Risk Management and Derivatives 2026
Shanghai Jiao Tong University
Summer School 2026
Final Exam: Topics 7–11
With Solutions

Reading time: 10 minutes. **Writing time:** 120 minutes.

Total marks: 100. **Calculators are permitted.**

Instructions: Answer **all** questions. In Section A, circle the letter of the single best answer. In Section B show your working; a correct number with no method earns little credit. Section C is a short essay. Interest rates are **continuously compounded** unless stated otherwise.

Structure:

- **Section A** – Multiple choice: 20 questions, 2.5 marks each (**50 marks**).
- **Section B** – One analytical problem in five parts (**30 marks**).
- **Section C** – One short essay (**20 marks**).

Useful formulas. Swap to the fixed payer $V = B_{\text{float}} - B_{\text{fixed}}$, with B_{float} = notional at a reset date; par swap rate $c^* = \dfrac{1 - DF(n)}{\sum_{t=1}^{n} DF(t)}$; tranche loss = $\min\big(\max(L - \text{attachment}, 0),\ \text{tranche size}\big)$; intrinsic value: call $\max(S-K,0)$, put $\max(K-S,0)$; time value = premium − intrinsic; bounds: $\max(S - Ke^{-rT}, 0) \le c \le S$, $\max(Ke^{-rT} - S, 0) \le p \le Ke^{-rT}$; put-call parity $c + Ke^{-rT} = p + S$; early exercise before an ex-dividend date can pay only if $D > K\big(1 - e^{-r\tau}\big)$; naked-call writer margin per share $= \max\big(0.20S + \text{premium} - \text{OTM amount},\ 0.10S + \text{premium}\big)$; bull call spread: max profit $= (K_2 - K_1) -$ net cost, breakeven $= K_1 +$ net cost; straddle breakevens $= K \pm$ total premium.

Section A: Multiple Choice (20 questions × 2.5 marks = 50 marks)

1. A firm enters a three-year interest-rate swap today as the fixed-rate payer. The fixed rate was set at the current market *par* swap rate. At the moment of signing, the value of the swap to the firm is:

(A) positive, because paying fixed is a bet that rates will rise.

(B) about zero, because the two legs have equal present value.

(C) negative, because a fixed payment is always a net cost.

(D) unknown until the first floating fixing is announced.

Solution. (B). The par swap rate is defined as the fixed rate that makes the present value of the fixed leg equal the present value of the floating leg, so a swap struck at par is born worth zero to both sides, like a forward struck at the forward price. It drifts into positive or negative value only as rates move afterwards. Options (A) and (C) confuse a position's direction with its value; (D) is wrong because valuation never needs a forecast of the floating path.

2. Firm A can borrow fixed at 3.8% or floating at SOFR +0.4%. Firm B can borrow fixed at 5.4% or floating at SOFR +1.0%. If each firm borrows where it is relatively best and they swap, the total gain available to be shared each year is:

(A) 0.6% per year.

(B) 1.6% per year.

(C) 2.2% per year.

(D) 1.0% per year.

Solution. (D). Total gain = (fixed-rate gap) − (floating-rate gap) = $(5.4 - 3.8) - (1.0 - 0.4) = 1.6 - 0.6 = 1.0\%$ per year. Firm A has comparative advantage in fixed, B in floating; each borrows there and they swap. Option (A) is only the floating gap, (B) only the fixed gap, and (C) adds the two gaps instead of differencing them.

3. A bank pays 5% fixed annually and receives SOFR on a notional of \$100 million. Two payment dates remain and a floating reset has just occurred. Discount factors are $DF(1) = 0.96$ and $DF(2) = 0.90$. Valuing the swap as two bonds, its value to the bank is:

(A) +\$0.7 million.

(B) −\$0.7 million.

(C) \$0, since swaps are always worth zero.

(D) +\$9.3 million.

