

Game Theory and Linear Programming

Math 20

December 1, 2005

Goals

- Convert a game theory problem into a linear programming problem.
- Use the simplex method to solve game theory problems

Notes

This is based on the November 28 and November 30 classes. See the references at the end for where I got some of these ideas.

1 Review of Game Theory so far

Remember that a standard zero-sum game involves two players (called R and C) and a payoff matrix A . If R makes choice i and C makes choice j , then C pays R a sum of a_{ij} dollars (this number could be negative, representing a payoff of R to C). What should each player choose?

We learned that if A has a **saddle point**, that is, an entry which is the smallest in its row and the largest in its column, represents a sort of equilibrium strategy. Such a game is called **strictly determined**. R and C should make this choice no matter how many times the game is played.

If there is no saddle point to the payoff matrix, the strategies can be interpreted probabilistically. That is, each player chooses a probability vector indicating the percentage of the time they will choose each strategy. This can mean the amount of resources spent in pursuing each strategy, or the frequency of the time they will choose each strategy over a number of repeated iterations of the game. R 's strategy is often represented by a row vector \mathbf{p} and C 's strategy by a column vector \mathbf{q} . Then the expected payoff is

$$E(\mathbf{p}, \mathbf{q}) = \mathbf{p}A\mathbf{q}.$$

In the case of 2×2 payoff matrices, there is a way to find the best strategy for R and C by explicitly finding the expected value and maximizing it. It turns out that R 's

optimal strategy is

$$\mathbf{p}^* = [p_1^* \quad p_2^*] = \left[\frac{a_{22} - a_{21}}{a_{11} + a_{22} - a_{12} - a_{21}} \quad \frac{a_{11} - a_{12}}{a_{11} + a_{22} - a_{12} - a_{21}} \right].$$

Likewise, C 's optimal strategy is

$$\mathbf{q}^* = \begin{bmatrix} q_1^* \\ q_2^* \end{bmatrix} = \begin{bmatrix} \frac{a_{22} - a_{12}}{a_{11} + a_{22} - a_{12} - a_{21}} \\ \frac{a_{11} - a_{21}}{a_{11} + a_{22} - a_{12} - a_{21}} \end{bmatrix}$$

The value of the game (the average payoff per game played if each play according to their optimal strategy) is

$$v = \frac{a_{11}a_{22} - a_{21}a_{12}}{a_{11} + a_{22} - a_{12} - a_{21}}.$$

But what do we do in the case of larger games which are not strictly determined?

2 Restatement of LP

A **standard linear programming problem** is to maximize the quantity

$$c_1x_1 + c_2x_2 + \dots + c_nx_n = \mathbf{c}^\top \mathbf{x}$$

subject to constraints

$$\begin{aligned} a_{11}x_1 + a_{12}x_2 + \dots + a_{1n}x_n &\leq b_1 \\ a_{21}x_1 + a_{22}x_2 + \dots + a_{2n}x_n &\leq b_2 \\ &\dots \\ a_{m1}x_1 + a_{m2}x_2 + \dots + a_{mn}x_n &\leq b_m \end{aligned}$$

or

$$A\mathbf{x} \leq \mathbf{b}.$$

We usually include the nonnegativity constraint $\mathbf{x} \geq \mathbf{0}$.

We learned the simplex method to solve problems of this form. We would like to use LP to solve Game Theory problems.

3 Restatement of GT problems as LP problems

This derivation is not something that needs to be memorized, but should be understood at least once.

Let's think about the problem from the column player's perspective. If she chooses strategy \mathbf{q} , and R knew it, he would choose \mathbf{p} to maximize the payoff $\mathbf{p}A\mathbf{q}$. Thus the

column player wants to minimize that quantity. That is, C 's objective is realized when the payoff is

$$E = \min_{\mathbf{q}} \max_{\mathbf{p}} \mathbf{p}A\mathbf{q}.$$

Here $\max_{\mathbf{p}}$ means to take the maximum over all of R 's strategies, that is, over all row probability vectors $\mathbf{p} \in \mathbb{R}^m$ where $\mathbf{p} \geq \mathbf{0}$ and $\sum \mathbf{p}_i = 1$.

