An application in game theory

Combinatorial Optimization – Group K

Classes 19, 20, 21

Spring 2023

A sample problem

Two players, Rowan and Colin play Rock-Paper-Scissors for money. The rules are: if a player wins with throwing Rock (against Scissors), then the loser pays \$3 to the winner; if a player wins with Paper (against a Rock), then the loser pays \$2 to the winner; if a player wins with Scissors (against a Paper), then the loser pays \$1 to the winner; finally, if they throw the same, then it is a draw, there is no payoff.

The game is summarized in the following table: the amounts shown always represent the money Rowan pays Colin (and, obviously, a negative entry means that it is actually Colin who pays the absolute value of that to Rowan).

		Colin		
		Rock	Paper	Scissors
Rowan	Rock	0	2	-3
	Paper	-2	0	1
	Scissors	3	-1	0

Is there a smart way of playing this game? Of course, unpredictability is essential: Colin is easily defeated if Rowan can somehow anticipate his next throw. Therefore Colin has to incorporate some indeterminacy in his play. Consider the following two options: Colin first decides to throw all three hand-signs with an equal probability of $1/3$ (strategy #1); then he thinks that throwing Scissors is not such a good idea, so he decides on throwing that only with a chance of $1/5$ and Rock and Paper both get a chance of $2/5$ (strategy #2). Which of these two strategies is better for him?

First consider strategy number #1. For the sake of analysis, assume that the game is repeatedly played many times. Out of the games in which Rowan throws Rock, in one third of the cases Colin also throws Rock and wins nothing, in one third of the cases he throws Paper and wins \$2 and in the remaining one third he throws Scissors and loses \$3; that is, the average (or *expected value*) of his winnings if Rowan throws Rock is $1/3 \cdot 0 + 1/3 \cdot 2 + 1/3 \cdot (-3) = -1/3$. Similarly, if Rowan throws Paper or Scissors, then Colin wins $1/3 \cdot (-2) + 1/3 \cdot 0 + 1/3 \cdot 1 = -1/3$ or $1/3 \cdot 3 + 1/3 \cdot (-1) + 1/3 \cdot 0 = 2/3$ on average, respectively.

Repeating the above calculation for strategy #2, the average of Colin's winnings is $2/5 \cdot 0 + 2/5 \cdot 2 + 1/5 \cdot (-3) = 1/5$ or $2/5 \cdot (-2) + 2/5 \cdot 0 + 1/5 \cdot 1 = -3/5$ or $2/5 \cdot 3 + 2/5 \cdot (-1) + 1/5 \cdot 0 = 4/5$ if Rowan throws Rock, Paper or Scissors, respectively. The above results are summarized in the following table.

		Colin's average gain	
		Strategy #1	Strategy #2
Rowan	Rock	$-1/3$	$1/5$
	Paper	$-1/3$	$-3/5$
	Scissors	$2/3$	$4/5$

How do these results help compare the two strategies? Obviously, Colin has no information on how Rowan decides his throws, so the only sensible idea is to compare worst-case scenarios: with strategy #1 Colin can lose an average of \$$1/3$ in the worst case (if Rowan throws Rock or Paper), while in case of strategy #2 he can lose an average of \$$3/5$ (if Rowan throws Paper). Therefore, if Colin wants to minimize his worst-case losses (or, in other words, maximize his minimum gain), then he should prefer strategy #1 to strategy #2. Strategy #1 guarantees for Colin that, no matter how Rowan plays, his average loss per game is not worse than $\$^1\!/\!_3$.

Of course, the best idea for Colin would be to drop both of the above strategies and look for one that maximizes his minimum gain over all possible strategies. Before we see how Colin can do that, we put this problem into a more general context.

Two-player, zero sum games

Game Theory is a rapidly improving field of mathematics with a wide range of applications mainly in economics, but also in political science, sociology and even biology. It deals with problems regarding how the decisions of individual agents, all following their own interests affect each other. In spite of its name, it is far from being some light amusement: it served as the ground for eight Nobel Prizes in economics in the past few decades.

