WEEK 2 · THE ANSWERS

What the Week 2 pack was testing, and what it answers to

2105623 Optimization of Chemical Processes · Semester 1 / 2026 · Issued Monday 24 August 2026

I am away on Monday 24 August, so the five minutes that would have opened Week 3 are written down here instead. Read this **after** you have submitted the Week 2 pack and **before** you start the Week 3 pack — Week 3 uses the same two ideas, on a graph instead of a time line.

Read it with your own submission next to you. The value is in the places where your answer and this one differ, and in whether you can say why.

Part A · Aurora Specialty Chemicals

A1 · the specification

Sets. $g \in G = \{\text{Standard, HighPurity}\}$, grades. $t \in T = \{1, 2, 3, 4\}$, months, ordered.

Parameters. $R = 640$ h/month, reactor availability. a_g, reactor hours per tonne: $a_{\text{Std}} = 1.0$, $a_{\text{HP}} = 1.4$ h/t. $d_{g,t}$, demand in t/month: Standard 180, 210, 260, 190; HighPurity 90, 120, 150, 110. $c_{g,t}$, production cost in USD/t — differs by grade and rises in month 3; *no numbers were given, and that is itself an ambiguity.* $h = 7$ USD/t/month, holding cost. $W = 120$ t, warehouse capacity. $I_{g,0}$, opening stock: Standard 30 t, HighPurity 0 t.

Decision variables. $x_{g,t} \geq 0$, production of grade g in month t, t/month. $I_{g,t} \geq 0$, closing stock of grade g at the end of month t, t.

Objective, in words. Minimise the total four-month cost in USD: production cost plus holding cost.

Constraints, in words.

R1. For each grade and month, opening stock plus production minus demand equals closing stock.

R2. In each month, the reactor hours consumed by both grades cannot exceed the hours available.

R3. In each month, the stock held cannot exceed the warehouse capacity.

R4. Production and stock cannot be negative.

Assumptions, numbered alongside. A1 inventory is measured at the *end* of each month. A2 the 120 t limit applies to the *total* of both grades. A3 demand must be met in full in its own month; no backlog. A4 production is continuous, not in fixed batches. A5 stock left at the end of month 4 is left free.

What was actually being marked. Units on every parameter and every variable, and constraints stated *in words*. A different assumption from mine is not an error — a specification that says "the warehouse limit is per grade" and *says so* has done the exercise correctly. What loses marks is a silent choice: a reader cannot tell whether you decided or simply did not notice.

A2 · the ambiguities — this was the part that mattered

1. **Is inventory measured at the start or the end of the period?** Both readings are defensible, the balance equation differs, and *the answer differs*, not just the notation.
2. **Is the 120 t warehouse limit per grade, or on the total?** The text says one warehouse, so the total — but the wording does not force it.
3. **Is the demand forecast a requirement or a maximum?** And if a requirement, is a shortfall forbidden outright, or backlogged at a penalty the scenario never states?
4. **Fractional tonnes, or fixed batches?** Nothing says. That question is Week 5.
5. **And a fifth, which few found:** nothing says what happens to stock left at the end of month 4. A model that leaves it free will end at zero; a supply contract might require otherwise.

A sixth, strictly: the production costs $c_{g,t}$ are described but never given. A specification that lists the symbol and its units, and notes that the values are missing, is complete. One that quietly invents numbers is not.

One thing worth noticing before Week 3. The reactor never binds. The hours required are 306, 378, 470 and 344 against 640 available in every month. So the whole Aurora problem is driven by the month-3 cost rise pushed against 7 USD/t/month of holding, capped by the 120 t warehouse — and the constraint that looks like the physical bottleneck is slack throughout. If you spotted that, you found the shape of the problem, and you have met the idea Week 8 will call complementary slackness.

Part B · the multiperiod sheet

Period	1	2	3	4	5	6
Demand, t	120	150	180	210	160	130
Production x_t, t	130	170	170	170	160	130
Closing inventory I_t, t	30	50	40	0	0	0
Backlog B_t, t	0	0	0	0	0	0
Capacity slack, t	40	0	0	0	10	40
Capacity dual, USD/t	**0**	**6**	**9**	**12**	**0**	**0**

B1 **Total cost 42,080.00 USD** = 41,360.00 production + 720.00 holding + 0 backlog. The workbook's Readme states this figure, so B1 was a checksum on your formulas, not the test.

