

This lecture (Fall 2016: this is an optional lecture)

Linear programming duality + robust linear programming

- Intuition behind the derivation of the dual
- Weak and strong duality theorems
 - Max-flow=Min-cut
- Primal/dual possibilities
- Interpreting the dual
- An application: robust linear programming

 TAs:
B. El Khadir,
G. Hall,
Z. Li,
K. Wang,
J. Ye,
J. Zhang

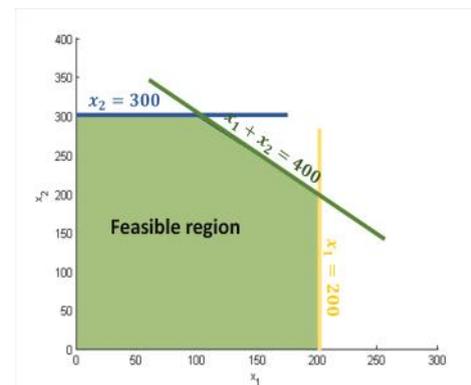
The idea behind duality

- For any linear program (LP), there is a closely related LP called the *dual*. The feasible and optimal solutions of the dual provide very useful information about the original (aka *primal*) LP.
- In particular, if the primal LP is a maximization problem, the dual can be used to find upper bounds on its optimal value. (Similarly, if the primal is a minimization problem, the dual gives lower bounds.) Note that this is useful in certifying optimality of a candidate solution to the primal. One does not get information of this type from feasible solutions to the primal alone.

Let's understand this concept through an example first.

Consider this LP (from [DPV08]):

$$\begin{aligned}
 &\max. x_1 + 6x_2 \\
 &\text{s.t. } x_1 \leq 200 \quad (1) \\
 &\quad x_2 \leq 300 \quad (2) \\
 &\quad x_1 + x_2 \leq 400 \quad (3) \\
 &\quad x_1, x_2 \geq 0
 \end{aligned}$$



- Somebody comes to you and claims $\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} 100 \\ 300 \end{pmatrix}$ with objective value 1900 is optimal.
- How can we check his claim? Well, suppose we combine the constraints to produce new "valid inequalities" in a way that they upper bound the objective function when evaluated on the feasible set. For example,

$$(1) + 6(2) \Rightarrow x_1 + 6x_2 \leq 2000.$$
- This means that it is impossible for the objective function to take value larger than 2000 when evaluated at a feasible point (why?). Can we bring down this upper bound further?

Lec12p2, ORF363/COS323

- What if we try

$$0(1) + 5(2) + 1(3) \Rightarrow x_1 + 6x_2 \leq 1900.$$

This indeed shows that $\begin{pmatrix} x_1 \\ x_2 \end{pmatrix}$ is optimal. The coefficients $(0,5,1)$ are called the "*dual multipliers*".

Let's systematize what we did in this example. Start by introducing a multiplier for each constraint:

$$\begin{array}{rcl} x_1 \leq 200 & (1) & y_1 \\ x_2 \leq 300 & (2) & y_2 \\ x_1 + x_2 \leq 400 & (3) & y_3 \end{array}$$

- First, we need $y_1, y_2, y_3 \geq 0$ to preserve the inequalities after multiplication.
- After we multiply and add, we obtain a new valid inequality of the form

$$(y_1 + y_3)x_1 + (y_2 + y_3)x_2 \leq 200y_1 + 300y_2 + 400y_3.$$

We need the left hand side (and hence the right hand side) to be an upper bound on the objective function $x_1 + 6x_2$. This can be achieved by enforcing

$$\begin{array}{l} y_1 + y_3 \geq 1 \\ y_2 + y_3 \geq 6. \end{array}$$

Indeed, this implies our desired upper bound as $x_i \geq 0$ for $i = 1, 2$.

