Week 5: Modeling IV
Logical and Discrete Decisions, Big-M versus Convex Hull

2105623 Optimization of Chemical Processes

Assoc. Prof. Dr. Soorathep Kheawhom

Department of Chemical Engineering, Chulalongkorn University

Semester 1/2026

Session plan, three hours

The shape of today

Time	Block	Duration
0:00 to 0:10	Opening: recap of Week 4, one question from HW1, today's objectives	10 min
0:10 to 0:40	**Lecture 1**: binaries as modeling devices, logic as linear constraints	30 min
0:40 to 1:00	**Activity 1**: logic to linear, board race (pairs, paper)	20 min
1:00 to 1:30	**Lecture 2**: big-M, how to choose it, and its two failures	30 min
1:30 to 1:45	Break	15 min
1:45 to 2:10	**Activity 2**: big-M roulette (computer, groups of 3 to 4)	25 min
2:10 to 2:35	**Lecture 3**: convex hull versus big-M, formulation strength	25 min
2:35 to 2:55	**Activity 3**: tighten it (groups of 3 to 4)	20 min
2:55 to 3:00	Wrap-up, takeaways, what is due next week	5 min

Bring: the notebook `w05-modeling4-logical-discrete.ipynb`, the workbook `w05-fixed-charge.xlsx`, OpenSolver with the CBC engine, and one sheet of paper per pair.

Learning objectives

After this session you should be able to

1. Use a binary variable as a modeling device: yes/no decision, fixed charge, minimum run rate, semicontinuous variable. CLO 1
2. Translate a propositional statement into linear constraints and verify the translation against its truth table. CLO 1
3. Choose a big-M that is valid and tight, and state the two distinct numerical failures caused by an oversized M. CLO 1, CLO 4
4. Write the convex hull (disaggregated) form of a disjunction and explain why its relaxation is the tightest possible. CLO 1
5. Measure formulation strength as an integrality gap and use it, rather than the row count, to predict solution effort. CLO 4

Reference: Williams, Ch. 9 and 10; Rao, Ch. 10. Notebook: `w05-modeling4-logical-discrete.ipynb`. Workbook: `w05-fixed-charge.xlsx`.

Where we are

Weeks 2 to 4

- Linear models, multiperiod planning, inventory.
- Networks, transportation, superstructures.
- Blending, pooling, quality specifications.
- Everything continuous: *how much*.

This week

- Decisions that are *whether*, not *how much*.
- Install or not, run or not, at least k of these, this only if that.

The scope boundary

This week is about **formulating** discrete decisions and about how the formulation changes the linear relaxation. How a solver *searches* the resulting model (branch and bound, cutting planes, node selection) is Week 11.

The one sentence

Two formulations of the same discrete problem give the same answer and can differ by orders of magnitude in the effort needed to prove it.

Four patterns cover most process models

A continuous variable answers *how much*. A binary answers *whether*.

The vocabulary

Pattern	Meaning	Algebra
Yes/no decision	$z = 1$ if the unit is installed, the stream lined up, the contract signed	$z \in \{0, 1\}$, price it in the objective
Fixed charge	cost F paid if and only if the activity happens at all	Fz in the objective, plus $u \leq Mz$
Minimum run rate	a pump that runs at all must run above L	$u \geq Lz$
Semicontinuous	the two above, together	$u \in \{0\} \cup [L, M]$

Why the link constraint is not optional

Without $u \leq Mz$ the solver sets $z = 0$, pays no fixed charge, and still uses $u > 0$. The binary must be *wired* to the continuous variable or it decorates the model instead of constraining it.

The semicontinuous set is a disjunction, not an interval

$$u \in \{0\} \cup [L, M]$$

is the smallest nonconvex set that appears in process models.

- It is **not** $0 \le u \le M$: the band $(0, L)$ is forbidden.
- It is **not** convex: $u = 0$ and $u = M$ are feasible, their midpoint is not.
- The binary is exactly what makes it expressible in a mixed-integer *linear* program:

$$Lz \le u \le Mz, \qquad z \in \{0, 1\}.$$

- Relaxing z to $[0, 1]$ sweeps out the shaded wedge, which is the convex hull of the two pieces.