This seems hard. We have to solve a maximization problem for every \mathbf{q} , then find the one of these which minimizes. Luckily, linearity saves us.

3.1 From the continuous to the discrete

Lemma. *Regardless of \mathbf{q} , we have*

$$\max_{\mathbf{p}} \mathbf{p}A\mathbf{q} = \max_{1 \leq i \leq m} \mathbf{e}_i^T A\mathbf{q}$$

Here \mathbf{e}_i^T is the probability vector represents the pure strategy of going only with choice i .

Proof. Since each pure strategy vector is a probability vector, we must have

$$\max_{\mathbf{p}} \mathbf{p}A\mathbf{q} \geq \max_{1 \leq i \leq m} \mathbf{e}_i^T A\mathbf{q}$$

(the maximum over a larger set must be at least as big). On the other hand, let \mathbf{q} be C 's strategy. Let the quantity on the right be maximized when $i = i_0$. Let \mathbf{p} be any strategy for R . Notice that $\mathbf{p} = \sum_i p_i \mathbf{e}_i^T$. So

$$\begin{aligned} E(\mathbf{p}, \mathbf{q}) &= \mathbf{p}A\mathbf{q} = \sum_{i=1}^m p_i \mathbf{e}_i^T A\mathbf{q} \\ &\leq \sum_{i=1}^m p_i \mathbf{e}_{i_0}^T A\mathbf{q} \\ &= \left(\sum_{i=1}^m p_i \right) \mathbf{e}_{i_0}^T A\mathbf{q} \\ &= \mathbf{e}_{i_0}^T A\mathbf{q}. \end{aligned}$$

Thus

$$\max_{\mathbf{p}} \mathbf{p}A\mathbf{q} \leq \mathbf{e}_{i_0}^T A\mathbf{q}.$$



The next step is to introduce a new variable v representing the value of this inner maximization. Our objective is to minimize it. Saying it's the maximum of all payoffs from pure strategies is the same as saying

$$v \geq \mathbf{e}_i^T A\mathbf{q}$$

for all i . So we finally have something that looks like an LP problem! We want to choose \mathbf{q} and v which maximize v subject to the constraints

$$\begin{aligned} v &\geq \mathbf{e}_i^T \mathbf{A} \mathbf{q} & i = 1, 2, \dots, m \\ q_j &\geq 0 & j = 1, 2, \dots, n \\ \sum_{j=1}^n q_j &= 1 \end{aligned}$$

3.2 Standardization

There are a couple of problems with this problem. First, we don't know how to handle equalities with the simplex method. I suppose we could introduce two inequalities $\sum_{j=1}^n q_j \leq 1$ and $\sum_{j=1}^n q_j \geq 1$. But that seems inefficient. Secondly, it's not yet in standard form, which is the only problem we know how to solve with the simplex method.

We'll solve both of these in one fell swoop by making a transformation. First, the strategies are not changed if we add a constant to every value of A . Such an addition would be the equivalent of two games: "Give me \$5, then we'll play this game..." The column player will just be trying to minimize something which is at least 5. So temporarily at least, we can assume that all $a_{ij} \geq 0$. This means $v \geq 0$, too, and we can scale the individual entries by it. That is, let

$$x_j = \frac{q_j}{v}.$$

Since we know $v > 0$, we still have $\mathbf{x} \geq \mathbf{0}$. Now

$$\sum_{j=1}^n x_j = \frac{1}{v} \sum_{j=1}^n q_j = \frac{1}{v}.$$

So our problem is now to choose $\mathbf{x} \geq \mathbf{0}$ which maximizes $\sum_j x_j$. The constraints now take the form

$$v \geq \mathbf{e}_i^T \mathbf{A} \mathbf{q} \iff 1 \geq \mathbf{e}_i^T \mathbf{A} \mathbf{x},$$

for all i . Another way to write this is

$$\mathbf{A} \mathbf{x} \leq \mathbf{1},$$

where $\mathbf{1}$ is the vector consisting of all ones. Finally, we have transformed the game theory problem into an LP problem in standard form, that we know how to solve with the simplex method.

