CS 590.2

Linear Programming Duality, Normal Form Games, and Minimax Theorem

Yu Cheng

Linear Programming Duality

Example linear program

- We make reproductions of two paintings

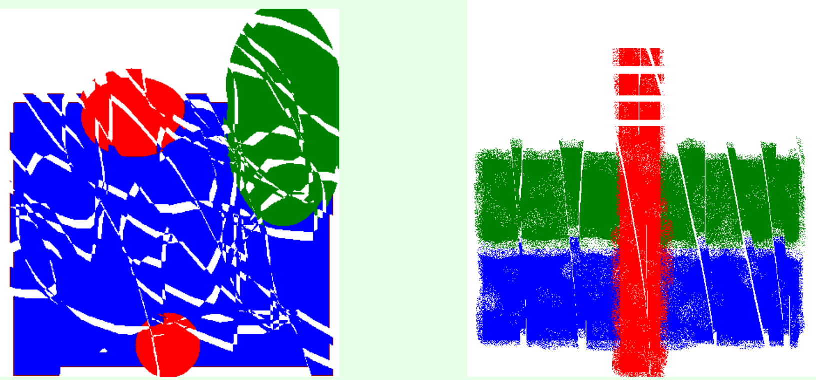

- Painting 1 sells for \$30, painting 2 sells for \$20
- Painting 1 requires 4 units of blue, 1 green, 1 red
- Painting 2 requires 2 blue, 2 green, 1 red
- We have 16 units blue, 8 green, 5 red

