15.053 Thursday, March 3

- 2-person 0-sum (or constant sum) game theory
  - 2-dimensional
  - multi-dimensional

  - Comments on first midterm:
    - practice test will be on line
    - coverage: every lecture prior to game theory
    - quiz difficulty < midterm difficulty < homework difficulty
    - closed book, no calculator

Game theory is a very broad topic

- 6.972 Game Theory and Equilibrium Analysis
- 14.12 Economic Applications of Game Theory
  – Analysis of strategic behavior in multi-person economic settings.
- 14.122 Microeconomic Theory II
  – Introduction to game theory.
- 14.126 Game Theory
  – Optimal decisions of economic agents depend on expectations of other agents' actions.
- 14.147 Topics in Game Theory
- 17.881 & 17.882 Game Theory and Political Theory
  – Introduces students to the rudiments of game theory within political science.

From Marilyn Vos Savant's column.

Say you're in a public library, and a beautiful stranger strikes up a conversation with you. She says: 'Let's show pennies to each other, either heads or tails. If we both show heads, I pay you $3. If we both show tails, I pay you $1. If they don't match, you pay me $2.'

At this point, she is shushed. You think: 'With both heads 1/4 of the time, I get $3. And with both tails 1/4 of the time, I get $1. So 1/2 of the time, I get $4. And with no matches 1/2 of the time, she gets $4. So it's a fair game.' As the game is quiet, you can play in the library.

But should you? Should she?

submitted by Edward Spellman to Ask Marilyn on 3/31/02

Marilyn Vos Savant has a weekly column in Parade. She has the highest recorded IQ on record.

2-person 0-sum Game Theory

- Two people make decisions at the same time.
- The payoff depends on the joint decisions.
- Whatever one person wins the other person loses.
- Talk with your neighbor to see if you can guess if one person has an advantage in the previous game, and if so, who does? We will return to this game at the end of the lecture.
- Incidentally, Marilyn vos Savant answered the question incorrectly.

2-person 0-sum game theory

Person R chooses a row: either 1, 2, or 3

Person C chooses a column: either 1, 2, or 3

-2	1	2
2	-1	0
1	0	-2

This matrix is the payoff matrix for player R. (And player C gets the negative.)

e.g., R chooses row 3; C chooses column 1

R gets 1; C gets –1 (zero sum)

Some more examples of payoffs

R chooses 2, C chooses 3

R gets 0; C gets 0 (zero sum)

-2	1	2
2	-1	0
1	0	-2

R chooses row 3; C chooses column 3

R gets -2; C gets +2 (zero sum)

Next: 2 volunteers

Player R puts out 1, 2 or 3 fingers

Player C simultaneously puts out 1, 2, or 3 fingers

We will run the game for 5 trials.

-2	1	2
2	-1	0
1	0	-2

R

Total

R tries to maximize his or her total

C tries to minimize R's total.

Next: Play the game with your partner (If you don't have one, then watch)

Player R puts out 1, 2 or 3 fingers

Player C simultaneously puts out 1, 2, or 3 fingers

-2	1	2
2	-1	0
1	0	-2

Run the game for 5 trials.

R

Total	

R tries to maximize his or her total

C tries to minimize R's total.

R	C	Tie

Who has the advantage: R or C?

Suppose that R and C are both brilliant players and they play a VERY LONG TIME.

-2	1	2
2	-1	0
1	0	-2

We will find a the best guaranteed payoff to R using linear programming.

Will R’s payoff be positive in the long run, or will it be negative, or will it converge to 0?

Can R guarantee an expected return independent of what C does? (a lower bound on the opt for R).

Suppose R chooses row j. What can R guarantee?

-2	1	2
2	-1	0
1	0	-2

What row offers R the best guaranteed return?

A strategy that consists of selecting the same row over and over again is a "pure strategy."
R can guarantee a payoff of at least –1.

Random (mixed) strategies

Suppose we permit R to choose a random strategy. What can R guarantee.

-2	1	2
2	-1	0
1	0	-2

Suppose R will flip a coin, and chooses: Row 1 if Heads, and Row 3 if tails.

That is, .5 for Row 1
.5 for Row 3

					Prob.
.5 × 2		-2	1	2	.5
+ 0 × 0		2	-1	0	0
+ .5 × -2		1	0	-2	.5
Expected Payoff		-.5	.5	0	

So, with a random strategy R guarantees at least -.5 regardless of what column C chooses.

Another example: Player R randomizes between row 1 and row 2.

					Prob.
.5 × -2	-2	1	2		.5
+ .5 × 2	2	-1	0		.5
+ 0 × 1	1	0	-2		0
Expected Payoff					

Exercise. Determine the expected payoffs. What can R guarantee?

Optimizing for player R

- Whatever random strategy (x_1, x_2, x_3) that R chooses, we can quickly compute the payoff for each column.

- We will let (x_1, x_2, x_3) be decision variables, and we will write an LP that guarantees the maximum payoff.

What is R's best random strategy?

Prob.
x_1
x_2
x_3

$x_1 + x_2 + x_3 = 1$

Expected Payoff	A	B	C	Prob.
$x_1 \times -2$	-2	1	2	x_1
$+ \; x_2 \times 2$	2	-1	0	x_2
$+ \; x_3 \times 1$	1	0	-2	x_3

$P_A = -2x_1 + 2x_2 + x_3$

$P_B = x_1 - x_2$

$P_C = 2x_1 - 2x_3$

R's strategic problem, as an optimization problem.

	-2	1	2	x_1
	2	-1	0	x_2
	1	0	-2	x_3
Expected Payoff	A	B	C	

$$
\begin{array}{lll}
\text{Maximize} & \min(P_A, P_B, P_C) & \\
& P_A = -2x_1 + 2x_2 + x_3 & \\
& P_B = x_1 - x_2 & \\
& P_C = 2x_1 - 2x_3 & \\
& x_1 + x_2 + x_3 = 1 & \\
& x_1, x_2, x_3 \geq 0 &
\end{array}
$$

R's strategic problem, as an LP

Maximize $\min(P_A, P_B, P_C)$

$$P_A = -2x_1 + 2x_2 + x_3$$
$$P_B = x_1 - x_2$$
$$P_C = 2x_1 - 2x_3$$
$$x_1 + x_2 + x_3 = 1$$
$$x_1, x_2, x_3 \geq 0$$

Maximize z (the payoff to R)

A: $z \leq -2x_1 + 2x_2 + x_3$

B: $z \leq x_1 - x_2$

C: $z \leq 2x_1 - 2x_3$

$$x_1 + x_2 + x_3 = 1$$
$$x_1, x_2, x_3 \geq 0$$

2-person 0-sum game

The Row Player's LP, in general

$x_1 \times a_{11}$

$+ \; x_2 \times a_{21}$

…

$+ \; x_n \times a_{n1}$

Expected Payoff

a_{11}	a_{12}	a_{13}	…	a_{1m}	x_1
a_{21}	a_{22}	a_{23}	…	a_{2m}	x_2
	…				
a_{n1}	a_{n2}	a_{n3}	…	a_{nm}	x_n
P_1	P_2	P_3	…	P_m	

Maximize $\quad z \quad$ (the payoff to x)

P_j: $\quad z \leq a_{1j} x_1 + a_{2j} x_2 + \ldots + a_{nj} x_n$ for all j

$$x_1 + x_2 + \ldots + x_n = 1$$

$$x_j \geq 0 \text{ for all } j$$

An *optimal* random strategy for R.

