

FINM 36700 - Final Exam

Portfolio and Risk Management

Winter 2026

- This exam contains 100 multiple choice questions.
- Questions marked with † are [**Select All That Apply**] — choose all correct answers (there may be 0, 1, or multiple correct answers).
- No calculators or computers are allowed.
- All calculations are designed for mental math.

Questions

1. A student claims that diversification only works when assets are negatively correlated. This claim is:
 - Correct — portfolio variance can only decrease below the average when $\rho < 0$
 - Incorrect — diversification reduces portfolio volatility for any $\rho < 1$
 - Correct — the equally-weighted portfolio formula shows variance increases with positive ρ
 - Incorrect — diversification works only when $\rho = 0$
2. An equally-weighted portfolio of n stocks, each with volatility $\sigma = 20\%$ and average pairwise correlation $\rho = 0.04$, has portfolio volatility that converges to approximately:
 - 0% — all risk is diversified away
 - About 12% — only systematic risk remains
 - 20% — diversification has no effect
 - About 4% — average covariance remains even in a very large portfolio
3. A stock has daily volatility of 0.5%. Its approximate annualized volatility (assuming 250 trading days) is closest to:
 - 500%
 - 32%
 - 2%
 - 8%
4. A student ranks assets by individual Sharpe ratio and predicts the highest-Sharpe asset will receive the largest weight in the tangency portfolio. This prediction is:
 - Always correct — the tangency portfolio maximizes Sharpe by loading ρ on high-Sharpe assets
 - Incorrect — tangency weights depend on the full covariance structure, not just individual Sharpe ratios
 - Correct only if all pairwise correlations are equal
 - Correct only if the risk-free rate is zero
5. In the S&P 500 constituent analysis, increasing the asset universe from 5 to 442 stocks caused the condition number of the covariance matrix to increase from ~ 4 to $\sim 83,000$. The practical consequence is:
 - The efficient frontier shifts inward, reducing the maximum achievable Sharpe ratio
 - The optimizer can no longer find a solution because the problem becomes non-convex
 - Small estimation errors in covariances produce large, unstable swings in optimal weights
 - The covariance matrix becomes negative definite

6. In the HMC case, Private Equity appears to offer a large return premium over Domestic Equity with similar volatility and low correlation — a 'too good to be true' profile. The professor explained that the most likely cause of this apparent dominance is:
- Private Equity managers generate genuine alpha through superior selection
 - Survivorship bias inflates average Private Equity returns
 - Stale pricing of illiquid PE assets artificially smooths reported volatility and correlation
 - The sample period was unusually favorable for PE
7. An investor with access to a risk-free asset and the tangency portfolio wants to achieve a lower expected return than the tangency portfolio itself while remaining efficient. Under the Capital Market Line framework, the investor should:
- Search for assets not on the efficient frontier
 - Lever up by borrowing at the risk-free rate and investing more than 100% in the tangency portfolio
 - Hold the global minimum variance portfolio instead
 - Lend at the risk-free rate and invest less than 100% in the tangency portfolio
8. A bank calculates daily VaR and needs to report 100-day VaR to regulators. The standard convention is to scale daily VaR by:
- 10 (because $\sqrt{100} = 10$ under the square-root-of-time rule)
 - $\sqrt{10} \approx 3.2$ (square-root-of-time rule)
 - $10^2 = 100$ (variance scaling)
 - $\log(10) \approx 2.3$
9. A bank holds two trading desks. Each desk's standalone 5% VaR is \$10M. The combined desk's 5% VaR is \$22M. This violates which coherence property?
- Positive homogeneity
 - Translation invariance
 - Subadditivity
 - Monotonicity
10. You backtest a 5% VaR model and observe that actual losses exceed the VaR threshold on 1% of trading days. This model is:
- Conservative — it overestimates risk
 - Anticonservative — it underestimates risk, which is dangerous for capital adequacy
 - Well-calibrated — fewer breaches just mean the realized sample happened to be calmer than usual
 - Uninformative — hit rates cannot evaluate VaR models
11. A portfolio has daily mean return of 0% and daily volatility of 0.1%. Using the normal distribution ($z_{0.05} = -1.645$), the 5% daily VaR is approximately:

- -0.16%
 - -1.6%
 - -5%
 - -16%
12. † [Select All That Apply] Which of the following are weaknesses of the expanding-window volatility estimator for conditional VaR?
- Each new observation receives vanishing weight ($1/t$), so it responds slowly to regime changes
 - It cannot be computed with fewer than 60 observations
 - During a crisis, the estimate lags behind the true spike in volatility
 - It assigns negative weights to recent observations
13. The EWMA volatility model assigns weights λ^i to past squared returns (with λ close to 1), so that recent observations receive more weight than distant ones. Compared to a 60-day rolling window, EWMA's primary advantage is:
- It assigns equal weight to all observations in the window
 - It produces a smoother volatility estimate because large past shocks decay geometrically rather than dropping off abruptly
 - It always produces lower volatility estimates
 - It requires fewer parameters to estimate via MLE
14. A stock has very high standalone volatility but zero correlation with the market. Under the CAPM, its expected excess return is:
- High — investors demand compensation for bearing high volatility
 - Negative — uncorrelated assets are penalized
 - Equal to the market premium — all stocks earn the same expected return
 - Zero — only systematic risk (market covariance) is compensated
15. A time-series regression of stock excess returns on market excess returns yields $\alpha = 0.03\%$ per month (t-stat = 0.2). Under the CAPM framework, this alpha indicates:
- The CAPM is not rejected for this stock because alpha is economically tiny and statistically indistinguishable from zero
 - The stock is mispriced — it earns a return not explained by its market beta
 - The regression R-squared must be very high
 - The stock's beta is greater than 1
16. Under the CAPM, the Sharpe ratio of any individual asset i is related to the market Sharpe ratio by $SR_i = \rho_{i,m} \times SR_m$. This implies that no individual asset can have a Sharpe ratio exceeding the market's because:
- Individual assets always have higher volatility than the market
 - The CAPM assumes all assets have the same expected return

- Diversification always increases expected returns
 - $-\rho_{i,m} \leq 1$, so the scaling factor is at most 1
17. A stock has volatility $\sigma_i = 30\%$, the market has volatility $\sigma_m = 15\%$, and their correlation is $\rho = 0.6$. The stock's market beta is:
- 0.3
 - 0.6
 - 1.2
 - 3.0
18. In the DFA case, running the two-stage CAPM test on 25 size-value sorted portfolios produces a low cross-sectional R-squared and a positive intercept ($\gamma \neq 0$). This result most directly motivates:
- Rejecting all factor models as useless
 - Using a higher-frequency dataset to improve beta estimation
 - Replacing beta with volatility as the risk measure
 - Adding size (SMB) and value (HML) factors to better explain cross-sectional return variation
19. The risk-free rate is 2% and the market excess return is 5%. A stock with $\beta = 0.2$ has a CAPM-predicted expected return of:
- 9%
 - 11%
 - 6%
 - 3%
20. In the 49-industry CAPM regressions, the Software sector has a low R-squared but statistically significant alpha, while the Retail sector has a high R-squared but alpha near zero. The professor stated: 'Software had the worst R-squared... Retail looks great. Great R-squared. What does the CAPM say about that? Nothing.' Which statement best reflects this framing?
- The CAPM works better for Software because its beta is higher
 - The CAPM works better for Retail because its R-squared is higher, showing a tighter fit
 - The CAPM works better for Retail because its alpha is near zero — pricing accuracy, not R-squared, determines whether a factor model 'works'
 - Neither regression tells us whether the CAPM works
21. The lecture says a momentum strategy built from the top decile has a much higher mean return than one built from the top half. What is the intended lesson?
- Any broader portfolio must be mismeasured because means should be identical
 - Momentum only works if the long side is held without any short side
 - The top half construction is invalid because it uses too many stocks
 - Extreme sorts can amplify a small underlying autocorrelation signal