- Finally, we want to get the best possible upper bound which means that we want to minimize $200y_1 + 300y_2 + 400y_3$.
- Altogether, this gives us the following linear programming problem

$$\begin{array}{l} \min. 200y_1 + 300y_2 + 400y_3 \\ \text{s.t. } y_1 + y_3 \geq 1 \\ y_2 + y_3 \geq 6 \\ y_1, y_2, y_3 \geq 0 \end{array}$$

This problem is called the dual LP!

The optimal solution to the primal is $\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} 100 \\ 300 \end{pmatrix}$ with optimal value 1900.

The optimal solution to the dual is $\begin{pmatrix} y_1 \\ y_2 \\ y_3 \end{pmatrix} = \begin{pmatrix} 0 \\ 5 \\ 1 \end{pmatrix}$ with optimal value 1900. Bingo!

They match!

More generally, if the primal LP (P) is of the form

$$\begin{aligned} \max. \quad & c^T x \\ \text{s.t.} \quad & Ax \leq b \\ & x \geq 0 \end{aligned}$$

where $A \in \mathbb{R}^{m \times n}$, $b \in \mathbb{R}^m$ and $c \in \mathbb{R}^n$, then the dual LP (D) will be of the form

$$\begin{aligned} \min. \quad & b^T y \\ \text{s.t.} \quad & A^T y \geq c \\ & y \geq 0. \end{aligned}$$

Weak duality and strong duality

- The following are the two fundamental theorems of duality. We only prove the weak version.

Weak duality: If x is feasible for (P) and y is feasible for (D), then $b^T y \geq c^T x$.

Strong duality: If (P) has a finite optimal value, then so does (D) and the two optimal values coincide.

Proof of weak duality:

$$\begin{aligned} b^T y = y^T b &\geq y^T Ax = x^T A^T y \geq x^T c = c^T x. \\ &\quad \uparrow \qquad \qquad \qquad \uparrow \\ & y \geq 0, Ax \leq b \quad x \geq 0, A^T y \geq c \qquad \qquad \square \end{aligned}$$

- The Primal/Dual pair can appear in many other forms, e.g., in standard form. Duality theorems hold regardless.

$$\begin{array}{ll} \text{(P)} & \min. \quad c^T x \\ & Ax = b \\ & x \geq 0 \end{array} \qquad \qquad \begin{array}{ll} \text{(D)} & \max. \quad y^T b \\ & A^T y \leq c \end{array}$$

Proof of weak duality in this form:

$$\begin{aligned} c^T x = x^T c &\geq x^T A^T y = y^T Ax = y^T b \\ &\quad \uparrow \qquad \qquad \qquad \uparrow \\ A^T y \leq c, x \geq 0 & \qquad \qquad Ax = b \end{aligned}$$

□

Lec12p4, ORF363/COS323

- Figure 7.11 from [DPV08] gives a general recipe for constructing the dual from the primal.

Figure 7.11 In the most general case of linear programming, we have a set I of inequalities and a set E of equalities (a total of $m = |I| + |E|$ constraints) over n variables, of which a subset N are constrained to be nonnegative. The dual has $m = |I| + |E|$ variables, of which only those corresponding to I have nonnegativity constraints.

Primal LP:

$$\begin{aligned} \max \quad & c_1x_1 + \cdots + c_nx_n \\ \text{s.t.} \quad & a_{i1}x_1 + \cdots + a_{in}x_n \leq b_i \quad \text{for } i \in I \\ & a_{i1}x_1 + \cdots + a_{in}x_n = b_i \quad \text{for } i \in E \\ & x_j \geq 0 \quad \text{for } j \in N \end{aligned}$$

Dual LP:

$$\begin{aligned} \min \quad & b_1y_1 + \cdots + b_my_m \\ \text{s.t.} \quad & a_{1j}y_1 + \cdots + a_{mj}y_m \geq c_j \quad \text{for } j \in N \\ & a_{1j}y_1 + \cdots + a_{mj}y_m = c_j \quad \text{for } j \notin N \\ & y_i \geq 0 \quad \text{for } i \in I \end{aligned}$$

Observation: The dual of the dual is the primal.

