delta_m$$

since $V_I = \max_{u \in \Delta_m} v_I(u)$. Thus, $V_I > d$.
Now, let $s \in S$ such that $s_i = z$, $\forall i = 1 \dots m$. Then:

$$u^\top s = \sum_i u_i z = z < d$$

i.e. $z < d < V_I$. Now, since this is true for any $z < V_{II}$, in particularly any sufficiently close to V_{II} we have:

$$V_{II} \leq V_I$$

and hence $V_{II} = V_I$. $\square$.

2.8 Equilibrium Theorem

Another technique to solve zero-sum games is using the Equilibrium Theorem (Prop. 2.5.3 in [1]):

Theorem 2.4. Consider a game with a payoff matrix $A \in \mathbb{R}^{m \times n}$ and value V. Let $p = (p_1, \dots, p_m)$ and $q = (q_1, \dots, q_n)$ be optimal (or safe) mixed strategies. Then:

$$p\top Ae_j = \sum_{i=1}^{m} p_i a_{ij} = V \qquad \text{for all } j \text{ for which } q_j > 0$$

For the cases on which $q_j = 0$ we have:

$$p^\top Ae_j = \sum_{i=1}^{m} p_i a_{ij} \geq V \qquad \text{for all } j \text{ for which } q_j = 0$$

That is, for player I, if player II column is active in his optimal mixed strategy, then for PI, the payoff of his mixed strategy over any active column is the same.

Similarly, for PII it holds that:

$$e_i^\top Aq = \sum_{j=1}^{n} a_{ij} q_j = V \qquad \text{for all } i \text{ for which } p_i > 0$$

For the cases on which $p_i = 0$ we have:

$$e_i^\top Aq = \sum_{j=1}^{n} a_{ij} q_j \leq V \qquad \text{for all } i \text{ for which } p_i = 0$$

That is, for player II, if player I row is active in his optimal mixed strategy, then for PII, the payoff of his mixed strategy over any active column is the same.

Proof [**2**]**:** Suppose there is a k such that $p_k > 0$ and $\sum_{j=1}^{n} a_{kj} q_j \neq V$. Since q is an optimal strategy, it must hold that $\sum_{j=1}^{n} a_{kj} q_j < V$. However, when both players play their optimal strategies, it is true that:

$$V = \sum_{i=1}^{m} p_i \left(\sum_{j=1}^{n} a_{ij} q_j \right) < \sum_{i=1}^{m} p_i (V) = V$$

The inequality is strict since it strict for the k-th term of the sum. This contradiction proves the second statement. The first statement follows analogously. □

Example 2.8. Consider the following payoff matrix:

$$A = \begin{pmatrix} 8 & 3 & 4 & 1 \\ 4 & 7 & 1 & 7 \\ 0 & 3 & 8 & 5 \end{pmatrix}$$

In this game there is no chance to easily use domination. The idea to solve it is by trying which strategies will be active. In this case, a hint is that $q_2 = 0$. That, is the second column is not active. With that we can write the following system of equations:

$$8p_1 + 4p_2 = V \tag{1}$$

$$4p_1 + p_2 + 8p_3 = V \tag{2}$$

$$p_1 + 7p_2 + 5p_3 = V \tag{3}$$

$$p_1 + p_2 + p_3 = 1 \tag{4}$$

Subtracting (1) - (2) and (2) - (3) we got:

$$4p_1 + 3p_2 - 8p_3 = 0 \tag{5}$$

$$3p_1 - 6p_2 + 3p_3 = 0 \tag{6}$$

Then the system of equations using (4), (5) and (6) can be written as:

$$\begin{bmatrix} 1 & 1 & 1 \\ 4 & 3 & -8 \\ 3 & -6 & 3 \end{bmatrix} p = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}$$

Then we can pivot this matrix. Subtract 4 times the row 1 to row 2, and subtract 3 times the row 1 to row 3:

$$\begin{bmatrix} 1 & 1 & 1 \\ 0 & -1 & -12 \\ 0 & -9 & 0 \end{bmatrix} p = \begin{bmatrix} 1 \\ -4 \\ -3 \end{bmatrix}$$

From the third row is clear that $p_2 = 1/3$. From the second row we have that:

$$-p_2 - 12p_3 = -4 \ \rightarrow \ p_3 = \frac{4 - p_2}{12} = \frac{11}{36}$$

and hence:

$$p_1 = 1 - p_2 - p_3 = 1 - \frac{1}{3} - \frac{11}{36} = \frac{36 - 12 - 11}{36} = \frac{13}{36}$$

The value of the game assuming that $q_2 = 0$ is:

$$V = 8p_1 + 4p_2 = \frac{8 \cdot 13}{36} + \frac{4 \cdot 1}{3} = \frac{26}{9} + \frac{4}{3} = \frac{26 + 12}{9} = \frac{38}{9}$$

Checking for the second column we have:

$$3p_1 + 7p_2 + 3p_3 = \frac{13}{3} = \frac{39}{9} \geq \frac{38}{9} = V$$

Thus, this is a valid solution. In summary:

$$p^* = \begin{pmatrix} 13/36 \\ 1/3 \\ 11/36 \end{pmatrix} \qquad V = \frac{38}{9}$$

For PII, using the same argument and similar system of equations, the optimal strategy can be found as:

$$q^* = \begin{pmatrix} 1/3 \\ 0 \\ 8/27 \\ 10/27 \end{pmatrix}$$

Note that this game was already covered in Example 2.3 using linear programming. Both methods yield the same results.

3 General Sum Games

We now study general sum games, on which players can have different payoffs depending on their actions and hence not necessarily the payoff sum to zero.

3.1 2×2 Two-person General Sum Games

This games are specified by a Bi-matrix payoff for each player on which both want to maximize. In this scenario, both players want to maximize their respective payoff.

Example 3.1. Consider the following payoff bi-matrix:

$$\begin{bmatrix} (2,0) & (1,3) \\ (0,1) & (3,2) \end{bmatrix}$$

we can write this a two matrices:

$$A = \begin{bmatrix} 2 & 1 \\ 0 & 3 \end{bmatrix} \qquad B = \begin{bmatrix} 0 & 3 \\ 1 & 2 \end{bmatrix}$$

Observe each entry a_{ij} represents the payoff for PI and b_{ij} the payoff for PII. Our first analysis for these games is using arrows to analyze player preferences. Vertical arrows are used to analyze PI preferences (i.e. choosing which row) and horizontal arrows to analyze PII preferences:

$$\begin{bmatrix} (2,0) & \rightarrow & (1,3) \\ \uparrow & & \downarrow \\ (0,1) & \rightarrow & \boxed{(3,2)} \end{bmatrix}$$

From this we observe that $(3,2)$ is a pure Nash equilibrium, since at that point no player has incentive to modify its own strategy unilaterally.

3.1.1 Finding Mixed Nash Equilibrium

The key idea is to find a strategy $p = (p, 1-p)$ for PI that makes PII indifferent for any of his strategies.

Example 3.2. Consider the game with the following payoff matrices:

$$A = \begin{bmatrix} -1 & 1 \\ 2 & 0 \end{bmatrix} \qquad B = \begin{bmatrix} 0 & 3 \\ 2 & 1 \end{bmatrix}$$

Then for $p = (p, 1-p)$ we have:

$$p^\top B = (p \quad 1-p) \begin{bmatrix} 0 & 3 \\ 2 & 1 \end{bmatrix} = \begin{bmatrix} 2(1-p) \\ 3p + (1-p) \end{bmatrix}^\top$$

and we make each column equal to each other:

$$2 - p = 3p + 1 - p \rightarrow 2 - 2p = 1 + 2p \rightarrow p = 1/4$$

Equivalently for PII, he can pick $q = (q, 1-q)$ to make PI indifferent:

$$Aq = \begin{bmatrix} -1 & 1 \\ 2 & 0 \end{bmatrix} \begin{bmatrix} q \\ 1-q \end{bmatrix} = \begin{bmatrix} -2 + (1-q) \\ 2q \end{bmatrix}$$

Equating rows we have:

$$-q + (1-q) = 2q \rightarrow q = 1/4$$

Then, a mixed Nash Equilibrium is given by:

$$p^* = \begin{bmatrix} 1/4 \\ 3/4 \end{bmatrix}, \qquad q^* = \begin{bmatrix} 1/4 \\ 3/4 \end{bmatrix}$$

With that the expected payoff for PI is:

$$p^{*\top} Aq* = 1/2$$

and for PII:

$$p^{*\top} Bq* = 3/2$$

3.1.2 Nash Equilibrium

We define a Nash equilibrium as follows:

$(p^*, q^*) \in \Delta_m \times \Delta_n$ is a Nash Equilibrium if:

1. $(p^{*\top})Aq^* \geq p^\top Aq^*,\ \forall p \in \Delta_m$
2. $(p^*)^\top Bq^* \geq (p^*)^\top Bq,\ \forall q \in \Delta_n$

that is, if PII is playing q^*, PI has no incentive to change his strategy p^*. Equivalently, if PI is playing p^*, PII has no incentive to modify from his strategy q^*.

We say that if $A = B^\top$, the game is symmetric with symmetric Nash Equilibria.

3.1.3 Pareto Optimality

An outcome is non-Pareto optimal (or Pareto inferior) if there is another outcome that gives both players better payoff, or at least one stay the same and the other improves. To find we use the convex hull of the strategies.

Example 3.3. Consider the following bimatrix:

$$\begin{bmatrix} (2,3) & (3,5) \\ (2,2) & (1,3) \end{bmatrix}$$

Its convex hull is given by:

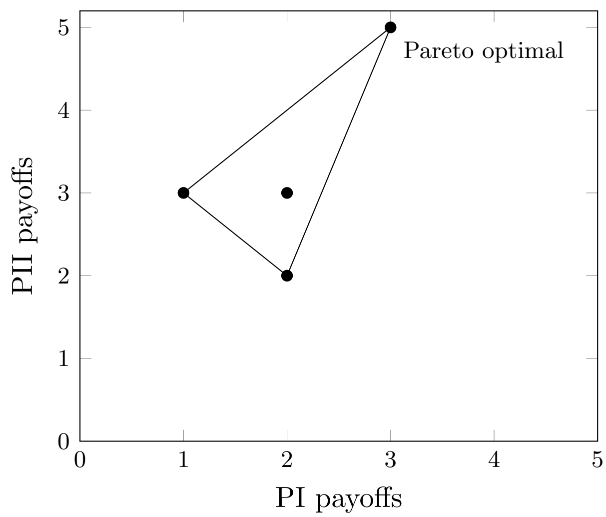

Figure 3.1: Convex hull of example 3.3.

It is clear that (3,5) is Pareto optimal (and also a Nash Equilibrium), since all the possible payoffs given by the convex hull is the better for each player.

Example 3.4. Consider the following bi-matrix:

$$\begin{bmatrix} (2,3) & (1,5) \\ (2,2) & (0,3) \end{bmatrix}$$

Its convex hull is given by:

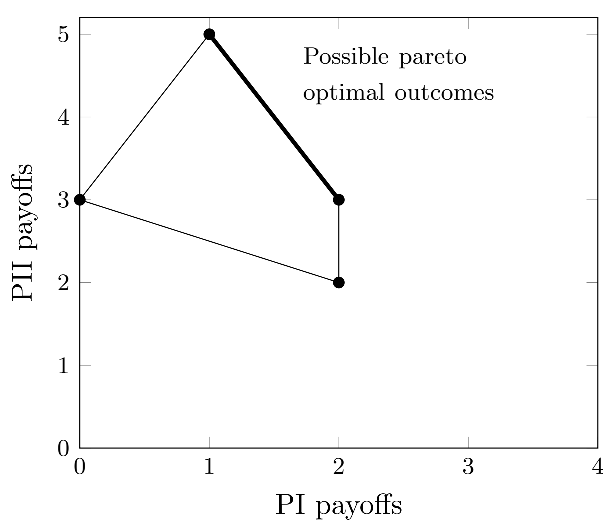

Figure 3.2: Convex hull of example 3.4.

In this case the line between $(1,5)$ and $(2,3)$ are all possible pareto optimal outcomes. Nevertheless, not all of those outcomes can be achieved, and will depend on the specific problem.

3.1.4 Battle of the sexes (1957)

In this game the set-up is as follows: a husband wants to watch a prize fight (P) while his wife wants to go to the ballet (B). The payoff matrix is the following one:

	Husband B	Husband P
Wife B	(4,1)	(0,0)
Wife P	(0,0)	(1,4)

Table 3.1: Battle of the sexes payoff matrix.

By setting the arrows, it is clear than both (B,B) or (P,P) are pure Nash Equilibria since:

$$\begin{bmatrix} \boxed{(4,1)} & \leftarrow & (0,0) \\ \uparrow & & \downarrow \\ (0,0) & \rightarrow & \boxed{(1,4)} \end{bmatrix}$$

To find mixed pure strategies we define:

$$Aq \rightarrow 4q = 1 - q \rightarrow q = 1/5 \qquad p^\top B \rightarrow p = 4(1-p) \rightarrow p = 4/5$$

Then, a mixed Nash Equilibria is given by $p^* = (4/5, 1/5)$ and $q^* = (1/4, 4/5)$ with payoffs:

$$\frac{4}{5} = p^{*\top} A q^* = p{*}^\top B q^*$$

With that we have the following convex hull:

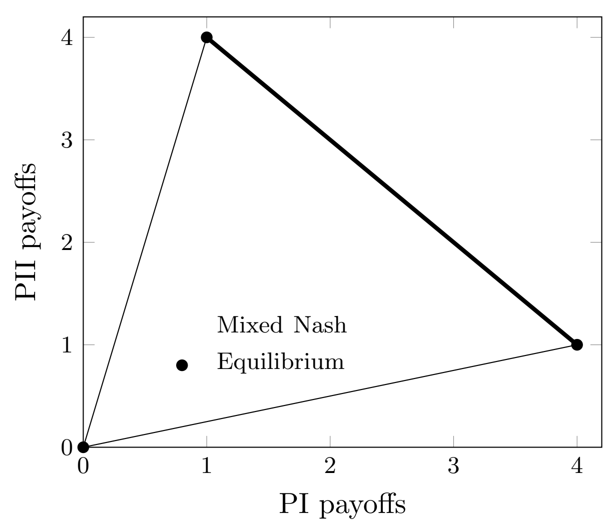

Figure 3.3: Convex hull of battle of sexes.

Sadly, we see that the mixed Nash equilibrium is worse for both wife and husband, so it is pareto inferior.

3.1.5 Game of chicken or the Cuban-Missile crisis

In this game there are two options for each player, to persist or to swerve. The classical game of chicken, consider two chicken running into each other, and both of them can persist on running (and hence crushing) or one (or both) swerve. The analogy is with the Cuban-missile crisis. Both the US and USSR were on the coast of Cuba with their floats. The US want to block the coast and persist to avoid the entrance of the USSR. We say that if both persist then a nuclear war occurs. The payoff bi-matrix is set up as follows:

	USSR Swerve	USSR Persist
US Swerve	(3,3)	(2,4)
US Persist	(4,2)	(1,1)

Table 3.2: Cuban-missile crisis payoff matrix.

