Linear Programming Duality

In this lecture we discuss the general notion of Linear Programming *Duality*, a powerful tool that can allow us to solve some linear programs easier, gain theoretical insights into the properties of a linear program, and has many more applications that we might see later in the course. We will show how duality connects to some topics we have already seen, like minimax optimal strategies in zero-sum games.

Objectives of this lecture

In this lecture, we will

- Motivate and define the idea of the dual of a linear program
- See how to convert a linear program into its dual program
- Learn some powerful theorems that tell us about the behavior of a linear program and its dual
- See how duality can teach us about minimax optimal strategies for zero-sum games

1 The Dual Program

Consider the following LP which is written in standard form.

$$\begin{aligned} \text{maximize } & 2x_1 + 3x_2 \\ \text{s.t. } & 4x_1 + 8x_2 \le 12 \\ & 2x_1 + x_2 \le 3 \\ & 3x_1 + 2x_2 \le 4 \\ & x_1, x_2 \ge 0 \end{aligned} \tag{1}$$

Here it is in a diagram which shows each constraint, the feasible region shaded in gray, and the objective direction as a red arrow.

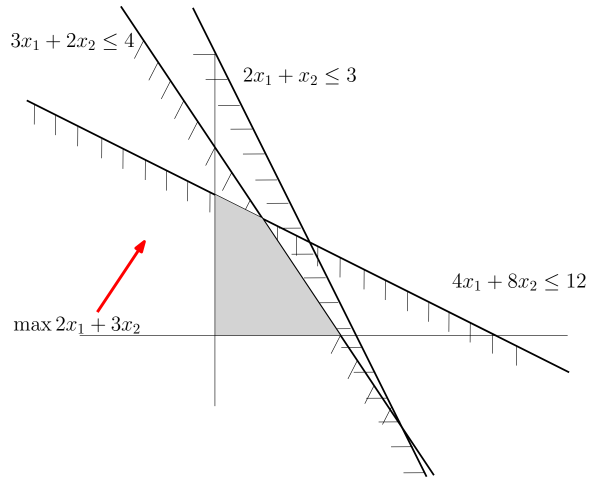

Rather than try to solve the LP using an algorithm directly, we are going to do an experiment. Lets see if we can figure out some bounds on the objective value, and use those to hone in on the optimal value. How can we bound the objective? Well the only other information that we have are the constraints, so lets use them! Lets refer to the optimal objective value as OPT.

- First, since $x_1, x_2 \geq 0$, we can notice that the left-hand equation of the first constraint ($4x_1 + 8x_2$) must be bigger than the objective ($2x_1 + 3x_2$), in other words, we can write

$$\underbrace{2x_1 + 3x_2}_{\text{objective function}} \leq \underbrace{4x_1 + 8x_2}_{\text{first constraint}} \leq 12$$

  Note that the left-hand side is the objective function and the right hand side is the first constraint. So we therefore know that OPT ≤ 12. That's a start. Can we get a tighter bound?

- Yes, the first constraint is more than double the objective function, so we can write a tighter bound by using just half of the constraint

$$\underbrace{2x_1 + 3x_2}_{\text{objective function}} \leq \underbrace{\tfrac{1}{2}(4x_1 + 8x_2) \leq 6}_{\text{half of the first constraint}}$$

  This gives us a bound of OPT ≤ 6. Can we do better? Maybe by combining multiple constraints!

- We want to combine some constraints such that we get as close to a coefficient of 2 for x_1 and a coefficient of 3 for x_2. By inspection, we can see that if we add the first and second constraint, we will have $6x_1 + 9x_2$, which is exactly three times our objective function, so lets try one third of that combination.

$$\underbrace{2x_1 + 3x_2}_{\text{objective function}} \leq \underbrace{\tfrac{1}{3}((4x_1 + 8x_2) + (2x_1 + x_2)) \leq 5}_{\text{one third of the first two constraints}}$$

Note that we get 5 on the right-hand side because we sum the right-hand sides of the original constraints to get $12+3=15$, then take one third of it. So we know that OPT ≤ 5.

