FINM 36700 — Midterm Solutions

Portfolio and Risk Management

Autumn 2026

Name: _____

UChicago ID: _____

- This exam is **closed book** and **closed notes**. No calculator or other electronics.
- Questions marked † **[Select All That Apply]** may have zero, one, or more correct choices. Mark every correct choice, and leave the question blank if none is correct. Full credit requires exactly the correct set; otherwise the question earns zero.
- All other questions have exactly one correct choice.

Question type	**Questions**	**Points each**	**Points**
Single-select	35	2	70
Multiple-select (†)	27	2	54
Total	**62**		**124**

Returns and performance

Exhibit 1: Two fund histories Hypothetical annualized sample statistics over the same dates. Means are arithmetic total returns. The risk-free return is a constant 3% per year; borrowing and lending use that rate. Maximum drawdown is measured over the complete sample, not annualized.

Fund	Mean	Volatility	Maximum drawdown
A	9%	10%	-38%
B	12%	20%	-24%

1. Using Exhibit 1, an investor scales Fund A with borrowing to match Fund B's 12% mean total return. Ignoring borrowing constraints, what annualized volatility does that scaled investment have?

- ✗ **A.** 13.3%
- ✓ **B.** 15%
- ✗ **C.** 18%
- ✗ **D.** 20%
- ✗ **E.** 30%

B: A has a 6% excess mean. Matching B's 9% excess mean requires 1.5 times A and a -0.5 cash weight, so volatility is 15%. A scales by the ratio of total means, 12/9, and ignores that the borrowed third costs 3%; that position's total mean is only 11%. C implies 1.8 times A and a 13.8% total mean, so it misses the stated target. D copies Fund B's risk. E triples A, giving a 21% total mean.

2. † **[Select All That Apply]** Which conclusions about Exhibit 1 are correct?

- ✗ **A.** Fund B has the higher Sharpe ratio because its total mean is higher.
- ✓ **B.** Fund A has the higher Sharpe ratio once the 3% risk-free return is subtracted.
- ✗ **C.** An equal-weight mix of A and B has volatility above 15%.
- ✓ **D.** The maximum-drawdown comparison and the Sharpe-ratio comparison favor different funds.
- ✓ **E.** Subtracting the 3% risk-free return changes each fund's mean but not its volatility.

B: Sharpe ratios are (9-3)/10=0.60 and (12-3)/20=0.45. D: B has the smaller historical drawdown. E: the risk-free return is a constant, so subtracting it shifts each mean by 3% and leaves each volatility where it was. A ranks the funds by total mean alone; the Sharpe ratio subtracts the 3% risk-free return and divides by volatility, giving Fund B the lower ratio. C: with correlation below one the mix's volatility is below 15%, and it equals 15% only if the correlation is one.

Exhibit 2: Two investment paths Hypothetical wealth from an initial investment of 100, with no deposits, withdrawals, fees, or distributions outside the account. The intervals have equal length.

Date	Path A	Path B
0	100.00	100.00
1	120.00	120.00
2	96.00	144.00
3	115.20	115.20
4	92.16	92.16

3. † **[Select All That Apply]** Which statistics are identical for Paths A and B in Exhibit 2?

- ✓ **A.** The arithmetic mean of the four simple returns.
- ✓ **B.** The standard deviation of the four simple returns.
- ✓ **C.** The sum of the four log returns.
- ✗ **D.** The maximum drawdown.
- ✓ **E.** The worst single-interval simple return.

A, B, and E: both paths contain two +20% and two -20% returns, so any statistic of the unordered set agrees, including the worst interval at -20%. The largest dollar drops differ, 24 against 28.8, but losses are measured in returns. C: both finish at the same wealth, so cumulative log returns agree. D: the ordering changes the running peak; the maxima are 23.2% and 36%.

4. Over one year, a portfolio earns a 6% nominal total return, a Treasury bill earns 4%, and inflation is 5%. Which pair correctly describes the portfolio's excess return and its approximate real return?

- ✓ **A.** 2% excess; 1% real.
- ✗ **B.** 1% excess; 2% real.
- ✗ **C.** 2% excess; 6% real.
- ✗ **D.** 6% excess; 1% real.
- ✗ **E.** 2% excess; -3% real.

A: excess return is 6%-4%=2%; real return is 1.06/1.05-1, approximately 1%. B interchanges the benchmarks. C and D leave a total return unadjusted. E subtracts inflation from the excess return instead of from the nominal return; the real return deflates the 6% by inflation alone, and the T-bill plays no part in it.

5. A performance report lists a monthly mean excess return of 1%, a monthly standard deviation of 4%, and an annualized Sharpe ratio of 3.0. Under the usual no-serial-correlation annualization, which error produced the 3.0?

- ✗ **A.** Both the mean and the standard deviation were annualized by 12.
- ✗ **B.** The mean was annualized by 12 and the standard deviation by the square root of 12.
- ✗ **C.** Both the mean and the standard deviation were annualized by the square root of 12.
- ✗ **D.** The mean was annualized by the square root of 12 and the standard deviation by 12.
- ✓ **E.** The mean was annualized by 12, but the standard deviation was left monthly.

E: 12 times 1% is 12%, and 12% divided by the monthly 4% is 3.0. The correct annualized ratio is the monthly 0.25 times the square root of 12, about 0.87. A and C scale numerator and denominator alike and leave the ratio at 0.25. B is the correct annualization and gives 0.87, not 3.0. D gives about 0.07.

6. † **[Select All That Apply]** Which statements about analyzing a stock index's history are correct?

- ✓ **A.** Returns are analyzed rather than price levels because a price series is not stationary: its sample mean and variance do not estimate stable quantities.
- ✓ **B.** Weekly index returns are close to serially uncorrelated, so last week's return says little about the sign of this week's.
- ✓ **C.** The clearest failure of the IID assumption is in volatility: calm and turbulent periods cluster, so the variance is not constant over time.
- ✗ **D.** Normality is assumed because daily returns support it; the assumption breaks down only when returns are aggregated to monthly frequency.
- ✗ **E.** Converting prices to returns removes the clustering of volatility, so a single historical variance describes every period equally well.

ABC and AB both receive full credit. C is optional because it says the variance is not constant over time without naming the conditional variance: a stationary return series with volatility clustering, such as a GARCH process, has a changing conditional variance and a constant unconditional variance, so a reader who takes the variance to mean the unconditional variance correctly rejects C. The lectures used variance as shorthand for the conditional variance, and a correct econometric reading is not penalized to match that shorthand. A and B must still be selected and D and E must be omitted. The gray check marks optional choice C.

A, B, and C are the lecture's sequence: stationarity is needed, so returns rather than prices; independence is wished for, and returns show little serial correlation; but the conditional variance moves in clusters, which is where IID fails. C describes that conditional variance: a stationary return series with volatility clustering, such as a GARCH process, has a changing conditional variance and a constant unconditional variance. D reverses the frequencies: normality fails most at daily frequency and is assumed for convenience, not because the data support it. E: the return transformation fixes the stationarity of the level, not the clustering of volatility, so no single variance describes calm and turbulent periods alike.

7. Where does the normal model for an equity index's returns fail most clearly?

- ✓ **A.** In the tails, and more at daily than at monthly frequency: the most extreme standardized daily moves lie far beyond anything a normal distribution allows, while monthly extremes are much milder.
- ✗ **B.** At every quantile about equally, because the sample mean and variance do not characterize the data at any horizon.
- ✗ **C.** In the center of the distribution: the histogram of monthly index returns is visibly asymmetric and flat around its mean.
- ✗ **D.** At monthly more than daily frequency, because aggregating daily returns over a month compounds the extreme days.
- ✗ **E.** Only in the positive tail, since large gains are rarer under normality than large losses of the same size.

A: the center of the histogram is close to bell-shaped; the violation is in the tails, and it is far larger at daily frequency than at monthly (in the class example, daily S&P 500 z-scores reached about plus or minus 10 while monthly extremes were near plus or minus 3; the exact figures depend on the index and sample). C puts the failure in the wrong place. B overstates it. D reverses the frequencies; aggregation moves the distribution toward normal. E: the failures appear in both tails, and for equities the lower tail is the heavier. A's phrase "beyond anything a normal distribution allows" is loose: the normal has unbounded support, so a ten-standard-deviation day is possible under it, with probability far too small to expect even one such day in centuries of data. The point is that such days occur in the data far more often than the normal model predicts. The key is unchanged because no other option places the failure correctly.

8. † **[Select All That Apply]** Which of these return streams would you expect to be positively skewed?

- ✗ **A.** A diversified equity index, which tends to rise gradually and fall sharply.
- ✓ **B.** A venture portfolio of twenty start-ups in which a few large successes carry the fund.
- ✗ **C.** A strategy that sells insurance against market crashes and collects premiums in calm years.
- ✓ **D.** A long position in far out-of-the-money call options that usually expire worthless.
- ✓ **E.** Bitcoin, because its volatility is several times that of an equity index.

BD and BDE both receive full credit. E is optional because the key assumed Bitcoin's returns are roughly symmetric, which the Week 1 table shows at daily frequency (skewness 0.04) but not at the monthly frequency the homework uses for skewness: in the same data Bitcoin's monthly simple returns have skewness +0.69 while its log returns are close to symmetric, a pattern consistent with compounding turning high volatility into positive skewness of simple returns. E does not name that mechanism; the accommodation rests on the unspecified frequency and on the key having generalized from daily evidence. The question does not fix the frequency, and a correct econometric reading is not penalized to match the lecture's daily-frequency shorthand. B and D must still be selected and A and C must be omitted. The gray check marks optional choice E.

B and D are lottery-like payoffs: many small losses and a few large gains. A is the "up the stairs, down the elevator" pattern, negatively skewed. C collects small premiums and takes occasional large losses, also negatively skewed. E: at daily frequency, the lecture's frame, Bitcoin's returns are close to symmetric (skewness 0.04 in the Week 1 table) and its high volatility is a statement about scale rather than shape. At monthly frequency the same data show skewness of +0.69 with log returns close to symmetric, the pattern that compounding a high-volatility process produces, so at that frequency E's stream is positively skewed.

9. Over the same sample a high-yield bond ETF has a Sharpe ratio of 0.30 and an information ratio of -0.46 relative to the equity market. What does the combination tell an investor who already holds the market?

- ✗ **A.** The ETF improves the portfolio because its Sharpe ratio is positive; once the market is already held, the information ratio carries no additional information.
- ✓ **B.** Its sample alpha on the market is negative, so a small long addition lowers a market holding's in-sample Sharpe ratio even though the ETF's own excess mean is positive.
- ✗ **C.** The two ratios are inconsistent: a positive Sharpe ratio implies a positive alpha against a benchmark that itself has a positive premium.
- ✗ **D.** The information ratio is negative only because it uses total rather than excess returns in its numerator, so it understates the ETF's performance.
- ✗ **E.** The ETF is a useful hedge for the market, because a negative information ratio means its residual return moves against the market's return.

