LP Duality: An Example

→ LP Duality is a way of transforming every LP into an equivalent "dual" form.

→ allows short "certificates" of optimality.

Example. Consider the LP

$$\max\ 4x+3y:$$
$$x \geq 0$$
$$y \geq 0$$
$$(1)\quad y \leq 5$$
$$(2)\quad x+y \leq 8$$
$$(3)\quad 2x+y \leq 12.$$

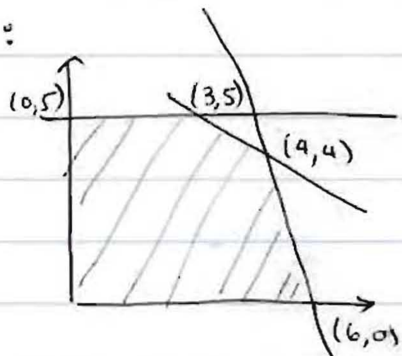

$4\cdot 4+3\cdot 4 \leq 28$

Note that $(x,y)=(4,4)$ gives a solution of value.

Claim. This is optimal; no solution has $4x+3y>28$.

Proof. If $(\hat{x},\hat{y})$ is feasible,

$$2\times(2)\quad 2(x+y)\leq 8\times 2$$
$$(3)\quad 2x+y \leq 12$$
$$2\times(2)+(3)\quad 4x+3y \leq 16+12=28.$$

In general, we can <u>always</u> prove optimality of an LP by combining multiples of constraints in this way. In this example the multipliers were

$$z_1=0,\ z_2=2,\ z_3=1.$$

It is not obvious how to find the <u>best</u> multipliers. In fact finding the best multipliers is another LP:

because we want the <u>best</u> upper bound

$$\min\ \boxed{5z_1+8z_2+12z_3}:$$

value of the upper bound we obtain.

$$z_1, z_2, z_3 \geq 0$$
$$z_2+2z_3 \geq 4$$
$$z_1+z_2+z_3 \geq 3$$

This is the <u>dual</u> LP.

∘ holds for all LPs
∘ proof: "Fourier-Motzkin elim, Farkas' Lemma"

LP Duality Theorem (strong version). The primal LP and dual LP have the <u>same</u> optimal value (if both are feasible).

Zero-Sum Games and Maxminimization

- in this section we'll prove a special case of Nash's theorem
  - every zero-sum 2-player finite static game has a mixed Nash equilibrium.

Def^n A game with vN-M preferences is zero-sum if for all outcomes a,

$$\sum_{i=1}^{n} u_i(a) = 0.$$

Remark Sometimes a game, although it doesn't satisfy the above condition, can be made equivalent to one. Simplest example is a constant-sum game. In homework you'll develop this idea.

Remark For a 2-player zero-sum game we get
$u_2(a) = -u_1(a)$ $\forall$ outcomes a
so it is typical to write only player 1's utility.

E.g. Rock-Paper-Scissors

p1 \ p2	R	P	S
R	0	-1	1
P	1	0	-1
S	-1	1	0

(u_1 shown, $u_2 = -u_1$)

Q. What's special about strat $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ for p1? A. Guarantees no loss.

Let's look at a less familiar example

E.g. "Matching Coins" (Chvátal? 1981? chp. in Math. Rec.)
- two players each privately pick a nickel, dime, or quarter
- both are revealed.
- if they match, p1 wins them, else p2 wins them

p1 \ p2	N	D	Q
N	5	-5	-5
D	-10	10	-10
Q	-25	-25	25

(u_1 shown, $u_2 = -u_1$)

Maxminimization

Q. In Matching Coins can player 2 guarantee some gain?
A. Yes. E.g., if they pick ~~they can~~ make their choice ~~randomly,~~ uniformly at random.

	N	D	Q	$\frac{1}{3}N + \frac{1}{3}D + \frac{1}{3}Q$
N	5	-5	5	5/3
D	-10	10	-10	10/3
Q	-25	-25	25	25/3.

u_2 shown

So whether player 1 uses a deterministic or fixed strategy, player 2 obtains at least 5/3 ¢ in expectation.
(give explanation)

Q. What does this tell us about the best player 1 can do?
A. There is no mixed strategy that can guarantee player 1 any expected utility better than −5/3 ¢.

So p1 is going to lose in this game over the long run, no matter what he does. Can he limit his losses? Consider what happens if he also picks uniformly randomly:

	N	D	Q
N	5	-5	-5
D	-10	10	-10
Q	25	-25	25
$\frac{1}{3}N+\frac{1}{3}D+\frac{1}{3}Q$	-10	$-\frac{20}{3}$	$\frac{10}{3}$

u_1 shown

Hence the strat $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ guarantees a profit of at least −10¢ for player 1, no matter what player 2 does.

But we obviously haven't figured everything out, since there is a gap between
"−10¢ for player 1" and "+5/3 ¢ for player 2"

Idea: instead of using a uniformly random choice for each player, pick "the best" mixed strategy for each player.

In fact, for each player, finding the strategy which guarantees them the most expected utility ~~is called~~ can be done by linear programming.

E.g. If player 1 uses mixed strategy $\alpha_1 = (p_N, p_D, p_Q)$ we get

	N	D	Q	α_1
N	5	-5	-5	$5p_N - 5p_D - 5p_Q$
D	-10	10	-10	$-10p_N + 10p_D - 10p_Q$
Q	-25	-25	25	$-25p_N - 25p_D + 25p_Q$

and the value that α guarantees them, no matter what player 2 does, is $\min\{5p_N - 5p_D - 5p_Q, \ldots\}$.