In each of these cases we take a positive linear combination of the constraints, looking for better and better bounds on the maximum possible value of $2x_1+3x_2$. Why positive? Because if we multiply by a negative value, the sign of the inequality changes.

1.1 The tightest possible bound

After playing with this experiment for a bit, the natural question that arises is how do we find the tightest lower bound that can be achieved with this method, i.e., by writing down a a linear combination of the constraints? This is just another algorithmic problem, and we can systematically solve it, by letting y_1, y_2, y_3 be the (unknown) coefficients of our linear combination. So our goal is to write the combination like so

$$y_1(4x_1+8x_2)+y_2(2x_1+x_2)+y_3(3x_1+2x_2) \leq 12y_1+3y_2+4y_3,$$

such that the value on the right-hand side is as small as possible. This sounds like just another linear program! To bound the original objective function, we require that the coefficients of x_1 add up to at least 2, and the coefficients of x_2 add up to at least 3. We can write these requirements down as a linear program.

$$\begin{aligned} \text{minimize}\ \ & 12y_1+3y_2+4y_3 \\ \text{s.t.}\ \ & 4y_1+2y_2+3y_3 \geq 2 \\ & 8y_1+y_2+2y_3 \geq 3 \\ & y_1, y_2, y_3 \geq 0 \end{aligned} \tag{2}$$

This is indeed an LP! We refer to this LP (2) as the "**dual**" and the original LP (1) as the "**primal**". We designed the dual to serve as a method of constructing an upper bound on the optimal value of the primal, so if y is a feasible solution for the dual and x is a feasible solution for the primal, then $2x_1+3x_2 \leq 12y_1+3y_2+4y_3$.

This serves as an upper bound, but what happens if we make it tight? If we can find two feasible solutions $\mathbf{x}$ and $\mathbf{y}$, that make these equal, then we know we have found the provably optimal values of these LPs. In this case the feasible solutions $x_1=\frac{1}{2}, x_2=\frac{5}{4}$ and $y_1=\frac{5}{16}, y_2=0, y_3=\frac{1}{4}$ give us a value and matching upper bound of 4.75, which therefore must be the optimal value.

1.2 An alternate motivating example – the carpenter

Suppose you are a humble carpenter; you spend your days making tables, chairs, and shelves, all out of wood, nails, and paint. Each item you make requires a specific amount of each of the three materials and can be sold at the market for a specific price. Given the amount of material you have, you would like to determine the best way to use them to make the most money.

Item	Wood	Nails	Paint	Sale Price
Table	8	20	5	\$50
Chair	4	15	3	\$30
Shelf	3	5	3	\$20
Stock	100	300	80	

Ignoring rounding issues, since we don't want to deal with integrality, we can write this problem as a linear program. This leads us to the following. Lets define the variables x, y, z to be the number of tables, chairs, and shelves, respectively, that we should make.

$$\begin{aligned} \text{maximize } & 50x + 30y + 20z \\ \text{subject to } & 8x + 4y + 3z \le 100 \\ & 20x + 15y + 5z \le 300 \\ & 5x + 3y + 3z \le 80 \\ & x, y, z \ge 0 \end{aligned}$$

If we use a computer program to solve this, we will find that an optimal solution is $x \approx 1.82$, $y \approx 14.55$, $z \approx 9.09$, with an objective value of \$709.09.

Along comes a merchant A merchant approaches you, the humble carpenter, with the prospect of buying your materials, i.e., all of your wood, nails, and paint. You consider selling them at a fair price, but what should that be? You'd like to make sure that you sell them for at least as much as you could make if you turned them into furniture, but you know that the merchant will try to get the lowest price from you, so you'll settle for that.