maximize $3x + 2y$

subject to

$$4x + 2y \leq 16$$

$$x + 2y \leq 8$$

$$x + y \leq 5$$

$$x \geq 0$$

$$y \geq 0$$

Solving the linear program graphically

maximize $3x + 2y$

subject to

$4x + 2y \leq 16$

$x + 2y \leq 8$

$x + y \leq 5$

$x \geq 0$

$y \geq 0$

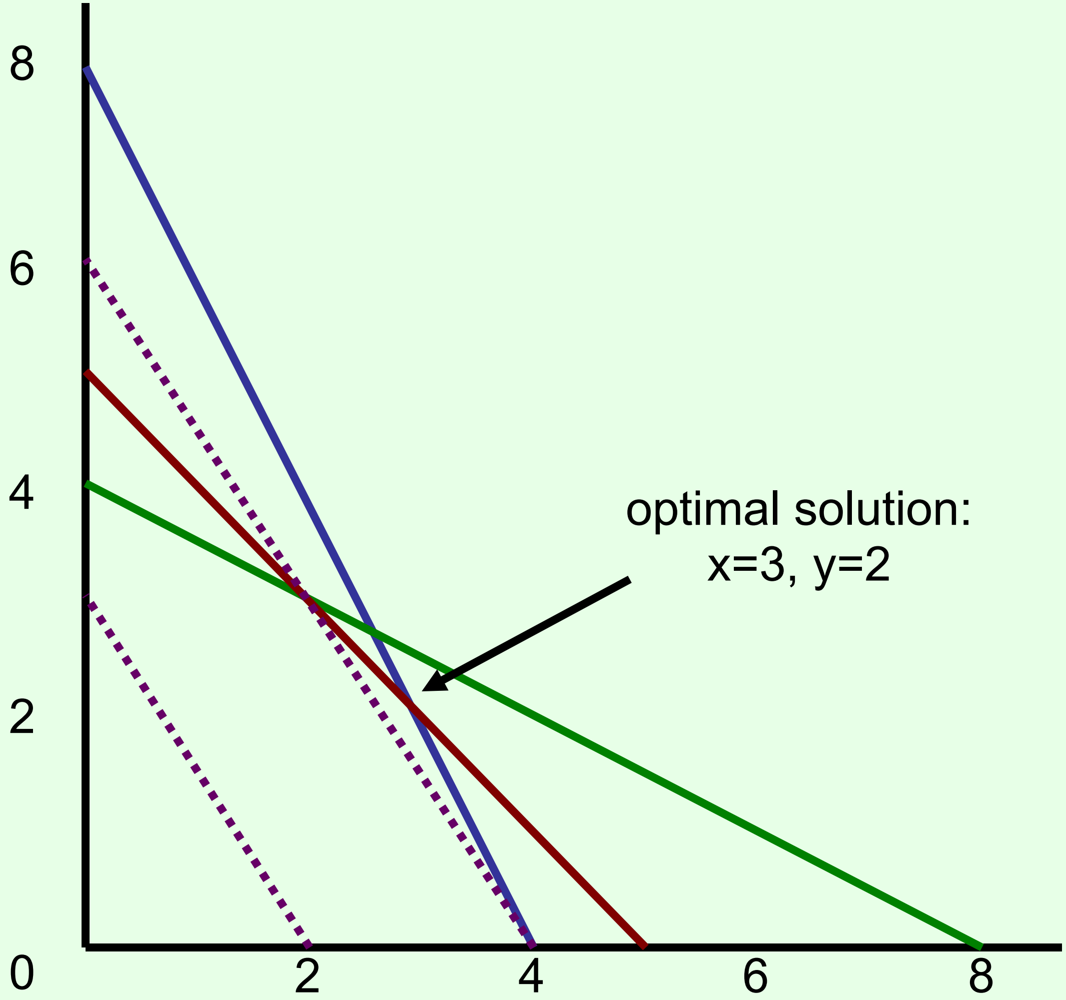

Proving optimality

maximize $3x + 2y$

subject to

$4x + 2y \leq 16$

$x + 2y \leq 8$

$x + y \leq 5$

$x \geq 0$

$y \geq 0$

Recall: optimal solution: $x=3$, $y=2$

Solution value = 9+4 = 13

How do we prove this is optimal (without the picture)?

Proving optimality…

maximize $3x + 2y$

subject to

$4x + 2y \leq 16$

$x + 2y \leq 8$

$x + y \leq 5$

$x \geq 0$

$y \geq 0$

We can rewrite the blue constraint as

$2x + y \leq 8$

If we add the red constraint

$x + y \leq 5$

we get

$3x + 2y \leq 13$

Matching upper bound!

(Really, we added .5 times the blue constraint to 1 times the red constraint)

Linear combinations of constraints

maximize $3x + 2y$

subject to

$4x + 2y \leq 16$

$x + 2y \leq 8$

$x + y \leq 5$

$x \geq 0$

$y \geq 0$

$b(4x + 2y \leq 16) +$

$g(x + 2y \leq 8) +$

$r(x + y \leq 5)$

$=$

$(4b + g + r)x +$

$(2b + 2g + r)y \leq$

$16b + 8g + 5r$

$4b + g + r$ must be at least 3

$2b + 2g + r$ must be at least 2

Given this, minimize $16b + 8g + 5r$

Using LP for getting the best upper bound on an LP

maximize $3x + 2y$

subject to

$4x + 2y \leq 16$

$x + 2y \leq 8$

$x + y \leq 5$

$x \geq 0$

$y \geq 0$

minimize $16b + 8g + 5r$

subject to

$4b + g + r \geq 3$

$2b + 2g + r \geq 2$

$b \geq 0$

$g \geq 0$

$r \geq 0$

the dual of the original program

- Duality theorem: any linear program has the same optimal value as its dual!

Another View

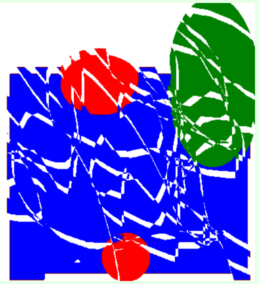

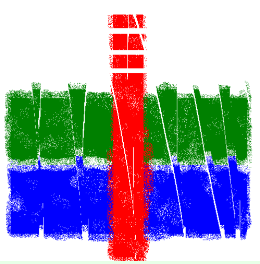

- Painting 1: 4 blue, 1 green, 1 red, sells for \$30
- Painting 2: 2 blue, 2 green, 1 red, sells for \$20
- We have 16 units blue, 8 green, 5 red

- Suppose Vince wants to buy paints from us.
- Pay \$b for a unit of blue, \$g for green, \$r for red.
- We can choose to sell the paints, or produce paintings and sell the paintings, or both.