Theorem. Consider a game with payoff matrix A , where each entry of A is positive. The column player's optimal strategy \mathbf{q} is $\frac{\mathbf{x}}{x_1 + \dots + x_n}$, where $\mathbf{x} \geq \mathbf{0}$ satisfies the LP problem of maximizing $x_1 + \dots + x_n$ subject to the constraints $\mathbf{A} \mathbf{x} \leq \mathbf{1}$.

If our payoff matrix doesn't satisfy the nonnegativity constraint, we can just add a constant.

4 Example

Let's play rock-paper-scissors. The payoff matrix is

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

We can add 2 to everything to make

$$\tilde{A} = \begin{bmatrix} 2 & 1 & 3 \\ 3 & 2 & 1 \\ 1 & 3 & 2 \end{bmatrix}.$$

The problem is to maximize $x_1 + x_2 + x_3$ subject to the constraints

$$\begin{aligned} 2x_1 + x_2 + 3x_3 &\leq 1 \\ 3x_1 + 2x_2 + x_3 &\leq 1 \\ x_1 + 3x_3 + 2x_3 &\leq 1. \end{aligned}$$

We introduce slack variables $y_1, y_2,$ and $y_3,$ so the constraints now become

$$\begin{aligned} 2x_1 + x_2 + 3x_3 + y_1 &= 1 \\ 3x_1 + 2x_2 + x_3 + y_2 &= 1 \\ x_1 + 3x_3 + 2x_3 + y_3 &= 1. \end{aligned}$$

An easy initial basic solution is to let $\mathbf{x} = \mathbf{0}$ and $\mathbf{y} = \mathbf{1}$.¹ The initial tableau is therefore

	x_1	x_2	x_3	y_1	y_2	y_3	z	value
y_1	2	1	3	1	0	0	0	1
y_2	3	2	1	0	1	0	0	1
y_3	1	3	2	0	0	1	0	1
z	-1	-1	-1	0	0	0	1	0

Which should be the entering variable? The coefficients in the bottom row are all the same, so let's just pick one, x_1 . To find the departing variable, we look at the ratios $\frac{1}{2}, \frac{1}{3},$ and $\frac{1}{1}$. So y_2 is the departing variable.

We scale row 2 by $\frac{1}{3}$:

	x_1	x_2	x_3	y_1	y_2	y_3	z	value
y_1	2	1	3	1	0	0	0	1
y_2	1	$\frac{2}{3}$	$\frac{1}{3}$	0	$\frac{1}{3}$	0	0	$\frac{1}{3}$
y_3	1	3	2	0	0	1	0	1
z	-1	-1	-1	0	0	0	1	0

¹T

Then we use row operations to zero out the rest of column one:

	x_1	x_2	x_3	y_1	y_2	y_3	z	value
y_1	0	$-\frac{1}{3}$	$\frac{7}{3}$	1	$-\frac{2}{3}$	0	0	$\frac{1}{3}$
x_1	1	$\frac{2}{3}$	$\frac{1}{3}$	0	$\frac{1}{3}$	0	0	$\frac{1}{3}$
y_3	0	$\frac{7}{3}$	$\frac{5}{3}$	0	$-\frac{1}{3}$	1	0	$\frac{2}{3}$
z	0	$-\frac{1}{3}$	$-\frac{2}{3}$	0	$\frac{1}{3}$	0	1	$\frac{1}{3}$