If we denote by w, s, p, the price of wood, nails, and paint, respectively, we can model this problem as another LP as follows:

$$\begin{aligned} \text{minimize } & 100w + 300s + 80p \\ \text{subject to } & 8w + 20s + 5p \ge 50 \\ & 4w + 15s + 3p \ge 30 \\ & 3w + 5s + 3p \ge 20 \\ & w, s, p \ge 0 \end{aligned}$$

This LP has an optimal solution of $w \approx 2.73$, $s \approx 0.73$, $p \approx 2.73$, and an objective value of \$709.09. What a coincidence! The objective value is the same as the previous LP. Perhaps that should not be so surprising... we were not willing to sell the materials for less than we could turn them into items, so we would expect it to be *at least as much*, but the fact that they are *exactly equal* is not quite so obvious in advance...

The structure of this pair of LPs is very special, and if we look closely at them we will see that they are made up of the same exact ingredients, just laid out a little differently. Since the first LP is in standard form, we can write it in matrix form with

$$A = \begin{bmatrix} 8 & 4 & 3 \\ 20 & 15 & 5 \\ 5 & 3 & 3 \end{bmatrix}, \qquad \mathbf{b} = \begin{bmatrix} 100 \\ 300 \\ 80 \end{bmatrix}, \qquad \mathbf{c} = \begin{bmatrix} 50 \\ 30 \\ 20 \end{bmatrix}, \qquad \mathbf{x} = \begin{bmatrix} x \\ y \\ z \end{bmatrix}$$

This way, the LP can be written as

$$\begin{aligned} \text{maximize } & \mathbf{c}^T \mathbf{x} \\ \text{subject to } & A\mathbf{x} \le \mathbf{b} \\ & \mathbf{x} \ge \mathbf{0}. \end{aligned} \tag{3}$$

Looking closely at the pricing LP, we see that if we define the variable vector $\mathbf{y} = [w \;\; s \;\; p]^T$, then it can be written as

$$\begin{aligned} \text{minimize } & \mathbf{b}^T \mathbf{y} \\ \text{subject to } & A^T \mathbf{y} \geq \mathbf{c} \\ & \mathbf{y} \geq \mathbf{0}. \end{aligned} \tag{4}$$

We happened to derive this particular LP by intuition, but in fact, given *any LP in standard form*, one could apply this transformation to obtain this second program. It turns out to be a wildly useful and powerful concept, so it has a name – its called the dual program!

2 A General Formulation of the Dual

Definition 1: The dual of a linear program

The dual of the standard form LP (3) is

$$\begin{aligned} \text{minimize } & \mathbf{b}^T \mathbf{y} \\ \text{subject to } & A^T \mathbf{y} \geq \mathbf{c} \\ & \mathbf{y} \geq \mathbf{0}. \end{aligned}$$

The original standard form LP (3) is referred to as the *primal problem*.

And if you take the dual of (4) , what do you think you will get back? You'll get (3). *The dual of the dual is the primal.* Because of this, which program we refer to as the primal and which we refer to as the dual is just a matter of convention, it is completely symmetric. Think about how you would actually take the dual of the dual as an exercise. Since the dual as written is not in standard form, it would need to first be converted to standard form.

2.1 The Theorems

Our intuitive derivation of the dual program as a pricing problem for the carpenter implied that the value of the dual solutions should be at least as large as the profit the carpenter could make from turning the materials into furniture, i.e., it should always give at least as large of a value as the primal problem.