Proof: We consider the following forms of the primal and the dual (the same result is true no matter what form you work with):

$$\begin{array}{ll} \text{(P)} & \begin{aligned} \max. \quad & c^T x \\ \text{s.t.} \quad & Ax \leq b \\ & x \geq 0 \end{aligned} \\ \text{(D)} & \begin{aligned} \min. \quad & b^T y \\ \text{s.t.} \quad & A^T y \geq c \\ & y \geq 0 \end{aligned} \end{array}$$

Using simple transformations, we find that (D) is equivalent to (D') where

$$\begin{aligned} \text{(D')} \quad & \begin{aligned} -\max. \quad & -b^T y \\ \text{s.t.} \quad & -A^T y \leq -c \\ & y \geq 0 \end{aligned} \end{aligned}$$

As (D') is in the "form" of (P), we can easily take its dual to obtain:

$$\begin{aligned} & \begin{aligned} -\min. \quad & -c^T z \\ \text{s.t.} \quad & -Az \geq -b \\ & z \geq 0 \end{aligned} \end{aligned}$$

Notice that this last LP is equivalent to (P).

Primal/dual possibilities

Again, we consider the following forms of the primal and dual.

$$\begin{array}{ll}
 \text{(P)} & \max. c^T x \\
 & \text{s.t. } Ax \leq b \\
 & \quad x \geq 0 \\
 \text{(D)} & \min. b^T y \\
 & \text{s.t. } A^T y \geq c \\
 & \quad y \geq 0
 \end{array}$$

We know from the previous lecture that an LP can have three possibilities: either it has a finite optimum, or it is unbounded, or it is infeasible. Here are the possibilities that we can have when we consider a primal/dual pair:

		PRIMAL		
		Finite optimal	Unbounded	Infeasible
DUAL	Finite optimal	Possible	Impossible	Impossible
	Unbounded	Impossible	Impossible	Possible
	Infeasible	Impossible	Possible	Possible

Using only strong and weak duality, can you explain each entry of the table?

Figure 7.9 By design, dual feasible values \geq primal feasible values. The duality theorem tells us that moreover their optima coincide.

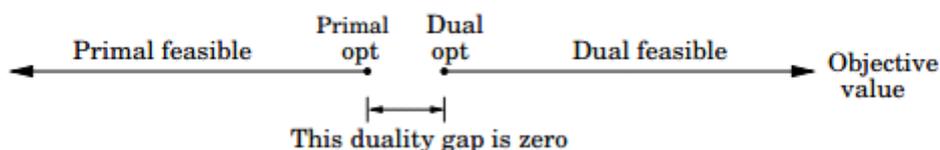


Image credit [DPV08]

An interpretation of the dual in a classic example

The example we will be looking at here is taken from [Wri05].

- A student wants to purchase a snack from a bakery to meet certain dietary requirements by choosing the best combination of brownies and cheesecake. (The brownies are not from Amsterdam or else the decision would be obvious...)
- The characteristics of each product are given in the following table.

Lec12p6, ORF363/COS323

- The student is following some new diet trend which requires her to eat at least 6oz of chocolate, 8oz of cream cheese, and 10oz of sugar.
- Her goal is to satisfy these requirements at minimal cost.

Ingredients needed				
	3 oz	2 oz	2 oz	50 cts
	0 oz	4 oz	5 oz	80 cts
Requirements	6 oz	10 oz	8 oz	

This problem (also known as the diet problem) can be solved using LP. Here, the decision variables are x_1 , the amount of brownies, and x_2 , the amount of cheesecake that the student decides to purchase.