22. The professor used a coin-flip analogy to explain momentum. A coin that comes up heads 50.5% of the time represents:
- The small but positive alpha of a momentum strategy
 - The tiny autocorrelation in individual stock returns, where diversification across many bets makes the edge reliable
 - The probability that a momentum strategy beats the market in any given year
 - The fraction of stocks that are winners minus losers
23. The professor explained that naive monthly-rebalanced momentum strategies fail in practice because:
- The signal decays — big winners in January are rarely big winners in February, causing 100% monthly turnover
 - Short selling is prohibited for most stocks in the momentum universe
 - The momentum premium has been negative since 2000
 - Monthly data contains too much noise to identify momentum
24. The professor says that if you left the class knowing only three factors, they should be market, value, and momentum. If you knew a fourth, which one did he say to add?
- SMB
 - CMA
 - The risk-free rate
 - Profitability
25. A momentum strategy rebalances monthly with 100% turnover (the entire portfolio is replaced each month). If each trade costs approximately 50 basis points (0.50%), the annual transaction cost drag is approximately:
- 0.5%
 - 3%
 - 6%
 - 12%
26. According to the professor, most practitioners diverge from the Fama-French 5-factor model by:
- Dropping value and keeping investment
 - Keeping profitability, using momentum, and dropping the investment factor
 - Replacing all FF factors with PCA-based statistical factors
 - Using only the market factor with leverage
27. In the week 6 lecture, the professor asks: what if a small stock behaves like a big stock? Under the Fama-French factor-pricing logic, what should NOT determine whether it earns a size premium?

- Its accounting label as a small-cap stock
 - Its beta or behavior with respect to the size factor
 - Its recent realized return over the last month
 - Whether it is in the Russell 2000 index
28. The practical solution to momentum's turnover problem, as described by the professor, is:
- Use weekly instead of monthly rebalancing
 - Use a 12-month ranking window instead of a 1-month window
 - Only trade the long side
 - Implement momentum with sector ETFs instead of individual stocks
29. Regarding the momentum premium's behavior in recent decades, the professor observed that:
- Momentum has been consistently profitable every decade since publication
 - Momentum performed well in the first ~15 years after publication but poorly in the most recent ~15 years
 - Momentum was only profitable among micro-cap stocks after publication
 - Momentum returns have become more volatile but maintained the same mean
30. According to the professor, the fourth most important factor after market, value, and momentum is:
- Size (SMB)
 - Investment (CMA)
 - Profitability (RMW)
 - Low volatility (BAB)
31. According to the professor, what happened to the size factor premium after it became widely known in the 1980s?
- It increased as more investors tried to exploit it
 - It remained stable but became more volatile
 - Its Sharpe ratio dropped off a cliff
 - It reversed sign and became negative
32. A researcher claims that momentum is entirely a small-stock phenomenon and does not exist among large-cap stocks. Based on the evidence discussed in class, this claim is:
- Correct: momentum is only significant among small-cap stocks
 - Incorrect: momentum is positive among both large and small stocks, though stronger in small stocks
 - Incorrect: momentum is actually stronger in large stocks
 - Correct: large-cap stocks are too efficiently priced for momentum to exist
33. The professor described a company that sells products to small firms. Even though the company itself is large, it might still earn the size premium because:

- Its revenues are correlated with the small-cap sector, giving it a high SMB beta
 - Its book-to-market ratio is similar to small firms
 - It is included in the small-cap index by construction
 - Its dividend yield matches that of small firms
34. For the same factor model applied to a set of test assets, the cross-sectional MAE is 1.2% and the time-series MAE is 2.5%. Is this result consistent with theory?
- No: cross-sectional errors should always be larger than time-series errors
 - Yes: the cross-sectional test has it easier, so its MAE must be less than or equal to the time-series MAE
 - No: the two error measures should be identical for any correctly specified model
 - Yes, but only if the model has at least three factors
35. A long-only momentum strategy that buys the top decile of past winners has a correlation to the overall market of approximately:
- -30%
 - 10%
 - 50%
 - 90%
36. A student claims: 'Alpha should always be zero in a good regression.' This statement is:
- Correct for factor pricing regressions only – in forecasting, alpha is just a re-leveling constant and irrelevant
 - Correct for both factor pricing and return forecasting regressions
 - Correct for forecasting regressions only – in pricing, alpha has no special meaning
 - Incorrect in all contexts – alpha is never exactly zero in practice
37. An analyst reports a forecasting R-squared of 0.01% for predicting monthly equity returns. The professor's response to this finding would be:
- This is a useful result that should be traded aggressively
 - This is too small to be economically meaningful — you would go bankrupt before the edge materializes
 - This is a strong result for monthly return forecasting
 - The R-squared is irrelevant; only the alpha matters for forecasting
38. Why does the course emphasize constructing factors as long-short portfolios rather than as long-only baskets?
- Long-short construction helps strip out broad market exposure so the factor is not just a disguised market bet
 - Long-only factors cannot be used in regressions
 - Long-short factors always have higher mean returns than long-only portfolios