We can formally prove that it indeed always does just that. This fact is called *weak duality.*

Theorem 1: Weak Duality

If $\mathbf{x}$ is a feasible solution to the primal (3) and $\mathbf{y}$ is a feasible solution to the dual (4) then

$$\mathbf{c}^T \mathbf{x} \leq \mathbf{b}^T \mathbf{y}.$$

Proof. This follows by applying the constraints of the primal and dual LPs in (3) and (4) and the fact that $\mathbf{x} \geq 0$ and $\mathbf{y} \geq 0$. Since $A^T\mathbf{y} \geq \mathbf{c}$, we can plug this into the objective $\mathbf{c}^T\mathbf{x}$ and get

$$\mathbf{c}^T\mathbf{x} \leq (A^T y)^T\mathbf{x} = (\mathbf{y}^T A)\mathbf{x}$$

Now we can move the brackets (associativity), and use the fact that $A\mathbf{x} \leq \mathbf{b}$, to get

$$(\mathbf{y}^T A)\mathbf{x} = \mathbf{y}^T(A\mathbf{x}) \leq \mathbf{y}^T\mathbf{b} = \mathbf{b}^T\mathbf{y}.$$

□

The amazing (and deep) result here is to show that the dual actually gives not just an upper bound on the primal, but, assuming some mild conditions, it perfectly equals the primal!

Theorem 2: Strong Duality Theorem

Suppose the primal LP (3) is feasible (i.e., it has at least one solution) and bounded (i.e., the optimal value is not ∞). Then the dual LP (4) is also feasible and bounded. Moreover, if $\mathbf{x}^*$ is the optimal primal solution, and $\mathbf{y}^*$ is the optimal dual solution, then

$$\mathbf{c}^T\mathbf{x}^* = \mathbf{b}^T\mathbf{y}^*.$$

In other words, the maximum of the primal equals the minimum of the dual.

We will not prove Theorem 2 in this course, since the proof is a bit long, though it isn't too difficult (feel free to look it up if interested). Why is this useful? If I wanted to prove to you that $\mathbf{x}^*$ was an optimal solution to the primal, I could give you the solution $\mathbf{y}^*$, and you could check that $\mathbf{x}^*$ was feasible for the primal, $\mathbf{y}^*$ feasible for the dual, and they have equal objective function values.

This relationship is like in the case of s-t flows: the max flow equals the minimum cut. Or like in the case of zero-sum games: the payoff for the optimal strategy of the row player equals the (negative) of the payoff of the optimal strategy of the column player. Indeed, both these things are just special cases of strong duality!

2.2 Using duality to determine feasibility and boundedness

In addition to helping us bound feasible solutions to our LPs, duality can also be used as a tool to determine when certain programs are feasible or infeasible, or perhaps show that they are bounded or unbounded.

- If the primal is feasible and bounded, strong duality says the dual is also feasible and bounded.
- Suppose the primal (maximization) problem is unbounded. What can duality tell us? Weak duality says $\mathbf{c}^T x \leq \mathbf{b}^T y$... If there existed any feasible $\mathbf{y}$ for the dual, this would imply that the primal is bounded, and hence by the contrapositive, if the primal is unbounded, then the dual *must be infeasible*.

- By the exact same logic (reversed), if the dual is unbounded, since the primal is a lower bound on the dual, the primal must be infeasible.
- Can both the primal and dual be unbounded? No, because as the two previous points show, if one of them is unbounded, then the other is infeasible, and if a program is infeasible, it certainly can not be unbounded.

We can use these facts to represent all of the possible situations in a table like so:

		Dual		
		Inf	**F&B**	**Unb**
Primal	**Inf**	✓	X	✓
	F&B	X	✓	X
	Unb	✓	X	X

Here, **Inf** means infeasible, **F&B** means feasible and bounded, and **Unb** means unbounded. The only scenario that duality does not cover for us is the top-left cell. You can figure that out as an exercise.

Remark: Usefulness

This table has some very useful implications. If we have an LP for some problem, we might want to prove conditions on when it is feasible or infeasible. Directly proving that the LP is infeasible might be too difficult. Instead, if we can write the dual program and give a proof that the dual is unbounded, then we have indirectly proven that the primal is infeasible! A useful trick.

