INTERNATIONAL ECONOMICS OLYMPIAD

2026 Economics Open Questions

Solutions and Marking Schemes

Open Question 1. "The Model That Could Not Cross the Border"

Frontier AI systems are dual-use technologies: they may help firms improve productivity, but at the same time they may help foreign users discover and exploit vulnerabilities or attempt to replicate some of the model's capabilities or extract sensitive information from it.

Country A's firm AlphaAI owns the best AI model and sells monitored access to entities in Country B (one unit = one standardized access contract during the policy period).

Country B's inverse demand is $P = 100 - Q$. AlphaAI's marginal cost is 20 and it is a foreign monopolist in Country B.

The government of A believes exports create an expected national-security external cost $E(Q) = 8 \cdot Q + Q^2/2$ (where $8 \cdot Q$ captures an ordinary misuse risk, while $Q^2/2$ is due to risks rising disproportionately with scale, e.g. jailbreaks).

Country A defines welfare as: AlphaAI profit + government license revenue − external cost.

(a) (8 points) Suppose exports are unregulated. With unregulated exports, find AlphaAI's profit, the external cost, and Country A's welfare.

Is a complete export ban better for Country A compared to unregulated exports? Explain.

(b) (8 points) Instead of a ban, suppose Country A requires AlphaAI to buy one export license for every unit of access sold abroad. Find the per-unit license fee that maximizes Country A's welfare. Compare this policy with both unregulated exports and a complete ban on profit, government revenue, welfare, and the external cost.

(c) (4 points) Instead, Country A imposes an export quota equal to the quantity found in part (b). Consider two methods of allocating the quota rights:
(i) the rights are given to AlphaAI for free;
(ii) AlphaAI pays the government, for each unit of quota used, the per-unit fee found in part (b).

Under each method, calculate AlphaAI's profit, government revenue, Country A's welfare, and the external cost. Identify which parameters differ between the two methods and which remain unchanged.

(d) (5 points) Alternatively AlphaAI sells a copy of the model weights outright to a foreign firm, for a one-time payment of 900. There would be no license revenue afterwards. Country A's government estimates that this sale would create an external cost of 700.

Would AlphaAI prefer this to the licensing like in (b)? Would Country A prefer it? Explain.

(e) (5 points) After Country A restricts access, Country B builds its own substitute AI model. Give one potential benefit and two potential costs/risks of this response for Country A.

Solution

(a) AlphaAI maximizes $\pi = (100 - Q)Q - 20Q$. MR = MC gives $100 - 2Q = 20$, so $Q = 40$, $P = 60$.

- Profit: $(60 - 20)\cdot 40 = 1600$
- External cost: $E(40) = 8\cdot 40 + 40^2/2 = 320 + 800 = 1120$
- Welfare: $1600 + 0 - 1120 = 480$

Under a complete ban, Q = 0, so profit, revenue, and external cost are all zero: W_ban = 0. Since 480 > 0, unregulated exports are better than a ban: even at the monopoly quantity, the profit captured by the domestic firm exceeds the security externality. (Note B's consumer surplus is irrelevant — Country A's welfare definition excludes foreign consumers.)

(b) With a fee t, AlphaAI's effective marginal cost is $20 + t$, so $100 - 2Q = 20 + t$ gives the reaction $Q(t) = 40 - t/2$, i.e., $t = 80 - 2Q$.

License revenue tQ is a transfer within Country A, so A's welfare in terms of Q is:

$$W(Q) = (P - 20)Q - E(Q) = (80 - Q)Q - 8Q - Q^2/2 = 72Q - (3/2)Q^2$$

Maximizing: $72 - 3Q = 0 \rightarrow Q^* = 24$, hence $t^* = 80 - 48 = 32$, and $P = 76$.

(Equivalent valid method: because the fee is a pure internal transfer and the firm already maximizes the profit component of A's welfare, the optimal fee is exactly Pigouvian at the optimum: $t = E'(Q) = 8 + Q$ solved with the firm's maximization $80 - 2Q = t$ gives $Q = 24$, $t = 32$.)

Outcomes at Q = 24:

	Unregulated	License fee t = 32	Ban
Profit	1600	$(76 - 20 - 32)\cdot 24 = 576$	0
Gov. revenue	0	$32\cdot 24 = 768$	0

External cost	1120	$8 \cdot 24 + 288 = 480$	0
Welfare	480	$576 + 768 - 480 = 864$	0

The license dominates both alternatives on welfare ($864 > 480 > 0$). It cuts the external cost by more than half relative to laissez-faire while preserving a positive, monitored export flow; the ban minimizes the externality but destroys all surplus. Profit falls from 1600 to 576 — the firm loses, the country gains.

(c) The quota binds (unconstrained optimum is $40 > 24$), so under both methods $Q = 24$, $P = 76$, $E = 480$ — identical to (b).

(i) Free allocation: $\pi = (76 - 20) \cdot 24 = 1344$;

government revenue = 0; $W = 1344 - 480 = 864$.

(ii) Paying 32 per unit used: $\pi = 1344 - 768 = 576$;

revenue = 768;

$W = 576 + 768 - 480 = 864$ (identical to (b) in every respect).

What differs: only the division of the quota rent of 768 between the firm and the government (profit and government revenue).

What is unchanged: quantity, price, external cost, and total welfare — the allocation method is purely distributional because the quota, not the payment, determines behavior at the margin.

(d) AlphaAI's view. The relevant baseline is the licensing regime of (b), under which AlphaAI earns 576. The sale pays $900 > 576$, so AlphaAI prefers to sell the weights.

Country A's view. The 900 accrues to a domestic firm and counts in A's welfare; there is no license revenue; the external cost is 700:

$W_{sale} = 900 + 0 - 700 = 200 < 864 = W_{licensing}$.

Country A prefers licensing. The interests diverge because AlphaAI internalizes neither the increase in the external cost (700 vs 480) nor the loss of government license revenue (768).

A structural point worth credit: selling weights is irreversible and unmonitored — the "monitored access" structure that made controlled exporting relatively safe disappears, which is exactly why E jumps despite the transaction being privately profitable.

(e) One potential benefit: the security externality from A's own exports shrinks or disappears — B's users no longer probe, jailbreak, or attempt capability extraction from AlphaAI's model, so the risk channel $E(Q)$ tied to A's model is closed (and A retains its frontier weights domestically).

Two potential costs/risks:

1. Loss of revenue and market position: AlphaAI's profit and A's license revenue from market B may fall substantially or disappear if Country B substitutes away from AlphaAI, and B's substitute may compete with AlphaAI in third-country markets, eroding A's technological and commercial lead over time.
2. Loss of oversight and leverage: the risk does not vanish, it migrates. B's model operates entirely outside A's monitoring, licensing, and safety controls, so misuse risk may end up higher than under monitored access — and A loses the policy lever (access conditionality) it previously held over B.

(Other creditable answers: B's model may be built with weaker safety standards; restriction accelerates B's indigenous capability development, shortening A's lead; reduced visibility into B's capability level.)

Marking Scheme — Open Question 1. "The Model That Could Not Cross the Border" (30 points)

General provisions: any correct alternative method earns full credit (in (b), the Pigouvian route $t = E'(Q)$ combined with the firm's response is fully equivalent to direct welfare maximization). Error carried forward applies throughout: downstream values consistent with the contestant's own earlier numbers earn credit, unless the error qualitatively simplified the logic (e.g., eliminated the quadratic risk term or the trade-off itself), in which case the affected criterion is capped at half. A bare correct number with no visible method earns at most half of the corresponding criterion. Half-points are permitted wherever the ladder implies them.

Part (a) — 8 points

#	Pts	Criterion	Evidence and partial credit
a1	3	AlphaAI's unregulated monopoly profit equals 1600.	Full credit requires a correct monopoly-profit calculation. $Q = 40$ and $P = 60$ are supporting steps and are not separately required. Correct condition with an arithmetic slip in Q, P, or profit: 2. Profit function or FOC set up but not solved: 1.

