Industrial Economics

TA Session 2 — Worked Solutions

Athens University of Economics and Business
Academic Year 2025–2026

Professor: Chrysovalantou Milliou `cmilliou@aueb.gr`
Teaching Assistant: Angel Koulas `koulas@aueb.gr`

About these solutions

This document contains worked solutions to the exercises covered in TA Session 2 (21 May 2026), drawn from **Problem Set III** (Dynamic Oligopoly) and **Problem Set IV** (Collusion). The solutions emphasise both the algebraic steps and the underlying economic intuition, with highlight boxes for the key results.

Contents

PS III, Problem 2 — Stackelberg in Prices, Homogeneous Product

Problem Statement

Sequential price-setting duopoly. Demand $D(p) = 1 - p$; the lowest-price firm captures the whole market and ties are split. *(i)* Same cost function $C(q_i) = c\,q_i$ for both firms, $0 < c < 1$; firm 1 sets p_1 first. *(ii)* $c_1 < c_2$; find the equilibrium prices when $c_2 > p_1^m$ and when $c_2 < p_1^m$, where p_1^m is firm 1's monopoly price.[a]

---

[a] The original problem statement reads "when $c_1 > p^m$ as well as when $c_1 < p^m$". With $c_1 < 1$ and $p^m = (1+c_1)/2$, the condition $c_1 > p^m$ would require $c_1 > 1$, which contradicts the assumption. The natural reading — consistent with the standard small-vs-large asymmetry analysis of Bertrand pricing with asymmetric costs — is to compare c_2 (the rival's marginal cost) with p_1^m (firm 1's monopoly price). We solve it under that reading.

Part (i) — Symmetric costs

We solve by backward induction.

Stage 2. Firm 2 observes p_1 and chooses p_2. Its profit is

$$\pi_2(p_2 \mid p_1) = \begin{cases} (p_2 - c)(1 - p_2) & \text{if } p_2 < p_1, \\ \frac{1}{2}(p_2 - c)(1 - p_2) & \text{if } p_2 = p_1, \\ 0 & \text{if } p_2 > p_1. \end{cases}$$

Two observations:

- For any $p_1 > c$, firm 2 strictly prefers to undercut, since $(p_1 - \varepsilon - c)(1 - p_1 + \varepsilon) > 0$ for ε small.
- For $p_1 \le c$, firm 2 has no profitable response below c (would incur a loss); it stays at $p_2 \ge c$ and earns 0.

Stage 1. Firm 1 anticipates firm 2's reaction.

- $p_1 > c \;\Rightarrow$ firm 2 undercuts, firm 1 sells nothing, $\pi_1 = 0$.
- $p_1 = c \;\Rightarrow$ both firms earn 0 whether firm 2 matches or stays out.
- $p_1 < c \;\Rightarrow$ firm 1 takes a loss.

Hence firm 1 weakly prefers $p_1 = c$. The unique Subgame Perfect Nash equilibrium outcome is

$$p_1^* = p_2^* = c, \qquad \pi_1^* = \pi_2^* = 0.$$

Key Result

With homogeneous products and symmetric costs, sequential price setting yields the *same* outcome as simultaneous Bertrand: $p^* = c$ and zero profits. The Bertrand paradox persists; moving first confers no commitment value.

Economic Intuition

Commitment helps the leader only when it can change the rival's optimal response. Under homogeneous Bertrand, the follower's response to any price above c is always "undercut", so the leader has nothing useful to commit to. As we will see in Q3, introducing differentiation breaks this logic.

Part (ii) — Asymmetric costs ($c_1 < c_2$)

Firm 1's (unconstrained) monopoly price is found from $\max_p (p - c_1)(1 - p)$:

$$p_1^m = \frac{1 + c_1}{2}.$$

The relevant comparison is between p_1^m and the rival's marginal cost c_2.

