What the Week 3 pack was testing

2105623 Optimization of Chemical Processes · Semester 1 / 2026 · Issued Monday 31 August 2026

You did Week 3 with nobody in the room, and the pack asked you for three **predictions** written before checking — B3, B5 and C3. This document answers them. Read it with your own submission beside you; the value is where the two differ.

Two of the three had an answer I could not have guessed you would find, and a good number of you did. Those are named below.

Part A · the node balance

The balance at hub $H1$, in the convention outflow minus inflow:

$$f_{H1,H2} + f_{H1,G1} + f_{H1,G2} + f_{H1,G3} + f_{H1,G4} \; - \; \big(f_{R1,H1} + f_{R2,H1} + f_{H2,H1}\big) \; = \; 0.$$

Eight nodes, fifteen arcs, $\sum_n b_n = 830 - 830 = 0$ — and it *must* be zero because every column of the incidence matrix has exactly one $+1$ and one -1, so the rows sum to the zero vector.

A4 Your diagram almost certainly looked different from your neighbour's and from the one on slide 13. Layout, ordering and crossings are not the model. **The model is the incidence structure** — which arcs exist, and which nodes each one joins.

A5 The reverse arc $H2 \to H1$ is a **separate arc with its own non-negative flow**. Treating it as negative flow on $H1 \to H2$ is a different and weaker model, and it throws away the sign structure the whole method rests on. The other common slip was writing b_{H1} as a supply because the hub is drawn in the middle: a hub holds no stock, so $b = 0$.

If you did the optional solve: the network costs **19,570.00 USD/week**, and the inter-terminal arc $H1 \to H2$ carries **40 t**. The arc most pairs forgot to draw is one the optimum actually uses.

Part B · transportation

x_{ij}, t/week	$G1$	$G2$	$G3$	$G4$	u_i, USD/t
$R1$	**300**	**150**	0	0	**0**
$R2$	0	**110**	**270**	0	-5
$R3$	0	0	**60**	**240**	-2
v_j, USD/t	18	24	27	23	

B1–B2 **23,570.00 USD/week**, and the dual objective agrees to the cent, so strong duality holds. Every shipment came out a whole number of tonnes from a pure LP with no integer variable: that is **total unimodularity** of the node-arc incidence matrix, and it is a property of the *matrix*, not of the costs.

B3 The three answers are **0, 5 and 2** USD/t for $R1$, $R2$ and $R3$. Raising $R2$ to 381 t gives 23,565.00; raising $R3$ to 301 t gives 23,568.00; raising $R1$ to 451 t gives **23,570.00 — no change at all**.

$R1$ is the one to sit with. Its supply row *binds* — it ships all 450 t it has — and it is still worth nothing. That is a **degenerate vertex**: $R1$ is the marginal plant, exhausted, but the last tonne it ships is worth precisely what it costs, so relaxing its limit buys you nothing. Slide 9 printed this dual; if you had read it, the mark was for writing all three numbers down before re-solving anyway.

B4 With $R2$ cut to 200 t the model is **infeasible** — 950 t of supply against 1,130 t of demand, and demand written with $\geq$. With a dummy plant at a uniform penalty M the answer is $19{,}610.00 + 180M$, and **the dummy serves $G3$, all 180 t of it, for every value of M.**

Why $G3$, and why it does not depend on M: with a uniform penalty the destination is decided entirely by the real network, and $G3$ carries the largest demand dual, $v_3 = 27$ USD/t. It is the most expensive customer to serve and therefore *the cheapest to disappoint*. The dummy did not merely repair the model — it identified something.

B5 Both statements are true, and this is the question the pack was built around. The sensitivity report prices one more tonne at $G3$ *in the model as written*, holding everything else free to adjust. But the instance is balanced — 1,130 t against 1,130 t — and every supply row binds, so **there is no tonne anywhere to move**. Raising d_3 to 331 t leaves the feasible region empty.

A shadow price says nothing about whether a change keeps a problem feasible. It is valid only over a range, and here that range is empty on the upward side.

The honest reading of v_j on a balanced instance: it is the cost of one more tonne at G_j *paired with* one more tonne of supply somewhere. Paired with $R1$, the marginal plant, that cost is exactly $18, 24, 27, 23$ — the v_j themselves. Paired with $R2$ it is $13, 19, 22, 18$.

