


Quantitative Methods



Learning Module 1

Rates and Returns



LOS: Interpret interest rates as required rates of return, discount rates, or opportunity costs and explain an interest rate as the sum of a real risk-free rate and premiums that compensate investors for bearing distinct types of risk.

LOS: Calculate and interpret different approaches to return measurement over time and describe their appropriate uses.

LOS: Compare the money-weighted and time-weighted rates of return and evaluate the performance of portfolios based on these measures.

LOS: Calculate and interpret annualized return measures and continuously compounded returns and describe their appropriate uses.

LOS: Calculate and interpret major return measures and describe their appropriate uses.

Interest Rates and the Time Value of Money



LOS: Interpret interest rates as required rates of return, discount rates, or opportunity costs and explain an interest rate as the sum of a real risk-free rate and premiums that compensate investors for bearing distinct types of risk.

Interest rates can be thought of in three ways:

- The minimum rate of return required to accept a payment at a later date
- The discount rate that must be applied to a future cash flow to determine its present value (PV)
- The opportunity cost of spending the money today instead of saving it and earning a return

Interest rates are determined by the demand and supply of funds. They comprise the real risk-free rate plus compensation for bearing different types of risks:

- The **real risk-free rate** is the single-period return on a risk-free security assuming zero inflation. With no inflation, every dollar holds its purchasing power, so this rate reflects the trade-off between current versus future consumption.
- An **inflation premium** is added to the real risk-free rate to reflect the expected loss in purchasing power. The real risk-free rate + the inflation premium = the nominal risk-free rate.
- The **default risk premium** compensates investors for the risk that the borrower might fail to make promised payments in full in a timely manner.

- The **liquidity premium** compensates investors for any difficulty in converting their holdings readily into cash at their fair value. Securities that trade infrequently or with low volumes require a higher liquidity premium than those that trade frequently or with high volumes.
- The **maturity premium** compensates investors for the greater sensitivity of the market values of long-term debt instruments to changes in interest rates.

Rates of Return



LOS: Calculate and interpret different approaches to return measurement over time and describe their appropriate uses.

Holding Period Return

The **holding period return** is simply the return earned on an investment over a single specified period of time.

Exhibit 1 Holding period return

$$\text{HPR} = \frac{P_{\text{End}} - P_{\text{Beg}} + I}{P_{\text{Beg}}}$$

Price at the end of the period
 ↑
 P_{End} - P_{Beg} + I → Income
 ↓
 P_{Beg}
 Price at the beginning of the period

Note: This formula assumes that the income (eg, interest or dividend) is paid *at the end* of the period. If the income is paid any time *before* the end of the period, the return earned by investing the income for the remainder of the period would also have to be accounted for, leading to a *higher* holding period return.

Holding period returns may also be calculated for more than one period by *compounding* single-period returns:

$$R = [(1 + R_1) \times (1 + R_2) \times \dots \times (1 + R_n)] - 1$$

where:

R_1, R_2, \dots, R_n are subperiod returns

Arithmetic or Mean Return

One way to determine the periodic return over more than one holding period is to calculate the arithmetic or mean return, which is a simple average of all holding period returns:

$$\bar{R}_i = \frac{R_{i1} + R_{i2} \dots + R_{iT-1} + R_{iT}}{T} = \frac{1}{T} \sum_{t=1}^T R_{iT}$$

where:

\bar{R}_i = Arithmetic or mean return

i = Asset

R_{it} = Return in period t

T = Total number of periods

Arithmetic return is easy to calculate and has known statistical properties such as standard deviation, which is used to evaluate the dispersion of observed returns. However, the arithmetic mean return is *biased upward* because it assumes that the amount invested at the beginning of each period is the same. This bias is particularly severe if holding period returns are a *mix of both positive and negative returns*.

Geometric Mean Return

The **geometric mean return** accounts for compounding of returns and does not assume that the amount invested in each period is the same. The geometric mean is less than the arithmetic mean (due to the effects of compounding) unless there is no variation in returns, in which case they are equal. A geometric mean return provides a more accurate representation of the growth in portfolio value.



Note: The geometric mean reflects a “buy-and-hold” strategy, whereas the arithmetic mean reflects a constant dollar invested at the beginning of each time period.

The geometric mean return is calculated as:

$$R = \{(1 + R_1) \times (1 + R_2) \times \dots \times (1 + R_n)\}^{1/n} - 1$$

The Harmonic Mean

The harmonic mean is a *weighted average*. It reduces the tendency of extreme outliers to pull the mean in that direction. It is useful in dollar cost averaging or other applications in which a *ratio* is applied to a fixed quantity to obtain a *variable* number of units (eg, P/E ratios). The harmonic mean of a set of observations X_1, X_2, \dots, X_n is calculated as shown in Exhibit 2.

Exhibit 2 Harmonic Mean

$$\bar{X}_H = \frac{n}{\left(\frac{1}{X_1} + \frac{1}{X_2} + \dots + \frac{1}{X_n}\right)} = \frac{n}{\sum_{i=1}^n \frac{1}{X_i}}$$

Unless all observations are equal, the harmonic mean will be less than the geometric mean, which is, in turn, less than the arithmetic mean. In fact:

$$\text{Arithmetic mean} \times \text{Harmonic mean} = (\text{Geometric mean})^2$$

To summarize, use the arithmetic mean if the data are “well-behaved” (ie, no outliers), and use the geometric mean to represent the compound growth rate. To suppress outliers you can use the harmonic mean but alternatives to the harmonic mean are:

- the **trimmed mean**, which removes a defined percentage of the largest and smallest values from a set of data before calculating the mean, or
- the **winsorized mean**, which is calculated after replacing outlier observations with the next nearest observations.

Money-Weighted and Time-Weighted Return



LOS: Compare the money-weighted and time-weighted rates of return and evaluate the performance of portfolios based on these measures.

Money-Weighted Returns

Unlike the return measures previously discussed, the **money-weighted return** considers the amount of money invested in each period to be equivalent to the internal rate of return (IRR) (ie, the discount rate at which the sum of PVs of cash flows will equal zero). From the individual investor’s point of view, amounts invested are considered as cash outflows. Funds distributed to or withdrawn by the investor and the pool of money left at the end of the last period are considered as cash inflows.

