

Level I of the CFA® Exam

Study Notes - Quantitative Methods

Offered by AnalystPrep

Last Updated: Oct 1, 2026

Table of Contents

1	- Returns of Financial Assets and Instruments	3
2	- Types of Financial Returns	37
3	- Benchmarking Returns	55
4	- The Time Value of Money in Finance	90
5	- Statistical Characteristics of Asset Returns	137
6	- Statistical Distributions for Financial Asset Prices and Returns	200
7	- Estimation and Hypothesis Testing	274
8	- The Return and Risk of a Financial Portfolio	328
9	- Simulation of Financial Asset Prices and Returns	394
10	- Applications of Simple Linear Regression in Finance	413
11	- Introduction to Financial Data Science	495

Learning Module 1: Returns of Financial Assets and Instruments

LOS 1a: Describe, compare, and interpret returns

Introduction to Financial Returns

A **return** measures the change in an investor's economic position over a specified period. At first glance, the idea seems simple: compare what the investment is worth at the end of the period with what it was worth at the beginning. In practice, however, investments produce different combinations of price changes, cash distributions, and contractual payoffs. A useful return measure must therefore match both the instrument being analyzed and the period over which performance is being assessed.

For example, a growth company's shares may generate most of their return through an increase in market price. A mature dividend-paying company may provide a meaningful portion of return through cash dividends. A bond may combine a changing market price with periodic coupon payments. The same return framework can also be applied to newer investment themes, such as digital-asset funds, energy-transition infrastructure, and technology-focused exchange-traded funds, even though their cash-flow patterns may differ.

The central question throughout this learning outcome is not merely, "What number did the investment earn?" It is also, "What produced the return, over what period was it measured, and how should it be interpreted?"

Financial Assets, Financial Instruments, and Financial Indicators

A **financial asset** is a claim or store of financial value held by an investor. Cash, equity ownership, debt claims, and hybrid securities are common examples. An equity investment represents an ownership interest, while a bond represents a claim on the issuer's promised payments.

A **financial instrument** is the standardized and tradable form through which a financial claim

is issued or transferred. Shares, bonds, options, futures, and other contracts package particular cash-flow or payoff patterns into securities that can be bought and sold.

A **financial indicator** is an observable measure of value or market conditions rather than a directly tradable asset. Interest rates, exchange rates, and market indexes are examples. An indicator does not itself make a cash distribution, although a security or derivative may be designed to provide exposure to movements in that indicator.

Table 1: Financial Assets, Financial Instruments, and Financial Indicators			
Category	Main Characteristic	Common Examples	Return Implication
Financial asset	Represents financial value or a claim	Cash, equity, debt, hybrid security	May produce price changes, distributions, or both
Financial instrument	Standardizes and packages a financial claim or payoff	Share, bond, option, futures contract	Return depends on the instrument's contractual cash flows and market value
Financial indicator	Measures value or market conditions but is not itself a cash-flow-producing asset	Exchange rate, interest rate, market index	An investor generally obtains exposure through another instrument

Table 1 highlights why the nature of the investment matters. A market index may rise by 8%, but an investor cannot receive that index return without holding an instrument that tracks or replicates it. Similarly, a bond's coupon payment and price movement must both be considered when evaluating the investor's performance.

How an Investment Generates Return

Figure 1: How an Investment Generates Total Return




Figure 1 separates the two basic sources of investment return. The first is the **change in market value** between the beginning and end of the holding period. The second is any **cash distribution** received while the investment is held. Combining the two gives total return. This decomposition is useful because two investments can have the same total return while relying on very different sources. One may depend almost entirely on price appreciation, while another may deliver a large part of its return through dividends or coupon income.

Price Return

The price return, also called the **capital appreciation return**, measures the percentage change in the asset's market price over the holding period:

$$r_{\text{price}} = \frac{P_1 - P_0}{P_0}$$

Where:

- r_{price} is the price return over the holding period
- P_0 is the price at the beginning of the holding period
- P_1 is the price at the end of the holding period

A positive price return indicates appreciation; a negative price return indicates depreciation. The price return excludes dividends, coupons, and other cash distributions.