Stage 2 (firm 2's best response). Identical reasoning as in part (i): firm 2 wants to undercut p_1 only if it can do so without operating below its own MC. Specifically:

- if $p_1 > c_2$, firm 2 undercuts to $p_1 - \varepsilon$;
- if $p_1 \le c_2$, firm 2 sets any $p_2 \ge c_2$ and earns 0.

Stage 1 (firm 1's choice). Two cases:

Case A — "Large asymmetry" ($c_2 \ge p_1^m$). Firm 1's monopoly price p_1^m already lies (weakly) below c_2, so firm 2 cannot profitably undercut it. Firm 1 acts as an unconstrained monopolist:

$$p_1^* = p_1^m = \frac{1 + c_1}{2}, \qquad p_2^* \ge c_2, \qquad \pi_1^* = \frac{(1 - c_1)^2}{4}, \qquad \pi_2^* = 0.$$

Case B — "Small asymmetry" ($c_2 < p_1^m$). At $p_1 = p_1^m$ firm 2 would profitably undercut and capture the market. Firm 1 therefore *limit-prices*: it sets a price just below c_2 so that firm 2 cannot afford to undercut without incurring a loss. Formally, $p_1^* = c_2 - \varepsilon$ with $\varepsilon \to 0^+$, and we summarise the limit:

$$p_1^* \nearrow c_2, \qquad p_2^* = c_2, \qquad \pi_1^* = (c_2 - c_1)(1 - c_2), \qquad \pi_2^* = 0.$$

(Strictly speaking, the open-set issue means there is no exact optimum; the supremum profit $(c_2 - c_1)(1 - c_2)$ is approached as $p_1 \to c_2^-$. This is the standard treatment in the literature.)

Key Result

Under asymmetric costs with firm 1 efficient, the leader either limit-prices at $\approx c_2$ (small asymmetry) or charges its own monopoly price (large asymmetry). In both cases firm 2 earns zero, and the outcome is the same as in the simultaneous Bertrand game with asymmetric costs.

Common Mistake to Avoid

Students sometimes set $p_1^* = c_1$ (Bertrand-symmetric instinct) or $p_1^* = p_1^m$ unconditionally. Both are wrong. The whole point of the asymmetric case is that the constraint imposed by firm 2's marginal cost can bind *strictly* (small asymmetry) or be slack (large asymmetry); compare c_2 with p_1^m before choosing the equilibrium expression.

PS III, Problem 3 — Stackelberg in Prices, Differentiated Products

---

Problem Statement

Two firms produce differentiated goods at common MC $= 0$. Demands: $q_i(p_i, p_j) = 10 - p_i + g\,p_j$, $0 < g < 1$. *(i)* Simultaneous price-setting. *(ii)* Sequential: firm 1 sets p_1 first; compare with (i) and decide whether firms prefer to move first or second.

Part (i) — Simultaneous Bertrand equilibrium

Firm i chooses p_i to maximise $\pi_i = p_i \cdot q_i = p_i(10 - p_i + g\,p_j)$ (since MC $= 0$). The FOC is

$$\frac{\partial \pi_i}{\partial p_i} = 10 - 2p_i + g\,p_j = 0 \implies p_i^{BR}(p_j) = \frac{10 + g\,p_j}{2}.$$

The best-response slope is $g/2 > 0$, so **prices are strategic complements**.

By symmetry, $p_1^* = p_2^* = p^*$ satisfies $p^* = (10 + g\,p^*)/2$, giving

$$p^* = \frac{10}{2-g}, \qquad q^* = 10 - p^* + g\,p^* = \frac{10}{2-g}, \qquad \pi^* = p^*\,q^* = \frac{100}{(2-g)^2}.$$

Key Result

Simultaneous Bertrand with differentiated goods: $p_i^* = q_i^* = 10/(2-g)$, $\pi_i^* = 100/(2-g)^2$.