A drawback of the money-weighted return is that it does not allow for return comparisons between different individuals or different investment opportunities. For example, two investors in the same mutual fund could have different money-weighted returns if they invested varying amounts in different periods over time.



Example 1 Money-weighted rate of return

Calculate

The investor’s *money-weighted rate of return* on the investment.

Solution

Step 1 Determine the timing and nature of the cash flow. Inflows are positive cash flows, whereas outflows are negative cash flows.

Timing of cash flows (in \$)			
$t = 0$	Purchase of first share	–	50
$t = 1$	Dividend from first share	+	1
	Purchase of second share	–	60
	Net cash flow at $t = 1$	–	59
$t = 2$	Dividends from both shares	+	2
	Proceeds from selling shares	+	130
	Net cash flow at $t = 2$	+	132

Step 2 Equate the PV of cash outflows to the PV of cash inflows (PV of outflows = PV of inflows).

$$50 + \frac{59}{(1+R)^1} = \frac{132}{(1+R)^2}$$

Step 3 Calculate the value of R to find the money-weighted return.

There are two ways to calculate R . 1) Trial and error or 2) use of the IRR function on a financial calculator. The following are the keystrokes using the TI BA II Plus calculator:

TI BA II Plus Calculator		
Keystrokes	Explanation	Display
[CF] [2nd] [CE C]	Clear CF memory registers	CF = 0.0000
50 [+ / -] [ENTER]	Initial cash outlay	CF0 = - 50.0000
[↓] 59 [+/-][ENTER]	Period 1 cash flow	C01 = -59.0000
1 [ENTER]	Frequency	F01 = 1
[↓] [↓] 132 [ENTER]	Period 2 cash flow	C02 = 132.0000
[IRR] [CPT]	Calculate IRR	IRR = 13.8612

The money-weighted rate of return (ie, the IRR) is approximately 13.86%.



Example 2 Money-weighted rate of return

Problem

An analyst gathered the following information regarding a mutual fund's returns over 5 years:

Year	Assets under management at beginning of the year (\$ millions)	Net return
1	40	25%
2	35	10%
3	55	- 10%
4	70	5%
5	30	20%

Calculate

1. Holding period return for the 5-year period.
2. Arithmetic mean return.
3. Geometric mean return.
4. Money-weighted annual return.

Solution

1. Holding period return = $[(1.25) \times (1.10) \times (0.90) \times (1.05) \times (1.20)] - 1 \approx 55.93\%$
2. Arithmetic mean return = $(0.25 + 0.10 - 0.10 + 0.05 + 0.20) / 5 = 10\%$
3. Geometric mean return = $\{[(1.25) \times (1.10) \times (0.90) \times (1.05) \times (1.20)]^{1/5}\} - 1 \approx 9.29\%$
4. To calculate the annual money-weighted return, we need to determine cash inflows and outflows.

Year	1	2	3	4	5
Balance at the beginning of year	40	35	55	70	30
Investment return of the year	25%	10%	-10%	5%	20%
Investment gain/loss	10	3.5	-5.5	3.5	6
Balance at the end of year	35	55	70	30	0
Net investment or withdrawal = (Beginning balance + Gain or loss) – Ending balance	15	-16.5	-20.5	43.5	36

The following cash flows are used to calculate the money-weighted rate of return:

$$CF_0 = -40; CF_1 = 15; CF_2 = -16.5; CF_3 = -20.5; CF_4 = 43.5; CF_5 = 36$$

$$0 = -40 + \frac{15}{(1+R)^1} + \frac{-16.5}{(1+R)^2} + \frac{-20.5}{(1+R)^3} + \frac{43.5}{(1+R)^4} + \frac{36}{(1+R)^5}$$

$$R = \text{IRR} \approx 7.97\%$$

Time-Weighted Returns

The **time-weighted rate of return** measures the compounded rate of growth of an investment over a stated measurement period. This method averages the holding period returns over time. The standard in the investment management industry is to express returns on a time-weighted basis. Whereas the money-weighted return is affected by cash withdrawals or contributions to the portfolio, which may be at the discretion of the investor, the time-weighted return is not and therefore is considered a better measure of performance.

Calculating Time-Weighted Returns

To calculate an exact time-weighted rate of return on a portfolio, the steps follow:

- Price the portfolio immediately *prior* to any significant addition or withdrawal of funds. Break the overall evaluation period into subperiods based on the dates of cash inflows and outflows.
- Calculate the holding period return on the portfolio for each subperiod.
 - For example, if a portfolio is valued daily over the course of a year, the time-weighted rate of return can be calculated as:

$$(1 + R_1) \times (1 + R_2) \times \dots \times (1 + R_{365}) - 1$$

where:

R = Daily holding period returns

- Link or compound holding period returns to obtain an annual rate of return for the year (the time-weighted rate of return for the year).

If we have annual returns data, the annualized time-weighted return can be calculated as the geometric mean of N annual returns, as follows:

$$R_{\text{TW}} = \{[(1 + R_1) \times (1 + R_2) \times \dots \times (1 + R_N)]^{1/N}\} - 1$$

Comparing money-weighted and time-weighted returns, also note that:

- if funds are deposited into the investment portfolio prior to a period of superior performance, the money-weighted return will be *greater* than the time-weighted return, and
- if funds are deposited into the investment portfolio just before a period of relatively poor performance, the money-weighted return will be *less than* the time-weighted return.

If the portfolio experiences frequent additions and/or withdrawals of funds, the time-weighted rate of return can be estimated by valuing the portfolio at frequent, regular intervals, particularly if additions and withdrawals are unrelated to market movements. The more frequent the valuation, the greater the accuracy. For many investment products and portfolios, daily valuation is commonplace.



Example 3 Time-weighted return

Problem

An investor purchases a share for \$50 today.

At the end of the year, she purchases another share for \$60.

At the end of Year 2, she sells the shares for \$65 each.

At the end of each year in the holding period, she also receives \$1 per share as dividend.

Calculate

The time-weighted rate of return.