We can still improve this: x_3 is the entering variable and y_1 is the departing variable. The new tableau is

	x_1	x_2	x_3	y_1	y_2	y_3	z	value
x_3	0	$-\frac{1}{7}$	1	$\frac{3}{7}$	$-\frac{2}{7}$	0	0	$\frac{1}{7}$
x_1	1	$\frac{5}{7}$	0	$-\frac{1}{7}$	$\frac{5}{7}$	0	0	$\frac{5}{7}$
y_3	0	$\frac{18}{7}$	0	$-\frac{5}{7}$	$\frac{7}{7}$	1	0	$\frac{3}{7}$
z	0	$-\frac{2}{7}$	0	$\frac{2}{7}$	$\frac{1}{7}$	0	1	$\frac{3}{7}$

Finally, entering x_2 and departing y_3 gives

	x_1	x_2	x_3	y_1	y_2	y_3	z	value
x_3	0	0	1	$\frac{7}{18}$	$-\frac{5}{18}$	$\frac{1}{18}$	0	$\frac{1}{6}$
x_1	1	0	0	$\frac{1}{18}$	$\frac{7}{18}$	$-\frac{5}{18}$	0	$\frac{1}{6}$
x_2	0	1	0	$-\frac{5}{18}$	$\frac{1}{18}$	$\frac{7}{18}$	0	$\frac{1}{6}$
z	0	0	0	$\frac{1}{6}$	$\frac{1}{6}$	$\frac{1}{6}$	1	$\frac{1}{2}$

So the x variables have values $x_1 = \frac{1}{6}$, $x_2 = \frac{1}{6}$, $x_3 = \frac{1}{6}$. Furthermore $z = x_1 + x_2 + x_3 = \frac{1}{2}$, so $v = \frac{1}{z} = 2$. This also means that $p_1 = \frac{1}{3}$, $p_2 = \frac{1}{3}$, and $p_3 = \frac{1}{3}$. So the optimal strategy is to do each thing the same number of times.

Let's remember what it is we were trying to solve: This was a game with payoff matrix \tilde{A} . The same strategy applied to the original payoff matrix A results in a payoff of zero. Thus iterations of rock, paper, scissors should result in a payoff of zero. This means the game is fair.

5 Review of the dual problem of LP

Each problem in standard form has a dual problem. We'll explain it again, by example. In the above, suppose the decider is the baker with n products and m ingredients. Each a_{ij} represents the amount of ingredient i needed to make product j . Each b_i is the total amount of ingredient i on hand. Each c_j represents the profit for selling one product j . The decision variable x_j represents the amount of each product the baker will make.

The factory (who is trying to dissuade the baker from selling his products) wants to offer the baker a price y_i for each ingredient i . Obviously the factory wants to minimize the amount paid out, which is $y_1 b_1 + \dots + y_m b_m$. If the baker chooses bundle \mathbf{x} of products, he uses $A\mathbf{x}$ for ingredients (the entry in position i of this vector is the amount of ingredient i he will use). If he were to sell these ingredients to the factory instead of making them into products, he would make $\mathbf{y}^\top A\mathbf{x}$ rather than $\mathbf{c}^\top \mathbf{x}$. The only way this could happen for all bundles \mathbf{x} is if

$$\mathbf{y}^\top A \geq \mathbf{c}^\top,$$

or (taking transposes)

$$A^T \mathbf{y} \geq \mathbf{c}.$$

Thus the dual LP problem to

$$\text{“maximize } \mathbf{c}^T \mathbf{x} \text{ subject to } A \mathbf{x} \leq \mathbf{b} \text{ and } \mathbf{x} \geq \mathbf{0}\text{”}$$

is

$$\text{“minimize } \mathbf{b}^T \mathbf{y} \text{ subject to } A^T \mathbf{y} \geq \mathbf{c} \text{ and } \mathbf{y} \geq \mathbf{0}\text{”}.$$

The original problem is sometimes called the primal problem. The value y_i is sometimes called the **shadow price** of resource i . It might be interpreted as the value to the baker of this ingredient.