Capital Distribution Return

The **capital distribution return** measures cash income received during the holding period relative to the initial investment:

$$r_{\text{distribution}} = \frac{\text{Inc}}{P_0}$$

Where:

- $r_{\text{distribution}}$ is the capital distribution return
- Inc is the cash income received during the holding period
- P_0 is the beginning price or initial amount invested

For equity, the capital distribution return is commonly called the **dividend yield**. For a bond, the annual coupon divided by the bond's beginning market price is commonly called the **current yield**. The denominator is the market price paid, not necessarily the security's face value.

Total Return

The **total return** combines the price return and the capital distribution return:

$$r = r_{\text{price}} + r_{\text{distribution}} = \frac{P_1 - P_0 + \text{Inc}}{P_0}$$

Where:

- r is the total holding-period return
- r_{price} is the price return
- $r_{\text{distribution}}$ is the capital distribution return
- P_0 is the beginning price
- P_1 is the ending price
- Inc is the cash distribution received during the period

The total return is the appropriate measure when the investor receives both a change in market value and cash income. Focusing on only one component can produce a misleading assessment. A bond's price might fall slightly during a year, for example, while coupon income allows the investor to earn a positive total return.

Example 1: Calculating an Equity Investment's Total Return

Suppose we wanted to calculate the one-year total return on shares of an AI-infrastructure company. An investor buys the shares for USD 84.00, receives a dividend of USD 1.80, and sells the shares one year later for USD 91.50. This is how to approach the problem.

Start by quoting the formula:

$$r = \frac{P_1 - P_0 + \text{Inc}}{P_0}$$

Where:

- P_0 =USD 84.00, the purchase price
- P_1 =USD 91.50, the sale price
- Inc=USD 1.80, the dividend received
- r is the total return to be calculated

First, calculate the price return:

$$r_{\text{price}} = \frac{91.50 - 84.00}{84.00} = 8.93\%$$

Next, calculate the dividend yield:

$$r_{\text{distribution}} = \frac{1.80}{84.00} = 2.14\%$$

Finally, combine the two components:

$$r = 8.93\% + 2.14\% = 11.07\%$$

The investor earned a total return of **11.07%**. Most of the return came from the increase in the share price, while the dividend provided an additional 2.14 percentage points.

Example 2: Interpreting a Bond's Price Return and Current Yield

Suppose we wanted to calculate the total return on a corporate bond purchased at 97.60 per 100 of face value. One year later, the bond is priced at 99.20 and has paid a coupon of 4.20 per 100 of face value. This is how to approach the problem.

The formula is given as:

$$r = \frac{P_1 - P_0 + \text{Inc}}{P_0}$$

Where:

- $P_0 = 97.60$, the bond's initial market price
- $P_1 = 99.20$, the bond's ending market price
- $\text{Inc} = 4.20$, the coupon received
- r is the bond's total return

The price return is:

$$r_{\text{price}} = \frac{99.20 - 97.60}{97.60} = 1.64\%$$

The current yield is:

$$r_{\text{distribution}} = \frac{4.20}{97.60} = 4.30\%$$

The total return is:

$$r = 1.64\% + 4.30\% = 5.94\%$$

The bond earned a total return of **5.94%**. The example also shows why the coupon rate and

current yield should not be treated as identical. The coupon is stated relative to the face value, whereas the current yield is based on the bond's market price.

Expected and Actual Returns

An **expected return**, also called an **ex ante return**, is estimated before the investment outcome is known. It may be based on valuation models, historical evidence, economic forecasts, or scenario analysis. Expected return is a decision-making input, not a guarantee.

An **actual return**, also called an **ex post return**, is measured after the holding period has ended. It reflects what the investor actually experienced. Unexpected changes in interest rates, inflation, regulation, business conditions, or market sentiment can cause actual return to differ materially from expected return.

The distinction is central to investment risk. Investors commit capital based on expected returns, but wealth changes with realized outcomes. When the range of possible actual returns is wide, investors generally require greater compensation for bearing uncertainty.

Realized and Unrealized Returns

An **unrealized return** arises when the market value of an investment changes while the investor continues to hold it. It is often described as a paper gain or paper loss because the position has not yet been sold.

A **realized return** occurs when the investor sells the investment or receives a cash distribution. Dividends and coupon payments are normally realized when received. A price gain or loss becomes realized when the security is sold.

This distinction matters for performance reporting, liquidity planning, and taxation. A portfolio may show a large unrealized gain, but that gain can change before the position is sold. Conversely, a realized distribution has already increased the investor's cash holdings.

