



Linear Programming and Duality in Practice

Optimisation and OR 01 | Formulating LPs, the graphical method, simplex vs interior point vs first-order (PDLP), the dual as resource prices, complementary slackness, reduced costs, sensitivity ranges, the transportation problem; PuLP, Pyomo, SciPy and HiGHS code

#LinearProgramming #Duality #ShadowPrices #HiGHS #PuLP

What an LP is

A workshop deciding its weekly production		A linear programme
How many of each product to make	≈	Decision variables $x, y \geq 0$
Profit per unit of each product	≈	Objective coefficients c
Hours available on each machine	≈	Right-hand sides b
Hours each product uses per machine	≈	Constraint coefficients A
Machines fully used at the best plan	≈	Binding constraints
What one more machine hour is worth	≈	Shadow price (dual value)

WHERE THE ANALOGY BREAKS

Linearity assumes constant returns: no volume discounts, no set-up costs, fractional units allowed. Set-ups and whole units need integers (sheet 02).

A linear programme read as a workshop plan

Before Modelling and reading LPs, can you explain...

	Not yet	Partly	Yes
MUST KNOW Solving 2×2 linear systems Learn: maths 01	☐	☐	☐
MUST KNOW Matrix-vector notation $Ax \leq b$ Learn: maths 03	☐	☐	☐
MUST KNOW Convex sets Learn: maths 07	☐	☐	☐
HELPFUL Lagrangian duality, KKT Learn: maths 17	☐	☐	☐
HELPFUL Python and pandas Learn: programming	☐	☐	☐


Readiness; each item recalled in one line where used

Item	Formula
Primal	$\max_x c^\top x \text{ s.t. } Ax \leq b, x \geq 0$
Dual	$\min_y b^\top y \text{ s.t. } A^\top y \geq c, y \geq 0$
Weak duality	$c^\top x \leq b^\top y$ for any feasible x, y
Strong duality	$c^\top x^* = b^\top y^*$ when an optimum exists
Complementary slackness	$y_i^*(b_i - a_i^\top x^*) = 0, \quad x_j^*((A^\top y^*)_j - c_j) = 0$
Reduced cost	$\bar{c}_j = c_j - (A^\top y^*)_j$

where x = primal decision vector (n variables) (units of each activity); c = profit (or cost) per unit of each activity (EUR per unit); $A, a_i = m \times n$ matrix of resource use; a_i its row i (resource per unit); b = resource availabilities (right-hand sides) (resource units); y = dual vector: one price per constraint (EUR per resource unit); x^*, y^* = optimal primal and dual solutions (-)

Standard primal-dual pair (canonical form) and the three theorems used in practice

Solve and read a 2-variable LP



optimum (2, 6) unique vertex	z* (max) 36 3x + 5y
------------------------------------	---------------------------

constraint	a · x*	b	slack	shadow
plant 1 (h)	2	≤ 4	2	0
plant 2 (h)	12	≤ 12	0	binding 1.5
plant 3 (h)	18	≤ 18	0	binding 1

plus $x \geq 0, y \geq 0$
5 feasible vertices of 10 line pairs

Optimum at the vertex (2, 6): $z^* = 3x + 5y = 36$

$\max z = 3x + 5y$ subject to $a_{1i}x + a_{2i}y (\leq, \geq, =) b_i$; the optimum of a bounded feasible LP lies at a vertex
vertex = solution of two constraint lines (Cramer: $D = a_{1i}a_{2j} - a_{1j}a_{2i} \neq 0$), kept when it meets every constraint
shadow price $y_i = \Delta z^* / \Delta b_i \approx (z^*(b_i + \delta) - z^*(b_i)) / \delta$, $\delta = 0.0001 \cdot \max(1, |b_i|)$; slack > 0 gives $y_i = 0$

where x, y = decision variables; a_{1i}, a_{2i}, b_i = coefficients and bound of constraint i ; z^* = optimal objective value; slack = $b_i - a_{1i}x^* - a_{2i}y^*$ ($\leq$) or $a_{1i}x^* + a_{2i}y^* - b_i$ ($\geq$); y_i = objective gain per extra unit of b_i

Product mix: x, y = batches of products 1 and 2 per week, profit 3 and 5 kEUR per batch; plant hours limit production.
The optimum is the vertex (2, 6) with profit 36 kEUR; plants 2 and 3 bind

PROBLEM Primal: $\max 3x + 5y$ s.t. $x \leq 4, 2y \leq 12, 3x + 2y \leq 18, x, y \geq 0$. Find the dual, its optimum and interpret it.