It is clear that both (S,P) and (P,S) are pure Nash Equilibrium, and was exactly what happened. The US were to persist no matters what, so at the end the USSR preferred to swerve.

3.1.6 Computing pareto optimal candidates

Consider the following bi-matrix:

$$\begin{bmatrix} (0,0) & (1,0) \\ (0,1) & (0,0) \end{bmatrix}$$

it is clear that both (0,1) and (1,0) are pure Nash equilibria. Then a point on the line from these points are Pareto optimal candidates. This line is described by $x + y = 1$, so any point in this line can be described as the points $(t, 1-t)$ for $t \in [0,1]$. Let $p = (p, 1-p)$ and $q = (1-q, q)$. Observe that the payoffs are then:

$$(p^\top Aq, p^\top Bq) = (pq, (1-p)(1-q)) = (t, 1-t)$$

Solving this system only yields $p = q = 0$ or $p = q = 1$ that implies that only the pure Nash equilibria are pareto optimal.

This example showcase that those points are only candidates to be pareto optimal, but in practice it is feasible that no probabilities can yield any point there.

In addition, there can be other mixed pareto optimal candidates outside those two, but finding those may not be trivial.

3.1.7 Freeloader problem

A public service is worth \$3 to each player at a cost of \$4 if at least one player volunteer towards the cost. If both volunteer they share the cost of the service:

	chip-in		freeload
chip-in	$(3-2,3-2)$	$\rightarrow$	$(3-4,3)$
	$\downarrow$		$\downarrow$
freeload	$(3,3-4)$	$\rightarrow$	$\boxed{(0,0)}$

Table 3.3: Freeloader problem.

Thus, (0,0) is a pure Nash Equilibrium. We then try to find mixed Nash equilibrium:

$$Aq \rightarrow q-(1-q)=3q \rightarrow q=-1$$

then there is no other solution and thus (0,0) is the only Nash equilibrium (sadly).

3.1.8 Antelope Hunting

Consider the Antelope hunting problem, on which two cheetahs (PI and PII) must choose to go to two antelopes, one large (ℓ) and one small (s). If they both go for the same, then they must share the antelope. With that the payoff matrix can be written as follow:

	ℓ	s
ℓ	$(\ell/2,\ell/2)$	(ℓ,s)
s	(s,ℓ)	$(s/2,s/2)$

Table 3.4: Antelopes hunting payoff matrix.

It is clear that if $\ell > 2S \rightarrow \ell/2 > s$ and then there is a Nash equilibrium for both cheetahs to just go to the large one. Now if $s < \ell < 2s$, there a two pure Nash equilibria: (s,ℓ) and (ℓ,s). To find a mixed one:

$$Aq \rightarrow \frac{\ell}{2}q + L(1-q) = sq + \frac{s}{2}(1-q) \rightarrow L - \frac{s}{2]} = \left(\frac{\ell+s}{2}\right) q \rightarrow q = \frac{2\ell-s}{l+s} = q$$

and since $A = B^{\top}$, we have

$$p^* = q^* = \begin{pmatrix} \dfrac{2\ell-s}{\ell+s} \\ \\ \dfrac{2s-\ell}{\ell+s} \end{pmatrix}$$

3.2 Pure Nash Equilibrium for two-players larger games

Consider the following bi-matrix:

$$\begin{bmatrix} (5,2) & (3,0) & (\boxed{8},1) & (\boxed{2},\overline{3}) \\ (6,3) & (\boxed{5},\overline{4}) & (7,\overline{4}) & (1,1) \\ (\boxed{7},5) & (4,6) & (6,\overline{8}) & (0,2) \end{bmatrix}$$

In the previous it is boxed the maximum on each column in the first entry, that is the best strategy for PI given a fixed strategy of PII. There is an overline on the maximum of each row in its second entry, that is the best strategy for PII given a fixed strategy of PI. With that, it is clear than there are two pure Nash Equilibria (5,4) and (2,3).

3.3 $k-$ person games

We now explore the games with more than two players.

- Each player has a set of pure strategies S_i.
- Payoffs are described by utility functions for each player:

$$u_i : S_1 \times S_2 \times \dots S_k \to \mathbb{R}$$

  i.e. $u_i(s_1, s_2, \dots, s_k) \in \mathbb{R}$

Example 3.5. Consider the $2 \times 2 \times 2$ zero sum game described by two tri-matrices. Each player has 2 possible actions A and B. We can write the payoff matrices as follow:

$$\text{PIII} = A : \begin{bmatrix} \nwarrow(1,1,-2) & \rightarrow & (-4,3,1)\nearrow \\ \downarrow & & \uparrow \\ \nearrow\boxed{(2,-4,2)} & \leftarrow & (-5,-5,10)\nwarrow \end{bmatrix}$$

$$\text{PIII} = B : \begin{bmatrix} \searrow\boxed{(3,-2,-1)} & \leftarrow & (-6,-6,12)\swarrow \\ \uparrow & & \downarrow \\ \swarrow(2,2,-4) & \rightarrow & (-2,3,-1)\searrow \end{bmatrix}$$

In this case, vertical arrows are used to denote preferences of PI, horizontal arrows are used to denote preferences of PII while diagonal arrow (into or out of the matrix) are used to denote preferences of PIII. As can be seen, there are two pure Nash equilibria, given by (B, A, A) and (A, A, B).

Now we explore mixed strategies for coalitions between players. For example, assume that PII and PIII form a coalition of the form:

	PII,PIII $= (A, A)$	PII,PIII $= (A, B)$	PII,PIII $= (B, A)$	PII,PIII $= (B, B)$
PI: A	1	3	-4	-6
PI: B	2	2	-5	-2

Table 3.5: Game for PI if PII and PIII form a coalition.

Note that (B, A) dominates both (A, A) and (A, B). Thus this is a 2×2 zero sum game that can be solved by equating payoffs:

$$V = -4.4, \qquad \text{PI} = \begin{bmatrix} 3/5 \\ 2/5 \end{bmatrix}, \qquad \text{PII-PIII} = \begin{bmatrix} 0 \\ 0 \\ 4/5 \\ 1/5 \end{bmatrix}$$

Similarly for PI and PIII coalition:

	PI,PIII $= (A, A)$	PI,PIII $= (A, B)$	PI,PIII $= (B, A)$	PI,PIII $= (B, B)$
PII: A	1	-2	$\boxed{-4}$	2
PII: B	3	-6	-5	3

Table 3.6: Game for PII if PI and PIII form a coalition.

on which we note that $A, (B, A)$ is a saddle point and hence:

$$V = -4, \qquad \text{PII} = \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \qquad \text{PII-PIII} = \begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \end{bmatrix}$$

And finally for PI and PII coalition:

	PI,PII $= (A, A)$	PI,PII $= (A, B)$	PI,PII $= (B, A)$	PI,PII $= (B, B)$
PIII: A	-2	1	2	10
PIII: B	-1	12	-4	-1

Table 3.7: Game for PIII if PI and PII form a coalition.

with solution:

$$V = -1.43, \qquad \text{PIII} = \begin{bmatrix} 3/7 \\ 4/7 \end{bmatrix}, \qquad \text{PI-PII} = \begin{bmatrix} 6/7 \\ 0 \\ 1/7 \\ 0 \end{bmatrix}$$

Now, in the first case of PII-PIII form the coalition. How they should split in the NE?

Using the probabilities let's compute the expected payoffs of the first case $p = (3/5, 2/5)$ and $q = (0, 0, 4/5, 1/5)$ for all players:

$$\begin{aligned} \mathbb{E}[(\text{PI,PII,PIII})] = & \frac{3}{5} \cdot 0 \cdot \underbrace{(1, 1, 2)}_{(A,A,A)} + \frac{3}{5} \cdot 0 \cdot \underbrace{(3, -2, -1)}_{A,A,B} + \frac{3}{5} \cdot \frac{4}{5} \cdot \underbrace{(-4, 3, 1)}_{(A,B,A)} + \cdot 35 \cdot \frac{1}{5} \cdot \underbrace{(-6, -6, 12)}_{(A,B,B)} \\ & + \frac{2}{5} \cdot 0 \cdot \underbrace{(2, -4, 2)}_{(B,A,A)} + \frac{2}{5} \cdot 0 \cdot \underbrace{(2, 2, -4)}_{(B,A,B)} + \frac{2}{5} \cdot \frac{4}{5} \cdot \underbrace{(-5, -5, 10)}_{(B,B,A)} + \frac{2}{5} \cdot \frac{1}{5} \cdot \underbrace{(-2, -3, -1)}_{(B,B,B)} \\ = & (-4.4, -0.64, 5.04) \end{aligned}$$

And similarly for PI-PIII coalition:

$$\mathbb{E}[\text{payoffs}] = (2, -4, 2)$$

and for PI-PII coaltion:

$$\mathbb{E}[\text{payoffs}] = (2.12, -0.69, -1.43)$$

We then we want to explore fully mixed equilibria. The following sections will explore how to compute those.

Notation: Let us define the simplex for player i, $i \in \{1, 2, \ldots, k\}$, denoted as $\Delta_{|S_i|}$, where S_i is the set of pure strategies for player i. As mentioned before, each player has a utility function $u_i : S_1 \times \ldots \times S_k \to \mathbb{R}$. We denote the strategy played for player i as s_i, and the strategy of all other players as:

$$s_{-i} = (s_1, s_2, \ldots, s_{i-1}, s_{i+1}, \ldots, s_k)$$

i.e. the set of all strategies without the strategy of PI. Abusing some notation we can write the utility function of player i as:

$$u_i(s_i, s_{-i}) \in \mathbb{R}$$

3.3.1 Pure Nash Equilibria

We say that $s^* = (s_1^*, s_2^*, \ldots, s_k^*)$ is a Pure Nash Equilibria (PNE) if for every i the following is satisfied:

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

That is, for each player playing the strategy s_i^* is the best strategy given that all other players are playing s_{-i}^*. Equivalently, we said that no player can find unilaterally a better strategy.

3.3.2 Mixed Nash Equilibria

Let $p_i \in \Delta_{|S_i|}$ a discrete probability distribution over the set of pure strategies S_i of player i. That is, the j-th entry in $p_i(s_j)$, is the respective probability of the strategy $s_i(j)$ (the j-th strategy of player i). We then define the expected payoff as follows:

$$\bar{u}_i(p_1, p_2, \ldots, p_k) = \bar{u}_i(p_i, p_{-i}) = \sum_{\substack{(s_1,\ldots,s_k) \\ \in S_1 \times \ldots \times S_k}} u_i(s_1, \ldots s_k) p_1(s_1) p_2(s_2) \ldots p_k(s_k)$$

where the sum goes over all permutations over pure strategies. Then, a Mixed Nash Equilibrium (MNE), $p^* = (p_1^*, p_2^*, \ldots, p_k^*)$ where each $p_i^* \in \Delta_{|S_i|}$ for every $i \in \{1, \ldots, k\}$, satisfies that:

$$\bar{u}_i(p_i^*, p_{-i}^*) \geq \bar{u}_i(p_i, p_{-i}^*), \quad \forall p_i \in \Delta_{|S_i|}, \forall i \in \{1, \ldots, k\}$$

3.3.3 Tragedy of the commons

The tragedy of the commons are cases or examples when each player behaving rationally individually can have a significant worse result than considering collaboration.

Consider 10 families herding goats on 1 square mile of common land. The goats graze on grass and return buckets of milk. A goat that grazes a fraction a of land returns b buckets of milk, where:

$$b = e^{1 - \frac{1}{10a}}$$

Note that if $a = 1/10$, then $b = 1$. A social planner hired to optimize the total number of goats, where each goat grazes on land $a = \frac{1}{N}$. Given the production function, the optimal number is 10, that is, one goat per family. However, families can freely decide how many goats to keep g, ignoring the social planner recommendation. So for each family the function to optimize without considering other families decision:

$$b = g e^{1-(g+G)/10}$$

Optimizing over g yields $g = 10$ for each family, and hence $G = 100$. With that the result is that each family receives:

$$b = 10e^{1-110/10} = 0.012$$

buckets of milk.

3.3.4 Nash's Theorem

The following theorem states the Nash's Theorem:

Theorem 3.1. Every finite general-sum game has a Nash equilibrium.

The proof involves the usage of Kakutani fixed-point theorem or Brouwer fixed-point theorem. Both proofs, can be found on Wikipedia.

3.3.5 Computing mixed Nash Equilibria

This result is referred as Lemma 4.3.7 in [1]:

Theorem 3.2. Let T_i be set of active pure strategies por player i under mixed strategy p_i:

$$T_i = \{s_i \in S_i : \; p_i(s_i) > 0\}$$

Then p_i^* is a mixed Nash Equilibrium if and only if $\forall i = 1, \ldots, k$ there exists c_i such that:

- $\forall s_i \in T_i$ it holds that:
$$\bar{u}_i(s_i, p_{-i}^*) = c_i$$
- $\forall s_i \notin T_i$ it holds that:
$$\bar{u}_i(s_i, p_{-i}^*) \leq c_i$$

Example 3.6. Consider three firms that can purify or pollute. Their respective payoffs are the following:

$$\text{PIII} = \text{pur} : \begin{bmatrix} (1,1,1) & (1,0,1) \\ (0,1,1) & (3,3,4) \end{bmatrix}$$

$$\text{PIII} = \text{pol} : \begin{bmatrix} (1,1,0) & (4,3,3) \\ (3,4,3) & (3,3,3) \end{bmatrix}$$

Drawing the arrows it can be seen that if two firms purifying and one firm polluting is a pure NE. In addition, everybody polluting is also a pure NE. That gives 4 pure NE. In addition, consider that one firm decides to purify all the time. This yields a bi-matrix for game for the other players. For instance consider that firm III decides to purify all the time. Then the following bi-matrix game for PI and PII occurs:

$$\begin{bmatrix} (1,1) & (1,0) \\ (0,1) & (3,3) \end{bmatrix}$$

that can be solved by equating payoff of a player strategy in the other payoff matrix. With that:

$$\begin{bmatrix} 1 & 1 \\ 0 & 3 \end{bmatrix} p_2 \to p_2 + (1 - p_2) = 3(1 - p_2) \to p_2 = \frac{2}{3}$$

and since it is a symmetric game:

$$p_1^* = p_2^* = \begin{bmatrix} 2/3 \\ 1/3 \end{bmatrix}$$

With that we have 3 partially mixed NE, depending on which firm decides to purify all the time.

