ORF 307: Lecture 12

Linear Programming: Chapter 11: Game Theory

Robert J. Vanderbei

April 2, 2019

Slides last edited on April 1, 2019

Game Theory

John Nash = A Beautiful Mind

Rock-Paper-Scissors

A two person game.

Rules.
At the count of three declare one of:

Rock Paper Scissors

Winner Selection. Identical selection is a draw. Otherwise:

- Rock dulls Scissors
- Paper covers Rock
- Scissors cuts Paper

Check out Sam Kass' version: Rock, Paper, Scissors, Lizard, Spock

It was featured on The Big Bang Theory.

Payoff Matrix

Payoffs are *from* row player *to* column player:

$$A = \begin{array}{c} \\ R \\ P \\ S \end{array} \begin{array}{c} \begin{array}{ccc} R & P & S \end{array} \\ \begin{bmatrix} 0 & 1 & -1 \\ -1 & 0 & 1 \\ 1 & -1 & 0 \end{bmatrix} \end{array}$$

Note: Any *deterministic* strategy employed by either player can be defeated systematically by the other player.

Two-Person Zero-Sum Games

Given: $m \times n$ matrix A.

- *Row player* selects a *strategy* $i \in \{1, ..., m\}$.
- *Column player* selects a *strategy* $j \in \{1, ..., n\}$.
- Row player pays column player a_{ij} dollars.

Note: The rows of A represent deterministic strategies for row player, while columns of A represent deterministic strategies for column player.

Deterministic strategies can be (and usually are) bad.

Randomized Strategies.

- Suppose row player picks i with probability y_i.
- Suppose column player picks j with probability x_j.

Throughout, $x = \begin{bmatrix} x_1 & x_2 & \cdots & x_n \end{bmatrix}^T$ and $y = \begin{bmatrix} y_1 & y_2 & \cdots & y_m \end{bmatrix}^T$ will denote *stochastic vectors*:

$$\begin{aligned} x_j &\geq 0, \quad j = 1, 2, \ldots, n \\ \sum_j x_j &= 1 \end{aligned} \qquad\qquad \begin{aligned} y_i &\geq 0, \quad i = 1, 2, \ldots, m \\ \sum_i y_i &= 1 \end{aligned}$$

If row player uses random strategy y and column player uses x, then *expected payoff* from row player to column player is

$$\sum_i \sum_j y_i a_{ij} x_j = y^T A x$$

Column Player's Analysis

Suppose column player were to adopt strategy x.

Then, row player's best defense is to use strategy y that minimizes $y^T A x$:

$$\min_y y^T A x$$

And so column player should choose that x which maximizes these possibilities:

$$\max_x \min_y y^T A x$$

Quiz

What's the solution to this problem:

$$\begin{array}{llllllllllll} \text{minimize} & 3y_1 & + & 6y_2 & + & 2y_3 & + & 18y_4 & + & 7y_5 & & \\ \text{subject to:} & y_1 & + & y_2 & + & y_3 & + & y_4 & + & y_5 & = & 1 \\ & & & & & & & & & y_i & \geq & 0, \qquad i = 1, 2, 3, 4, 5 \end{array}$$

Solving Max-Min Problems as LPs

Inner optimization is easy:

$$\min_y y^T Ax = \min_i e_i^T Ax$$

(e_i denotes the vector that's all zeros except for a one in the i-th position—that is, deterministic strategy i).

Note: Reduced a minimization over a *continuum* to one over a *finite set*.

We have:

$$\max \left(\min_i \; e_i^T Ax\right)$$

$$\begin{aligned} \sum_j x_j &= 1, \\ x_j &\geq 0, \qquad j = 1, 2, \ldots, n. \end{aligned}$$

Reduction to a Linear Programming Problem

Introduce a scalar variable v representing the value of the inner minimization:

$$
\begin{aligned}
\max \;\; v & \\
v \;&\leq\; e_i^T A x, \qquad && i = 1, 2, \dots, m, \\
\sum_j x_j \;&=\; 1, && \\
x_j \;&\geq\; 0, && j = 1, 2, \dots, n.
\end{aligned}
$$

Writing in pure matrix-vector notation:

$$
\begin{aligned}
\max \;\; v & \\
ve - Ax \;&\leq\; 0 \\
e^T x \;&=\; 1 \\
x \;&\geq\; 0
\end{aligned}
$$

(e without a subscript denotes the vector of all ones).