			Prob.
-2	1	2	7/18
2	-1	0	5/18
1	0	-2	1/3

Expected Payoff			
	1/9	1/9	1/9

The optimal payoff to R for Game 2R is 1/9.

So, R can guarantee at least 1/9 for Game 1.

On the payoff guarantee for R

- **R can guarantee a payoff of 1/9.**
  - **But R does not need to reveal a strategy to C.**
  - **Can R do better?**

- **We will later show that C can guarantee that R gets at most 1/9.**

- **But first …**

A 2-dimensional Example

0	4
2	1

The payoff matrix

Game 2R: the row player declares his or her randomized strategy as follows:

p = prob of selecting row 1
1-p = prob of selecting row 2

For each fixed value of p, the row player R can determine payoff for each column.

In 2-dimensions, graph the payoff as a function of p, and then choose the best value of p.

Graphing the payoff function: Step 1

			Prob
p × 0	0	4	p
+ (1-p) × 2	2	1	1-p

Payoff for Column 1: 2 – 2p

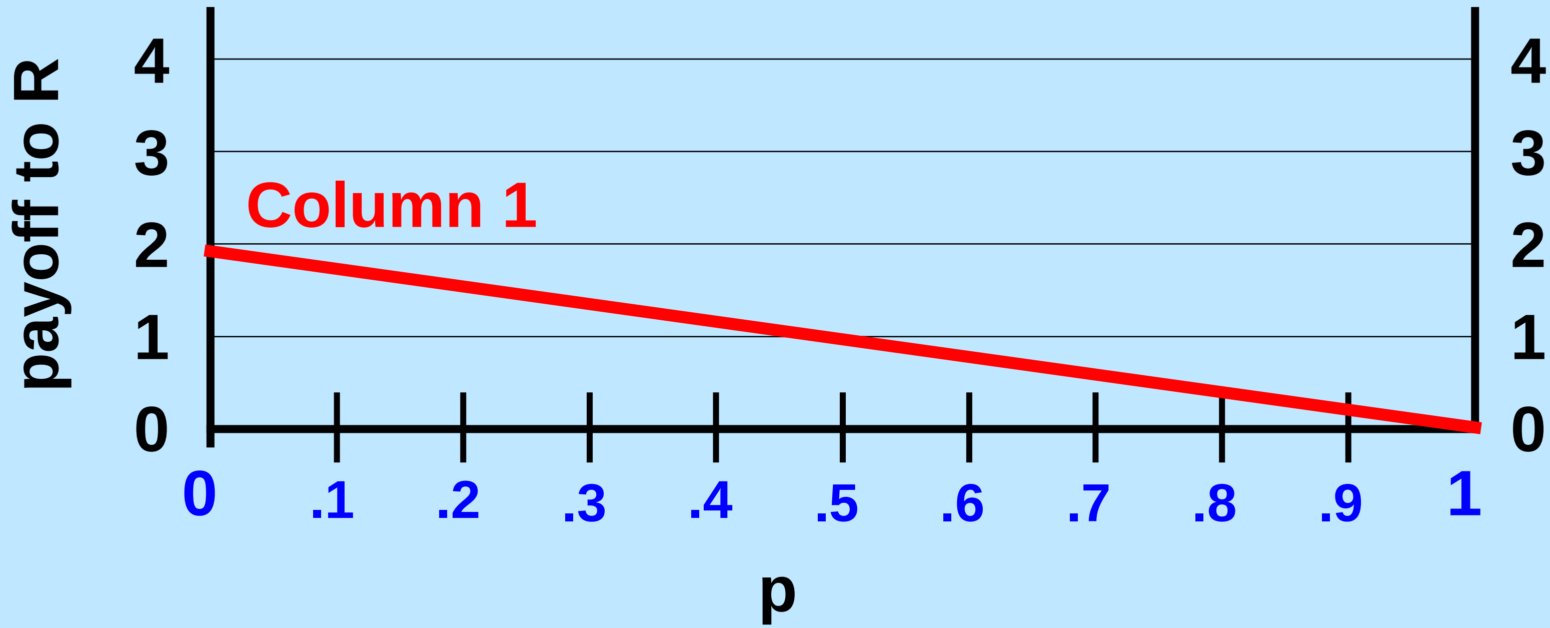

Graphing the payoff function: Step 2

			Prob
$p \times 4$	0	4	p
$+ (1-p) \times 1$	2	1	1-p

Payoff for Column 2: 1 + 3p

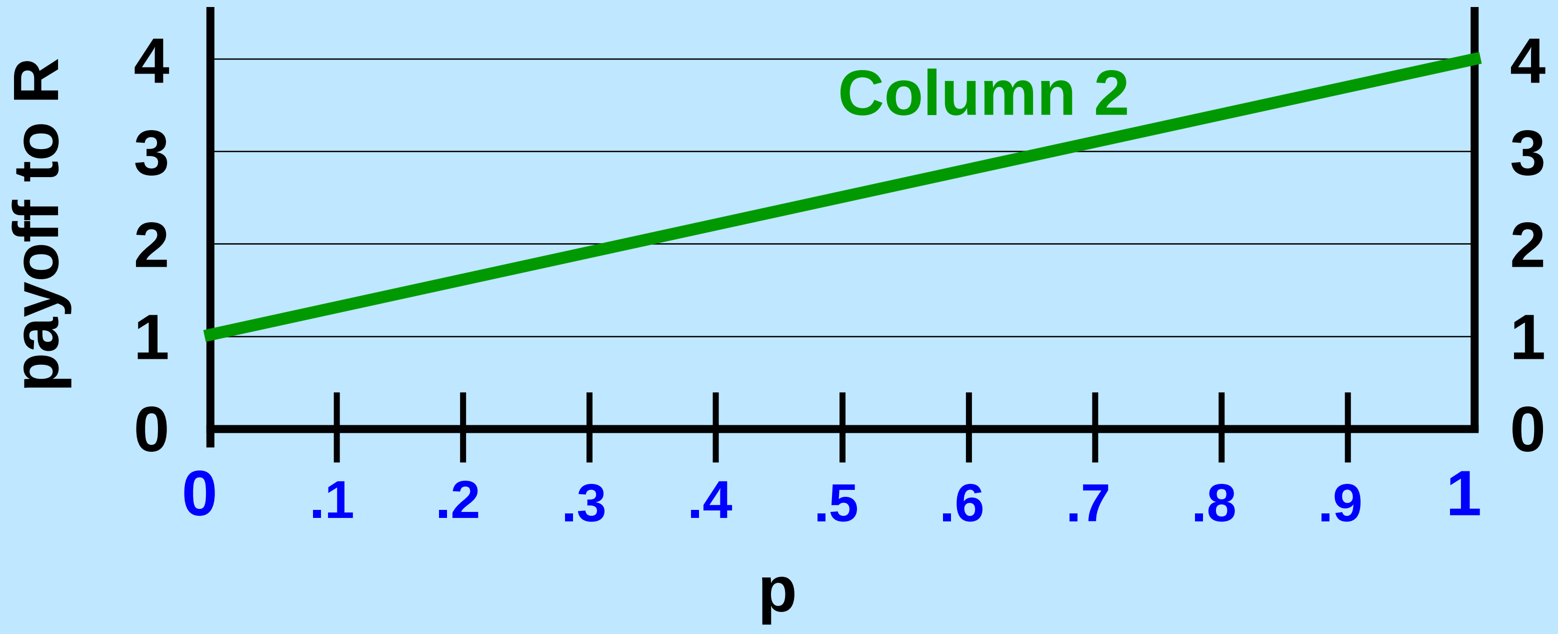

Graphing the payoff function: Step 3

0	4
2	1

Combine the two lines.

The blue line is the best that R can guarantee.

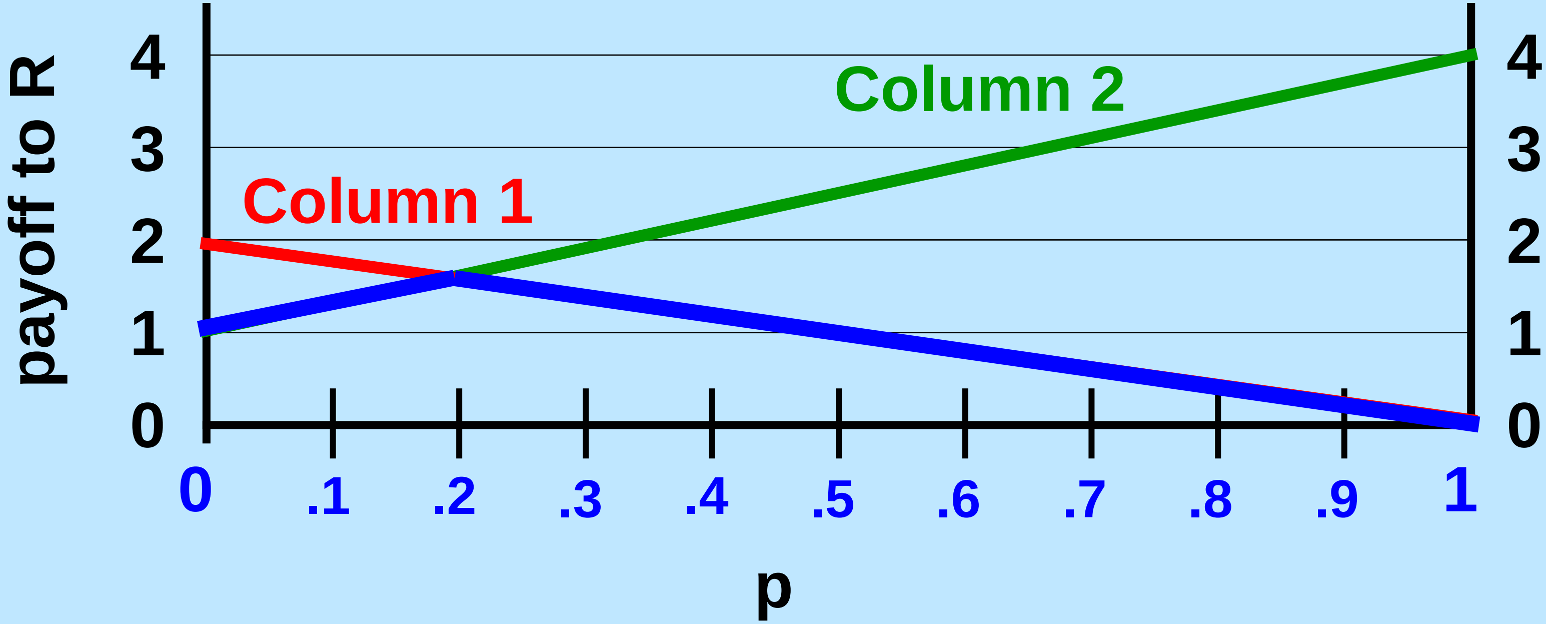

Step 4. Player R chooses the best p.

0	4
2	1

The maximum guarantee can be chosen by selecting the best value of p.

