CS599: Convex and Combinatorial Optimization
Fall 2013
Lecture 3: Linear Programming Duality II

Instructor: Shaddin Dughmi

Announcements

- Today: wrap up linear programming
- Readings on website

Outline

Weak and Strong Duality

Primal LP

$$\begin{aligned} \text{maximize} \quad & c^\intercal x \\ \text{subject to} \quad & Ax \le b \\ & x \ge 0 \end{aligned}$$

Dual LP

$$\begin{aligned} \text{minimize} \quad & b^\intercal y \\ \text{subject to} \quad & A^\intercal y \ge c \\ & y \ge 0 \end{aligned}$$

Theorem (Weak Duality)

$OPT(primal) \le OPT(dual)$.

Theorem (Strong Duality)

$OPT(primal) = OPT(dual)$.

Informal Proof of Strong Duality

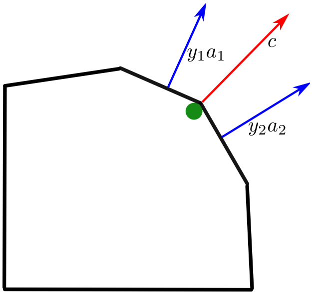

- Recall the physical interpretation of duality

Informal Proof of Strong Duality

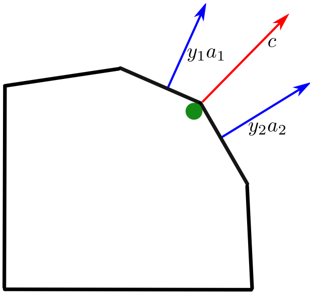

- Recall the physical interpretation of duality
- When ball is stationary at x, we expect force c to be neutralized only by constraints that are tight. i.e. force multipliers $y \geq 0$ s.t.
  - $y^\intercal A = c$
  - $y_i(b_i - a_i x) = 0$

Informal Proof of Strong Duality

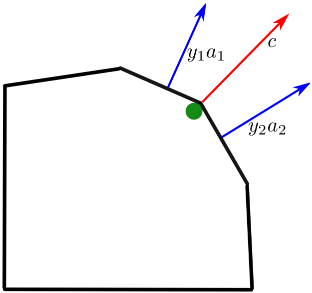

- Recall the physical interpretation of duality
- When ball is stationary at x, we expect force c to be neutralized only by constraints that are tight. i.e. force multipliers $y \geq 0$ s.t.
  - $y^\intercal A = c$
  - $y_i(b_i - a_i x) = 0$

$$y^\intercal b - c^\intercal x = y^\intercal b - y^T A x = \sum_i y_i(b_i - a_i x) = 0$$

We found a primal and dual solution that are equal in value!

Outline

Separating Hyperplane Theorem

If $A, B \subseteq \mathbb{R}^n$ are disjoint convex sets, then there is a hyperplane separating them. That is, there is $a \in \mathbb{R}^n$ and $b \in \mathbb{R}$ such that $a^\intercal x \leq b$ for every $x \in A$ and $a^\intercal y \geq b$ for every $y \in B$.

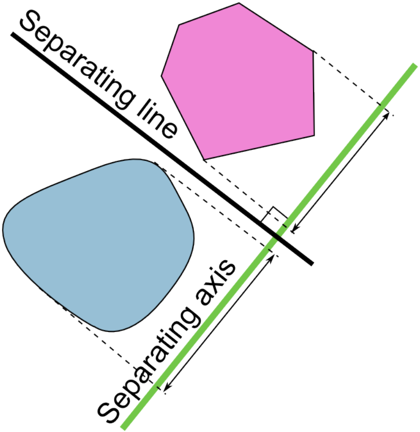

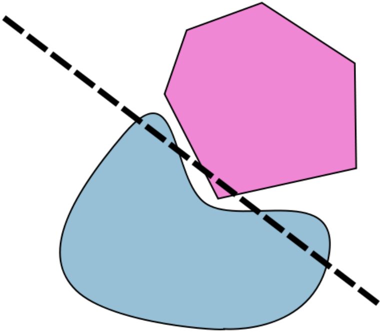

Definition

A convex cone is a convex subset of $\mathbb{R}^n$ which is closed under nonnegative scaling and convex combinations.