Finally, in Block Matrix Form

$$\max \begin{bmatrix} 0 \\ 1 \end{bmatrix}^T \begin{bmatrix} x \\ v \end{bmatrix}$$

$$\begin{bmatrix} -A & e \\ e^T & 0 \end{bmatrix} \begin{bmatrix} x \\ v \end{bmatrix} \begin{matrix} \leq \\ = \end{matrix} \begin{bmatrix} 0 \\ 1 \end{bmatrix}$$

$$x \geq 0$$

$$v \quad \text{free}$$

Row Player's Perspective

Similarly, row player seeks y^* attaining:

$$\min_y \max_x y^T A x$$

which is equivalent to:

$$\begin{aligned} \min u & \\ ue - A^T y &\geq 0 \\ e^T y &= 1 \\ y &\geq 0 \end{aligned}$$

Row Player's Problem in Block-Matrix Form

$$\min \begin{bmatrix} 0 \\ 1 \end{bmatrix}^T \begin{bmatrix} y \\ u \end{bmatrix}$$

$$\begin{bmatrix} -A^T & e \\ e^T & 0 \end{bmatrix} \begin{bmatrix} y \\ u \end{bmatrix} \begin{matrix} \geq \\ = \end{matrix} \begin{bmatrix} 0 \\ 1 \end{bmatrix}$$

$$y \geq 0$$

$$u \;\; \text{free}$$

Note: Row player's problem is dual to column player's:

$$\max \begin{bmatrix} 0 \\ 1 \end{bmatrix}^T \begin{bmatrix} x \\ v \end{bmatrix}$$

$$\begin{bmatrix} -A & e \\ e^T & 0 \end{bmatrix} \begin{bmatrix} x \\ v \end{bmatrix} \begin{matrix} \leq \\ = \end{matrix} \begin{bmatrix} 0 \\ 1 \end{bmatrix}$$

$$x \geq 0$$

$$v \;\; \text{free}$$

$$\min \begin{bmatrix} 0 \\ 1 \end{bmatrix}^T \begin{bmatrix} y \\ u \end{bmatrix}$$

$$\begin{bmatrix} -A^T & e \\ e^T & 0 \end{bmatrix} \begin{bmatrix} y \\ u \end{bmatrix} \begin{matrix} \geq \\ = \end{matrix} \begin{bmatrix} 0 \\ 1 \end{bmatrix}$$

$$y \geq 0$$

$$u \;\; \text{free}$$

MiniMax Theorem

Theorem.

Let x^* *denote column player's solution to her max–min problem.*
Let y^* *denote row player's solution to his min–max problem.*
Then

$$\max_x \; {y^*}^T A x = \min_y \; y^T A x^*.$$

Proof. From *Strong Duality Theorem*, we have

$$u^* = v^*.$$

Also,

$$\begin{aligned} v^* &= \min_i e_i^T A x^* = \min_y y^T A x^* \\ u^* &= \max_j {y^*}^T A e_j = \max_x {y^*}^T A x \end{aligned}$$

QED

AMPL Model

```
set ROWS;
set COLS;
param A {ROWS,COLS} default 0;

var x{COLS} >= 0;
var v;

maximize zot: v;

subject to ineqs {i in ROWS}:
    sum{j in COLS} -A[i,j] * x[j] + v <= 0;

subject to equal:
    sum{j in COLS} x[j] = 1;
```

AMPL Data

```
data;
set ROWS := P S R;
set COLS := P S R;
param A: P   S   R:=
      P   0   1  -2
      S  -3   0   4
      R   5  -6   0
    ;

solve;
printf {j in COLS}: "    %3s %10.7f \n", j, 102*x[j];
printf {i in ROWS}: "    %3s %10.7f \n", i, 102*ineqs[i];
printf: "Value = %10.7f \n", 102*v;
```

