

Two-Player Games

Suppose that Alice and Bob are playing a game. During one round of this game, Alice can make one of m possible plays, and Bob can make one of n possible plays. If Alice chooses play i and Bob chooses play j , the payoff to Alice is given by a_{ij} (note here that a_{ij} can be negative, in which case Bob benefits from the round of play). We call the $m \times n$ matrix $A := (a_{ij})$ the *payoff matrix* for the game.

A *pure strategy* for one of the players is a strategy in which the player always selects the same play. Note that a pure strategy is not necessarily a good approach to playing a game. For instance, if Alice decides to always make play 3, and Bob discovers this strategy, he can also choose a pure strategy that minimizes his losses by always choosing play $\operatorname{argmin}\{a_{3j} : j = 1, 2, \dots, n\}$.

A different approach for the players to take into this game is to pick each possible play with some predetermined probability. That is, Alice will select play i with probability x_i , and Bob will select play j with probability y_j . With this strategy, the expected payoff for Alice is given by

$$E[\text{payoff}] = \sum_{i=1}^m \sum_{j=1}^n x_i a_{ij} y_j = \mathbf{x}^T \mathbf{A} \mathbf{y},$$

where $\mathbf{x} := (x_1, x_2, \dots, x_m)$ and $\mathbf{y} = (y_1, y_2, \dots, y_n)$.

If she is in any way competitive, Alice will want to somehow maximize her expected payoff $\mathbf{x}^T \mathbf{A} \mathbf{y}$. But what will Bob do? Well, if Bob can somehow discover Alice's strategy \mathbf{x} , he will try to select his strategy \mathbf{y} such that $\mathbf{x}^T \mathbf{A} \mathbf{y}$ is minimized. More formally, for a fixed \mathbf{x} Bob will go about solving the following linear program.

$$\begin{aligned} & \text{minimize} && \mathbf{x}^T \mathbf{A} \mathbf{y} \\ & \text{subject to} && \sum_{j=1}^n y_j = 1 \\ & && y_j \geq 0 \text{ for all } j = 1, 2, \dots, n \end{aligned}$$

Claim 0.1. *For a fixed \mathbf{x} , we have that*

$$\min_{\mathbf{y}} \mathbf{x}^T \mathbf{A} \mathbf{y} = \min_j \sum_{i=1}^m a_{ij} x_i,$$

where \mathbf{y} is such that $\sum_{j=1}^n y_j = 1$ and $\mathbf{y} \geq 0$, and j ranges from 1 to n .

Proof. Fix \mathbf{x} . We have that $\min_{\mathbf{y}} \mathbf{x}^T \mathbf{A} \mathbf{y} \leq \sum_{i=1}^m a_{ij} x_i$ for any $j = 1, 2, \dots, n$, since $\sum_{i=1}^m a_{ij} x_i = \mathbf{x}^T \mathbf{A} \boldsymbol{\epsilon}_j$, where $\boldsymbol{\epsilon}_j$ is the j^{th} standard basis vector. Thus, $\min_{\mathbf{y}} \mathbf{x}^T \mathbf{A} \mathbf{y} \leq \min_j \sum_{i=1}^m a_{ij} x_i$.

For the other inequality, consider

$$\min_{\mathbf{y}} \mathbf{x}^T \mathbf{A} \mathbf{y} = \min_{\mathbf{y}} \sum_{j=1}^n y_j \left(\sum_{i=1}^m x_i a_{ij} \right).$$

Notice that each term of the form $\sum_{i=1}^m x_i a_{ij}$ must be greater than or equal to $\min_j \sum_{i=1}^m x_i a_{ij}$, which is independent of j . Thus, we have

$$\begin{aligned} \min_{\mathbf{y}} \mathbf{x}^T \mathbf{A} \mathbf{y} &\geq \min_{\mathbf{y}} \sum_{j=1}^n y_j \left(\min_j \sum_{i=1}^m x_i a_{ij} \right) \\ &= \min_{\mathbf{y}} \left(\min_j \sum_{i=1}^m x_i a_{ij} \right) \sum_{j=1}^n y_j \\ &= \min_{\mathbf{y}} \left(\min_j \sum_{i=1}^m x_i a_{ij} \right) \cdot 1 \\ &= \min_j \sum_{i=1}^m x_i a_{ij}. \end{aligned}$$

□

A consequence of Claim 0.1 is that if Bob knows Alice's strategy \mathbf{x} , he has a pure strategy to counter it. Alice is aware of this, so she must select her strategy \mathbf{x} with the supposition that eventually during the game Bob will discover \mathbf{x} and counter with a better strategy. So, Alice must solve the problem

$$\max_{\mathbf{x}} \min_{\mathbf{y}} \mathbf{x}^T \mathbf{A} \mathbf{y},$$

which by Claim 0.1 is equivalent to solving

$$\begin{aligned} &\text{maximize } \min_j \sum_{i=1}^m a_{ij} x_i \\ &\text{subject to } \sum_{i=1}^m x_i = 1 \\ &\mathbf{x} \geq 0. \end{aligned}$$

This problem is now independent of Bob's strategy \mathbf{y} .