Finally, we are interested in computing fully mixed strategies. Assuming that all strategies are active we can state the following:

$$u_1(\text{purify}, p_2^*, p_3^*) = u_1(\text{pollute}, p_2^*, p_3^*) = c_1$$

that is:

$$\begin{aligned} 1p_2p_3 + 1(1 - p_2)p_3 + 1p_2(1 - p_3) + 4(1 - p_2)(1 - p_3) = \\ 0p_2p_3 + 3(1 - p_2)p_3 + 3p_2(1 - p_3) + 3(1 - p_2)(1 - p_3) \end{aligned}$$

Simplifying terms:

$$1 = 3(p_2 + p_3 - 2p_2p_3)$$

Similarly, posing the same equations for the other cases:

$$\begin{aligned} 3(p_1 + p_3 - 2p_1p_3) = 1 \\ 3(p_1 + p_2 - 2p_1p_2) = 1 \end{aligned}$$

Subtracting two equations:

$$(p_1 - p_3)(1 - 2p_2) = 0$$

There are two possible cases here. The first one is to set $p_2 = 1/2$, the first third equation has:

$$3(p_1 + 1/2 - p_1) = 1 \to 3/2 = 1$$

that is a contradiction. Now, if we set $p_1 = p_3$ and replacing in the second equation we got:

$$2p_1 - p_1^2 = \frac{1}{3} \to 6p_1^2 - 6p_1 + 1 = 0$$

Subtracting different equations we have that $p_1 = p_2 = p_3$, and solving the quadratic equation we have:

$$p_1 = p_2 = p_3 = \frac{3 \pm \sqrt{3}}{6}$$

on which both solutions are valid. Hence, we have 2 fully mixed NE. In summary, this game has 4 pure NE, 3 partially mixed NE and 2 fully mixed NE.

3.4 Evolutionarily Stable Strategies

Evolutionarily Stable Strategies (ESS) refers to particular Nash equilibria that satisfies that the usage of a pure strategy does not change the current Nash equilibria in the long-run.

Consider a symmetric, two-player game with n pure strategies for each player. Consider that a mixed strategy represents the proportions of each type within the population. Let p be the symmetric mixed Nash equilibrium in this population and now consider that invaders participate with strategy q. For p to be an ESS, we require that the payoff of playing p does better in expectation than playing q in this society. If we say that these invaders are proportion ε of the new population, then the new probability distribution (or population proportion) is given by:

$$r = (1 - \varepsilon)p + \varepsilon q$$

Now, note that the expected payoff for a player that plays p is:

$$\mathbb{E}[\text{payoff of playing } p] = pAr = (1 - \varepsilon)pAp + \varepsilon pAq$$

while the expected of a invader that plays q is:

$$\mathbb{E}[\text{payoff of playing } q] = qAr = (1 - \varepsilon)qAp + \varepsilon qAq$$

Then, if for small enough ε, we have that:

$$\mathbb{E}[\text{payoff of playing } p] > \mathbb{E}[\text{payoff of playing } q]$$

then, we say that p is an ESS, since it will implies that in the long run, the invaders will start playing p, since it achieves a better payoff than playing q. Formally, this can be written as:

Definition 3.1. A mixed strategy $p \in \Delta_n$ is an evolutionarily stable strategy (ESS) if por any pure 'mutant' strategy s it holds that:

i. $s^\top Ap \leq p^\top Ap$.

ii. If $s^\top Ap = p^\top Ap$, then $s^\top As < p^\top As$.

That is, criterion (i.) is always satisfied if p is a mixed Nash equilibrium. Moreover, if the pure strategy s has a probability different than zero in p, then (i.) holds with equality. Thus, for p to be an ESS, it has at least to be a mixed Nash equilibrium. Now, criterion (ii.) says that, if s performs as well than playing p in the original population, it must perform worse when playing against the invaders that play s. This condition ensures that the invaders that are playing s, prefers to simply plays p since it is the same against p, and better against s. This will imply that in the long-run all invaders will start playing p, making p evolutionarily stable.

Example 3.7 (Example 7.1.3 in [1])**.** Consider the rock paper scissors game, a zero-sum game with the following payoff matrix:

	PII: Rock	PII: Paper	PII: Scissors
PI: Rock	0	-1	1
PI: Paper	1	0	-1
PI: Scissors	-1	1	0

Table 3.8: Rock-Paper-Scissors payoff matrix.

Using Theorem 3.2 is clear that $p = (1/3, 1/3, 1/3)$ is a mixed symmetric Nash equilibrium. Now, observe the following, the expected payoff of playing p against p has a value of zero, since $p^\top Ap = 0$. Since all strategies are active, it holds that $s^\top Ap = p^\top As = 0$. Now observe that, $s^\top As = 0$, since using the same strategy against itself also yields zero payoff. This implies that $s^\top As \not< p^\top As$ (actually is equal), and hence p is not an ESS. What we expect to happen is a cyclic behavior, on which a population with many rocks will be taken over by paper, which in turn will be invaded by scissors, and so forth.

Example 3.8 (Example 7.1.4 in [1])**.** Consider the following symmetric payoff bi-matrix with strategies a and b:

$$\begin{bmatrix} (10,10) & (0,0) \\ (0,0) & (5,5) \end{bmatrix}$$

In this case both pure strategies (a, a) and (b, b) are ESS, while $p = (1/3, 2/3)$ is not. However, even though (b, b) is an ESS, if a sufficiently large of a's invade, then 'stable' population will shift to being entirely composed of a's. This is done by, computing the fraction of population ε that will play a. Consider the case that the population is playing (b, b), and assume that a proportion of ε will play a. That is the population is now given by $r = (1 - \varepsilon)b + \varepsilon a$. Then, the expected payoffs is given by:

$$\mathbb{E}[\text{payoff of playing } b] = bAr = (1 - \varepsilon)bAb + \varepsilon bAa = 5(1 - \varepsilon)$$

while

$$\mathbb{E}[\text{payoff of playing } a] = aAr = (1 - \varepsilon)aAb + \varepsilon aAa = 10\varepsilon$$

Now, if the payoff of playing a is better than playing b, then a will dominate, that occurs when:

$$10\varepsilon > 5(1 - \varepsilon) \to 15\varepsilon > 5 \to \varepsilon > \frac{1}{3}$$

3.5 Correlated Equilibria

Consider the game of chicken (or the cuban-missile crisis). The following payoff matrix define the game:

	PII: Cross	PII: Yield
PI: Cross	$(-10, -10)$	$(5, 0)$
PI: Yield	$(0, 5)$	$(-1, -1)$

Table 3.9: Game of chicken payoff matrix.

Observe that $(0, 5)$ and $(5, 0)$ are pure symmetric Nash Equilibria, while the mixed one can be computed by equating payoffs over the other player strategies:

$$Aq \to -10q + 5(1 - q) = -1(1 - q) \to 6 - 6q = 10q \to q = \frac{3}{8}$$

Since $A = B^\top$, then this is a symmetric game, that implies $p^* = q^* = (3/8, 5/8)$. With that, the expected payoff is given by:

$$p^\top Aq = [3/8 \;\; 5/8] \begin{bmatrix} -10 & 5 \\ 0 & -1 \end{bmatrix} \begin{bmatrix} 3/8 \\ 5/8 \end{bmatrix} = -\frac{5}{8}$$

Now consider traffic lights. These can be thought as a randomization device or third party that tells each player what to do, deciding the probability of playing pure strategies for them. Essentially, it defines a joint distribution on pure strategies, on which there may not be independence between players. For example, consider the following joint distribution:

	PII: Cross	PII: Yield
PI: Cross	0	0.55
PI: Yield	0.4	0.05

Table 3.10: Joint distribution probabilities for players.

This can be thought as traffic lights. The third party tells to the players that the probability of both crossing (i.e. both having green light) is zero. The probability of PI crossing and PII yielding is 0.55, the probability of PI yielding and PII crossing is 0.4, while finally, the probability of both yielding is 0.05 (i.e. both with red light).

The first thing to note is that the joint distribution is not the product of independent strategies for each player. Recall that two random variables X and Y are independent if they satisfy:

$$\mathbb{P}[X = x|Y = y] = \mathbb{P}[X = x], \qquad \forall x, y$$

or equivalently

$$\mathbb{P}[X = x, Y = y] = \mathbb{P}[X = x]\mathbb{P}[Y = y], \qquad \forall x, y$$

Now, to compute probabilities if we only have the joint distribution, we can simply say:

$$\mathbb{P}[X = x] = \sum_y \mathbb{P}[X = x, Y = y] = \sum_y \mathbb{P}[X = x \mid Y = y]\mathbb{P}[Y = y]$$

and

$$\mathbb{P}[Y = y] = \sum_x \mathbb{P}[X = x, Y = y] = \sum_x \mathbb{P}[Y = y \mid X = x]\mathbb{P}[X = x]$$

With this it is clear that these probabilities are not independent. Denote $\mathcal{R}$ as PI (row player) and $\mathcal{C}$ as PII (column player), with c is cross and y is yield:

$$\mathbb{P}[\mathcal{R} = c] = 0 + 0.55 = 0.55$$

and

$$\mathbb{P}[\mathcal{C} = c] = 0 + 0.4 = 0.4$$

that makes clear that:

$$\mathbb{P}[\mathcal{R} = c, \mathcal{C} = c] = 0 \neq \mathbb{P}[\mathcal{R} = c]\mathbb{P}[\mathcal{C} = c] = 0.55 \cdot 0.4$$

And for example:

$$\mathbb{P}[\text{PI told to cross (green)} \mid \text{PII told to stop (red)}] = \frac{\mathbb{P}[\mathcal{C} = c, \mathcal{R} = y]}{\mathbb{P}[\mathcal{C} = y]} = \frac{0.55}{0.55 + 0.05} = \frac{11}{12}$$

Definition 3.2. Suppose PI has m available actions and PII has n available actions. A mixed strategy pair (p, q) that we can write as a pair of random variables $(\mathcal{R}, \mathcal{C})$, on which $\mathcal{R} \sim p \in \Delta_m$ and $\mathcal{C} \sim q \in \Delta_n$, satisfies the following:

$$\mathbb{P}[\mathcal{R} = i, \mathcal{C} = j] = p_i q_j$$

i.e. a mixed strategy are independent joint probability distribution. Then, for a bi-matrix game $M = (A, B)$, (p, q) will be a Nash equilibrium if

- If $\mathbb{P}(R = i) > 0$ (i.e. strategy i is active) then:

  $$\mathbb{E}[A_{i,\mathcal{C}}] \geq \mathbb{E}[A_{\ell,\mathcal{C}}]$$

  for all pairs (i, ℓ), $1 \leq i, \ell \leq m$. That is, the expected payoff (on the probability distribution $\mathcal{C} \sim q \in \Delta_n$ of PII) for PI for active strategies must be as good for any other strategy. This is equivalent of the equilibrium theorem for zero-sum games. Essentially says that for PII, their probability distribution q is a NE if it satisfies that for active strategies of PI that are active has to be at least as good as any other pure strategy of PI.

- Similarly, if $\mathbb{P}[\mathcal{C} = j] > 0$ (i.e. strategy j is active for PII) then:

  $$\mathbb{E}[B_{\mathcal{R},j}] \geq \mathbb{E}[B_{\mathcal{R},k}]$$

  for all pairs (j, k) for $1 \leq j, k \leq n$.

Definition 3.3. A correlated strategy is a joint distribution on the strategy pairs:

$$p = [p_{ij}] \in \mathbb{R}^{m \times n}$$

with $0 \leq p_{i,j} \leq 1$ and $\sum_{i,j} p_{i,j} = 1$, where:

$$p_{ij} = \mathbb{P}[\mathcal{R} = i, \mathcal{C} = j]$$

A correlated equilibrium is a joint distribution $p_{\mathcal{R},C}(i, j)$ where neither player has incentive to deviate given their signal and the action induced on the other player. That is, when PI (PII) is told to play row i (column j) the cannot gain from disobeying **given** the conditional probability probability that signal induces on other players actions. This can be written as follows (Definition 7.2.3 in [1]):

- If $\mathbb{P}[\mathcal{R} = i] > 0$, then:

  $$\mathbb{E}[A_{i,\mathcal{C}} \mid \mathcal{R} = i] \geq \mathbb{E}[A_{\ell,\mathcal{C}} \mid \mathcal{R} = i]$$

  for all pairs $1 \leq i, \ell \leq m$. That is, the expected payoff for PI on the distribution of PII, given that PI will play i is better than any other pure strategy, **given** the conditional probability that signal induces on PI.

- Similarly, if $\mathbb{P}[\mathcal{C} = j] > 0$ then:

  $$\mathbb{E}[B_{\mathcal{R},j} \mid \mathcal{C} = j] \geq \mathbb{E}[B_{\mathcal{R},k} \mid \mathcal{C} = j]$$

  for all pairs $1 \leq j, k \neq n$.

Now, the expected payoff can be computed as follows:

$$\begin{aligned}\mathbb{E}[A_{i,\mathcal{C}} \mid \mathcal{R} = i] &= \sum_{j=1}^{n} a_{i,j}\mathbb{P}[\mathcal{C} = j \mid \mathcal{R} = i] \\ &= \sum_{j=1}^{n} a_{i,j}\frac{\mathbb{P}[\mathcal{C} = j, \mathcal{R} = i]}{\mathbb{P}[\mathcal{R} = i]} \\ &= \sum_{j=1}^{n} a_{i,j}\left(\frac{p_{i,j}}{\sum_{k=1}^{n} p_{i,k}}\right)\end{aligned}$$

Then, the condition $\mathbb{E}[A_{i,\mathcal{C}} \mid \mathcal{R} = i] \geq \mathbb{E}[A_{i,\mathcal{C}} \mid \mathcal{R} = i]$ can be written as:

$$\sum_{j=1}^{n} a_{i,j}\left(\frac{p_{i,j}}{\sum_{k=1}^{n} p_{i,k}}\right) \geq \sum_{j=1}^{n} a_{\ell,j}\left(\frac{p_{i,j}}{\sum_{k=1}^{n} p_{i,k}}\right)$$

for all pairs (i, ℓ). Observe that the probability distribution does not change, since it is always conditional that PI will play i. In addition, note that the denominator $\sum_k p_{i,k}$ does not depend on the sum and can be cancelled:

$$\sum_{j=1}^{n} a_{ij}p_{ij} \geq \sum_{j=1}^{n} a_{\ell j}p_{ij} \;\rightarrow\; \sum_{j=1}^{n}(a_{ij} - a_{\ell j})p_{ij} \geq 0$$

for all (i, ℓ) pairs for $1 \leq i, \ell \leq m$. Since we do not count a pair with itself strategy, all pairs give a total of $m(m-1)$ constraints. This can be written compactly in the following form. Let a_i be the i-th row of the A matrix and p_i be the i-th row of the $[p_{ij}]$ joint distribution matrix. Then, the condition is equivalently to:

$$(a_i - a_\ell)p_i^\top \geq 0$$

Similarly for PII:

$$\sum_{i=1}^{m}(b_{ij} - b_{ik})p_{ij} \geq 0$$

for all (j, k) pairs for $1 \leq j, k \leq n$, that give a total of $n(n-1)$ constraints. Let b_j be the j-th column of the B matrix and q_j the j-th column of the $[p_{ij}]$ joint distribution matrix. Then, the condition can be written compactly as:

$$(b_j - b_k)^\top q_j \geq 0$$

Example 3.9 (Game of Chicken)**.** Consider the game with payoff bi-matrix (already defined in Table 3.9):

$$A = \begin{bmatrix} -10 & 5 \\ 0 & -1 \end{bmatrix} \qquad B = \begin{bmatrix} -10 & 0 \\ 5 & -1 \end{bmatrix}$$