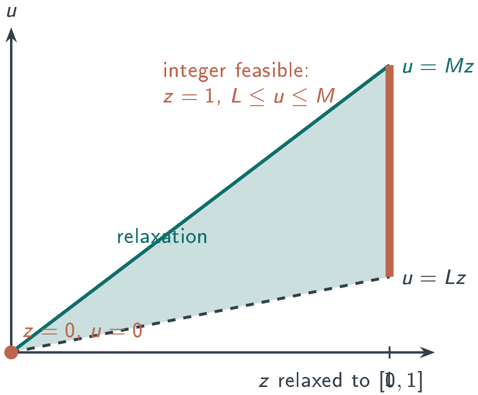

Worked example: refinery blending with operating reality

Five component streams are blended into two gasoline grades (the Week 4 instance). Three operating realities were ignored there.

Stream	cost USD/t	avail. t	fixed F_c USD	min run L_c t
Reformate	760	3200	12 000	800
FCC naphtha	730	4000	9 000	1000
Alkylate	805	2000	15 000	500
Butane	510	900	4 000	200
Straight run	695	2600	7 000	600

The three realities

1. Lining up a stream needs a tank, a pump and a transfer line: fixed charge F_c whether 1 t or 3 000 t moves.
2. A running transfer line has a minimum rate L_c; below it the pump cavitates and the metering is unreliable.
3. The tank farm can line up at most four streams, and butane can move only while the straight-run splitter runs (shared vapor recovery header).

$$\max \sum_g p_g y_g - \sum_c c_c u_c - \sum_c F_c z_c \quad \text{s.t.} \quad L_c z_c \le u_c \le \text{avail}_c\, z_c, \quad z_c \in \{0,1\}$$

plus the Week 4 blending rows: grade balances, volume windows, octane, vapor pressure and sulfur specifications.

What the discrete structure costs

Margin, USD

A. LP, no discrete decisions	413 439.75
B. + fixed charges and minimum runs	342 092.31
cost of the discrete structure	71 347.44
of which fixed charges paid	47 000.00
of which minimum run rates	24 347.44

Stream	L_c	LP u_c	MILP u_c
Reformate	800	3200.0	3200.0
FCC naphtha	1000	4000.0	4000.0
Alkylate	500	833.0	1533.0
Butane	200	574.5	633.9
Straight run	600	473.1	600.0

Read the last column

Straight run was used at 473.1 t in the LP. The minimum run rate of 600 t forces the plant either to move 600 t of a stream it did not want that much of, or to shut the line down entirely. The optimizer chose to run it, and paid for the difference. **The minimum run rate costs one third as much again as all five fixed charges combined.**

Translation table, part 1: propositions over binaries

Let $z_a, z_b, z_c \in \{0, 1\}$ be 1 when the corresponding proposition is true.

Exact and dimensionless: no constant is required

Statement in words	Logic	Linear constraint
a implies b	$a \Rightarrow b$	$z_a \le z_b$
a and b are equivalent	$a \Leftrightarrow b$	$z_a = z_b$
at least one of a, b	$a \vee b$	$z_a + z_b \ge 1$
not both	$\neg(a \wedge b)$	$z_a + z_b \le 1$
exactly one of a set S		$\sum_{j \in S} z_j = 1$
at least k of S		$\sum_{j \in S} z_j \ge k$
at most k of S		$\sum_{j \in S} z_j \le k$
a and b together imply c	$a \wedge b \Rightarrow c$	$z_a + z_b - 1 \le z_c$
a implies b or c	$a \Rightarrow b \vee c$	$z_a \le z_b + z_c$

Every propositional statement over finitely many binaries can be written this way. These nine are the ones worth memorizing.

Translation table, part 2, and how to check a translation

Definitional equivalences need three rows each

Statement	Linear constraints
$c \Leftrightarrow a \wedge b$	$z_c \le z_a$, $z_c \le z_b$, $z_c \ge z_a + z_b - 1$
$c \Leftrightarrow a \vee b$	$z_c \ge z_a$, $z_c \ge z_b$, $z_c \le z_a + z_b$

Where continuous variables enter, and M appears

Statement	Linear constraints
if a then $u \le U_1$, else $u \le U_2$	$u \le U_1 z_a + U_2(1 - z_a)$
if a then $g(x) \le 0$	$g(x) \le M\,(1 - z_a)$

Everything above the alert block is exact. Everything inside it depends on a numerical constant, which is Lecture 2.