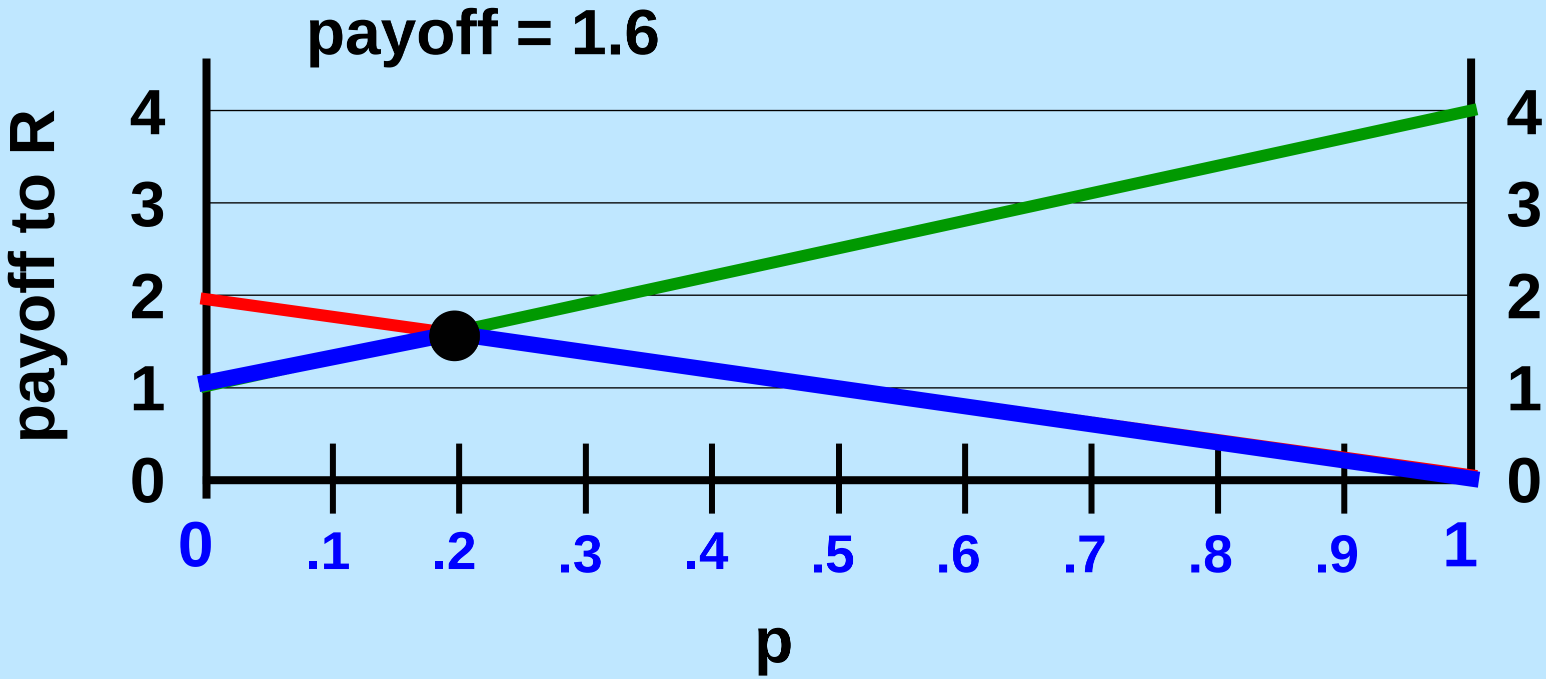

The column player's viewpoint

Prob	q	1-q
	0	4
	2	1

$q \times 0 + (1\text{-}q) \times 4$

Consider the best that the column player can guarantee.