$$b \geq 0$$
$$g \geq 0$$
$$r \geq 0$$

$$4b + g + r \geq 3$$
$$2b + 2g + r \geq 2$$

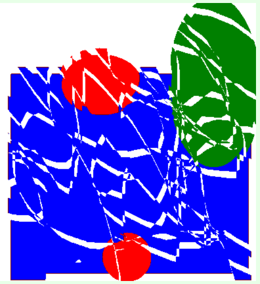

Another View

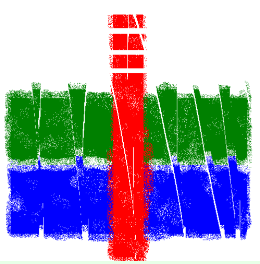

- Vince pays \$(16b + 8g + 5r) in total.

- We have 16 units blue, 8 green, 5 red

  - Suppose Vince wants to buy paints from us.
  - Pay \$b for a unit of blue, \$g for green, \$r for red.
  - We can choose to sell the paints, or produce paintings and sell the paintings, or both.

$$b \geq 0$$
$$g \geq 0$$
$$r \geq 0$$

$$4b + g + r \geq 3$$
$$2b + 2g + r \geq 2$$

Using LP for getting the best upper bound on an LP

maximize $3x + 2y$	*minimize* $16b + 8g + 5r$
subject to	*subject to*
$4x + 2y \leq 16$	$4b + g + r \geq 3$
$x + 2y \leq 8$	$2b + 2g + r \geq 2$
$x + y \leq 5$	$b \geq 0$
$x \geq 0$	$g \geq 0$
$y \geq 0$	$r \geq 0$
primal	dual

Duality

- Weak duality:

  Optimal value of primal ≤ Optimal value of dual
  - when primal is maximize(...) and dual is minimize(…)

- We can make $13 if we produce paintings

  Vince should pay at least as much

- Any upper bound we get from the dual should be at least the optimal value of the primal

Duality

- Strong Duality

  Optimal value of primal = Optimal value of dual

- We can make $13 if we produce paintings

  Vince should pay at least as much

  Vince is a good negotiator and can buy all the paints with $13.

- Any upper bound we get from the dual should be at least the optimal value of the primal

  Optimal dual solution gives a **tight** upper bound

Using LP for getting the best upper bound on an LP

maximize $3x + 2y$

subject to

$4x + 2y \le 16$

$x + 2y \le 8$

$x + y \le 5$

$x \ge 0$

$y \ge 0$

primal

minimize $16b + 8g + 5r$

subject to

$4b + g + r \ge 3$

$2b + 2g + r \ge 2$

$b \ge 0$

$g \ge 0$

$r \ge 0$

dual

Normal-Form Games

Rock-paper-scissors

Column player aka.
player 2
(simultaneously)
chooses a column

		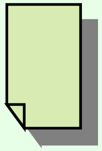	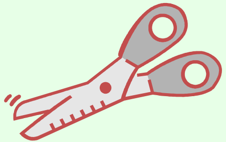
	0, 0	-1, 1	1, -1
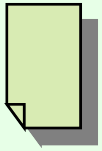	1, -1	0, 0	-1, 1
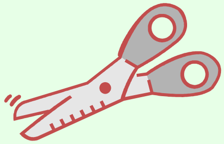	-1, 1	1, -1	0, 0

Row player
aka. player 1
chooses a row

A row or column is called an action or (pure) strategy

Row player's utility is always listed first, column player's second

Zero-sum game: the utilities in each entry sum to 0 (or a constant)
Three-player game would be a 3D table with 3 utilities per entry, etc.