3 Example: Zero-Sum Games

Consider a 2-player zero-sum game defined by an n-by-m payoff matrix R for the row player. To simplify things a bit, let's assume that all entries in R are positive (this is without loss of generality since as pre-processing we can always translate values by a constant and this will just change the game's value to the row player by that constant). The *lower bound* for the row player was defined to be

$$\mathsf{lb}^* = \max_{\mathbf{p}} \min_j \sum_i p_i R_{ij},$$

We can solve for the lower bound by writing it as an LP.

Variables: $p_1,\dots,p_n$ and v.
Objective: Maximize v.
Constraints:

- $p_i \ge 0$ for all $1 \le i \le n$,
- $\sum_{i=1}^{n} p_i = 1.$ (the p_i form a probability distribution)
- $\sum_{i=1}^{n} p_i R_{ij} \ge v$ for all columns $1 \le j \le m$

To apply our techniques for duality, we need to put it in standard form. To do so, we can do the following:

- we can replace $\sum_i p_i = 1$ with $\sum_i p_i \le 1$ since we said that all entries in R are positive, so the maximum will occur with $\sum_i p_i = 1$,
- since all entries in R are positive, we can also safely add in the constraint $v \ge 0$,
- we can also rewrite the third set of constraints as $v - \sum_i p_i R_{ij} \le 0$.

This then gives us an LP in the form of (3) with

$$\mathbf{x} = \begin{bmatrix} v \\ p_1 \\ p_2 \\ \dots \\ p_n \end{bmatrix}, \mathbf{c} = \begin{bmatrix} 1 \\ 0 \\ 0 \\ \dots \\ 0 \end{bmatrix}, \mathbf{b} = \begin{bmatrix} 0 \\ 0 \\ \dots \\ 0 \\ 1 \end{bmatrix}, \text{and } A = \left[\begin{array}{c|ccc} 1 & & & \\ 1 & & -R^T & \\ \dots & & & \\ 1 & & & \\ \hline 0 & 1 & \dots & 1 \end{array}\right].$$

i.e., maximizing $\mathbf{c}^T\mathbf{x}$ subject to $A\mathbf{x} \le \mathbf{b}$ and $\mathbf{x} \ge \mathbf{0}$.

We can now write the dual, following (4). Let $\mathbf{y}^T = (q_1, q_2, \dots, q_m, v')$. We now are asking to minimize $\mathbf{b}^T\mathbf{y}$ subject to $A^T\mathbf{y} \ge \mathbf{c}$ and $\mathbf{y} \ge \mathbf{0}$. Writing out the objective function, we get $\mathbf{b}\cdot\mathbf{y} = [0\ 0\ \cdots\ 0\ 1]\cdot[q_1\ q_2\ \cdots\ q_m\ v'] = v'$, so the objective is just minimize v'. If we transpose A, we get

$$A^T = \left[\begin{array}{ccc|c} 1 & \cdots & 1 & 0 \\ \hline & & & 1 \\ & -R & & \vdots \\ & & & 1 \end{array}\right]$$

Now we can write out the constraints, we have

1. $q_1 + \ldots + q_m \ge 1$,
2. $-q_1 R_{i1} - q_2 R_{i2} - \ldots - q_m R_{im} + q_{m+1} \ge 0$ for all rows i,

Since the objective is to minimize, and all R entries are positive, we can make a similar argument to the primal case that $q_1 + \ldots + q_m = 1$, i.e., this constraint must be tight because increasing any of the q values would only further increase the objective. Some algebra turns the second constraint into $q_1 R_{i1} + q_2 R_{i2} + \ldots + q_m R_{im} \le v'$ for all rows i, so we obtain the LP:

Variables: $q_1, \ldots, q_m$ and v'.
Objective: Minimize v'.
Constraints:

- $q_i \geq 0$ for all $1 \leq i \leq m$,
- $\sum_{i=1}^{m} q_i = 1$.
- $\sum_{j=1}^{m} q_j R_{ij} \leq v'$ for all rows $1 \leq i \leq n$

This LP look an awful lot similar to the primal LP, which was computing lb^* for the row player. What is this LP saying? We can interpret v' as being the value of the game to the row player once again, and $q_1, \ldots, q_m$ as the randomized strategy of the *column player* this time, and we want to find a randomized strategy for the column player that minimizes v' subject to the constraint that the row player gets *at most* v' no matter what row he plays. In other words, we've just found an LP for the *upper bound* ub^* to the row player!