B: the information ratio is the Sharpe ratio of the part of the return the market does not explain; a negative value means a negative sample alpha, and at the market-only position the sign of the Sharpe-ratio change from a small long addition is the sign of that alpha. Making the ETF available does not lower the best attainable ratio, since zero or short exposure remains feasible; it is the long addition that hurts. A uses the wrong metric for an addition to an existing holding. C: a positive Sharpe ratio and a negative alpha coexist whenever the market exposure delivers more than the asset earns. D: the ratio is alpha over residual volatility, both from excess returns. E: the residual is uncorrelated with the market by construction. The information ratio takes its sign from alpha, which measures mean return beyond the market contribution; a negative ratio does not imply negative market co-movement.

Portfolio allocation

Exhibit 3: Two assets Hypothetical annual excess returns on two assets, with parameters treated as known. Returns are in excess of a risk-free asset that can be held long or short without limit.

Asset	Mean excess return	Volatility	Correlation with the other asset
1	8%	20%	0.5
2	4%	20%	0.5

10. In Exhibit 3, without computing it exactly, which range contains the volatility of a portfolio with 50% in each asset, and why?

✗ **A.** Exactly 20%, because both assets have 20% volatility and the weights sum to one.
✗ **B.** Exactly 10%, because each asset's 20% volatility is halved by its 50% weight.
✓ **C.** Between 14.1% and 20%, because the correlation of 0.5 lies between zero and one.
✗ **D.** Below 14.1%, because a correlation of 0.5 removes more than half of the variance.
✗ **E.** Exactly 14.1%, because each 50% weight adds a quarter of its asset's variance.

C: with equal volatilities the 50/50 variance is 0.02 times one plus the correlation: 0.02 at correlation zero (14.1%), 0.04 at correlation one (20%), and 0.03 here (17.3%). A treats volatility as linear in the weights, which holds only at correlation one. B halves a single asset's volatility. D has the direction wrong: diversification removes less variance as the correlation rises toward one. E drops the covariance term; the two quarter-variances alone give 0.02, which is the correlation-zero case, and the positive correlation adds 0.01 more.

11. Regress asset 2's excess return on asset 1's excess return in Exhibit 3. What is asset 2's alpha, and what does it imply for asset 2's weight in the tangency portfolio?

✗ **A.** Alpha is zero, but asset 2 receives a positive weight because its Sharpe ratio is positive.
✗ **B.** Alpha is 4%, and asset 2 receives a positive weight.
✗ **C.** Alpha is -4%, and asset 2 is shorted.
✗ **D.** Alpha is zero, and asset 2 receives the same weight as asset 1 because the volatilities are equal.
✓ **E.** Alpha is zero, and asset 2 receives zero weight.

E: beta is 0.5(20/20)=0.5, so alpha is 4%-0.5(8%)=0. An asset with zero alpha relative to an asset already held adds nothing at the margin. Directly, $\Sigma^{-1}\mu$ is proportional to (0.08-0.5(0.04), -0.5(0.08)+0.04)=(0.06, 0). A ranks by stand-alone Sharpe ratio; a positive ratio earns no weight when asset 1 already delivers everything asset 2 offers. B is the raw mean. C has the sign of a short that a zero alpha does not justify. D confuses equal volatility with equal usefulness.

12. † **[Select All That Apply]** An investor holds \$1 of asset 2 in Exhibit 3 and hedges the exposure to asset 1 with the regression hedge. Which statements are correct?

- ✓ **A.** The hedge shorts \$0.50 of asset 1 per \$1 of asset 2.
- ✓ **B.** The hedged position's expected excess return equals asset 2's alpha on asset 1.
- ✗ **C.** The hedged position has zero volatility.
- ✓ **D.** A dollar-for-dollar hedge, short \$1.00 of asset 1, would leave the position with negative exposure to asset 1.
- ✗ **E.** The hedge removes half of asset 2's variance, since the correlation with asset 1 is 0.5.

A: the hedge ratio is $\rho\sigma_2/\sigma_1 = 0.5$, the regression beta. B: the hedged mean is the alpha, which is zero here. D: exposure would be 0.5-1=-0.5. C: the residual variance remains; the hedged volatility is 20% times the square root of 0.75, about 17%. E confuses the correlation with its square: the hedge removes R^2, a quarter of the variance, not half.

Exhibit 4: Harvard ETF allocation exercise Monthly ETF excess returns, February 2011–July 2026; SHV is subtracted and QAI is excluded. Each allocation is scaled to a 1% monthly sample mean excess return. The table annualizes means by 12 and volatilities by the square root of 12. All rows are evaluated using the same original sample covariance matrix. Shorting and borrowing are allowed. Before scaling, the equal-weight portfolio (10% in each ETF) has an annualized mean excess return of 4.3% and volatility 9.9%.

Allocation	Mean excess return	Volatility
Tangency	12.0%	8.1%
Equal weight	12.0%	27.5%
Inverse variance	12.0%	32.6%
Regularized	12.0%	10.5%

13. Why is the tangency portfolio's lower volatility than the regularized portfolio in Exhibit 4 unsurprising?

- ✗ **A.** The regularized portfolio is required by its construction to hold more risky assets than the tangency portfolio.
- ✗ **B.** The table evaluates each portfolio with its own covariance estimate rather than a common one.
- ✗ **C.** Shrinking the covariances raises the regularized portfolio's realized sample volatility by construction.
- ✗ **D.** The sample covariance matrix was estimated accurately enough that shrinking it only discarded information.
- ✓ **E.** Tangency weights maximize the in-sample Sharpe ratio; at a common mean, nothing has lower in-sample volatility.

E is the in-sample optimization comparison at a common mean. A is not a constraint in the setup. B contradicts the common evaluation matrix. C confuses an estimation input with realized returns. D reads an in-sample win as evidence about estimation accuracy; the tangency portfolio beats every alternative in the sample that produced it, however noisy the estimate.

14. † **[Select All That Apply]** Which interpretations of Exhibit 4 are correct?

- ✓ **A.** The common 12% mean is imposed by rescaling, not evidence that the methods forecast the same return.
- ✗ **B.** The inverse-variance rule uses the cross-asset covariances when assigning its relative weights across assets.
- ✓ **C.** The equal-weight row holds equal risky weights that sum to more than one after scaling; borrowing finances the difference.
- ✓ **D.** At the common mean, inverse variance has the lowest Sharpe ratio of the four allocations.
- ✓ **E.** Inverse variance tilts toward the lower-volatility ETFs, yet at the common mean it has more volatility than equal weight.

A: the target is imposed. C: scaling equal weights up to a 12% mean needs about 277% in risky assets, financed by borrowing. D: 12/32.6 is the smallest ratio. E: before scaling, inverse variance has a sample mean of 2.18% and volatility of 5.93%, versus 4.33% and 9.94% for equal weight. Reaching the 12% target requires about 5.50 times the inverse-variance portfolio and 2.77 times equal weight. The scaling reverses their volatility ranking: 32.6% versus 27.5%. This is a result for this sample, not a general ranking of the rules. B describes a covariance-aware rule, not inverse stand-alone variance.

15. For an unconstrained risky portfolio, $w \propto \Sigma^{-1}\mu^e$. If all expected excess returns are multiplied by the same positive constant, with covariance unchanged, what happens to the risky weights normalized to sum to one? Assume the original normalization is well defined.

- ✗ **A.** Every normalized weight is multiplied by that constant.
- ✗ **B.** They become equal weights.
- ✗ **C.** Every short position becomes a long position.
- ✓ **D.** They are unchanged.
- ✗ **E.** They tilt further toward the highest-mean assets.

D: the common multiplier cancels when the vector is divided by its sum. A describes the unnormalized vector. B: a common multiplier preserves relative weights and cannot equalize unequal ones. C would require a sign change, whereas the multiplier is positive. E reads the wider gaps between means as a stronger tilt; the ratios between the means, which fix the direction of the weight vector, are unchanged.

Exhibit 5: A covariance estimate before and after shrinkage Hypothetical annual covariance estimates for assets X and Y, using returns as decimals. The rule is $\widetilde{\Sigma} = (\Sigma + \text{diag}(\Sigma))/2$. Expected excess returns are unchanged.

Entry	Original	Regularized
Variance of X	0.040	0.040
Variance of Y	0.010	0.010
Covariance of X and Y	0.012	0.006

16. † **[Select All That Apply]** An optimizer produces large offsetting long and short positions. Which statements about regularizing its covariance input with the rule in Exhibit 5 are correct?

- ✓ **A.** It shrinks each covariance toward zero and leaves each variance unchanged.
- ✓ **B.** For a position long \$1 of X and short \$1 of Y, it raises the estimated portfolio variance, because it shrinks a positive covariance that was offsetting part of that position's risk.
- ✗ **C.** It is equivalent to prohibiting short positions.
- ✗ **D.** It multiplies every optimal weight by one half.
- ✗ **E.** It would still improve the weights if the covariance matrix were known rather than estimated.

A is the rule. B: the long-short variance 0.04+0.01-2(0.012)=0.026 becomes 0.05-0.012=0.038; shrinking a positive covariance lowers the variance of a long-long position and raises it for a long-short one. C: shorts remain allowed; Exhibit 4's regularized portfolio still has shorts. D: weights change relative to one another, not by a common factor. E: shrinkage trades bias for lower estimation variance; with a known covariance matrix there is no estimation error to reduce, and only the bias remains.

Exhibit 6: Harvard TIPS sensitivity The same monthly excess returns as Exhibit 4: February 2011–July 2026, 186 months, ten ETFs with QAI excluded. TIP is the TIPS ETF. Its annualized mean excess return is 1.3%, its volatility 5.0%, and its Sharpe ratio 0.26; its correlation is 0.76 with IEF and 0.39 with SPY. The standard error of TIP's monthly sample mean is about 0.11%. Tangency weights are normalized to sum to one; shorting is allowed.

Case	TIP weight in the tangency portfolio	Tangency Sharpe ratio, annualized
Baseline sample	-19%	1.48
TIP removed from the investment set	none	1.47
TIP's monthly mean raised by 0.12%	+126%	1.57

17. † **[Select All That Apply]** Which conclusions does Exhibit 6 support?

- ✗ **A.** Because TIPS' standalone Sharpe ratio is below the tangency portfolio's, the optimizer shorts them.
- ✓ **B.** Removing TIPS barely changes the best attainable Sharpe ratio, so TIPS add little to the estimated opportunity set.
- ✓ **C.** A change in the TIPS mean of about one standard error moves the TIPS weight from a short to a large long position.
- ✗ **D.** TIPS are nearly uncorrelated with the other assets, which is why removing them changes the Sharpe ratio so little.
- ✓ **E.** TIP is shorted in the baseline even though its sample mean excess return is positive.