Verify, do not remember

Enumerate all 2^n assignments, evaluate the proposition, evaluate the inequalities, confirm they agree on every row. All ten translations in the notebook pass on every assignment.

The control case

Writing $a \Rightarrow b$ as $z_a + z_b \le 1$ fails the same harness: it is the translation of *not both*, and it forbids the assignment $(1, 1)$ that the implication permits.

A wrong translation is silent. The model still solves, still reports `optimal`, and answers a different question. That is the Week 6 theme arriving early.

Activity 1 | 20 min | pairs, paper then board

Task

Translate all eight English statements into **linear** constraints in binary variables. Use z_A, z_B, z_C, z_D for the units, z_1, z_2, z_3 for the technologies, u_C for the run rate of unit C, and z_W for the warehouse lease. At 12 minutes, eight pairs are called to the board, one statement each.

1. If unit A is built, unit B must also be built.
2. At most two of the four units may be built.
3. Exactly one of the three technologies must be chosen.
4. If unit A is built, unit B must not be built.
5. Unit C runs at a rate of at least 20 t per day, or it does not run at all.
6. If total production exceeds 100 t, the second warehouse must be leased.
7. Unit D may be built only if both A and B are built.
8. At least one of A, B, C must be built, and if all three are built a supervisor must be hired.

Hand back: one sheet per pair, all eight translations. No products of binaries are permitted.

[ANSWER] Activity 1: the eight translations

	Statement	Linear constraints
1	A implies B	$z_A \leq z_B$
2	at most two of four	$z_A + z_B + z_C + z_D \leq 2$
3	exactly one technology	$z_1 + z_2 + z_3 = 1$
4	A implies not B	$z_A + z_B \leq 1$
5	run at least 20 t/d or not at all	$20\, z_C \;\leq\; u_C \;\leq\; U_C\, z_C, \quad z_C \in \{0,1\}$, with U_C the physical capacity of C
6	production over 100 t forces the lease	$\sum_j p_j \leq 100 + M\, z_W, \quad M = P^{\max} - 100$
7	D only if both A and B	$z_D \leq z_A$ **and** $z_D \leq z_B$
8	at least one of A, B, C; all three force a supervisor	$z_A + z_B + z_C \geq 1$ **and** $z_S \geq z_A + z_B + z_C - 2$

The two that generate argument

5 and 6 are the only ones that need a constant. 5 is the semicontinuous pattern; 6 is the bridge into Lecture 2. Note that 6 is a one-way implication: it does not forbid leasing the warehouse at low production.

The two commonest errors

7 written as $z_D \leq z_A + z_B$ is wrong: it admits $z_D = 1$ with only A built. The pair $z_D \leq z_A$, $z_D \leq z_B$ is correct; the single row $2z_D \leq z_A + z_B$ is also correct but has a weaker relaxation.

8 written as $z_S \geq z_A z_B z_C$ is not linear. Verified: all eight forms above agree with their truth tables on every assignment.

First, what the logic costs on the blend

Two conditions on the tank farm:

$$\sum_c z_c \leq 4 \quad \text{(cardinality)}, \qquad z_{\text{Butane}} \leq z_{\text{Straight run}} \quad \text{(implication)}.$$

Model	margin, USD	streams on
B. fixed charge + min run	342 092.31	all five
C1. + cardinality only	335 539.13	Ref, FCC, Alk, But
C2. + implication only	342 092.31	all five
C3. + both	281 361.41	Ref, FCC, But, SR

Cost of cardinality alone	6 553.18
Cost of implication alone	0.00
Cost of both together	60 730.89

Restrictions interact, and the interaction is not additive

The implication alone costs nothing, because the unrestricted optimum already runs straight run. Once cardinality removes straight run from the four-stream solution, the implication also removes butane, and the two together cost 60 731 USD, an order of magnitude more than the sum of their individual costs. Total cost of discrete reality: $413\,439.75 - 281\,361.41 = 132\,078.33$ USD, **31.9 percent of the LP margin**.

The big-M pattern: valid and tight are different words

A constraint that applies only when a binary is 1:

$$g(x) \;\le\; M\,(1-z), \qquad z \in \{0,1\}.$$

- $z = 1$: the row reads $g(x) \le 0$ and is enforced.
- $z = 0$: the row reads $g(x) \le M$ and must be non-binding.