Notice that the fact that the maximum value of v in the primal is equal to the minimum value of v' in the dual follows from strong duality. Therefore, the minimax theorem is a corollary to the strong duality theorem!

Corollary 1: Minimax Theorem

Given a finite 2-player zero-sum game with payoff matrices $R = -C$,

$$\mathsf{lb}^* = \max_{\mathbf{p}} \min_{\mathbf{q}} V_R(\mathbf{p}, \mathbf{q}) = \min_{\mathbf{q}} \max_{\mathbf{p}} V_R(\mathbf{p}, \mathbf{q}) = \mathsf{ub}^*.$$

This common value is called the value of the game.

Proof. Follows from strong duality and the argument above, where we showed that the dual problem to computing lb^* is a linear program that computes ub^*. □

4 The Geometric Intuition for Strong Duality

Optional content — Not required knowledge for the exams

To give a geometric view of the strong duality theorem, consider an LP of the following form:

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

For concreteness, let's take the following 2-dimensional LP:

$$\begin{aligned} \text{maximize } & x_2 \\ \text{subject to } & -x_1 + 2x_2 \le 3 \\ & x_1 + x_2 \le 2 \\ & -2x_1 + x_2 \le 4 \\ & x_1, x_2 \ge 0 \end{aligned}$$

If $\mathbf{c} := (0,1)$, then the objective function wants to maximize $\mathbf{c} \cdot \mathbf{x}$, i.e., to go as far up in the vertical direction as possible. As we have already argued before, the optimal point $\mathbf{x}^*$ must be obtained at the intersection of two constraints for this 2-dimensional problem (n tight constraints for n dimensions). In this case, these happen to be the first two constraints.

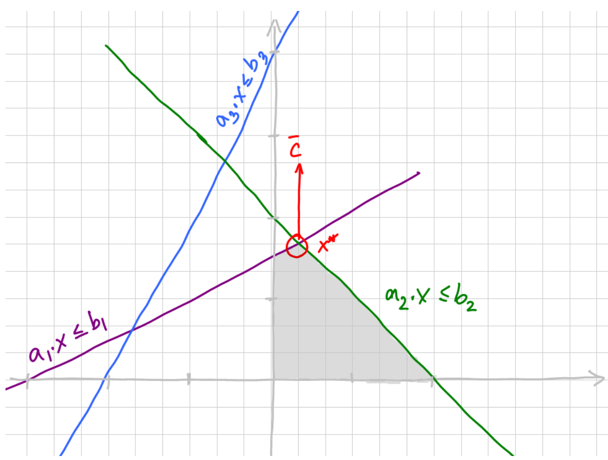

If $\mathbf{a}_1 = (-1,2)$, $b_1 = 3$ and $\mathbf{a}_2 = (1,1)$, $b_2 = 2$, then $\mathbf{x}^*$ is the (unique) point $\mathbf{x}$ satisfying both $\mathbf{a}_1 \cdot \mathbf{x} = b_1$ and $\mathbf{a}_2 \cdot \mathbf{x} = b_2$. Indeed, we're being held down by these two constraints. Geometrically, this means that $\mathbf{c} = (0,1)$ lies "between" these the vectors $\mathbf{a}_1$ and $\mathbf{a}_2$ that are normal (perpendicular) to these constraints.