Payoff for Row 1: 4 – 4q

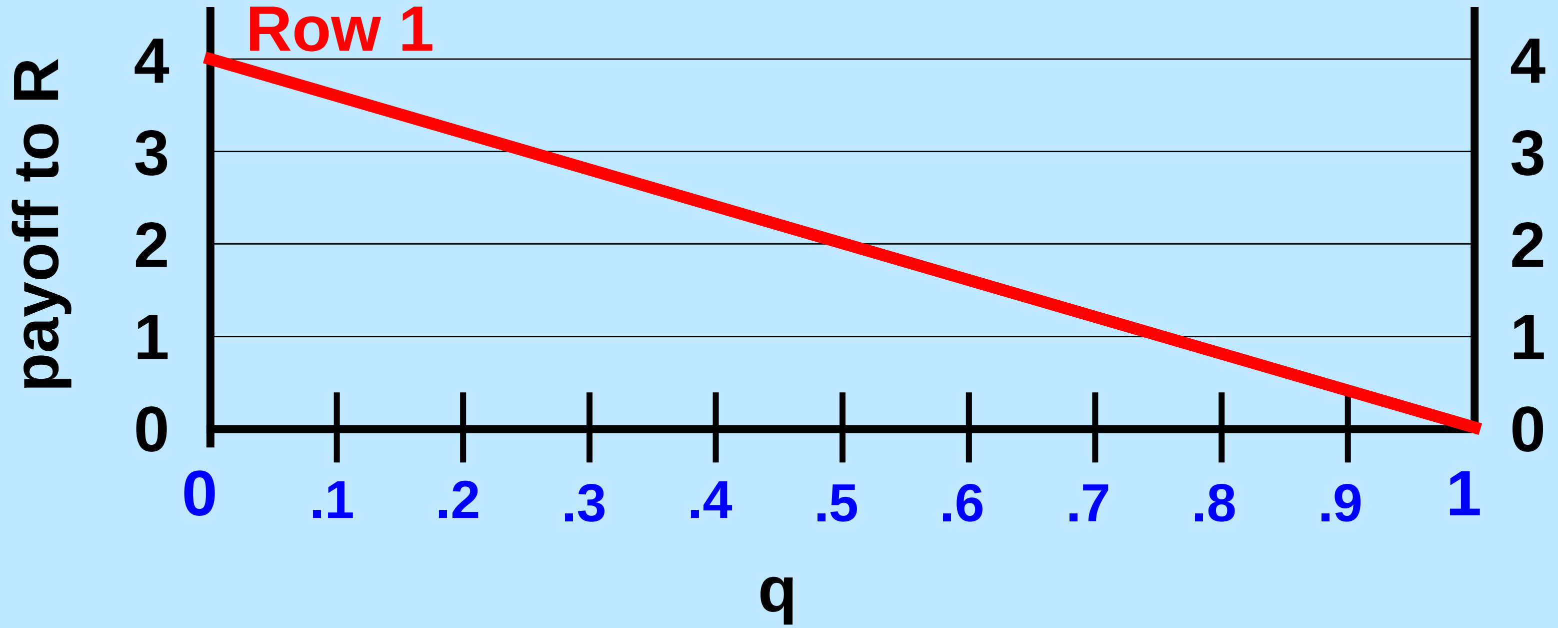

The column player's viewpoint

Prob	q	1-q
	0	4
	2	1

$q \times 2 + (1\text{-}q) \times 1$

Payoff for Row 2: 1 + q

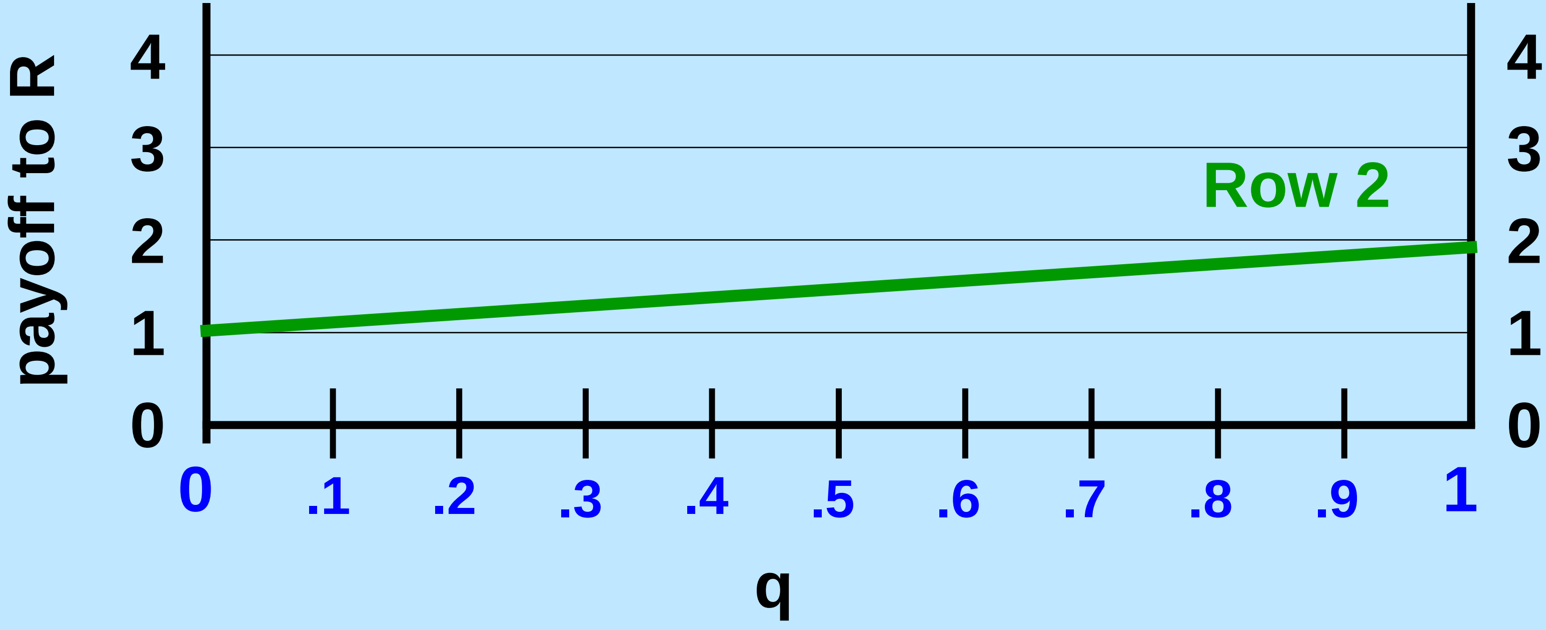

Graphing the payoff function: Step 3

0	4
2	1

Combine the two lines.

For each q, the column player can guarantee a payoff to R of at most the blue line.

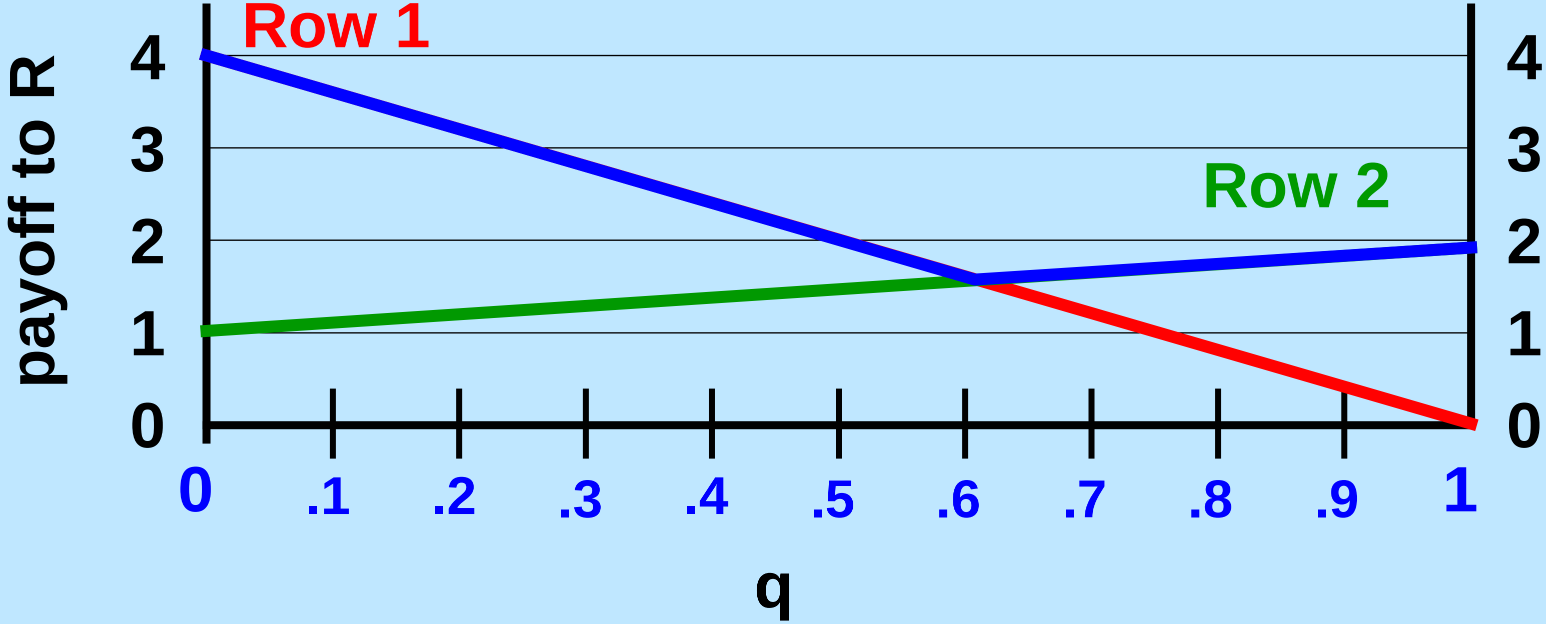

Determining the payoff

0	4
2	1

Player C looks at the payoff function and chooses the value of q that minimizes the payoff to R.