Solution

Step 1 Break down the cash flows into two holding periods based on timing:

	<u>Holding period 1</u>	<u>Holding period 2</u>
Beginning price	50	120 (2 shares)
Dividends received	1	2 (1 per share)
Ending price	60	130 (2 shares)

Step 2 Calculate the holding-period return (HPR) for each period.

$$\text{HPR}_1 = [(60 + 1) / 50] - 1 = 22\%$$

$$\text{HPR}_2 = [(130 + 2) / 120] - 1 = 10\%$$

Step 3 Finally, calculate the compounded annual rate that would produce the same return as the investment over the two-year period.

$$(1 + \text{Time-weighted rate of return})^2 = (1 + \text{HPR}_1)(1 + \text{HPR}_2) = (1.22)(1.10)$$

$$\text{Time-weighted rate of return} = [(1.22)(1.10)]^{0.5} - 1 \approx 15.84\%.$$

Note that in Examples 1 and 3, the time-weighted return is greater than the money-weighted return. This is because the money-weighted return gave more weight to the return in the second year because the value of the portfolio was greater in the second year. Because the portfolio return was less in the second year, the money-weighted return was less overall.

Annualized Return



LOS: Calculate and interpret annualized return measures and continuously compounded returns and describe their appropriate uses.

If an investment has a term of less than one year, the return on the investment is annualized to enable comparisons across investment instruments with different maturities. So, for example, daily, monthly, or quarterly returns are converted to annualized returns.

Non-Annual Compounding

Interest is often paid less than annually. To work with interest payments paid more than once a year, the PV formula is changed to:

$$PV = FV_N + \left(1 + \frac{R_S}{m}\right)^{-mN}$$

where:

m = Number of compounding periods per year

R_S = Quoted annual interest rate (ie, the periodic rate times the number of compounding periods)

N = Number of years

For example, if interest payments are made monthly, the interest rate is not quoted on a monthly basis but is the stated annual interest rate. This rate must be “unannualized” to bring compatibility between the periodic interest rate and the number of compounding periods, N .



Example 4 Present value of a lump sum with quarterly compounding

Problem

A pension fund must make a lump-sum payment of €10 million seven years from today. The fund's manager wants to invest in an insurance-wrapped guaranteed annuity to meet this obligation. The current interest rate is 5.25%, compounded quarterly.

Calculate

The amount that the fund's manager should invest today.

Solution

Using the formula: $PV = FV_N + \left(1 + \frac{R_S}{m}\right)^{-mN}$

$$FV_N = €10,000,000$$

$$R_S = 5.25\% = 0.0525$$

$$m = 4$$

$$R_S / m = 0.0525 / 4 = 0.013125$$

$$N = 7$$

$$mN = 4(7) = 28$$

$$PV = FV_N + \left(1 + \frac{R_S}{m}\right)^{-mN}$$

$$PV = 10,000,000(1.013125)^{-28}$$

$$PV = 10,000,000(0.69412086)$$

$$PV = 6,941,208.60$$

The fund's manager will need to fund the annuity today with €6,941,208.60.

Annualizing Returns

To annualize any return, the return for the period must be compounded by the number of periods in a year.

- A monthly return is compounded 12 times, a weekly return is compounded 52 times, and a quarterly return is compounded 4 times, etc.
- Daily returns are normally compounded 365 times. For a nonstandard number of days, compounding is done by the ratio of 365 to the number of days.
- Annualized returns are calculated as:

$$R_{\text{annual}} = (1 + R_{\text{period}})^n - 1$$

where:

R = Return on investment

n = Number of periods in a year

The assumption is that the returns earned over these short investment horizons can be replicated over the year. However, this is not always possible.

**Example 5 Annualized returns****Problem**

An analyst obtained the following rates of return for three investments:

- Investment 1 offers a 7.3% return in 120 days.
- Investment 2 offers a 6.2% return in 16 weeks.
- Investment 3 offers a 5.5% return in 4 months.

Calculate

The *annualized* rates of return for these investments.

Solution

1. Investment 1:

$$R = (1 + 0.073)^{365 / 120} - 1 \approx 23.90\%$$

2. Investment 2:

$$R = (1 + 0.062)^{52 / 16} - 1 \approx 21.59\%$$

3. Investment 3:

$$R = (1 + 0.055)^{12 / 4} - 1 \approx 17.42\%$$

Continuously Compounded Returns

In continuous compounding, the number of compounding periods for a holding period is infinite. A continuously compounded return is the natural logarithm (log base e) of one plus the return of the holding period. Equivalently, it is the natural logarithm of the ending price over the beginning price.

For example, if the one-month holding period return for a bank deposit account is 0.003333, the equivalent one-month continuously compounded return is:

$$\ln(1 + 0.003333) = \ln 1.003333 = 0.003327$$

So, \$1 invested for one month at 0.3327% continuously compounded gives \$1.003333 for the holding period.

Other Return Measures and Their Applications



LOS: Calculate and interpret major return measures and describe their appropriate uses.

Gross and Net Returns

Gross returns are calculated before deductions for management expenses, custodial fees, taxes, and other expenses not directly linked to the generation of returns. Note that trading expenses, including commissions, are accounted for in the calculation of gross returns. Gross returns are an appropriate measure for evaluating portfolio performance.

Net returns deduct all managerial and administrative expenses that reduce an investor's return. Investors are primarily concerned with net returns.

Pre-Tax and After-Tax Nominal Returns

Pre-tax nominal returns do not adjust for taxes or inflation. Unless otherwise stated, assume that stated returns are pre-tax nominal returns.

After-tax nominal returns account for taxes. Most investors are concerned with returns on an after-tax basis. However, tax rates are not uniform for all investors either domestically or internationally.

Real Returns

Nominal returns are the real risk-free rate plus a premium for risk and a premium for inflation.

$$(1 + \text{Nominal risk-free rate}) = (1 + \text{Real risk-free rate}) (1 + \text{Inflation premium})$$

Investors calculate the real return because it is useful in comparing returns:

- across time periods because inflation rates may vary over time, and
- among countries when returns are expressed in local currencies and inflation rates vary between countries.

The after-tax real return is also used as an investor's benchmark return because it is the required compensation for postponing consumption and assuming risk after paying taxes on investment returns.

Frequently, the real risk-free return and the risk premium are combined to arrive at the real "risky" rate, which is simply referred to as the real return, or:

$$(1 + \text{Real return}) = \frac{(1 + \text{Real risk-free rate})(1 + \text{Risk premium})}{(1 + \text{Inflation premium})}$$

Real after-tax returns are not usually calculated by investment managers because it is difficult to estimate a general tax rate that is applicable to all investors.