Definition

The convex cone generated by vectors $u_1, \ldots, u_m \in \mathbb{R}^n$ is the set of all nonnegative-weighted sums of these vectors (also known as conic combinations).

$$Cone(u_1, \ldots, u_m) = \left\{ \sum_{i=1}^{m} \alpha_i u_i : \alpha_i \geq 0 \; \forall i \right\}$$

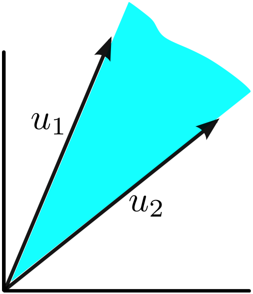

The following follows from the separating hyperplane Theorem.

Farkas' Lemma

Let $\mathcal{C}$ be the convex cone generated by vectors $u_1, \ldots, u_m \in \mathbb{R}^n$, and let $w \in \mathbb{R}^n$. Exactly one of the following is true:

- $w \in \mathcal{C}$
- There is $z \in \mathbb{R}^n$ such that $z \cdot u_i \leq 0$ for all i, and $z \cdot w \geq 0$.

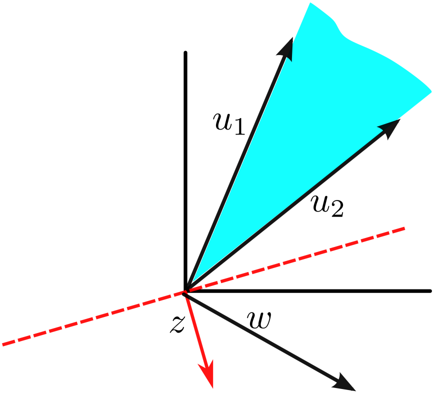

Equivalently: Theorem of the Alternative

One of the following is true, where $U = [u_1, \ldots, u_m]$

- The system $Uy = w$, $y \geq 0$ has a solution
- The system $U^\intercal z \leq 0$, $z^\intercal w \geq 0$ has a solution.

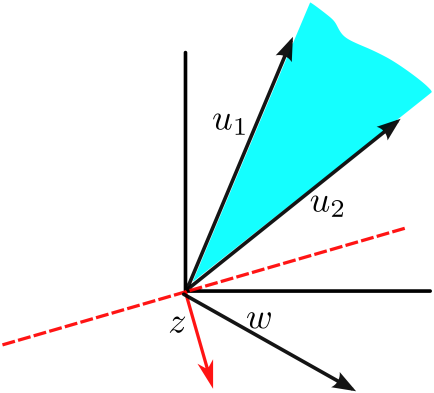

Formal Proof of Strong Duality

Primal LP

maximize $c^\intercal x$
subject to $Ax \leq b$

Dual LP

minimize $b^\intercal y$
subject to $A^\intercal y = c$
$y \geq 0$

Given v, by Farkas' Lemma one of the following is true

1. The system $\begin{pmatrix} A^\intercal \\ b^\intercal \end{pmatrix} y = \begin{pmatrix} c \\ v \end{pmatrix}$, $y \geq 0$ has a solution.
   - $OPT(dual) \leq v$
2. The system $(A; b)\, z \leq 0$, $z^\intercal \begin{pmatrix} c \\ v \end{pmatrix} \geq 0$ has a solution.
   - Let $z = \begin{pmatrix} z_1 \\ z_2 \end{pmatrix}$, where $z_1 \in \mathbb{R}^n$ and $z_2 \in \mathbb{R}$
   - Setting $x = -z_1/z_2$ gives $Ax \leq b$, $c^T x \geq v$.
   - $OPT(primal) \geq v$

Outline

Complementary Slackness

Primal LP

$$\begin{array}{ll} \text{maximize} & c^\intercal x \\ \text{subject to} & Ax \leq b \\ & x \geq 0 \end{array}$$

Dual LP

$$\begin{array}{ll} \text{minimize} & y^\intercal b \\ \text{subject to} & A^\intercal y \geq c \\ & y \geq 0 \end{array}$$