			Perfect-competition pricing $P = 20$, $Q = 80$, or any non-monopoly setup: 0 here; downstream by ECF but a3 capped at 2.
a2	2	External cost $E(40) = 1120$ and welfare $W = 480$ under A's definition	Both terms of E computed (320 + 800) and $W = 1600 + 0 - 1120$. One of the two values wrong or E missing the $Q^2/2$ term with W consistent: 1. Using marginal external cost $E'(Q) = 8 + Q$ in place of $E(Q)$: 0, downstream by ECF. Adding B's consumer surplus or a DWL term to W: 0 for the welfare half (the definition is given in the stem), so at most 1.
a3	3	Correct ban preference with a relevant economic explanation.	Full credit requires the conclusion that unregulated exports are preferable and a relevant explanation: the profit captured by the domestic firm exceeds the security externality, so a ban destroys more surplus than the risk it removes; Explicitly calculating $W_{ban} = 0$ or using the specific profit-versus-externality wording is not separately required. Correct ranking with $W_{ban} = 0$ shown but no economic content: 2. Ranking asserted without computing the ban outcome, or explanation resting on foreign consumers: 1. Ranking contradicting the contestant's own numbers: 0. Ranking internally consistent with erroneous own W: graded on the same ladder via ECF.

Traps (a):

welfare as total surplus including B's consumers — loses the welfare half of a2 and caps a3 at 1;
comparing profit and external cost verbally without computing W — a2 at most 1, a3 at most 1;
E vs E' confusion — a2 = 0, a3 by ECF.

Part (b) — 8 points

#	Pts	Criterion	Evidence and partial credit

b1	3	Correct structure: firm's response to the fee and the welfare objective with the transfer logic	$Q(t) = 40 - t/2$ (or $t = 80 - 2Q$) and $W(Q) = (P - 20)Q - E(Q) = 72Q - (3/2)Q^2$, with tQ recognized as an internal transfer explicitly or evidently from the algebra. Equivalent: Pigouvian condition $t = E'(Q)$ at the optimum combined with the firm's FOC. Firm response correct but naive Pigouvian fee at the wrong quantity ($t = E'(40) = 48$): 2. Firm response correct, objective wrong but a trade-off preserved: 1. License revenue omitted from welfare (corner at $t = 0$) or government maximizing tQ: 1 for the firm response only; the rest of the part by ECF with b3 capped at 1.
b2	2	$Q^* = 24$ and $t^* = 32$ ($P = 76$ accepted as evidence of t)	Both values from a structurally correct objective. One value wrong or an arithmetic slip: 1. Both by ECF from a wrong-but-trade-off-preserving objective: graded against the contestant's own algebra on the same ladder.
b3	3	Outcomes and three-way comparison: profit 576, revenue 768, external cost 480, welfare 864; ranking 864 > 480 > 0 with substance	Full credit: all four values (profit net of the fee: $(76 - 20 - 32)\cdot 24$) and the ranking with at least one substantive observation (external cost reduced vs. laissez-faire; the firm loses while the country gains; the ban minimizes E but destroys all surplus). One value wrong, or ranking with no comparison on the other requested parameters: 2. Gross-of-fee profit 1344 counted together with revenue 768 (double-counting the transfer): 1, since profit and welfare are both wrong. Values only, no comparison at all: 1. All by ECF against the contestant's own Q and t.

Traps (b): naive Pigouvian $t = 48$, $Q = 16$ — b1 = 2, b2 and b3 by ECF, so up to 6 for the part;

revenue-maximizing fee $t = 40$ — b1 = 1, rest by ECF with b3 capped at 1; subtracting the fee from welfare as if it were a social cost — same cap as revenue omission;

solving the planner's problem while ignoring that the firm chooses Q given t earns full credit only if the correct fixed point $Q = 24$, $t = 32$ is reached, in which case the method is equivalent.

Part (c) — 4 points

#	Pts	Criterion	Evidence and partial credit
c1	2	Both allocation methods computed at Q = 24: (i) π = 1344, revenue 0, W = 864; (ii) π = 576, revenue 768, W = 864	Quota bindingness (40 > 24) may be shown or simply used; using Q = 24 with P = 76, E = 480 suffices as evidence. All six values correct: 2. One method fully correct, the other with an error: 1. Firm treated as unconstrained (Q = 40 despite the quota): 0, and c2 by ECF only if the contestant's comparison is internally coherent. Values by ECF from the contestant's own (b) figures.
c2	2	Correct identification of which requested outcomes differ and which remain unchanged; quantity, price, external cost, and welfare are unchanged, because the quota, not the payment, determines behavior at the margin	Full credit requires stating that profit and government revenue differ, while external cost and welfare remain unchanged. The pure-transfer mechanism and discussion of quantity or price are supplementary, not required. Correct sorting without the reason: 1. Claim that welfare differs between methods, or that the payment in (ii) is a deadweight cost: 0.

Traps (c): asserting method (i) is better for Country A because the firm keeps more — c2 = 0, the transfer is internal either way.

Part (d) — 5 points

#	Pts	Criterion	Evidence and partial credit
d1	2	Firm's choice: baseline is licensing profit 576 (own (b) figure accepted), 900 > 576, AlphaAI prefers the sale	Full credit: correct baseline and conclusion. Baseline taken as unregulated profit 1600 with the comparison then done correctly on that (wrong) basis: 1 — the stem specifies "the licensing like in (b)". Conclusion inconsistent with the contestant's own comparison: 0.

d2	3	Country A's preference with a relevant explanation — the firm internalizes neither the higher external cost (700 vs 480) nor the lost license revenue (768); equally creditable: the sale is irreversible and unmonitored, eliminating the safety structure that made access exports controllable	Full credit requires the correct conclusion that Country A prefers licensing and a relevant numerical or qualitative explanation. Computing sale welfare of 200 is sufficient but not mandatory. A separate explanation of firm-country divergence, irreversibility, or monitoring is supplementary. Welfare and comparison correct, no explanation: 2. Welfare structure correct with one element misplaced (e.g., 900 omitted) and the rest consistent: 2. Comparison of external costs only (700 vs 480) with an explanation but no welfare computation: 1. Netting the 900 out of A's welfare "because the money comes from abroad": 0 for the welfare component (the inflow accrues to a domestic firm), explanation still creditable: 1.

Traps (d): "AlphaAI is greedy" or arguments via foreign consumers — no explanation credit; concluding for Country A from the external-cost comparison alone — d2 at most 1.

Part (e) — 5 points

Only the first benefit and the first two costs presented are graded, in order of appearance. Each item requires identification of an effect relevant to Country A and a causal mechanism.

The two costs must operate through distinct channels; two restatements of lost revenue or deand count once.

#	Pts	Criterion	Evidence and partial credit
e1	2	One benefit for Country A in the post-entry situation with a	Creditable: the externality channel tied to A's model closes — B's users no longer probe, jailbreak, or extract capabilities from AlphaAI's model, so the E(Q) risk falls and A retains its frontier weights; or reduced dependence of A's security posture on monitoring foreign usage. Relevant effect named without a mechanism: 1.

		causal mechanism	Benefit accruing to B (technological autonomy), or generic "competition is good" with no channel to A's welfare: 0.
e2	3	Two distinct costs/risks for Country A, each with mechanism	Creditable channels: permanent loss of profit and license revenue, plus possible competition from B's model in third markets eroding A's lead; risk migration and lost leverage — B's model operates outside A's monitoring and licensing, total misuse risk may rise, and A loses access conditionality as a policy instrument; weaker safety standards in B's model; accelerated indigenous capability development; reduced visibility into B's capabilities; AlphaAI's established brand, reputation, or first-mover advantage may reduce initial customer switching after Country B enters because consumers are already familiar with and trust AlphaAI. Both items with mechanisms: 3. One full item with mechanism: 2. Two labels, or one full item only partially explained: 1. Items framed from B's perspective, or duplicating one channel twice: the duplicate/misframed item earns 0. Contradictory items (the externality both disappears and rises, channels not distinguished) are not evaluated..