To find all correlated equilibria we define:

$$[p_{ij}] = \begin{bmatrix} p_{11} & p_{12} \\ p_{21} & p_{22} \end{bmatrix}$$

and then we wrote the conditions:

- For PI: $i = 1, \ell = 2$. We have:

$$\begin{aligned}(a_1 - a_2)p_1^\top &= ([-10 \quad 5] - [0 \quad -1]) \begin{bmatrix} p_{11} \\ p_{12} \end{bmatrix} \\ &= [-10 - 0 \quad 5 - (-1)] \begin{bmatrix} p_{11} \\ p_{12} \end{bmatrix} \\ &= (-10 - 0)p_{11} + (5 - (-1))p_{12} \\ &= -10p_{11} + 6p_{12} \geq 0\end{aligned}$$

- For PI: $i = 2, \ell = 1$. We have:

$$(a_2 - a_1)p_2^\top = (0 - (-10))p_{21} + (-1 - 5)p_{22} = 10p_{21} - 6p_{22} \geq 0$$

- For PII: $j = 1, k = 2$. We have:

$$\begin{aligned}(b_1 - b_2)^\top q_1 &= \left(\begin{bmatrix} -10 \\ 5 \end{bmatrix} - \begin{bmatrix} 0 \\ -1 \end{bmatrix} \right)^\top \begin{bmatrix} p_{11} \\ p_{21} \end{bmatrix} \\ &= [-10 \quad 6] \begin{bmatrix} p_{11} \\ p_{21} \end{bmatrix} \\ &= -10p_{11} + 6p_{21} \geq 0\end{aligned}$$

- For PII: $j = 2, k = 1$. We have:

$$(b_2 - b_1)^\top q_2 = 10p_{12} - 6p_{22} \geq 0$$

and the conditions that $p_{ij} \geq 0$ and $\sum_{ij} p_{ij} = 1$. From $i = 1, \ell = 2$ and $j = 1, k = 2$ we can write:

$$p_{11} \leq \frac{3}{5}p_{12} \qquad \wedge \qquad p_{11} \leq \frac{3}{5}p_{21}$$

and from $i = 2, \ell = 1$ and $j = 2, k = 1$ we have:

$$p_{22} \leq \frac{5}{3}p_{21} \qquad \wedge \qquad p_{22} \leq \frac{5}{3}p_{12}$$

All p_{ij} that satisfies those condition are correlated equilibria. For example, by setting $p_{11} = p_{22} = 0$ and $p_{21} = p_{12} = 0.5$ satisfies it and is a valid correlated equilibria. Finally, we will check for the following joint distribution:

$$P = \begin{bmatrix} 0 & 1/3 \\ 1/3 & 1/3 \end{bmatrix}$$

on which the first row and column is cross (or go), while second row and column is yield (or stop) satisfies directly the conditions. This is obviously satisfied since $p_{22} \leq (5/3) \min\{p_{12}, p_{21}\}$. Now, we want to compute the expected payoffs of obeying the signals. For that, we need to compute directly the probabilities:

$$\begin{aligned}\mathbb{P}(\mathcal{C} = \text{yield} \mid \mathcal{R} = \text{yield}) &= \frac{\mathbb{P}(\mathcal{C}, \mathcal{R} = \text{yield})}{\mathbb{P}(\mathcal{R} = \text{yield})} = \frac{1/3}{1/3 + 1/3} = \frac{1}{2} \\ \mathbb{P}(\mathcal{C} = \text{cross} \mid \mathcal{R} = \text{yield}) &= \frac{\mathbb{P}(\mathcal{C} = \text{cross}, \mathcal{R} = \text{yield})}{\mathbb{P}(\mathcal{R} = \text{yield})} = \frac{1/3}{1/3 + 1/3} = \frac{1}{2}\end{aligned}$$

Now, what is the expected payoff for PI, given the marginal probabilities for PII, conditioned that PI is seeing a stop signal. On this case, the probabilities for PII are given by $\tilde{q} = (1/2, 1/2)^\top$. Now, we compute the expected payoffs for PI as follow. The payoff of obeying the signal (i.e. to yield) is computed as:

$$\mathbb{E}(A_{\text{yield},\tilde{q}} \mid \mathcal{R} = \text{yield}) = p_{\text{yield}}^\top A\tilde{q} = [0 \quad 1] \begin{bmatrix} -10 & 5 \\ 0 & -1 \end{bmatrix} \begin{bmatrix} 1/2 \\ 1/2 \end{bmatrix} = [0 \quad -1] \begin{bmatrix} 1/2 \\ 1/2 \end{bmatrix} = -\frac{1}{2}$$

and the expected payoff of neglecting the signal:

$$\mathbb{E}(A_{\text{cross},\tilde{q}} \mid \mathcal{R} = \text{yield}) = p_{\text{cross}}^\top A\tilde{q} = [1 \;\; 0] \begin{bmatrix} -10 & 5 \\ 0 & -1 \end{bmatrix} \begin{bmatrix} 1/2 \\ 1/2 \end{bmatrix} = [-10 \;\; 5] \begin{bmatrix} 1/2 \\ 1/2 \end{bmatrix} = -\frac{5}{2}$$

We observe that $-5/2 < -1/2$, and hence it is more convenient for PI to obey the signal of yield. Similarly, we can compute the payoff of crossing when the signal for PI is to cross. Again, is better for PI to obey the signal:

$$\mathbb{P}(\mathcal{C} = \text{yield} \mid \mathcal{R} = \text{cross}) = \frac{\mathbb{P}(\mathcal{C} = \text{yield}, \mathcal{R} = \text{cross})}{\mathbb{P}(\mathcal{R} = \text{cross})} = \frac{1/3}{1/3 + 0} = 1$$
$$\mathbb{P}(\mathcal{C} = \text{cross} \mid \mathcal{R} = \text{cross}) = 1 - \mathbb{P}(\mathcal{C} = \text{yield} \mid \mathcal{R} = 1 - 1 = 0$$

and hence $\hat{q} = (0, 1)^\top$:

$$\mathbb{E}(A_{\text{yield},\hat{q}} \mid \mathcal{R} = \text{cross}) = 0 + (-1) \cdot 1 = -1$$
$$\mathbb{E}(A_{\text{cross},\hat{q}} \mid \mathcal{R} = \text{cross}) = 0 + (5) \cdot 1 = 5$$

and hence since $5 > -1$, it is convenient for PI to cross when the signal is cross, showcasing that is preferable to obey the signal. Similarly, for PII we can show the same result, that is more convenient to obey the signal.

4 Social Choice and Mechanism Design

The key question to answer is how do we aggregate the preferences of individuals in a society? In practice, when there are only two options to choose from, **majority rule** can be used to select a choice that is preferred for more than half of the population. However, when more than two options are available, inconsistent results can occur, assuming that voters are *rational* (i.e. transitivity of votes).

Example 4.1. Consider a society with 100 voters and three candidates. Assume that their ranking preferences are as follow:

$$40: \begin{bmatrix} A \\ C \\ B \end{bmatrix} \qquad 35: \begin{bmatrix} B \\ A \\ C \end{bmatrix} \qquad 25: \begin{bmatrix} C \\ B \\ A \end{bmatrix}$$

And considering that votes are transitive, and hence consider the following pairwise cases:

- A vs B: B wins with 60 votes against 40 votes.
- A vs C: A wins with 75 votes against 25 votes.
- B vs C: C wins with 65 votes against 35 votes.

and hence, the following paradox occurs:

$$B \succ A \succ C \succ B$$

To start we define the following terms. In this case there is a finite set of candidates $\mathcal{A} = \{1, 2, \ldots, m\}$, such that $m > 2$ (the case of $m = 2$ is a boring case) and there are n voters.

Definition 4.1 (Preference relation)**.** Each voter is represented by a **preference relation** $\succ$ on the set of candidates $\mathcal{A}$. This preference relation $\succ$ is complete, that is $\forall a, b \in \mathcal{A}$, either $a \succ b$ or $b \succ a$, and is also **transitive**, that is if $a \succ b$ and $b \succ c$, then this implies that $a \succ c$. We denote $\succ_i$ to denote voter's i's preference relation: $a \succ_i b$ if voter i prefers candidate a over candidate b

Definition 4.2 (Preference profile)**.** We define the **preference profile** as the n-tuple $(\succ_1, \succ_2, \ldots \succ_n)$ that describes the preference relation of all n voters.

Definition 4.3 (Voting rule)**.** A **voting rule** f maps each **preference profile** $\pi = (\succ_1, \ldots, \succ_n)$ to an element of $\mathcal{A}$, the winner of the elections.

Definition 4.4 (Ranking rule)**.** A **ranking rule** R maps each **preference profile** $\pi = (\succ_1, \ldots, \succ_n)$ to a **social ranking** $\triangleright$. This is a complete and transitive preference relation $\triangleright = R(\pi)$. In this case $a \triangleright b$ means that a is preferred over b in the social ranking.

Note that a voting rule is straightforward obtained from a ranking rule by simply picking the top ranked candidate. On the other hand, a ranking rule can be obtained from a voting rule in the following way. Apply the vote rule to obtain the top candidate. Remove that candidate, and apply the voting rule to the remaining candidates, until no candidates (or only 1 candidate) remains.

4.1 Desired properties of voting and ranking rules

In a nutshell, we are interested in which properties should a "fair" voting f or ranking rule R have. Some of these properties are listed next, but we will see that most ranking rules will violate some of the following properties:

1. **Unanimity:** If every voter i prefers candidate a over candidate b, then R should rank a higher than b ($a \triangleright b$).

2. **Arrow's Independence of Irrelevant Alternatives (IIA):** For any two candidates a and b, the preference between a and b in the social ranking depends only on the voters' preferences between a and b. That is, if $\pi = \{\succ_i\}$ and $\pi' = \{\succ'_i\}$ are two profiles for which each voter has the same preference between and a and b, then $a \triangleright b$ implies $a \triangleright' b$.

3. **Anonymity (or symmetry)**: The identities of the voters should not affect the results, or equivalently, that if the preference ordering of voters are permuted, the social ranking should not change. This is true in most voting systems. However, for electoral colleges in the US, this is not satisfied, on which permuting voters between states could easily change the result of a president election.

4. **Monotonicity:** If a voter moves only up a candidate, then it should not decrease in the social ranking.

5. **Condorcet Criterion:** If a candidate beats all other candidates in one-to-one contests, then he should win the election. Similarly, the **Condorcet Loser Criterion** establish that a candidate that loses to all other candidates in one-to-one contests, should not win the election.

We say that a ranking rule that satisfies a "fairness criteria" is a ranking rule that satisfies both *unanimity* and *IIA*.

4.2 Voting and ranking mechanisms

4.2.1 Plurality voting

In *plurality voting* each voter chooses his/her favorite candidate, and the candidate with most votes wins the election. If each voter submits a ranking-ordering of the candidates, then the candidate with most first-place votes will win the election. In the case of runoff, it is plurality with elimination, on which a second election with only two candidates must be conducted to find the winner (that is expensive to do).

For the induced ranking of plurality voting, you remove the first winner, and assign the voters of that winner to their second choice and again compute the winner without the removed candidate. You continue the process until all candidates are ranked.

4.2.2 Plurality ranking

In this case, you rank all the candidates based on the order of top votes. That is, you only count the top votes and sort them to decide the social ranking.

4.2.3 Borda count (ranking rule)

Each voter must rank its m candidates. Then candidate at the top receive m points, the second one receives $m-1$ votes, and the bottom one receives one point. Adding up all the voters, the candidate with highest total points wins.

4.2.4 Condorcet rule (voting rule)

Each voter must rank its m candidates. The winner of pairwise (one-to-one) elections should win (by moving votes given the voters' ranking). In practice, that is whichever candidate that would win in (most) head-to-head contest, with greatest difference.

4.2.5 Instant runoff voting rule (IRV)

Each voter must rank its m candidates. Discard the loser with plurality (i.e. the one with lowest top preferences) and give votes to second preferences. Keep going until only one candidate remains.

Example 4.2. Consider the following preference relation of 100 voters:

$$40 : \begin{bmatrix} a \\ b \\ c \end{bmatrix} \qquad 25 : \begin{bmatrix} b \\ c \\ a \end{bmatrix} \qquad 35 : \begin{bmatrix} c \\ b \\ a \end{bmatrix}$$

Observe the following:

- By plurality, a wins with 40 votes. We have a social ranking of $a \triangleright c \triangleright b$.
- By Borda count, the total votes are:

$$\begin{aligned} a &: \ 40 \cdot 3 + 25 \cdot 1 + 35 \cdot 1 = 180 \\ b &: \ 40 \cdot 2 + 25 \cdot 3 + 35 \cdot 2 = 225 \\ c &: \ 40 \cdot 1 + 25 \cdot 2 + 35 \cdot 3 = 195 \end{aligned}$$

  and hence b wins with 225 votes. The social ranking is $b \triangleright c \triangleright a$.
- With instant runoff voting we discard b first, that has only 25 top votes. That gives 25 votes to c, that discard a. With that, the social ranking is $c \triangleright a \triangleright b$.
- With Condorcet rule we have that b beats a (60-40), that c beats a (60-40) and that b beats c (65-35). This makes b the Condorcet winner, and a social ranking of $b \triangleright c \triangleright a$.

Example 4.3. Consider the following preference relation of 100 voters:

$$45 : \begin{bmatrix} a \\ b \\ c \end{bmatrix} \qquad 30 : \begin{bmatrix} b \\ a \\ c \end{bmatrix} \qquad 25 : \begin{bmatrix} c \\ b \\ a \end{bmatrix}$$

Observe that by plurality we have $a \triangleright b \triangleright c$. Now, using instant runoff voting (IRV), we remove c, and we have:

$$45 : \begin{bmatrix} a \\ b \end{bmatrix} \qquad 55 : \begin{bmatrix} b \\ a \end{bmatrix}$$

that showcase that $b \rhd a \rhd c$ with IRV.

Now, consider that 10 people that put a on top, move c to the top without changing the relative ranking of a and b:

$$35 : \begin{bmatrix} a \\ b \\ c \end{bmatrix} \qquad 10 : \begin{bmatrix} c \\ a \\ b \end{bmatrix} \qquad 30 : \begin{bmatrix} b \\ a \\ c \end{bmatrix} \qquad 25 : \begin{bmatrix} c \\ b \\ a \end{bmatrix}$$

Now, observe that IRV remove b first, and then a beats c with 65 votes against 35. That is, with only moving the irrelevant alternative c, we have changed the social preference between a and b using IRV, that showcase that IRV does not satisfy the independence of irrelevant alternatives (IIA).