For Alice to maximize the quantity $\min_j \sum_{i=1}^m a_{ij}x_i$, she could try to maximize a lower bound on the set $\{\sum_{i=1}^m a_{ij}x_i : j = 1, 2, \dots, n\}$. This approach yields an equivalent problem:

$$\begin{aligned}
 & \text{maximize } z \\
 & \text{subject to } \sum_{i=1}^m x_i = 1 \\
 & \quad z - \sum_{i=1}^m a_{ij}x_i \leq 0 \text{ for all } j \\
 & \quad x_i \geq 0 \text{ for all } i \\
 & \quad z \text{ free.}
 \end{aligned}$$

This is now a linear program that Alice can solve to determine her best strategy.

Consider the dual of the linear program above:

$$\begin{aligned}
 & \text{minimize } w \\
 & \text{subject to } \sum_{j=1}^n y_j = 1 \\
 & \quad w - \sum_{j=1}^n a_{ij}y_j \geq 0 \text{ for all } i = 1, 2, \dots, m \\
 & \quad y_j \geq 0 \text{ for all } j = 1, 2, \dots, n \\
 & \quad w \text{ free.}
 \end{aligned}$$

This dual linear program determines Bob's optimal strategy. Here's why: first, minimizing w is equivalent to maximizing $-w$ over the same region. Also, we may multiply certain constraints defining the feasible set by -1 to obtain an equivalent representation. Thus, the linear program above is equivalent to

$$\begin{aligned}
& \text{maximize } (-w) \\
& \text{subject to } \sum_{j=1}^n y_j = 1 \\
& \quad -w - \sum_{j=1}^n (-a_{ij})y_j \leq 0 \text{ for all } i = 1, 2, \dots, m \\
& \quad y_j \geq 0 \text{ for all } j = 1, 2, \dots, n \\
& \quad w \text{ free.}
\end{aligned}$$

We may now replace $-w$ by a new dummy variable t . Since $-w$ is free, t is also free. Thus, we have

$$\begin{aligned}
& \text{maximize } t \\
& \text{subject to } \sum_{j=1}^n y_j = 1 \\
& \quad t - \sum_{j=1}^n (-a_{ij})y_j \leq 0 \text{ for all } i = 1, 2, \dots, m \\
& \quad y_j \geq 0 \text{ for all } j = 1, 2, \dots, n \\
& \quad t \text{ free.}
\end{aligned}$$

Recall that a_{ij} was defined to be the payoff to Alice. Thus, the quantity $-a_{ij}$ is the payoff to Bob. With this in mind, the linear program above is in the exact form we derived earlier to determine a player's optimal strategy. We conclude that the problems of optimizing Alice's strategy and optimizing Bob's strategy are duals of each other. Strong duality then implies that the optimal cost is the same for both problems.

An Example: The Game of Mora. The game works as follows: each player chooses whether to hold 1 or 2 coins behind his or her back. In addition to this choice of 1 or 2 coins, each player also fixes a guess of how many coins his or her opponent is holding. Both players then reveal how many coins they are holding. If both made correct guesses, then the payoff to each player is 0. Similarly, if neither guess was correct the payoff to both is 0. If one guesses correctly while the other doesn't, the loser forfeits the coins he or she was holding to the winner.

Suppose that Alice and Bob are playing the Game of Mora. Find optimal strategies for both Alice and Bob, and determine the expected payoff to each player.