Part C · transshipment, and the best question in the pack

25,850.00 USD/week. If your flow pattern disagreed with slide 10, **you were not wrong**: this network has several different optimal shipping plans at exactly that cost, and OpenSolver's CBC finds a different one from the slide — one in which the direct lane $R1 \to G1$ carries 50 t. Three arcs can be saturated in some optimal plan, not two. The total and the node balances are the check; the pattern is not.

saturated arc	capacity +1 t	capacity −1 t	what it is
$R1 \to H1$	cost −10.00	cost +10.00	a real price, symmetric
$R2 \to H2$	cost ±0.00	cost ±0.00	full, and worth nothing either way
$R3 \to H2$	cost ±0.00	cost +**5.00**	**one-sided**

C3 and C4 — the answer, and why it is worth more than the arithmetic.

Only $R1 \to H1$ behaves the way a price is supposed to: widen it and you save 10 USD/t, narrow it and you pay 10 USD/t. $R3 \to H2$ **does not.** Widening it is worth *nothing*; narrowing it costs *5 USD/t.* Its shadow price is **one-sided** — a right derivative of 0 against a left derivative of 5 — and no single number describes it.

Why. $R3$ has 300 t to send and exactly 300 t of pipe. A wider pipe carries nothing extra, because there is no more material behind it; a narrower one strands a tonne that must then travel the dearer way. *Saturation is a statement about the flow. A shadow price is a statement about the optimum.* They coincide often enough that people stop distinguishing them, and this instance is where that habit fails.

So OpenSolver may have printed 0 for that arc, or it may have printed 5, and both are correct. The optimal capacity dual there is the whole interval [0, 5]. Whatever it printed, your re-solve settled what actually happens — which is why the pack made you re-solve rather than read.

C5 **Ten arcs** carry a zero reduced cost. Seven of them are the arcs carrying flow strictly between their bounds, and those *must* be zero — that is complementary slackness, a check rather than a discovery. The three worth reading are: $R3 \to H2$, zero and *saturated* (allowed: a zero reduced cost permits any flow between the bounds, the bounds included — and it is the algebra behind C4); and $R1 \to G1$ and $H2 \to G2$, zero and *empty*, which is the dual side of the non-uniqueness above.

And if your list had nine rows rather than ten, that is also right. **One of the nine node potentials is not determined by the problem** — π_{R3} can be anything from 23 to 28 — so two correct answers can differ. What is marked is whether your eighteen rows are consistent with the potentials *you* used.

A second thing several of you asked: the notebook printed nine potentials that were all 32 lower than the sensitivity report's. Both are right. Potentials are fixed only **up to an added constant**, because only the difference $\pi_i - \pi_j$ ever appears. Compare reduced costs, never potentials.

Parts D and E

D1 The LP returns a single path because total unimodularity forces an integral basic optimal solution, and with $b = +1$ at the origin and -1 at the destination the only integral feasible flows *are* paths. Nothing told it the answer had to be a path. *Shortest path $R1 \to G4$ =* ***29 h****, via $H1$ and $H2$.*

D2 Maximum flow **1,100 t** against 1,130 t of demand, the whole 30 t shortfall on $G3$. Only widening $H1 \to G3$ or $H2 \to G3$ recovers it — and the half of the answer that matters is the other half: **every arc not in the cut is worth exactly zero at the margin.** A debottlenecking study that reports only the throughput number has thrown the answer away.

D3 The 35 percent share row has coefficients 1 and −0.35 spread across arcs that share no node, so it is not part of the incidence pattern, the matrix stops being totally unimodular, and fractional vertices reappear: **26,195.00** USD/week with flows at exact half tonnes. The model has changed complexity class, from LP to MILP.

E4–E5 There is no split fraction in the linear superstructure and there does not need to be one: the split is *recovered after the solve* (100/115 to PSA, 15/115 to MEM), and writing it as a ratio would introduce a division and make the model nonlinear for nothing. And the MILP switches on electrolysis — the route the LP refused to use at all — because once ATR and the membrane have to be paid for as capital, dropping both is worth more than the running-cost penalty of covering the residual 20.8 t/day. **A route with a high running cost can still be correct if it avoids a large capital commitment.**

What Week 3 was really for

Three of its questions asked you to commit to a number before you could check it, and two of those numbers were designed to be wrong in an instructive way. That is not a trick. It is the only way to find out what you actually believe about a model, and it is the habit that separates an engineer who uses an optimizer from one who is used by it.

Today's session is the same lesson at a higher temperature. In Week 3 a shadow price was one-sided. Today a solver will report `optimal` twice, on the same model with the same data, and give two different answers — and both status words will be honest.