Leveraged Return

The leveraged return is calculated when an investor uses leverage (by either borrowing money or using derivative contracts) to invest in a security. Leverage enhances returns but also magnifies losses.



Example 6 Calculation of special returns

Problem

Continuing from Example 2, suppose that the mutual fund spends a fixed amount of \$600,000 every year on expenses that are *unrelated* to the manager's performance.

Year	Assets under management at beginning of the year (\$ millions)	Net return
1	40	25%
2	35	10%
3	55	-10%
4	70	5%
5	30	20%

Assume that an investor faces a **tax rate of 25.00%** and that the **inflation rate is 3.00%**,

Calculate

1. The annual **gross** return for the fund in Year 1.
2. The **after-tax net** return for the investor in Year 2.

Assume that all gains are realized at the end of the year and that taxes are paid immediately at that time.

3. The expected **after-tax real** return for the investor in Year 5.
4. The **net** return earned by investors in the fund over the 5-year period.

Solution

1. The fixed expenses of \$600,000 would make the gross return greater than net return by 1.50% ($600,000 / 40,000,000$).

Therefore, gross return for Year 1 would equal 26.50% (1.50% + 25.00%).

2. After-tax return (Year 2) = $10\% \times (1 - 0.25) = 7.50\%$
3. After-tax return (Year 5) = $20\% \times (1 - 0.25) = 15.00\%$

$$\begin{aligned}\text{After-tax real return (Year 5)} &= \frac{(1 + 0.15)}{(1 + 0.03)} - 1 \\ &\approx 1.1165 - 1 \approx 0.1165 \approx 11.65\%\end{aligned}$$

4. The HPY for the fund over the 5-year period is calculated after considering all direct and indirect expenses. The net return is 55.93% (same answer as shown in Example 2).

Learning Module 2

The Time Value of Money in Finance



LOS: Calculate and interpret the present value (PV) of fixed-income and equity instruments based on expected future cash flows.

LOS: Calculate and interpret the implied return of fixed-income instruments and the required return and implied growth of equity instruments given the present value (PV) and cash flows.

LOS: Explain the cash flow additivity principle, its importance for the no-arbitrage condition, and its use in calculating implied forward interest rates, forward exchange rates, and option values.

Time Value of Money in Fixed Income and Equity



LOS: Calculate and interpret the present value (PV) of fixed-income and equity instruments based on expected future cash flows.

Fixed-Income Instruments and the Time Value of Money

Fixed-income instruments are bonds or debt that are characterized by cash flows that occur at regular intervals. Three main cash flow patterns can occur:

- **Discount**, where the only cash flows are:
 - the lender's upfront payment (investment) in the instrument, and
 - the borrower's repayment of the face value when the instrument matures.
- **Periodic interest**, where the investor receives regular interest payments on the principal amount plus a final payment that incorporates both the last interest payment and the principal amount.
- **Level payments**, where each payment is an identical amount. Payments include both interest and principal, and the amount of each payment is typically structured so that principal has been entirely repaid when the instrument matures.

Discount Instruments

The only cash flows for a discount instrument are the amount that the investor pays today (present value, or PV) and the face value that the investor receives at maturity (future value, or FV). What links the PV and the FV is the **discount factor**. For bonds, the discount factor is the bond's yield to maturity (r). The formula for PV of a discount bond is:

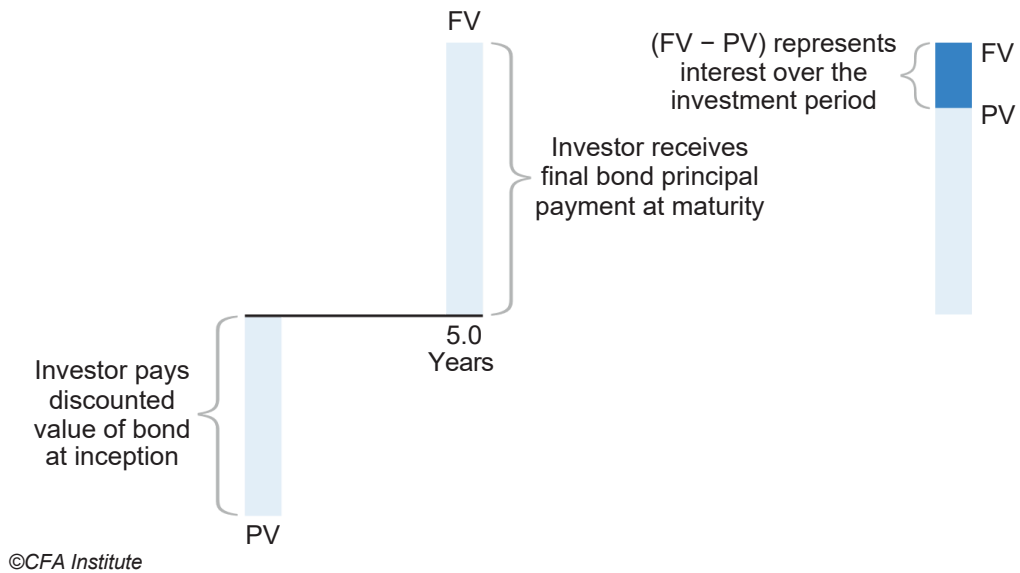
$$PV = FV_t / (1 + r)^t$$

where:

t = The number of periods to maturity

The timeline for cash flows of a discount bond is shown in Exhibit 1.