B: 1.48 against 1.47. C: 0.12% is about 1.1 standard errors and flips the position. E: TIP's mean excess return is 1.3%, yet the optimizer shorts it; the short reflects its covariances with the assets held long, not a negative expected return. A: every asset's standalone Sharpe ratio is below the tangency ratio; that says nothing about the sign of its weight. D is backwards: TIPS add little because they are highly correlated with IEF (0.76), so the frontier already contains most of what they offer.

18. In Exhibit 6 the TIP weight moves from -19% to +126% while the tangency Sharpe ratio moves only from 1.48 to 1.57. Which statement explains the contrast?

✗ **A.** The Sharpe ratio is insensitive to the mean inputs, so only the weights carry the information that the data hold about TIPS.

✓ **B.** The inverse covariance matrix turns a small change in one mean into large offsetting positions among correlated assets, while the attainable Sharpe ratio moves little.

✗ **C.** The weight change is an artifact of normalizing the weights to sum to one, which magnifies small differences between the cases.

✗ **D.** The weight swing means the covariance estimate must have changed between the cases, since changing the means alone cannot move weights that much.

✗ **E.** The optimizer treats the 0.12% as certain, while the Sharpe ratio discounts it for the sampling error in TIP's mean.

B: the weights solve a system with the inverse covariance matrix, which amplifies small changes in the mean vector into large offsetting positions when assets are highly correlated (TIP correlates 0.76 with IEF, the kind of overlap that makes the amplification large), while the attainable Sharpe ratio is a smooth function of the inputs and moves only from 1.48 to 1.57. The frontier barely moves while the recipe flips; the exhibit illustrates the sensitivity rather than isolating one pair as its sole cause. A: the Sharpe ratio does respond to the mean, by 0.09 here. C: both normalization sums are positive here, and TIP's unnormalized position itself changes sign, so normalization does not create the swing. D: only the TIP mean changed; the covariance matrix was held fixed. E: both numbers come from the same point estimates; nothing in either calculation discounts for the standard error of TIP's mean.

19. An analyst estimates the inputs of a mean-variance optimization from 50 years of monthly index returns; the annual volatility is about 16%. Which statement about the precision of the inputs is correct?

✗ **A.** The volatility is the hard input: with 600 monthly observations its standard error is larger than that of the mean return.

✗ **B.** Both inputs are estimated to within a fraction of a percent per year, because 600 monthly observations is a large sample by the usual standards.

✗ **C.** The expected return is estimated well because the sample is long, so disagreement among investors about it reflects views on the future rather than on the data.

✓ **D.** The expected return is the hard input: its standard error is about 16% over the square root of 50 years, roughly 2 points per year, large relative to plausible premia.

✗ **E.** Neither input is usefully estimated from history, because price levels are not stationary and returns inherit that problem.

D: the standard error of an average annual return falls with the square root of the number of years, 16% over about 7, so even half a century leaves a two-point band around a premium that may itself be five or six points. The volatility is estimated far more precisely from the same history. A reverses the two. B and C overstate the precision of the mean. E: returns are close enough to stationary for the volatility to be estimated well; the problem is specific to the mean.

20. Mean-variance investors share the same estimates of means and covariances and can borrow and lend at the risk-free rate, but differ in risk aversion. How do their portfolios differ?

- ✓ **A.** They hold the same risky portfolio in different amounts, combining it with more or less of the risk-free asset.
- ✗ **B.** They hold different risky portfolios: the more risk-averse tilt toward lower-volatility assets within their risky holdings.
- ✗ **C.** They hold the same total amount in risky assets but choose different points along the risky-only frontier.
- ✗ **D.** They differ only in leverage; a more risk-averse investor shorts the tangency portfolio to reduce risk.
- ✗ **E.** They cannot be compared, because the tangency portfolio is defined only for a given level of risk aversion.

A is two-fund separation: with a risk-free asset the tangency portfolio is the same for everyone who shares the inputs, and risk aversion sets only the dose. B and C describe choices along the risky-only frontier, which the risk-free asset makes dominated. D: a risk-averse investor holds less of the tangency portfolio and more cash, not a short position. E: the tangency portfolio depends on the inputs, not on risk aversion.

Exhibit 7: Unconstrained stock allocation Selected rows from the assigned 10-stock-plus-SPY optimization, using weekly returns from July 10, 2015–June 26, 2026. As permitted in that exercise, the risk-free rate is set to zero. The full weight vector sums to 100%; its sum of absolute weights is 1,007.5%. Ratios are annualized using 52 weeks. Values below are rounded; omitted stocks account for the remaining positions. Correlation with SPY: BRK/B 0.70, NVDA 0.64, LLY 0.34. BRK/B's correlations with the other nine stocks run from 0.20 to 0.41. BRK/B has the lowest annualized mean of the ten stocks (13.5%); its ratio of 0.716 is the second lowest after META (0.649).

Asset	Optimized weight	Mean/volatility
BRK/B	169.6%	0.716
LLY	90.5%	1.011
NVDA	84.6%	1.420
SPY	-453.8%	0.859

21. Use Exhibit 7's SPY weight at the beginning of a hypothetical week in which SPY gains 1%. What is the SPY position's contribution to the full portfolio's return, expressed relative to the investor's beginning capital? Ignore financing and trading costs.

- ✗ **A.** Approximately +4.54 percentage points.
- ✗ **B.** Approximately -0.0454 percentage point.
- ✗ **C.** Approximately -0.45 percentage point.
- ✗ **D.** Approximately -1.00 percentage point.
- ✓ **E.** Approximately -4.54 percentage points.

E: the SPY weight is -4.538 per unit of capital, so its contribution is -4.538 times 1%. A reverses the sign of the short position. B divides the weight by 100 twice. C divides the SPY weight by the gross exposure of 1,007.5%, stating the contribution relative to gross exposure rather than to capital. D assumes only a 100% short. The other positions' contributions are needed for the full portfolio return.

22. † **[Select All That Apply]** Which readings of Exhibit 7 are correct?

- ✗ **A.** A 100% net weight means the portfolio uses no leverage.
- ✓ **B.** The gross exposure is about ten times the investor's capital despite a 100% net allocation.
- ✗ **C.** NVDA has the largest weight of the ten stocks because it has the largest ratio.
- ✓ **D.** The BRK/B and NVDA comparison shows that standalone ratios alone do not determine optimized weights.
- ✗ **E.** SPY is shorted because it has the lowest mean/volatility ratio of the eleven assets.

B uses the sum of absolute weights. D compares a larger BRK/B weight with a lower stand-alone ratio. A confuses net and gross. C is contradicted by BRK/B and LLY. E: SPY's 0.859 is above BRK/B's 0.716 and META's 0.649; the short reflects SPY's covariances with the stocks, not a low ratio.

23. Why does the optimizer in Exhibit 7 hold about 170% in BRK/B, the stock with the second-lowest mean/volatility ratio of the ten?

- ✗ **A.** BRK/B has the highest mean return among the lower-volatility stocks in the set, which the optimizer rewards.
- ✓ **B.** BRK/B's relatively high correlation with SPY helps offset the market short, while its weaker correlations with the other stocks add diversification.
- ✗ **C.** The optimizer weights stocks by inverse variance, and BRK/B has the lowest variance of the ten stocks.
- ✗ **D.** BRK/B's ratio is understated because its returns are total rather than excess returns, unlike the others.
- ✗ **E.** With SPY shorted about 454%, BRK/B's weight is the residual that brings the net allocation back to 100%.

B is the mechanism in the assigned solution: correlation 0.70 with SPY and 0.20 to 0.41 with the other stocks. The full optimizer uses every mean and covariance, so these correlations support the mechanism rather than reproduce the 170% weight. A: BRK/B's mean (13.5%) is the lowest of the ten. C describes a different rule. D applies to every series equally. E: the weight vector is proportional to the inverse covariance times the means; the 100% sum sets a common scale for every weight and does not make any one stock the plug.

24. † **[Select All That Apply]** A portfolio's unconstrained optimum shorts domestic equities. A long-only optimization on the same inputs instead assigns domestic equities zero. Which conclusions are correct?

- ✓ **A.** The zero weight is consistent with a binding lower bound.
- ✗ **B.** The equity asset has a negative expected excess return.
- ✗ **C.** Tightening allocation bounds leaves the other assets' weights unchanged when their expected returns are unchanged.
- ✗ **D.** The equity asset has zero covariance with the remaining assets, which is why its constrained weight is zero.
- ✓ **E.** On the same inputs, the long-only portfolio's in-sample Sharpe ratio is at most the unconstrained optimum's.

A follows from the imposed no-short boundary. E: a bound can only shrink the feasible set, so the constrained in-sample optimum is no better than the unconstrained one. B ignores relative returns and covariance. C: the other weights are re-optimized when the feasible holdings change. D mistakes zero holdings for zero co-movement.

Decomposition, hedging, and replication

Exhibit 8: A one-factor decomposition Hypothetical OLS regression with an intercept: $r_F^e = \alpha + \beta r_M^e + \epsilon$. All moments use the same observations. Returns are excess returns. Means and volatilities below are annualized consistently; the sample residual is orthogonal to the market.

Quantity	Value
Fund mean / volatility	14% / 20%
Market mean / volatility	6% / 8%
Market beta, β	2.0
Centered R^2	0.64

25. † **[Select All That Apply]** Which statements follow from Exhibit 8's in-sample OLS decomposition?

- ✓ **A.** The fund's alpha is 2% per year.
- ✓ **B.** Market exposure accounts for 64% of the fund's variance, not 64% of its standard deviation.
- ✓ **C.** A beta-neutral fund-minus-market strategy retains residual risk.
- ✗ **D.** The alpha can be earned by itself, without holding the residual risk.
- ✗ **E.** The fund's information ratio is 2% over 20%, or 0.10.

A: 14%-2(6%)=2%. B is the centered R-squared variance identity. C: residual variance is 36% of total variance. D mistakes a regression intercept for an investable security. E divides alpha by total volatility; the information ratio uses residual volatility, 2% over 12%, about 0.17.

26. An investor who already holds the market portfolio considers adding the fund in Exhibit 8. Treat the exhibit's sample moments as expected inputs, and allow the market and risk-free positions to be adjusted freely. Which statement is correct?