We solved in small cases the primal problem and the dual problem separately. The big deal is:

Theorem. *If either the primal problem or the dual problem has an optimal solution with finite objective value, then the other problem also has an optimal solution. Moreover, the objective values of the two problems are the same.*

This means that if \mathbf{x} is the solution to the primal problem and \mathbf{y} is the solution to the dual problem, then $\mathbf{c}^T \mathbf{x} = \mathbf{b}^T \mathbf{y}$.

Also, in the simplex method, the solution to the dual problem is computed for free. This is basically because the original slack variables are the decision variables for the dual problem. The values of these variables for the dual problem appear in the bottom row of the final tableau, under the names of those variables.

6 The row player’s problem

Now let’s think about the problem from the column player’s perspective.² If he chooses strategy \mathbf{p} , and C knew it, he would choose \mathbf{q} to minimize the payoff $\mathbf{p}A\mathbf{q}$. Thus the row player wants to maximize that quantity. That is, R ’s objective is realized when the payoff is

$$E = \max_{\mathbf{p}} \min_{\mathbf{q}} \mathbf{p}A\mathbf{q}.$$

Here $\max_{\mathbf{p}}$ means to take the maximum over all of R ’s strategies, that is, over all row probability vectors $\mathbf{p} \in \mathbb{R}^m$ where $\mathbf{p} \geq \mathbf{0}$ and $\sum \mathbf{p}_i = 1$.

Linearity saves us again: the inner minimization is the same as if we had only considered the pure strategies.

Lemma. *Regardless of \mathbf{p} , we have*

$$\min_{\mathbf{q}} \mathbf{p}A\mathbf{q} = \min_{1 \leq j \leq n} \mathbf{p}A\mathbf{e}_j$$

Here \mathbf{e}_j is the probability vector represents the pure strategy of going only with choice j .

²Some of this section is a cut-and-paste, search-and-replace of Section 3.

Proof. Exercise. We did a very similar thing in Section 3. ■

The next step is to introduce a new variable v' representing the value of this inner minimization. Our objective is to maximize it. Saying it's the minimum of all payoffs from pure strategies is the same as saying

$$v' \leq \mathbf{p}A\mathbf{e}_j$$

for all j . Again, we have something that looks like an LP problem! We want to choose \mathbf{p} and v' which maximize v' subject to the constraints

$$\begin{aligned} v' &\leq \mathbf{p}A\mathbf{e}_j & j = 1, 2, \dots, n \\ p_i &\geq 0 & i = 1, 2, \dots, m \\ \sum_{i=1}^m p_i &= 1 \end{aligned}$$

6.1 Standardization

As before, we can standardize this by renaming

$$\mathbf{y} = \frac{1}{v'}\mathbf{p}^\top$$

(this makes \mathbf{y} a column vector). Then

$$\sum_{i=1}^m y_i = \frac{1}{v'}$$

So maximizing v' is the same as minimizing $\mathbf{1}^\top \mathbf{y}$. Likewise, the equations of constraint become $v' \leq (v'\mathbf{y}^\top)A\mathbf{e}_j$ for all j , or $\mathbf{y}^\top A \geq \mathbf{1}$, or (taking transposes) $A^\top \mathbf{y} \geq \mathbf{1}$. If all the entries of A are positive, we may assume that v' is positive, so the constraints $\mathbf{p} \geq \mathbf{0}$ are satisfied if and only if $\mathbf{y} \geq \mathbf{0}$.

Thus:

Theorem. Consider a game with payoff matrix A , where each entry of A is positive. The row player's optimal strategy \mathbf{p} is $\frac{\mathbf{y}^\top}{y_1 + \dots + y_n}$, where $\mathbf{y} \geq \mathbf{0}$ satisfies the LP problem of minimizing $y_1 + \dots + y_n = \mathbf{1}^\top \mathbf{y}$ subject to the constraints $A^\top \mathbf{y} \geq \mathbf{1}$.