Complementary Slackness

Primal LP

$$\begin{array}{ll} \text{maximize} & c^\intercal x \\ \text{subject to} & Ax \leq b \\ & x \geq 0 \end{array}$$

Dual LP

$$\begin{array}{ll} \text{minimize} & y^\intercal b \\ \text{subject to} & A^\intercal y \geq c \\ & y \geq 0 \end{array}$$

- Let $s_i = (b - Ax)_i$ be the i'th primal slack variable
- Let $t_j = (A^\intercal y - c)_j$ be the j'th dual slack variable

Complementary Slackness

Primal LP

$$\begin{aligned} \text{maximize} \quad & c^\intercal x \\ \text{subject to} \quad & Ax \leq b \\ & x \geq 0 \end{aligned}$$

Dual LP

$$\begin{aligned} \text{minimize} \quad & y^\intercal b \\ \text{subject to} \quad & A^\intercal y \geq c \\ & y \geq 0 \end{aligned}$$

- Let $s_i = (b - Ax)_i$ be the i'th primal slack variable
- Let $t_j = (A^\intercal y - c)_j$ be the j'th dual slack variable

Complementary Slackness

x and y are optimal if and only if

- $x_j t_j = 0$ for all $j = 1, \ldots, n$
- $y_i s_i = 0$ for all $i = 1, \ldots, m$

	x_1	x_2	x_3	x_4	
y_1	a_{11}	a_{12}	a_{13}	a_{14}	b_1
y_2	a_{21}	a_{22}	a_{23}	a_{24}	b_2
y_3	a_{31}	a_{32}	a_{33}	a_{34}	b_3
	c_1	c_2	c_3	c_4	

Interpretation of Complementary Slackness

Economic Interpretation

Given an optimal primal production vector x and optimal dual offer prices y,

- Facility produces only products for which it is indifferent between sale and production.
- Only raw materials that are binding constraints on production are priced greater than 0

Interpretation of Complementary Slackness

Physical Interpretation

Only walls adjacent to the balls equilibrium position push back on it.

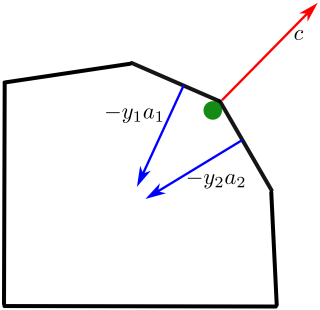

Proof of Complementary Slackness

Primal LP

$$\begin{array}{ll} \text{maximize} & c^{\mathsf{T}} x \\ \text{subject to} & Ax \leq b \\ & x \geq 0 \end{array}$$

Dual LP

$$\begin{array}{ll} \text{minimize} & y^{\mathsf{T}} b \\ \text{subject to} & A^{\mathsf{T}} y \geq c \\ & y \geq 0 \end{array}$$

Proof of Complementary Slackness

Primal LP

$$\begin{aligned} \text{maximize} \quad & c^\intercal x \\ \text{subject to} \quad & Ax + s = b \\ & x \geq 0 \\ & s \geq 0 \end{aligned}$$

Dual LP

$$\begin{aligned} \text{minimize} \quad & y^\intercal b \\ \text{subject to} \quad & A^\intercal y - t = c \\ & y \geq 0 \\ & t \geq 0 \end{aligned}$$

- Can equivalently rewrite LP using slack variables

Proof of Complementary Slackness

Primal LP

maximize	$c^\intercal x$
subject to	$Ax + s = b$
	$x \geq 0$
	$s \geq 0$

Dual LP

minimize	$y^\intercal b$
subject to	$A^\intercal y - t = c$
	$y \geq 0$
	$t \geq 0$

- Can equivalently rewrite LP using slack variables

$$y^\intercal b - c^\intercal x = y^\intercal (Ax + s) - (y^\intercal A - t^\intercal)x = y^\intercal s + t^\intercal x$$

Proof of Complementary Slackness

Primal LP

$$\begin{aligned} \text{maximize} \quad & c^\intercal x \\ \text{subject to} \quad & Ax + s = b \\ & x \geq 0 \\ & s \geq 0 \end{aligned}$$

Dual LP

$$\begin{aligned} \text{minimize} \quad & y^\intercal b \\ \text{subject to} \quad & A^\intercal y - t = c \\ & y \geq 0 \\ & t \geq 0 \end{aligned}$$

- Can equivalently rewrite LP using slack variables

$$y^\intercal b - c^\intercal x = y^\intercal (Ax + s) - (y^\intercal A - t^\intercal) x = y^\intercal s + t^\intercal x$$

Gap between primal and dual objectives is 0 if and only if complementary slackness holds.