- Long-short construction guarantees zero volatility
39. In the AQR Momentum case, comparing momentum returns in small stocks versus large stocks (1994-2024), the evidence suggests:
- Momentum is equally profitable in small and large stocks
 - Small-stock momentum has a higher Sharpe (36.9%) than large-stock momentum (11.4%), but large-stock momentum still earns a positive mean return
 - Large-stock momentum has a higher Sharpe ratio than small-stock momentum
 - Momentum is only profitable in small stocks; large-stock momentum has a negative mean return
40. † According to the professor, the two key ingredients that make momentum strategies profitable despite tiny individual stock autocorrelation are: (Select all that apply.)
- Trading only the extreme winners and losers (not the middle of the distribution)
 - Using leverage of at least 3:1
 - Diversifying across many stocks to reduce noise
 - Holding positions for multiple years to let autocorrelation compound
41. In the GMO/forecasting lecture, why is it a mistake to set a trading threshold using the full future sample rather than information available at time t ?
- It makes the strategy too conservative
 - It introduces leakage and can generate phenomenal fake results
 - It eliminates the forecast intercept
 - It forces the strategy beta to equal one
42. GMO nearly went out of business in the late 1990s. According to lecture, what caused this?
- GMO used too much leverage and suffered margin calls
 - GMO's value-oriented strategy underperformed during the growth-dominated tech bubble
 - GMO's risk models failed to account for correlation spikes
 - GMO was caught in an accounting scandal
43. DFA's investment philosophy is best described as tilting toward which rewarded risk exposures?
- Momentum beta and profitability beta
 - Size beta and value beta
 - Dividend yield and earnings growth
 - Market timing and volatility targeting
44. What fundamentally distinguishes GMO's investment approach from DFA and AQR?
- GMO uses leverage while DFA and AQR do not
 - GMO forecasts where returns will go; DFA and AQR identify exposures with equilibrium high returns

- GMO trades at higher frequency than DFA and AQR
 - GMO invests only in fixed income while DFA and AQR invest in equities
45. The E/P tilt strategy stayed invested through the 2008 financial crisis but was largely out of the market from 2014 onward. What explains this pattern?
- The strategy's risk model detected elevated volatility in 2014 and triggered a risk-off signal
 - Earnings-price ratios remained below the expanding median after 2014, signaling poor valuations for equities
 - The strategy's momentum signal turned negative after 2014
 - Transaction costs made rebalancing prohibitively expensive after 2014
46. A student implements a forecasting strategy by computing $r_{\text{strategy},t} = w_t \times r_t$. What error has the student made?
- The student should use excess returns instead of raw returns
 - The student is multiplying a time- t weight by time- t returns instead of time- $(t+1)$ returns
 - The student forgot to subtract transaction costs
 - The student should be using log returns
47. The professor used the analogy of a "voting machine" versus a "weighing machine." Which interpretation is correct?
- In the short run, markets weigh fundamentals; in the long run, markets vote on sentiment
 - In the short run, markets are a voting machine driven by sentiment; in the long run, they are a weighing machine driven by fundamentals
 - The voting machine refers to shareholder proxy votes affecting stock prices
 - The weighing machine refers to market-cap weighting in index construction
48. The lecture notes that a forecast-target correlation around 24% in the classroom exercise looks suspiciously high. What is the professor's immediate concern?
- The benchmark is too weak because it uses the expanding mean
 - The forecast is probably perfect and should be trusted without further checks
 - There may be some data leakage or other backtest contamination
 - Correlation is mathematically impossible for forecasting returns
49. A forecaster correctly identifies that a stock is overvalued based on fundamentals. In the short run, the stock continues to rise. Which concept from lecture best explains this outcome?
- The forecaster's model is wrong because prices are efficient
 - Short-run prices reflect whether others recognize the same mispricing, not whether the forecast is correct
 - The stock must have a higher beta than the forecaster assumed
 - Survivorship bias is inflating the stock's apparent return