Exhibit 1 5-year discount bond



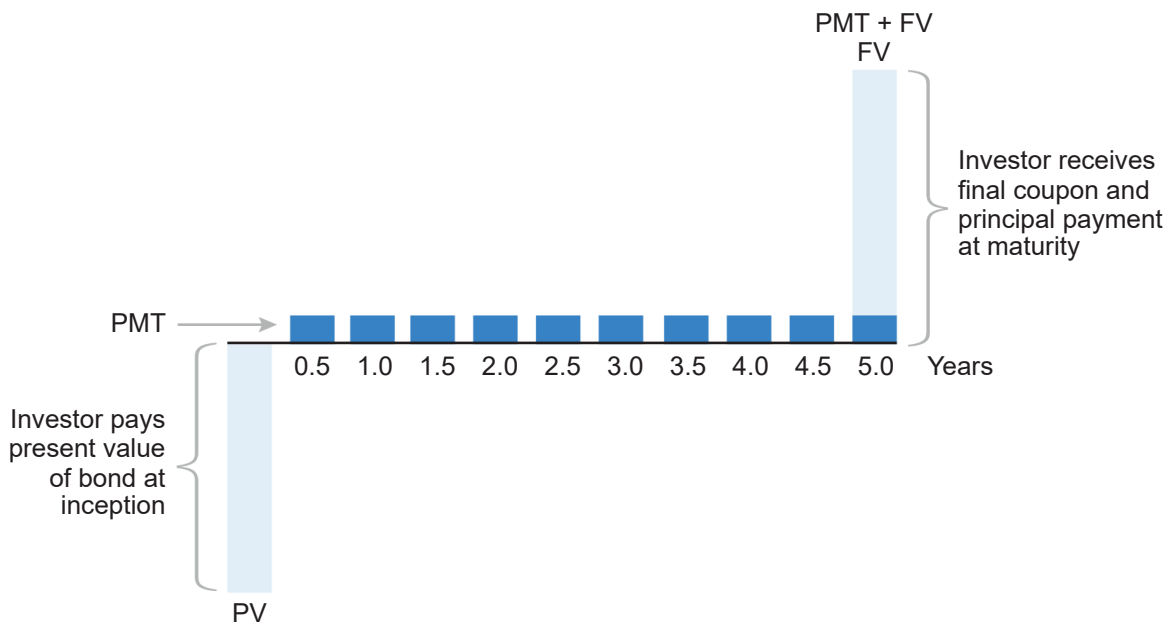
Coupon Instrument

Coupon bonds make regular interest payments (“coupons”) to investors. In most cases, principal is paid with the last interest payment at maturity. The price of the bond is the sum of the present values of all coupons and principal. As with discount bonds, the discount factor is the bond’s YTM. For a bond that pays coupons annually, the present value is calculated as:

$$PV_{\text{coupon bond}} = \frac{C_1}{(1+r)^1} + \frac{C_2}{(1+r)^2} + \dots + \frac{C_t + FV_t}{(1+r)^t}$$

If the bond in Exhibit 1 paid a coupon, the cash flow timeline would be as depicted in Exhibit 2.

Exhibit 2 Coupon bond cash flows



The bond in Exhibit 2 pays coupons over five years, but each coupon occurs every six months. From a time value of money perspective, the frequency of payments can have a significant effect on the present value of the instrument. Example 1 compares the present value of two bonds that are identical in all respects except that one pays coupons annually, whereas the other pays every six months.



Example 1 Coupon bond cash flows

Problem

A three-year bond with face value of 100 has a 10% coupon, and the current YTM is 7%.

Calculate

Calculate the PV of the bond, assuming (a) the coupon is paid annually and (b) the coupon is paid semiannually.

Solution

For an annual-pay bond, the PV is calculated as:

$$\frac{10}{(1 + 0.07)^1} + \frac{10}{(1 + 0.07)^2} + \frac{10 + 100}{(1 + 0.07)^3} = 107.87$$

If this bond pays coupons semiannually, the numbers must be adjusted to account for the greater frequency of payments:

- The coupon is 10 for the year, but since it is paid every six months, each coupon payment is 5
- YTM is expressed as an annual percentage, so the semiannual YTM is 7% / 2, or 3.5%

The PV of the semiannual-pay bond is calculated as:

$$\frac{5}{\left(1 + \frac{0.07}{2}\right)^1} + \frac{5}{\left(1 + \frac{0.07}{2}\right)^2} + \dots + \frac{5}{\left(1 + \frac{0.07}{2}\right)^5} + \frac{5 + 100}{\left(1 + \frac{0.07}{2}\right)^6} = 107.99$$

The semiannual bond, despite being identical to the annual bond in all respects except for coupon frequency, has a higher PV. Why?


One assumption in fixed-income pricing is that coupon payments are reinvested at the YTM. More frequent cash flows mean more frequent reinvestment opportunities. The semiannual bond investor in Example 1 will generate slightly more total income over the three-year life of the bond, since money will be reinvested 5 times in that period, as opposed to only twice for the annual bond.

Annuities

An annuity is a fixed-income instrument that features identical payments made at fixed intervals. Unlike coupon bonds, each payment includes both interest and principal. The payments are structured such that, at the maturity of the instrument, the principal has been repaid in full (ie, $FV_t = 0$). This process is known as **amortization**. A common type of annuity is a home mortgage loan.

Exhibit 3 shows the cash flows for an annuity.

Exhibit 3 Cash flows for an annuity



©CFA Institute

Equity Instruments and the Time Value of Money

Equity investments, in addition to representing ownership shares in a company, entitle investors to receive discretionary cash flows such as dividends. Similar to fixed-income instruments, the value of equity can be estimated using a discount rate to discount expected future cash flows. However, the calculated value cannot be determined with certainty, since one of the component cash flows is the amount to be received when the instrument is sold in the future.

One way to value a company's equity is by discounting expected cash flows using the cost of equity (r). Common assumptions associated with valuing equity through discounting cash flows are:

- **Constant dividends:** an investor pays a fixed price (PV) for a preferred or common share in exchange for fixed periodic dividends
- **Constant dividend growth rate:** an investor pays an initial price (PV) for a share of stock and received an initial dividend in one period $D_{(t+1)}$, which is expected to grow at a constant growth rate of g
- **Changing dividend growth rate:** an investor pays an initial price (PV) for a share of stock and receives an initial dividend $D_{(t+1)}$. The dividend is expected to grow at a rate that changes over time as the company moves from a period of high initial growth to longer-term more stable growth.

Constant Dividends

If a common share or preferred share is expected to pay a constant dividend, the formula to discount this infinite series of constant cash flows is:

$$PV = \sum_{i=1}^{\infty} \frac{D_t}{(1+r)^i} = \frac{D_t}{r}$$

where:

D_t = Dividend at time t

r = Required rate of return



Example 2 Value of stock paying constant dividend

Problem

A company is expected to pay regular dividends of EUR 3.7 per share in perpetuity. What is the value of a share of this company if the appropriate discount rate is 12%?