payoff = 1.6

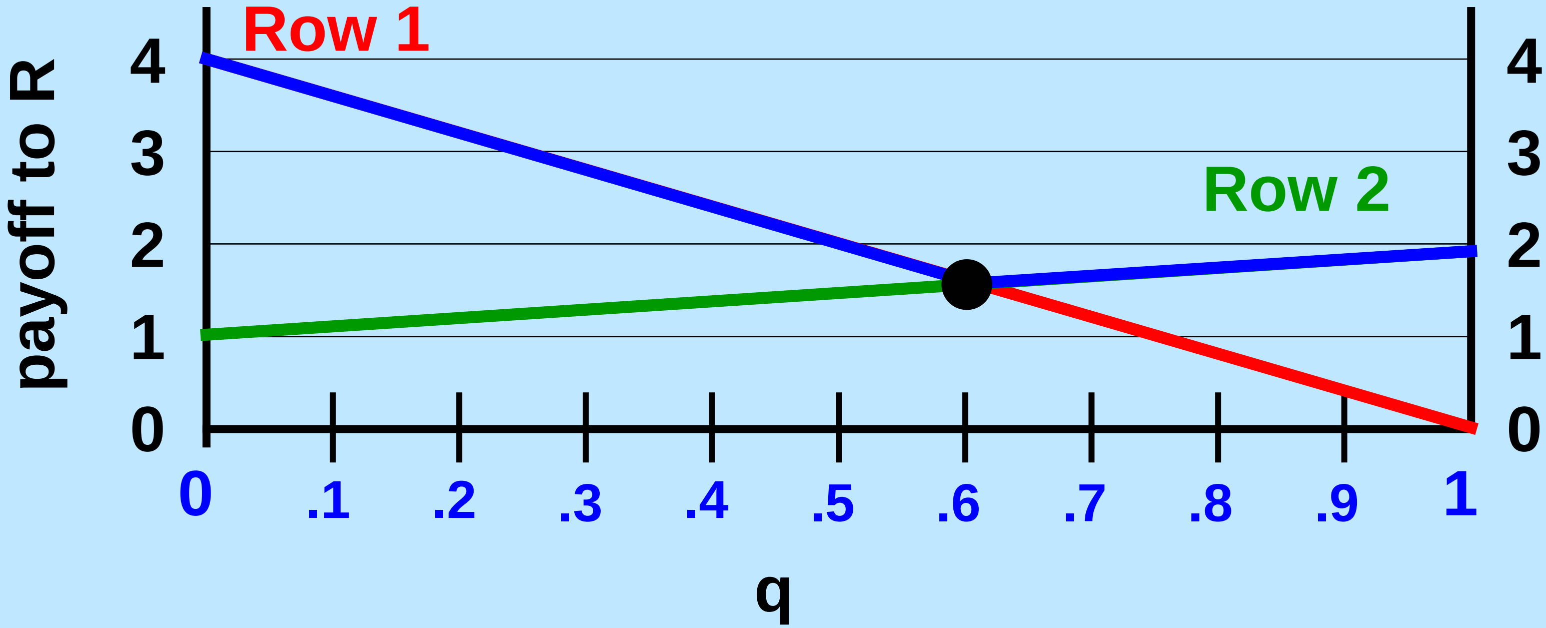

Amazing Result

- **The lower bound for the value of the game for R and the upper bound are the same for all 2 person zero sum games.**
- **That is, linear programming will give you the "value" of the game. This is the best the Row player can guarantee, and the column player can guarantee that no more is obtained.**
- **Von Neumann and Morgenstern**

cepa.newschool.edu/het/ profiles/morgenst.htm

Summary so far

- Game 2R: Player R chooses a random strategy and announces it. Then player C goes next.
  - lower bound on the payoff to R for the original game
  - LP-based approach
  - graph based approach
- Game 2C: Player C chooses a random strategy and player R goes next.
  - This gives us an upper bound on the payoff to R for the original game
  - ***The upper and lower bounds will be the same***

Comments on announcing strategies

- **In reality, players do not announce strategies.**
- **But the point is this: you can choose a mixed strategy that yields a maximum payoff P*, and your opponent can choose a mixed strategy that guarantees you earn no more than P*.**
- **Playing against an opponent who is not perfect, you may do even better. But you can't do worse if you choose your best mixed strategy.**

C chooses a random (mixed) strategy

Exp. payoff	Prob.	y_1	y_2	y_3
$-2\,y_1 + y_2 + 2\,y_3$		-2	1	2
$2\,y_1 - y_2$		2	-1	0
$y_1 - 2\,y_3$		1	0	-2

Minimize v (the payoff to R)

A: $v \geq -2\,y_1 + y_2 + 2\,x_3$

B: $v \geq 2\,y_1 - y_2$

C: $v \geq y_1 - 2\,y_3$

$y_1 + y_2 + y_3 = 1$

$y_1\,,\ y_2\,,\ y_3 \geq 0$

The best strategy for the column player for Game 2C

Prob.	1/3	5/9	1/9	Exp. payoff
	-2	1	2	1/9
	2	-1	0	1/9
	1	0	-2	1/9

So, if C plays the optimal strategy for Game 2C, then R gets at most 1/9.

A Fundamental Theorem of 0-sum Game Theory

- ***Game 2R and Game 2C always produce the same optimal payoff to the R player. Let us call this optimal payoff P*.***

- **The optimal mixed strategy for R always guarantees a payoff of at least P*, regardless of what C does.**

- **The optimal mixed strategy for C always guarantees a payoff to R of at most P*, regardless of what R does.**

- **This is a special case of Linear Programming Duality.**

The Fundamental Theorem in 2-dimensions.

- **We return to the 2-dimensional example**
- **Review: In Game 2R, the Row Player could guarantee a payoff of 1.6 by setting p = .2**
- **We will see that in Game 2C, the Column Player can guarantee that the Row Player gets no more than 1.6**

From Marilyn Vos Savant's column.

Say you're in a public library, and a beautiful stranger strikes up a conversation with you. She says: 'Let's show pennies to each other, either heads or tails. If we both show heads, I pay you $3. If we both show tails, I pay you $1. If they don't match, you pay me $2.'

At this point, she is shushed. You think: 'With both heads 1/4 of the time, I get $3. And with both tails 1/4 of the time, I get $1. So 1/2 of the time, I get $4. And with no matches 1/2 of the time, she gets $4. So it's a fair game.' As the game is quiet, you can play in the library.

But should you? Should she?

Determining the optimal strategy

		B. S. H	B. S. T	Prob
H		3	-2	p
T		-2	1	1-p

You are the row player. What is your best randomized strategy?