Definition 4.5 (Dictatorship)**.** We say that a ranking rule is a **dictatorship** if there is a voter v whose preferences are reproduced in the social outcome. That is, for every pair (a, b) of candidates we have:

$$a \rhd b \;\Leftrightarrow\; a \succ_v b$$

4.3 Arrow's Impossibility Theorem

Theorem 4.1 (Arrow's Impossibility Theorem)**.** The theorem is stated as follows (for a set with more than 2 candidates): "Any ranking rule that satisfies a fairness criteria (i.e. both unanimity and independence of irrelevant alternatives), then that ranking rule must be a dictatorship".

The proof relies on several definitions on lemmas that we will state.

Let i be a voter and $\beta \in \mathcal{A}$ a candidate. Let π be a preference profile. We define the following two preference profiles:

- $\pi^+(i, \beta)$: a preference profile that is identical to π except that voter i puts β on top, i.e. $\beta \succ_i a, \ \forall a \in \mathcal{A}$.
- $\pi^-(i, \beta)$: a preference profile that is identical to π except that voter i puts β on the bottom, i.e. $a \succ_i \beta, \ \forall a \in \mathcal{A}$.

β is called a "extremal" for voter i in the profiles $\pi^+(i, \beta)$ and $\pi^-(i, \beta)$ (since it is on the extremes of the preference relation).

Definition 4.6 (β-pivotal)**.** A voter i is called β-pivotal (for a ranking rule R) if:

1. $\exists \pi$ such that β is extremal for all voters in π.
2. $\pi^+(i, \beta)$ leads to $\beta \rhd a, \ \forall a \in \mathcal{A} - \{\beta\}$ (i.e. β is on top of the social ranking).
3. $\pi^-(i, \beta)$ leads to $a \rhd \beta, \ \forall a \in \mathcal{A} - \{\beta\}$ (i.e. β is on the bottom of the social ranking).

Lemma 4.1 (Extremal lemma)**.** If every voter has a candidate β either at the top of bottom of their ranking (that is β is extremal for all voters) then β is extremal in $R(\pi)$, where R satisfies the fairness criteria (unanimity and IIA).

Proof: Suppose not, so β is extremal for all voters but is not extremal on the social ranking $R(\pi)$. This implies that $\exists a, c \in \mathcal{A}$ with $a \neq \beta$ and $c \neq \beta$, such that $a \rhd \beta$ and $\beta \rhd c$, since β is not extremal in $R(\pi)$. We use the following notation to denote voters: $\{_\beta$ represent a preference relation of a voter that has β on the bottom, while $\{^\beta$ a preference relation that has β on top. Then π is of the form:

$$\pi: \quad \underset{1}{\{^\beta}, \ \underset{2}{\{^\beta}, \ \underset{3}{\{_\beta}, \ \dots, \underset{n}{\{_\beta}$$

that is, β is extremal for every voter. Then, create a new preference profile π' from π in which every voter moves c above a, but not above β, so β is still extremal for every voter in π' and does not change individual preferences of the pair (c, β) and (a, β). By IIA, since we do not change the relative rankings of (a, β) and (c, β) we must have that $R(\pi)$ and $R(\pi')$ have the same social preferences between these pairs, that is $a \rhd' \beta$ and $\beta \rhd c$. By transitivity, we then have $a \rhd' c$. However, by unanimity, since c is above a for every voter we must have that $c \rhd' a \ \rightarrow\leftarrow$. This is a contradiction, and therefore β is extremal on $R(\pi)$. $\square$

Lemma 4.2 (Existence of a β-pivotal voter)**.** For every candidate β, there exists a β-pivotal voter v_β. That is, there exists a preference profile π such that β is extremal for all voters. To find such preference profile, start with a preference profile π_0 that has β on the bottom for all voters:

$$\pi_0: \quad \underset{1}{\{_\beta}, \ \underset{2}{\{_\beta}, \ \underset{3}{\{_\beta}, \ \dots, \underset{n}{\{_\beta}$$

One by one, let each voter flip β from the bottom to the top, i.e.:

$$\pi_0 \to \pi_1: \quad \underset{1}{\{^\beta}, \ \underset{2}{\{_\beta}, \ \underset{3}{\{_\beta}, \ \dots, \underset{n}{\{_\beta} \qquad\qquad \pi_1 \to \pi_2: \quad \underset{1}{\{^\beta}, \ \underset{2}{\{^\beta}, \ \underset{3}{\{_\beta}, \ \dots, \underset{n}{\{_\beta}$$

until all voters have flipped their preference on candidate β:

$$\pi_n: \quad \underset{1}{\{^\beta}, \ \underset{2}{\{^\beta}, \ \underset{3}{\{^\beta}, \ \dots, \underset{n}{\{^\beta}$$

By the Extremal lemma (Lemma 4.1), $\forall \pi_k, \ 0 \le k \le n$, $R(\pi_k)$ will have β on the top or bottom. Let v_β be the first voter such that when v_β put β on top of his own preference relation, it flips β from bottom to top in the society, then v_β is the β-pivotal voter. $\square$

Proof of Arrow's Impossibility Theorem: We will show that v_β, the β-pivotal voter, is a dictator on all a, c candidate pairs, for all $a, c \in \mathcal{A} - \{\beta\}$. Since v_β is β-pivotal, then $\exists \pi$ a preference profile such that β is extremal (for all voters) and:

- $\pi^+(i, \beta)$ implies that $R(\pi^+)$ has β on top of the social ranking, i.e. $\beta \rhd^+ a, \ \forall a \in \mathcal{A} - \{\beta\}$.
- $\pi^-(i, \beta)$ implies that $R(\pi^-)$ has β on the bottom of the social ranking, i.e. $a \rhd^- \beta, \ \forall a \in \mathcal{A} - \{\beta\}$.

For the proof, construct π' from π^+ by moving a above β for v_β, keeping everything else same. That is in $R(\pi^+)$, β is on top of the social ranking, while since in $R(\pi^-)$ we have moved β from the top on the preference relation of v_β, to the previous position of a.

Thus, in π', v_β has a preference relation of the form $a \succ_{v_\beta} \beta \succ_{v_\beta} c$. Now, let everyone else arbitrarily rearrange a and c relative ranking leaving β alone (that is extremal on those voters). By IIA, this should not affect a and β's relative social ranking. To recap, we have the following:

$$\pi^+: \quad \underset{1}{\{_\beta}, \ \underset{2}{\{_\beta}, \ \dots, \ \underset{v_\beta}{\{^\beta}, \ \dots, \underset{n}{\{^\beta} \ \Rightarrow \ R(\pi^+): \quad \underset{\rhd}{\{^\beta}$$

$$\pi^-: \quad \underset{1}{\{_\beta}, \ \underset{2}{\{_\beta}, \ \dots, \ \underset{v_\beta}{\{_\beta}, \ \dots, \underset{n}{\{^\beta} \ \Rightarrow \ R(\pi^-): \quad \underset{\rhd}{\{_\beta}$$

Note that in π', v_β put a on top of β. This means that the preference between a and β is like π^-. This implies that β cannot be on top of the social ranking on π', since $a \rhd' \beta$. However, all (β, c) pairwise relative rankings has the same relation as $R(\pi^+)$. Then:

- In $R(\pi')$ we have $a \rhd' \beta \rhd' c$ by IIA, and hence $a \rhd' c$.
- We had v_β move a so $a \succ_{v_\beta} c$. In addition, we have also has everyone else arbitrarily rearranged a and c, but for IIA and v_β being β-pivotal, still resulted in $a \rhd c$.

The two previous conditions showcase that v_β **dictates** social relative ranking for $a, c \in \mathcal{A} - \{\beta\}$. That is, v_β is a dictator of the ranking rule R. $\square$

To finalize, we are interested in finding if there is another dictator. Suppose that $x \neq \beta$ is some other candidate and let v_x be the x-pivotal voter of that candidate. Let $a \neq x, \beta$ another candidate such that:

$$\pi^+(v_\beta, \beta) \to \beta \rhd a \qquad\qquad \pi^-(v_\beta, \beta) \to a \rhd \beta$$

However, we know that v_x dictates the social preferences for pair (a, β), and hence this is not possible. Thus, $v_x = v_\beta$ and there is only one dictator. $\square$

4.4 Gibbard-Satterthwaite Theorem

We start with the following definition:

Definition 4.7 (Strategy-proof)**.** We say that a voting rule f is "strategy-proof" (or non manipulable) if for all π preference profiles, for all candidates $A, B \in \mathcal{A}$, and for all voters i, the following holds: if $A \succ_i B$ and $f(\pi) = B$, then if π' only differs from π in voter i's ranking, then $f(\pi') \neq A$. That is, if voter i prefers A over B, but B still wins, there is nothing voter i can do by modifying his vote to make A wins.

Theorem 4.2 (Gibbard-Satterthwaite Theorem)**.** Let f be a "strategy-proof" (or non-manipulable) voting rule, then f is a dictatorship. That is, there is a voter i such that for every profile π, voter i's highest ranked candidate is equal to $f(\pi)$.

5 Stable Matching

This problem was introduced by David Gale and Lloyd Shapley in 1962 [5]. Consider two colleges A and B, and two students α and β that are applying to the colleges. We say that student α prefers A over B, while student β prefers B over A. However, A prefers β, while B prefers α. It is clear that no pair-up could satisfy all preferences, but in this problem we are interested in what is called an **unstable matching**.

Suppose that α is admitted in A and β is admitted in B, but now β prefers A and still A prefers β over α. In such case, both student β and college A can discuss and finally admit β in A, letting α go, that is a better result for **both** the student and college. In such case, we say that the matching $\alpha A, \beta B$ is an unstable matching, since both A and β prefer each other.

The problem easily extend to the problem as matching men and women in a marriage. The problem is known as the stable marriage problem. Typically we denote men using the greek alphabet $\alpha, \beta, \gamma, \ldots$ and women using capital letters $A, B, C, \ldots$

Example 5.1. Consider the example provided in [5], on which we have three men α, β, γ and three women A, B, C with the following ranking:

α	β	γ	A	B	C
A	B	C	β	γ	α
B	C	A	γ	α	β
C	A	B	α	β	γ

Table 5.1: Ranking preference of both men and women.

We can write this as a pairwise matrix, with the first entry to be the row preference (men) and the second entry to be the column preference (women).

	A	B	C
α	$(1,3)$	$(2,2)$	$(3,1)$
β	$(3,1)$	$(1,3)$	$(2,2)$
γ	$(2,2)$	$(3,1)$	$(1,3)$

Table 5.2: Preference of both men and women in a pairwise matrix.

For a stable marriage problem with n men and n women, there are $n!$ possible matchings. In this case $3! = 6$, and we will analyze the stability of each possible matching.

- Consider the matching given by: $\alpha A, \beta B, \gamma C$. Observe that all men (α, β, γ) have their top preferences, so they will never would want to change, independently of what other women wants. Thus, this is a stable matching.

- Consider the matching given by $\alpha A, \beta C, \gamma B$. Observe that A is with her third preference (men α), while γ is also with his third preference (women B). Thus, both A and γ would prefer to match between them, achieving a better matching for both of them. Thus, the proposed is an unstable matching, because of $A\gamma$ will form.

– Consider the matching given by $\alpha B, \beta A, \gamma C$. Observe that β is with his third preference (women A) while C is with her third preference (men γ). Thus, both β and C will prefer to match between them. Thus, this is an unstable matching, because βC will occur.

– Consider the matching given by $\alpha B, \beta C, \gamma A$. Everyone is with his second best choice, so no change will occur. Hence, this is a stable matching.

– Consider the matching given by $\alpha C, \beta A, \gamma B$. In this case, all women (A, B, C) are with their top preferences, so they will never would want to change, independently of what other men wants. Thus, this is a stable matching.

– Finally, consider the matching given by $\alpha C, \beta B, \gamma A$. In this case, α is with his third choice, while B is with her third choice β. Thus, the matching αB will occur, showcasing that this is an unstable matching.

Example 5.2. Consider a variation of the stable marriage problem, that is the stable roommate problem. That is, there is an even number of people $2n$ that want to divide up into pairs of roommates. Similarly, we say that a set of pairing is stable if there are no two people who are not roommates and who mutually prefer each other, with respect to their assigned roommate. In [5] it is showcased that there does not always exist a stable pairing in this situation. The proposed example is as follows:

Consider 4 men α, β, γ and δ. We have that α ranks β first, while β ranks γ first and γ ranks α first. In addition, all α, β, γ rank δ at their last position. Observe that, independently of δ preferences, there cannot be an stable pairing, since whoever is paired with δ will want to move out, and there will be one person willing to accept him:

– If α is paired with δ, then γ will accept α, since he is his top preference.

– If β is paired with δ, then α will accept β, since he is his top preference.

– If γ is paired with δ, then β will accept γ, since he is his top preference.

Thus, there cannot be a stable pairing in this example.

5.1 The Gale-Shapley proposing algorithm

We are interested in answering if always will exist a stable set of marriages for every preference profile, and in particular if we are able to find it. David Gale and Lloyd Shapley propose that algorithm in [5] as follows (this is the man-proposing algorithm, the woman-proposing algorithm simply differs by making women to propose instead of man):

1. In the first round, each man proposes to his favorite woman. Then, most likely there will be women with multiple proposal and some women without proposal (it is clear if a woman has more than one proposal, that implies that some woman does not receive a proposal).

   Then, the women tentatively accept their top preference of all the proposal they receive and reject the others. At the end of the round there will be tentative marriages.

2. In the second round, all the men that were rejected propose to their second best choice, and now women that has more than one proposal (considering their previous one if they were in a tentative marriage) can reject and choose their best option.

3. Repeat (2) until there are no more rejections. There are at most $n(n-1)+1$ proposals. Hence, since it is a finite number, then the algorithm must terminate.

Theorem 5.1. The Gale-Shapley (G-S) man-proposing algorithm produces a stable matching.

Proof: We know that the algorithm terminates when there are no more rejections, and since there a finite number of proposal, the algorithm must end. Let's denote as M the terminal matching using the G-S algorithm. Suppose, that there is a male $a \in M$ and a female $A \in M$ such that a prefers A than the current partner or a is single (in the case of more men than women). Then, A must be rejected a in some previous rounds for someone she likes better. Thus, M is stable. □

Definition 5.1. A woman A is called **attainable** for a man α if exists a stable matching M in which they are paired.

The following properties hold for the G-S proposing algorithm.

1. In G-S man-proposing algorithm, every man is paired with his most preferred attainable woman.

   Similarly, in G-S woman-proposing algorithm every woman is paired with her most preferred attainable men.

2. In G-S man-proposing algorithm, every woman in paired with her least preferred attainable man.

   Similarly, in G-S woman-proposing algorithm every man is paired with his least preferred attainable woman.

3. In the case of both men-proposing algorithm and woman-proposing algorithm achieve the same matching, then both are men-optimal and men-worse at the same time, that implies that there is a unique stable matching.