Solution. (A). At a reset date the floating leg is worth par: $B_{\text{float}} = 100$. The fixed leg is a 5% bond: $B_{\text{fixed}} = 5 \times (0.96 + 0.90) + 100 \times 0.90 = 9.3 + 90 = 99.3$. To the fixed payer $V = B_{\text{float}} - B_{\text{fixed}} = 100 - 99.3 = +\$0.7$ million: the bank is paying 5% when the par rate implied by these discount factors is about 5.38%, so its side is the cheap one. Option (B) values the receiver's side; (C) is true only at inception at par; (D) is the present value of the coupons alone.

4. In a fixed-for-fixed *currency* swap, unlike a single-currency interest-rate swap:

(A) principal is never exchanged, exactly as in the single-currency case, so only the coupons differ.

(B) its value is driven by the domestic yield curve alone, since the foreign leg's payments are fixed.

(C) principals are exchanged at the start and at maturity, so the exchange rate and both curves drive its value.

(D) counterparty credit risk is eliminated, because the two currency legs offset each other exactly.

Solution. (C). In a currency swap the principals in the two currencies are exchanged at the start and re-exchanged at maturity, so the swap's value moves with the exchange rate and with both interest-rate curves. The familiar "the notional is only a scaling figure" intuition is specific to single-currency swaps, where identical currency amounts would simply cancel. Because the final principal exchange is large, credit exposure is typically bigger, not smaller (D).

5. A \$100 million mortgage pool is tranched into \$80 million senior, \$15 million mezzanine, and \$5 million equity, with losses applied bottom up. Defaults produce \$14 million of losses on the pool. The mezzanine tranche loses:

(A) 0% of its principal.

(B) 14% of its principal.

(C) 60% of its principal.

(D) 100% of its principal.

Solution. (C). The equity tranche absorbs the first \$5 million and is wiped out. The next $14 - 5 =$ \$9 million falls on the mezzanine tranche: $9/15 = 60\%$ of its principal. The senior tranche, attaching at \$20 million, is untouched. Option (B) applies the pool's loss rate to the tranche, forgetting the waterfall; (A) and (D) put the mezzanine at the wrong place in the queue.

6. Two loan pools have the *same* average default rate. In Pool X the loans default nearly independently; in Pool Y they are highly correlated and tend to default together. The senior tranche is:

(A) equally safe in both, because the average default rate is equal.

(B) riskier in Pool Y, because clustered defaults can reach the senior layer.

(C) safer in Pool Y, because correlation spreads defaults over time.

(D) risk-free in both, because senior tranches carry AAA ratings.

Solution. (B). Correlation does not change the pool's expected loss; it changes the *shape* of the loss distribution. With independent loans, extreme outcomes are rare and losses almost never climb past the junior layers. With correlated loans, defaults cluster: bad states are bad everywhere at once, and pool losses can burn through equity and mezzanine into the senior tranche. The senior investor is short exactly this tail.

7. A tiny pool holds two loans; each defaults with probability 10%. A "senior" piece loses money only if *both* loans default. Moving from independent defaults to perfectly correlated defaults, the probability that the senior piece loses rises:

(A) from 1% to 10%.

(B) from 10% to 20%.

(C) from 5% to 10%.

(D) not at all; it stays at 10%.

Solution. (A). Independent: $P(\text{both}) = 0.10 \times 0.10 = 1\%$. Perfectly correlated: the loans default together or not at all, so $P(\text{both}) = 10\%$. The senior piece is ten times riskier, while the pool's total expected loss is unchanged in both cases. This two-loan example is the whole credit crisis in miniature: correlation moves risk *up* the capital structure. Option (B) adds probabilities; (C) averages them; (D) assumes correlation is irrelevant to a joint event.

8. A structurer pools the mezzanine tranches of one hundred mortgage securitisations from one hundred different cities, re-tranches the pool, and wins a AAA rating for the new senior layer by assuming the inputs default nearly independently. House prices then fall across the whole country at once. The senior layer suffers because:

(A) under stress the waterfall reverses and pays the new pool's equity layer ahead of its senior layer.