Recovering Primal from Dual

- Will encounter LPs where the dual is easier to solve than primal
- Complementary slackness allows us to recover the primal optimal from the dual optimal, and vice versa.

Recovering Primal from Dual

- Will encounter LPs where the dual is easier to solve than primal
- Complementary slackness allows us to recover the primal optimal from the dual optimal, and vice versa.
  - Assuming non-degeneracy: Every vertex of primal [dual] is the solution of exactly n [m] tight constraints.

Recovering Primal from Dual

- Will encounter LPs where the dual is easier to solve than primal
- Complementary slackness allows us to recover the primal optimal from the dual optimal, and vice versa.
  - Assuming non-degeneracy: Every vertex of primal [dual] is the solution of exactly n [m] tight constraints.

Primal LP
(n variables, $m+n$ constraints)

$$\begin{aligned} \text{maximize} \quad & c^\intercal x \\ \text{subject to} \quad & Ax \leq b \\ & x \geq 0 \end{aligned}$$

Dual LP
(m variables, $m+n$ constraints)

$$\begin{aligned} \text{minimize} \quad & y^\intercal b \\ \text{subject to} \quad & A^\intercal y \geq c \\ & y \geq 0 \end{aligned}$$

Recovering Primal from Dual

- Will encounter LPs where the dual is easier to solve than primal
- Complementary slackness allows us to recover the primal optimal from the dual optimal, and vice versa.
  - Assuming non-degeneracy: Every vertex of primal [dual] is the solution of exactly n [m] tight constraints.

Primal LP
(n variables, $m+n$ constraints)

$$\begin{aligned} \text{maximize} \quad & c^\intercal x \\ \text{subject to} \quad & Ax \leq b \\ & x \geq 0 \end{aligned}$$

Dual LP
(m variables, $m+n$ constraints)

$$\begin{aligned} \text{minimize} \quad & y^\intercal b \\ \text{subject to} \quad & A^\intercal y \geq c \\ & y \geq 0 \end{aligned}$$

- Let y be dual optimal. By non-degeneracy:
  - Exactly m of the $m+n$ dual constraints are tight at y
  - Exactly n dual constraints are loose

Recovering Primal from Dual

- Will encounter LPs where the dual is easier to solve than primal
- Complementary slackness allows us to recover the primal optimal from the dual optimal, and vice versa.
  - Assuming non-degeneracy: Every vertex of primal [dual] is the solution of exactly n [m] tight constraints.

Primal LP
(n variables, $m+n$ constraints)

$$\begin{aligned} \text{maximize} \quad & c^\intercal x \\ \text{subject to} \quad & Ax \leq b \\ & x \geq 0 \end{aligned}$$

Dual LP
(m variables, $m+n$ constraints)

$$\begin{aligned} \text{minimize} \quad & y^\intercal b \\ \text{subject to} \quad & A^\intercal y \geq c \\ & y \geq 0 \end{aligned}$$

- Let y be dual optimal. By non-degeneracy:
  - Exactly m of the $m+n$ dual constraints are tight at y
  - Exactly n dual constraints are loose
- n loose dual constraints impose n tight primal constraints

Recovering Primal from Dual

- Will encounter LPs where the dual is easier to solve than primal
- Complementary slackness allows us to recover the primal optimal from the dual optimal, and vice versa.
  - Assuming non-degeneracy: Every vertex of primal [dual] is the solution of exactly n [m] tight constraints.