50. A quant team's return forecasts have a negative OOS R-squared. Management considers firing the team. According to the lecture, why might this be a mistake?
- Negative OOS R-squared always results from insufficient data and will improve with more observations
 - OOS R-squared is always negative for equity forecasts, so this is normal
 - The team should simply multiply their forecasts by -1 to flip the sign of R-squared
 - The forecasts may still rank assets correctly — negative OOS R-squared reflects level/scale miscalibration, not sorting ability
51. When computing a threshold (e.g., a median) for a trading signal, why must the threshold be an expanding-window statistic rather than a full-sample statistic?
- Expanding windows reduce estimation error
 - Expanding windows increase the strategy's Sharpe ratio
 - Full-sample statistics are computationally more expensive
 - Full-sample statistics introduce look-ahead bias by using future data
52. In the E/P tilt strategy discussed in lecture, when the E/P ratio is above its expanding median, the portfolio allocation is:
- 50% equities, 50% risk-free
 - 100% equities
 - 150% equities (borrowing 50% at the risk-free rate)
 - 200% equities (borrowing 100% at the risk-free rate)
53. According to a formal theorem discussed in lecture, how should the variability of an optimal forecast compare to the variability of the quantity being forecast?
- The forecast should be more variable than the target to capture tail events
 - The forecast should have the same variability as the target
 - The forecast should be less variable than the target
 - There is no theoretical relationship between forecast and target variability
54. When constructing a forecasting-based trading strategy, portfolio weights at time t must be a function of:
- Returns realized at time t
 - Information available at time t only
 - Returns realized at time $t + 1$
 - The full-sample average of the signal
55. The professor called this his 'trick question': In the forecasting regression $\tilde{r}_{t+1} = \alpha + \beta x_t + \epsilon_{t+1}$, the professor asks whether we want α to be zero. The answer is:
- Yes – a nonzero alpha indicates model misspecification, just like in factor pricing
 - Yes – alpha should be zero for the forecast to be unbiased

- No – in forecasting, we only care whether beta is nonzero; alpha is just a re-leveling constant that does not affect forecast quality
 - No – we want alpha to be as large as possible to maximize expected returns
56. A fund forecasts returns for 500 individual stocks using separate models for each stock. According to lecture, what fundamental problem arises?
- The fund will have too many positions to manage operationally
 - Individual stock forecasts from separate models will be internally inconsistent, potentially creating near-arbitrage opportunities against the fund's own positions
 - Individual stock models always have lower R-squared than market models
 - Forecasting 500 stocks requires at least 500 factors
57. How does the lecture suggest resolving the internal consistency problem when forecasting individual stock returns?
- Use the same model specification for every stock
 - Average the forecasts across all stocks to remove inconsistencies
 - Forecast a small number of factor returns and derive stock-level expected returns through factor loadings
 - Discard any stock that is highly correlated with other stocks
58. According to the classic Carhart finding on mutual fund persistence, a fund with strong recent performance:
- Is very likely to continue outperforming
 - Faces roughly a coin flip on whether it will outperform in the next period
 - Is guaranteed to mean-revert and underperform
 - Will outperform if it has positive momentum exposure
59. The GMO lecture reports that GMWAX is about 80% correlated with SPY, about 86% with its benchmark, and has only a modest information ratio. What is the best allocator-level conclusion?
- GMO is basically unrelated to broad equities
 - GMO's alpha is so strong that correlation information becomes irrelevant
 - The benchmark must be wrong because SPY correlation is not exactly 100%
 - GMO behaves a lot like equities or its benchmark, with only a modest extra piece beyond that
60. In the forecasting SPY exercise, the DP (dividend-price ratio) forecast regression has an in-sample R-squared of approximately 0.68%. When this forecast is used out-of-sample with an expanding window, the most likely outcome for the OOS R-squared is:
- Similar to in-sample, around 0.5-1.0%
 - Higher than in-sample, around 2-3%
 - Approximately zero or negative