(B) defaults in one city trigger cross-default clauses written into the other ninety-nine securitisations.

(C) the new senior layer was sized larger than mortgage collateral can support in any national downturn.

(D) the hundred inputs share one national housing factor, so they default together and never diversify.

Solution. (D). Every input mezzanine tranche is a bet on the same thing: national house prices. Geographic labels give no diversification against a common factor, so when the factor moves, all one hundred inputs are hit simultaneously and the re-securitised "AAA" layer behaves like a single concentrated bet. The waterfall never reverses (A), cross-default clauses across separate deals do not exist (B), and the problem was correlation, not sizing (C).

9. A stock trades at $S = \$48$. A call with strike $K = \$45$ is quoted at \$4.50. The call's *time value* is:

(A) \$3.00.

(B) \$1.50.

(C) \$4.50.

(D) \$0, because an in-the-money option has no time value.

Solution. (B). Intrinsic value $= \max(S-K, 0) = 48-45 = \$3.00$; time value = premium − intrinsic $= 4.50-3.00 = \$1.50$. Every live option's premium splits into these two parts, and a live call's time value is positive whether it is in or out of the money, which rules out (D). Option (A) is the intrinsic part; (C) is the whole premium.

10. A trader *writes* one put contract (100 shares) with strike $K = \$50$ for a premium of \$3.80 per share and holds to expiry. The stock finishes at $S_T = \$44$. The writer's total profit or loss is:

(A) a loss of \$220.

(B) a gain of \$380.

(C) a loss of \$600.

(D) a gain of \$220.

Solution. (A). The put finishes $50 - 44 = \$6.00$ in the money, so the writer pays out \$6.00 per share against \$3.80 collected: $3.80 - 6.00 = -\$2.20$ per share, or $-\$220$ on the contract. Option (B) keeps the premium and forgets the assignment; (C) is the payoff with the premium forgotten; (D) has the right size and the wrong sign. Writer profit is always the mirror image of the holder's.

11. An investor holds one call contract (100 shares), strike \$80, on a stock trading at \$88. The company announces a 2-for-1 stock split, after which the shares trade near \$44. Under the standard exchange adjustment, the investor now holds:

(A) a worthless call, since the post-split price of \$44 sits below the \$80 strike.

(B) no position at all, since outstanding contracts are cancelled when a stock splits.

(C) a call with strike \$40 on the same 100 shares, worth half what it was before.

(D) a call with strike \$40 on 200 shares, with the position's value fully preserved.

Solution. (D). On an n-for-m split the strike is scaled by m/n and the contract size by n/m: here the strike halves to \$40 and the contract covers 200 shares. Intrinsic value before, $100 \times (88 - 80) = \$800$; after, $200 \times (44 - 40) = \$800$. The adjustment exists precisely so that a corporate action changes neither side's wealth. Option (C) adjusts the strike but not the size, which would halve the position's value.

12. You wrote an American call, strike \$35, on a stock now at \$40. The stock pays a \$1.50 dividend and goes ex-dividend tomorrow; your call has \$0.40 of time value left. Which statement is correct?

(A) You decide when you are assigned, so you can safely wait until the expiry date.

(B) You cannot be assigned before expiry, since exercise settles only at maturity.

(C) Assignment tonight is likely, since the dividend to be captured exceeds the remaining time value.

(D) The large dividend protects the writer, making early exercise unattractive and assignment unlikely.

Solution. (C). Early exercise of an American call can pay just before an ex-dividend date when the dividend to be captured exceeds the time value surrendered: here $\$1.50 > \$0.40$. Rational holders exercise tonight, and the clearing house assigns writers at random; the writer has no say in the timing (A). American options can be exercised on any day (B), and a bigger dividend makes early exercise more attractive, not less (D).