- ✗ **A.** The fund cannot improve on holding the market alone, because its own Sharpe ratio (0.70) is below the market's Sharpe ratio (0.75).
- ✗ **B.** Hedging out the fund's market exposure leaves a position whose Sharpe ratio equals the fund's own Sharpe ratio of 0.70, since hedging removes only risk.
- ✗ **C.** The fund improves the portfolio only if its R-squared on the market is below one half, so that more than half of its risk is residual.
- ✓ **D.** A small long position in the fund, with the market position adjusted, raises the attainable Sharpe ratio, because the fund's alpha on the market is positive.
- ✗ **E.** The fund improves the portfolio only if its information ratio exceeds the market's Sharpe ratio of 0.75, which it does not.

D: a positive alpha means the fund's expected excess return exceeds what its market exposure delivers, so a small addition improves on the market alone even though the fund's own Sharpe ratio (14/20 = 0.70) is below the market's (6/8 = 0.75). A compares stand-alone ratios; the relevant quantity for an addition is the alpha, or the information ratio of 2% over 12%. B: the hedged position holds alpha plus residual, so its Sharpe ratio is the information ratio, about 0.17. C: the sign of alpha, not the R-squared, decides whether the addition helps. E: any positive information ratio helps; the best combination's squared Sharpe ratio is the market's squared Sharpe ratio plus the fund's squared information ratio, so the information ratio need not exceed 0.75.

Exhibit 9: ProShares case: the hedge-fund series Monthly returns from the assigned ProShares workbook, August 2011–August 2025: the 169 months in which every series is observed (HFRI's September 2025 value is missing). Returns are used as given in the assignment, not converted to excess returns. Means, volatilities, and alphas are annualized; the Sharpe ratio is mean divided by volatility. Each series is regressed on SPY with an intercept. MLFM-ES is the tradable Merrill Lynch factor index that HDG tracks; the correlation between HDG and MLFM-ES is 0.99.

Series	Mean	Volatility	Sharpe	Beta on SPY	Alpha on SPY	R^2	Corr. with HFRI
HFRI	5.1%	5.9%	0.87	0.35	0.2%	0.72	1.00
MLFM-ES	3.6%	5.5%	0.65	0.34	-1.3%	0.79	0.90
HDG	2.6%	5.7%	0.45	0.35	-2.4%	0.77	0.89
QAI	2.8%	5.0%	0.56	0.30	-1.5%	0.75	0.86
SPY	14.3%	14.4%	1.00	1.00	0.0%	1.00	0.85

27. Every hedge-fund series in Exhibit 9 has a SPY beta near 0.35 and an R-squared near 0.75. HFRI's Sharpe ratio is 0.87 and HDG's is 0.45. Which statement best accounts for the gap?

✗ **A.** HDG carries more market risk per dollar invested than HFRI does, which lowers its ratio.

✓ **B.** HDG's alpha on SPY is negative while HFRI's is near zero; the gap is in mean return, not systematic risk.

✗ **C.** HDG's volatility is much higher than HFRI's, which lowers its ratio for the same mean.

✗ **D.** HDG is weakly correlated with HFRI, so it tracks a different set of risks with a lower premium.

✗ **E.** HDG's higher R-squared (0.77 against 0.72) leaves less residual return, the source of HFRI's extra mean.

B: both series have beta 0.35; the gap appears as HDG's negative alpha, 2.4% per year below what its SPY exposure delivers, while HFRI's alpha is near zero. The table does not allocate that gap among fees and other causes. A reads a lower Sharpe ratio as more risk. C: the volatilities are 5.7% and 5.9%. D: the correlation is 0.89. E treats residual variance as a source of mean return; with an intercept the residual has zero mean, and the 0.05 difference in R-squared says nothing about the 2.5-point gap in means.

Exhibit 10: ProShares case: replicating HFRI Regression of HFRI on the six Merrill factors over the same 169 months as Exhibit 9, with and without an intercept. Coefficients are the fitted positions per dollar of HFRI; factor positions exclude the fitted intercept. For this exercise, treat the supplied T-bill proxy as an investable cash-return series. The intercept is annualized. Tracking error is the annualized volatility of the residual.

Quantity	With intercept	Without intercept
Coefficient on SPY	0.04	0.06
Coefficient on the 3-month T-bill factor	0.32	0.65
Coefficient on EEM	0.09	0.08
Coefficient on EFA	0.07	0.07
Coefficient on EUO	0.03	0.03
Coefficient on IWM	0.15	0.14
Sum of absolute coefficients	0.70	1.03
Intercept, annualized	1.4%	none
R^2	0.84	—
Mean of the fitted series, annualized	5.1%	4.4%
Correlation of the fitted series with HFRI	0.92	0.92
Tracking error, annualized	2.3%	2.3%

28. † **[Select All That Apply]** Which statements about the replication in Exhibit 10 are correct?

- ✓ **A.** The with-intercept factor positions have a gross size of about 0.70 per dollar of HFRI, so they can be held without leverage or large short positions.
- ✓ **B.** Removing the intercept leaves the fitted series' correlation with HFRI at about 0.92 at the printed precision but lowers its mean below HFRI's mean.
- ✓ **C.** With an intercept, the fitted series (intercept included) has the same sample mean as HFRI by construction.
- ✗ **D.** A tracking error of 2.3% per year bounds the replication's annual return within 2.3 percentage points of HFRI's.
- ✓ **E.** The factor positions alone, excluding the intercept, earned about 3.7% per year over the sample.

ABCE and ABC both receive full credit. E is optional because Exhibit 10 shows two regressions and E, unlike A, B, and C, does not say which one it describes: the with-intercept factor positions earned 5.1% less the 1.4% intercept, or 3.7%, while the no-intercept positions are the whole fit and earned 4.4%. Both Week 3 lectures and the Winter practice solutions presented the no-intercept fit as the tradable positions, so reading E against that column and omitting it is not penalized. A, B, and C must still be selected and D must be omitted. The gray check marks optional choice E.

A answers the assignment's question about realistic position sizes. B is the §3.1 finding: correlation 0.92 either way, fitted mean 4.4% against 5.1%. C: least squares with an intercept gives zero-mean residuals. E: in the with-intercept regression, the fitted mean of 5.1% includes the 1.4% intercept, so the factor positions account for 3.7%; without an intercept the factor positions are the whole fit and earned 4.4%. D: tracking error is the standard deviation of the residual, not a bound.

29. When the intercept is removed in Exhibit 10, the T-bill coefficient rises from 0.32 to 0.65 while the other five coefficients barely move. Why?

- ✗ **A.** The T-bill factor is the regressor most highly correlated with HFRI, so it gains weight.
- ✗ **B.** Without an intercept, least squares forces the coefficients to sum to one, and the T-bill fills the gap.
- ✗ **C.** Removing the intercept raises the fitted mean, and a larger T-bill position offsets the increase.
- ✓ **D.** The T-bill return has very low volatility relative to the other factors, so it absorbs the level the intercept had carried.
- ✗ **E.** The T-bill factor is the most collinear of the six regressors, so its coefficient is the least stable across specifications.

D: without a constant, the low-volatility T-bill regressor absorbs part of the target's mean level, so its coefficient rises; the fitted mean still falls from 5.1% to 4.4%. An intercept is a return term, not a cash weight; what a low-volatility bill series can imitate is the return on a cash position. A: the T-bill series has almost no correlation with HFRI's monthly variation. B: the sum of 1.03 is a coincidence, not a constraint. C: the fitted mean falls. E: the T-bill series is nearly constant, so it is close to collinear with the intercept, not with the other five factors; that is the mechanism in D, not instability from an ill-conditioned set of factors.

30. A replicator must decide whether to include an intercept when fitting tradable factor positions to a target's returns. When is it better to leave the intercept out?

- ✓ **A.** When the target's in-sample mean is trusted to persist, so the positions should match the level of returns as well as the co-movement.
- ✗ **B.** When the target's in-sample mean is not trusted to persist, because dropping the intercept stops the positions from chasing that mean.
- ✗ **C.** If the R-squared of the regression with an intercept is below 0.9, because the fit then needs the extra freedom.
- ✗ **D.** If the regression's residuals have a nonzero sample mean, because an intercept is then needed to absorb it.
- ✓ **E.** When the replication must be tradable, since an intercept is not a position that can be held.

Either A alone or E alone receives full credit. A is the intended answer. E is also accepted because the initial Week 3 tradability argument could reasonably have led to that choice. The gray check marks the accepted alternative.

A: including an intercept lets the betas fit co-movement while the intercept absorbs the mean difference. Omitting it makes the positions also try to match the target's sample mean, which is useful when that mean is trusted to persist. B reverses this mechanism. C: the R-squared cutoff does not decide the specification. D: estimating an intercept makes the sample residual mean zero. E is the initial tradability argument discussed in Week 3. The further distinction is that the beta positions can be held after either regression; estimating an intercept does not require investing in it. Keeping the intercept can therefore be appropriate when the sample mean is too uncertain to guide the positions.

Exhibit 11: Decomposing Apollo returns Weekly excess returns, January 9, 2015–October 3, 2025, from the assigned private-equity data. Both OLS regressions include an intercept and use the same 561 observations. SPY, IEF, and HYG are traded ETF returns. Alpha and residual volatility are annualized using 52 weeks. This is a descriptive decomposition.

Model	Alpha	SPY beta	IEF beta	HYG beta	R^2	Residual vol.
SPY only	4.9%	1.480	—	—	0.466	26.9%
Three ETFs	10.0%	1.066	-0.704	0.931	0.497	26.1%

31. † **[Select All That Apply]** Which interpretations of Exhibit 11 are correct?

- ✓ **A.** The three-ETF regression explains somewhat more of APO's weekly variation in this sample.
- ✗ **B.** Adding regressors moves the fitted intercept toward zero.
- ✓ **C.** The SPY coefficient changes because it is now a partial exposure, holding the other included factors fixed.
- ✗ **D.** The larger intercept means APO's manager added more value in the three-factor sample.
- ✓ **E.** The SPY-only beta of 1.480 is APO's correlation with SPY scaled by APO's volatility over SPY's.

A compares 0.466 and 0.497 on the same observations. C describes the change from a one-variable to a partial coefficient. E: with one regressor the slope is the correlation times the ratio of volatilities, which is how 1.480 exceeds one with a correlation near 0.68. B is contradicted here and is not an OLS property. D: both regressions describe the same return history.

32. † **[Select All That Apply]** An investor holds \$1 of APO and wants to remove the fitted three-ETF exposures in Exhibit 11. Adjust the risk-free position to keep net investment at \$1. Ignore trading costs and estimation error. Which statements are correct?

- ✗ **A.** The IEF hedge is a short position because APO's estimated IEF coefficient is negative.
- ✓ **B.** The hedge shorts SPY and HYG and goes long IEF, in the displayed beta magnitudes.
- ✗ **C.** The hedge eliminates the residual variation as well as the fitted ETF component.
- ✓ **D.** The historical hedged excess mean per original \$1 of APO is approximately the annual alpha.
- ✗ **E.** Each hedge ratio should come from a separate one-factor regression, such as 1.480 for SPY.