Primal LP
(n variables, $m + n$ constraints)

$$\begin{array}{ll} \text{maximize} & c^\intercal x \\ \text{subject to} & Ax \leq b \\ & x \geq 0 \end{array}$$

Dual LP
(m variables, $m + n$ constraints)

$$\begin{array}{ll} \text{minimize} & y^\intercal b \\ \text{subject to} & A^\intercal y \geq c \\ & y \geq 0 \end{array}$$

- Let y be dual optimal. By non-degeneracy:
  - Exactly m of the $m + n$ dual constraints are tight at y
  - Exactly n dual constraints are loose
- n loose dual constraints impose n tight primal constraints
  - Assuming non-degeneracy, solving the linear equation yields a unique primal optimum solution x.

Sensitivity Analysis

Primal LP

$$\begin{aligned} \text{maximize} \quad & c^\intercal x \\ \text{subject to} \quad & Ax \leq b \\ & x \geq 0 \end{aligned}$$

Dual LP

$$\begin{aligned} \text{minimize} \quad & y^\intercal b \\ \text{subject to} \quad & A^\intercal y \geq c \\ & y \geq 0 \end{aligned}$$

Sometimes, we want to examine how the optimal value of our LP changes with its parameters c and b

Sensitivity Analysis

Primal LP

$$\begin{aligned} \text{maximize} \quad & c^\intercal x \\ \text{subject to} \quad & Ax \leq b \\ & x \geq 0 \end{aligned}$$

Dual LP

$$\begin{aligned} \text{minimize} \quad & y^\intercal b \\ \text{subject to} \quad & A^\intercal y \geq c \\ & y \geq 0 \end{aligned}$$

Sometimes, we want to examine how the optimal value of our LP changes with its parameters c and b

Sensitivity Analysis

Let $OPT = OPT(A, c, b)$ be the optimal value of the above LP. Let x and y be the primal and dual optima.

- $\frac{\partial OPT}{\partial c_j} = x_j$ when x is the unique primal optimum.
- $\frac{\partial OPT}{\partial b_i} = y_i$ when y is the unique dual optimum.

Sensitivity Analysis

Primal LP

$$\begin{aligned} \text{maximize} \quad & c^\intercal x \\ \text{subject to} \quad & Ax \leq b \\ & x \geq 0 \end{aligned}$$

Dual LP

$$\begin{aligned} \text{minimize} \quad & y^\intercal b \\ \text{subject to} \quad & A^\intercal y \geq c \\ & y \geq 0 \end{aligned}$$

Sometimes, we want to examine how the optimal value of our LP changes with its parameters c and b

Economic Interpretation of Sensitivity Analysis

- A small increase δ in c_j increases profit by $\delta \cdot x_j$
- A small increase δ in b_i increases profit by $\delta \cdot y_i$
  - y_i measures the "marginal value" of resource i for production

Outline

Shortest Path

Given a directed network $G = (V, E)$ where edge e has length $\ell_e \in \mathbb{R}_+$, find the minimum cost path from s to t.

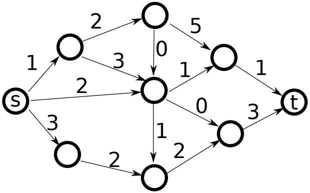

Shortest Path

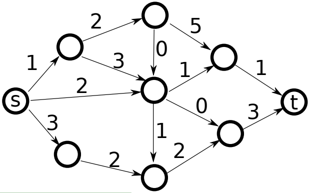

Primal LP

min $\sum_{e \in E} \ell_e x_e$
s.t.

$$\sum_{e \to v} x_e - \sum_{v \to e} x_e = \delta_v, \quad \forall v \in V.$$
$$x_e \geq 0, \quad \forall e \in E.$$

Dual LP

max $y_t - y_s$
s.t.

$$y_v - y_u \leq \ell_e, \quad \forall (u, v) \in E.$$

Where $\delta_v = -1$ if $v = s$, 1 if $v = t$, and 0 otherwise.

Shortest Path

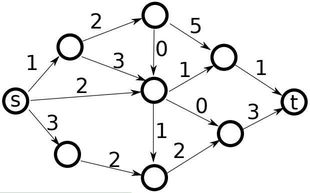

Primal LP

min $\sum_{e \in E} \ell_e x_e$
s.t.

$$\sum_{e \to v} x_e - \sum_{v \to e} x_e = \delta_v, \quad \forall v \in V.$$
$$x_e \geq 0, \quad \forall e \in E.$$

Dual LP

max $y_t - y_s$
s.t.

$$y_v - y_u \leq \ell_e, \quad \forall (u,v) \in E.$$

Where $\delta_v = -1$ if $v = s$, 1 if $v = t$, and 0 otherwise.