Part (ii) — Sequential Bertrand

Stage 2 (follower). Firm 2 observes p_1 and chooses p_2 via the same FOC as before, but treating p_1 as given:

$$p_2^{BR}(p_1) = \frac{10 + g\,p_1}{2} = 5 + \tfrac{g}{2}\,p_1.$$

Stage 1 (leader). Firm 1 substitutes the follower's BR into its own demand:

$$q_1 = 10 - p_1 + g\,p_2^{BR}(p_1) = 10 - p_1 + g\Big(5 + \tfrac{g}{2}\,p_1\Big) = (10 + 5g) - p_1\Big(1 - \tfrac{g^2}{2}\Big).$$

Profit (MC $= 0$): $\pi_1 = p_1\Big[(10 + 5g) - p_1(1 - g^2/2)\Big]$. FOC:

$$(10 + 5g) - 2p_1\Big(1 - \tfrac{g^2}{2}\Big) = 0 \implies p_1^S = \frac{10 + 5g}{2 - g^2} = \frac{5(2+g)}{2 - g^2}.$$

Then

$$p_2^S = 5 + \tfrac{g}{2}\,p_1^S = \frac{5\big(4 + 2g - g^2\big)}{2(2 - g^2)}.$$

Equilibrium quantities. Direct substitution (and using the identity $4+2g-2g^2-g^3 = (2-g^2)(2+g)$):

$$q_1^S = \frac{5(2+g)}{2}, \qquad q_2^S = p_2^S = \frac{5(4+2g-g^2)}{2(2-g^2)}.$$

Equilibrium profits.

$$\pi_1^S = p_1^S\, q_1^S = \frac{25(2+g)^2}{2(2-g^2)}, \qquad \pi_2^S = (p_2^S)^2 = \frac{25(4+2g-g^2)^2}{4(2-g^2)^2}.$$

Key Result

Sequential equilibrium prices: $p_1^S = 5(2+g)/(2-g^2)$ and $p_2^S = 5(4+2g-g^2)/[2(2-g^2)]$. Both exceed the simultaneous price $p^* = 10/(2-g)$, but the leader sets a *higher* price than the follower.

Comparison and the second-mover advantage

Numerical check at $g = 1$ (a clean case):

	Simultaneous	Stackelberg leader	Stackelberg follower
Price	10	15	12.5
Quantity	10	7.5	12.5
Profit	100	112.5	156.25

Both firms gain from the Stackelberg structure relative to simultaneous moves (because the leader's commitment pushes both prices up), but the follower's gain is strictly larger.

Key Result

Second-mover advantage: when prices are strategic complements ($g > 0$), the follower's profit exceeds the leader's. Both firms prefer to move *second*.

Economic Intuition

The strategic-substitutes / strategic-complements dichotomy fixes the leader/follower ranking:

- Cournot Stackelberg (quantities): BR slope $< 0 \Rightarrow$ leader gains. Committing to a large quantity forces the follower to retract.
- Bertrand Stackelberg with differentiation (prices): BR slope $> 0 \Rightarrow$ *follower* gains. The leader's high price pulls the follower's price up too, but the follower can undercut just enough to grab market share while still selling at a higher price than under simultaneous play.

PS IV, Problem 1 — Tacit Collusion in Prices (Three Firms & Mergers)

Problem Statement

Three firms compete in prices over $t \to \infty$. Common MC $= 6$; demand $P = 20 - Q$; discount factor δ. *(i)* Static Nash vs. collusion. *(ii)* Condition for tacit collusion. *(iii)* Symmetric merger of firms 1 and 2 (MC $= 6$ for the merged firm). *(iv)* Same merger but with $\text{MC}_M = 0$; collusion at the price a monopolist with MC $= 6$ would charge.

Part (i) — Static Nash and the collusive price

Static Nash. Standard symmetric Bertrand with $n = 3$ firms and common MC $= 6$: $p^N = 6$, $\pi_i^N = 0$.

Collusion at the monopoly price. Maximise $(p-6)(20-p)$. FOC: $20 - 2p + 6 = 0 \Rightarrow p^m = 13$. Then $Q^m = 7$ and

$$\pi^m = (13-6) \times 7 = 49.$$

Sharing the market equally, $\pi_i^c = \pi^m/3 = 49/3$.