Solution:

Let $[i, j]$ denote a play in which the player is holding i coins and guesses that his or her opponent is holding j coins. With Alice's possible plays represented by rows and Bob's by the columns, the payoff matrix for Alice is given by

	[1, 1]	[1, 2]	[2, 1]	[2, 2]
[1, 1]	0	1	-1	0
[1, 2]	-1	0	0	2
[2, 1]	1	0	0	-2
[2, 2]	0	-2	2	0

An optimal strategy for Alice is given by solving the following linear program:

$$\text{maximize } x_5$$

subject to

$$\begin{array}{rccccrcr}
 & & x_2 & -x_3 & & +x_5 & \leq & 0 \\
 -x_1 & & & & +2x_4 & +x_5 & \leq & 0 \\
 x_1 & & & & -2x_4 & +x_5 & \leq & 0 \\
 & -2x_2 & +2x_3 & & & +x_5 & \leq & 0 \\
 x_1 & +x_2 & +x_3 & +x_4 & & & = & 1
 \end{array}$$

$$\begin{array}{l}
 x_i \geq 0 \text{ for } i = 1, 2, 3, 4 \\
 x_5 \text{ free}
 \end{array}$$

One solution to this linear program is $(0.667, 0, 0, 0.333, 0)$, which translates to an optimal strategy for Alice being to play $[1, 1]$ with probability

0.667 and to play $[2, 2]$ with probability 0.333. That the value of x_5 in this solution is 0 implies that Alice's maximum expected payoff is 0.

An optimal strategy for Bob is given by solving the following linear program:

$$\begin{array}{rcll}
 & & & \text{minimize } y_5 \\
 \text{subject to} & & & \\
 & -y_2 & +y_3 & +y_5 \geq 0 \\
 y_1 & & & -2y_4 + y_5 \geq 0 \\
 -y_1 & & & +2y_4 + y_5 \geq 0 \\
 & 2y_2 & -2y_3 & +y_5 \geq 0 \\
 y_1 & +y_2 & +y_3 & +y_4 = 1 \\
 & & & y_i \geq 0 \text{ for } i = 1, 2, 3, 4 \\
 & & & y_5 \text{ free}
 \end{array}$$

Solving this tells us that $(0.667, 0, 0, 0.333)$ is also an optimal strategy for Bob, and Bob's optimal expected payoff is also 0.

Example: The Number Game. This game works as follows. Alice and Bob both choose an integer from the set $\{1, 2, 3, 4\}$. If Bob's guess is strictly less than Alice's, Bob wins as many dollars as his choice. Otherwise, Alice wins exactly one dollar. To make the game interesting, there are two exceptions to the above rules. If both Bob and Alice choose 1, Alice wins 4 dollars. Finally, if Bob chooses 4 and Alice chooses 1, Bob wins 3 dollars. Calculate the optimal strategies for Alice and Bob, and determine the expected payoff to each player.

Solution:

With Alice's possible plays represented by rows and Bob's by the columns, the payoff matrix for Alice is given by

	1	2	3	4	
1	4	1	1	-3	
2	-1	1	1	1	.
3	-1	-2	1	1	
4	-1	-2	-3	1	

Alice's optimal strategy can be computed by solving the following linear program (the construction of this linear program is slightly more delicate than it was in the previous example, since the payoff matrix is not skew-symmetric).

maximize x_5

subject to

$$\begin{array}{rcccccc}
 -4x_1 & +x_2 & +x_3 & +x_4 & +x_5 & \leq & 0 \\
 -x_1 & -x_2 & +2x_3 & +2x_4 & +x_5 & \leq & 0 \\
 -x_1 & -x_2 & -x_3 & 3x_4 & +x_5 & \leq & 0 \\
 3x_1 & -x_2 & -x_3 & -x_4 & +x_5 & \leq & 0 \\
 x_1 & +x_2 & +x_3 & +x_4 & & = & 1
 \end{array}$$

$$\begin{array}{l}
 x_i \geq 0 \text{ for } i = 1, 2, 3, 4 \\
 x_5 \text{ free.}
 \end{array}$$

Solving this linear program tells shows that an optimal strategy for Alice is given by the vector $(0.222, 0.778, 0, 0)$.

Bob's optimal strategy is found by solving

minimize y_5

subject to

$$\begin{array}{rcccccc}
-4y_1 & -y_2 & -y_3 & +3y_4 & +y_5 & \geq & 0 \\
y_1 & -y_2 & -y_3 & -y_4 & +y_5 & \geq & 0 \\
y_1 & +2y_2 & -y_3 & -2y_4 & +y_5 & \geq & 0 \\
y_1 & 2y_2 & 3y_3 & -y_4 & +y_5 & \geq & 0 \\
y_1 & +y_2 & +y_3 & y_4 & & = & 1
\end{array}$$

$$\begin{array}{l}
y_i \geq 0 \text{ for } i = 1, 2, 3, 4 \\
y_5 \text{ free}
\end{array}$$

Solving this tells us that an optimal strategy for Bob is given by (0.444, 0, 0, 0.556).

In both cases, we see that when Alice and Bob utilize their optimal strategies, the game tends to favor Alice - her long-term expected payoff is about \$0.11 per play.