Table 2: Comparison of Important Return Distinctions

Return	First Concept	Second Concept	Key Interpretation
Expected versus actual	Forecast before the outcome	Observed after the outcome	Investment decisions use expectations, but wealth changes according to actual results
Unrealized versus realized	Market-value change on an open position	Gain, loss, or distribution already crystallized	Unrealized performance can change before sale
Price versus distribution	Change in market price	Cash income received	Both components may be needed to evaluate total return
Single-period versus multi-period	Performance over one interval	Performance across several linked intervals	Multi-period analysis must account for compounding

Table 2 brings together distinctions that candidates often confuse. The concepts are related but not interchangeable. An expected price gain, for instance, is neither an actual return nor a realized return until the relevant events occur.

Single-Period and Multi-Period Returns

A **single-period return** measures performance over one clearly defined interval, such as one day, one month, one quarter, or one year. The beginning and ending values must correspond to that same interval.

A **multi-period return** summarizes performance over several consecutive intervals. Because each period's ending value becomes the next period's starting value, returns compound through time. This compounding effect is why multi-period performance cannot generally be summarized by simply adding or averaging returns without considering the purpose of the calculation.

For practical calculations, analysts commonly use the closing price on the final trading day of one period as the beginning price for the next period. Monthly returns, for example, compare the

closing price at the end of the previous month with the closing price at the end of the current month.

Arithmetic Mean Return

The **arithmetic mean return** is the simple average of periodic returns:

$$\bar{r}_i = \frac{1}{T} \sum_{t=1}^T r_{i,t}$$

Where:

- r_i is the arithmetic mean return for asset i
- T is the number of holding periods
- $r_{i,t}$ is the return on asset i during period t
- The summation operator indicates that the periodic returns are added from period 1 through period T

The arithmetic mean describes the average one-period outcome when each observation receives equal weight. It is often useful for estimating an average return for a single future period, but it does not measure the compound growth rate actually earned over several periods.

Geometric Mean Return

The **geometric mean return** is the constant periodic rate that compounds to the same ending wealth as the observed sequence of returns:

$$\bar{r}_{G,i} = \left[\prod_{t=1}^T (1 + r_{i,t}) \right]^{\frac{1}{T}} - 1$$

Where:

- $\bar{r}_{G,i}$ = geometric mean return for asset i

- T = number of holding periods
- $r_{i,t}$ = return during period t
- $\prod$ denotes the product operator, indicating that the periodic growth factors are multiplied.
- $1 + r_{i,t}$ = growth factor for period t

The geometric mean is the more appropriate measure of historical compound performance because it accounts for the effect of compounding. It recognizes that losses reduce the capital base on which subsequent returns are earned.

Example 3: Comparing Arithmetic and Geometric Mean Returns

Suppose we wanted to summarize the annual performance of a cybersecurity fund that earned 24% in Year 1, lost 12% in Year 2, and earned 9% in Year 3. We want both the average one-year return and the compound annual growth rate. This is how to approach the problem.

Start by quoting the formulas:

$$\bar{r} = \frac{1}{T} \sum_{t=1}^T r_t$$

$$\bar{r}_G = \left[\prod_{t=1}^T (1 + r_t) \right]^{\frac{1}{T}} - 1$$

Where:

- $T = 3$, the number of annual observations
- $r_1 = 24\%$, $r_2 = -12\%$, and $r_3 = 9\%$
- $\bar{r}$ is the arithmetic mean return
- $\bar{r}_G$ is the geometric mean return

The arithmetic mean is:

$$\bar{r} = \frac{24\% - 12\% + 9\%}{3} = 7.00\%$$

The geometric mean is:

$$\bar{r}_G = [(1.24)(0.88)(1.09)]^{\frac{1}{3}} - 1 = 5.95\%$$

The arithmetic mean is **7.00%**, while the geometric mean is **5.95%**. The geometric mean is lower because the sequence contains variable returns, including a loss. It is the appropriate measure of the fund's annualized compound growth over the three-year period.

The Arithmetic-Geometric Mean Relationship

When periodic returns are identical, the arithmetic and geometric means are equal. When returns vary, the geometric mean is lower. The difference between the two measures becomes wider as return variability increases.