The primal problem is then:

$$\begin{aligned}
 & \min_{x_1, x_2} 50x_1 + 80x_2 \\
 & \text{s.t. } 3x_1 \geq 6 \\
 & \quad 2x_1 + 4x_2 \geq 10 \\
 & \quad 2x_1 + 5x_2 \geq 8 \\
 & \quad x_1, x_2 \geq 0
 \end{aligned}$$

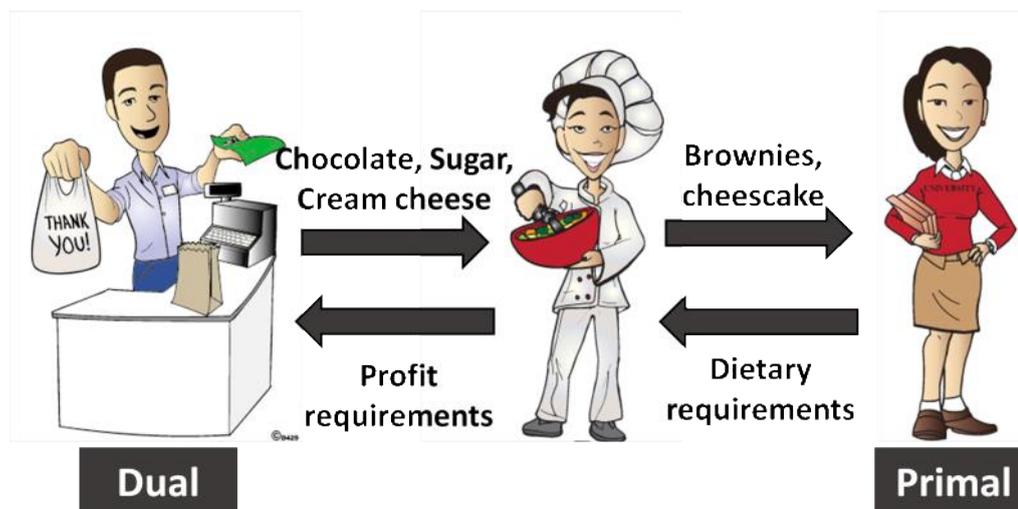
The dual problem is given by

$$\begin{aligned}
 & \max 6y_1 + 10y_2 + 8y_3 \\
 & \text{s.t. } y_1, y_2, y_3 \geq 0 \\
 & \quad 3y_1 + 2y_2 + 2y_3 \leq 50 \\
 & \quad 4y_2 + 5y_3 \leq 80
 \end{aligned}$$

For this example, the dual has a nice interpretation.
(In fact, the dual almost always has a nice interpretation.)

Lec12p7, ORF363/COS323

Let's consider the problem from the point of view of a grocery store that provides the baker with the required ingredients. We denote by y_1, y_2, y_3 the prices of chocolate, sugar, and cream cheese respectively. These prices are nonnegative and the store wants to know how it should set them.



Figures from <http://myjollyfamily.com/>

- The grocery store knows that the baker will only buy the ingredients if she is sure of making a profit on the item she sells. In other words, the sum of the cost of the ingredients should not exceed the price of the product. This leads to the constraints for the store:

$$\begin{aligned}3y_1 + 2y_2 + 2y_3 &\leq 50 \\4y_2 + 5y_3 &\leq 80\end{aligned}$$

- If these constraints are satisfied, the grocery store knows that the baker will buy at least 6oz of chocolate, 10oz of sugar, and 8oz of cream cheese to satisfy the student's requirements. Hence, it should set the prices of these ingredients in a way that maximizes profit:

$$\max. 6y_1 + 10y_2 + 8y_3$$

As you see, the optimization problem that the grocery store has to solve is the dual problem.

Applications of duality

Max flow / Min cut

Max Flow

Recall that we covered this problem in our first lecture.