Hence, what player 1 wants to solve is:

- what values of p_N, p_D, p_Q, subject to $p_N \geq 0, p_D \geq 0, p_Q \geq 0$ and $p_N + p_D + p_Q = 1$, maximizes
$$\min\{5p_N - 5p_D - 5p_Q, -10p_N + 10p_D - 10p_Q, -25p_N - 25p_D + 25p_Q\}$$

- hence called <u>maxminimization</u>.

- we're going to write generic LP to handle <u>any</u> 2-player zero-sum game

- Let $A_1 = \{1, 2, \ldots, m\}$ and $A_2 = \{1, 2, \ldots, n\}$

- LP for Player 1 is

$$\begin{aligned} &\max \quad v && \leftarrow \text{"value" guaranteed} \\ &\text{s.t.} \quad p_i \geq 0 && \text{for } i = 1, \ldots, m \\ (LP_1) \quad &\sum_{i=1}^{m} p_i = 1 && \\ &v \leq \sum_{i=1}^{m} p_i \cdot u_1(i, j) && \text{for } j = 1, \ldots, n. \end{aligned}$$

- E.g., solving the matching coins we got optimal solution

$$p_N^* = \frac{10}{17} \quad p_D^* = \frac{5}{17} \quad p_Q^* = \frac{2}{17} \quad v^* = -\frac{50}{17}$$

so we know that player 1 can protect himself from losing more than $-50/17$ ¢ per game.

On the other hand, we can do the same for player 2.

$$(LP_2) \qquad \begin{array}{ll} \max & u \\ \text{s.t.} & q_j \geq 0 \quad \text{for } j = 1, \dots, n \\ & \sum_{j=1}^{n} q_j = 1 \\ & u \leq \sum_{j=1}^{n} q_j \cdot u_2(i,j) \quad \text{for } i = 1, \dots, m. \end{array}$$

and we get the optimal solution for player 2

$$q_N^* = \frac{7}{34} \quad q_D^* = 12/34 \quad q_Q^* = 15/34 \quad u^* = +\frac{50}{17}.$$

Interesting: $u^* = -v^*$, the gap disappeared!

Q: (guess) what does this tell us about the mixed action profile (p^*, q^*)?

A: it's a NE!

Proof. By the LP, any (pure or mixed) response of player 1 to q^* results in utility at least u^* for player 2. Since $u_1 = -u_2$, any response of player 1 to q^* results in utility at most $-u^*$ for player 1. But $-u^* = v^*$, which player 1 obtains in (p^*, q^*), hence no incentive to deviate.
Similarly player 2 has no incentive to deviate. ☒

Claim. If $u^* = -v^*$ then (p^*, q^*) is a NE.

The "interesting" fact is actually always true! As a result, we get a special case of Nash's theorem:

- every zero-sum 2-player finite game has a Nash eq.

In order to prove it we need to complete the following claim:

<u>Claim.</u> If u^* is the optimal value of LP_1 and v^* is the optimal value of LP_2, then $u^* = -v^*$.

<u>Proof.</u> We're going to use strong LP duality.
First, mechanical step, compute dual

LP_1 max v

$$p_i \geq 0 \qquad i = 1..m$$

$$\sum_{i=1}^{m} p_i = 1 \qquad \leftarrow \text{define var } z$$

$$v - \sum_{i=1}^{m} p_i \cdot u_1(i,j) \leq 0 \qquad j = 1..n \qquad \leftarrow \text{define var } y_j$$

LP_1^* OBJ = min z

$$y_j \geq 0 \qquad j = 1..n$$

$$\sum_{i=1}^{n} y_j = 1 \qquad \leftarrow \overset{\text{unbounded}}{\text{from}} \text{ var } v$$

$$z - \sum_{i=1}^{n} y_j u_1(i,j) \geq 0 \qquad \leftarrow \text{from var } p_i$$

Now a couple more tricks shows LP_1^* and LP_2 are the same:

- change "$u_1(i,j)$" to $-u_2(i,j)$
- replace "$-z$" with $\tilde{z}$
- change "OBJ = min z" to "OBJ = $-\max \tilde{z}$"

then we have LP_1^* and LP_2 are the same under relabeling

$$\tilde{z} \leftrightarrow u$$

$$y_j \leftrightarrow q_j,$$

except that objective in LP_1^* is "$-\max \tilde{z}$" instead of "max u". Hence optimal values of LP_1^* and LP_2 add to zero.

By strong duality, LP_1 and LP_1^* have the same objective value (CHECK: both are feasible).
So optimal values of LP_1 & LP_2 add to zero $\Rightarrow u^* = -v^*$. $\boxtimes$

Remark: looking back at things we could say, with respect to $LP_{1,2}$

mixed strat for player 1 $\Leftrightarrow$ primal solution
value of " " " " " $\equiv$ amt. guaranteed to p1
mixed strat for player 2 $\Leftrightarrow$ dual solution
value of " " " " " $\equiv$ -amt. guaranteed to p2.

and then LP strong duality says
max value of primal sol^n = min value of dual sol^n.

Claim + Claim $\Rightarrow$ Th^m (1926, von Neumann), zero sum 2p finite strategic games have an NE $\boxtimes$