The definition

Any M large enough to make the row non-binding when $z = 0$ is **valid**: every valid M gives the same integer optimum. Only the smallest such value is **tight**. The tightest valid value is

$$M^\star \;=\; \max\,\{\, g(x) \;:\; x \text{ feasible for the other alternatives} \,\},$$

which is itself an optimization problem, usually a small LP.

What never to do

Pick a round number such as 10^6 because it looks big enough. Two different things go wrong, and they fail in different ways.

Three recipes for M, in decreasing order of quality

1. **Solve the little LP.** $M^\star = \max\{g(x)\}$ over the rest of the feasible set. Exact, row by row, and what a good modeling layer does automatically.
2. **Use variable bounds.** If $g(x) = a^\top x - b$ and $x^L \le x \le x^U$, then

$$M \;=\; \sum_i \max(a_i x_i^U,\; a_i x_i^L) \;-\; b$$

   is valid. Cheap, no LP, usually loose by a factor of a few.
3. **Use a physical cap already in the model.** In the blending problem the availability limit caps the stream usage, so $M_c = \mathrm{avail}_c$ is both valid and tight for $u_c \le M_c z_c$, and L_c is already tight for $u_c \ge L_c z_c$.

A useful habit

Write M as a named expression in the units of the row it appears in (`M_c = avail[c]`), never as a bare literal. A big-M with units attached is a big-M somebody can check. In the workbook it gets its own visible cell, `Model!$C$36`.

Failure 1: the relaxation collapses

Relaxing z to $[0, 1]$ in $u \leq Mz$ gives $z \geq u/M$. A large M lets the relaxation buy the stream for a small fraction of its fixed charge.

Blending model with logic, same 34 rows in every case

Formulation	MILP optimum	LP bound	gap, USD	$\sum_c z_c$ in LP
tight $M_c = \text{avail}_c$	281 361.41	379 170.50	97 809.08	3.6932
loose $M = 50\,000$	281 361.41	411 575.46	130 214.04	0.1836
very loose $M = 10^7$	281 361.41	413 430.43	132 069.01	0.0009

Read the last column

With $M = 10^7$ the relaxation switches on all five streams for a total z-mass of 0.0009, so it pays 0.09 percent of the fixed charges and its bound is essentially the pure LP value. **A loose M does not change the answer. It changes how long the search takes to prove it**, because the bound is what branch and bound must close. You will measure this yourself in Activity 2.

Failure 2: the integrality tolerance becomes a hole in the model

Solvers accept z as integral within an absolute tolerance, typically $\varepsilon = 10^{-6}$. A value $z = 10^{-6}$ is *reported as* $z = 0$, while $u \leq Mz$ still permits $u \leq M\varepsilon$.

M	flow allowed, t	charge paid, USD
3 200	0.0032	0.012
50 000	0.05	0.012
10^7	10	0.012
10^{10}	10 000	0.012

Made concrete

Set $M = 10^{10}$ and fix every $z_c = 10^{-6}$, which the solver reports as off:

- margin obtained: 413 439.70 USD,
- fixed charges actually paid: 0.047 USD,
- true MILP optimum: 342 092.31 USD.

Whose fault

The plan moves ten thousand tonnes through five lines the solution claims are switched off, and it looks 71 347 USD better than the true optimum. The model is not wrong in exact arithmetic. It is wrong in floating-point arithmetic, which is the only kind a solver has. $M \times \varepsilon$ **is a modeling decision, not a solver defect.**

Break, 15 minutes

Next: **Activity 2**, big-M roulette, at the computer, in groups of 3 to 4.

Open `w05-fixed-charge.xlsx` and check that OpenSolver finds CBC.

Before you leave the room: form your group of 3 to 4 and collect your assigned value of M from the front. Every group solves the *same* model with a *different* M.

The Activity 2 instance: a zinc-air battery materials plant

Three intermediates must be supplied; six candidate units are offered. This is the workbook model, and it is solved with Pyomo in section 6 of the notebook.