- Around 5%, due to the longer estimation window
61. LTCM had approximately \$8 billion in equity capital. Their convergence trades earned spreads of roughly 1% when they converged. At 12x leverage, the expected return on equity from a convergence trade was approximately:
- 1%
 - 4%
 - 12%
 - 50%
62. Why did 'flight to quality' create a common-factor problem for LTCM's many spread trades?
- Because it made all risky assets and low-quality instruments cheap at the same time
 - Because many seemingly unrelated positions were effectively short lower-quality assets and long higher-quality ones
 - Because it only affected Russian bonds and nothing else in the portfolio
 - Because it eliminated all correlations across trades
63. The lecture's Royal Dutch/Shell example makes what asymmetry point about convergence trades?
- Potential profit from convergence is capped by the initial spread, but divergence can be much larger
 - Potential profit is unlimited if the spread keeps widening
 - Losses are capped at the initial spread because the firms are the same company
 - The trade cannot lose because the assets have the same underlying business
64. The lecture explains that market timing a single security (SPY) using fundamental signals can produce a strategy with only 45% R-squared against passive SPY. This illustrates that:
- The signals have no forecasting power
 - Dynamic positioning in a single asset can create a return stream that behaves like an entirely different investment
 - SPY is not a valid benchmark for timing strategies
 - Regression R-squared is meaningless for timing strategies
65. The lecture emphasizes that correlation between forecast and target may be superior to out-of-sample R-squared for evaluating forecasts because:
- Correlation is always higher than R-squared
 - Correlation is invariant to additive shifts in the forecast level
 - Correlation accounts for non-linear relationships
 - Correlation penalizes overfitting more harshly
66. In the lecture, GMO is distinguished from DFA and AQR primarily because GMO:
- Invests only in fixed income

- Relies on equilibrium factor risk premia
 - Actively forecasts returns using fundamentals like earnings and dividends
 - Uses high-frequency trading strategies
67. The out-of-sample R-squared metric compares a model's forecast errors to:
- The in-sample R-squared of the regression
 - A benchmark forecast such as the expanding mean
 - The volatility of the target variable
 - The Sharpe ratio of the passive strategy
68. In the lecture, when discussing the earnings-price ratio tilt strategy, the instructor notes that since 2014 the strategy has mostly been out of the market. Why?
- Earnings collapsed during this period
 - Earnings-price ratios were below their expanding median, indicating expensive valuations
 - Interest rates were too high for equities to be attractive
 - The strategy's lookback window was too short
69. In the lecture, the instructor describes a practical experience where a quant research team had negative out-of-sample R-squared but their forecasts still generated positive alpha. The key insight was:
- Their forecasts had high R-squared in-sample
 - Their forecasts did a good job of sorting bulls from bears, but were poorly scaled
 - They were using too few signals
 - Their forecasts were perfectly correlated with the market
70. In the Forecasting with Fundamentals exercise, why does sector-hedging tend to reduce the apparent alpha of a dividend-yield carry strategy?
- Because much of the strategy's return is actually coming from sector tilts rather than pure stock-level carry
 - Because sector hedging mechanically forces the alpha to zero in every regression
 - Because dividend yield cannot be used as a forecasting variable
 - Because the strategy becomes a pure risk-free arbitrage after hedging
71. The lecture uses the analogy of a "voting machine vs weighing machine" in the context of GMO's investment philosophy. A convergence trader like LTCM faces the related risk that:
- Fundamentals never determine prices in financial markets
 - Even if you are right about fundamental value, the market may not recognize it before you run out of capital
 - Leverage is always beneficial for convergence trades
 - Flight-to-quality events are unpredictable and therefore unhedgeable

72. In the LTCM case, extending the regression to include a quadratic market term (SPY-squared), the coefficient on SPY-squared was approximately -2.16. This indicates LTCM behaved as if it was:
- Long market options (positive gamma)
 - Short market options (negative gamma)
 - Delta-neutral with no volatility exposure
 - Long a straddle on the market
73. The professor argued that LTCM's convergence trades are 'a win-win – if I only used my own money.' If the spread converges, you profit. If the spread diverges, the trade looks even better. The professor then asked: what breaks this logic?
- The counterparty might default, eliminating the position entirely
 - The spread can diverge without bound because the underlying assets are fundamentally different
 - With leverage, a widening spread triggers margin calls that force you to close the position at the worst possible time – exactly when the trade looks best
 - Transaction costs on spread trades are too high to capture the small convergence premium
74. The LTCM case lists four funding risks. Which of the following is NOT one of them?
- Collateral haircuts
 - Repo maturity
 - Equity redemption
 - Dividend tax risk
75. The Royal Dutch/Shell trade was an example of:
- A momentum strategy
 - An equity relative value trade
 - A swap spread convergence trade
 - A fixed-rate mortgage carry trade
76. In the TA review's collar example, a portfolio that is long SPY and buys a 3% OTM put while selling a 3% OTM call will have monthly returns bounded between:
- -3% and +3%
 - -5% and +5%
 - 0% and +3%
 - -3% and unlimited upside
77. In the LTCM case, the piecewise regression of LTCM net excess returns on SPY reveals that LTCM is:
- Long the call-like factor and long the put-like factor
 - Short the call-like factor and short the put-like factor