Matching pennies (~penalty kick)

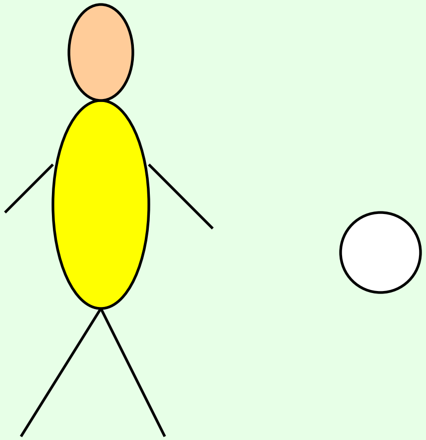

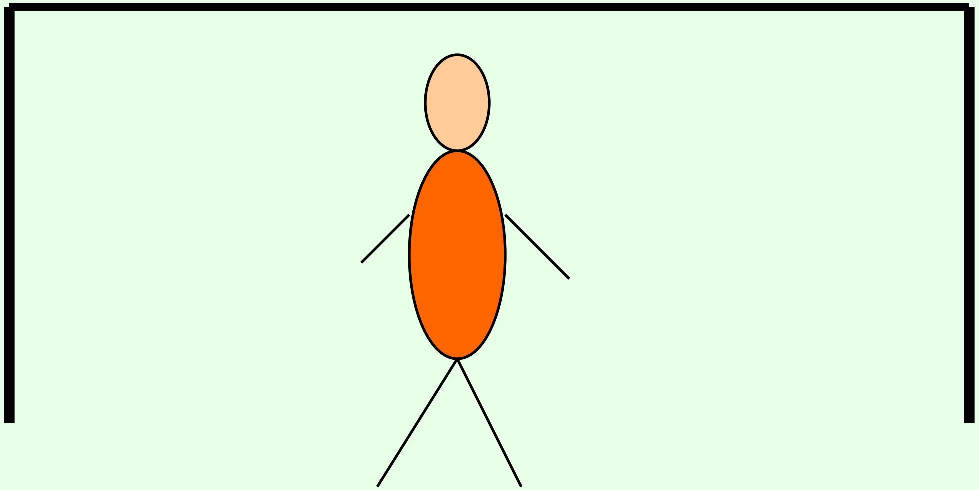

	L	R
L	1, -1	-1, 1
R	-1, 1	1, -1

Two-player zero-sum games

- In a zero-sum game, payoffs in each entry sum to zero
  - … or to a constant: recall that we can subtract a constant from anyone’s utility function without affecting their behavior
- What the one player gains, the other player loses