Traps (e): "B's model might be dangerous" with no link to why this is worse for A than monitored access — label credit only;

a third, stronger cost listed after two weaker ones — not regarded, the first two stand.

a second, stronger benefit listed after weaker one — not regarded, the first stands.

Open Question 2. Mutual funds, fees, and closet indexing

A mutual fund pools money from many investors and invests it on their behalf. Some mutual funds are passively managed: they simply track a benchmark, such as a stock index. Other mutual funds are actively managed: the fund manager tries to beat the benchmark by choosing investments.

Mutual fund managers are often paid a management fee equal to a percentage of the fund's assets under management, or AUM.

Suppose investors are risk neutral. They care only about expected return, not about risk. Each investor can always put money into a passive benchmark fund that earns zero excess return. Therefore, an active mutual fund can attract money only if investors expect its net excess return, after management fees, to be at least zero. If expected net excess return is positive, investors are willing to supply unlimited funds to the mutual fund. If expected net excess return is negative, no investor is willing to invest. In equilibrium, if the fund receives any money, its expected net excess return must be exactly zero.

Suppose a fund manager has skill, but there are decreasing returns to scale. If the manager actively manages x billion dollars, the fund generates total gross excess return: $G(x) = 0.08x - 0.005x2$ (billions per period).

(a) (12 points) Suppose the manager freely sets a fee rate and actively invests all money.

(i) If the fund has x billion dollars under management, what is the maximum fee rate $f(x)$ that investors are willing to pay?

(ii) Which x maximises the manager’s fee revenue $f(x)\cdot x$?

(iii) What fee rate does the manager charge, and how much money does the manager earn in one period?

(b) (12 points) Now suppose regulation or market convention fixes the fee at 2 percent of AUM. The manager can actively manage only x billion dollars and put the rest of the fund into the passive benchmark. This is called closet indexing: the fund charges active-management fees on all AUM, while some money is invested passively. Find:

(i) If the manager actively manages x billion dollars, what is the maximum total AUM $Q(x)$ that the fund can attract?

(ii) Which x maximises the manager’s fee revenue?

(iii) How much total AUM does the fund attract?

(iv) How much money is invested passively?

(v) How much does the manager earn in fees?

(c) (6 points) Suppose a fund had very high excess returns during the last period, so a rational investor infers the manager is highly skilled.

Should the investor expect positive net excess returns next period? Explain.

Solution

(a) (i) Investors supply funds until expected net excess return is zero. With AUM x and fee rate f, net excess return is $G(x) - f\cdot x$, so the maximum acceptable fee rate satisfies $f(x)\cdot x = G(x)$:

$f(x) = G(x)/x = 0.08 - 0.005x$.

(ii) Fee revenue is $f(x)\cdot x = G(x) = 0.08x - 0.005x^2$. The manager therefore maximizes the gross excess return itself: $G'(x) = 0.08 - 0.01x = 0$ gives $x = 8$ (billion). Second-order condition holds since G is concave.

(iii) $f(8) = 0.08 - 0.04 = 0.04$, i.e., a 4% fee rate. Earnings: $f \cdot x = 0.04 \cdot 8 = 0.32$ billion, i.e., \$320 million per period. Note the interpretation: the zero-net-return condition means the manager extracts the entire value of her skill, $G(8) = 0.32$, as fees; investors get exactly the passive benchmark return.

(b) (i) The fund actively manages x and puts $Q - x$ into the passive benchmark, which contributes zero excess return. Gross excess return is still $G(x)$, but fees are charged on all AUM: $0.02Q$. The participation condition $G(x) - 0.02Q = 0$ gives:

$Q(x) = G(x)/0.02 = 50 \cdot (0.08x - 0.005x^2) = 4x - 0.25x^2$.

(ii) Fee revenue is $0.02 \cdot Q(x) = G(x)$ — identical to the objective in (a). Maximization gives $x = 8$ again.

(iii) $Q(8) = 32 - 16 = 16$ billion total AUM.

(iv) Passively invested: $Q - x = 16 - 8 = 8$ billion, i.e., half the fund is closet-indexed.

(v) Fees: $0.02 \cdot 16 = 0.32$ billion — \$320 million, exactly as in (a).

The economic point: the fee cap does not reduce the manager's income. Because investors' participation constraint ties total fees to total gross excess return (fees = $G(x)$ in equilibrium either way), the manager who cannot raise the fee rate instead inflates the base — attracting excess AUM and parking it in the index. Closet indexing is the equilibrium response to fee regulation, not a scam that fools investors: investors still earn exactly zero net, they merely hold a diluted product. The distortion is real but distributional-neutral here; what changes is fund size (16 vs 8), not who gets what.

(c) No. The investor should expect a net excess return of exactly zero. Inferring high skill from past returns is rational, but skill is a scarce resource whose rents accrue to its owner — the manager — not to investors. Upon the good signal, capital flows into the fund; because of decreasing returns to scale, gross excess return per dollar falls (and/or the manager raises the fee or dilutes with passive holdings) until the marginal investor again earns zero net. Past performance predicts fund flows and manager compensation, not future net returns to investors. The absence of net-of-fee return persistence in the data is consistent with managerial skill, because competition among investors, not lack of skill, drives net returns to zero.

Marking Scheme (30 points)
Global rules

- **Binding ECF (error carried forward) rule:** Penalize an error at the criterion testing that error, but do not automatically zero later ECF-authorized criteria. Grade each later operation using the contestant's own stated, coherent values or model; full downstream credit is allowed when that operation is correctly and consistently performed.
- A bare final number without visible setup earns only the answer point of that criterion, never method points.
- Units: values may be given in billions (0.32) or millions/dollars (320 million); either is accepted if internally consistent. Fee rates as 0.04 or 4% equally accepted.

Part (a) — 12 points

#	Pts	Criterion	Accepted evidence and partial credit
a1	2	Zero-net-return participation condition stated and used	$G(x) - f \cdot x = 0$, or "fee revenue equals gross excess return", or net return per dollar $G(x)/x - f = 0$. Stating "investors need net ≥ 0" is full points if later one use equality to find equilibrium, partial credit 1pt if not
a2	2	$f(x) = 0.08 - 0.005x$	2 pts for correct expression; ECF (error carried forward) from a1. Writing $f(x) = G(x)$ (level, not rate): 0 here, see trap T-a1
a3	2	Recognition that fee revenue $f(x) \cdot x = G(x)$, or equivalent correct objective	Explicit simplification, or maximizing $f(x) \cdot x$ written out correctly. Maximizing the correct expression without noticing the equivalence: full credit
a4	2	$x = 8$ from a correct FOC	$0.08 - 0.01x = 0$; ECF from a3. FOC set up correctly with arithmetic slip: 1 pt Give 1pt partial credit if x=8 stated at some point, but then incorrect x is given.
a5	2	f = 4%	From their f(x) at their x; ECF
a6	2	Earnings = 0.32 billion ($320m)	ECF; must be $f \cdot x$ or $G(x)$ at their values

Common traps (a):

- T-a1: Confusing the fee rate with the fee level—writing ($f(x)=0.08x-0.005x^2$)—gives **a2 = 0**.
  If the contestant then maximizes
  ($f(x)x=(0.08x-0.005x^2)x$), the objective is incorrect, so **a3 = 0**. This downstream ECF applies even though the original rate-versus-level error is conceptual. Grade a4–a6 using ECF when the calculations are internally consistent. If a4 got no credit so far, then award: 1 point if ($x=8$) is stated or derived anywhere before or alongside the incorrect final value.
- T-a2: Maximizing G(x) net of some invented cost, or maximizing investors' return: wrong objective, a3–a4 = 0.
- T-a3: Setting f equal to the marginal product $G'(x) = 0.08 - 0.01x$: a2 = 0; downstream by ECF (their x will differ; check consistency).
- T-a4: x in wrong units (8000 vs 8): treat as a units slip, one point off the first affected criterion, ECF thereafter.