Example 5.3. Consider the following characters from Emma, by Jane Austen. In the women we have Emma Woodhouse (EW), Jane Fairfax (JF), Harriet Smith (HS) and Augusta Hawkins (AH), while on men we have George Knightley (gk), Frank Churchill (fc), Philip Elton (pe) and Robert Martin (rm). Their ranking preferences are as follow:

	EW	JF	HS	AH
gk	$(1,2)$	$(2,2)$	$(3,1)$	$(4,1)$
fc	$(2,1)$	$(1,1)$	$(3,4)$	$(4,3)$
pe	$(1,3)$	$(3,3)$	$(4,2)$	$(2,2)$
rm	$(3,4)$	$(2,4)$	$(1,3)$	$(4,4)$

Table 5.3: Preference of both men and women in Emma.

We will produce a stable matching using the man-proposing algorithm:

- Stage 1: In this case we have following proposals. Boxed proposal means tentative matching:

EW	JF	HS	AH
gk	fc	rm	
pe			

Table 5.4: Proposals in stage 1.

- Stage 2: Since Philip (pe) was rejected in stage 1, he propose to his second best choice AH:

EW	JF	HS	AH
gk	fc	rm	pe

Table 5.5: Proposals in stage 2.

With this, we are done since there aren't rejections in stage 2. The matching is given by: EW|gw, JF|fc, HS|rm and AH|pe.

Now, we explore the woman-proposing algorithm:

– Stage 1:

gk	fc	pe	rm
HS	JF		
AH	EW		

Table 5.6: Proposals in stage 1.

– Stage 2: Since AH and EW were rejected in stage 1, they propose to their second best choice:

gk	fc	pe	rm
HS	JF	AH	
EW			

Table 5.7: Proposals in stage 2.

– Now, HS was rejected in stage 2 (but was not rejected in stage 1), so she proposes to her second best choice:

gk	fc	pe	rm
EW	JF	AH	
		HS	

Table 5.8: Proposals in stage 3.

and is rejected again.

– Finally, HS proposes to her third choice:

gk	fc	pe	rm
EW	JF	AH	*HS*

Table 5.9: Proposals in stage 4.

Without rejections, we stop the algorithm. We observe that we obtain the same matching as before: EW|gw, JF|fc, HS|rm and AH|pe. This implies that this is a unique stable matching.

Example 5.4. Consider the following example of a stable marriage problem. Another algorithm proposed by Donald E. Knuth in 1997, called successive divorces can be used to find a stable matching. The key idea is that you start with a guess of a possible matching and start doing successive divorces (on which both parts prefer each other). The ranking matrix is as follows:

	A	B	C	D
a	$(4,1)$	$(2,2)$	$(1,4)$	$(3,2)$
b	$(2,2)$	$(1,4)$	$(3,2)$	$(4,1)$
c	$(3,4)$	$(1,1)$	$(4,1)$	$(2,3)$
d	$(2,3)$	$(4,3)$	$(1,3)$	$(3,4)$

Table 5.10: Preference ranking.

- Now, we start with the following pair-up Aa, Bb, Cc, Dd. Observe that for a, A is his worst choice, and similarly for B, on which b is her worst choice. Thus, they match together, breaking the current couple.
- Now, the pair is Ab, Ba, Cc, Dd. Observe that for c, C is his worst choice, and for D, d is her worst choice, so they prefer to change.
- Now, the pair is Ab, Ba, Cd, Dc. Now observe that for B, a is his second choice, while for c, D is his second choice. But, both of them prefer to be match together, and hence this match occurs.
- Now, the pair Ab, Bc, Cd, Da. Note that this is a stable matching.

This algorithm can get stuck in a loop, and so it may not end all the time, depending on the initial guess.

6 Auctions

In simple terms, auctions are mechanisms to sell and buy goods when you really don't know the value to fix for the item being sold. That is, the auction will be used to figure out the value of the item, by setting the price dynamically.

In this section, we will focus on single item auctions, that is, there is only one seller and multiple buyers. In addition, we will use indistinguishable the terms players and bidders (typically to avoid phrases like *"bidder i bid his bid b_i"*).

Definition 6.1 (Walkaway price)**.** We say that the value or *walkaway price* of a buyer i is v_i, if they would not pay more than v_i for the item.

We will assume that all buyers are *rational* players. That is, they want to maximize their utility, computed by $v - p$, where p is the price paid by the buyer for item, given the auction rules and knowledge of other players' bid.

Definition 6.2 (Private value)**.** We say that a *private value* is independent of other people valuations. That is, for example, a fanatic of Harry Potter, could value way more a first edition of "Harry Potter and the Philosopher's Stone" than a person who is not interested in Harry Potter.

Definition 6.3 (Common value)**.** We say that an item has a *common value* for all the bidders, but they may have different estimates of that common value. That is, for example, an offshore oil lease, on which people may have different estimates on how much oil is there, or different estimates on what will be the future price of oil. However, if the information were perfect, then the value will be the same for all bidders (for now, we ignore that there are some elements of private value if a company have more efficient drilling methods or techniques used for extracting the oil).

6.1 Types of Auctions

6.1.1 English auction

An English auction is the traditional auction, an ascending open bid auction for an item. The highest bidder wins the item. In this type of auction, the winner pays the second highest valuation, assuming continuous increments of bids. Observe that, the winner only needs to bid ϵ above the second highest valuation to win the bid, that of course is below the winner private value (assuming that they are not exactly the same valuation).

The **optimal strategy** for this type of auction is to bid in continuous increments until your private value. If your private value is the current price, then you simply walk away, or you will win with a price below your value.

A variation of this type is the Japanese auction (or ascending clock auction). It proceeds in the following way:

- An initial price is displayed. This is usually a low price - it may be either 0 or the seller's reserve price.
- All buyers that are interested in buying the item at the displayed price enter the auction arena.

- The displayed price increases continuously, or by small discrete steps (e.g. one cent per second).
- Each buyer may exit the arena at any moment.
- No exiting buyer is allowed to re-enter the arena.
- When a single buyer remains in the arena, the auction stops. The remaining buyer wins the item and pays the displayed price.

6.1.2 Sealed bid first-price

Each bidder submits their bid in a private way (typically a closed letter), and the highest bidder wins, and pays their bid. In this type of auctions, for common value items, there is an important aspect on what is called *winner's curse*. That is, if you over-value the item and submit a really high bid, you could be end up paying way more than the other people think what is the value of that item (and most likely the real value of the item).

In this case, the **optimal strategy** on how much to bid for each player depends on what think other players will bid, that makes it complicated to solve.

6.1.3 Vickrey auction or sealed bid second-price

This is similar than the sealed bid first-price, on which each bidder submits their bid in a private way (typically a closed letter), and the highest bidder wins. However, the winner pays the second highest bid.

The **optimal strategy** for this game is to submit their private value. Thus, the Vickrey auction encourages truthful bids, since it is the dominant strategy to maximize the utility. To show this, consider that bidder i has a private value of v_i. Let us consider the following cases:

- What happens if v_i submits a lower bid $b_i = v_i - \delta$, instead of b_i? Consider that for example another bidder j submit a bid of $b_j = v_i - \delta/2$. In such case, bidder j wins the auction, and bidder i ends up with zero utility. But, if he would simply bid $b_i = v_i$, could win the auction with positive utility $u_i = v_i - (v_i - \delta/2) = \delta/2$. Thus, it is convenient to simply bid your own private value.
- Similarly it can be showed that submitting higher bids than your private value is worse, since you can end-up with negative utility.

6.1.4 Dutch auction

Is the reverse case of the Japanese auction, on which it starts with a high price that falls until first bidder raises their hand and wins. In here the strategy is to bid less than your true value.

Example 6.1 (ACME example in [6]). Suppose you want to buy a firm, that you will be able to increase its value by 50% and sell it. Sadly, you don't know the current value, but believe it is uniformly distributed between $2 to $12 million dollars. You can make a single "take-it-or-leave-it" bid. The owners will accept your bid if it is greater than the current value. How much should you bid, considering that you don't want to lose money in expectation?

The key in this example is that some information is revealed if your bid is accepted. Suppose you bid x, and your bid is accepted. Then, you know that the value of the company is at least x, if not your bid would not be accepted. That means, now the distribution of the value is $U[2, x]$. Then, the expected value of the company is $(2 + x)/2$. Thus, to find the bid we simply equate:

$$\frac{2+x}{2} \cdot 1.5 = x \ \rightarrow \ x = 6$$

Example 6.2 (k-unit Vickrey auction)**.** Consider a k-unit Vickrey auction. That is, there are n bidders and k items (with $n > k$). Each player submit a sealed bid, and the first k bids win one item, and have to pay the $k + 1$ higher bid. We will show that this auction is truthful.

Let's assume that each player bids their private value v_i. Without loss of generality, assume that $v_1 \geq v_2 \geq \ldots \geq v_k \geq v_{k+1} \geq \ldots \geq v_n$. We will show that for any player, there is no incentive to not be truthful. There are two scenarios:

1. Consider the winners of the auction, i.e. $i = 1, 2, \ldots, k$. For any winner i, if he goes higher of his private value, then no change occurs. The only way he can change its utility is by bidding below v_{k+1}, but that would imply that he is no longer a winner and will incur in less utility than before. Thus, a winner has no incentive to deviate from his private value (we say that this is a weak dominant strategy, since there are other bids that also yield the same utility).

2. Now, consider the losers of the auction, i.e. $j = k + 1, \ldots, n$. Note that reducing its bid does not change anything, since he would still lose, so he should at least bid v_j to hope to win.

   Now, in order to increase its utility he would have to bid at least above v_k. That would imply, that now he is a winner, but the price of the auction will be v_k that is higher or equal than v_j, resulting in a negative (or zero) utility for bidder j. Thus, there is also no incentive in overbidding.

6.2 Single item auctions

We say that a direct single item auction $\mathcal{A}$ with n bidders is a mapping that assigns a winner to a vector of bids, or bidding profile: $b = (b_1, b_2, \ldots, b_n)$.

Definition 6.4 (Private values)**.** We write as $V_1, V_2, \ldots, V_n$ the independent random variables (not necessarily identically) that denotes the private values of each bidder i. We write v_i as one (possible) realization of V_i.

For every bidder i, we denote as F_i, the distribution of V_i, that is common knowledge for each bidder. In addition, each bidder knows his own realization of their private value v_i (but only knows the distribution of the other players).

Definition 6.5 (Bidding strategy)**.** A bidding strategy $\beta_i : [0, \infty) \to [0, \infty)$ is a function such that for all bidders i and for all possible private values $v_i \geq 0$, β_i specifies a bid $b_i = \beta_i(v_i)$.

Definition 6.6 (Allocation rule)**.** An allocation rule of $\mathcal{A}$ is denoted by $\alpha^{\mathcal{A}}[b]$, and it is a vector of indicator functions:

$$\alpha^{\mathcal{A}}[b] = (\alpha_1[b], \alpha_2[b], \ldots, \alpha_n[b])$$

on which:

$$\alpha_i[b] = \begin{cases} 1 & \text{if } i \text{ wins with bid } b_i \text{ and other bids } b_{-i} \\ 0 & \text{otherwise} \end{cases}$$

Basically an allocation rule is a vector of zeros and ones that determines which bidder won the auction. In the case of single item auctions, only one entry can have a one, while all other entries must have zeros.

Figure 6.1 describes the process in terms for bidder i, that only knows his own valuation and other bidders distribution:

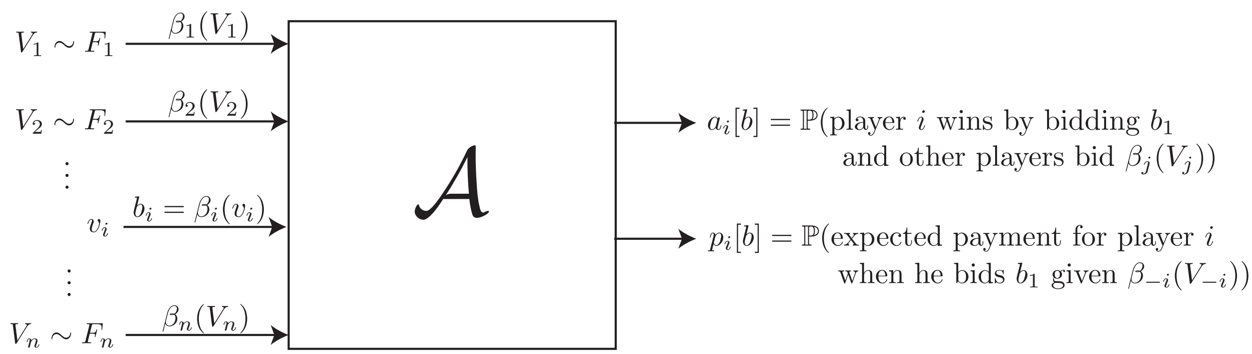

Figure 6.1: Description of auction procedure for bidder i.

In particular, the probability that bidder i wins when he bids b_i and other players bid $\beta_{-i}(V_{-i})$ can be computed using the indicator function:

$$a_i[b] = \mathbb{E}[\alpha_i(b_i, \beta_{-i}(V_{-i}))]$$

Finally, the expected utility for bidder i can be computed as:

$$u_i[b_i \mid v_i] = v_i a_i[b] - p_i[b]$$

Note that here we are abusing some notation since $u_i[b_i \mid v_i]$ will depend on other players' bids b_{-i}, but we do not denote this explicitly here. It is also possible to denote $u_i[b \mid v_i]$ to explicitly denote that all bids play a role in the utility for bidder i.

Definition 6.7 (Bayes–Nash equilibrium)**.** We say that a set of bidding strategies $(\beta_1, \dots, \beta_n)$ is a Bayes–Nash equilibrium (BNE) if $b_i \mapsto u_i[b \mid v_i]$ is maximized at $b = \beta_i(v_i)$ for all bidders i, for all bids b and private values v_i. That is:

$$u_i[\beta_i(v_i) \mid v_i] \geq u_i[b \mid v_i] \quad \forall i, b, v_i$$

Example 6.3 (2 players sealed-bid first-price auction)**.** Consider two players bidding via sealed-bid first-price auction, on which V_1 and V_2 are their private values. Each player knows that the other player's private value is drawn from a Unif$[0, 1]$ distribution. We will show that the best response from each player is to bid half their private value.

Our first approach consider that each player bids $\beta_i(V_i)$, where $\beta_i : [0, 1] \to [0, 1]$. Let's compute the expected gain of PI (value v_1 and bid b_1)) and suppose that player 2 bids $V_2/2$ (that is unknown to PI). We have:

$$\begin{aligned}
\text{expected utility} &= \mathbb{E}[(v_1 - b_1)\mathbb{1}_{\{b_1 \geq b_2\}} \mid V_2] \\
&= (v_1 - b_1)\mathbb{E}[\mathbb{1}_{\{b_1 \geq b_2\}} \mid V_2] \\
&= (v_1 - b_1)\mathbb{P}[b_1 \geq b_2 \mid V_2] \\
&= (v_1 - b_1)\mathbb{P}\left[b_1 \geq \frac{V_2}{2}\right] \\
&= (v_1 - b_1)\mathbb{P}[V_2 \leq 2b_1] \\
&= (v_1 - b_1)(2b_1) \\
&= 2b_1 v_1 - 2b_1^2
\end{aligned}$$

Now, to compute the best b_1 we simply take the derivative and set it up to zero:

$$f'(b_1) = 2v_1 - 4b_1^2 = 0 \to b_1 = \frac{v_1}{2}$$

that is indeed a maximum since $f''(b_1) = -4$. By symmetry, you can see that if PII knew that PI is bidding half their private value, then the best thing for PII is to also bid half of his private value. Thus, $\beta_i(v_i) = v_i/2$ for $i = 1, 2$ is a Bayes–Nash equilibrium.