Unit	fixed kUSD/yr	hours h/yr	min thru. t/yr
U1 mixer-extruder	180	3500	500
U2 roll press	240	3400	450
U3 slot-die coater	300	3600	600
U4 spray dryer	150	3200	350
U5 calciner	260	3300	550
U6 calender line	210	3400	400

Rates r_{up}, t/h (dash = not possible):

Unit	anode	cathode	separator
U1	0.42	–	–
U2	0.30	0.55	–
U3	–	0.62	0.48
U4	0.36	–	0.40
U5	–	0.50	0.30
U6	–	0.44	0.52

min $\sum_u F_u z_u + \sum_{(u,p)} c_{up} x_{up}/1000$

s.t.	$\sum_u x_{up} = d_p$	demand
	$\sum_p x_{up}/r_{up} \le H_u$	hours
	$\sum_p x_{up} \le M z_u$	**on/off link, the big-M**
	$\sum_p x_{up} \ge L_u z_u$	minimum throughput
	$z_{U2} + z_{U6} = 1$	competing licenses
	$z_{U4} \le z_{U3}$	shared solvent recovery
	$\sum_u z_u \le 4$	commissioning crew

Demand, t/yr: anode paste 1800, cathode sheet 2400, separator film 1500.

The single number you will change is the big-M in cell `Model!$C$36`. It appears in the on/off link and nowhere else.

Activity 2 | 25 min | groups of 3 to 4, computer

Task: big-M roulette

Every group solves the **same** model in `w05-fixed-charge.xlsx`. Complete the yellow cells, then set `Model!$C$36` to **your group's value of** M, and record three numbers: (1) the **LP relaxation** objective, obtained by replacing the `bin` constraint on `$H$40:$H$45` by ≤ 1 and re-solving; (2) the **integer optimum** with `bin` restored; (3) the **solve time** reported by OpenSolver.

Group	multiple	M, t/yr
1	$M^\star$, tightest	2 232
2	$2M^\star$	4 464
3	$10M^\star$	22 320
4	$10^3 M^\star$	2 232 000
5	$10^6 M^\star$	2 232 000 000

Group	M	LP bound	integer opt.	time, s
1	2 232	________	________	_____
2	4 464	________	________	_____
3	22 320	________	________	_____
4	2.232×10^6	________	________	_____
5	2.232×10^9	________	________	_____

This table is copied onto the board and filled in from the floor.

Hand back: three numbers per group, called out for the board table. Predict, before you solve, which of the three columns will differ between groups.

[ANSWER] Activity 2: the board table, filled in

Computed in `w05-modeling4-logical-discrete.ipynb` and reproduced in the workbook

Group	M, t/yr	LP bound, kUSD/yr	integer optimum	gap %	CBC time
1	2 232 ($M^\star$)	2 203.148	2 533.226	13.030	< 0.05 s
2	4 464	1 974.134	2 533.226	22.070	< 0.05 s
3	22 320	1 879.765	2 533.226	25.796	< 0.05 s
4	2 232 000	1 844.009	2 533.226	27.207	< 0.05 s
5	2 232 000 000	1 843.620	2 533.226	27.222	< 0.05 s

What the board shows

1. The integer optimum is 2 533.226 kUSD/yr in every row: every valid M describes the same set of integer points.
2. The LP bound degrades monotonically, by 359.5 kUSD across the table, and then saturates, because it has already collapsed to the no-fixed-charge LP.
3. Solve time does *not* move. Six binaries is too small to time. On the 20-unit instance in section 3 of the notebook, loose M takes 4.28 s and tight M takes 0.078 s, a factor of 55: the bound is the cause, the time is the symptom.

Two instructive wrong answers

A group that never changed `$C$36` reports the shipped default $M = 2\,500$, LP bound 2 143.300, gap 15.39 percent.

A *per-unit* M_u (1470, 1870, 2232, 1280, 1650, 1726.8) beats every row of the table: LP bound 2 403.382, gap 5.126 percent. One number for six different units is already loose, and the disaggregated form of Lecture 3 goes further still.

A disjunction, written two ways

$$\bigvee_k \left[\, Y_k \;;\; A_k x \le b_k \,\right], \qquad \text{exactly one } Y_k \text{ true.}$$

Big-M

$$A_k x \le b_k + M_k(1 - z_k), \quad \textstyle\sum_k z_k = 1$$

- Compact: one copy of x, one row per original row.
- Weak: a fractional z relaxes *every* alternative at once, by a large amount.

Convex hull (disaggregated, Balas)

$$x = \textstyle\sum_k x_k, \quad A_k x_k \le b_k z_k,$$
$$0 \le x_k \le x^U z_k, \quad \textstyle\sum_k z_k = 1$$

- One copy of the continuous variables per alternative.
- With z relaxed to $[0,1]$, the projection onto x is exactly the convex hull of the union of the alternatives.