Part (b) — 12 points

#	Pts	Criterion	Accepted evidence and partial credit
b1	3	Participation condition with fees on all AUM but returns from active share only: $G(x) - 0.02Q = 0$	3 pts for the correct condition; 1 pt only if fees are applied to x instead of Q (this is the central conceptual step of the part)
b2	2	$Q(x) = 50 \cdot G(x) = 4x - 0.25x^2$	ECF from b1
b3	2	Objective $0.02 \cdot Q(x) = G(x)$; $x = 8$	1 pt for a correct objective, 1 pt for stating x = 8, even without derivation; ECF. Recognizing equivalence with (a) and citing it: full credit without re-derivation
b4	2	$Q = 16$	ECF from b2–b3
b5	1	Passive amount = $Q - x = 8$	ECF
b6	2	Fees = $0.02 \cdot 16 = 0.32$; earnings equal (a)	2 pt for the number (ECF)

Common traps (b):

- T-b1: Charging the 2% fee only on actively managed money (net condition $G(x) - 0.02x = 0$). This deletes the closet-indexing mechanism entirely: b1 max 1 pt, b2 = 0, but grade b3–b6 by ECF on their (degenerate) structure.
- T-b2: Requiring zero net return per actively managed dollar rather than per dollar of AUM: same family as T-b1, same treatment.
- T-b3: Assuming the passive share also earns 0.08 or some positive excess return: b1 = 0; downstream ECF.
- T-b4: Maximizing Q instead of fee revenue (gives x = 8 here anyway, since 0.02Q is monotone in Q): the objective happens to be an affine transformation, so the FOC coincides. Award b3 in full only if the stated objective is fee revenue or the contestant notes maximizing Q is equivalent because fees are proportional to Q; a bare "maximize AUM" with no link to fees: 1 pt of b3.

Part (c) — 6 points

#	Pts	Criterion	Accepted evidence and partial credit
c1	2	Correct answer: no — expected net excess return is zero	An unambiguous "no/zero" tied to the next period. "No" with no or incorrect reasoning at all: 1 pt
c2	3	Flow mechanism: inferred skill attracts inflows; decreasing returns to scale push gross return per dollar down until net return is again zero	Both elements needed for 3 pts: (i) capital inflow in response to the signal, (ii) decreasing returns (or fee/dilution adjustment) as the equalizing force.

c3	1	Incidence of skill rents: the manager, not investors, captures the value of skill (via larger fee base or higher fee); equivalently, past performance predicts flows/manager income, not investor net returns	A clear statement of who captures the rent from skill. Restating c2 in different words without the incidence point: 0

Common traps (c):

- T-c1: "Yes, skill persists, so returns persist" — conflates gross and net returns. c1 = 0; c2–c3 still gradable if the contestant correctly describes the flow mechanism but draws the wrong conclusion from it (rare; typically all three are lost).
- T-c2: "No, because past performance was luck / mean reversion" — contradicts the stem, which stipulates the inference of skill is rational. c1 max 1 pt; c2–c3 = 0 (wrong mechanism).
- T-c3: "No, because of the EMH / markets are efficient" with no fund-flow mechanism: c1 earned, c2 max 1 pt (an equilibrating force is gestured at but not the one in this model), c3 = 0 unless rent incidence is stated.

Total: 12 + 12 + 6 = 30.

Open Question 3. Designed here, assembled there

A finished chip requires one design blueprint (D) and one manufactured component (C): chips = min(D, C). Each worker performs one activity, and may divide working time between activities. Per period:

- an engineer produces 2 blueprints or 1 component;
- a machinist produces 1 blueprint or 3 components.

There are two countries in the world: Northia and Fabrica. Northia has 60 engineers and 20 machinists. Fabrica has 20 engineers and 60 machinists. Countries consume only finished chips, assembled domestically.

(a) *(6 points)* Each country lives in autarky. Find the maximum number of chips each country can produce. Explain why the opportunity cost of blueprints in terms of components is increasing in each country, even though every individual worker has constant productivity.

(b) *(7 points)* Blueprints and components can now be traded freely, but labor cannot migrate. Find the maximum world chip output and describe who does what. Is each country fully specialized? Explain why or why not.

(c) *(9 points)* Northia's parliament, worried that "we design brilliantly but have lost the ability to make things," caps component imports, forcing some engineers into manufacturing. The government announces: "The policy created new jobs." Evaluate three claims without detailed calculation:

(i) the claim that employment was increased;
(ii) a Fabrica newspaper's claim that Northia's quota will *help* Fabrica's design sector;
(iii) Northia's claim that the quota is insurance: "if trade is ever cut off completely, we will be better prepared." This is not true *within this model*, but in reality, this argument can make sense. What real-world feature, absent from the model, makes this true?

(d) *(8 points)* The production function chips = min(D, C) assumes one blueprint is "used up" per chip. In the real semiconductor industry, a single design is copied onto millions of chips at almost zero extra cost, while manufacturing capacity requires enormous investment. Suppose we replaced the assumption with: *one blueprint is enough for any number of components, but designing a blueprint requires 10 engineers*. Without calculations, explain:

(i) how the world's optimal allocation of workers between design and manufacturing would change;
(ii) why, under this technology, a country specializing in design might capture a share of world income far larger than its share of world labor.

Model Solution

(a) According to the workers' advantages, engineers design and machinists manufacture. To achieve D = C, we may have to shift some of the workers towards their less productive activity.

Northia: from (120 blueprints, 60 components), shift x engineers to components: 120 − 2x = 60 + x → x = 20, **80 chips**.

Fabrica: 40 + m = 180 − 3m → m = 35, **75 chips**.

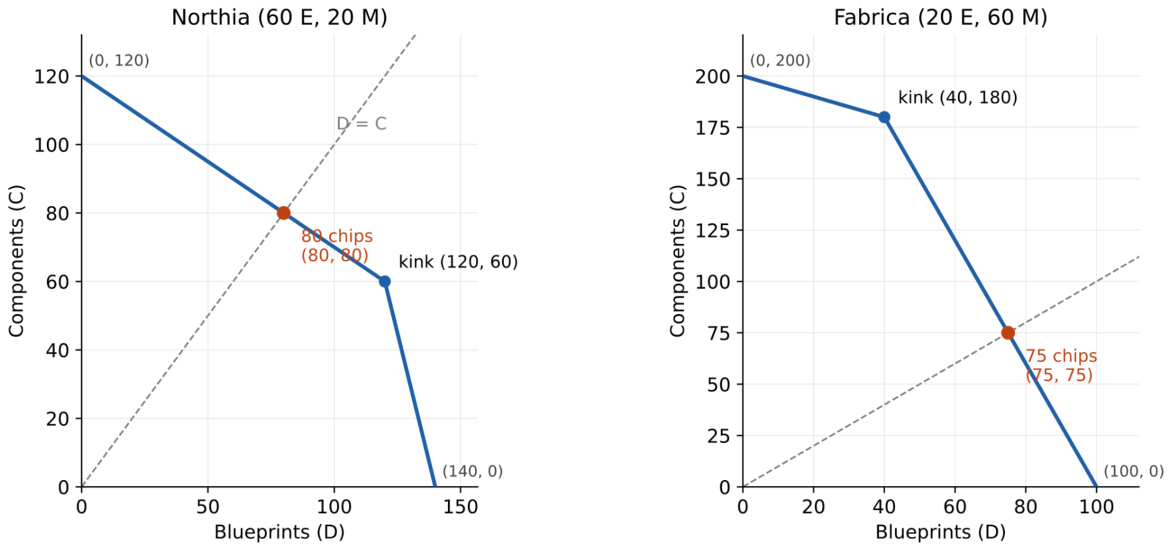

Increasing opportunity cost: blueprints are supplied first by engineers (0.5 component per blueprint), then by machinists (3 components per blueprint) — the PPF is kinked because *resources* are heterogeneous, not because any worker's productivity diminishes.