AMPL Output

```
ampl gamethy.mod
LOQO: optimal solution (12 iterations)
primal objective -0.1568627451
  dual objective -0.1568627451
      P 40.0000000
      S 36.0000000
      R 26.0000000
      P 62.0000000
      S 27.0000000
      R 13.0000000
Value = -16.0000000
```

Dual of Problems in General Form (Review)

Consider:

$$\begin{aligned} &\max c^T x \\ Ax &= b \\ x &\geq 0 \end{aligned}$$

Rewrite equality constraints as pairs of inequalities:

$$\begin{aligned} &\max c^T x \\ Ax &\leq b \\ -Ax &\leq -b \\ x &\geq 0 \end{aligned}$$

Put into block-matrix form:

$$\begin{aligned} &\max c^T x \\ \begin{bmatrix} A \\ -A \end{bmatrix} x &\begin{matrix} \leq \\ \leq \end{matrix} \begin{bmatrix} b \\ -b \end{bmatrix} \\ x &\geq 0 \end{aligned}$$

Dual is:

$$\begin{aligned} \min \begin{bmatrix} b \\ -b \end{bmatrix}^T \begin{bmatrix} y^+ \\ y^- \end{bmatrix} & \\ \begin{bmatrix} A^T & -A^T \end{bmatrix} \begin{bmatrix} y^+ \\ y^- \end{bmatrix} &\geq c \\ y^+, y^- &\geq 0 \end{aligned}$$

Which is equivalent to:

$$\begin{aligned} &\min b^T (y^+ - y^-) \\ A^T (y^+ - y^-) &\geq c \\ y^+, y^- &\geq 0 \end{aligned}$$

Finally, letting $y = y^+ - y^-$, we get

$$\begin{aligned} &\min b^T y \\ A^T y &\geq c \\ y & \quad \text{free.} \end{aligned}$$

Summary

- Equality constraints $\Longrightarrow$ free variables in dual.
- Inequality constraints $\Longrightarrow$ nonnegative variables in dual.

Corollary:

- Free variables $\Longrightarrow$ equality constraints in dual.
- Nonnegative variables $\Longrightarrow$ inequality constraints in dual.

A Real-World Example

The Ultra-Conservative Investor

Consider again some historical investment data ($S_j(t)$):