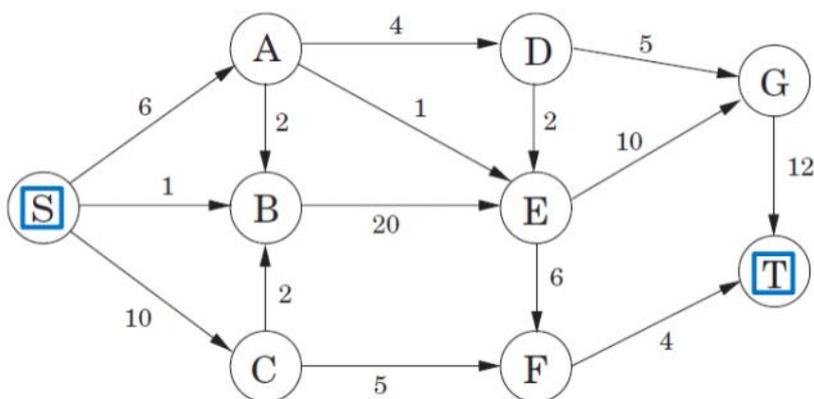


Image credit: [DPV08]

- The goal was to ship as much oil as possible from S to T.
- The amount of oil shipped on each edge could not exceed the capacity of the edge.
- For each node (except S and T), we must have flow in = flow out.

This problem is called the max-flow problem. It can be written as an LP.

$x_{SA}, x_{AD}, x_{BE}, \dots, x_{GT}$ ← Decision variables

max. $x_{SA} + x_{SB} + x_{SC}$ ← Objective function

s.t.

$x_{SA}, x_{AD}, x_{BE}, \dots, x_{GT} \geq 0$

$x_{SA} \leq 6, x_{AB} \leq 2, x_{EG} \leq 10, \dots, x_{GT} \leq 12$

$$\begin{cases} x_{SA} = x_{AD} + x_{AB} + x_{AE} \\ x_{SC} = x_{CB} + x_{CF} \\ \vdots \\ x_{CF} + x_{EF} = x_{FT} \end{cases}$$

← Constraints

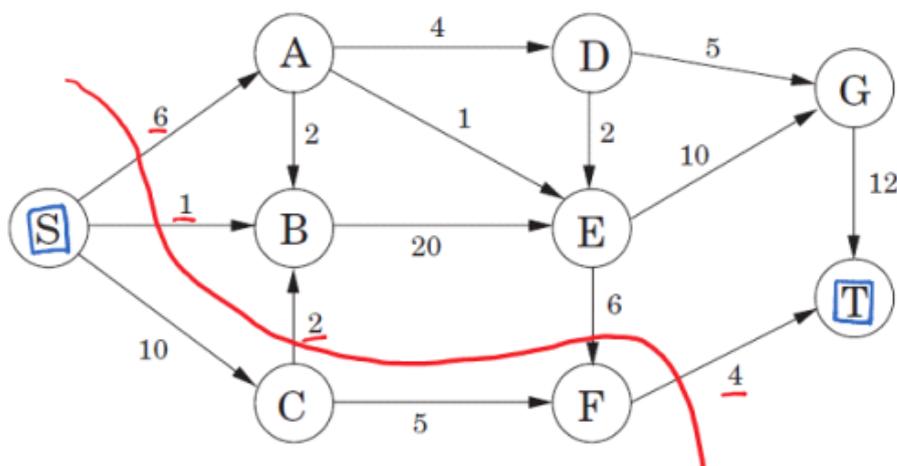
What is min-cut?

- A *cut* is a partition of the nodes of a graph into two disjoint non-empty subsets. An *S-T cut* is a cut that has node S on one side and node T on the other.
- The *value of an S-T cut* is the sum of the weights on the edges that cross the cut from the subset including S to the subset including T (the edges going in the opposite direction are not included).
- The *(S-T) min-cut problem* is the problem of finding an S-T cut with minimum value.

Example: We consider the example given on the previous page lecture.

The cut depicted below corresponds to the sets {S,C,F} and {A,B,D,E,G,T}.

The value of the cut is $6 + 1 + 2 + 4 = 13$. This in fact must be the value of the min-cut as it matches the value of the max-flow! (See below.)