(b) Pool the workforce: 80 engineers design (160 blueprints), 80 machinists manufacture (240 components); move m machinists to design: 160 + m = 240 − 3m → m = 20 → **180 chips** (world gain 25 over autarky's 155). All engineers design; 20 machinists must also design. The location of the 20 machinists is not unique. They may come from Fabrica, Northia, or both. Country-level specialization depends on which machinists are selected.

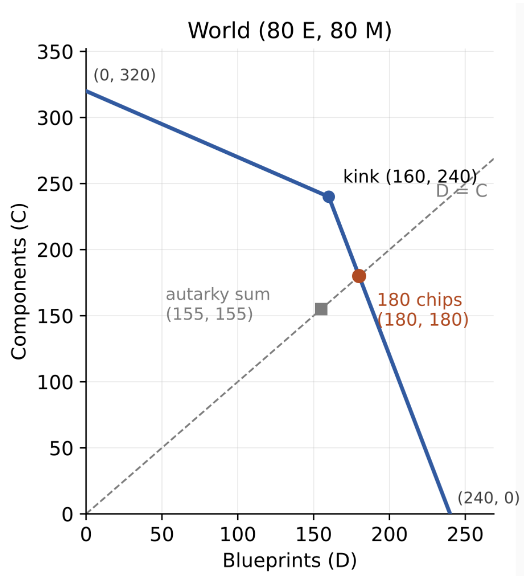

(c) (i) False. Both countries have full employment throughout; the quota *relocates* engineers into manufacturing, where each produces 1 component instead of 2 blueprints — output falls, no employment is created.
(ii) True. Fabrica now receives fewer blueprints and must shift additional machinists into design; Northia's protectionism is an involuntary subsidy to Fabrica's "design sovereignty."
(iii) In the real world, reallocation is slow and capacity may be irreversible (factories take years to build; designers do not become production technicians overnight). So, if the trade is cut off, it may turn out to be difficult to reorganize production quickly.

(d) (i) Design becomes a *fixed input*: the world needs only enough designers to produce one blueprint (or one per chip model), and virtually all labor should move into manufacturing — components, not blueprints, become the binding constraint on output. The very notion of a design–manufacturing PPF trade-off collapses: design is a one-time cost, manufacturing scales with labor.
(ii) A replicable blueprint is a non-rival input with increasing returns: whoever owns the design can license it for every chip produced worldwide, so design income is proportional to *global output*, while manufacturing income is proportional to labor employed. A small design country can therefore earn a share of world income far exceeding its share of world labor — which is roughly the position of US chip designers relative to East Asian manufacturers.

Marking Scheme (30 points)

Part (a) — 6 points

#	Pts	Criterion	Evidence and partial credit
a1	2	1 pt correct output (80) 1 pt valid method (equation, worker allocation whose D=C, PPF graph)	1 point: correct output. 1 point: any valid justification that the output is both feasible and maximal: an equation, worker allocation, PPF, inequality, optimization argument, marginal argument, or any equivalent reasoning.
a2	2	1 pt correct output (75) 1 pt valid method (equation, worker allocation whose D=C, PPF graph)	Same as a1

a3	2	2 points: cost ranking with ordering (engineers supply blueprints first at 0.5 C, then machinists at 3 C) OR the switch/exhaustion or kink logic.	Both opportunity costs correctly identified without the ordering/exhaustion logic: 1 Exact numerical opportunity costs, the word *kink*, and an explicit contrast with diminishing productivity are not required. Attributing the rising cost to diminishing returns, fatigue, or falling productivity = 0.

Traps (a): pooling workers as homogeneous — a1 = a2 = 0, a3 = 0 unless the script separately invokes the two worker types; assigning machinists to design before engineers — method points lost, answer point too if output falls short of 80/75.

Part (b) — 7 points

#	Pts	Criterion	Evidence and partial credit
b1	3	3 points: the contestant gives an output-maximizing world allocation equivalent to: • all 80 engineers design; • 20 machinists design; • 60 machinists manufacture. The 20 designing machinists may come from either country or both.	Evidence: stated or implicit in the balancing equation, or graph 2 points: the correct worker types and direction of allocation are identified, but one quantity contains an arithmetic or accounting slip. 1 point: substantial partial progress toward a feasible optimal allocation.
b2	2	1 point: maximum world output is 180 chips. 1 point: any valid justification that 180 is maximal, including balancing blueprints and components, a complete optimal allocation, a PPF or linear-programming argument, or marginal opportunity-cost reasoning.	Shifting engineers into components (reduces the scarce good): 0. Arithmetic slip: 1 pt

b3	2	1 point: a country-level specialization conclusion consistent with the contestant’s valid allocation. 1 point: an explanation based on what each country produces, or a correct recognition that the location of the 20 designing machinists is not unique.with the contestant’s allocation, for example: *“Northia fully specializes in design while Fabrica produces both goods”; “Northia and Fabrica both produce blueprints, while only Fabrica manufactures”; “both countries produce both goods”*; another allocation using exactly 20 machinists in design and 60 in manufacturing.	A country-specialization statement inconsistent with the contestant’s own allocation loses 1 point only. A bare correct judgment (not fully specialized / consistent country statement) without mechanism earns the country-level point (1), not the worker-type point.

Traps (b): Output 160 with no machinist reallocation: b2 = 0; under b3, award only any valid comparison point. Shifting engineers from design into components at the world optimum: b2 = 0 because this reduces the scarce input. Labor migration: do not award credit for a mechanism relying on migration; separately grade any correct goods-trade allocation. Incorrect country-specialization label: affects b4 only unless it also changes the allocation or output. A correct equivalent allocation without the canonical equation remains eligible for full b1–b3 credit.

Part (c) — 9 points, 3 per claim.

Part (c) is graded on verbal reasoning; calculations are acceptable evidence but not required.

Only the first evaluation of each claim and the first feature in (iii) are graded. Credit may be awarded for a mechanism stated in ordinary language even if the contestant does not use the rubric’s terminology. A conclusion alone earns only the judgment point.

#	Pts	Criterion	Evidence and partial credit
c1	3	1 point: clearly rejects the claim that total employment increased. 2 points: gives a coherent explanation that workers were already employed and the quota merely reallocates existing workers between activities rather than creating additional employment.	“Manufacturing employment rises while design employment falls” earns the first point only unless the answer also establishes prior full employment or unchanged total employment.

		A concise statement that clearly conveys both prior employment and reassignment earns full credit. Falling output, engineers' relative productivity, and the change from 2 blueprints to 1 component are valid supporting observations but are not required.	
c2	3	1 point: concludes that Fabrica's design sector expands. 1 point: explains that Northia supplies or exports fewer blueprints, or that Fabrica can rely less on Northian designs. 1 point: explains that Fabrica must devote more domestic workers or resources to producing blueprints.	A contestant who answers "true, but Fabrica's overall income falls" earns full credit — the welfare caveat is correct. Do not require claims about education, investment, innovation, or long-run skill formation; these are neither necessary nor sufficient unless linked to the model's reduced blueprint trade.
c3	3	1 point: identifies any relevant real-world friction or dynamic feature under which current production affects future productive capability or speed of adjustment. 2 points: explains how that feature makes maintaining domestic production today reduce the cost, difficulty, or delay of responding to a future trade cutoff. Examples include construction lags, specific capital, retraining, learning-by-doing, supplier relationships, capability decay, and irreversible investment, but the list is non-exhaustive and none of these terms is required.	Feature named without the mechanism linking current production to future capability: 2 of 3. Static features (transport costs) or non-arguments (pride, deficits): 0. "The model says otherwise" without the reason: 0

Part (d) — 8 points, 4 per subpart

Part (d) is graded on verbal reasoning; calculations are acceptable evidence but not required.