GIVEN

- Optimum found above: $x^* = 2, y^* = 6$, profit $z^* = 36$ kEUR
- Dual variables u_1, u_2, u_3 = value of one extra hour in plants 1, 2, 3 (kEUR per hour)

1 Write the dual

$\min 4u_1 + 12u_2 + 18u_3$ s.t. $u_1 + 3u_3 \geq 3$ (product 1), $2u_2 + 2u_3 \geq 5$ (product 2), $u \geq 0$

Why: one dual constraint per primal variable: hours must be worth at least the profit they make

2 Use complementary slackness

plant 1 uses 2 of 4 h $\rightarrow$ slack $2 > 0 \rightarrow u_1 = 0$

Why: an unused resource is worth nothing at the margin

3 $x^*, y^* > 0 \rightarrow$ both dual constraints tight

$3u_3 = 3 \rightarrow u_3 = 1; 2u_2 + 2 = 5 \rightarrow u_2 = 1.5$

Why: a product made at the optimum has zero reduced cost

4 Check strong duality

$4 \cdot 0 + 12 \cdot 1.5 + 18 \cdot 1 = 18 + 18 = 36$

Why: equals the primal profit: both are optimal

ANSWER Shadow prices 0, 1.5 and 1 kEUR per hour for plants 1, 2, 3; pay at most 1.5 kEUR for an extra plant-2 hour

CHECK ☒ Valid only within the right-hand-side range that keeps the same binding set; beyond it the price changes (solvers report the range)

Writing and checking the dual of the product-mix LP

Shadow price

per constraint: objective gain per unit of right-hand side

vs

Reduced cost

per variable: how far its coefficient is from making it worth using

Tell: Am I changing a resource or a product's profit?

Slack

unused resource in the primal solution

vs

Dual value

price of the resource; zero when slack > 0

Tell: Quantity left or value of one more unit?

Infeasible

no point satisfies all constraints

vs

Unbounded

objective improves forever (missing constraint)

Tell: Is the region empty or open in the improving direction?

Degenerate optimum

more constraints binding than variables; duals not unique


vs

Alternative optima

objective parallel to a binding edge; many best x

Tell: Is the dual or the primal non-unique?

LP outputs practitioners confuse



optimum (3, 1) unique vertex	z^* (min) 9 $2x + 3y$
------------------------------------	-------------------------------

constraint	a · x*	b	slack	shadow
nutrient A	4	≥ 4	0	binding 1.5
nutrient B	6	≥ 6	0	binding 0.5

plus $x \geq 0, y \geq 0$
3 feasible vertices of 6 line pairs

Optimum at the vertex (3, 1): $z^* = 2x + 3y = 9$

$\min z = 2x + 3y$ subject to $a_{11}x + a_{12}y (\leq, \geq, =) b_1$; the optimum of a bounded feasible LP lies at a vertex
vertex = solution of two constraint lines (Cramer: $D = a_{11}a_{22} - a_{12}a_{21} \neq 0$), kept when it meets every constraint
shadow price $y_i = \Delta z^* / \Delta b_i \approx (z^*(b_i + \delta) - z^*(b_i)) / \delta$, $\delta = 0.0001 \cdot \max(1, |b_i|)$; slack > 0 gives $y_i = 0$


where x, y = decision variables; a_{11}, a_{12}, b_1 = coefficients and bound of constraint i ; z^* = optimal objective value; slack = $b_i - a_i \cdot x^*$ ($\leq$) or $a_i \cdot x^* - b_i$ ($\geq$); y_i = objective gain per extra unit of b_i

Minimisation with $\geq$ constraints (a feed blend): x, y = kg of two ingredients costing 2 and 3 EUR/kg, nutrient A needs $x + y \geq 4$, nutrient B needs $x + 3y \geq 6$. Optimum (3, 1) at 9 EUR; the region is unbounded upward but the minimum exists

⚙ Algorithms and solvers

Family	How it moves	Strength	Typical use
Primal / dual simplex (1947)	vertex to adjacent vertex	exact vertex, fast warm start	re-solves inside MIP branch and bound; small-medium LPs
Interior point / barrier (1984)	through the interior along the central path	few iterations, large sparse LPs	big LPs; then crossover to a vertex
First-order PDLF (2021)	primal-dual gradient steps, matrix-vector products only	huge LPs, GPU-friendly, low memory	very large LPs at 1e-4 to 1e-6 accuracy

Three LP algorithm families, all inside modern solvers (as of 2026)



Transportation LP solved with HiGHS: plants P1 (30 t) and P2 (40 t) supply C1 (20 t), C2 (25 t), C3 (25 t). Unit costs in EUR/t on the lanes; optimal flows shown, total cost 275 EUR

