Math 3272: Linear Programming[1] | Mikhail Lavrov

Lecture 18: Matrix games

October 18, 2022 | *Kennesaw State University*

1 Introduction to games

1.1 Matrix games

We are going to consider two-player games where the players simultaneously pick strategies, and are rewarded with a **payoff** based on their choices. A classic example is the prisoner's dilemma (which we're only going to briefly mention).

Here, the two players (Alice and Bob) are suspected of a bank robbery (2 years in prison), but arrested only for some minor thing like tax evasion (1 year in prison). The prosecutor offers each of them a bargain: testify against the other, and the tax evasion charges are dropped. As a result:

- If Alice and Bob both stay silent, each spends 1 year in prison.
- If Alice testifies against Bob, Alice goes free and Bob spends 3 years in prison (and vice versa).
- If Alice and Bob both testify against each other, they both spend 2 years in prison.

We can represent spending t years in prison by a payoff of $-t$. (The goal is always to maximize your payoff, so bad outcomes get negative payoffs.)

We can represent the prisoner's dilemma (and more generally, any such game) by a payoff matrix. Assign each of Alice's strategies a row, and each of Bob's strategies a column. At the intersection of the row and the column, record Alice and Bob's respective payoffs:

	Bob stays silent	Bob testifies
Alice stays silent	$(-1, -1)$	$(-3, 0)$
Alice testifies	$(0, -3)$	$(-2, -2)$

It is best for both players to stay silent, rather than both testify, so they should agree not to testify. However, each individual player is better of testifying, no matter what the other does, so they should betray that agreement. However, that leaves both players worse off. This weird behavior gives the prisoner's dilemma a rich and complicated dynamic...

1.2 Zero-sum games

...which we're going to ignore entirely, because in this class we will only fully analyze zero-sum games.

[1] This document comes from the Math 3272 course webpage: https://facultyweb.kennesaw.edu/mlavrov/courses/3272-fall-2022.php

A **zero-sum game** is one in which any hope of cooperation between the players is eliminated, because Alice's payoff is the negative of Bob's payoff. (They sum to 0.) Whatever outcome helps one player, hurts the other player equally.

Though a lot of the concepts we will introduce make sense for general matrix games, we will mostly rely on the assumption that Alice should expect Bob to make the choice that's worst for Alice, and vice versa. This is a bad assumption in general: Bob wants to make the choice that's best for Bob, whether or not that hurts Alice! But in the case of zero-sum games, the two are equivalent.

In any case, our goal in analyzing these games will be to determine what Alice and Bob's optimal strategies are, and what the resulting payoff is for both players.

2 Strategies

We will look at some examples of zero-sum games to illustrate a few cases where we can find the optimal strategies easily, and the general case which is more complicated.

2.1 Dominated strategies

Consider the following game, called "higher number". It's a pretty stupid game. Alice and Bob each hold up one, two, or three fingers. The player holding up fewer fingers gives $1 to the other player. The payoff matrix for the two players is:

	Bob: one	Bob: two	Bob: three
Alice: one	$(0,0)$	$(-1,1)$	$(-1,1)$
Alice: two	$(1,-1)$	$(0,0)$	$(-1,1)$
Alice: three	$(1,-1)$	$(1,-1)$	$(0,0)$

It is immediate to see that both players want to hold up more fingers rather than fewer. The point of this example is just to introduce two terms that help in analyzing matrix games.

- We say that one strategy of a player **dominates** another strategy if, no matter what the other player does, the first strategy is better than (or at least as good as) the second.

  Here, "two" dominates "one" and "three" dominates both "one" and "two", for both players: no matter what the other player does, you can never go wrong by holding up more fingers.

- If a strategy dominates every other strategy, we call it a **dominant strategy**.

  Here, "three" is the dominant strategy for both players.

If we can identify a dominant strategy for a game, there is nothing left to analyze—at least in the case of zero-sum games. A player with an optimal strategy should always play it. With that assumption, the only thing left for the other player to do is to pick the best response to that strategy.

Even if there is no dominant strategy, it makes sense to eliminate from consideration any strategy that's dominated by another strategy. After all, that other strategy is never worse. This can help us simplify the problem and make the payoff matrix smaller.