For the trigger-strategy analysis, the relevant per-period payoffs for each firm are

$$\pi^c = \tfrac{49}{3}, \qquad \pi^d \approx 49, \qquad \pi^p = 0.$$

(π^d is computed by undercutting to $p^m - \varepsilon$ and serving the whole market.)

Part (ii) — Critical discount factor (3 firms)

The trigger-strategy incentive constraint is

$$\frac{\pi^c}{1-\delta} \;\geq\; \pi^d + \frac{\delta\,\pi^p}{1-\delta},$$

which rearranges to the general formula

$$\delta \;\geq\; \frac{\pi^d - \pi^c}{\pi^d - \pi^p}.$$

Plugging in:

$$\delta \;\geq\; \frac{49 - 49/3}{49 - 0} = \frac{98/3}{49} = \frac{2}{3}.$$

Key Result

Tacit collusion is sustainable iff $\delta \geq 2/3$. With n symmetric firms, this generalises to $\delta \geq 1 - 1/n$. More firms make collusion harder.

Part (iii) — Symmetric merger

After the merger of firms 1 and 2 at common MC $= 6$, we have $n = 2$ symmetric firms. The collusion price and the monopoly profit are unchanged ($p^m = 13$, $\pi^m = 49$). The relevant per-period payoffs are now

$$\pi^c = 49/2 = 24.5, \qquad \pi^d \approx 49, \qquad \pi^p = 0,$$

giving

$$\delta \;\geq\; \frac{49 - 24.5}{49 - 0} = \frac{1}{2}.$$

Key Result

The merger of two symmetric firms lowers the critical discount factor from 2/3 to 1/2: collusion becomes *easier*. Each collusive firm now keeps a larger share of the monopoly pie, so the temptation to undercut is smaller.

Part (iv) — Merger with cost advantage

Now the merged firm (call it M) has $\text{MC}_M = 0$ while the outsider (firm 3, call it O) has $\text{MC}_O = 6$. The collusion price stays at $p^m = 13$ (defined for MC $= 6$), with the market split 50/50.

Per-period collusion profits. At $p = 13$, each firm sells $Q^m/2 = 3.5$:

$$\pi^c_M = 13 \times 3.5 = 45.5, \qquad \pi^c_O = (13 - 6) \times 3.5 = 24.5.$$

One-shot deviation payoffs. Each firm could undercut to $13 - \varepsilon$ and serve $Q^m = 7$:

$$\pi^d_M \approx 13 \times 7 = 91, \qquad \pi^d_O \approx (13 - 6) \times 7 = 49.$$

Punishment payoffs. The static Nash of asymmetric Bertrand with $\text{MC}_M = 0$ and $\text{MC}_O = 6$: since firm M's unconstrained monopoly price $(20+0)/2 = 10$ exceeds $\text{MC}_O = 6$, the asymmetry is *small*. Equilibrium is the limit-pricing outcome: $p_M \approx 6 - \varepsilon$, M captures the entire market $Q = 14$, and

$$\pi^p_M = 6 \times 14 = 84, \qquad \pi^p_O = 0.$$

The merged firm's IC binds. Firm M's payoff from sticking with collusion forever is $45.5/(1-\delta)$. From deviating once and then reverting to static Nash: $91 + \delta \cdot 84/(1-\delta)$. The IC is

$$\frac{45.5}{1-\delta} \;\geq\; 91 + \frac{84\,\delta}{1-\delta} \iff 45.5 \;\geq\; 91 - 7\delta \iff \delta \geq 6.5,$$

which is impossible since $\delta \leq 1$.

Key Result

The merger that gives the merged firm a cost advantage makes tacit collusion *impossible* (the merged firm's IC cannot be satisfied for any $\delta \in [0, 1]$). Compared with part (iii) where collusion was sustainable for $\delta \geq 1/2$, the cost asymmetry destroys collusion entirely.

Economic Intuition

Two forces work against firm M:

- *The deviation prize is large*: undercutting captures the entire monopoly margin of 13 on every unit.
- *The punishment is mild*: in the punishment phase, M still earns a healthy 84 per period from limit-pricing the rival out. There is barely any difference between "cooperating" and "competing" for M.