B reverses every fitted exposure: short 1.066 SPY, long 0.704 IEF, short 0.931 HYG. D follows from zero sample-mean residuals with an intercept. A has the hedge sign wrong. C confuses systematic exposure removal with total risk removal. E: a joint hedge uses the multiple-regression betas; one-factor slopes such as 1.480 ignore the co-movement among the three ETFs.

Exhibit 12: A rolling replication record Hypothetical two-factor version of the ProShares exercise. At the end of month 60, a regression with an intercept is fitted using months 1–60. Its coefficients are held fixed for month 61. The displayed month-61 returns become known only at the end of month 61. Both factors and the target are excess returns; ignore trading costs.

Quantity	Value
Estimated monthly intercept	0.10%
Estimated beta on factor 1 / factor 2	0.40 / 0.60
Realized month-61 factor 1 / factor 2 return	5.00% / -1.00%
Realized month-61 target return	1.50%

33. Does using the realized month-61 factor returns in Exhibit 12 make this an invalid out-of-sample replication evaluation?

✗ **A.** Yes. An out-of-sample test must use the factors' prior sample means in place of their realized returns.

✗ **B.** Yes. Using any outcome that was unknown when the positions were chosen makes the evaluation in-sample.

✓ **C.** No. The coefficients were fixed before month 61; the realized factor returns are what those positions earned.

✗ **D.** No. The month-61 target return may enter the estimation once the procedure is labeled a rolling model.

✗ **E.** No. Any evaluation that uses realized returns is out of sample, whether or not the coefficients were fixed before the month began.

C separates choosing positions from evaluating their subsequent returns. A substitutes a return forecast for a replicating portfolio. B would prevent any realized backtest. D uses the target outcome to estimate its own weights, contrary to the setup. E: realized returns do not make a test out of sample; what does is that the positions were fixed before the month began, so a fit that used month 61 in estimation would be in-sample even though it uses realized returns.

34. In Exhibit 12, the factor positions earn 1.40% in month 61 and the target earns 1.50%. An analyst adds the estimated 0.10% intercept to the positions' return and reports that the replication matched the target exactly. Which assessment is correct?

- ✗ **A.** Correct: because the factor weights sum to one, a cash position finances the intercept and delivers it.
- ✗ **B.** Correct: fixing the coefficients before month 61 makes the intercept an earned return for that month, like the factor positions.
- ✓ **C.** Incorrect: the fitted intercept is not earned by any traded position, so the positions trail the target by 0.10 percentage points.
- ✗ **D.** Incorrect: a positive intercept means the positions beat the target on average, so the month-61 shortfall is noise rather than a gap.
- ✗ **E.** Correct: with an intercept the fitted series matches the target's mean by construction, so adding it recovers the target each month.

C: the intercept is a fitted constant, not a payoff; only the factor positions earn returns, and they trail the target by 0.10 points in month 61. A treats the intercept as a funded cash position; the weights summing to one is incidental. B confuses the timing that makes the evaluation out of sample with investability. D has the sign backwards: a positive intercept means the target exceeded the fitted positions on average. E: the mean identity holds on average over the estimation window, not month by month, and even on average the intercept is a fitted gap, not a return the positions collect.

35. † **[Select All That Apply]** A factor covariance model uses $\Sigma_r = B\Sigma_f B^\top + D$, where D is a diagonal matrix of positive residual variances. Which statements are correct?

- ✓ **A.** The model removes cross-asset residual covariances but retains each asset's residual variance.
- ✗ **B.** The model assumes the total covariance matrix of asset returns is diagonal.
- ✓ **C.** The common factors can produce cross-asset covariance even when D is diagonal.
- ✓ **D.** A diagonal residual covariance matrix makes the residuals statistically independent of one another.
- ✓ **E.** With the factors fixed, the parameter count grows with the number of assets, not with its square.

ACE and ACDE both receive full credit. A diagonal residual covariance matrix does not by itself imply statistical independence. The Week 3 lecture explicitly set aside that distinction for course purposes, so selecting D is not penalized. A, C, and E must still be selected and B must be omitted. The gray check marks optional choice D.

A, C, and E are correct. A and C describe the two covariance components: the factors supply all cross-asset covariance, and D supplies each asset's own residual variance. E: B has one row per asset, D one entry per asset, and the factor covariance does not grow with the number of assets, whereas an unrestricted covariance matrix has an entry for every pair. B ignores shared factor risk. D confuses zero linear covariance with independence; joint normality or another sufficient assumption would be needed.

36. Suppose 100 assets each have beta one to the same factor. Their residuals have equal, positive, finite variance, are mutually uncorrelated, and are uncorrelated with the factor. As equally weighted holdings expand to more such assets, what happens to portfolio risk?

- ✗ **A.** Both factor and residual variance fall in proportion to the number of assets.
- ✗ **B.** Factor variance falls, while residual variance remains unchanged.
- ✗ **C.** Neither component changes, since every asset has the same beta.
- ✓ **D.** Residual variance declines, while the common factor component remains.
- ✗ **E.** Neither falls, since diversifying residuals needs negative correlation.

D: the variance of the average residual is the per-asset residual variance divided by N, whereas the portfolio factor beta stays one. A diversifies the common shock incorrectly. B reverses the components. C overlooks the averaging of mutually uncorrelated residuals. E: any correlation below one diversifies; zero correlation among the residuals is enough, and negative correlation is not required.

37. An analyst proposes a factor covariance model with a diagonal residual covariance matrix, $\Sigma_r = B\Sigma_f B^\top + D$, and treats the high R^2 of each asset's factor regression as sufficient validation of the diagonal restriction. Which assessment is correct?

- ✗ **A.** The restriction holds by construction: least squares makes each asset's residuals uncorrelated with the factors, so nothing common remains in them.
- ✗ **B.** High R-squared for every asset implies that the residuals are nearly uncorrelated across assets, because little variance is left for them to share.
- ✗ **C.** The restriction can be checked by confirming that each asset's residual variance is small relative to its total variance, which is the same information as R-squared.
- ✗ **D.** If the factors are uncorrelated with one another, the residuals must be uncorrelated across assets as well, so the restriction holds automatically.
- ✓ **E.** High R-squared is not sufficient: the restriction concerns the correlations among the residuals across assets, which must be examined directly.

E: R-squared is a per-asset statement about explained variance; the diagonal restriction is a cross-asset statement about residual correlations, which can be large when a common influence is missing from the factors. A: least squares makes each asset's residual orthogonal to the included factors, not to the other assets' residuals; an omitted common influence sits in all of them. B: residuals that are small can still move together. C restates R-squared. D: the correlation structure of the factors says nothing about what the residuals share.

Tail risk and dynamic risk

Exhibit 13: Risk of a four-stock portfolio Weekly returns, July 10, 2015–June 26, 2026, from the assigned exercise. The equal-weight portfolio holds 25% each of AAPL, META, NVDA, and TSLA. The cash-replacement portfolio holds 25% each of AAPL, META, NVDA, and zero-return cash. Volatility is annualized using 52 weeks. VaR is the empirical fifth percentile of weekly returns; CVaR is the mean at or below that percentile. For the equal-weight portfolio, a normal distribution with its sample weekly mean (0.75%) and weekly volatility (4.33%) implies a 5% VaR of -6.37% and a 5% CVaR of -8.17%.

Investment	Annual vol.	Weekly 5% VaR	Weekly 5% CVaR
AAPL	27.8%	-5.49%	-8.30%
META	35.7%	-7.23%	-10.64%
NVDA	45.4%	-8.58%	-11.53%
TSLA	58.0%	-11.70%	-14.78%
Equal weight	31.2%	-6.07%	-8.42%
25% cash replacement	21.7%	-4.23%	-6.00%

38. A second analyst removes TSLA from Exhibit 13's equal-weight portfolio and spreads its 25% equally over AAPL, META, and NVDA, leaving no cash. How does that portfolio's annualized volatility compare with the 21.7% of the cash-replacement portfolio?

✗ **A.** It equals 21.7%, because the same three stocks are held in the same proportions to one another.

✗ **B.** It equals the average of the three stocks' volatilities, 36.3%, because TSLA no longer diversifies them.

✓ **C.** It is 4/3 of 21.7%, because every risky weight is scaled up by 4/3 and cash carries no risk.

✗ **D.** It cannot be ranked against 21.7% without the pairwise correlations among the three remaining stocks, which the exhibit omits.

✗ **E.** It is below 21.7%, because the 25% now adds diversification across three stocks instead of sitting idle in cash.

C: the three risky weights rise from 25% to one third each, a common factor of 4/3, and scaling every risky weight by a common factor scales volatility by that factor, to about 28.9%. A ignores that the risky weights rise. B ignores the diversification among the three stocks, which the 21.7% already reflects. D: those correlations are already embedded in the 21.7%. E: cash contributes no variance and no covariance, so moving its 25% into the three stocks adds risk; spreading a weight across more risky holdings is not diversification when it replaces a riskless one.

39. † **[Select All That Apply]** Which statements about Exhibit 13 are correct?

- ✓ **A.** The equal-weight portfolio's volatility (31.2%) is below the average of the four stocks' volatilities (41.7%); for equal weights this holds whenever the correlations are below one.
- ✗ **B.** The equal-weight portfolio's VaR (-6.07%) is smaller in magnitude than the average of the four stocks' VaRs (-8.25%); VaR is sub-additive, so this holds for any portfolio of the four.
- ✓ **C.** Replacing TSLA with cash lowers volatility by 9.5 percentage points, less than one quarter of TSLA's standalone volatility.
- ✗ **D.** TSLA contributed one quarter of its standalone VaR, about 2.9 percentage points, to the equal-weight portfolio's VaR.
- ✗ **E.** The equal-weight portfolio's CVaR (-8.42%) is smaller in magnitude than the average of the four stocks' CVaRs (-11.31%), but CVaR, like VaR, can violate sub-additivity.

A is sub-additivity of volatility. C: one quarter of 58.0% is 14.5%; the actual reduction is 31.2-21.7. B: the inequality holds here, but VaR is not sub-additive in general; CVaR is. D: a position's contribution to portfolio VaR is not its weight times its standalone VaR; the joint distribution of the four returns determines it. E: the inequality holds here and the second clause is false: CVaR is sub-additive, so for nonnegative weights summing to one a portfolio's CVaR is never worse than the same weighted average of its components' CVaRs; VaR carries no such guarantee.

40. For Exhibit 13's equal-weight portfolio, the note gives the 5% VaR and CVaR implied by a normal distribution with the portfolio's sample mean and volatility, -6.37% and -8.17%, against the empirical -6.07% and -8.42%. Which conclusion does the comparison support?