Now, let's consider the problem again from a slighty different point of view. Again, two bidders in a sealed-bid first-price auction, on which $V_i \sim U[0,1]$. Let $\beta_1 = \beta_2 = \beta$ (symmetric bid strategies) be a BNE. We will assume that β is differentiable and strictly increasing. We will study the problem from the perspective of bidder 1. Bidder 1 knows his own private value v_1, and knows that bidder 2 is using the same bidding strategy $\beta(V_2)$.

Let player 1 bid be $b_1 \in [\beta(0), \beta(1)] = [0, \beta(1)]$. Note that we assume that $\beta(0) = 0$, since if a player value the item at zero, then it will not bid a value higher than zero. Now, define the bid as $b_1 = \beta(w)$ on which $w \neq v_1$. The expected utility for bidder 1 is:

$$\begin{aligned}
\text{expected utility} &= v_1 a_1[b_1] - p_1[b_1] \\
&= v_1\mathbb{P}[b_1 > b_2] - b_1\mathbb{P}[b_1 > b_2] \\
&= (v_1 - b_1)\mathbb{P}[b_1 > b_2] \\
&= (v_1 - b_1)\mathbb{P}[\beta(w) > \beta(V_2)] \\
&= (v_1 - b_1)\mathbb{P}[w > V_2] \\
&= (v_1 - b_1)\mathbb{P}[V_2 < w] \\
&= (v_1 - b_1)w \\
&= (v_1 - \beta(w))w
\end{aligned}$$

If β is a BNE bidding strategy, then the expected utility is maximized for PI using his private value $\beta(v_1)$. Thus, $w \mapsto u_i(w \mid v_1)$ is maximized when $w = v_1$.

Let $f(w) = (v_1 - \beta(w))w$ and so:

$$\frac{\partial f}{\partial w} = v_1 - \beta(w) - \beta'(w)w$$

If β is a BNE, then f is maximized at $w = v_1$:

$$0 = v_1 - \beta(v_1) - \beta'(v_1)v_1$$

This defines a ordinary differential equation for β as:

$$\beta'(v_1)v_1 + \beta(v_1) = v_1 \to \frac{d}{dv_1}(v_1\beta(v_1)) = v_1$$

Integrating at both sides from 0 to v_1 we have:

$$\int_0^{v_1} (x\beta(x))dx = \int_0^{v_1} xdx \;\to\; v_1\beta(v_1) - 0 = \frac{v_1^2}{2} - 0 \;\to\; \beta(v_1) = \frac{v_1}{2}$$

that showcases that indeed bidding half the private value is a BNE.

Example 6.4 (Revenue to the seller)**.** Consider example 6.3, the 2 player sealed-bid first-price auction. We are interested in computing the expected revenue for the seller assuming that both players use their BNE bidding strategies. Then, the expected payoff is given by:

$$\mathbb{E}\left[\max\left(\frac{V_1}{2}, \frac{V_2}{2}\right)\right] = \frac{1}{2}\mathbb{E}[\max(V_1, V_2)]$$

Let $Y = \max(V_1, V_2)$. To compute its CDF we write:

$$\begin{aligned} F_Y(y) &= \mathbb{P}[Y \leq y] \\ &= \mathbb{P}[\max(V_1, V_2) \leq y] \\ &= \mathbb{P}[V_1 \leq y, V_2 \leq y] \\ &= \mathbb{P}[V_1 \leq y]\mathbb{P}[V_2 \leq y] \\ &= y^2 \end{aligned}$$

for $y \in [0, 1]$. Thus, $f_Y(y) = F_Y'(y) = 2y$, and its expected value:

$$\mathbb{E}[Y] = \mathbb{E}[\max(V_1, V_2)] = \int_0^1 y \cdot 2ydy = \frac{2}{3}$$

and hence the expected utility is given by:

$$\frac{1}{2}\mathbb{E}[\max(V_1, V_2)] = \frac{1}{2} \cdot \frac{2}{3} = \frac{1}{3}$$

Example 6.5 (Revenue to the seller: Vickrey auction)**.** Consider the same case as example 6.3, but now they participate in a sealed-bid second price (Vickrey) auction. We know that this type of auction is truthful, so bidding their true value is the BNE bidding strategy for the players. Now, for the seller, he will receive the second price as a payment, that implies that the expected utility for him is given by:

$$\mathbb{E}[\min(V_1, V_2)]$$

Let $Z = \min(V_1, V_2)$. We can compute the complementary CDF as:

$$\begin{aligned} 1 - F_Z(z) &= \mathbb{P}[Z > z] \\ &= \mathbb{P}[\min(V_1, V_2) > z] \\ &= \mathbb{P}[V_1 > z, V_2 > z] \\ &= \mathbb{P}[V_1 > z]\mathbb{P}[V_2 > z] \\ &= (1 - z)^2 \end{aligned}$$

and hence $F_Z(z) = 1-(1-z)^2$, from $z \in [0,1]$. To compute its density we simply take the derivative:

$$f_Z(z) = F_Z'(z) = 2(1-z)$$

Thus, the expected utility is:

$$\mathbb{E}[Z] = \int_0^1 z \cdot 2(1-z)dz = 1 - \frac{2}{3} = \frac{1}{3}$$

Theorem 6.1 (Vickrey auction is truthful)**.** The sealed-bid second-price auction is truthful for all players i and for any fixed set of bids of other players. That is, bidders i's utility is maximized by bidding their true value v_i.

Proof: Suppose the maximum of other bids is m (for $j \neq i$). Then:

$$u_i(\beta(v_i) \mid v_i) \leq \max(v_i - m, 0)$$

that is, the maximum of winning the auction by paying m or losing the auction. We have two cases:

- If $v_i \geq m$, is better to bid v_i to have positive utility (although it does not change the utility).
- If $v_i \leq m$, is also better to bid v_i to avoid negative utility. □

Remark: Recall the case of 2 bidders $V_1, V_2 \sim U[0,1]$ with realized values v_1, v_2. Consider two auctions, 1st price sealed-bid and 2nd price sealed-bid from the point view of bidder 1:

- For the first price case we have:

$$\mathbb{P}[\text{of winning}] = \mathbb{P}\left[\frac{V_2}{2} \leq \frac{v_1}{2}\right] = \mathbb{P}[V_2 \leq v_1] = v_1$$

$$\text{Conditional expectation of payment} = \frac{v_1}{2} \text{ (given that bidder 1 wins)}$$

$$\mathbb{E}[\text{payment}] = v_1 \times \frac{v_1}{2} = \frac{v_1^2}{2}$$

- For the second price case we have:

$$\mathbb{P}[\text{of winning}] = \mathbb{P}[V_2 \leq v_1] = v_1$$

$$\text{Conditional expectation of payment} = \mathbb{E}[V_2 \mid V_2 \leq v_1] = \mathbb{E}[\underbrace{Z}_{\sim U[0,v_1]}] = \frac{v_1}{2}$$

$$\mathbb{E}[\text{payment}] = v_1 \times \frac{v_1}{2} = \frac{v_1^2}{2}$$

Then the expected payment for the seller is given by:

$$\begin{aligned}
\mathbb{E}[\text{payment}] &= \mathbb{E}[\text{payment of P1}] + \mathbb{E}[\text{payment of P2}] \\
&= 2\mathbb{E}[\text{payment of P1}] \\
&= 2\mathbb{E}\left[\frac{V_1^2}{2}\right] \\
&= \mathbb{E}[V_1^2] \\
&= \int_0^1 V_1^2 \cdot 1 \, dV_1 = \frac{V_1^3}{3}\bigg|_0^1 = \frac{1}{3}
\end{aligned}$$

6.3 Revenue equivalence

In order find Bayes-Nash equilibrium bidding strategies, by instead looking at functions of bids b, we look at allocation probabilities, expected payments and expected payments as functions of alternatives valuations w (and not v). That is, a bidder i will use his bidding strategy β_i, but instead of bidding $\beta(v_i)$, he will bid as if it private value was w, and bid $\beta_i(w)$. We will use the following notation to denote probabilities of using alternative valuations:

$$\begin{aligned}
&a_i(w) := a_i[\beta_i(w)] = \mathbb{P}[\text{bidder } i \text{ wins by bidding } \beta_i(w) \text{ and others bid } \beta_{-i}(V_{-i})] \\
&p_i(w) := p_i[\beta_i(w)] = \text{Expected payment of bidder } i \\
&u_i(w \mid v_i) := u_i[\beta_i(w) \mid v_1] = v_i a_i(w) - p_i(w)
\end{aligned}$$

We say that $(\beta_1, \beta_2, \ldots, \beta_n)$ is a BNE only if:

$$u_i(v_i \mid v_i) \geq u_i(w \mid v_1)$$

for all bidders i, for all alternative valuations w and for all private values v_i.

Aside: Order Statistics

Consider n independent and identically distributed (iid) random variables $X_1, X_2, \ldots, X_n \sim F, f$. Let $U := \max\{X_1, X_2, \ldots, X_n\}$. We are interested in computing the cumulative distribution function (CDF) $F_U(u)$ and probability density function (pdf) $f_U(u)$. Note that:

$$\begin{aligned}
F_U(u) &= \mathbb{P}[U \leq u] \\
&= \mathbb{P}[X_1 \leq u, X_2 \leq u, \ldots, X_n] \\
&= \prod_{i=1}^{n} \mathbb{P}(X_i \leq u) \\
&= F(u)^n
\end{aligned}$$

and hence:

$$f_U(u) = \frac{d}{du} F_U(u) = n f(u) F(u)^{n-1}$$

Let $V := \min\{X_1, X_2, \ldots, X_n\}$. In this case the CDF is given by:

$$\begin{aligned}
1 - F_V(v) &= \mathbb{P}[V \geq v] \\
&= \mathbb{P}[X_1 \geq v, X_2 \geq v, \ldots, X_n \geq v] \\
&= \prod_{i=1}^{n} (1 - F(v)) \\
&= (1 - F(v))^n
\end{aligned}$$

Thus, $F_V(v) = 1 - (1 - F(v))^n$ and hence:

$$f_V(v) = n f(v) (1 - F(v))^{n-1}$$

In general, we can compute the k-th order statistics from n random variables $X_1, X_2, \ldots, X_n \sim F, f$, denoted as:

$$X_{(1)} < X_{(2)} < \ldots < X_{(k)} < \ldots < X_{(n)}$$

Observe that the n-th order statistics is the maximum case, while 1-st order statistics is the minimum case. If $n = 2m+1$ is odd, then we say that $X_{(m)}$ is the median of the X_i random variables. The density function of the k-th order statistics can be computed as follow:

$$f_k(x) = \frac{n!}{(k-1)!(n-k)!} f(x)[F(x)]^{k-1}[1 - F(x)]^{n-k}$$

Example 6.6. Consider $X_1, \ldots, X_n \sim U[0, 1]$, then:

$$\begin{aligned} f_k(x) &= \frac{n!}{(n-k)!(k-1)!} x^{k-1}(1-x)^{n-k} \\ &= \frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)} x^{a-1}(1-x)^{b-1},\ x \in [0,1] \\ &\sim B(a,b) \ \leftarrow \text{ Beta distribution} \end{aligned}$$

on which $a = k$ and $b = n - k + 1$.

Now, let's return to the setting for the remainder of the section, on which $V_i \sim F$ iid, $\beta_i = \beta$, $\forall i$ and β is strictly increasing. In such case we have that $a_i(w) = a(w)$, $p_i(w) = p(w)$ and $u_i(w \mid v_i) = u(w \mid v_i)$. Then, from the point of view of bidder i:

$$\begin{aligned} a(w) &= \mathbb{P}[\beta(w) > \beta(V_j), \forall j \neq i] \\ &= \mathbb{P}[\beta(w) > \beta(\max_{j \neq i} V_j)] \\ &= \mathbb{P}[w > \max_{j \neq i} V_j] \\ &= \mathbb{P}[\max_{j \neq i} V_j < w] \\ &= F(w)^{n-1} \end{aligned}$$

Then, the expected utility for bidder i by using its alternative valuation w is given by:

$$u(w \mid v_i) = v_i a(w) - p(w) = v_i F(w)^{n-1} - p(w)$$

on which $p(w)$ is the expected payment, that should depend on the auction type? If $(\beta, \ldots, \beta)$ is a BNE, then:

$$u(v_i \mid v_i) \geq u(w \mid v_i),\ \forall w, v_i$$

That implies that the derivative of $u(w \mid v_i)$ with respect to w vanishes at $w = v_i$. Then:

$$v_i a'(w) - p'(w) = 0 \ \Big|_{w = v_i}$$

Thus:

$$p'(v_i) = v_i a'(v_1)$$

Integrating both sides from 0 to v_i, and using w as dummy variable we have:

$$\rightarrow \int_0^{v_i} p'(w) = \int_0^{v_i} wa'(w)dw \quad \text{integrating by parts}$$
$$\rightarrow p(v_i) - p(0) = v_1 a(v_1) - \int_0^{v_i} a(w)dw$$
$$\rightarrow p(v_i) = v_i a(v_i) - \int_0^{v_i} a(w)dw$$

Observe that the expected payment only depend on $a(v)$, i.e. the allocation rule. That is, if two different auctions are used with the same allocation rule (e.g. 1st price and 2nd price use the same allocation rule of highest bid wins) then the expected payments are the same. The figure is as follow:

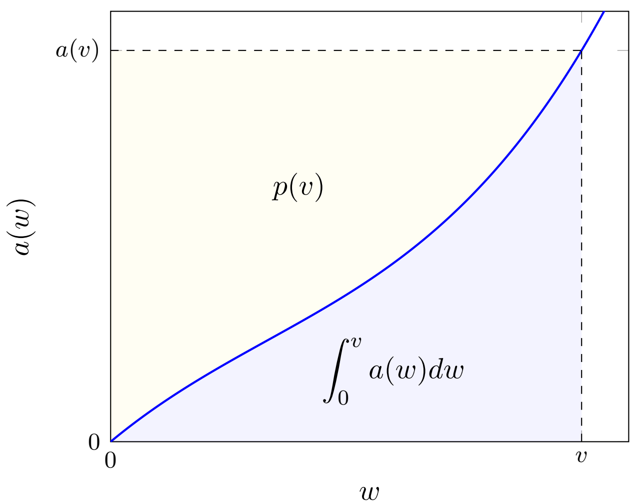

Figure 6.2: Diagram that showcases the expected payment of bidders. The area filled in yellow represent the expected payment of bidder i.