#	Pts	Criterion	Evidence and partial credit
d1	4	Allowed scores: 0, 2, 4 only 2 points: allocation conclusion. Only a small fixed design team—enough to produce the required blueprint or blueprints—is needed, while substantially more or nearly all remaining labor should be allocated to manufacturing. 2 points: explanation. One blueprint can be used for any number of components, so additional design labor does not increase the number of chips that can be produced once the blueprint exists, whereas additional manufacturing capacity does.	A bare "fewer designers and more manufacturers" earns 2 points. The exact terms *fixed input*, *binding constraint*, and *PPF trade-off collapses* are not required. The exact statement "10 engineers" is not required if the contestant clearly identifies a small fixed design team.
d2	4	1 point: explains that the same blueprint can be reused or copied for many chips at negligible additional design cost. 1 point: explains or implies that the economic value or revenue associated with the design can scale across worldwide production or many manufacturers. 1 point: explains how the design country captures that value — for example through ownership, licensing, royalties, patents, sale of access, bargaining power, or market power. 1 point: connects this scalability to the conclusion that the country's share of world income can substantially exceed its share of world labor.	Real-world illustration (US designers vs East Asian manufacturers, ARM/Nvidia) supports but does not replace the mechanism: example alone, 1 pt total. Technical terms such as *non-rival*, *increasing returns*, and *licensing* are not required. Do not require an explicit statement that manufacturing income is proportional to labor, or that the designer is paid for every individual chip.

Total: 6 + 7 + 9 + 8 = 30.

Open Question 4. Utility maximization, social preferences, and joint contributions

Anil has won £100 in the lottery and is deciding how to split it between himself (x, the amount he keeps) and his friend Bala (y, the amount he gives). Suppose Anil has the utility function $u(x,\ y)\ =\ x\ -\ \delta\cdot(y\ -\ x)^2$, where $\delta \geq 0$ represents Anil's degree of inequality aversion.

(a) (6 points) Suppose you observe Anil giving £37.5 (out of £100) to Bala and keeping the rest for himself. What value of δ is consistent with this behaviour?

(b) (6 points) How does the amount Anil gives to Bala vary with δ? Refer to Anil's utility function when explaining your answer. (You can give a mathematical and/or an intuitive explanation.)

(c) (6 points) Now suppose that Anil becomes purely selfish and plays a one-round game with Bala, Cindy and David (4 people total).

- Each player is given £20 and chooses a contribution to a common pot
- At the end of the round, every player receives k = 0.2 times the total amount in the pot plus the money they chose not to contribute.
- If a player contributes less than some expected amount C, they are punished: an amount P deducted from their money at the end of the round.

Suppose P = £6. What is the largest minimum contribution that Anil would be willing to make (that is, find maximum C that he will agree to contribute)? Assume that he meets the requirement when indifferent.

(d) (6 points) If the share of the total amount in the pot changes to k = 0.5, what is the largest minimum contribution C Anil would be willing to contribute if the punishment is P = £5?

(e) (6 points) Before the round, the four players jointly agree on a single punishment size P and a single required minimum contribution C that will apply to all of them. Once play begins each player is again purely selfish. Anticipating this behaviour, the group picks P and C to maximise the combined end-of-round wealth of all four players. If indifferent between different levels of P, they choose the smallest. Determine the P and C the group should agree on when (i) k = 0.2 and (ii) k = 0.5, and give the resulting combined wealth in each case.

Model Solution

(a) With x + y = 100, u = (100 − y) − δ(2y − 100)². FOC: −1 − 4δ(2y − 100) = 0, so y = 50 − 1/(8δ); equivalently x − y = 1/(4δ). SOC holds (u" = −8δ < 0). Observed y = 37.5 gives 1/(8δ) = 12.5, so δ = 0.01. Check: x − y = 25 = 1/(4·0.01).

(b) From y(δ) = 50 − 1/(8δ): dy/dδ = 1/(8δ²) > 0 — the gift rises with δ. Mechanism: the benefit of keeping money is linear (constant marginal utility 1), the inequality penalty is quadratic with marginal penalty 4δ·(gap); higher δ raises the price of any gap, so the optimal gap $1/(4\delta)$ shrinks. Boundaries: for $\delta \le 1/400$ the corner $y = 0$ applies (weak aversion gives nothing); as $\delta \to \infty$, $y \to 50$ but never exceeds it — at equality the marginal penalty is zero while a kept pound still yields 1, so Anil always keeps more than half.

(c) Comply: $20 - C + 0.2(C + O)$, where O is others' contributions. The best violation is contributing 0, since each contributed pound returns 0.2 at cost 1; violation payoff: $20 + 0.2 \cdot O - 6$. The $0.2 \cdot O$ cancels, so compliance requires $0.8C \le 6$: C = £7.50 (indifferent, hence complies).

(d) Net private cost per pound is now $1 - 0.5 = 0.5$, punishment 5. Indifference: $0.5C = 5$, so C = £10. Relative to (c), the higher k cheapens compliance while the lower P weakens the threat; on net the enforceable minimum rises.

(e) With common compliance, each contributed pound yields 4k to the group at cost 1, so $W = 80 + 4C(4k - 1)$; no punishment is paid in equilibrium, so P matters only via enforceability, which generalizes (c)–(d) to $C \le P/(1 - k)$, plus the cap $C \le 20$.

(i) k = 0.2: $4k - 1 < 0$, contributions destroy wealth. C = 0, and by the tiebreak P = 0; W = £80.

(ii) k = 0.5: $4k - 1 = 1 > 0$, W increases in C, so C = 20 (endowment cap). Enforceability needs $P \ge 0.5 \cdot 20 = 10$; smallest is P = 10. Pot = 80, each gets 40: W = £160.

The two k values sit on opposite sides of the threshold $4k = 1$ separating a wasteful pot from a genuine public good; punishment is worth using only in the second case, and then at the minimal self-enforcing size.

Marking Scheme (30 points)

Global rules

- ECF: a numerical slip is penalized once; later steps earn credit if the method is correct on the contestant's own values. ECF does not rescue conceptual errors (wrong deviation benchmark, wrong payoff structure). A conceptual error repeated in (c) and (d) loses points in both.
- A bare final answer with no setup earns only the answer point of the criterion.
- δ as 0.01 or 1/100, amounts with or without £, either variable used — all accepted.

Part (a) — 6 points

#	Pts	Criterion	Evidence and partial credit

a1	1	Budget constraint imposed (one-variable problem)	Substitution or tangency setup
a2	2	Correct FOC	$-1 - 4\delta(2y - 100) = 0$ or $1 = 4\delta(x - y)$. Sign slip: 1 pt
a3	1	Observed values inserted ($y = 37.5$ / $x = 62.5$ / gap 25)	ECF
a4	2	$\delta = 0.01$	ECF; arithmetic slip: 1 pt. Negative δ as final answer: 0

Traps (a): treating the observation as indifference to (50, 50) rather than an optimum — a2 = 0, downstream only if internally consistent; missing chain-rule factor 2 ($\delta = 0.02$ or 0.005) — a2 = 1 pt, a4 = 1 pt by ECF; maximizing without the constraint — a1 = a2 = 0, nothing recoverable.