Table 3: Return Variability and the Arithmetic-Geometric Mean Gap

Three-Year Return Path	Arithmetic Mean	Geometric Mean
10%, 10%, 10%	10.00%	10.00%
30%, -10%, 10%	10.00%	8.77%
50%, -30%, 10%	10.00%	4.92%

Every return path in Table 3 has the same arithmetic mean, but the compound outcomes differ sharply. Greater variability produces a lower geometric mean because losses and gains are not symmetric in wealth terms. A 50% loss, for example, requires a 100% gain to restore the original value.

This relationship is sometimes described as **volatility drag**. It does not mean that volatility is a separate cash expense. Rather, it describes how uneven returns reduce compound growth relative to the arithmetic average.

Annualizing Returns

Returns measured over different periods should be placed on a common time basis before being

compared. If a periodic return is assumed to repeat and compound for c periods per year, the annualized return is:

$$r_{\text{annualized}} = (1 + r_{\text{period}})^c - 1$$

Where:

- $r_{\text{annualized}}$ is the annualized return
- r_{period} is the return earned over one measurement period
- c is the number of such periods in one year

The inverse relationship converts an annual return into an equivalent periodic return:

$$r_{\text{period}} = (1 + r_{\text{annualized}})^{\frac{1}{c}} - 1$$

Where:

- r_{period} is the periodic return
- $r_{\text{annualized}}$ is the annual return
- c is the number of compounding periods in one year

Annualization improves comparability, but it also introduces an important assumption: the observed periodic return can be repeated at the same rate. This assumption is particularly questionable when a short period contains an unusually strong rally, a sharp selloff, or a one-time event.

Example 4: Annualizing a Four-Month Return

Suppose we wanted to compare a four-month return of 5.2% on an energy-storage fund with annual returns reported for other funds. We therefore need to annualize the four-month result. This is how to approach the problem.

Start by quoting the formula:

$$r_{\text{annualized}} = (1 + r_{\text{period}})^c - 1$$

Where:

- $r_{\text{period}} = 5.2\%$, the return earned over four months.
- $c = 3$, because three four-month periods fit into one year.
- $r_{\text{annualized}}$ is the return to be calculated.

Substituting the information:

$$r_{\text{annualized}} = (1.052)^3 - 1 = 16.43\%$$

The annualized return is **16.43%**. This result is useful for comparison, but it should not be interpreted as a forecast that the fund will necessarily earn 16.43% over the next twelve months. It assumes that the four-month performance can be repeated twice more at the same compound rate.

Continuously Compounded Returns

A **continuously compounded price return**, also called a **logarithmic return**, is the natural logarithm of the ending-to-beginning price ratio. If cash distributions are part of the return, the analyst should use an appropriately adjusted total-return series.

$$\tilde{r}_{t,t+1} = \ln\left(\frac{P_{t+1}}{P_t}\right) = \ln(1 + r_{t,t+1})$$

Where:

- $\tilde{r}_{t,t+1}$ is the continuously compounded return from time t to time $t + 1$.
- P_t is the price at time t .
- P_{t+1} is the price at time $t + 1$.
- $r_{t,t+1}$ is the corresponding simple holding-period return.

- $\ln(\cdot)$ is the natural logarithm.

The continuously compounded return can be converted back to a simple return:

$$r_{t,t+1} = e^{\tilde{r}_{t,t+1}} - 1$$

Where:

- e is the base of the natural logarithm.
- $\tilde{r}_{t,t+1}$ is the continuously compounded return.
- $r_{t,t+1}$ is the simple holding-period return.

The most useful property of continuously compounded returns is that they are **additive through time**:

$$\tilde{r}_{0,T} = \sum_{t=0}^{T-1} \tilde{r}_{t,t+1} = \ln\left(\frac{P_T}{P_0}\right)$$

Where:

- $\tilde{r}_{0,T}$ is the continuously compounded return over the complete period.
- $\tilde{r}_{t,t+1}$ is the continuously compounded return for each subperiod.
- P_0 is the initial price.
- P_T is the final price.
- The summation operator indicates that the subperiod logarithmic returns are added.

Simple returns must be compounded by multiplying growth factors. Log returns can instead be added, which is convenient in time-series analysis and quantitative modeling.