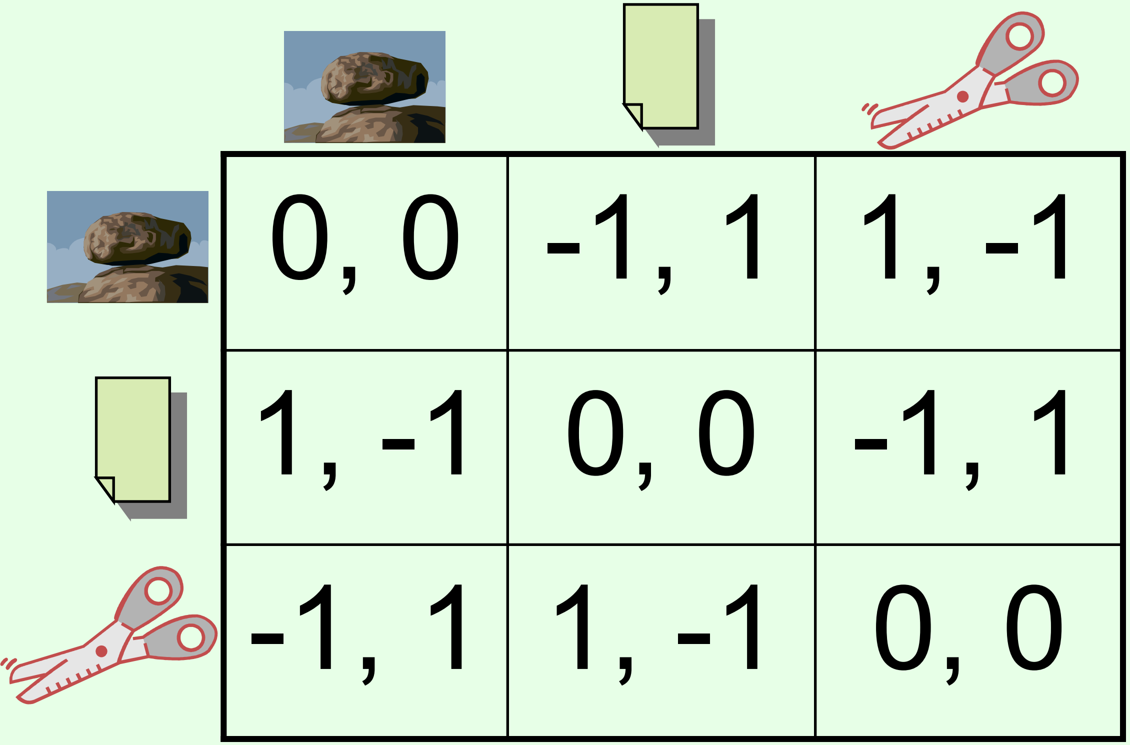

	🪨	📄	✂
🪨	0, 0	-1, 1	1, -1
📄	1, -1	0, 0	-1, 1
✂	-1, 1	1, -1	0, 0

Note: a general-sum k-player game can be modeled as a zero-sum (k+1)-player game by adding a dummy player absorbing the remaining utility, so zero-sum games with 3 or more players have to deal with the difficulties of general-sum games; this is why we focus on 2-player zero-sum games here.

Mixed strategies

- Mixed strategy for player i = probability distribution over player i's (pure) strategies
- E.g.,1/3  , 1/3 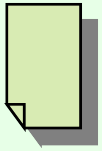 , 1/3 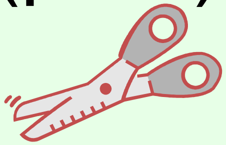

- If we go second:
  - Suppose we know the opponent's mixed strategy, but not his coin flips.
  - What is the best strategy for us to play?

- If we go first:
  - Assume opponent knows our **mixed** strategy (but not our coin flips) and he plays his best-response.
  - What is the best mixed strategy?

Best-response strategies

- Opponent plays rock 50% of the time and scissors 50%
  - Rock gives .5*0 + .5*1 = .5
  - Paper gives .5*1 + .5*(-1) = 0
  - Scissors gives .5*(-1) + .5*0 = -.5

- So the best response to this opponent strategy is to (always) play rock
- There is always some pure strategy that is a best response
  - Suppose you have a mixed strategy that is a best response; then every one of the pure strategies that that mixed strategy places positive probability on must also be a best response

How to play matching pennies

		Them	
		L	R
Us	L	1, -1	-1, 1
	R	-1, 1	1, -1

- Assume opponent knows our **mixed** strategy
- If we play L 60%, R 40%:
  - opponent will play R
  - we get .6*(-1) + .4*(1) = -.2
- What's optimal for us? What about rock-paper-scissors?

Matching pennies with a sensitive target

		Them	
		L	R
Us	L	1, -1	-1, 1
	R	-2, 2	1, -1

- If we play 50% L, 50% R, opponent will attack L
  - We get .5*(1) + .5*(-2) = -.5
- What if we play 55% L, 45% R?
- Opponent has choice between
  - L: gives them .55*(-1) + .45*(2) = .35
  - R: gives them .55*(1) + .45*(-1) = .1
- We get -.35 > -.5

Matching pennies with a sensitive target

Them

		L	R
Us	L	1, -1	-1, 1
	R	-2, 2	1, -1

- What if we play 60% L, 40% R?
- Opponent has choice between
  - L: gives them .6*(-1) + .4*(2) = .2
  - R: gives them .6*(1) + .4*(-1) = .2
- We get -.2 either way
- This is the maximin strategy
  - Maximizes our minimum utility

Let’s change roles

		Them	
		L	R
Us	L	1, -1	-1, 1
	R	-2, 2	1, -1

- Suppose **we** know **their** strategy
- If they play 50% L, 50% R,
  - We play L, we get .5*(1)+.5*(-1) = 0
- If they play 40% L, 60% R,
  - If we play L, we get .4*(1)+.6*(-1) = -.2
  - If we play R, we get .4*(-2)+.6*(1) = -.2
- This is the minimax strategy

von Neumann’s minimax theorem [1928]: maximin value = minimax value (~LP duality)