Payoff for Column 1:
3p + -2(1-p) = -2 + 5p

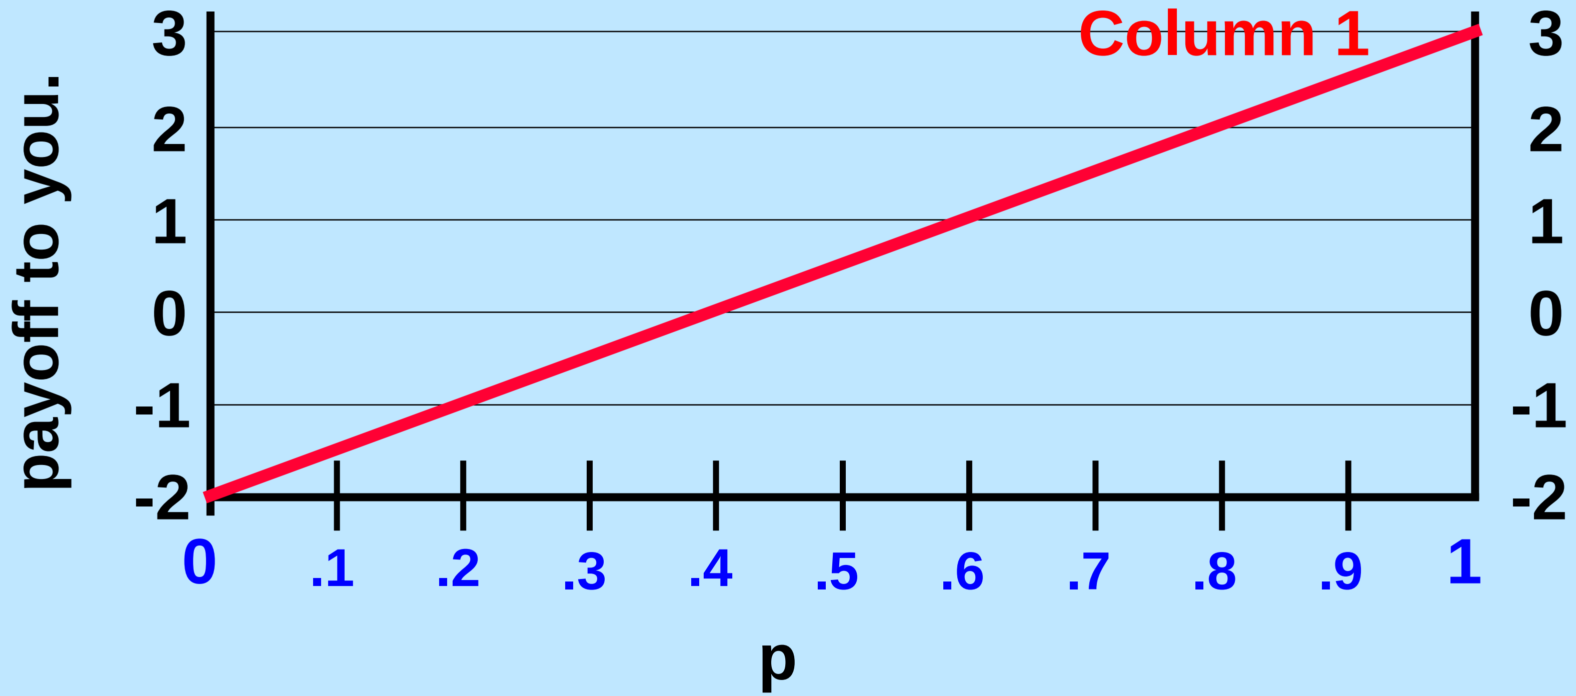

Determining the optimal strategy

	B. S. H	B. S. T	Prob
H	3	-2	p
T	-2	1	1-p

Payoff for Column 2:
-2p + 1(1-p) = 1 - 3p

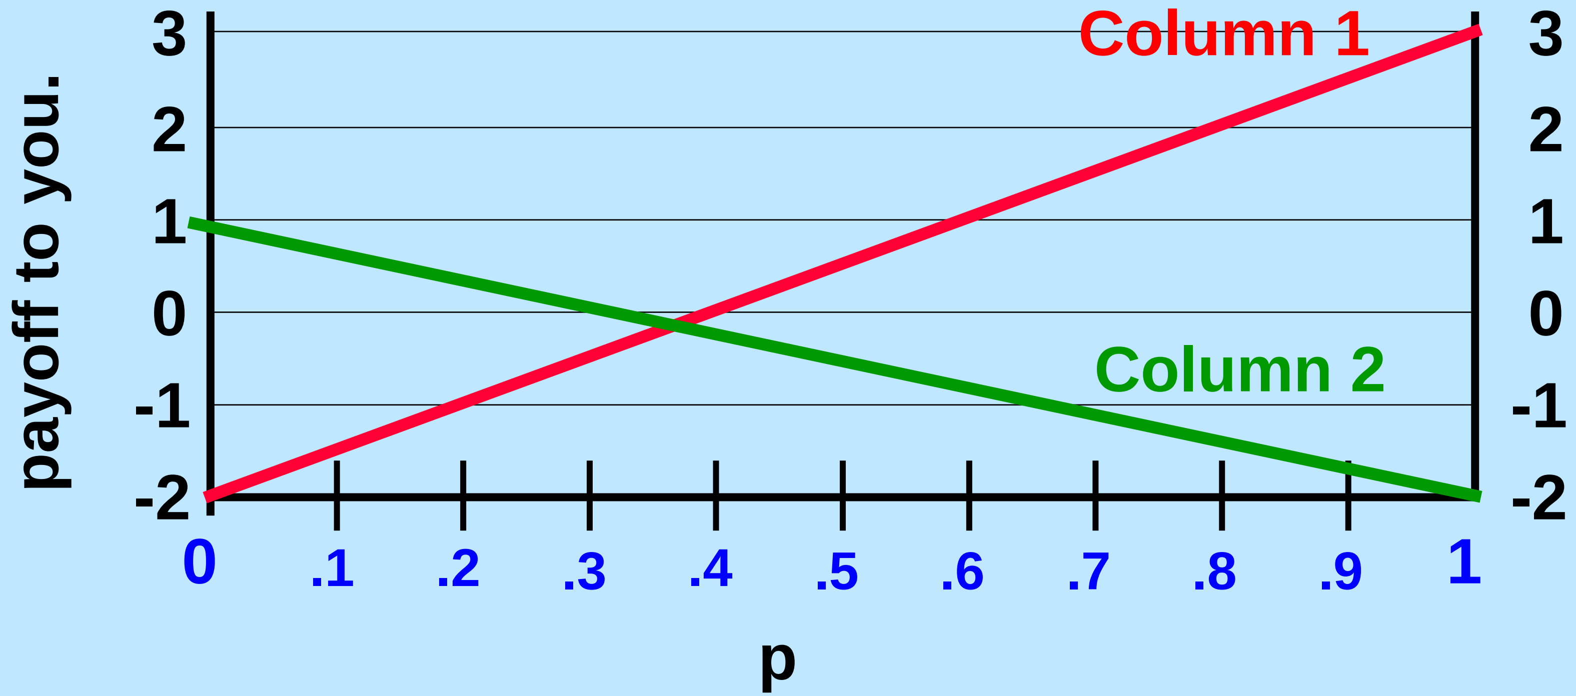

Determining the optimal strategy

	B. S. H	B. S. T	Prob
H	3	-2	p
T	-2	1	1-p

The blue line is the guaranteed payoff.

Player R chooses the value of p that maximizes the payoff.