Example 5: Calculating and Combining Continuously Compounded Returns

Suppose we wanted to calculate the six-month continuously compounded return on a technology-services share whose price rises from USD 58.00 to USD 63.50. The price is USD 60.50 after the first three months. We also want to confirm that the two three-month log returns add to the full six-month log return. This is how to approach the problem.

Start by quoting the formulas:

$$\tilde{r}_{0,2} = \ln\left(\frac{P_2}{P_0}\right)$$

$$\tilde{r}_{0,2} = \tilde{r}_{0,1} + \tilde{r}_{1,2}$$

Where:

- P_0 = USD 58.00, the beginning price.
- P_1 = USD 60.50, the price after three months.
- P_2 = USD 63.50, the price after six months.
- $\tilde{r}_{0,1}$ is the first three-month continuously compounded return.
- $\tilde{r}_{1,2}$ is the second three-month continuously compounded return.
- $\tilde{r}_{0,2}$ is the full six-month continuously compounded return.

The direct six-month calculation is:

$$\tilde{r}_{0,2} = \ln\left(\frac{63.50}{58.00}\right) = 9.06\%$$

The two subperiod returns are:

$$\tilde{r}_{0,1} = \ln\left(\frac{60.50}{58.00}\right) = 4.22\%$$

$$\tilde{r}_{1,2} = \ln\left(\frac{63.50}{60.50}\right) = 4.84\%$$

Adding the subperiod logarithmic returns:

$$4.22\% + 4.84\% = 9.06\%$$

The result matches the direct calculation. This confirms the additive property of continuously compounded returns.

Interpretation and Evaluation

Several practical principles follow:

- A return should always be linked to a clearly defined holding period.
- Total return should include cash distributions when they are part of the investor's payoff.
- Arithmetic mean return describes an average one-period observation; geometric mean return describes compound growth.
- Annualized returns improve comparability but can exaggerate short-term performance when the repetition assumption is unrealistic.
- Continuously compounded returns are useful because they are additive through time.
- Expected, actual, realized, and unrealized returns answer different questions and should not be used interchangeably.

Question

An equity fund reports annual total returns of 16.0%, -8.0%, and 12.0% over three consecutive years. Its geometric mean annual return is closest to:

- A. 5.40%
- B. 6.13%
- C. 6.67%

Answer Explanation

The correct answer is **B**.

Start by quoting the formula:

$$\bar{r}_G = \left[\prod_{t=1}^T (1 + r_t) \right]^{\frac{1}{T}} - 1$$

Where:

- $\bar{r}_G$ is the geometric mean annual return.
- $T = 3$ is the number of annual observations.
- $r_1 = 16.0\%$, $r_2 = -8.0\%$, and $r_3 = 12.0\%$ are the annual returns.
- The product operator indicates that the annual growth factors are multiplied.

Substituting the returns:

$$\bar{r}_G = [(1.16)(0.92)(1.12)]^{\frac{1}{3}} - 1 = 6.13\%$$

The arithmetic mean does not capture the compound effect of the return sequence.

LOS 1b: describe, compare, and interpret required rates of return, risk-free rates, risk premia, and inflation

Introduction to Common Return Measures

The first learning outcome dealt with the mechanics of return measurement: price return, distribution return, total return, arithmetic and geometric averages, annualization, and continuously compounded returns. This learning outcome turns to a related but different question: **what return is required for an investor to commit capital?**

An investor can consume today or delay consumption by investing. Delaying consumption has an opportunity cost. For the investment to be acceptable, the investor must expect a return that compensates for time, expected inflation, uncertainty, costs, taxes, and, in many cases, leverage. The **required rate of return** is therefore not just a number inserted into a model. It is the minimum compensation an investor demands for bearing the relevant risks and postponing consumption.

The same investment can look attractive or unattractive depending on the measure used. A fund may report a strong gross return, while the investor receives a materially lower net return after fees. A bond may offer a positive nominal yield, while the real return becomes weak if inflation is high. An equity strategy may perform well before tax, but after-tax performance may be less impressive if returns are realized frequently. A leveraged investment can magnify gains, but it can also deepen losses when the portfolio return falls below the borrowing cost.

The Required Rate of Return

The **required rate of return** is the minimum rate of return an investor must receive to accept an investment. It reflects the compensation required for delaying consumption and bearing risk.