13. A European call has strike $K = \$48$ on a non-dividend stock at $S = \$50$, with $r = 6\%$ and $T = 0.5$ years, so $Ke^{-rT} = \$46.58$. A dealer quotes the call at \$1.80. With no option model at all, you can conclude:

(A) the quote is below the no-arbitrage floor $S - Ke^{-rT} = \$3.42$, so a riskless profit is available.

(B) the quote is plausible, since the call is only \$2.00 in the money.

(C) the quote is too high, since a call can never be worth more than its intrinsic value.

(D) nothing can be concluded, since an option's value depends on the stock's volatility.

Solution. (A). The lower bound $c \geq \max(S - Ke^{-rT}, 0) = 50 - 46.58 = \$3.42$ needs no model: buying the call and shorting the stock against it locks a profit whenever the call trades below this floor. At \$1.80 the call is at least \$1.62 too cheap. Note the floor exceeds intrinsic value (\$2.00) because the strike is paid later; (C) inverts the bound, and (D) forgets that bounds, unlike prices, are volatility-free.

14. For a stock with $S = \$50$, $K = \$48$ and $Ke^{-rT} = \$46.58$, the market quotes the European call at $c = \$5.90$ and the European put at $p = \$2.00$. Checking put-call parity:

(A) the put side is dear by about \$0.48.

(B) the call side is dear by about \$0.48.

(C) parity holds to within rounding, so no trade is available.

(D) the call side is dear by \$3.90, the difference $c - p$.

Solution. (B). Parity requires $c + Ke^{-rT} = p + S$. Left side: $5.90 + 46.58 = \$52.48$. Right side: $2.00 + 50.00 = \$52.00$. The call side is dear by \$0.48, so the conversion (sell the call, buy the put, buy the stock, borrow Ke^{-rT}) banks \$0.48 per share today and owes nothing at expiry. Option (D) compares raw premiums, but $c - p$ should equal $S - Ke^{-rT} = \$3.42$, not zero.

15. An American call has strike $K = \$50$ and $r = 6\%$. The stock is about to go ex-dividend with $D = \$2.00$; after the ex-dividend date, $\tau = 0.25$ years remain to expiry. Using the early-exercise test $D > K(1 - e^{-r\tau})$:

(A) the critical dividend is $\approx$ \$0.74, but an American call should never be exercised early whatever the dividend.

(B) the critical dividend is $\approx$ \$1.48, and \$2.00 exceeds it, so early exercise may be optimal.

(C) the critical dividend is $\approx$ \$2.91, and \$2.00 is below it, so early exercise is not optimal.

(D) the critical dividend is $\approx$ \$0.74, and \$2.00 exceeds it, so exercising just before the ex-dividend date may be optimal.

Solution. (D). $K(1-e^{-r\tau}) = 50 \times (1-e^{-0.06\times0.25}) = 50 \times 0.01489 = \$0.74$. The dividend of \$2.00 comfortably exceeds the interest saved by paying the strike late, so this is exactly the one exception to the "never exercise an American call early" rule, which holds only for non-dividend payers (A). Options (B) and (C) compute the critical value over the wrong horizon (six months and one year instead of the three months that remain after the date).

16. Higher volatility raises the value of *both* calls and puts because:

(A) higher volatility raises the expected return that the stock must pay its holders.

(B) it makes finishing in the money more likely, and that probability is what the premium measures.

(C) the holder's downside is capped at the premium while larger favourable moves pay in full.

(D) it raises the present value of the strike, which appears on both sides of parity.

Solution. (C). An option holder keeps the whole benefit of large favourable moves but loses at most the premium on unfavourable ones. Widening the distribution of outcomes therefore adds more value on the good tail than it costs on the bad tail, for calls and puts alike. Volatility says nothing about the stock's expected return (A). Option (B) is a half-truth: value depends on how *far* in the money the option can finish, not merely on whether it does, and for an option already in the money more volatility can even lower that probability. The discounted strike depends on r and T, not on volatility (D).