B2/B3 **The prediction: the cost falls by 12.00 USD**, to 42,068.00. That is the period-4 capacity dual. *Why:* the marginal tonne delivered in period 4 currently has to originate in period 1 — the only early period with spare capacity — and be held for three months, so it costs $42 + 3(6) = 60$ USD/t, against 48 USD/t if it could simply be made in period 4. The difference is 12. The same accounting gives 9 for period 3 and 6 for period 2.

second half **Raising period 1 changes nothing.** Its dual is zero because period 1 already has 40 t of unused capacity, and relaxing a constraint that is not binding cannot change the optimum. The cost stays at 42,080.00 exactly.

B4 Pyomo returns the same 42,080.00. If your sheet and the notebook disagreed, the sheet was wrong — and the question was whether you said which one you trusted.

The two predictions I saw most often, and what each one means.

6 or 9 instead of 12. The right mechanism, the wrong period. That is a good answer: you were pricing the marginal tonne, you just traced it to the wrong origin.

48, or 42. This is the important confusion, and it is worth naming. 48 is what a tonne *costs to make* in period 4; 42 is what it costs in period 1. Neither is what one *more* tonne of capacity is worth, because relaxing the capacity does not create a tonne out of nothing — it lets a tonne be made in a cheaper place and stop being carried. The quantity you want is always a *difference* between two plans, never a price read off the data.

A blank B2 and a correct B3. You checked first and wrote the prediction afterwards. No marks are lost this time, but the habit is the whole point of the exercise, and it will not be optional in Week 8.

Two errors worth checking your own sheet for. A balance row shifted by one period gives a number *close to* 42,080 rather than obviously wrong, which makes it hard to catch — test the period-1 balance against the opening inventory of 20 t. And inventory allowed to go negative, because non-negativity was never set, gives a total that is too low.

Part C · the three questions

C1 **End-of-period:** closing stock = opening stock + production − demand, and the holding charge is applied to the closing stock. **Start-of-period:** the stock that is charged is the one held at the beginning of the period, so the same physical plan is costed one period differently, and the optimal plan itself can shift. The full answer needed both statements *and* the observation that the answer changes, not only the notation.

C2 A backlog variable lets demand be met late at a stated penalty instead of not at all. Without it, and with demand written as a hard equality, a period whose demand exceeds everything that can be produced or carried in makes the whole model **infeasible**: the solver reports failure rather than the cheapest compromise, and the planner learns nothing about what the shortfall would have cost.

C3 Something must be said about stock at the end of the horizon, or the model will push difficult demand past the end of it — running stock and backlog down to whatever is cheapest in the last period, and leaving the plant empty or owing. In the workbook that statement is the condition $B_6 = 0$.

Part D · the HW1 pattern list

For each HW1 question, the data needed and the pattern used. A reasonable list: **P1** multiperiod production — *balance* and *capacity*. **P2** transshipment — *index set*, *balance* at the ports, *capacity* at the refineries. **P3** superstructure — *linking*, through the shared feed and the blended purity. **P4** classification only, no model. **P5** the same model as P2, written over an arc set instead of a matrix.

If your list named different patterns but named them for stated reasons, it is right. The purpose of Part D was that HW1 was opened before the weekend, and the pattern vocabulary of Week 2 is the vocabulary Week 3 extends.

What carries into Week 3

Two things, and you have now met both.

One. The number you predicted in B2 — what one more unit of a scarce resource is worth — has a name, and Week 8 will give it one. You used it six weeks early. In Week 3 the same number reappears as a *node potential*: what one more tonne is worth *at a particular place* rather than *in a particular month*.

Two. A constraint that does not bind is never worth anything — period 1 has spare capacity, so its dual is zero. The converse, which the Week 2 data never put to the test, is the more useful half: *a constraint that does bind is not automatically worth money either.* Week 3 puts that to the test directly, and one of its questions asks you to predict the answer before you check it.

Week 3, Monday 24 August. I am away and there is no session. The Week 3 self-study pack replaces it; it is on the course page with the slides, the workbook and the notebook. www.skhgroup.net/teaching/2105623/ The Week 2 pack told you to bring software problems "to the start of Week 3" — **there is no Week 3 to bring them to**, so e-mail them to me instead, the same day, and there will be an installation clinic at the start of Week 4. **HW1 is due Monday 31 August, at the start of Week 4**, and it covers Weeks 1 to 3.