The big observation is this:

Theorem. The row player's LP problem is the dual of the column player's LP problem.

The fact that the initial problem and its dual have the same solution proves for us the fundamental theorem of zero-sum games.

7 Examples Again

Go back to rock-paper-scissors. The final tableau turned out to be:

	x_1	x_2	x_3	y_1	y_2	y_3	z	value
x_3	0	0	1	$\frac{7}{18}$	$-\frac{5}{18}$	$\frac{1}{18}$	0	$\frac{1}{6}$
x_1	1	0	0	$\frac{1}{18}$	$\frac{7}{18}$	$-\frac{5}{18}$	0	$\frac{1}{6}$
x_2	0	1	0	$-\frac{5}{18}$	$\frac{1}{18}$	$\frac{7}{18}$	0	$\frac{1}{6}$
z	0	0	0	$\frac{1}{6}$	$\frac{1}{6}$	$\frac{1}{6}$	1	$\frac{1}{2}$

Looking at the bottom row, the shadow prices (solution to the dual problem) are

$$y_1 = y_2 = y_3 = \frac{1}{6}.$$

This means $v' = \frac{1}{3/6} = 2$ and

$$p_1 = p_2 = p_3 = \frac{1}{3}.$$

Not surprisingly, the row player's strategy is the same. There is a sort of symmetry to the problem (see the problem set).

Consider a new game: players R and C each choose a number 1, 2, or 3. If they choose the same thing, C pays R that amount. If they choose differently, R pays C the amount that C has chosen. What should each do?

The payoff matrix is

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

We can alter this by adding 4 to everything.

$$\tilde{A} = \begin{bmatrix} 5 & 2 & 1 \\ 3 & 6 & 1 \\ 3 & 2 & 7 \end{bmatrix}$$

So the primal problem (which will result in the column player's strategy) is choose \mathbf{x} maximizing $\mathbf{1}^T \mathbf{x}$ subject to $\mathbf{x} \geq \mathbf{0}$ and $\tilde{A}\mathbf{x} \leq \mathbf{1}$.

We introduce slack variables \mathbf{y} and the constraints become $\tilde{A}\mathbf{x} + \mathbf{y} = \mathbf{1}$. The initial basic solution is $\mathbf{x} = \mathbf{0}$ and $\mathbf{y} = \mathbf{1}$. Thus our initial tableau is

	x_1	x_2	x_3	y_1	y_2	y_3	z	value
y_1	5	2	1	1	0	0	0	1
y_2	3	6	1	0	1	0	0	1
y_3	3	2	7	0	0	1	0	1
z	-1	-1	-1	0	0	0	1	0

Looking at the bottom row; we can choose any of the decision variables to be the entering one. Let's pick x_1 . Of the θ -ratios (quotient of the entry in the last column

and the entry in the column of the new entering variable), $\frac{1}{5}$ is the smallest. So y_1 is the departing variable. The new tableau is:

	x_1	x_2	x_3	y_1	y_2	y_3	z	value
x_1	1	$\frac{2}{5}$	$\frac{1}{5}$	$\frac{1}{5}$	0	0	0	$\frac{1}{5}$
y_2	0	$\frac{24}{5}$	$\frac{2}{5}$	$-\frac{3}{5}$	1	0	0	$\frac{3}{5}$
y_3	0	$\frac{4}{5}$	$\frac{32}{5}$	$-\frac{3}{5}$	0	1	0	$\frac{3}{5}$
z	0	$-\frac{3}{5}$	$-\frac{4}{5}$	$\frac{1}{5}$	0	0	1	$\frac{1}{5}$