From the issuer's perspective, the required return is closely related to the cost of raising capital. A company issuing bonds, for example, must offer a yield high enough to attract lenders. A company issuing equity must offer a sufficient expected return to persuade investors to hold

residual ownership claims rather than safer alternatives.

A useful way to think about a required return is to separate it into a lower-risk base rate and additional compensation for specific risks. That base rate is the **risk-free rate**, and the additional compensation is a set of **risk premia**.

Figure 1: Components of the Required Rate of Return




Figure 1 shows the logic of the required return. The nominal risk-free rate reflects compensation for delaying consumption and expected inflation.

Risk premia are added because most investments are not risk-free. The exact premia depend on the instrument. Equity investors may require compensation for market risk, company size, or valuation characteristics. Bond investors may require compensation for default risk, maturity risk, or liquidity risk.

Risk-Free Rates and Inflation

The **risk-free rate** is the return on an investment with no default or reinvestment risk over the relevant investment horizon. In

practice, analysts often use short-term government securities issued by highly creditworthy sovereigns as approximations because these instruments have low default risk and limited sensitivity to interest rate changes.

The risk-free rate used in a model should match the investment's **currency** and **investment**

horizon. A

short-term government bill in one currency is not a universal risk-free rate for every investment. If a British investor evaluates a U.S. dollar bond, the U.S. risk-free rate is generally more appropriate than a local-currency rate. Similarly, if an analyst values a long-term project, using a very short-term rate without adjustment can create a maturity mismatch.

The **nominal risk-free rate** includes compensation for expected inflation. The **real risk-free rate** removes the effect of inflation and reflects compensation for postponing real consumption.

The exact relationship between the nominal return and the real return is:

$$r_{\text{real}} = \frac{1 + r_{\text{nominal}}}{1 + r_{\text{infl}}} - 1$$

Where:

- r_{real} = real return.
- r_{nominal} = nominal return.
- infl = inflation rate over the same period.

When interest rates and inflation are relatively low, the real return is often approximated as:

$$r_{\text{real}} \approx r_{\text{nominal}} - r_{\text{infl}}$$

Where:

- r_{real} = approximate real return.
- r_{nominal} = nominal return.
- infl = inflation rate.

The approximation is convenient but not exact. The approximation error is usually small when nominal returns and inflation are modest. However, the

difference becomes more meaningful when inflation is high, when returns are large, or when the effect compounds over a long investment horizon.

Example 1: Calculating a Real Risk-Free Rate

Suppose we want to calculate the real risk-free rate when a one-year nominal government bill yields 8.4% and expected inflation over the same period is 5.1%. We want the exact real return rather than the quick approximation.

Solution

$$r_{\text{real}} = \frac{1 + r_{\text{nominal}}}{1 + r_{\text{infl}}} - 1$$

Where:

- $r_{\text{nominal}} = 8.4\%$, the one-year nominal risk-free rate.
- $r_{\text{infl}} = 5.1\%$, the expected inflation rate.
- r_{real} is the real risk-free rate to be calculated.

Substitute the information:

$$r_{\text{real}} = \frac{1 + 0.084}{1 + 0.051} - 1 = 3.14\%$$

The exact real risk-free rate is **3.14%**.

The approximation would be:

$$r_{\text{real}} \approx 8.4\% - 5.1\% = 3.3\%$$

The approximation is close, but it slightly overstates the exact real return. Although the difference appears small over a single year, it can become meaningful when returns compound over many years or when inflation is relatively high.

Risk Premium

A **risk premium** is the additional return that investors expect for bearing a specific source of risk. Risk premia are not fixed constants. They vary over time, across asset classes, and among individual securities. During periods of market stress, risk premia often increase, whereas they tend to decline when investors are more willing to accept risk.

For a risky asset, the required return can be expressed approximately as the sum of the real risk-free rate, expected inflation, and the relevant risk premia:

$$r_i \approx r_{rf,real} + infl + \sum_{j=1}^n f_{ij}r_{p,j}$$

Where:

- r_i is the required return on security i
- $r_{rf,real}$ is the real risk-free rate
- $infl$ is the inflation rate
- n is the number of relevant risk premia
- f_{ij} is security i 's exposure to risk factor j
- $r_{p,j}$ is the risk premium associated with risk factor j
- $\sum_{j=1}^n f_{ij}r_{p,j}$ is the total compensation for the relevant risk exposures

This decomposition is an approximation, but it provides a useful framework for understanding what risks are being priced. A higher required return should not automatically be viewed as attractive. Instead, it may simply reflect greater investment risk, lower liquidity, longer maturity, weaker credit quality, or more uncertain future cash flows.