2.2 Saddle points

Suppose that Alice and Bob are two generals. Alice is defending a city, which has two gates: a north gate and a south gate. Alice can choose to defend just one gate, or split her forces and defend both. Meanwhile, Bob can either send his entire army to attack one gate in force, or split his forces and and send raiding groups against both gates.

This will be a zero-sum game, so we'll think of the possible outcomes as one player winning or losing "points" from the other. The possible outcomes are, in order:

- If Bob attacks an undefended gate, he captures the city. (Alice loses 5 points to Bob.)
- If Bob sends raiding groups and finds an undefended gate, that raiding group pillages supplies from the city. (Alice loses 2 points to Bob.)
- If Bob sends raiding groups against two gates, and both are partially defended, Bob's soldiers retreat with minimal losses. (No points.)
- If Bob attacks a partially defended gate, Alice holds him off and Bob suffers casualties. (Alice wins 1 point from Bob.)
- If Bob sends his entire army against Alice's, he takes heavy losses and the siege is broken. (Alice wins 5 points from Bob.)

In the form of a payoff matrix, we have:

	Bob: attack north	Bob: raid	Bob: attack south
Alice: defend north	$(5, -5)$	$(-2, 2)$	$(-5, 5)$
Alice: split forces	$(1, -1)$	$(0, 0)$	$(1, -1)$
Alice: defend south	$(-5, 5)$	$(-2, 2)$	$(5, -5)$

There is no dominant strategy in this game: for each player, both of the "all-in" strategies that focus on one gate have the highest reward, but also the highest risk. To analyze this game, we can make the following observations:

- If Alice splits her forces, she is certain not to lose points. She might even win a few points if Bob attacks in force.
- If Bob sends raiding groups, he is also certain not to lose points (and so Alice is certain not to win any points). If one of the gates is unprotected, Bob might even win some points.

This means that for Alice, splitting her forces is an optimal strategy. On the one hand, it guarantees her at least 0 points. On the other hand, Bob has a counter-strategy that prevents Alice from gaining any positive amount of points, so she can't possibly do better. For Bob, sending raiding groups is an optimal strategy for the same reason. (Either one of the generals can even announce the strategy they take in advance: it will not make any difference.)

In general, we call an outcome like "Alice splits her forces and Bob sends raiding groups" a **saddle point**. A saddle point in a zero-sum game is an outcome that's the worst outcome for Alice in its row, but the best outcome for Alice in its column. Whenever there is a saddle point, either player can guarantee an outcome at least as good as the saddle point by choosing its row or its column as a strategy.

It is a coincidence that the saddle point gives 0 points to both players in this example. We could modify the problem by having Bob give Alice one extra point in each outcome (to model the idea that Bob's army is running low on supplies, while Alice is in a well-stocked city with months of food). The same outcome would stay a saddle point, because relative comparisons between two outcomes come out the same way. However, the saddle point would now give payoffs of $(1, -1)$ to Alice and Bob.

2.3 Mixed strategies

Finally, we will consider the "odd-even game". In this game, Alice and Bob each hold up either 1 or 2 fingers. They add up the number of fingers held up; Alice wins if it is odd, and Bob wins if it is even. Additionally, Alice's choice determines the stakes: if Alice holds up 1 finger, she wins or loses \$1, nad if Alice holds up 2 fingers, she wins or loses \$2. The payoff matrix is

	Bob: 1	Bob: 2
Alice: 1	$(-1, 1)$	$(1, -1)$
Alice: 2	$(2, -2)$	$(-2, 2)$

This game has neither a dominant strategy for either player, nor a saddle point. What can we do?