Minimax Theorem

Minimax theorem [von Neumann 1928]

- Maximin utility: $\max_{\sigma_i} \min_{s_{-i}} u_i(\sigma_i, s_{-i})$
- Minimax utility: $\min_{\sigma_{-i}} \max_{s_i} u_i(s_i, \sigma_{-i})$

Notation:
σ_i denotes a
mixed strategy,
s_i denotes a
pure strategy

- Minimax theorem:

  $\max_{\sigma_i} \min_{s_{-i}} u_i(\sigma_i, s_{-i}) = \min_{\sigma_{-i}} \max_{s_i} u_i(s_i, \sigma_{-i})$

- Minimax theorem does not hold with pure strategies only (example?)

Solving for minimax strategies using linear programming

- maximize $\mathbf{u_i}$
- subject to

$\Sigma_{s_i} \mathbf{p_{s_i}} = 1$

for any s_{-i}, $\Sigma_{s_i} \mathbf{p_{s_i}} u_i(s_i, s_{-i}) \geq \mathbf{u_i}$

Can also convert linear programs to two-player zero-sum games, so they are equivalent

LP duality ~ minimax theorem

	r	b
x	1, -1	-1, 1
y	-2, 2	1, -1

We play a mixed strategy (x, y)

If opponent plays left column: x - 2y
If opponent plays right column: -x + y

maximize v
subject to
x - 2y >= v
-x + y >= v
x + y = 1
x, y >= 0

LP duality ~ minimax theorem

maximize v
subject to
x - 2y >= v
-x + y >= v
x + y = 1
x, y >= 0

r(x - 2y) +
b(-x + y)
=
(r - b)x +
(-2r + b)y
>= (r + b)v

minimize u
subject to
r - b <= u
-2r + b <= u
r + b = 1
r, b >= 0

(r + b)v <= (r - b)x + (-2r + b)y

When r + b = 1, r - b <= u, and -2r + b <= u

v <= ux + uy = u

LP duality ~ minimax theorem

	r	b
x	1, -1	-1, 1
y	-2, 2	1, -1

maximize v
subject to
x - 2y >= v
-x + y >= v
x + y = 1
x, y >= 0

minimize u
subject to
r - b <= u
-2r + b <= u
r + b = 1
r, b >= 0

LP duality ~ minimax theorem

	r	b
x	1, -1	-1, 1
y	-2, 2	1, -1

maximize v

subject to

$x - 2y >= v$

$-x + y >= v$

$x + y = 1$

$x, y >= 0$

← maximin

=

minimax →

minimize u

subject to

$r - b <= u$

$-2r + b <= u$

$r + b = 1$

$r, b >= 0$

General-sum games

- You could still play a minimax strategy in general-sum games
  - I.e., pretend that the opponent is only trying to hurt you
- But this is not rational:

0, 0	3, 1
1, 0	2, 1

- If Column was trying to hurt Row, Column would play Left, so Row should play Down
- In reality, Column will play Right (strictly dominant), so Row should play Up
- Is there a better generalization of minimax strategies in zero-sum games to general-sum games?

Nash equilibrium

[Nash 50]

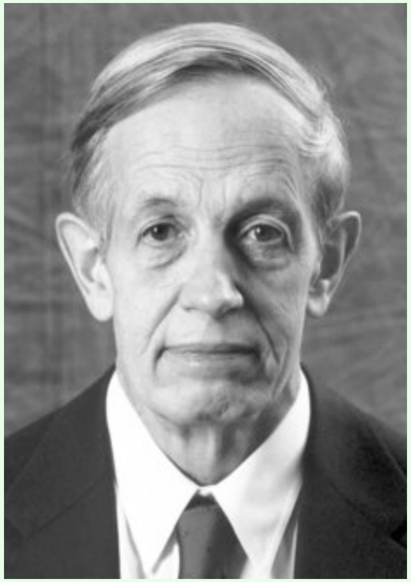

- One mixed strategy for each player
- Every player knows the mixed strategies of the other players
- No player has incentive to deviate