Solution

Solving for PV using the formula previously given, we substitute EUR 3.7 for D_t and 0.12 for r . We get $3.7 / 0.12 = \text{EUR } 30.08$.

Dividends with Constant Growth Rate

When we assume that dividends will grow at a constant growth rate into perpetuity, the formula to discount that infinite series of cash flows is the following:

$$PV = \sum_{i=1}^{\infty} \frac{D_t(1+g)^i}{(1+r)^i} = \frac{D_{(t+1)}}{(r-g)}$$

where:

g = Assumed growth rate

For this formula to work, the required rate of return, r , must be greater than the growth rate, g . In the long term, g will be smaller than r . We will explore the short-term scenarios where g is larger than r later in this same learning module.



Example 3 Value of stock paying a constantly growing dividend

Problem

A Japanese company is expected to pay a dividend of JPY 150 yen at the end of the current period. That dividend is expected to grow by 3% in perpetuity and the required rate of return is 10%.

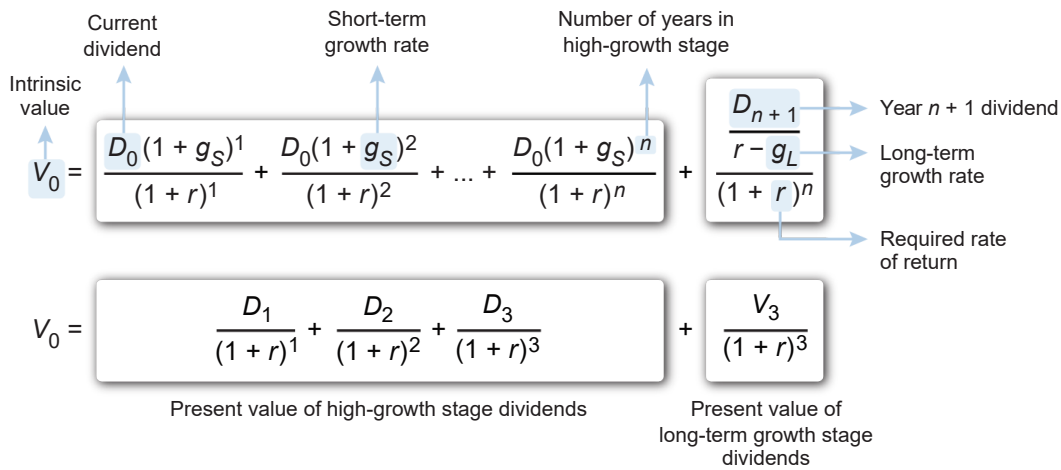
Solution:

Substituting 150 for $D_{(t+1)}$, 3% for g , and 10% for r , we get $150 / (0.10 - 0.03) = \text{JPY } 2,142.56$.

Dividends with Changing Growth Rates

Where dividends are expected to grow at a high rate for a short time and then at a lower rate (necessarily lower than the discount rate) thereafter, we first forecast dividends during the high-growth stage, then we forecast dividends using the lower long-term (eg, terminal) growth rate. The value of the low-growth dividends is as of the end of the high-growth stage; this value is then discounted back to today. The sum of the present values for both the high- and low-growth dividends comprises the estimated price of the stock. This estimation can use the formula given before. Subsequently, we discount all the cash flows at the appropriate discount rate. Exhibit 4 demonstrates how to derive the value of a share using a multi-stage dividend discount model.

Exhibit 4 Two-stage dividend discount model



Example 4 Multi-stage dividend discount model

Problem

A Brazilian company is currently undergoing a three-year period of 20% growth, subsequently the growth rate is expected to level off to 3%. The current dividend is BRL 28.33, and the appropriate discount rate is 16%.

What is the value of a share of its stock?

Solution:

We first need to estimate the dividends for the next three years, as well as estimate the price of the stock at the end of the third year.

$$D_1 = 28.33 \times (1 + 0.20) = 34$$

$$D_2 = 34 \times (1 + 0.20) = 40.8$$

$$D_3 = 40.8 \times (1 + 0.20) = 48.96$$

Next, calculate the present value for each of the first three dividends:

$$PV_1 = \frac{34}{1 + 0.16} \approx 29.31$$

$$PV_2 = \frac{40.8}{(1 + 0.16)^2} \approx 30.32$$

$$PV_3 = \frac{48.96}{(1 + 0.16)^3} \approx 31.37$$

The sum of those present values is approximately BRL 91.

Next, calculate the present value of the dividend stream for the long-term growth stage. The long-term growth stage begins at the end of Year 3, so the first dividend in this stage is paid at the end of Year 4:

$$D_4 = 48.96 \times (1 + 0.03) \approx 50.43$$

Since it is assumed that the dividend will grow at 3% indefinitely, the present value of this stream is:

$$PV_{\text{long term}} = \frac{D_{\text{long term}}}{r - g} = \frac{50.43}{0.16 - 0.03} \approx \text{BRL } 387.91$$

The value of a share is the sum of the present values of both dividend streams:

$$91 + 387.91 \approx 468.91$$

Implied Return and Growth



LOS: Calculate and interpret the implied return of fixed-income instruments and the required return and implied growth of equity instruments given the present value (PV) and cash flows.

Implied Return for Fixed-Income Instruments

In fixed-income instruments, if we observe the present value (price) and the contractual cash flow payments, then the discount rate (yield-to-maturity) is simply a measure of implied return.

The formula to calculate PV for a discount instrument,

$$PV = \frac{FV_t}{(1 + r)^t}$$

Can be rearranged to derive r :

$$(1 + r) = \left(\frac{FV_t}{PV} \right)^{1/t}$$

$$r = \left(\frac{FV_t}{PV} \right)^{1/t} - 1$$

With coupon bonds, the YTM is the discount rate used for all cash flows, regardless of timing.

Equity Instruments, Implied Return, and Implied Growth

As noted previously, the present value (PV) of equity instruments is a function not only of the cash flows and the discount rate, but also of the implied cash flow growth.