The hull form is the **tightest linear relaxation any formulation of that disjunction can have**.

The instance: a reactor section in one of two configurations

x_1 = feed rate, x_2 = product rate.

Configuration A, small continuous unit

$10 \le x_1 \le 40, \quad 0.3\,x_1 \le x_2 \le 0.5\,x_1,$
annualized fixed cost 120 kUSD.

Configuration B, large unit

$60 \le x_1 \le 90, \quad 0.5\,x_1 \le x_2 \le 0.7\,x_1,$
annualized fixed cost 260 kUSD.

Operating cost 1.8 kUSD per unit feed and 3.2 kUSD per unit product in either configuration; the section must deliver at least 25 units of product.

Row-wise tightest M, by LP

A row	$M^\star$	B row	$M^\star$
$x_1 \ge 10$	0	$x_1 \ge 60$	50
$x_1 \le 40$	50	$x_1 \le 90$	0
$x_2 \le 0.5x_1$	18	$x_2 \le 0.7x_1$	0
$x_2 \ge 0.3x_1$	0	$x_2 \ge 0.5x_1$	8

A zero means the row is already implied by the other alternative, so it needs no relaxation at all. Four of the eight rows are in that class, which no round-number M would ever discover.

The two relaxation regions, drawn

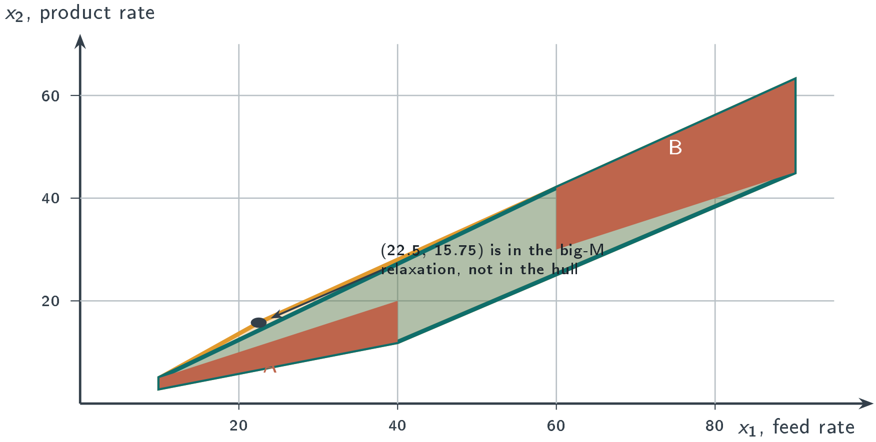

■ true feasible set (A ∪ B) ■ convex hull, area 1100 ■ tightest big-M, area 1137.5

The same picture as numbers

Four formulations of one disjunction, same integer optimum 464.0

Formulation	rows	cols	LP bound	gap %	relaxation area
big-M, $M = 500$	10	4	207.000	55.39	9600.00
big-M, $M = 120$ (bounds)	10	4	229.167	50.61	9133.33
big-M, row-wise tightest	10	4	298.744	35.62	1137.50
convex hull (disaggregated)	16	8	298.744	35.62	1100.00

- The hull region is contained in every big-M region, and the ordering of the areas matches the ordering of the bounds.
- Going from a round $M = 500$ to the row-wise tightest M costs nothing in rows and buys 91.7 units of bound. **Tightening M is the cheapest improvement available.**
- On this instance the hull and the tightest big-M give the same bound in the direction the objective points, even though the hull region is strictly smaller. Domination guarantees "never worse", not "always strictly better".

Formulation strength as a measurable quantity

Definition

$$\text{integrality gap} = \frac{z_{\text{MILP}} - z_{\text{LP}}}{z_{\text{MILP}}} \ (\text{min}), \qquad \frac{z_{\text{LP}} - z_{\text{MILP}}}{z_{\text{MILP}}} \ (\text{max})$$

- Two formulations of the same discrete problem have the same integer optimum *by definition*, as Activity 2 demonstrated. They differ only in their LP relaxation.
- F_1 **dominates** F_2 if the feasible set of its relaxation is contained in that of F_2 on every instance. Domination implies a gap that is never worse.
- The convex hull formulation dominates every big-M formulation of the same disjunction, because its projection is the smallest convex set containing the alternatives.