Interpretation of Dual

Stretch s and t as far apart as possible, subject to edge lengths.

Maximum Weighted Bipartite Matching

Set B of buyers, and set G of goods. Buyer i has value w_{ij} for good j, and interested in at most one good. Find maximum value assignment of goods to buyers.

Maximum Weighted Bipartite Matching

Primal LP

$$\max \sum_{i,j} w_{ij} x_{ij}$$

s.t.

$$\begin{aligned}
\sum_{j \in G} x_{ij} &\le 1, && \forall i \in B. \\
\sum_{i \in B} x_{ij} &\le 1, && \forall j \in G. \\
x_{ij} &\ge 0, && \forall i \in B, j \in G.
\end{aligned}$$

Dual LP

$$\min \sum_{i \in B} u_i + \sum_{j \in G} p_j$$

s.t.

$$\begin{aligned}
u_i + p_j &\ge w_{ij}, && \forall i \in B, j \in G. \\
u_i &\ge 0, && \forall i \in B. \\
p_j &\ge 0, && \forall j \in G.
\end{aligned}$$

Maximum Weighted Bipartite Matching

Primal LP

max $\sum_{i,j} w_{ij} x_{ij}$

s.t.

$$\begin{aligned} \sum_{j \in G} x_{ij} &\le 1, && \forall i \in B. \\ \sum_{i \in B} x_{ij} &\le 1, && \forall j \in G. \\ x_{ij} &\ge 0, && \forall i \in B, j \in G. \end{aligned}$$

Dual LP

min $\sum_{i \in B} u_i + \sum_{j \in G} p_j$

s.t.

$$\begin{aligned} u_i + p_j &\ge w_{ij}, && \forall i \in B, j \in G. \\ u_i &\ge 0, && \forall i \in B. \\ p_j &\ge 0, && \forall j \in G. \end{aligned}$$

Interpretation of Dual

- p_j is price of good j
- u_i is utility of buyer i
- Complementary Slackness: each buyer grabs his favorite good given prices

2-Player Zero-Sum Games

Rock-Paper-Scissors

	R	P	S
R	0	1	-1
P	-1	0	1
S	1	-1	0

- Two players, row and column
- Game described by matrix A
- When row player plays pure strategy i and column player plays pure strategy j, row player pays column player A_{ij}

2-Player Zero-Sum Games

Rock-Paper-Scissors

	R	P	S
R	0	1	-1
P	-1	0	1
S	1	-1	0

- Two players, row and column
- Game described by matrix A
- When row player plays pure strategy i and column player plays pure strategy j, row player pays column player A_{ij}
- Mixed Strategy: distribution over pure strategies

2-Player Zero-Sum Games

Rock-Paper-Scissors

	R	P	S
R	0	1	-1
P	-1	0	1
S	1	-1	0

- Two players, row and column
- Game described by matrix A
- When row player plays pure strategy i and column player plays pure strategy j, row player pays column player A_{ij}
- Mixed Strategy: distribution over pure strategies
- Assume players know each other's mixed strategies but not coin flips

2-Player Zero-Sum Games

- Assume row player moves first with distribution $y \in \Delta_m$
  - Loss as a function of column's strategy given by $y^\intercal A$
  - A best response by column is pure strategy j maximizing $(y^\intercal A)_j$

	x_1	x_2	x_3	x_4
y_1	a_{11}	a_{12}	a_{13}	a_{14}
y_2	a_{21}	a_{22}	a_{23}	a_{24}
y_3	a_{31}	a_{32}	a_{33}	a_{34}

2-Player Zero-Sum Games

- Assume row player moves first with distribution $y \in \Delta_m$
  - Loss as a function of column's strategy given by $y^\intercal A$
  - A best response by column is pure strategy j maximizing $(y^\intercal A)_j$