If we start with the constant growth model (Gordon growth model) to solve for r we get:

$$PV_t = \frac{D_{t+1}}{(r - g)}$$

$$r - g = \frac{D_{t+1}}{PV_t}$$

$$r = \frac{D_{t+1}}{PV_t} + g$$

Hence the cost of equity is simply the forward dividend yield plus the rate of growth. Similarly, the rate of growth is the cost of equity less the forward dividend yield.

$$g = r - \frac{D_{t+1}}{PV_t}$$

Implied P/E Ratio

Starting from the constant growth (Gordon growth) formula:

$$PV_t = \frac{D_{t+1}}{(r - g)}$$

We can also derive a formula for the **forward P/E ratio** by simply dividing both sides by the EPS.

$$\frac{PV_t}{E_{t+1}} = \frac{\frac{D_{t+1}}{E_{t+1}}}{(r - g)}$$

Likewise, the formula for the **P/E ratio** can also be derived from the constant growth formula:

$$\frac{PV_t}{E_{t+1}} = \frac{\frac{D_t}{E_t} (1 + g)}{(r - g)}$$

A dividend/share divided by an earnings per share from the same period is called a **dividend payout ratio**. Given a dividend payout ratio and a P/E ratio, one can now easily solve for either the required rate of return or the implied dividend growth rate.

Once the implied growth rate has been calculated, it can be compared to the company's expected growth rate or its historical growth rate.

Cash Flow Additivity



LOS: Explain the cash flow additivity principle, its importance for the no-arbitrage condition, and its use in calculating implied forward interest rates, forward exchange rates, and option values.

Under the **cash flow additivity principle**, the present value of any future cash flow stream is equal to the sum of the present value of the cash flows. This principle is useful in ensuring that the market prices reflect the condition of no arbitrage.

As an example of the cash flow additivity principal, assume we have the following two investment opportunities:

Period	Project A	Project B
Initial investment	-100	-100
Period 1	50	10
Period 2	50	40
Period 3	50	70
Period 4	50	95

If we have to select between only one of these two investment opportunities and the relevant discount rate is 11.11%, we would first have to determine whether either has a positive NPV and then select the project with the highest NPV.

We could also calculate the present value of the difference in cash flows and discount that difference, which would give us the same difference as the difference in NPVs calculated separately.

NPV using a financial calculator for Project A is \$54.76 million:

- $CF_0 = -100$
- $CF_1 = 50, F_1 = 4$
- $i = 11.11$

Project B's NPV is also \$54.76 million:

- $CF_0 = -100$
- $CF_1 = 10$
- $CF_2 = 40$
- $CF_3 = 70$
- $CF_4 = 95$
- $i = 11.11$

This can be verified by subtracting Project B's cash flows from Project A's and entering those into the financial calculator.

Period	Project A – Project B
Initial investment	$(-100 - [-100]) = 0$
Period 1	$50 - 10 = 40$
Period 2	$50 - 40 = 10$
Period 3	$50 - 70 = -20$
Period 4	$50 - 95 = -45$

- $CF_0 = 0$
- $CF_1 = 40$
- $CF_2 = 10$
- $CF_3 = -20$
- $CF_4 = -45$
- $i = 11.11$

Solving for NPV, we get 0.0, which is precisely the difference between both NPVs that were both equivalent at \$54.76 million. If the NPV of the difference in cash flows had been positive, Project A would be a more attractive investment, and, if the NPV of the difference in cash flows had been negative, Project B would be preferred.

Implied Forward Rates Using Cash Flow Additivity

If a risk-neutral investor wishes to invest MXN 1,000 for two years and the one- and two-year risk-neutral rates of interest are currently:

- One-year risk-free discount bond, $r = 7.8\%$
- Two-year risk-free discount bond, $r = 9.5\%$

She can accomplish this by either investing today in the two-year bond or by investing in the one-year bond today and, when that bond matures, invest in the prevailing one-year rate available at that time.

First Investment Strategy

If she were to invest MXN 1,000 today in the two-year risk-free discount bond at the end of Year 2, she would have $\text{MXN } 1,000 (1 + 9.5\%)^2 = \text{MXN } 1,119.03$.

Second Investment Strategy

If she were to invest MXN 1,000 today in the one-year risk-free discount bond she would have MXN 1,078 at the end of Year 1 and then invest in the prevailing rate at that point. The level of that rate would determine which strategy was preferable.

If we define the forward rate in one year's time for one year as ${}_1F_1$, then under the cash flow additivity principle, she would be indifferent between the two strategies under the following circumstance:

$$(1 + r_2)^2 = (1 + r_1)(1 + {}_1F_1)$$

We can rearrange the formula to solve for ${}_1F_1$ as follows:

$${}_1F_1 = \frac{(1 + r_2)^2}{(1 + r_1)} - 1$$

In this example the ${}_1F_1$ rate is $(1.095)^2 / 1.078 - 1 = 11.23\%$.

This example illustrates how forward rates should be set so that investors cannot earn riskless arbitrage profits.

Forward Exchange Rates Using No-Arbitrage Condition

Assume you have EUR 100,000 to invest for nine months. You can either invest in a riskless investment in Germany or in the United States. Let us assume that the current exchange rate is USD 1.14/EUR and that the nine-month risk-free rates in USD and Euros are 4.7% and 5.3%, respectively.

Investment Strategy 1

Invest the EUR 100,000 at the Euro rate of 5.3% for nine months.

- At time ($t = 0$), invest the EUR 100,000 at the 5.3% Euro rate for 9 months.
- At time ($t = T$), receive EUR 104,0055.06 ($100,000 e^{(0.053 \times 9/12)}$)

Investment Strategy 2

Convert EUR 100,000 to USD at the USD 1.14/EUR rate and then invest the USD 1,140,000 at the nine-month rate of 4.7% and then convert them back to EUR at the USD 1.13488/EUR forward rate for nine months set today.

- At $t = 0$, convert the EUR 100,000 to USD 114,000 and invest at the nine-month 4.7% rate.
- At $t = T$, the investor receives USD proceeds of 118,090.17 and converts them back to EUR at the forward rate set at $t = 0$ of USD 1.13488/EUR to receive EUR 104,055.06.

These strategies have the same payoff due to the nine-month forward rate of USD 1.13488. If the forward rate were higher or lower, it would be possible to earn a riskless arbitrage profit.