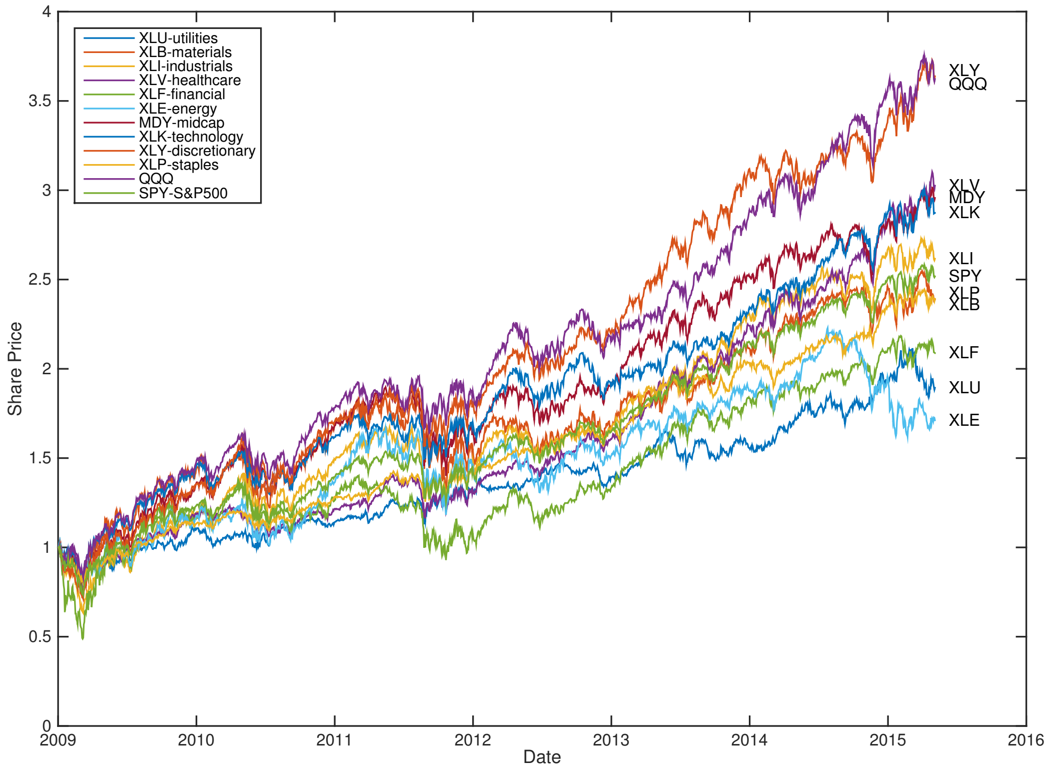

As before, we can let let $R_{t,j} = S_j(t)/S_j(t-1)$ and view R as a payoff matrix in a game between *Fate* and the *Investor*.

Fate's Conspiracy

The columns represent pure strategies for our conservative investor.
The rows represent how history might repeat itself.
Of course, for tomorrow, Fate won't just repeat a previous day's outcome but, rather, will present some mixture of these previous days.
Likewise, the investor won't put all of her money into one asset. Instead she will put a certain fraction into each.
Using this data in the game-theory AMPL model, we get the following mixed-strategy percentages for Fate and for the investor.

Investor's Optimal Asset Mix:

XLP	98.4
XLU	1.6

Mean Old Fate's Mix:

2011-08-08	55.9	$\Longleftarrow$ Black Monday (2011)
2011-08-10	44.1	

The value of the game is the investor's expected return, 96.2%, which is actually a loss of 3.8%.

The data can be download from here: http://finance.yahoo.com/q/hp?s=XLU
Here, XLU is just one of the funds of interest.

Starting From 2012...

To Ignore Black Monday (2011)