Why the gap and not the row count

Branch and bound stops when the incumbent and the best remaining bound meet, so the work is roughly a function of how far apart they start. An extra thousand rows that halve the gap are usually a bargain, and the number of rows is a poor proxy for difficulty.

Every formulation in this session, one measure

Smaller is better in the last column, and it is the only column that predicts effort

Instance	Formulation	rows	LP bound	MILP optimum	gap %
20 units, 40 orders	big-M, $M = 10\,000$	80	864.168	3548.408	75.6463
20 units, 40 orders	big-M, $M =$ capacity	60	3509.477	3548.408	1.0971
20 units, 40 orders	disaggregated (hull-derived)	860	3548.340	3548.408	0.0019
two-config. reactor	big-M, $M = 500$	10	207.000	464.000	55.3880
two-config. reactor	big-M, row-wise tightest	10	298.744	464.000	35.6150
two-config. reactor	convex hull (disaggregated)	16	298.744	464.000	35.6150
blending, fixed charges	tight $M_c = \text{avail}_c$	34	379 170.50	281 361.41	34.7628
blending, fixed charges	very loose $M = 10^7$	34	413 430.43	281 361.41	46.9393
zinc-air plant (Activity 2)	big-M on hours only	18	2402.903	2533.226	5.1450
zinc-air plant (Activity 2)	disaggregated (hull)	29	2412.448	2533.226	4.7680

The row count moves by a factor of 86 across this table while the gap moves by a factor of 40 000. They are not the same quantity and they are not correlated. On the 20-unit instance the loose M needs a separate capacity row as well, so it has *more* rows than the tight one and a bound four times worse.

Activity 3 | 20 min | groups of 3 to 4, paper

Task: tighten it

Each constraint below is written with a lazy $M = 10^6$. For each, derive the **tightest valid** M and give the one-line argument, an upper bound on the left-hand side over the rest of the feasible set. One of the three has **no finite valid** M: identify it and say what has to be added to the model.

(a) Zinc-air plant, on/off link for the slot-die coater U3: $x_{\mathrm{U3,cath}} + x_{\mathrm{U3,sep}} \le M\, z_{\mathrm{U3}}$.
U3 has 3 600 h/yr, makes cathode sheet at 0.62 t/h and separator film at 0.48 t/h, and separator demand is 1 500 t/yr.

(b) Refinery blend, the second tank-farm bay is leased only above 8 000 t of throughput:
$\sum_c u_c \le 8\,000 + M\, z_{\mathrm{lease}}$.
Availabilities total 12 700 t; the two grade volume windows are $[4000, 7000]$ and $[2500, 5000]$ t; the octane, vapor pressure and sulfur specifications also apply.

(c) Rented external storage $R_t \ge 0$ (t), usable only if the third-party contract is signed: $R_t \le M\, y_{\mathrm{contract}}$.
R_t appears in $I_t \le W + R_t$ and in the objective at 450 USD/(t month). It carries no upper bound anywhere in the model.

Hand back: three values of M, one sentence of justification each, one sheet per group.

[ANSWER] Activity 3: the three tightest values

(a) $M^\star = 2\,232$ t/yr, and the lazy M was 448 times too large

Maximize the left-hand side over the rest of the feasible set: U3 alone, 3 600 h/yr, arc bounds included. Cathode sheet is the faster product, 0.62 t/h, so all-cathode gives $3\,600 \times 0.62 = 2\,232$ t/yr; separator film is capped at 1 500 t/yr by demand and needs $1500/0.48 = 3\,125$ h to make less product, so no mix beats all-cathode. This is exactly the value computed in cell `Model!$C$37`.

(b) $M^\star = 2\,782.5$ t, and the two cheaper recipes are both loose

recipe	M, t	source of the bound
availability sum	4 700.0	$12\,700 - 8\,000$, ignores the volume windows
grade volume caps	4 000.0	$7\,000 + 5\,000 - 8\,000$, ignores the quality rows
LP over the feasible set	2 782.5	$\max \sum_c u_c = 10\,782.47$, so $M^\star = 10\,782.47 - 8\,000$

The blender cannot reach 12 000 t because the octane floor on Premium and the vapor pressure and sulfur caps limit how much of the cheap high-RVP and high-sulfur streams may enter.