17. A trader buys a call with strike \$50 for \$4.00 and writes a call with strike \$55 for \$2.00, same stock and expiry. The maximum profit per share at expiry is:

(A) \$5.00.

(B) \$3.00.

(C) \$2.00.

(D) $7.00.

Solution. (B). This is a bull call spread. Net cost $= 4.00 - 2.00 =$ $2.00. Above $55 both calls are exercised and the position pays the strike gap, $55 - 50 =$ $5.00; profit $= 5.00 - 2.00 =$ $3.00 per share. Option (A) is the payoff with the cost forgotten; (C) is the net cost itself (which is the maximum *loss*); (D) adds cost to width instead of subtracting.

18. A trader buys a straddle at strike $50: the call costs $3.00 and the put $2.50. The expiry breakeven prices are:

(A) $44.50 and $55.50.

(B) $47.00 and $53.00.

(C) $47.50 and $52.50.

(D) $44.50 only, since a straddle profits only from falls.

Solution. (A). Total premium $= 3.00 + 2.50 =$ $5.50; the straddle profits only once the stock has moved past the strike by more than the total premium, in either direction: 50 ± 5.50, so $44.50 and $55.50. Options (B) and (C) each use only one leg's premium, and (D) forgets the call side. The straddle is a bet on *distance*, not direction, and its price is the size of the move the market already expects.

19. An investor owns a share bought at $50 and writes a one-month call with strike $55 for a $2.00 premium (a covered call). The stock finishes the month at $62. Compared with simply holding the share, the covered call ended:

(A) $2.00 per share better, since the premium collected is pure additional income.

(B) exactly level, since the premium always compensates for the surrendered upside.

(C) $7.00 per share worse, since the writer loses the entire move above the strike.

(D) $5.00 per share worse, since $7.00 of upside was sold for a $2.00 premium.

Solution. (D). Covered call: the share is called away at $55, so the stock leg earns $55 - 50 =$ $5.00, plus the $2.00 premium: +$7.00. Plain holding earns $62 - 50 =$ $12.00. The covered call did $12 - 7 =$ $5.00 worse: the writer sold the upside tail above $55 for $2.00, and this month that tail was worth $7.00. Option (C) counts the lost upside but forgets the premium received; (A) and (B) are the seller's wishful accounting.

20. An investor owns a share at $50 and sets up a zero-cost collar: buying a put with strike $46 for $1.20 and writing a call with strike $56 for $1.20. Ignoring interest, the position's profit per share at expiry is:

(A) floored at $-\$4.00$ with unlimited upside above \$56.

(B) floored at $-\$5.20$ and capped at $+\$4.80$.

(C) floored at $-\$4.00$ and capped at $+\$6.00$.

(D) floored at \$0, since a zero-cost collar cannot lose money.

Solution. (C). The premiums cancel, so net premium is zero. Below \$46 the put sells the share at \$46: profit $46 - 50 = -\$4.00$, the floor. Above \$56 the share is called away at \$56: profit $56 - 50 = +\$6.00$, the cap. Between the strikes profit is simply $S_T - 50$. "Zero-cost" means zero *premium*, not zero risk (D): the collar is paid for by surrendering the upside above \$56, and it caps rather than removes the downside (A). Option (B) double-counts the premiums that in fact cancel.

Section B: Analytical Problem (30 marks)

You are on the equity-derivatives desk of a London bank. A client, Harcourt Capital, holds **100,000 shares** of the non-dividend US stock ATLAS, trading at $S = \$80$. Six-month ($T = 0.5$) European options on ATLAS are quoted per share below; each listed contract covers 100 shares. The riskless rate is $r = 4\%$ (continuously compounded); you may use $Ke^{-rT} = \$78.42$ for the \$80 strike.

strike K	call	put
72	–	2.10
80	7.90	5.80
84	–	7.80
90	2.90	–

(a) **[6 marks]** Split the \$84 put's premium and the \$80 call's premium into intrinsic value and time value, and state each option's moneyness. Where on a board is time value largest, and do these two quotes agree with that rule?