Row Moves First

min $\max_j (y^\intercal A)_j$
s.t.
$\sum_{i=1}^m y_i = 1$
$y \geq \vec{0}$

2-Player Zero-Sum Games

- Assume row player moves first with distribution $y \in \Delta_m$
  - Loss as a function of column's strategy given by $y^\intercal A$
  - A best response by column is pure strategy j maximizing $(y^\intercal A)_j$

Row Moves First

min u
s.t.
$u\vec{1} - y^\intercal A \geq \vec{0}$
$\sum_{i=1}^{m} y_i = 1$
$y \geq \vec{0}$

2-Player Zero-Sum Games

- Assume row player moves first with distribution $y \in \Delta_m$
  - Loss as a function of column's strategy given by $y^\intercal A$
  - A best response by column is pure strategy j maximizing $(y^\intercal A)_j$
  - Similarly when column moves first

Row Moves First

min u
s.t.
$u\vec{1} - y^\intercal A \geq \vec{0}$
$\sum_{i=1}^{m} y_i = 1$
$y \geq \vec{0}$

Column Moves First

max v
s.t.
$v\vec{1} - Ax \leq \vec{0}$
$\sum_{j=1}^{n} x_j = 1$
$x \geq \vec{0}$

2-Player Zero-Sum Games

- Assume row player moves first with distribution $y \in \Delta_m$
  - Loss as a function of column's strategy given by $y^\intercal A$
  - A best response by column is pure strategy j maximizing $(y^\intercal A)_j$
  - Similarly when column moves first

Row Moves First

min u
s.t.
$u\vec{1} - y^\intercal A \geq \vec{0}$
$\sum_{i=1}^{m} y_i = 1$
$y \geq \vec{0}$

Column Moves First

max v
s.t.
$v\vec{1} - Ax \leq \vec{0}$
$\sum_{j=1}^{n} x_j = 1$
$x \geq \vec{0}$

These two optimization problems are LP Duals!

Duality and Zero Sum Games

Weak Duality

- $u^* \geq v^*$
- Zero sum games have a second mover advantage

Duality and Zero Sum Games

Weak Duality

- $u^* \geq v^*$
- Zero sum games have a second mover advantage

Strong Duality (Minimax Theorem)

- $u^* = v^*$
- There is no second or first mover advantage in zero sum games with mixed strategies
- Each player can guarantee $u^* = v^*$ regardless of other's strategy.
- y^*, x^* are simultaneously best responses to each other (Nash Equilibrium)

Duality and Zero Sum Games

Weak Duality

- $u^* \geq v^*$
- Zero sum games have a second mover advantage

Strong Duality (Minimax Theorem)

- $u^* = v^*$
- There is no second or first mover advantage in zero sum games with mixed strategies
- Each player can guarantee $u^* = v^*$ regardless of other's strategy.
- y^*, x^* are simultaneously best responses to each other (Nash Equilibrium)

Complementary Slackness

x^* randomizes over pure best responses to y^*, and vice versa.

Saddle Point Interpretation

Consider the matching pennies game

	H	T
H	-1	1
T	1	-1

- Unique equilibrium: each player randomizes uniformly
- If row player deviates, he pays out more
- If column player deviates, he gets paid less

Saddle Point Interpretation

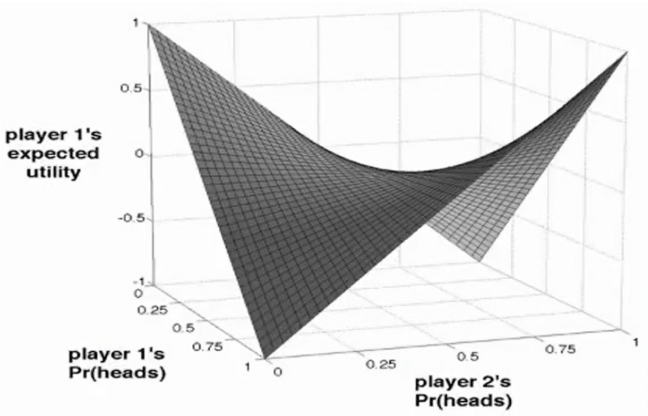

- Unique equilibrium: each player randomizes uniformly
- If row player deviates, he pays out more
- If column player deviates, he gets paid less

Next Lecture

- Begin Convex Optimization Background: Convex Sets