PuLP (CBC by default)	Pyomo + HiGHS (pip: highspy)
<pre>1 import pulp as pl 2 m = pl.LpProblem("mix", pl.LpMaximize) 3 x = pl.LpVariable("x", lowBound=0) 4 y = pl.LpVariable("y", lowBound=0) 5 m += 3*x + 5*y 6 m += x <= 4, "plant1" 7 m += 2*y <= 12, "plant2" 8 m += 3*x + 2*y <= 18, "plant3" 9 m.solve(pl.PULP_CBC_CMD(msg=False)) 10 print(pl.LpStatus[m.status], 11 pl.value(m.objective)) 11 print({n: c.pi for n, c in m.constraints.items()})</pre>	<pre>1 import pyomo.environ as pyo 2 m = pyo.ConcreteModel() 3 m.x = pyo.Var(domain=pyo.NonNegativeReals) 4 m.y = pyo.Var(domain=pyo.NonNegativeReals) 5 m.obj = pyo.Objective(expr=3*m.x + 5*m.y, 6 sense=pyo.maximize) 7 m.p1 = pyo.Constraint(expr=m.x <= 4) 8 m.p2 = pyo.Constraint(expr=2*m.y <= 12) 9 m.p3 = pyo.Constraint(expr=3*m.x + 2*m.y <= 18) 10 m.dual = pyo.Suffix(direction=pyo.Suffix.IMPORT) 11 pyo.SolverFactory("appsi_highs").solve(m) 12 print(pyo.value(m.obj), m.dual[m.p2], 13 m.dual[m.p3])</pre>
1 c.pi = dual value of each named constraint	2 ask the solver to return duals 3 solver plug-in name as of Pyomo 6.x

The product-mix LP in two modelling layers; both call an open-source solver

SciPy one-liner (HiGHS inside): minimise, so negate a maximisation

```
from scipy.optimize import linprog
r = linprog(c=[-3, -5], A_ub=[[1, 0], [0, 2], [3, 2]],
           b_ub=[4, 12, 18], method="highs")
print(r.x, -r.fun) # [2. 6.] 36.0
print(r.ineqlin.marginals) # [-0. -1.5 -1.] (signs: min form)
```

⚠ Duals of a MIP are not shadow prices

Once integer variables appear (sheet 02), the optimal value is a step function of b : no dual exists in the LP sense. Duals printed after fixing the integers and re-solving the LP are only local information. For pricing a resource in a MIP, re-solve with $b + \Delta$ and compare.

Practice and recap

Before trusting an LP in production

✓ baseline / mandatory □ recommended – verify in your context

- ✓ Units consistent in every row (hours vs minutes, t vs kg)
- ✓ Status checked: optimal, not just 'solved'
- ✓ Coefficients scaled to roughly $1e-3$.. $1e4$ (avoid $1e9$ next to $1e-6$) *bad scaling causes tolerance errors*
- **SHOULD** Sensitivity ranges read before quoting a shadow price
- **SHOULD** Slack or penalty variables on soft constraints so the model never returns infeasible
- **COULD** Warm start from yesterday's basis when re-solving daily

⊕ Why is the optimum of an LP at a vertex?

The objective is linear, so on any segment inside the feasible polyhedron it is monotone; moving along an improving direction ends on the boundary, then on an edge, then at a vertex. If an optimum exists, some vertex is optimal; several optimal vertices mean a whole optimal face (alternative optima). Simplex exploits this by only visiting vertices.

YOU MAY HEAR

"Simplex is exponential, so avoid it"

BUSTED

✓ ACTUALLY

It is the default for medium LPs and inside MIP

Worst cases (Klee-Minty) are artificial; dual simplex warm starts make it unbeatable for re-solves.

YOU MAY HEAR

"A shadow price holds for any change"

BUSTED

✓ ACTUALLY

Only inside its ranging interval

Outside, the binding set changes and the price jumps.

YOU MAY HEAR

"LP is too simple for real problems"

BUSTED

✓ ACTUALLY

Network flows, blending, transport and many planning models are pure LPs

And every MIP solver solves thousands of LPs per run.

YOU MAY HEAR

"Zero shadow price means the constraint is useless"


BUSTED

✓ ACTUALLY

It is not binding at this optimum

Remove it and another plan may break it.

Remember



Q1 Optimum and profit of $\max 3x + 5y$ with $x \leq 4$, $2y \leq 12$, $3x + 2y \leq 18$?

A (2, 6) with 36; plants 2 and 3 bind.

Q2 Shadow prices of the three plants in that LP?

A 0, 1.5 and 1 per hour; plant 1 has slack 2 so its price is 0.

Q3 Dual of $\max c \cdot x$, $Ax \leq b$, $x \geq 0$?

A $\min b \cdot y$, $A^T y \geq c$, $y \geq 0$.

Q4 State complementary slackness in words.

A Each constraint is binding or has a zero price; each used variable has zero reduced cost.

Q5 What does a reduced cost of -2 on an unused product mean (maximisation)?

A Its profit must rise by 2 per unit before making it pays off.

Q6 When prefer an interior-point method?

A Large sparse LPs solved once; simplex wins for warm-started re-solves.

Q7 What is PDLP good for?

A Very large LPs with matrix-vector products only, including on GPUs, at moderate accuracy.

Q8 Optimal cost of the transport example?

A 275 EUR: $P1 \rightarrow C1$ 20, $P1 \rightarrow C2$ 10, $P2 \rightarrow C2$ 15, $P2 \rightarrow C3$ 25 (t).

Q9 Feed blend $\min 2x + 3y$, $x + y \geq 4$, $x + 3y \geq 6$?

A (3, 1) at 9 EUR; duals 1.5 and 0.5 EUR per unit of nutrient.