Recall that the expected utility for the private value is given by $u(v \mid v) = va(v) - p(v)$ and hence:

$$u(v \mid v) = va(v) - p(v) = va(v) - \left(va(v) - \int_0^v a(w)dw \right) = \int_0^v a(w)dw$$

and hence, the expected utility for the bidder is given by the area filled in blue in Figure 6.2.

6.3.1 Revenue equivalence theorem

The following theorem is stated as 14.4.2 in [1]. Suppose that each agent private value V_i is drawn iid from the distribution F, on which $F \in [0, h]$, strictly increasing cumulative distribution. Consider any n-bidder single item auction in which the item is allocated to the highest bidder (i.e. same allocation rule) with $u_i(0) = 0, \ \forall i$.

Then, assuming bidders employ a symmetric bidding profile $\beta_i = \beta, \ \forall i$, with β strictly increasing in $[0, h]$. Then the following holds:

i. If $(\beta, \beta, \ldots, \beta)$ is a BNE, then for a bidder with private value v, the following is true:

$$a(v) = F(v)^{n-1} \qquad \text{and} \qquad p(v) = va(v) - \int_0^v a(w)dw$$

ii. On the other hand, if $a(v) = F(v)^{n-1}$ and $p(v) = va(v) - \int_0^v a(w)dw$, holds for $(\beta, \beta, \ldots, \beta)$, then for any i and for all $v, w \in [0, h]$, the following is true:

$$u(v \mid v) \geq u(w \mid v)$$

that implies that $(\beta, \beta, \ldots, \beta)$ is a BNE.

Proof: Note that part i., was already proven in the previous section. Now, for ii., we can show it visually. Consider the case of $w_1 \leq v$:

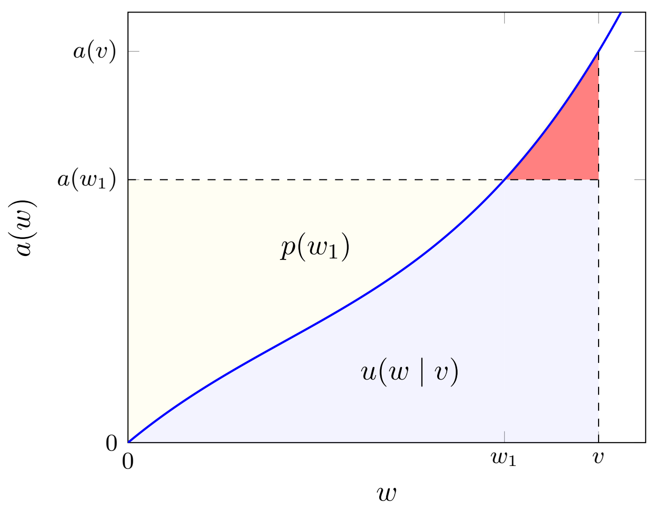

Figure 6.3: Diagram that showcases the expected payment of bidding w_1. The blue area depicts the expected utility. The red area showcases the missed utility of not using private value v, that proves that $u(w_1 \mid v) \leq u(v \mid v)$.

Indeed, note that:

$$u(v \mid v) = va(v) - p(v) = \int_0^v a(x)dx$$

and

$$\begin{aligned} u(w \mid v) &= va(w) - p(w) \\ &= va(w) - wa(w) + wa(w) - p(w) \\ &= (v - w)a(w) + wa(w) - p(w) \\ &= (v - w)a(w) + \int_0^w a(x)dx \end{aligned}$$

Then for the case of $w_1 < v$ we have:

$$\begin{aligned}
u(v \mid v) - u(w_1 \mid v) &= \int_0^v a(x)dx - (v - w_1)a(w_1) - \int_0^{w_1} a(x)dx \\
&= \int_{w_1}^v a(x)dx - \int_{w_1}^v a(w_1)dx \\
&= \int_w^v a(x) - a(w_1)dx \\
&\geq 0
\end{aligned}$$

since $a(x)$ is increasing. Then $u(v \mid v) \geq u(w_1 \mid v)$.

Now consider the case $w_2 > v$:

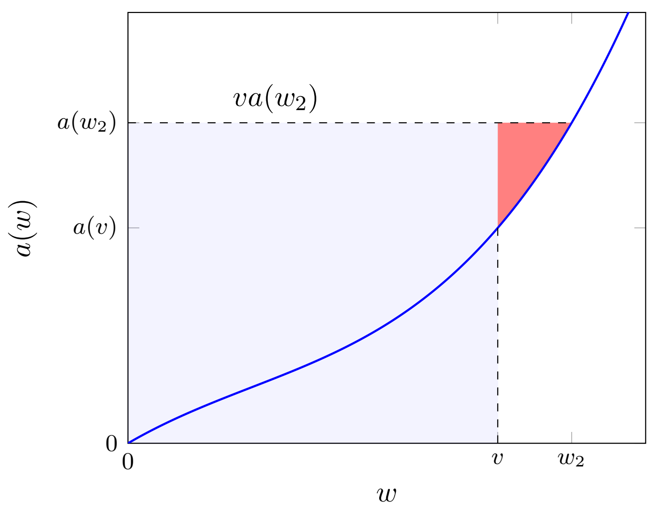

Figure 6.4: Diagram that showcases the expected payment of bidding w_2. The blue rectangle showcase the expected revenue given by $va(w_2)$ while the area between $a(w_2)$ and $a(w)$ subtracts that revenue. The red area is the extra payment that the bidder must pay for using an alternative value w_2 instead of v, that proves that $u(w_2 \mid v) \leq u(v \mid v)$.

on which we have:

$$\begin{aligned}
u(v \mid v) - u(w_2 \mid v) &= \int_0^v a(x)dx - (v - w_2)a(w_2) - \int_0^{w_2} a(x)dx \\
&= -\int_v^{w_2} a(x)dx + \int_v^{w_2} a(w_2)dx \\
&= \int_v^{w_2} a(w_2) - a(x)dx \\
&\geq 0
\end{aligned}$$

due to the increasing nature of $a(w)$. Then $u(v \mid v) \geq u(w_2 \mid v)$. $\square$

Corollary: (14.4.4 in [1]) Under the assumptions of the revenue equivalence theorem (14.4.2), then it holds that:

$$p(v) = F(v)^{n-1}\mathbb{E}\left[\max_{i\leq n-1} V_i \mid \max_{i\leq n-1} V_i \leq v\right]$$

This is obviously true, since the second-price auction is truthful, and hence the expected payment is simply given for the probability of winning $F(v)^{n-1}$ times the second max bid, conditioned that all other bids are lower than the private value (and hence the bid) v, i.e. the aforementioned expression.

6.3.2 Bidding strategy on first-price auctions

We will use the revenue equivalence theorem to obtain the BNE symmetric bidding strategy for all bidders. But first, we will compute it traditionally. Observe that in a first-price sealed-bid auction, you should bid the maximum of all other players, that is, to win the auction you should bid an ε higher than the maximum of all other players. In such case we have:

$$\beta(v) = \mathbb{E}\left[\max_{i\leq n-1} V_i \mid \max_{i\leq n-1} V_i \leq v\right]$$

To compute this, we will use the following property for the expectation of non-negative random variables:

$$\mathbb{E}[X] = \int_0^\infty \mathbb{P}(X \geq x)dx, \quad \text{for } X \text{ non-negative}$$

Let $Y := \max_{i\leq n-1} V_i$. With that we have:

$$\begin{aligned}
\beta(v) &= \mathbb{E}[Y \mid Y \leq v] \\
&= \int_0^v \mathbb{P}[Y \geq w \mid Y \leq v] \\
&= \int_0^v \frac{\mathbb{P}(w \leq Y \leq v}{\mathbb{P}(Y \leq v} dw \\
&= \frac{1}{F_Y(v)} \int_0^v [F_Y(v) - F_Y(w)]dw \\
&= \frac{1}{F(v)^{n-1}} \int_0^v F(v)^{n-1} - \underbrace{F(w)^{n-1}}_{a(w)}]dw \\
&= \int_0^v \left[1 - \left(\frac{F(w)}{F(v)}\right)^{n-1}\right] dw
\end{aligned}$$

Now, using the revenue equivalence theorem and $p(v) = \beta(v)a(v)$ we have:

$$\begin{aligned}
\beta(v) &= \frac{p(v)}{a(v)} \quad \text{for } a(v) > 0 \\
&= \frac{p(v)}{F(v)^{n-1}} \\
&= \frac{va(v) - \int_0^v a(w)dw}{F(v)^{n-1}} \\
&= v - \frac{1}{F(v)^{n-1}} \int_0^v a(w)dw \\
&= v - \frac{1}{F(v)^{n-1}} \int_0^v F(w)^{n-1}dw \\
&= \int_0^v \left[1 - \left(\frac{F(w)}{F(v)}\right)^{n-1}\right] dw
\end{aligned}$$

as expected.

Example 6.7. Consider $V_1, \ldots, V_n \sim F = U[0,1]$ and hence $F(x) = x$ for $x \in [0,1]$. We have that:

$$\beta(v) = \int_0^v 1 - \frac{w^{n-1}}{v^{n-1}} dw = v - \frac{1}{n}\frac{v^n}{v^{n-1}} = \frac{n-1}{n}v$$

i.e. for $n = 2$ bidders, we have that the BNE strategy is to bid 1/2 of your private value, and in case of $n = 3$ bidders, the BNE is to bid 2/3 of your private value.

Example 6.8. Consider two bidders such that $V_1, V_2 \sim \exp(1)$, on which $F(x) = 1 - e^{-x}$. Then:

$$\begin{aligned}
\beta(v) &= \int_0^v 1 - \frac{(1-e^{-w})^{2-1}}{(1-e^{-v})^{2-1}} dw \\
&= v - \frac{1}{1-e^{-v}} \int_0^v 1 - e^{-w} dw \\
&= v - \frac{1}{1-e^{-v}}(v + (e^{-v} - 1)) \\
&= v - \frac{v}{1-e^{-v}} + 1 \\
&= 1 + v\left(1 - \frac{1}{1-e^{-v}}\right) \\
&= 1 - \frac{ve^{-v}}{1-e^{-v}}
\end{aligned}$$

6.3.3 General presentation of revenue equivalence

To guarantee equlibirium, it's necessary and sufficient that $a(\cdot)$ is increasing, in the case of $V_1, V_2, \ldots, V_n \sim F$, i.i.d.

⇒ Necessity: If bidding strategies are BNE, then $a(\cdot)$ is increasing.

Assume $(\beta, \ldots, \beta)$ is a BNE, then the utility to bidder i is maximized at their private value:

$$u(v \mid v) \geq u(w \mid v), \quad \forall v, w \text{ and } \forall \text{ bidders}$$

Then:

$$\begin{aligned} & va(v) - p(v) \geq va(w) - p(w) \\ \rightarrow\ & v(a(v) - a(w)) \geq p(v) - p(w) \end{aligned} \tag{i}$$

Now, reverse the roles of v and w (so that now the private value is w):

$$\begin{aligned} & u(w \mid w) \geq u(v \mid w) \\ \rightarrow\ & w(a(w) - a(v) \geq -p(v) + p(w) \end{aligned} \tag{ii}$$

Adding (i) and (ii) we have:

$$(v - w)(a(v) - a(w)) \geq 0$$

that proves that $a(\cdot)$ is increasing.

$\Leftarrow$ Sufficiency: If $a(\cdot)$ is increasing then $(\beta, \ldots, \beta)$ is a BNE.

This was already showed in the revenue equivalence theorem in the previous subsection.

6.4 Characterization of Bayes-Nash equilbrium

Let $\mathcal{A}$ be a single item auction with n bidders, with independent $V_i \sim F_i$ (not necessary identical), strictly increasing and continuous on $[0, h_i]$, that is $F(0) = 0$, $F(h_i) = 1$, and note that h_1 could be infinite.

(a) If $(\beta_1, \beta_2, \ldots, \beta_n)$ is a BNE, then for each bidder i:

   i. $a_i(v)$ is monotone increasing.

   ii. The utility $u_i(v)$ is a convex function of v and can be computed up to an integration constant as:

   $$u_i(v) = \int_0^v a_i(x)dx + u_i(0)$$

   iii. The expected payment is determined by the allocation probabilities up to an integration constant $p_i(0)$ as:

   $$\begin{aligned} p_i(v) &= va_i(v) - \int_0^v a_i(x)dx - p_i(0) \\ &= \int_0^v x a_i'(x)dx - p_i(0) \end{aligned}$$

Aside: Convex functions
A real valued function $f : [a, b] \rightarrow \mathbb{R}$ is called convex if for any $x_1, x_2 \in [a, b]$ and $\alpha \in [0, 1]$, then:

$$f(\alpha x_1 + (1 - \alpha)x_2) \leq \alpha f(x_1) + (1 - \alpha)f(x_2)$$

that is, the line between two points of the function is always above the function itself.

Corollary: For two single item auctions $\mathcal{A}$ and $\mathcal{A}'$ with same allocation rule a_i, a'_i with same integration constant $p_i(0) = p'_i(0) = 0$ for every bidder, then we have that:

$$p_i(v_i) = p'_i(v_i)$$

for every bidder and private value. This implies that we have revenue equivalence in such case.

Remark: Be careful, that if V_i is not i.i.d., equilibrium may not be symmetric, and hence the first price auction may not allocate to highest bidder and then first price and second price auctions may not satisfy revenue equivalence.

6.5 Revelation Principle

Definition 6.8 (Bayes-Nash incentive compatible)**.** Call an auction $\mathcal{A}$ **Bayes-Nash incentive compatible** if bidding truthfully is a BNE for $\mathcal{A}$: $\beta_i(v_i) = v_i,\ \forall v, i$.

Definition 6.9 (Direct auction)**.** Call an auction $\mathcal{A}$ a **direct auction**, if every bidder only submit a single bid.

Then, the revelation principle is a follows:

Let $\mathcal{A}$ be a direct auction and $(\beta_1, \ldots, \beta_n)$ is a BNE for A. Then, there exists another direct auction $\mathcal{A}'$, which is Bayes-Nash incentive compatible, such that $\mathcal{A}$ and $\mathcal{A}'$ have the same winners and payments in equilibrium.

References

[1] A. R. Karlin and Y. Peres, *Game theory, alive.* American Mathematical Soc., 2017, vol. 101.

[2] T. S. Ferguson, *Game Theory.* UCLA, 2014.

[3] T. Basar and G. J. Olsder, *Dynamic noncooperative game theory.* Siam, 1999, vol. 23.

[4] S. Boyd and L. Vandenberghe, *Convex optimization.* Cambridge university press, 2004.

[5] D. Gale and L. S. Shapley, "College admissions and the stability of marriage," *The American Mathematical Monthly*, vol. 69, no. 1, pp. 9–15, 1962.

[6] A. K. Dixit and B. Nalebuff, *The art of strategy: a game theorist's guide to success in business & life.* WW Norton & Company, 2008.