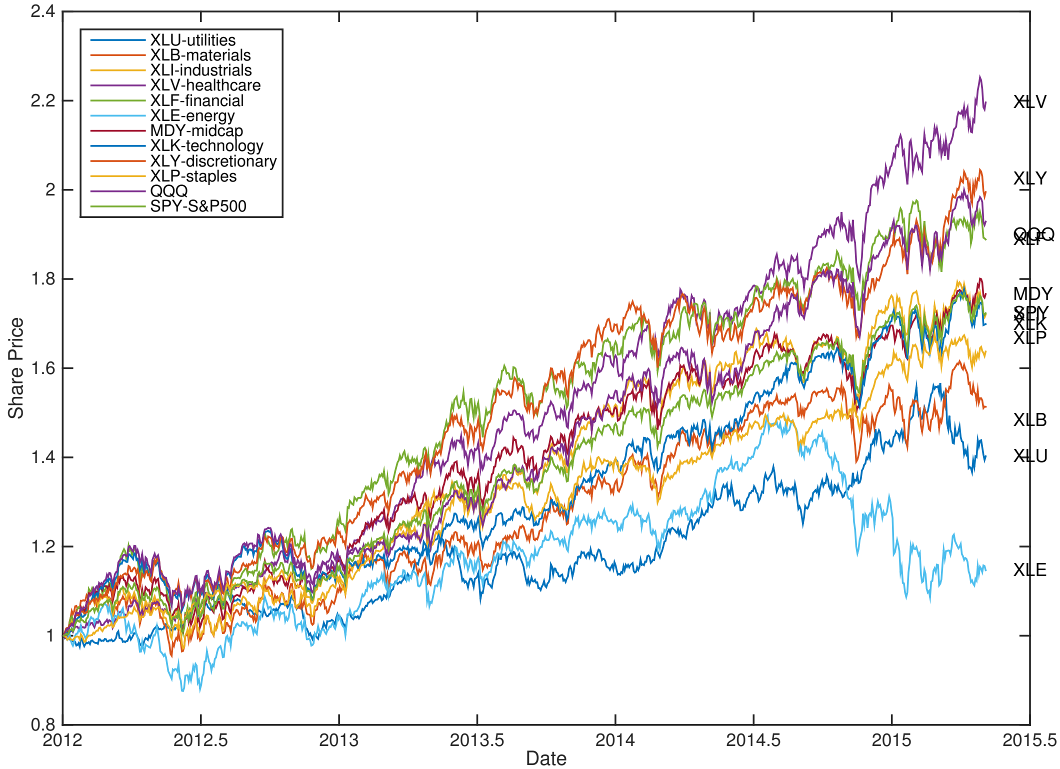

Fate's Conspiracy

Investor's Optimal Asset Mix:

XLK	75.5
XLV	15.9
XLU	6.2
XLB	2.2
XLI	0.2

Mean Old Fate's Mix:

2015-03-25	3.9
2014-04-10	1.7
2013-06-20	68.9
2012-11-07	13.9
2012-06-01	11.5

The value of the game is the investor's expected return, 97.7%, which is actually a loss of 2.3%.

Giving Fate Fewer Options

Thousands seemed unfair—How about 20...

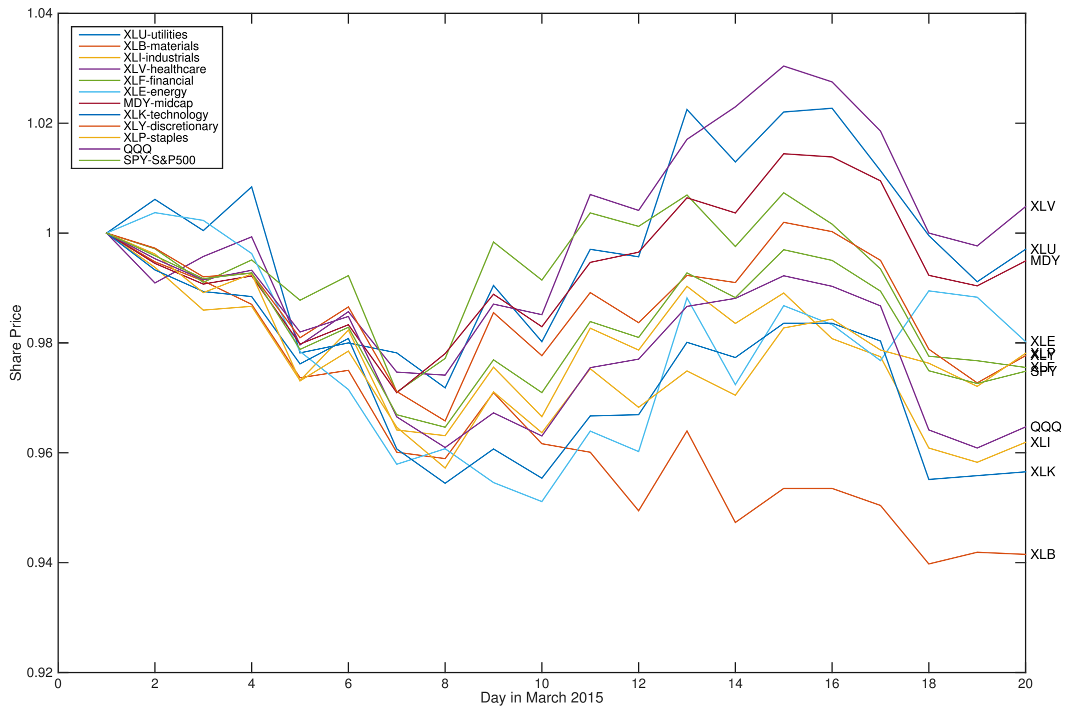

Fate's Conspiracy

Investor's Optimal Asset Mix:

MDY	83.7
XLE	13.2
XLF	3.2

Mean Old Fate's Mix:

2015-03-25	11.5
2015-03-10	33.5
2015-03-06	55.0

The value of the game is the investor's expected return, 98.7%, which is actually a loss of 1.3%.