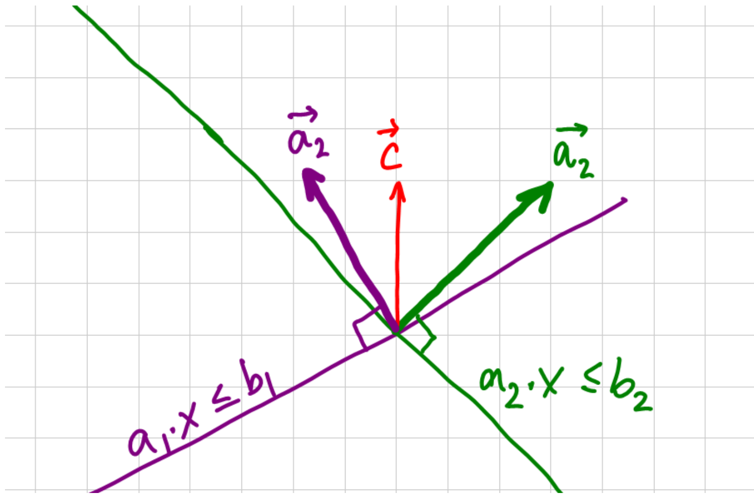

Consequently, $\mathbf{c}$ can be written as a positive linear combination of $\mathbf{a}_1$ and $\mathbf{a}_2$. (It "lies in the cone formed by $\mathbf{a}_1$ and $\mathbf{a}_2$.") I.e., for some positive values y_1 and y_2,

$$\mathbf{c} = y_1 \mathbf{a}_1 + y_2 \mathbf{a}_2.$$

Great. Now, take dot products on both sides with $\mathbf{x}^*$. We get

$$\begin{aligned}
\mathbf{c} \cdot \mathbf{x}^* &= (y_1\,\mathbf{a}_1 + y_2\,\mathbf{a}_2) \cdot \mathbf{x}^* \\
&= y_1(\mathbf{a}_1 \cdot \mathbf{x}^*) + y_2(\mathbf{a}_2 \cdot \mathbf{x}^*) \\
&= y_1 b_1 + y_2 b_2
\end{aligned}$$

Defining $\mathbf{y} = (y_1, y_2, 0, \ldots, 0)$, we get

$$\text{optimal value of primal} = \mathbf{c} \cdot \mathbf{x}^* = \mathbf{b} \cdot \mathbf{y} \geq \text{value of dual solution } \mathbf{y}.$$

The last inequality follows because

- the $\mathbf{y}$ we found satisfies $\mathbf{c} = y_1 \mathbf{a}_1 + y_2 \mathbf{a}_2 = \sum_i y_i \mathbf{a}_i = A^T \mathbf{y}$, and hence $\mathbf{y}$ satisfies the dual constraints $\mathbf{y}^T A \geq \mathbf{c}^T$ by construction.

In other words, $\mathbf{y}$ is a feasible solution to the dual, has value $\mathbf{b} \cdot \mathbf{y} \leq \mathbf{c} \cdot \mathbf{x}^*$. So the *optimal* dual value cannot be less. Combined with weak duality (which says that $\mathbf{c} \cdot \mathbf{x}^* \leq \mathbf{b} \cdot \mathbf{y}$), we get strong duality

$$\mathbf{c} \cdot \mathbf{x}^* = \mathbf{b} \cdot \mathbf{y}.$$

Above, we used that the optimal point was constrained by two of the inequalities (and that these were not the non-negativity constraints). The general proof is similar: for n dimensions, we just use that the optimal point is constrained by n tight inequalities, and hence $\mathbf{c}$ can be written as a positive combination of n of the constraints (possibly some of the non-negativity constraints too).

Exercises: Linear Programming Duality

Problem 1. Using the definitions we gave, show that the dual of the dual is the primal problem.

Problem 2. Find an LP that is infeasible such that its dual is also infeasible.

Problem 3. **(Hard - optional)** Prove the min-cut max-flow theorem using strong duality.