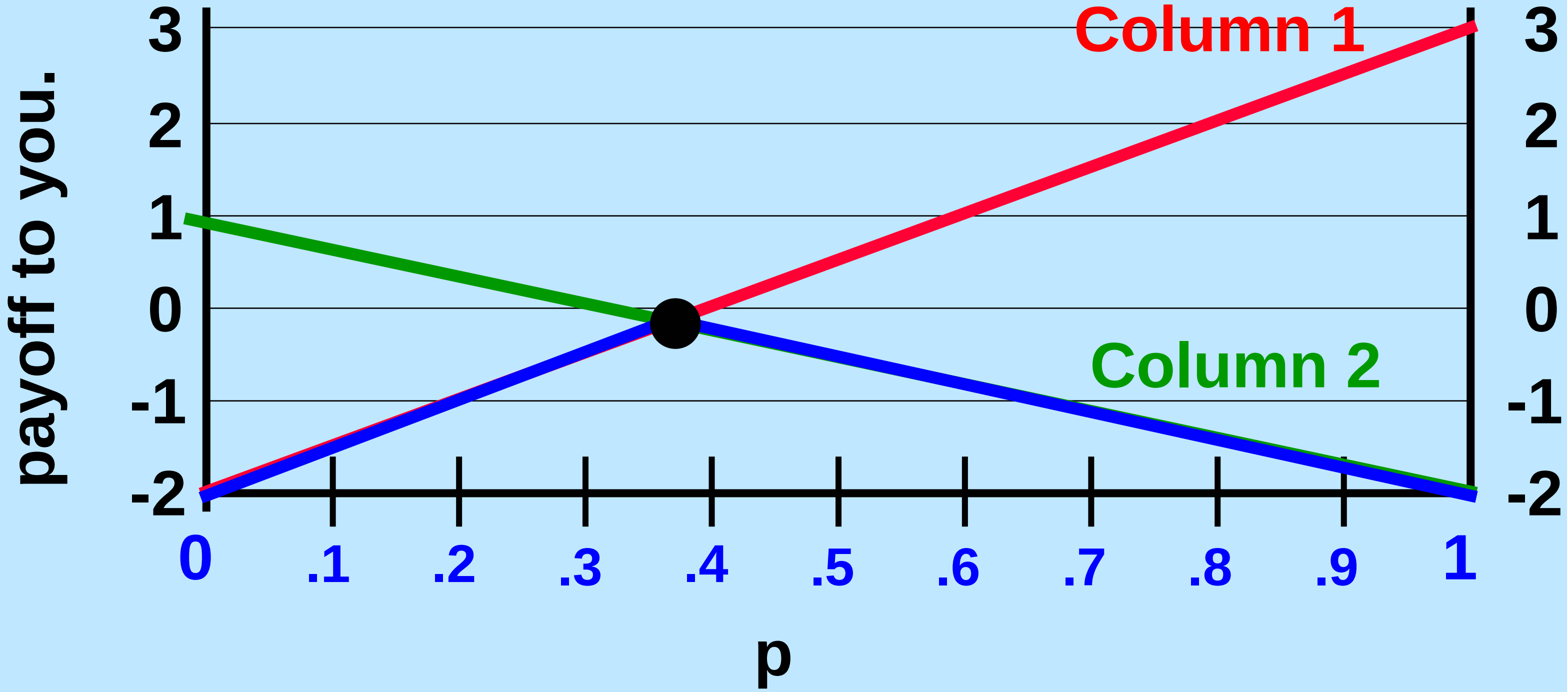

The payoff

Payoff for Column 1:
$= -2 + 5p$

Payoff for Column 2:
$= 1 - 3p$

The payoffs are the same when $p = 3/8$

optimal payoff to row player = $-1/8$

Determining the optimal strategy

	B. S. H	B. S. T
H	3	-2
T	-2	1

What is your best randomized strategy for the Beautiful Stranger?

Marilyn vos Savant chose q = 1/3, which would given the B.S. a payoff of 0.

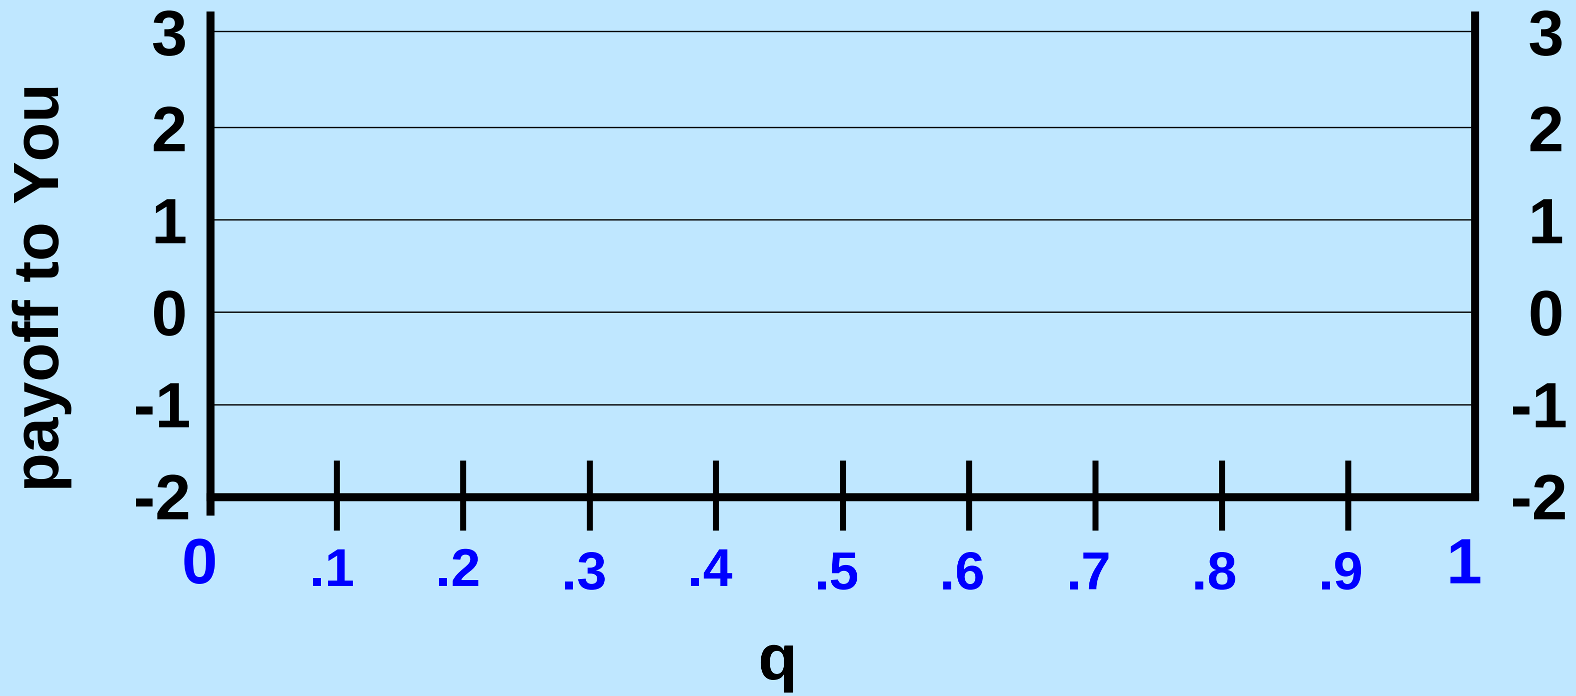

Summary

- 2-person 0 sum game theory
- Can be solved using LP
- 2-dimensional version can be solved graphically
- Fundamental theorem
- Marilyn vos Savant may be the person with the highest measured IQ, but she does not seem to know her game theory.