- Short the call-like factor and long the put-like factor
 - Long the call-like factor and short the put-like factor
78. In the TA review, LTCM's investment strategy is described as having a typical holding period of:
- Intraday to a few days
 - A few weeks
 - 6 months to 2 years
 - 5 to 10 years
79. † **[Select All That Apply]** LTCM's gross returns had annualized mean 24.4% and volatility 13.6%, yielding a Sharpe ratio of approximately 1.79. Over the same period, SPY had a Sharpe of approximately 1.36. Yet LTCM's net return skewness was -0.81 versus SPY's -0.41. Which combination correctly explains this pattern?
- LTCM's strategy was effectively short volatility, generating consistent small gains with occasional large losses
 - LTCM's leverage amplified both returns and tail losses
 - LTCM's strategy had purely linear exposure to equity markets
 - LTCM's returns were positively skewed due to convergence trade payoffs
80. † **[Select All That Apply]** The TA review identifies the fundamental problem with all of LTCM's trades. Which of the following captures the core issue?
- The mispricings they exploited were typically very small
 - Massive leverage was required to generate meaningful returns from small mispricings
 - Their trades had unlimited downside risk even without leverage
 - They could only trade in equity markets
81. LTCM returned approximately \$2 billion to outside investors before its collapse. According to the lecture discussion, why was this decision so consequential?
- It violated the fund's lockup provisions and triggered regulatory action
 - It increased the fund's leverage ratio dramatically (back to ~20x), amplifying losses when spreads diverged
 - It forced the fund to close out its most profitable positions
 - It signaled weakness to counterparties, triggering margin calls
82. What is Asness's main complaint in the lecture's lagged-beta example: that reported hedge-fund returns are wrong in level, or wrong in timing?
- Wrong in timing, because one month's reported return includes some of neighboring months
 - Wrong in level, because the average return must be negative
 - Wrong in both level and timing, and therefore unusable
 - Neither; the point is only about taxes

83. The week 9 lecture jokes that if you held SPY in a private-equity vehicle and only reported returns with a lag, what would happen to the reported volatility and correlation?
- Both would look better than they really are
 - Both would look worse than they really are
 - Volatility would rise but correlation would stay unchanged
 - Nothing would change because packaging cannot affect statistics
84. Why does diversifying across many hedge fund managers potentially fail to add value, according to the lecture?
- Hedge fund returns are perfectly correlated with each other
 - Combining many active long-short positions tends to reconstitute a passive market portfolio, but at much higher fees
 - Diversification reduces alpha without reducing beta
 - Hedge fund lockup periods prevent rebalancing
85. Why do lock-up provisions in hedge funds potentially benefit remaining investors, according to the lecture?
- They allow the fund to charge higher management fees
 - They prevent bank-run dynamics and fire-sale liquidation
 - They guarantee a minimum return over the lock-up period
 - They reduce the fund's tax obligations
86. A hedge fund manages \$100 million and charges 2% management fee plus 20% performance fee. In a year when the fund's gross return is 20% (above the high-water mark), the total fees paid to the manager are:
- \$2 million
 - \$4 million
 - \$6 million
 - \$20 million
87. A fund has iid annual returns with volatility 15% and a sample average return of 10% over $T = 5$ years. What is the approximate 95% confidence interval for the true mean return?
- 10% plus or minus 3%
 - 10% plus or minus 7%
 - 10% plus or minus 13%
 - 10% plus or minus 20%
88. [†] Which observations from the week 8 and week 9 materials support the course theme that labels can mislead about behavior? (Select all that apply.)
- GMO can be marketed as a distinct forecast-driven product and still end up highly correlated with broad equity benchmarks

- A private-equity reporting wrapper can make even liquid exposures look artificially smooth
- Return-based decomposition can reveal behavior even when the manager's narrative is vague
- A fund's stated label is enough to infer its true factor exposures without return data

89. Over 2000-2011, the GMO fund (GMWAX) and SPY show:

Fund	Cumulative Return	Ann. Sharpe
GMWAX	+25%	Positive
SPY	~0%	Near zero

Over 2012-2025, these results reverse: SPY dominates with a Sharpe near 90%, while GMWAX lags. The professor's key concern about GMO's contrarian approach is:

- Contrarian strategies always underperform in the long run
 - GMO's models are fundamentally wrong because they missed the tech rally
 - All value-oriented strategies should be avoided because the value premium has disappeared
 - If the periods of underperformance come first, the manager may lose capital (and clients) before the strategy pays off – this is career risk
90. What strike price does STCM use for its put options relative to the current market price?
- At-the-money ($K = S$)
 - 5% out of the money ($K = 0.95S$)
 - 20% out of the money ($K = 0.80S$)
 - 30% out of the money ($K = 0.70S$)
91. A hedge fund reports a Sharpe ratio of 1.0 based on 25 years of annual data. Using $SE(SR) = \sqrt{(1 + SR^2/2)/T}$, the standard error of this Sharpe ratio estimate is approximately:
- 0.04
 - 0.10
 - 0.25
 - 1.00
92. Two hedge funds each report a single +20% year. Fund A launched 8 months ago. Fund B has a 10-year history of mostly ordinary results. According to the week 9 lecture's Bayesian flow discussion, which fund should see the larger investor inflow from that one good year?
- Fund A, because one strong year updates beliefs more when the manager's skill is still highly uncertain
 - Fund B, because a long track record makes the new return more informative than for a young fund
 - Both should attract roughly the same inflow, because investors care only about the most recent return

- Fund B, because older funds face less redemption risk and therefore stronger return chasing
93. According to Carhart (1997), as discussed in lecture, which type of fund performance is persistent?
- Top-decile outperformance persists strongly
 - Middle-decile performance is the most stable
 - Poor performance (bottom decile) is persistent
 - Both top and bottom decile performance persist equally
94. What is the standard error formula for the Sharpe ratio, as presented in the TA review notebook?
- $SE(SR) = 1/\sqrt{T}$
 - $SE(SR) = SR/\sqrt{T}$
 - $SE(SR) = \sqrt{(1 + SR^2/2)/T}$
 - $SE(SR) = \sqrt{(1 + SR)/T}$
95. The professor's STCM (Short-Term Capital Management) simulation showed that conditional on surviving months, the fund had a Sharpe ratio of 150% and mean of 25%/year. But when 3 bust months were added back, the mean dropped to -27% and the Sharpe fell to -13%. The key lesson is:
- Tail-risk strategies should be evaluated using only the months when they are operational, since bust months are by definition extreme outliers
 - The Sharpe ratio is robust to outliers and should not change this dramatically from three data points
 - The STCM strategy is profitable as long as it avoids the three specific crisis months
 - Any investor being pitched a hedge fund is seeing conditional-on-survival statistics; the unconditional distribution – which includes the busts – tells a fundamentally different story
96. LTCM's convergence trades on bond spreads were leveraged approximately 30-50x. If a spread trade earns 30 basis points unleveraged and is leveraged 10x, what is the approximate leveraged return?
- 3%
 - 9%
 - 15%
 - 30%
97. A hedge fund has a 5-year track record with an observed annualized mean return of 10% and annualized volatility of 15%. Using the CLT approximation, the 95% confidence interval for the fund's true mean annual return is approximately:
- 7% to 13%

- 4% to 16%
 - 0% to 20%
 - -3% to 23%
98. An ETF markets itself as a value strategy, but a return-based factor decomposition shows little loading on value and much larger exposure to other risks. According to the course's week 9 takeaway, what should an allocator conclude?
- The label is decisive; if the prospectus says value, the fund should still be treated as value
 - Without seeing every holding, no inference about behavior is possible
 - Behavior matters more than marketing labels, and decomposition is the right tool for checking what the fund really does
 - Weak value loading proves the fund has generated true alpha
99. What does the phrase "you can't eat Sharpe ratio" mean in the context of hedge fund evaluation?
- The Sharpe ratio is unreliable for non-normal distributions
 - A high Sharpe ratio on a small asset base does not generate enough PnL to matter
 - Sharpe ratios cannot be compared across different funds
 - The Sharpe ratio ignores transaction costs and taxes
100. According to lecture, what was the Royal Dutch / Shell trade that LTCM engaged in?
- Betting that Shell's oil production would exceed Royal Dutch's
 - Exploiting a price discrepancy between two share classes of the same company that should converge
 - A currency arbitrage between Dutch guilder and British pound denominated shares
 - A merger arbitrage bet that Royal Dutch would acquire Shell