Link between min-cut and max-flow

- It can be shown that if we take the dual of the maximum flow LP, we get a formulation of the min S-T cut problem.
- Observe that the value of any S-T cut is clearly an upper bound on the maximum flow. (Indeed, any flow we can send from S to T has to go through the cut.) This result is in fact a statement of weak duality!
- There always exists a cut whose value matches the optimal flow. This is a consequence of strong duality.
- Note that the minimum cut and the maximum flow certify each other's optimality.

Zero-sum games

- One of the main theorems in game theory is about existence of equilibria in finite games. In the special case of the so-called "zero-sum games", this result is a direct corollary of strong duality in linear programming [DPV08, Sect. 7.5].
- We are skipping this beautiful application since it's covered in ORF307.

Robust linear optimization

"To be uncertain is to be uncomfortable, but to be certain is to be ridiculous."
-Chinese proverb
(from [BEN09])

- So far in this class we have assumed that an optimization is of the form

$$\begin{aligned} \min. & f(x) \\ & g_i(x) \leq 0, i = 1, \dots, n \\ & h_j(x) = 0, j = 1, \dots, m, \end{aligned}$$

where f, g_i, h_j are *exactly known*. In real life, this is most likely not the case; the objective and constraint functions are often not precisely known or at best known with some noise.

- *Robust optimization* is an important subfield of optimization that deals with uncertainty. Under this framework, the objective and constraint functions are only assumed to belong to certain sets in function space (called the "uncertainty sets"). The goal is to make a decision that is feasible no matter what the constraints turn out to be, and optimal for the worst-case objective function.
- We'll only be looking at robust linear programming, in fact a special form of it with polytopic uncertainty.

- A robust LP is a problem of the form :

$$\begin{aligned} \min. & c^T x \\ \text{s.t.} & a_i^T x \leq b_i, \forall a_i \in U_{a_i}, \forall b_i \in U_{b_i}, i = 1, \dots, m. \end{aligned}$$

where $U_{a_i} \subseteq \mathbb{R}^n$ and $U_{b_i} \subseteq \mathbb{R}$ are given uncertainty sets.

- Notice that with no loss of generality, we are assuming there is no uncertainty in the objective function. This is because of the following equivalence:

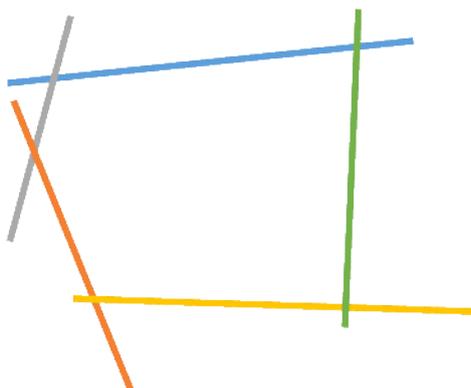
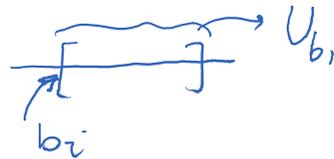
$$\begin{aligned} \min. x \max. c \in U_c c^T x \\ \text{s.t.} a_i^T x \leq b_i, \forall a_i \in U_{a_i}, \forall b_i \in U_{b_i}, i = 1, \dots, m \\ \Downarrow \\ \min. x, \alpha \alpha \\ c^T x \leq \alpha \forall c \in U_c \\ a_i^T x \leq b_i \forall a_i \in U_{a_i}, \forall b_i \in U_{b_i}, i = 1, \dots, m \end{aligned}$$

Robust LP with polytopic uncertainty

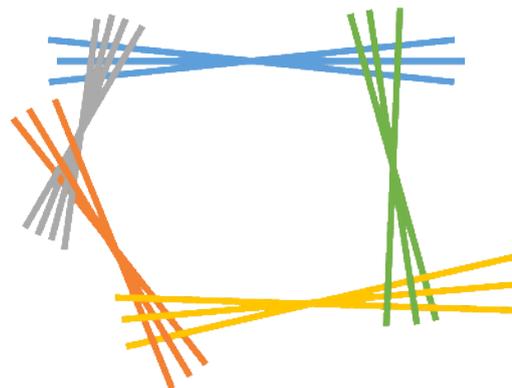
- This the special case where U_{a_i} and U_{b_i} are polyhedra.
- It means that $U_{a_i} = \{a_i \mid D_i a_i \leq d_i\}$ where $D_i \in \mathbb{R}^{k_i \times n}$ and $d_i \in \mathbb{R}^{k_i \times 1}$ are given to us as input. Similarly, each U_{b_i} is a given interval in \mathbb{R} .
- Clearly, we can get rid of the uncertainty in b_i because the worst-case scenario is achieved at the lower end of the interval. So our problem becomes:

$$\begin{aligned} (1) \min. & c^T x \\ \text{s.t.} & a_i^T x \leq b_i, \forall a_i \in U_{a_i}, i = 1, \dots, m, U_{a_i} = \{a_i \mid D_i a_i \leq d_i\} \end{aligned}$$

(with abuse of notation, we are reusing b_i to denote the lower end of the interval)



Feasible set with no uncertainty



Feasible set with polytopic uncertainty

Lec12p12, ORF363/COS323

- LP (1) can be equivalently written as

$$\begin{aligned} & \min_x c^T x \\ & \text{s.t.} \begin{bmatrix} \max_{a_i} a_i^T x \\ D_i a_i \leq d_i \end{bmatrix} \leq b_i, i = 1, \dots, m \end{aligned}$$

This is an LP with an LP in inside it. It doesn't seem like the type of thing we know how to solve! But duality will come to rescue! Let's take the dual of the inner LP.

- The inner LP has the form of the dual of an LP in standard form. This makes it easy to remember its dual (recall that the dual of the dual is the primal):

$$\begin{array}{ccc} \max_{a_i} a_i^T x & \xrightarrow{\text{take the dual}} & \min_{p_i \in \mathbb{R}^{k_i}} p_i^T d_i \\ D_i a_i \leq d_i & & D_i^T p_i = x \\ & & p_i \geq 0 \end{array}$$

- By strong duality, both problems have the same optimal value, so we can replace the inner LP by its dual and get

$$(2) \begin{aligned} & \min_x c^T x \\ & \text{s.t.} \begin{bmatrix} \min_{p_i \in \mathbb{R}^{k_i}} p_i^T d_i \\ D_i^T p_i = x \\ p_i \geq 0 \end{bmatrix} \leq b_i, i = 1, \dots, m \end{aligned}$$

- But now we are in business since we have two minimization problems inside each other and can combine them. The previous problem is equivalent to

$$(3) \begin{aligned} & \min_{x, p_i} c^T x \\ & \text{s.t.} \begin{aligned} & p_i^T d_i \leq b_i, i = 1, \dots, m \\ & D_i^T p_i = x, i = 1, \dots, m \\ & p_i \geq 0, i = 1 \dots m \end{aligned} \end{aligned}$$

Proof: (\Leftarrow) Suppose we have an optimal x, p for (3). Then x is also feasible for (2) and the objective values are the same.

(\Rightarrow) Suppose we have an optimal x for (2). As x is feasible for (2), there must exist p verifying the inner LP constraint. Hence, (x, p) would be feasible for (3) and would give the same optimal value.

LP (3) is a regular LP that we know how to solve. Duality has enabled us to solve a robust LP with polytopic uncertainty just by solving a regular LP!

Notes:

LP duality is covered in Chapter 7 of [DPV08]. The topic of robust linear programming is not covered in the book.

References:

- [DPV08] S. Dasgupta, C. Papadimitriou and U. Vazirani, Algorithms, McGraw Hills, 2008.
- [BT97] D. Bertsimas and J.N. Tsitsiklis, Introduction to Linear Optimization, Athena scientific, 1997.
- [Wri05] S. Wright, Notes on linear programming methods, University of Wisconsin, 2005.
- [BEN09] A. Ben-Tal, L. El-Ghaoui, A. Nemirovski, Robust Optimization, Princeton University Press, 2009.