Part (b) — 6 points

#	Pts	Criterion	Evidence and partial credit
b1	3	Direction: y increasing in δ, supported by an argument that refers to the utility function — mathematical or intuitive	Full credit for either route: (i) $dy/d\delta > 0$ or comparative statics on $y = 50 - 1/(8\delta)$; or (ii) a verbal argument grounded in the function, e.g. "a higher δ makes any gap between x and y more painful, so Anil narrows the gap by giving more." Direction claim with no supporting argument of either kind: 1 pt
b2	2	Mechanism: the trade-off inside u — constant marginal benefit of keeping money vs an inequality penalty whose weight is δ	Accept in mathematical form (marginal penalty $4\delta\cdot$(gap); optimal gap $1/(4\delta)$ shrinks) or in verbal form ("each kept pound always gives the same benefit, while the penalty for inequality grows with δ, so the tolerated gap shrinks"). One side of the trade-off only (benefit or penalty, in either form): 1 pt
b3	1	Boundary: corner $y = 0$ for small δ (threshold value not required)	Verbal recognition is enough; no value of δ needs to be provided

Traps (b): claiming y → 100 — 0 for b3; sign error making y decreasing — b1 = 0, grade the conclusion the contestant actually draws; answer that makes no reference to the utility function at all (a generic claim such as "more inequality-averse people give more" that could be written without reading u) — max 1 pt in b1, 0 in b2, since the stem explicitly requires referring to the function.

Part (c) — 6 points

#	Pts	Criterion	Evidence and partial credit
c1	2	Both payoff branches correct: comply $20 - C + 0.2(C + O)$; violate $20 + 0.2 \cdot O - P$	Any consistent O (it cancels). One branch right: 1 pt
c2	2	Best violation is 0, justified: marginal return $0.2 < 1$	Assumed without justification: 1 pt
c3	2	$0.8C = 6$, $C = 7.5$	ECF; arithmetic slip: 1 pt; strict inequality ("just under 7.5"): 1 pt

Traps (c): comparing with a marginal deviation instead of the optimal deviation to 0 (yields C = 20 or "no maximum") — c2 = c3 = 0, the central trap; violator loses pot share (violate payoff 20 − 6) — c1 = 1 pt, ECF downstream; P treated per-pound — c1 max 1 pt, ECF; bare 7.5 with no comparison — 1 pt in c3 only.

Part (d) — 6 points

#	Pts	Criterion	Evidence and partial credit
d1	2	Branches adapted: net cost 0.5 per pound, P = 5	Same standards as c1. Using the general rule $C = P/(1 - k)$ derived in (c): full credit for d1–d2
d2	2	$0.5C = 5$	ECF
d3	2	$C = 10$	ECF; arithmetic slip: 1 pt

Traps (d): reusing 0.8 (answer 6.25) — d1 = 0, ECF after; the T-c1 deviation error — same treatment, penalized independently; reusing P = 6 (answer 12) — one point off at d1, ECF.

Part (e) — 6 points

#	Pts	Criterion	Evidence and partial credit

e1	1	Group wealth $W = 80 + 4C(4k - 1)$, no punishment paid in compliant equilibrium	Any equivalent; subtracting 4P: 0, see traps
e2	1	Enforceability $C \le P/(1 - k)$ and cap $C \le 20$ recognized	Implicit but correct use: full credit
e3	2	$k = 0.2$: $C = 0$ because $4k < 1$ (1 pt); $P = 0$ by tiebreak, $W = 80$ (1 pt)	Positive P with $C = 0$: lose second point
e4	2	$k = 0.5$: C pushed to the enforceable/feasible maximum because $4k > 1$ (1 pt); $C = 20$, $P = 10$, $W = 160$ (1 pt, ECF on their constraint)	Anchoring on (d)'s values without maximizing: 0 unless the increasing-W logic appears (then 1 pt)

Traps (e): deducting punishment from equilibrium wealth — e1 = 0, ECF after, usually destroys e4; in (i), maximizing enforceable C instead of wealth — e3 = 0 (the part tests exactly this distinction); in (ii), $C = 20$ with $P = 5$ or no consistent P — second point of e4 = 0; ignoring the endowment cap ("unbounded C") — e2 = e4 = 0.

Total: 6 + 6 + 6 + 6 + 6 = 30.

OQ 5. AI data centres: Earth or orbit? (30 points)

Firm SpaceY must decide whether to locate an AI training facility on Earth or in low-Earth orbit. Consider sunlight in orbit is free and continuous, but every piece of orbital equipment must first be launched from Earth.

Assume the orbital facility is technically feasible and ignore every private cost not stated explicitly: computing hardware, cooling and heat rejection, communications, radiation shielding, maintenance, insurance, and end-of-life disposal. All quantities and prices below are hypothetical.

(a) (6 points) In choosing a location, is an AI training data centre more like a shopping mall or an aluminium smelter facility? Name the location principle at work, and state which location factors matter most and which barely matter for such a centre.

(b) (9 points) Electricity on Earth costs \$0.05/kWh. A centre draws power continuously, operating 8,000 hours per year for 5 years. Ignore discounting unless explicitly asked. On Earth the firm buys grid power. In orbit power is free, but supplying 1 kW of continuous power requires launching 10 kg of equipment, at a current launch price of \$1,500/kg.

(i) For a 1 kW load, compute the lifetime Earth electricity bill and the orbital launch cost.

(ii) At what launch price per kg does orbit just break even? By what factor must today's launch price fall to reach it?

(iii) Ignoring discounting was assumed above. Does that assumption favour Earth or orbit, and why?

(c) (8 points) Today a firm pays to launch, but pays nothing else. Identify two costs an orbital centre imposes on parties other than the firm and its customers, and state who bears each. Then name precisely the economic problem associated with these costs. Choose one of the two suggested costs and propose a corrective instrument and the single obstacle to enforcing it.

(d) (7 points) Some regions have abundant electricity that is effectively stuck (stranded) — remote hydro with no transmission line, associated gas which is burnt off — so its local value is near zero. Electricity is costly to transport and to store; a trained model is a small file (weights). Explain in what sense building an AI training centre at such a site lets a firm "export electricity" that otherwise would be difficult to sell.

Model Solution

(a) It is like an aluminium smelter, not a shopping mall. The principle is least-cost (input-oriented) location: locate where the dominant, hard-to-transport input is cheapest, because the output is cheap to ship relative to its value (see Weber’s location theory). A smelter sits next to cheap power and exports metal; a mall must sit next to customers because its "output" (retail service) cannot travel. An AI training centre's output — a trained model, data through fiber — moves at near-zero cost, while its dominant variable input, electricity, is expensive to move.

The centre consumes a very large quantity of electricity during model training. Electricity can be expensive to transmit and difficult to store. The output of the training process is a set of trained model parameters, commonly called **model weights**. Model weights are digital files and can be transferred through a communications network at a very low cost relative to their economic value.

Factors that matter most: cheap and reliable electricity, cooling conditions, land, adequate connectivity, regulatory/political stability. Factors that barely matter: proximity to customers or population centres, local market size, consumer footfall, large local labour pools (on-site staff is minimal). The question's orbital framing is the extreme case: the "site" has the cheapest possible energy and zero market access — which is irrelevant, exactly because the product is weightless.

(b) (i) Earth: 1 kW × 8,000 h × 5 yr = 40,000 kWh × \$0.05 = \$2,000. Orbit: 10 kg × \$1,500/kg = \$15,000.

(ii) Break-even: 10p = 2,000, so p = \$200/kg. Today's price must fall by a factor of 1,500/200 = 7.5 (or a fall by 86.7%).

(iii) Ignoring discounting makes the orbital option appear relatively more attractive. Therefore, it **favours orbit**.

The reason is that the two costs occur at different times:

- The orbital launch cost is paid at or near the beginning of the project.
- The Earth electricity costs are paid gradually over the five years of operation.

With a positive discount rate, a payment made in the future has a lower present value than an equal payment made today. Discounting would therefore reduce the present value of the Earth electricity bills.

By contrast, the orbital launch cost is paid upfront, so its present value remains 15,000. Discounting would therefore make Earth even cheaper relative to orbit. By ignoring discounting, the calculation overstates the present value of the Earth option and consequently makes orbit look more favourable than it would under a discounted comparison.