This is why competition policy looks suspiciously at mergers that create a strong *maverick*: such mergers may either destroy collusion (a pro-competitive effect) or, conversely, give the merging parties unilateral market power. The question depends on which effect dominates given market specifics.

PS IV, Problem 2 — Tacit Collusion in Prices with Asymmetric Costs

---

Problem Statement

Four firms compete in prices, $t \to \infty$. $c_1 = 4$, $c_2 = c_3 = c_4 = 5$. Demand $P = 20 - Q$; discount factor δ. Tacit collusion is at the price a monopolist with MC $= 4$ would set. *(i)* Firm 1's profit in (Nash, deviation, punishment) periods. *(ii)* Condition for tacit collusion.

Part (i) — Firm 1's three period payoffs

Static Nash (limit pricing by firm 1). Firm 1 has the lowest cost ($c_1 = 4$). Its unconstrained monopoly price is $(20 + 4)/2 = 12$, well above the rivals' cost of 5: small asymmetry. The unique equilibrium has firm 1 price just below $c_{-1} = 5$ and capture the whole market. At $p = 5$, $Q = 15$, so

$$\pi_1^p = (5 - 4) \times 15 = 15.$$

The other firms earn zero. We will use this as the per-period punishment payoff.

Collusion period. The collusion price is the monopoly price for MC $= 4$: $p^c = 12$, $Q^c = 8$. Sharing the market equally across 4 firms, each sells $8/4 = 2$. Firm 1's profit:

$$\pi_1^c = (12 - 4) \times 2 = 16.$$

Deviation period. Firm 1 undercuts to $12 - \varepsilon$ and serves the whole market $Q = 8$:

$$\pi_1^d = (12 - 4) \times 8 = 64.$$

Key Result

For firm 1: $\pi_1^c = 16$, $\pi_1^d = 64$, $\pi_1^p = 15$. Note that the punishment phase still pays firm 1 a positive rent (15), because firm 1 is the most efficient producer.

Part (ii) — Critical discount factor

Firm 1's IC:

$$\frac{16}{1-\delta} \geq 64 + \frac{15\,\delta}{1-\delta} \iff 16 \geq 64(1-\delta) + 15\delta \iff 16 \geq 64 - 49\delta.$$

Solving:

$$\delta \geq \frac{48}{49} \approx 0.980.$$

For the high-cost firms ($i \in \{2, 3, 4\}$): they earn $\pi_i^c = (12 - 5) \times 2 = 14$ during collusion, $\pi_i^d = (12 - 5) \times 8 = 56$ during a deviation, and $\pi_i^p = 0$ during punishment. Their IC is

$$\delta \geq \frac{56 - 14}{56 - 0} = \frac{3}{4}.$$

Since $48/49 > 3/4$, the binding constraint is firm 1's:

Key Result

Tacit collusion is sustainable iff $\delta \geq 48/49$. The low-cost firm 1 is the constraint that pins down the critical discount factor.

Economic Intuition

Two reasons firm 1's IC is so tight:

1. Firm 1's deviation profit ($\pi_1^d = 64$) is far above its collusive profit ($\pi_1^c = 16$): four times higher.
2. Firm 1's punishment profit ($\pi_1^p = 15$) is almost as high as its collusive profit (16): punishment is barely a punishment.

By the formula $\delta^* = (\pi^d - \pi^c)/(\pi^d - \pi^p)$, both the large numerator and the small denominator push δ^* close to 1.

Common Mistake to Avoid

A common error is to compute only the symmetric-firm IC and report $\delta \geq 3/4$. Always check *the most efficient firm's IC* when costs are asymmetric — it is typically the binding one.

PS IV, Problem 3 — Tacit Collusion in Quantities (Cournot)

Problem Statement

Two symmetric firms compete in quantities; demand $p(Q) = 1 - Q$ with constant MC c, $0 < c < 1$. Infinite horizon, discount factor δ. *(i)* Collusive output and profit per firm. *(ii)* Condition for tacit collusion.