✗ **A.** The normal model understates risk throughout the lower tail, because weekly stock returns are fat-tailed at every quantile, including the 5% cutoff shown here.

✓ **B.** At the 5% cutoff the normal model is the more extreme, and beyond it only slightly less so: at this level the normal translation of volatility serves this portfolio adequately.

✗ **C.** The normal model's more extreme cutoff shows that the sample mean, rather than the shape of the return distribution, drives the differences between the two sets.

✗ **D.** The sample must be negatively skewed, because its CVaR lies further beyond its VaR than the normal model's CVaR lies beyond its VaR.

✗ **E.** The close agreement at the 5% cutoff shows the portfolio's weekly returns are close to normal, so the same model can be trusted at the 1% level as well.

Every response receives full credit. The keyed option asks students to endorse a judgment, that the normal translation serves the portfolio adequately, which the question does not define: no tolerance or purpose is given, so a student who agreed with B's numerical clauses could still decline to endorse it and had no correct option left. B was the intended comparative reading. D does not follow: a larger gap between CVaR and VaR than the normal model's does not determine the sign of skewness.

B: the two sets differ by about 0.3 and 0.2 percentage points in opposite directions; at the 5% level a normal model attached to the sample volatility describes this portfolio about as well as the empirical quantile does, which is the lecture's point that the normal translation does well enough while its failures sit deeper in the tail and at higher frequency. "Adequately" is a judgment and the question supplies no tolerance or purpose, so every response receives full credit; B was the intended comparative reading. Two lessons stand: close VaR and CVaR estimates at the 5% level do not establish fitness for a purpose, and a larger gap between CVaR and VaR does not establish negative skewness. A: at the 5% cutoff the normal model is the more extreme, so it does not understate risk at every quantile. C: the mean shifts both normal numbers by the same amount and cannot explain differences that run in opposite directions. D: the gap between CVaR and VaR measures the average severity of returns beyond the lower-tail cutoff. It does not determine the sign of overall skewness. The sample skewness here is about zero (+0.03). E: agreement at the 5% quantile does not establish normality; the normal shape fails deeper in the tail, where the 1% quantile sits, and at higher frequency, which is where the lecture located its failures.

Exhibit 14: Backtest of the equal-weight portfolio's 5% VaR The equal-weight portfolio of Exhibit 13, July 10, 2015–June 26, 2026. Each week's 5% normal VaR threshold is $-1.65\sigma_t$, with σ_t estimated from returns through week $t-1$. An exception is a week whose realized return falls below that week's threshold. The final-week columns give each estimator's weekly volatility for the last week of the sample and the threshold it implies. The full-sample empirical 5% VaR of the same portfolio is -6.07%, and its full-sample weekly volatility is 4.3% (31% annualized). Volatility is the sample standard deviation of past returns, using at least two prior observations for the expanding window and the prior 26 weeks for the rolling window; the conditional mean is taken as zero.

Estimator	Weeks	Exceptions	Hit rate	Final-week vol.	Final-week 5% VaR
Expanding window	571	26	4.6%	4.3%	-7.1%
26-week rolling window	547	22	4.0%	3.2%	-5.3%

41. † **[Select All That Apply]** Which statements about the backtest in Exhibit 14 are correct?

- ✗ **A.** The expanding window's hit rate is closer to 5%, so it responds better to changes in volatility over time.
- ✓ **B.** Both dynamic models produced fewer exceptions than the nominal 5% rate.
- ✗ **C.** A hit rate below 5% means the model's threshold was too conservative in every week of the sample.
- ✓ **D.** The rolling model's final-week VaR (-5.3%) is less negative than the expanding model's (-7.1%) because the last 26 weeks were calmer than the sample as a whole.
- ✓ **E.** The hit rates alone do not show whether either model's exceptions arrived in runs, which is a separate failure to test for.

B: 4.6% and 4.0%. D: both thresholds are -1.65 times a volatility estimate; the rolling estimate uses only the last 26 weeks (3.2% per week) while the expanding estimate still reflects the whole sample (4.3%). E: a hit rate is an average over the sample; a model whose exceptions bunch in turbulent runs can have the right average and still fail, and the exhibit does not report when the exceptions occurred. A: frequency of exceptions and responsiveness are different properties; the expanding window is the least responsive estimator. C: an average hit rate says nothing about individual weeks.

Exhibit 15: Three forecasts for the same week Hypothetical forecasts of the same portfolio, all formed using information available through week $t-1$. Assume zero conditional mean and conditional normal returns, so that for a weekly volatility σ_t the 5% return-threshold VaR is $-1.65\sigma_t$. The realized return in week t is -5.50%.

Estimator	Forecast weekly vol.	5% VaR
Expanding window	2.00%	-3.30%
26-week rolling window	4.00%	-6.60%
EWMA	3.00%	-4.95%

42. † **[Select All That Apply]** Which statements about the forecasts in Exhibit 15 and the realized -5.50% return are correct?

- ✓ **A.** The expanding-window and EWMA forecasts each record an exception this week; the rolling forecast does not.
- ✗ **B.** The rolling model's clean week shows that it was the best-calibrated of the three forecasts.
- ✓ **C.** A single week's exception is consistent with a correctly calibrated 5% model, which should be exceeded about one week in twenty.
- ✗ **D.** The forecasts cannot be compared on the same week, because each estimator was formed from a different information set.
- ✗ **E.** A fair test would re-estimate each forecast including the -5.50% return itself, since that is the latest available information.

A: -5.50% is below -3.30% and below -4.95% but not below -6.60%. C: a single exception is possible under a calibrated 5% model and does not by itself establish miscalibration. B: one week cannot rank calibration; a backtest over many weeks can. D: all three use the same history through week t-1 and differ only in how they weight it; the expanding window gives every past week equal weight, the rolling window equal weight to the last 26 weeks and none to earlier ones, and EWMA declining weights. E: an exception is judged against the threshold set from information through week t-1; folding the realized return into the forecast it is meant to test is the backtesting error the lecture warned against.

43. An EWMA variance update is $\sigma^2_{t+1} = 0.9\sigma^2_t + 0.1r^2_t$, with zero mean assumed. Yesterday's variance estimate was 0.0004 and today's return was -10%. Which statement about the updated estimate is correct?

✗ **A.** It falls, because the factor 0.9 decays the previous estimate whatever the return.

✗ **B.** It is unchanged, because the assumed mean of zero makes a -10% return uninformative about variance.

✓ **C.** It rises by exactly as much as it would after a +10% return, because the return enters through its square.

✗ **D.** It rises by more than it would after a +10% return, because the update gives losses extra weight relative to gains of the same size.

✗ **E.** It rises by less than a 26-week rolling estimate would, because the update gives today's return only a 10% weight.

C: the update adds 0.1 times the squared return, 0.001, which is the same for -10% and +10%, and removes 0.1 times the old estimate, 0.00004; the estimate rises from 0.0004 to 0.00136. A: the decay is outweighed by the shock whenever the squared return exceeds the old estimate. B: with a zero mean the squared return is the variance signal. D: this update is symmetric; an asymmetric response to losses needs a different model. E: a 10% weight on the latest squared return is more than the 1/26 a rolling window gives it, so EWMA responds more, not less, to the shock.

44. In a backtest on weekly returns with fat tails and volatility clustering, a 5% VaR set at -1.65 times a 26-week rolling volatility produced a hit rate closer to 5%, with exceptions spread more evenly in time, than the empirical 5% quantile of all past returns. Which statement explains how the normal-based estimator can beat the distribution-free one here?

✗ **A.** The rolling window exploits the serial correlation of returns: last week's loss predicts this week's, so its threshold moves ahead of the loss while the all-history quantile cannot.

✗ **B.** It cannot, unless weekly returns are in fact close to normal; such a result must reflect sampling error in the backtest rather than a property of the estimators.

✗ **C.** A normal distribution fitted to the most recent 26 weeks is accurate in the tails, which is where a 5% VaR is measured and where an empirical quantile is noisiest.

✗ **D.** The two estimators produce the same thresholds on average over a long sample, so the ranking is a coincidence of the particular sample and carries no lesson.

✓ **E.** The rolling volatility responds to the current regime while the all-history quantile pools calm and turbulent periods; the normal shape need not be accurate for that responsiveness to dominate.

E is the lecture's lesson: a crude translation of a responsive volatility can beat a shape-free quantile that does not condition on the current regime, and the result does not require the normal shape to fit the tails. The ranking is an empirical finding for this kind of data, not a theorem that holds in every sample. A: liquid weekly returns are close to serially uncorrelated; what clusters is volatility, so the rolling estimate tracks the size of recent moves, not their sign, and a loss last week does not predict a loss this week. B wrongly makes normality a precondition. C: the normal model misfits the tails; it wins through responsiveness, not tail accuracy. D: similar average thresholds do not make the week-by-week thresholds or the timing of exceptions similar.

45. † **[Select All That Apply]** Use return thresholds, and assume the 5% VaR threshold is negative. Which statements about 5% VaR and CVaR are correct?

- ✓ **A.** CVaR is smaller in magnitude than VaR because it averages over a narrower set of outcomes.
- ✗ **B.** A correctly calibrated 5% VaR model produces no exceptions over a long backtest.
- ✗ **C.** Empirical VaR requires an assumption about the shape of the return distribution, while normal VaR does not.
- ✗ **D.** If weekly returns are serially uncorrelated, the 5% VaR of a two-week return is twice the one-week VaR.
- ✗ **E.** CVaR is the VaR recomputed conditional on the market being in a turbulent regime.

No choice and A alone both receive full credit. A is false as written: with a negative threshold, CVaR is at least as negative as VaR, so it is at least as large in magnitude. A alone is accepted because the stem's signed return thresholds invite reading "smaller" as "more negative," and the lecture quoted VaR as a positive number, so the exam's own framing created that confusion. B, C, D and E must still be omitted. The gray check marks the accepted alternative.

None is correct. A: with a negative threshold, CVaR averages the returns at or beyond the VaR cutoff, so it is at least as negative as VaR: a smaller number read with its sign, and a larger one in magnitude. B: a calibrated 5% model produces exceptions in about 5% of weeks; none at all over a long backtest is evidence of an overconservative threshold. C is reversed: the empirical quantile makes no shape assumption, the normal VaR assumes normality. D: under the zero-mean normal model the threshold scales with the square root of the horizon, by about 1.41, not 2. E: the conditioning in CVaR is on the return falling at or beyond the VaR cutoff, not on a market regime; it is the mean of that tail.

Long-horizon risk

Exhibit 16: A long-horizon model Hypothetical annual log return relative to the risk-free investment: $x_t = \ln(1 + r_{M,t}) - \ln(1 + r_{f,t})$. Assume independent, identically distributed normal x_t, known mean 0.04, and known standard deviation 0.20. Cumulative relative log return over h years is $S_h = \sum_{t=1}^{h} x_t$. Useful normal CDF values: $\Phi(-0.2) = 0.421$, $\Phi(-0.4) = 0.345$, $\Phi(-1) = 0.159$.