(b) **[6 marks]** Audit the \$80 strike against put-call parity. Compute the parity-fair call price from the put quote, state the size of the violation, and describe the four-leg conversion trade that exploits it, giving the cash banked today per share and showing that the position owes nothing at expiry whether S_T finishes above or below \$80.

(c) **[6 marks]** Instead of paying for protection, the client first proposes *writing* the \$90 call naked, one contract, at \$2.90. The broker margin per share is the greater of (i) $0.20S + \text{premium} - \text{OTM amount}$ and (ii) $0.10S + \text{premium}$. Compute both methods, the margin per contract, and the ratio of margin to premium received. What does this ratio tell the client about the trade?

(d) **[6 marks]** The client instead sets up a zero-premium-style collar on the whole holding: buy the \$72 put at \$2.10 and write the \$90 call at \$2.90 (per share, scaled to 100,000 shares). Compute the net premium, the floor profit and the cap profit per share, the breakeven price, and the worst-case and best-case outcomes in dollars on the full position.

(e) **[6 marks]** The \$8 million position is financed with a floating-rate loan, and the client asks you to swap it to fixed for two years (annual payments, annual compounding). Discount factors are $DF(1) = 0.97$ and $DF(2) = 0.93$. Compute the par swap rate, state the value of the swap at inception if it is struck at that rate, and explain in one sentence why the floating leg can be valued without forecasting SOFR.

Solution. (a) \$84 put: intrinsic $= \max(K - S, 0) = 84 - 80 = \$4.00$; time value $= 7.80 - 4.00 = \$3.80$; in the money. \$80 call: intrinsic $= \max(S - K, 0) = \$0$; time value $= 7.90 - 0 = \$7.90$; at the money. Time value is largest **at the money**, where the exercise decision is most uncertain, and the two quotes agree: the at-the-money call carries \$7.90 of time value against the in-the-money put's \$3.80.

(b) Parity: $c+Ke^{-rT} = p+S$, so the fair call $= p+S-Ke^{-rT} = 5.80+80.00-78.42 = \$7.38$. The quoted call at \$7.90 is **dear by \$0.52**. The **conversion**: sell the call (+7.90), buy the put (−5.80), buy the stock (−80.00), borrow Ke^{-rT} (+78.42); net cash today $= +\$0.52$ per share. At expiry the loan repays $78.42\,e^{0.02} = \$80.00$. If $S_T \geq 80$: the call is assigned, the stock is delivered at \$80, the put dies; the \$80 repays the loan. If $S_T < 80$: the put sells the stock at \$80, the call dies; again \$80 repays the loan. Terminal net is zero either way, so the \$0.52 (\$52 per contract) is riskless.

(c) The \$90 call is out of the money by $90 - 80 = \$10$. Method (i): $0.20 \times 80 + 2.90 - 10 = 16.00 + 2.90 - 10.00 = \$8.90$. Method (ii): $0.10 \times 80 + 2.90 = \$10.90$. The greater is \$10.90, so the margin is $10.90 \times 100 = \mathbf{\$1{,}090}$ per contract against \$290 of premium received: a ratio of about **3.8 to 1**. The client must tie up nearly four times the income just to open the position, and the requirement is recomputed as the market moves; the margin is the clearing system's price for an obligation whose loss has no ceiling.

(d) Net premium $= 2.90 - 2.10 = +\$0.80$ received per share. Floor ($S_T \leq 72$): the put sells at \$72, so profit $= 72 - 80 + 0.80 = -\$7.20$ per share. Cap ($S_T \geq 90$): called away at \$90, so profit $= 90 - 80 + 0.80 = +\$10.80$ per share. Between the strikes profit $= S_T - 80 + 0.80$, so breakeven at $S_T = \$79.20$. On 100,000 shares: worst case −\$720,000, best case +\$1,080,000. The floor is bought by selling the tail above \$90; here the sale even pays a small net premium.