The problem of *two player, zero sum games* was the starting point of modern game theory. Such a game is given by an $m \times n$ matrix A called the *payoff matrix*. We refer to the two players as the Row Player and the Column Player. The game consists of the Row Player choosing a row of A and, simultaneously, the Column Player choosing a column of A. After the choices are made, the Row Player pays the amount corresponding to the element of A in the intersection of the chosen row and column to the Column Player; again, a negative element means that it is the Column Player who pays. (The term „zero sum" refers to the fact that in each outcome of the game the gain of one player equals the losses of the other.)

Obviously, the above Rock-Paper-Scissors game is an example of a two-player, zero sum game: the Row Player is Rowan, the Column Player is Colin and the payoff matrix is the 3×3 matrix shown in the first table above.

A *mixed strategy* of the Column Player is an n-dimensional *probability vector* x; that is, a column vector with nonnegative entries that add up to 1. Its i-th entry measures the probability of the Column Player choosing the i-th column. For example, Strategies #1 and #2 above correspond to the mixed strategies $x_1 = (1/3, 1/3, 1/3)^\top$ and $x_2 = (2/5, 2/5, 1/5)^\top$, respectively for the Column Player in the Rock-Paper-Scissors game. Reviewing the above calculations on these two strategies, it is easy to see that the j-th entry of the vector $A \cdot x$ represents the *expected value* (that is, average) of the Column Player's winnings if the Row Player chooses the j-th row. (The columns of the second table above contain the vectors Ax_1 and Ax_2.)

As it was the case with the Rock-Paper-Scissors game, the Column Player is interested in maximizing his worst-case gain, that is, the minimum element of Ax over all probability vectors x. In other words, if $\min(v)$ denotes the minimum element of a vector v, then the Column Player wants to find

$$\max\left\{\min(Ax)\colon x \text{ is a probability vector}\right\}.$$

Analogously, if a mixed strategy of the Row Player is stored in the row vector y then the elements of yA measure the expected losses of the Row Player (assuming that the Column Player chooses the corresponding column). Since the Row Player is interested in minimizing his worst-case losses, the problem for him is to find

$$\min\left\{\max(yA)\colon y \text{ is a probability vector}\right\}.$$

The solution of two-player, zero-sum games

Consider the Rock-Paper-Scissors game again. As seen above, Colin's best mixed strategy is a column vector $x = (x_1, x_2, x_3)^\top$ for which $x_1, x_2, x_3 \geq 0$ and $x_1 + x_2 + x_3 = 1$ hold and $\min(Ax)$ is maximum possible (where A is the payoff matrix of the game). To make this into a linear program, we introduce a fourth variable μ that stands for the minimum element of Ax (that is, Colin's worst-case average gain with respect to x). For any given x the value of μ is the maximum of the lower bounds on all elements of Ax; that is, the maximum possible value of μ for which the inequalities $2x_2 - 3x_3 \geq \mu$, $-2x_1 + x_3 \geq \mu$ and $3x_1 - x_2 \geq \mu$ hold. All in all, finding Colin's best mixed strategy is equivalent to the following linear program.

$$\begin{aligned}
&\max \mu \\
&\text{subject to} \\
&2x_2 - 3x_3 - \mu \geq 0 \\
&-2x_1 + x_3 - \mu \geq 0 \\
&3x_1 - x_2 - \mu \geq 0 \\
&x_1 + x_2 + x_3 = 1 \\
&x_1 \geq 0, x_2 \geq 0, x_3 \geq 0
\end{aligned}$$

Of course, this linear program is easy to solve (with any LP solver), the optimum solution is: $x_1 = 1/6$, $x_2 = 1/2$, $x_3 = 1/3$ and $\mu = 0$. This means that the best idea for Colin is to roll a dice and throw Rock if the result is 6, throw Paper if the result is 5, 4 or 3 and throw Scissors if the result is 2 or 1. This way he can guarantee that, no matter how Rowan plays, his average gain is not worse than zero – and that is the best he can do. Evidently, since this game is symmetric for the two players, the analogous statement holds for Rowan. Of course, symmetry of the game suggests even by common sense that none of the two players can win money from the other on average, however, it is not at all obvious how to find the optimum mixed strategy that saves a player from a possible average loss.