Horizon h	Mean of cumulative log return	SD of cumulative log return
1 year	0.04	0.20
4 years	0.16	0.40
25 years	1.00	1.00

46. Under Exhibit 16's assumptions, what is the probability that stocks underperform the risk-free investment over four years?

- ✗ **A.** 3.1%
- ✗ **B.** 15.9%
- ✓ **C.** 34.5%
- ✗ **D.** 42.1%
- ✗ **E.** 65.5%

C: underperformance means the cumulative relative log return is below zero. Standardizing gives (0-0.16)/0.40=-0.4, so the probability is 0.345. A treats the four-year event as underperforming in each of the four years, 0.421 to the fourth power; the cumulative return can be negative without every year being negative, and positive with some negative years. B uses the 25-year horizon. D uses one year. E is the complementary probability.

47. † **[Select All That Apply]** Comparing one year with 25 years in Exhibit 16, which statements are correct?

- ✓ **A.** The standard deviation of the cumulative relative log return is larger at 25 years.
- ✓ **B.** The standard deviation of the annual-average relative log return is smaller at 25 years.
- ✓ **C.** The probability of underperforming the risk-free investment is smaller at 25 years.
- ✗ **D.** The cumulative relative log return has a smaller standard deviation at 25 years than at one year.
- ✗ **E.** An error in the assumed mean matters less at 25 years, since the average converges to the mean.

A: cumulative SD grows with the square root of h, from 0.20 to 1.00. B: average SD is 0.20 over the square root of h, from 0.20 to 0.04. C: the positive mean moves the standardized zero cutoff left. D contradicts A; it confuses the concentration of an average with the spread of a cumulative return. E: the average converges to whatever the true mean is, not to the assumed one; the standardized cutoff is -0.04 times the square root of h over 0.20, so an error in the assumed mean moves the shortfall probability more at 25 years, not less.

Exhibit 17: Evaluating a later return sample Hypothetical Barnstable-style calculation. At the beginning of a 25-year period, an analyst assumes annual stock log returns are independent normal draws with known mean 0.12 and standard deviation 0.20. At the end, the realized average annual stock log return is 0.04. These are stock log returns, not returns relative to bills.

Input	Value
Model mean, fixed before the period	0.12 per year
Model standard deviation	0.20 per year
Observed average over 25 years	0.04 per year

48. Using Exhibit 17, an analyst standardizes the observed 25-year average with the annual standard deviation, obtaining z = (0.04 - 0.12)/0.20 = -0.4, and concludes that the realized average was unremarkable under the model fixed before the period. Which statement is correct?

- ✓ **A.** Under the model a 25-year average has standard deviation 0.20/5 = 0.04, so z = -2: an average this low was a two-standard-deviation shortfall, not an unremarkable one.
- ✗ **B.** The analyst is right: with independent annual draws, the annual standard deviation describes the dispersion of the 25-year average as well as that of any single year.
- ✗ **C.** The analyst should have used the cumulative standard deviation of 1.00 instead, which makes the shortfall even less remarkable than the analyst's z = -0.4 suggests.
- ✗ **D.** No standardization is valid here, because the realized volatility of returns over the 25 years is not reported in the exhibit and the model's volatility may have been wrong.
- ✗ **E.** Under the model a 25-year average has standard deviation 0.20/25 = 0.008, so z = -10: an average this low was all but impossible under the model.

A: averaging 25 independent draws divides the standard deviation by 5; the z-score is -2, a shortfall the model gave a probability of about 2.3%. B ignores the averaging. C mixes the cumulative scale with an average: the cumulative shortfall is 25(0.04-0.12)=-2.00 against a cumulative standard deviation of 1.00, which is again z = -2. D: the calculation uses the model's standard deviation, which was fixed before the period. E divides the standard deviation by 25 rather than by its square root; it is the variance of the average that falls by a factor of 25, so its standard deviation is 0.04, not 0.008.

49. Suppose annual log stock returns relative to bills are independent normal draws with known mean zero and positive constant variance. How does the probability of stocks underperforming bills compare at one year and 25 years?

- ✗ **A.** It is lower at 25 years because the annual average return is more concentrated around zero.
- ✗ **B.** It is higher at 25 years because the cumulative log-return variance is 25 times larger.
- ✗ **C.** It cannot be compared without the numerical value of the positive variance and the horizon.
- ✓ **D.** It is 50% at both horizons because the cumulative distribution stays centered at zero.
- ✗ **E.** It is lower at 25 years because the 25-year Sharpe ratio is five times the one-year ratio.

D: the sum remains normal with mean zero and positive variance, so half the mass lies below zero. A and B ignore the unchanged standardized cutoff. C asks for a scale that does not change this probability. E multiplies a one-year Sharpe ratio of zero by five; the 25-year Sharpe ratio is also zero, and the cutoff stays at the center of the distribution.

50. † **[Select All That Apply]** A Barnstable-style calculation gives a 2% probability that stocks underperform bills over 30 years. The model assumes independent, identically distributed normal annual log returns relative to bills, with positive mean and positive variance, and treats the sample mean and volatility as known inputs. Which statements are correct?

- ✗ **A.** The 2% reflects the sampling error in the estimated mean and volatility as well as the return risk.
- ✗ **B.** The 2% is the probability that the 30-year cumulative stock return itself is negative.
- ✓ **C.** The 2% is conditional on the inputs; a lower true mean or a higher true volatility would raise it.
- ✗ **D.** The 2% means the shortfall relative to bills, when it does occur, is small.
- ✗ **E.** The 2% would be unchanged if annual returns were positively autocorrelated rather than independent.

C states what the number is. A: the inputs were treated as known. B: the event is underperforming bills, not losing money. D: the probability of the event and the size of the shortfall when it happens are different quantities; the dispersion of the cumulative return grows with the horizon. E: independence is what makes the cumulative variance grow in proportion to the horizon; positive autocorrelation adds the cross terms, the cumulative variance grows faster, and the probability rises.

51. The Barnstable-style argument that stocks are safe relative to bills over long horizons rests on more than the probability calculation. Which practical condition did the lecture stress as necessary before an investor can rely on it?

- ✗ **A.** Annual returns need to be normally distributed over the whole horizon, since the probability calculation relies on the normal distribution.
- ✗ **B.** The volatility estimate must come from the same historical sample as the mean estimate, so that the two inputs are consistent.
- ✓ **C.** The investor must be able to keep the capital committed through interim losses; the calculation says nothing about the path before the horizon ends.
- ✗ **D.** The investor must rebalance to a constant stock weight every year, because that is what makes returns independent across years.
- ✗ **E.** The risk-free rate must stay constant over the horizon, so that the return relative to bills has a known mean and variance.

C: a thirty-year investor who is forced to liquidate after a decade of underperformance never collects on the horizon; stocks trailed bills for roughly 1999 to 2009, and the argument is silent about that path. A: the calculation is an approximation; normality is not what makes or breaks the practical bet. B and E are not conditions of the argument. D: rebalancing does not create independence across years.

52. Investment X has annual arithmetic mean 8% and volatility 10%; Y has mean 10% and volatility 30%. Using $E[\log(1+r)] \approx E[r] - \sigma^2/2$ with returns in decimals gives approximate mean log returns of 7.5% for X and 5.5% for Y. Why does Y, with the higher arithmetic mean, have the lower mean log return?

- ✓ **A.** Y's variance drag exceeds X's by four points, more than Y's two-point advantage in arithmetic mean.
- ✗ **B.** Y's volatility is three times X's, so its mean log return is roughly one third of X's.
- ✗ **C.** The approximation subtracts half the volatility from the arithmetic mean, and Y's volatility is the larger.
- ✗ **D.** Y's lower figure arises only because the approximation ignores skewness, which would restore its advantage.
- ✗ **E.** The log correction removes sampling error in the arithmetic mean, and that error is larger for Y.

A: the drag is 0.045 for Y against 0.005 for X, a difference of four points that outweighs the two-point gap in arithmetic means. B treats the drag as proportional to volatility rather than to variance. C subtracts half the standard deviation instead of half the variance. D: the drag is a variance effect that is present whatever the skewness. E confuses the drag with sampling error; the approximation relates two population means of the same distribution and involves no estimate.

CAPM and DFA

Exhibit 18: CAPM regressions for size–value portfolios Monthly data, January 1981–August 2025, from the assigned DFA data file. Portfolio total returns have been converted to excess returns using RF. Separate OLS regressions on Mkt-RF include an intercept. The annualized sample market excess mean is 8.50%. LoBM denotes low book-to-market; HiBM denotes high book-to-market. Means and alphas are annualized; betas and R^2 are not.

Portfolio	Mean excess	Beta	Alpha	R^2
Small LoBM	1.17%	1.358	-10.37%	0.601
Small HiBM	12.73%	0.994	4.29%	0.548
Big LoBM	9.49%	0.995	1.03%	0.894

53. In Exhibit 18, which explanation accounts for the large negative alpha of Small LoBM?

- ✗ **A.** Its beta is negative, so the CAPM predicts a negative excess return for it.
- ✗ **B.** Its beta exceeds one, and the CAPM assigns a negative alpha to any portfolio with beta above one.
- ✓ **C.** Its 1.17% mean is far below the roughly 11.5% mean implied by its market exposure.
- ✗ **D.** The annualized market return should be subtracted once from every portfolio, regardless of beta.
- ✗ **E.** Its high volatility creates a variance drag on its mean that the CAPM-implied mean leaves out.

C: 1.358 times 8.50% is about 11.54%, and 1.17%-11.54% is about -10.37%. A reverses the beta sign. B: a beta above one raises the CAPM-implied mean; it does not set the sign of alpha. D imposes beta one on every asset. E imports the variance drag from the horizon chapter; the regression and the CAPM-implied mean are both in arithmetic excess returns, and the drag concerns the gap between arithmetic and log means, so it does not enter the alpha.

54. † **[Select All That Apply]** Which conclusions from Exhibit 18 are correct?

- ✗ **A.** Big LoBM's R-squared of 0.894 makes its population alpha zero to within rounding.
- ✓ **B.** Small HiBM and Big LoBM have almost the same estimated beta but sample mean excess returns 3.2 percentage points apart.
- ✗ **C.** Small LoBM's negative alpha means its average excess return over the sample was negative.
- ✓ **D.** The table reports estimated pricing errors; sampling uncertainty is needed to assess their statistical significance.
- ✓ **E.** The two LoBM alphas have opposite signs, so low book-to-market alone does not set the sign of the pricing error.