(e) Par swap rate:

$$c^* = \frac{1 - DF(2)}{DF(1) + DF(2)} = \frac{1 - 0.93}{0.97 + 0.93} = \frac{0.07}{1.90} = \mathbf{3.68}\% \text{ per year.}$$

Struck at c^*, the swap is worth **zero at inception**: the fixed leg's present value equals the floating leg's. The floating leg needs no SOFR forecast because a floating-rate note that pays the market rate is always worth **par at a reset date**: whatever rates do, the next coupon and the discounting move together, so $B_{\text{float}} =$ notional today.

Section C: Essay (20 marks)

Read the scenario and write a short essay of **about half a page** (roughly 5 to 8 sentences). Marks are for clear economic reasoning, not length.

In the mid-1980s, large pension funds adopted "portfolio insurance". Rather than buying put options on their equity holdings, which would have cost real premium, they ran a computer-driven strategy that promised the same result synthetically: as the market fell, the program sold stock-index futures to cut exposure; as it rose, the program bought them back. By October 1987, tens of billions of dollars were managed this way, by many institutions running essentially the same model. On Monday 19 October 1987 the US market gapped down at the open and fell about 20 per cent in one day. The programs all demanded huge futures sales at once, into a market with almost no buyers; the selling could not be executed anywhere near the model's assumed prices, and the promised floor failed exactly when it was needed.

Question [20 marks]. The funds believed they held the equivalent of a protective put. Using the tools of this course, explain how a synthetic put built by dynamic trading differs from a bought put, why the strategy failed precisely in the state of the world it was designed for, and what this episode teaches about when a hedger should pay the premium and buy the contract.

Solution. Model answer (replication is not a contract). A bought protective put is a **contract**: the premium is paid up front, and the writer is then obliged to deliver the floor whatever happens, however fast the market moves. A synthetic put is a **replication recipe**: it manufactures the put's payoff by continuously adjusting a futures position, selling as prices fall and buying as they rise, and it works only under its assumptions: prices move continuously rather than gapping, the futures market stays liquid enough to absorb the trades, and one fund's own selling does not move the price. In October 1987 all three assumptions failed together, and they failed **because** many institutions ran the same model: a falling market ordered every program to sell simultaneously, the selling pushed prices down further and triggered more model selling, and when the market gapped there was no sequence of trades to execute at the model's assumed prices. So the floor evaporated in exactly the state it was built for; the "saved" premium had not eliminated the cost of the tail but converted it into execution risk that came due all at once. The lesson: the premium on a real option is the price of transferring the tail to a writer in advance, and a hedger whose floor genuinely matters in a crash should pay it, because a replication strategy is, in effect, self-insurance that is cheapest in calm markets and fails in crowded, gapping ones.

Marking guide (20 marks).

- **8 marks** – the core distinction: bought put = contract, tail risk transferred to the writer for a known premium; synthetic put = dynamic replication (sell futures as prices fall, buy as they rise) that depends on continuous prices, liquid markets, and trades that do not move the price.

- **6 marks** – why it failed when needed: crowding and feedback, identical models selling into a falling market, prices gapping past the rebalancing points, so the hedge could not be executed at model prices and the floor failed in the exact state it was designed for.
- **4 marks** – the lesson: the premium is the price of the tail; avoiding it re-imports the tail as execution and liquidity risk; strategies must be stressed against gaps and against everyone else running the same trade; buy the contract when the floor truly matters.
- **2 marks** – clear, well-organised writing within the length limit.

Answers claiming the models were simply "wrong about direction," or that options themselves caused the crash, earn little: the failure was in the replication assumptions and the crowding, not in the view.