(c) Two external costs (others are acceptable — see scheme):

1. Space debris and collision risk. Each launch and each object in orbit raises collision probability for everyone else's satellites and can cascade (Kessler syndrome). Borne by other satellite operators, space agencies, and future users of orbit.
2. Launch emissions. Rocket exhaust deposits black carbon and alumina in the upper atmosphere and pollutes near launch sites. Borne by the general public — including people in other countries and future generations.

The economic problem, precisely named: negative externalities — the firm's private cost of launching excludes costs imposed on third parties, so orbit is overused relative to the social optimum. (For debris specifically, the sharper diagnosis is an open-access commons: orbital slots are unpriced and non-excludable, so the tragedy of the commons applies.)

Corrective instrument for debris: a Pigouvian launch/orbit fee equal to marginal expected debris damage, or a deposit–refund scheme (deposit refunded on verified deorbit). The single key obstacle: enforcement requires international coordination — no jurisdiction controls space (Coase theorem), so a unilateral fee simply shifts launches to non-participating flag states (leakage), and unpriced legacy debris is hard to attribute.

(d) Electricity is nearly non-tradable: transmission lines and storage are expensive, so power with no line to a market has a local value near zero. If the site cannot reach buyers, the local value of the electricity may therefore be close to zero. A trained model is the opposite extreme — a few hundred gigabytes of weights, transportable worldwide at essentially zero cost. Building the training centre at the stranded-energy site converts the immobile input into a weightless mobile output: the energy is exported indirectly, embodied in the valuable digital output, exactly as aluminium is "packaged electricity." The firm arbitrages the gap between the near-zero local price of energy and its value embodied in the model; the site's disadvantage (no transmission) is irrelevant because the product needs only a data link, not a power line. This is the same least-cost logic as (a) taken to its conclusion: when output transport cost is ~0, production migrates to wherever the immobile input is cheapest — flare gas fields, remote hydro, or, in the limit, orbit.

Marking Scheme (30 points)

Global rules

- ECF: a numerical slip is penalized once; later steps earn credit if the method is correct on the contestant's own values. ECF does not rescue conceptual errors (wrong break-even logic, wrong direction in (b)(iii)).
- A bare final number with no setup earns only the answer point.
- Parts (a), (c), (d) are graded on verbal reasoning; naming conventions ("Weber", "Pigouvian") are acceptable evidence but never required — a correct description of the concept earns full credit without the label.

Part (a) — 6 points

#	Pts	Criterion	Evidence and partial credit
a1	2	Correct analogy (smelter) with the explicit reasoning: output cheap to move relative to value, dominant input expensive to move	"Smelter" with no reason: 1 pt; explicit reasoning 1 pt. "Shopping mall": 0 for a1, and a2 unearnable

a2	2	Location principle: input-/resource-/energy-oriented least-cost location (Weber acceptable but not required); the contrast market-oriented vs input-oriented stated or clearly implied	Principle fully described/explained without a name: full credit. Only "locate where electricity is cheap" with no general principle: 1 pt
a3	2	Factors: at least two that matter (electricity price/reliability, cooling, land, connectivity, stability) and at least two that barely matter (customer proximity, local market size, footfall, large labor pool)	1 pt per factor type (if at least two discussed in each). Listing factors without sorting into matters/barely matters: 1 pt total

Traps (a): arguing connectivity means the centre must be near users — 0 within a3 for that item (latency matters for serving, not training; but a contestant who explicitly distinguishes training from inference and notes latency-sensitivity of the latter earns full credit); answering "both, depending" without committing — a1 max 1 pt.

Part (b) — 9 points

#	Pts	Criterion	Evidence and partial credit
b1	2	Earth bill \$2,000	1 pt for 40,000 kWh (8,000 × 5 × 1 kW), 1 pt for × 0.05. Using 8,760 h or omitting the 5 years: lose the first point, ECF after
b2	2	Orbital launch cost \$15,000	10 × 1,500. Bare correct number: 1 pt
b3	2	Break-even price \$200/kg	Equation 10p = (their Earth bill); ECF. Arithmetic slip on a correct equation: 1 pt
b4	1	Factor 7.5	1,500 / (their p); ECF. "Fall by 86.7%" accepted as equivalent
b5	2	Direction and reason in (iii): ignoring discounting favours orbit, because launch cost is upfront while electricity is paid over time; discounting would shrink the PV of Earth's bill only	1 pt direction, 1 pt reason. Direction without reason: 1 pt. "Favours Earth" with any reasoning: 0 (see trap)

Traps (b): comparing annual electricity ($400) with total launch cost — b1 loses a point, b3–b4 by ECF on their figures; break-even solved from the wrong side (launch price at which Earth breaks even, same equation, so usually harmless — grade the equation); (iii) answered "favours Earth because orbit's benefits come later" — 0 for b5: the cash-flow profiles are the opposite, orbit's cost is early and Earth's costs are late; asserting discounting is irrelevant because both are costs — 0 for b5.

Part (c) — 8 points

#	Pts	Criterion	Evidence and partial credit
c1	2	First external cost with its bearer	1 pt cost, 1 pt bearer. Creditable: debris/collision risk (other operators, future space users); launch emissions (public, other countries, future generations); light pollution for astronomy (researchers, observatories); re-entry/falling-object risk (people on the ground); launch-site noise and local pollution (nearby residents). Other examples accepted if correctly and logically argued (weak/questinable examples are not accepted)
c2	2	Second external cost, distinct in channel, with its bearer	Same standards. "Debris" and "collision risk" are one channel; emissions in the atmosphere and local launch-site pollution count as distinct
c3	2	Problem named precisely: negative externality / social cost above private cost; for debris, open-access commons / tragedy of the commons equally creditable	2 pts for negative externality (or tragedy of commons) with the wedge stated (private cost excludes third-party cost → overuse). Just the word "externalities" with no wedge: 1 pt. "Market failure" alone: 1 pt Free-riding can be accepted (if carefully and correctly explained).
c4	2	Corrective instrument for one chosen cost + the single key enforcement obstacle	1 pt instrument matched to the cost (Pigouvian launch fee/debris tax, deposit–refund for deorbit, property rights are properly specified and enforceable, tradable orbital permits (existence of the market for permits), liability rules, emission standards); 1 pt obstacle (no international jurisdiction / leakage to non-participating

			states / attribution of debris / measuring marginal damage). Instrument that doesn't address the chosen cost: 0 for both. Needs clear explanation of the obstacle's impact.

Traps (c): naming costs borne by the firm or its customers (insurance, hardware loss, service interruption to own users) — 0 for that item, the stem excludes them; "pollution" with no bearer — 1 pt of the 2; obstacle stated as generic "hard to enforce" without the mechanism — 0 for the obstacle point.

Part (d) — 7 points

#	Pts	Criterion	Evidence and partial credit
d1	2	Why the electricity is stranded: transport (transmission) and storage are costly, so local value ≈ 0	Either channel for 1 pt, both or an explicit "no buyer can be reached" statement for 2
d2	3	The embodiment mechanism: training converts immobile energy into model weights, a near-zero-transport-cost good; the energy is exported embodied in the output	3 pts for the full conversion logic; 2 pts if the contrast in transport costs is stated but "embodied export" is not articulated; 1 pt for "build the centre where power is cheap" with no export sense
d3	2	The economic payoff: the firm arbitrages the price gap between stranded local energy and its embodied value; the site's lack of transmission becomes irrelevant because only a data link is needed	1 pt for the arbitrage/value-unlocking point, 1 pt for why the usual disadvantage disappears. An apt analogy (aluminium as packaged electricity, flare-gas Bitcoin mining) is acceptable evidence for either point

Traps (d): claiming the firm literally sells electricity or transmits it as data — 0 for d2 (the point is embodiment, not transmission of energy); answering purely about cheap costs with no export interpretation — capped at 3 total across the part; asserting weights are valuable because of "data gravity" or storage of energy in the model — 0 for the affected criterion.

Total: 6 + 9 + 8 + 7 = 30.