(c) No finite valid M exists, and a bigger number is not the fix

R_t has no upper bound, so $\max R_t$ is $+\infty$ and no finite M leaves the row non-binding when $y = 1$. Whatever number you write becomes a silent, undocumented capacity limit on rented storage. **The fix is physical**: the contract offers at most 500 t per month, so add $R_t \leq 500$, and then $M^\star = 500$ t.

Practice checklist

When you add a binary to a model

1. State the disjunction in words first. If you cannot say what becomes impossible when $z = 1$, the constraint is not ready to be written.
2. Wire every binary to a continuous variable, in both directions if there is a minimum: $Lz \leq u \leq Mz$.
3. Give every M a name and a unit, derived from a bound already in the model.
4. Verify every logical translation on its truth table, not from memory.
5. Solve the LP relaxation and record the gap. It is one extra solve and it is the number that tells you whether the formulation is worth improving.
6. If the gap is large and the model is slow, disaggregate the disjunction before buying a faster computer.

Two smells

A declared parameter that no constraint references, and a binary whose relaxed value is 10^{-3} at the LP optimum. The first is a forgotten constraint; the second is a big M that is too large.

Summary

Takeaways

- A binary models a disjunction; $\{0\} \cup [L, M]$ is the smallest nonconvex set in process models.
- Every propositional statement over binaries is a set of linear inequalities, and the translation is checkable on a truth table.
- A big-M must be valid and tight. Too large costs the bound, and $M\varepsilon$ opens a hole in the model.
- The convex hull form is the tightest relaxation of a disjunction and dominates every big-M form.
- Formulation strength is the integrality gap, and it, not the row count, predicts effort.

Numbers established today

Blending LP / MILP margin	413 439.75 / 342 092.31
Fixed charges / min runs	47 000 / 24 347.44
Full discrete model	281 361.41 (−31.9%)
Big-M bound, tight / 10^7	379 171 / 413 430
$M = 10^{10}$, $z = \varepsilon$	margin 413 439.70
Hull / tight big-M area	1100 / 1137.5
20 units, gap loose/tight/hull	75.6 / 1.10 / 0.002 %
Activity 2 optimum	2533.226 kUSD/yr
A2 bound, $M^\star$ / $10^6 M^\star$	2203.1 / 1843.6

Before next week

Finish

- `w05-modeling4-logical-discrete.ipynb`, all seven sections, including the exercises. Sections 4 and 5 are the ones Activity 2 could not do in a spreadsheet.
- The workbook `w05-fixed-charge.xlsx`, if your group did not reach the integer optimum 2533.226 kUSD/yr.

Read

Williams Ch. 9 (fixed charges, logical conditions, semicontinuous variables) and Ch. 10 (formulation strength). Rao Ch. 10 for the integer-programming vocabulary.

Due next week

HW2 is due at the start of Week 6: blending, pooling, logical and discrete modeling, plus the model-audit task in which you receive a plausible but wrong generated model and must find and correct its errors.

Coming next: Week 6

A three-hour workshop. A language model drafts, the engineer verifies. You will audit three generated models that all solve to `optimal` and are all wrong, and you will quantify what each error would have cost. Bring a laptop with Pyomo and a solver working.

Further reading

- **Williams**, *Model Building in Mathematical Programming*, 5th ed., Ch. 9 (building integer models: fixed charges, logical conditions, semicontinuous variables) and Ch. 10 (formulation strength, the convex hull, and why reformulation beats computation).
- **Rao**, *Engineering Optimization*, 4th ed., Ch. 10, for the standard integer-programming vocabulary and worked formulations.
- **Balas** (1985), *Disjunctive programming and a hierarchy of relaxations for discrete optimization problems*, SIAM J. Alg. Disc. Meth. 6(3): the source of the disaggregated hull formulation used in Lecture 3.
- **Grossmann and Trespalacios** (2013), *Systematic modeling of discrete continuous optimization models through generalized disjunctive programming*, AIChE J. 59(9): GDP as the process systems engineering framing, and the automatic big-M and hull reformulations available in `pyomo.gdp`.
- **Edgar, Himmelblau and Lasdon**, *Optimization of Chemical Processes*, 2nd ed., Ch. 9, for process-scale mixed-integer applications.