Table 1: Examples of Common Risk Premia

Investment Type	Risk Premium	What the Premium Compensates For	Typical Interpretation
Equity	Market risk premium	Exposure to broad equity market risk	Compensation for holding risky equities rather than a risk-free asset
Equity	Size risk premium	Risks associated with smaller-capitalization companies	Smaller firms may have less diversified revenue and more volatile earnings
Equity	Value risk premium	Risks associated with high book-to-market or lower-valued stocks	The investor may be compensated for holding companies facing stress or lower growth expectations
Debt	Default risk premium	Possibility that the issuer fails to make promised payments	Weaker credit quality generally requires higher yield compensation
Debt	Maturity risk premium	Greater sensitivity to changes in interest rates over longer horizons	Longer-maturity bonds are usually more exposed to rate movements
Debt	Liquidity premium	Difficulty selling quickly without a price concession	Less liquid securities typically require additional compensation

Table 1 shows that risk premia are tied to specific economic risks. For example, a corporate bond may offer a higher expected return than a government bill because investors require compensation for default risk and liquidity risk. Likewise, a small-cap equity fund may require a higher expected return because its cash flows and market prices are generally more uncertain. The analyst's task is to identify which risks are present and determine whether the expected compensation is adequate for bearing those risks.

Example 2: Building Up a Required Return

Suppose we want to estimate the **required return** on a small-cap equity investment. The real risk-free rate is **1.6%**, expected inflation is **3.2%**, the market risk premium is **5.0%**, and the small-cap risk premium is **1.4%**. The investment has full exposure to both risk premia.

Start by quoting the formula:

$$r_i \approx r_{rf, \text{real}} + \text{infl} + \sum_{j=1}^n f_{ij} r_{p,j}$$

Where:

- $r_{rf,real} = 1.6\%$, the real risk-free rate
- $infl = 3.2\%$, expected inflation
- $r_{p,1} = 5.0\%$, the market risk premium
- $r_{p,2} = 1.4\%$, the small-cap risk premium
- $f_{i1} = 1.0$ and $f_{i2} = 1.0$, the investment's assumed exposures to the two risk premia
- r_i is the required return to be calculated

Substitute the inputs:

$$r_i \approx 1.6\% + 3.2\% + (1.0)(5.0\%) + (1.0)(1.4\%) = 11.2\%$$

The required return is **11.2%**. This does not mean the investment will actually earn 11.2%. Instead, it represents the minimum expected return an investor would require to compensate for delayed consumption, expected inflation, broad market risk, and the additional risk associated with small-cap stocks.

Gross Return and Net Return

Investors rarely keep all of the return generated by an investment portfolio. Part of the return is absorbed by transaction costs, management fees, administrative expenses, taxes, and other charges. As a result, analysts must distinguish carefully between **gross return** and **net return**.

Gross return is the return after deducting trading expenses and other costs directly associated with executing investment activity. It is commonly used to evaluate a manager's investment process because it focuses on portfolio performance before recurring management and administrative charges.

Net return is the return after deducting management fees, administrative expenses, and other costs ultimately borne by the investor.

For most investors, this is the more meaningful measure because it more closely represents the return actually received.

In simplified form:

$$r_{\text{gross}} = r_{\text{total}} - \text{TC}$$

Where:

- r_{gross} = gross return.
- r_{total} = total return before transaction costs.
- TC = transaction costs expressed as a percentage of assets.

The net return can be approximated as:

$$r_{\text{net}} = r_{\text{gross}} - f_{\text{MF}} - e_{\text{AE}}$$

Where:

- r_{net} = net return.
- r_{gross} = gross return.
- f_{MF} = management fee.
- e_{AE} = administrative expenses.