Obviously, the above solution method works for any two-player, zero-sum game. Finding the Column Player's best mixed strategy amounts to solving the linear program

$$\max\{\mu\colon Ax - \mu \cdot \mathbf{1} \geq 0, \mathbf{1}x = 1, x \geq 0\}, \tag{$*$}$$

where A is the payoff matrix of the game and $\mathbf{1}$ denotes an all-1 vector (both as an m-dimensional column vector and an n-dimensional row vector).

Neumann's Minimax Theorem

The above shows that the problem of two-player, zero-sum games is basically a subfield of linear programming. Therefore it is just a natural idea to apply the duality theorem on $(*)$ and see what it gives.

It will be convenient to apply the second form of duality that assumes the nonnegativity of all variables. The slight problem with that is that μ is not assumed to be nonnegative in $(*)$. To circumvent this, we replace μ with $\mu_1 - \mu_2$ in $(*)$ and thus the nonnegativity of the variables μ_1 and μ_2 can safely be assumed (since any real number can be written as the difference of two nonnegative numbers). Using this trick, the matrix form of $(*)$ becomes $\max\left\{\bar{c}\bar{x}\colon \bar{A}\bar{x} \leq \bar{b}, \bar{x} \geq 0\right\}$, where

$$\bar{A} = \left[\begin{array}{ccccc|c|c}
 & & & & & 1 & -1 \\
 & & & & & 1 & -1 \\
 & & -A & & & \vdots & \vdots \\
 & & & & & 1 & -1 \\
 & & & & & 1 & -1 \\
\hline
1 & 1 & \dots & 1 & 1 & 0 & 0 \\
\hline
-1 & -1 & \dots & -1 & -1 & 0 & 0
\end{array}\right], \quad
\left[\begin{array}{c}
0 \\ 0 \\ \vdots \\ 0 \\ 0 \\ \hline 1 \\ \hline -1
\end{array}\right] = \bar{b},$$

$$\bar{c} = \left[\begin{array}{ccccc|c|c}
0 & 0 & \dots & 0 & 0 & 1 & -1
\end{array}\right]$$

and $\bar{x}$ is obtained from x by adding the extra entries μ_1 and μ_2 as the last two coordinates. Then the dual of this program is $\min\left\{\bar{y}\bar{b}\colon \bar{y}\bar{A} \geq \bar{c}, \bar{y} \geq 0\right\}$. Let the last two coordinates of $\bar{y}$ be ν_1 and ν_2, and denote the rest of the vector $\bar{y}$ by y. Then the dual program becomes

$$\min\left\{\nu_1 - \nu_2\colon y(-A) + (\nu_1 - \nu_2) \cdot \mathbf{1} \geq 0, y\mathbf{1} = 1, y \geq 0\right\},$$

or, in an equivalent form (denoting $\nu_1 - \nu_2$ by ν)

$$\min\left\{\nu\colon yA - \nu\mathbf{1} \leq 0, y\mathbf{1} = 1, y \geq 0\right\}. \tag{$**$}$$

It leaps to the eye that (∗∗) is very similar to (∗). The system $yA - \nu\mathbf{1} \leq 0$ translates to saying that ν is an upper bound on every element of yA. Therefore (∗∗) amounts to looking for a probability vector y such that $\max(yA)$ is minimum possible; in other words, it is equivalent to the Row Player's problem. Since the maximum of (∗) is equal to the minimum of (∗∗) by the duality theorem, we get the following result.

Theorem (John von Neumann (János Neumann), 1928).

For every two-player, zero-sum game the maximum of the minimum expected gain of the Column Player is equal to the minimum of the maximum expected losses of the Row Player. In other words, there exist the optimum mixed strategies x and y for the Column Player and the Row Player, respectively and a common value μ (called the value *of the game) such that*

(1) no matter how the Row Player plays, x guarantees an expected gain of at least μ to the Column Player and

(2) no matter how the Column Player plays, y guarantees an expected loss of at most μ to the Row Player.

The above theorem is often referred to as the *Minimax Theorem.* The original proof of Neumann did not rely on the duality theorem of linear programming (since it did not exist then). Neumann is often quoted as saying: „As far as I can see, there could be no theory of games ... without that theorem... I felt that there was nothing worth publishing until the Minimax Theorem was proved".