Part (i) — Joint-monopoly output

Under collusion the cartel acts as a single monopolist:

$$\max_Q (1 - Q - c)\, Q \;\Rightarrow\; Q^m = \frac{1-c}{2}, \;\; p^m = \frac{1+c}{2}, \;\; \pi^m = \frac{(1-c)^2}{4}.$$

Each firm produces half:

$$q_i^c = \frac{Q^m}{2} = \frac{1-c}{4}, \qquad \pi_i^c = \frac{\pi^m}{2} = \frac{(1-c)^2}{8}.$$

For comparison, the static Cournot Nash gives $q_i^N = (1-c)/3$ and $\pi_i^N = (1-c)^2/9$. Collusion produces less and earns more.

Part (ii) — Critical discount factor

Deviation payoff. Firm i takes the rival's collusive output as fixed at $(1-c)/4$ and best-responds:

$$q_i^d = \frac{1 - c - q_j}{2} = \frac{1 - c - (1-c)/4}{2} = \frac{3(1-c)}{8}.$$

Total quantity then is $Q = q_i^d + q_j = 5(1-c)/8$ and the inverse demand gives $p = 1 - 5(1-c)/8 = (3+5c)/8$. Firm i's profit is

$$\pi_i^d = (p - c)\, q_i^d = \frac{3(1-c)}{8} \cdot \frac{3(1-c)}{8} = \frac{9(1-c)^2}{64}.$$

Punishment payoff. Reverting to the static Cournot Nash forever: $\pi_i^p = (1-c)^2/9$.

IC. Substituting into $\pi_i^c/(1-\delta) \geq \pi_i^d + \delta\, \pi_i^p/(1-\delta)$ and dividing through by $(1-c)^2$:

$$\frac{1/8}{1-\delta} \;\geq\; \frac{9}{64} + \frac{\delta/9}{1-\delta}.$$

Multiply both sides by $576\,(1-\delta)$ (LCM of 8, 64, 9):

$$72 \;\geq\; 9 \cdot 9\,(1-\delta) + 64\,\delta \;=\; 81 - 17\,\delta,$$

yielding

$$\delta \;\geq\; \frac{9}{17} \approx 0.529.$$

Key Result

Cournot tacit collusion is sustainable iff $\delta \geq 9/17$. Compared with the corresponding Bertrand condition $\delta \geq 1/2$ (part (iii) of Q1, $n = 2$), Cournot collusion is *harder* to sustain.

Economic Intuition

The Cournot vs. Bertrand contrast:

- In Bertrand, deviation captures the *entire* market, so $\pi^d/\pi^c = 2$ (with $n = 2$). The deviation gain is large.
- In Cournot, deviation only captures a third more output, so $\pi^d/\pi^c = 9/8$. The deviation gain is small.

One might expect Cournot collusion to be easier. But the punishment phase also pays positive profits under Cournot ($(1-c)^2/9 > 0$), whereas Bertrand punishment is zero. The weaker punishment threat dominates, making Cournot collusion harder.

Summary: the trigger-strategy toolkit

	Bertrand (homogeneous)	**Cournot (symmetric)**
Static Nash payoff π^p	0	$(1-c)^2/9$
Collusion payoff π^c	π^m/n	$\pi^m/2$
Deviation payoff π^d	$\approx \pi^m$	best response to $\pi^m/2n$
Critical δ^*	$1 - 1/n$	$9/17$ ($n = 2$)

General rule. The critical discount factor is $\delta^* = (\pi^d - \pi^c)/(\pi^d - \pi^p)$. Collusion is easier when (i) π^c is high (sharing makes the cartel attractive), (ii) π^d is low (undercutting is unappealing), and (iii) π^p is low (punishment is severe).

Asymmetry warning. When firms differ in costs, *always* check the IC of the most efficient firm: it has the biggest incentive to deviate and the smallest effective punishment. That firm pins down the critical δ.