Now x_3 is the entering variable and y_3 is the departing variable. The next tableau is:

	x_1	x_2	x_3	y_1	y_2	y_3	z	value
x_1	1	$\frac{3}{8}$	0	$\frac{7}{32}$	0	$-\frac{1}{32}$	0	$\frac{3}{16}$
y_2	0	$\frac{19}{4}$	0	$-\frac{9}{16}$	1	$-\frac{1}{16}$	0	$\frac{3}{8}$
x_3	0	$\frac{1}{8}$	1	$-\frac{3}{32}$	0	$\frac{5}{32}$	0	$\frac{1}{16}$
z	0	$-\frac{1}{2}$	0	$\frac{1}{8}$	0	$\frac{1}{8}$	1	$\frac{1}{4}$

Now x_2 is the entering variable and y_2 is the departing one. The next tableau is:

	x_1	x_2	x_3	y_1	y_2	y_3	z	value
x_1	1	0	0	$\frac{5}{19}$	$-\frac{3}{38}$	$-\frac{1}{38}$	0	$\frac{3}{19}$
x_2	0	1	0	$-\frac{9}{76}$	$\frac{4}{19}$	$-\frac{1}{76}$	0	$\frac{3}{38}$
x_3	0	0	1	$-\frac{3}{38}$	$-\frac{1}{38}$	$\frac{3}{19}$	0	$\frac{1}{19}$
z	0	0	0	$\frac{5}{76}$	$\frac{2}{19}$	$\frac{9}{76}$	1	$\frac{11}{38}$

We're done. We have

$$x_1 = \frac{3}{19}$$

$$x_2 = \frac{3}{38}$$

$$x_3 = \frac{1}{19}$$

$$z = \frac{11}{38}$$

Thus the column player's strategy is

$$q_1 = \frac{x_1}{z} = \frac{6}{11}$$

$$q_2 = \frac{x_2}{z} = \frac{3}{11}$$

$$q_3 = \frac{x_3}{z} = \frac{2}{11}$$

So the column player should pick one $\frac{6}{11}$ of the time, pick two $\frac{3}{11}$ of the time, and pick

three the rest of the time. The row player's strategy is also evident. We have

$$p_1 = \frac{5/76}{11/38} = \frac{5}{22}$$
$$p_2 = \frac{2/19}{11/38} = \frac{4}{11}$$
$$p_3 = \frac{9/76}{11/38} = \frac{9}{22}$$

The value v of the game is

$$\frac{1}{z} = \frac{38}{11} \approx 3.45455.$$

This is less than the 4 we added to everything. This means the value of the original game is negative, so the C player probably shouldn't be playing.

8 Extensions

There is a lot more to game theory than this. Consider

- Games with multiple players: You can treat this by pitting one player against an omniscient amalgam of the other players, who can choose each opponent's strategies to best compete against the one player. This would give each player's optimal strategy in turn.
- Games with uncertainty: In this case the strategies take into account a choice depending on an outcome. The payoff gets replaced by the expected value of the payoff. See the problem set for more about that.
- Non-zero sum games: in this case there might be different payoffs for each player which don't result in the other losing the same amount. Consider two people arrested as accomplices in a crime. If both deny their involvement, they might both be let go. But both realize that they can get a lighter sentence by confessing and implicating the accomplice. If you watch enough cop shows on TV, you know that usually both suspects confess, which is certainly not the "best" outcome for them.

9 References

- B. Kolman and D. Hill. *Linear Algebra: An Applied First Course* (Chapter 11).
- A.M. Glicksman. *An Introduction to Linear Programming and the Theory of Games*.
- R. D. Luce and H. Raiffa. *Games and Decisions: Introduction and Critical Survey*. Raiffa is emeritus of the Harvard Business School and one of the giants in decision theory.

- H.W. Kuhn. *Lectures on the Theory of Games*. I'm guessing the Kuhn is "the" Kuhn of the Kuhn-Tucker conditions, but I don't know.