If Bob knew what Alice would play, he could play the same move in response and win either \$1 or \$2. So there's no best move for Alice. Instead, one possible strategy for Alice is to flip a coin (invisibly from Bob) before making her move. If it's heads, Alice holds up 1 finger; if it's tails, Alice holds up 2 fingers. If Alice is doing this, we can compute the **expected payoff** for each strategy Bob chooses:

- If Bob holds up 1 finger, there is a $\frac{1}{2}$ chance that Alice will lose \$1 and a $\frac{1}{2}$ chance that she will win \$2. So the average amount of money Alice wins is $\frac{1}{2}(-1) + \frac{1}{2}(2) = \frac{1}{2}$: in expectation, she wins 50 cents.
- If Bob holds up 2 fingers, there is a $\frac{1}{2}$ chance that Alice will win \$1 and a $\frac{1}{2}$ chance that Alice will lose \$2. So the average amount of money Alice wins is $\frac{1}{2}(1) + \frac{1}{2}(-2) = -\frac{1}{2}$: in expectation, she loses 50 cents.

The coin-flip strategy has better worst-case behavior for Alice. If Alice is planning to hold up some number of fingers for certain, and Bob knows it, Bob can hold up the same number of fingers, and win either \$1 or \$2. But if Alice is planning to flip a coin, and Bob knows it, the best Bob can do is hold up 2 fingers, which guarantees him an average of 50 cents.

We can generalize this idea. Suppose that Alice and Bob are playing a matrix game where Alice has m choices numbered $1, 2, \dots, m$ and Bob has n choices numbered $1, 2, \dots, n$. We can record all of Alice's payoffs in an $m \times n$ matrix A, where a_{ij} is Alice's payoff when Alice chooses choice i and Bob chooses choice j. We can record Bob's payoffs in an $m \times n$ matrix B, though in the case of zero-sum games we will just have $B = -A$.

A **pure strategy** for Alice is the strategy of just picking one of the m choices she has. A **mixed strategy** involves picking between her options at random. There are many mixed strategies, each one represented by a **probability vector** $\mathbf{y} \in \mathbb{R}^m$. (The conditions for $\mathbf{y}$ to be a probability vector are that $\mathbf{y} \geq \mathbf{0}$ and that $y_1 + y_2 + \cdots + y_m = 1$.) We interpret a probability vector as saying that Alice picks choice i with probability y_i.

The calculation we did for Alice's expected payoffs (depending on Bob's actions) is generalized by the vector-matrix product $\mathbf{y}^\mathsf{T} A$. This is a row vector of length n in which the j^{th} entry is given by the sum

$$y_1 a_{1j} + y_2 A_{2j} + \cdots + y_m a_{mj}$$

which is exactly the calculation that tells us Alice's expected payoff when Alice plays the mixed strategy given by $\mathbf{y}$ and Bob picks choice j. It is the sum of the products

$$(\text{probability of outcome}) \times (\text{payoff from outcome})$$

over all outcomes in the j^{th} column.

Bob can also play a mixed strategy. Bob's mixed strategies can be described by probability vectors $\mathbf{x} \in \mathbb{R}^n$, since Bob has n options. The matrix-vector product $A\mathbf{x}$ gives us a column vector of Alice's expected payoffs when Bob plays this mixed strategy.

Finally, suppose that both players are playing mixed strategies. First, let's think about this in the odd-even game. Here, Alice's strategy is described by a probability vector $(y_1, y_2) \in \mathbb{R}^2$, and Bob's strategy is described by a probability vector $(x_1, x_2) \in \mathbb{R}^2$. Alice can get payoffs of:

- -1 with probability $y_1 x_1$;
- 1 with probability $y_1 x_2$;
- 2 with probability $y_2 x_1$;
- -2 with probability $y_2 x_2$.

Therefore Alice's expected payoff is $-y_1 x_1 + y_1 x_2 + 2y_2 x_1 - 2y_2 x_2$. This is exactly the value given by the product

$$\begin{bmatrix} y_1 & y_2 \end{bmatrix} \begin{bmatrix} -1 & 1 \\ 2 & -2 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}.$$

This generalizes: if the matrix of Alice's payoffs is the $m \times n$ matrix A, Alice's strategy is given by a probability vector $\mathbf{y} \in \mathbb{R}^m$, and Bob's strategy is given bby a probability vector $\mathbf{x} \in \mathbb{R}^n$, then Alice's expected payoff is given by $\mathbf{y}^\mathsf{T} A \mathbf{x}$.