B compares the nearly equal betas and the 3.24 percentage-point mean gap. D separates estimated magnitude from significance. E: Small LoBM's alpha is -10.37% and Big LoBM's is +1.03%, so the two low book-to-market portfolios are mispriced in opposite directions. A infers an alpha test from time-series fit. C: the mean excess return is +1.17%; the alpha is negative because the CAPM-implied mean is 11.5%.

Exhibit 19: The CAPM cross-sectional test on the 25 portfolios Results from the same data and sample as Exhibit 18, January 1981–August 2025 (536 months), for all 25 size–value portfolios; Exhibit 18 shows three of the 25 time-series regressions. The cross-sectional regression explains the 25 annualized sample mean excess returns with the 25 estimated betas: $\bar{r}_i^e = \eta + \lambda\widehat{\beta}_i + v_i$. The annualized sample market excess return is 8.5%. The 25 time-series R^2 values range from 0.55 to 0.89.

Diagnostic	Estimate
Cross-sectional intercept η, annualized	20.6%
Cross-sectional slope λ, annualized	-10.6%
Cross-sectional R^2	0.31
Range of the 25 estimated betas	0.88 to 1.36
Mean absolute time-series alpha, annualized	2.07%
Signed mean of the 25 time-series alphas, annualized	0.3%
Joint test of the 25 alphas, H	125.9
Degrees of freedom; p-value	25; below 0.001

55. Suppose the CAPM held exactly in population and were estimated without error, with the population market premium equal to the 8.5% sample mean. What would the cross-sectional regression in Exhibit 19 show?

✗ **A.** $\eta = 0$, $\lambda = 8.5\%$, and R^2 equal to the average time-series R^2.
✗ **B.** $\eta = 8.5\%$, $\lambda = 0$, and $R^2 = 1$.
✗ **C.** $\eta = 0$, $\lambda = 1$, and $R^2 = 1$.
✓ **D.** $\eta = 0$, $\lambda = 8.5\%$, and $R^2 = 1$.
✗ **E.** η equal to the risk-free rate, $\lambda = 8.5\%$, and $R^2 = 1$.

D: mean excess returns would lie exactly on a line through the origin with slope equal to the market premium. A confuses cross-sectional fit with time-series fit; with no pricing errors the cross-sectional R-squared is one whatever the time-series values are. B puts the premium in the intercept. C confuses the premium with a beta. E puts the risk-free rate in the intercept, as if the left-hand side were total returns; the regression uses excess returns, so the line passes through the origin.

56. † **[Select All That Apply]** Which statements about Exhibit 19 are correct?

✓ **A.** Across the 25 portfolios, higher market betas went with lower average excess returns.
✗ **B.** The signed mean alpha of 0.3% per year means the CAPM pricing errors are economically small.
✗ **C.** The joint-test p-value means the market return explains little of the portfolios' month-to-month variation.
✓ **D.** The time-series alphas and the cross-sectional fit point the same way: market beta does not explain these mean returns as the CAPM requires.
✗ **E.** The intercept of 20.6% is the average excess return earned by the lowest-beta portfolios in the sample.

A: the slope is negative. D is the conclusion of §3.4. B: positive and negative alphas offset in a signed mean; the mean absolute alpha is 2.07%. C: the joint test concerns intercepts, not R-squared; the time-series R-squared values are 0.55 to 0.89. E reads the intercept as a portfolio's return; it is the fitted value at a beta of zero, outside the 0.88 to 1.36 range of estimated betas, and the fitted value at the lowest estimated beta, 0.88, is about 20.6% minus 10.6% times 0.88, or 11%. This fitted value is distinct from the portfolio's observed sample mean.

57. † **[Select All That Apply]** Assume the population CAPM holds and the market expected excess return is positive. Two funds have the same positive market beta, but Fund X has greater idiosyncratic variance. Which statements are correct?

- ✗ **A.** X earns a higher expected excess return to compensate for the extra idiosyncratic variance.
- ✓ **B.** The two funds have the same expected excess return.
- ✗ **C.** The two funds have equal Sharpe ratios because their market betas agree.
- ✓ **D.** X has higher total variance when both residuals are orthogonal to the same market factor.
- ✓ **E.** X's excess return has the lower correlation with the market excess return.

B is the CAPM mean restriction. D adds a larger residual variance to the same systematic variance. E: with the same beta and more total variance, X's correlation with the market, beta times the market volatility over X's own volatility, is lower. A prices diversifiable risk. C confuses beta with total volatility; each fund's Sharpe ratio is its market correlation times the market's, so X has the lower Sharpe ratio.

Exhibit 20: DFA factor statistics by subsample Monthly factor returns from the assigned DFA workbook. Mkt-RF, SMB, and HML are already excess or long-short returns. Means and volatilities are annualized; the Sharpe ratio is mean divided by volatility. In 1981–2001 the correlation between Mkt-RF and HML was -0.53.

Factor	1926–1980: mean / vol / Sharpe	1981–2001: mean / vol / Sharpe	2002–Aug 2025: mean / vol / Sharpe
Mkt-RF	8.1% / 20.5% / 0.39	7.8% / 15.7% / 0.50	9.1% / 15.4% / 0.59
SMB	3.4% / 11.4% / 0.30	-0.2% / 11.7% / -0.02	0.8% / 8.8% / 0.09
HML	5.0% / 13.4% / 0.37	6.5% / 11.0% / 0.59	0.1% / 10.6% / 0.01

58. † **[Select All That Apply]** Which conclusions does Exhibit 20 support?

- ✗ **A.** Each of the three factors has a positive sample premium in each subsample.
- ✓ **B.** The size premium is smaller after 1980 than before 1980.
- ✓ **C.** Since 2002 the value premium has been close to zero, while the market premium has been the largest of the three in every subsample shown.
- ✗ **D.** Because HML's volatility fell from 13.4% to 11.0% between the first two subsamples, its Sharpe ratio fell.
- ✗ **E.** SMB's Sharpe ratio of -0.02 in 1981–2001 shows that small stocks were riskier than large stocks in that subsample.

A fails for SMB in 1981–2001. B: 3.4% falls to -0.2% and 0.8%. C: HML earned 0.1% since 2002; Mkt-RF is largest in all three columns. D: HML's Sharpe ratio rose from 0.37 to 0.59; its mean rose while its volatility fell. E: a Sharpe ratio takes its sign from the mean; SMB's -0.2% mean makes it negative, and SMB's 11.7% volatility was close to its 11.4% in the earlier subsample.

59. An investor holds the market over 1981–2001 and adds a sufficiently small positive exposure to the zero-cost HML long-short factor, keeping the market position fixed. Treat the printed 1981–2001 moments in Exhibit 20 as expected inputs. Which statement is correct?

- ✗ **A.** Expected excess return rises but variance rises too, because any added position adds risk.
- ✗ **B.** The addition cannot improve the portfolio because the market's Sharpe ratio was the higher one.
- ✗ **C.** The effect cannot be assessed without HML's beta on SMB and the SMB premium.
- ✓ **D.** Expected excess return rises and variance initially falls, because HML's premium was positive and its market correlation negative.
- ✗ **E.** Expected excess return is unchanged because HML is zero-cost, while variance falls because of the negative correlation.

D: a zero-cost addition with a positive premium raises expected excess return, and a negatively correlated addition lowers variance for small positions; with these moments the variance falls until the HML exposure reaches about 0.76 times the market position, rises beyond that, and returns to the market-only level at about 1.5 times. A: a small negatively correlated addition reduces variance. B: the market's Sharpe ratio (0.50) was below HML's (0.59), and the comparison is beside the point. C: the marginal effect depends on HML's covariance with the market, which is given. E: a zero-cost position needs no capital but still has an expected return; HML's 6.5% premium raises the portfolio's expected excess return.

60. † **[Select All That Apply]** Suppose the CAPM held exactly in population with a positive market expected excess return, every portfolio has a nonzero beta, and every residual volatility is positive. For the 25 time-series regressions of portfolio excess returns on the market excess return, which population statements would be true?

- ✓ **A.** Every alpha equals zero.
- ✓ **B.** Every Treynor ratio equals the market's expected excess return.
- ✓ **C.** Every information ratio equals zero.
- ✗ **D.** Portfolios with higher betas have higher Sharpe ratios.
- ✓ **E.** No portfolio's Sharpe ratio exceeds the market's.

A, B, C are the §3.2 answers: with alpha zero, mean excess return is beta times the market premium, so mean over beta is the market premium and alpha over residual volatility is zero. E: under the CAPM the market is the tangency portfolio; each portfolio's Sharpe ratio is its correlation with the market times the market's, and positive residual volatility keeps that correlation below one. D: beta alone does not determine the Sharpe ratio; two portfolios with the same beta and different residual volatilities have different correlations with the market.

61. Assume the CAPM holds, the risk-free rate is 3%, and the market's expected excess return is 6%. An asset with market beta -0.5 has an expected total return of:

- ✗ **A.** -3%
- ✓ **B.** 0%
- ✗ **C.** 1.5%
- ✗ **D.** 3%
- ✗ **E.** 6%

B: expected excess return is -0.5(6%)=-3%, so the total return is 3%-3%=0%. A reports the excess return as if it were total. C treats the 6% as the market's total return and subtracts the risk-free rate from it again, 3%-0.5(3%). D ignores the beta. E is the market premium. A negative-beta asset has an expected return below the risk-free rate, not necessarily a negative one.

62. † **[Select All That Apply]** Which statements correctly distinguish a joint alpha test from descriptive alpha summaries?

- ✓ **A.** For given alphas, residual covariances, and market Sharpe ratio, the joint test statistic grows in proportion to the sample length.
- ✗ **B.** An MAE on its own determines significance without sample size or uncertainty information.
- ✓ **C.** A joint test can account for dependence among estimated alphas through their sampling covariance.
- ✗ **D.** One portfolio with a significant individual alpha t-statistic makes the joint test reject.
- ✗ **E.** A joint rejection shows which of the 25 portfolios carry the nonzero alphas.

AC and C alone both receive full credit. A describes the exact sample-length scaling of the course H statistic, but the exam supplied its value without the formula and referred generically to a joint test. To avoid requiring recall of that formula property, omitting A is not penalized. C must be selected and B, D, and E must be omitted. The gray check marks optional choice A.

A refers to Exhibit 19's asymptotic H statistic: T times a quadratic form in the alphas, divided by one plus the squared market Sharpe ratio. Holding the other inputs fixed, H is proportional to T. The finite-sample GRS F statistic uses a different sample-size adjustment; the exact proportionality here refers to H. C describes the covariance-weighted vector test. B treats an economic magnitude as a standardized test. D: the joint test aggregates all 25 alphas and their covariance; one individually significant alpha among 25 need not produce a rejection. E: the joint test aggregates the 25 alphas into a single statistic; a rejection says the vector is not zero, not which elements are.