Table 2: Gross Return Versus Net Return

Return Measure	Costs Deducted	Main Use	Candidate Trap
Total return before costs	None, in the simplified presentation	Starting point for cost analysis	Treating this as the investor's realized return
Gross return	Transaction costs and trading-related expenses	Evaluating investment process or manager skill	Confusing gross return with the amount received by the investor
Net return	Management fees and administrative expenses, after gross return	Measuring the investor's return after recurring costs	Forgetting that several layers of costs may apply

Table 2 reinforces a practical point. Gross performance may be appropriate when comparing managers' investment skill, but net performance is usually more relevant to the investor's actual wealth accumulation. A strategy that looks attractive before costs can become ordinary after fees and expenses.

Example 3: Calculating Gross and Net Return

Suppose we wanted to calculate the gross and net return on a thematic equity fund. The fund earns a total return before costs of 12.8%. Transaction costs are 0.4%, management fees are 0.9%, and administrative expenses are 0.3%. This is how to approach the problem.

Start by quoting the formula:

$$r_{\text{gross}} = r_{\text{total}} - TC$$

$$r_{\text{net}} = r_{\text{gross}} - MF - AE$$

- $r_{\text{total}} = 12.8\%$, the total return before costs
- $TC = 0.4\%$, transaction costs
- $MF = 0.9\%$, the management fee
- $AE = 0.3\%$, administrative expenses
- r_{gross} and r_{net} are the returns to be calculated

The gross return is:

$$r_{\text{gross}} = 12.8\% - 0.4\% = 12.4\%$$

The net return is:

$$r_{\text{net}} = 12.4\% - 0.9\% - 0.3\% = 11.2\%$$

The investor's net return is **11.2%**, even though the underlying portfolio generated 12.8% before costs. The difference is not a modeling detail; it is a direct reduction in the investor's return.

Pre-Tax, After-Tax, Nominal, and Real Returns

A **pre-tax return** is measured before taxes. An **after-tax return** reflects the return remaining after taxes on realized gains and distributions. This distinction matters because investors are ultimately concerned with what they keep, not only what the investment earns before taxes.

Different components of return may be taxed differently. A realized price gain may be taxed as a capital gain, while dividends or interest may be taxed as investment income. When tax rates differ, the components should be taxed separately.

The after-tax nominal return can be written as:

$$r_{\text{after-tax}} = r_{\text{price}}(1 - \text{tax}_{\text{cg}}) + r_{\text{distribution}}(1 - \text{tax}_{\text{dist}})$$

Where:

- $r_{\text{after-tax}}$ is the after-tax nominal return
- r_{price} is the price return or capital gain return
- $r_{\text{distribution}}$ is the distribution return from dividends, coupons, or similar income
- tax_{cg} is the tax rate on realized capital gains
- tax_{dist} is the tax rate on distributions

Once the after-tax nominal return is known, the after-tax real return can be calculated by adjusting for inflation:

$$r_{\text{after-tax real}} = \frac{1 + r_{\text{after-tax}}}{1 + \text{infl}} - 1$$

Where:

- $r_{\text{after-tax real}}$ is the after-tax return adjusted for inflation
- $r_{\text{after-tax}}$ is the after-tax nominal return
- infl is the inflation rate

The after-tax real return is often the most meaningful measure for a taxable investor because it measures the increase in purchasing power after taxes. It is also difficult to generalize, because tax rates depend on the investor's circumstances, account type, holding period, and jurisdiction.

Example 4: Calculating an After-Tax Real Return

Suppose we wanted to calculate the after-tax real return on a dividend-paying share. The price return is 9.0%, the dividend yield is 3.0%, the capital gains tax rate is 20%, the dividend tax rate is 10%, and inflation is 4.0%. This is how to approach the problem.

Start by quoting the formulas:

$$r_{\text{after-tax}} = r_{\text{price}}(1 - \text{tax}_{\text{cg}}) + r_{\text{distribution}}(1 - \text{tax}_{\text{dist}})$$

$$r_{\text{after-tax real}} = \frac{1 + r_{\text{after-tax}}}{1 + \text{infl}} - 1$$

Where:

- $r_{\text{price}} = 9.0\%$, the price return.
- $r_{\text{distribution}} = 3.0\%$, the dividend yield.
- $\text{tax}_{\text{cg}} = 20\%$, the capital gains tax rate.
- $\text{tax}_{\text{dist}} = 10\%$, the tax rate on dividends.
- $\text{infl} = 4.0\%$, the inflation rate.
- $r_{\text{after-tax}}$ and $r_{\text{after-tax real}}$ are the returns to be calculated.