In this case, the forward exchange rate would be set using $1.14e^{(0.047 \times 9/12)} / e^{(0.053 \times 9/12)}$. That same formula simplified is $1.14e^{((0.047 - 0.053) \times 9/12)}$.

Option Pricing Using Cash Flow Additivity

Assume that an asset has a price of USD 45. The price might rise 25% in the next year (to USD 56.25) or fall 20% (to USD 36). If an investor wishes to sell a contract on the asset in which the buyer has the right, but not the obligation, to buy the asset for USD 50, we can use the principle of no arbitrage to determine the price at which to sell this contract.

Under these two possible price scenarios, in one year, if the stock rises to USD 56.25, the contract buyer will choose to buy it for USD 50, but if it falls to USD 36, the contract buyer will let the contract expire unexercised.

- At $t = 0$ the contact value is c_0 .
- At $t = 1$ the contract value will either be US 6.25 ($56.25 - 50$) or US 0.00.

Assume that, at $t = 0$, the contract is sold at a price of c_0 , and 0.3086 shares of the underlying asset are purchased.

Learning Module 2

Under both the price increase and decrease scenarios the value of the portfolio is:

- $V_0 = 0.3086 \times 45 - c_0$
- $V_1^u = 0.3086 \times 56.25 - 6.25 = 11.11$
- $V_1^d = 0.3086 \times 36 - 0.00 = 11.11$

Since the value of the replicating portfolio is 11.11 independently of whether the portfolio rises or falls, the payoffs can be discounted at the risk-free rate.

Assuming a risk-free rate of 5.2%:

$$V_0 = \frac{V_1^u}{(1+r)} = \frac{V_1^d}{(1+r)}$$

In this example the formula would be the following:

$$V_0 = 0.3086 \times 45 - c_0 = \frac{11.11}{1.052}$$

Given that $0.3086 \times 45 = 13.89$, then c_0 has to be $13.89 - 11.11 = 2.78$.

Learning Module 3

Statistical Measures of Asset Returns



LOS: Calculate, interpret, and evaluate measures of central tendency and location to address an investment problem.

LOS: Calculate, interpret, and evaluate measures of dispersion to address an investment problem.

LOS: Interpret and evaluate measures of skewness and kurtosis to address an investment problem.

LOS: Interpret correlation between two variables to address an investment problem.

Measures of Central Tendency and Location



LOS: Calculate, interpret, and evaluate measures of central tendency and location to address an investment problem.

A measure of central tendency (eg, a mean, median, or mode) shows where the data set is centered; a measure of location provides more description about the data set. A measure of central tendency is a parameter if it derives from population data and is a statistic if it derives from sample data.

Measures of Central Tendency

The Arithmetic Mean

The arithmetic mean is the sum of the observations divided by the number of observations. A sample mean is calculated as:

$$\bar{X} = \frac{X_1 + X_2 + \dots + X_n}{n}$$

where:

$X_{1,2,\dots,n}$ = Observation being measured (eg, individual equity return)

n = Number of observations

The arithmetic mean is sensitive to extreme values; that is, disproportionately large or small values will drag the mean away from its centrality. To help avoid the outlier problem, a trimmed mean includes only some percentage of the middle values of the data set, or a winsorized mean replaces outliers with the highest or lowest observation from a defined percentage of the middle data.

The Median

The median is the middle value of the data as ranked in ascending or descending order. The median for an odd number of variables is the observation in the position identified by $(n + 1) / 2$. The median for an even number of variables is the average of observations in the positions identified by $n / 2$ and $(n + 2) / 2$.

Regardless of which formula is used, an equal number of observations lie above and below the median, which is unaffected by outliers.



Example 1 Calculating the median

Problem

Calculate the median score of the following scores:

77, 90, 57, 85, 68, 31, 45, 86, 46, 98, 25, 10, 57, 67, 88, 77, 34, 89, 47, 77

Solution

First, arrange the scores in ascending or descending order:

10, 25, 31, 34, 45, 46, 47, 57, 57, 67, 68, 77, 77, 77, 85, 86, 88, 89, 90, 98

Since this is an even-numbered data set (20 observations), the median is the average of the tenth ($n / 2$) and eleventh $(n + 2) / 2$ observations. The median for this data set is therefore $67.5 (67 + 68) / 2$.

The Mode

The mode is the most frequently occurring value in a data set. For grouped data, the modal interval is the interval with the greatest number of observations. A data set may have no mode, one mode (ie, unimodal), or more than one mode.

Dealing with Outliers

There are three main ways to address outliers:

- Use the data without any adjustment.
 - Outliers may contain important information.
 - Determining what is an “outlier” may be subjective, so using all data points eliminates subjectivity.
- Delete all outliers. The **trimmed mean** calculates an arithmetic mean after excluding a given percentage of the highest and lowest values.
- Substitute values for extreme values. The **winsorized mean** substitutes one specified value for all values at or below a stated percentage of the lowest values and another specified value for all values at or above a stated percentage of the high values. For example, in a 90% winsorized mean, the lowest 5% of all values are set equal to the value that represents the lowest 5%, and the highest 5% of all values are set equal to the value at or below which 95% of all observations lie.

Measures of Location

Quantiles divide data into more groups: quartiles into four groups, quintiles into five groups, deciles into 10 groups, and percentiles into 100 groups. The y th percentile describes the value at or below which $y\%$ of observations occur. Cutoffs for the other quantiles may be stated in terms of percentages; for example, 75% of observations fall below the 4th quartile.

An interquartile range is the difference between the top of the third quartile and the bottom of the second quartile (ie, top of the first quartile) or $IQR = Q_3 - Q_1$.



Example 2 Calculating quartiles

Problem

1. Calculate the first quartile of a distribution that consists of the following asset returns: 10%, 23%, 13%, 17%, 19%, 5%, 4%.
2. If we include one more return observation of 10% in our data set, what is the new value of the first quartile?

Solution

1. First, arrange the data in ascending order:

4%, 5%, 10%, 13%, 17%, 19%, 23%

There are seven observations, so the first quartile = $(7 + 1) \times 25/100 = 2$. The first quartile is the second observation in the data set, or 5%.

One quarter, or 25%, of the observed returns lie below the second observation, which is 5%.

1. Once again, we begin by arranging the data in ascending order:

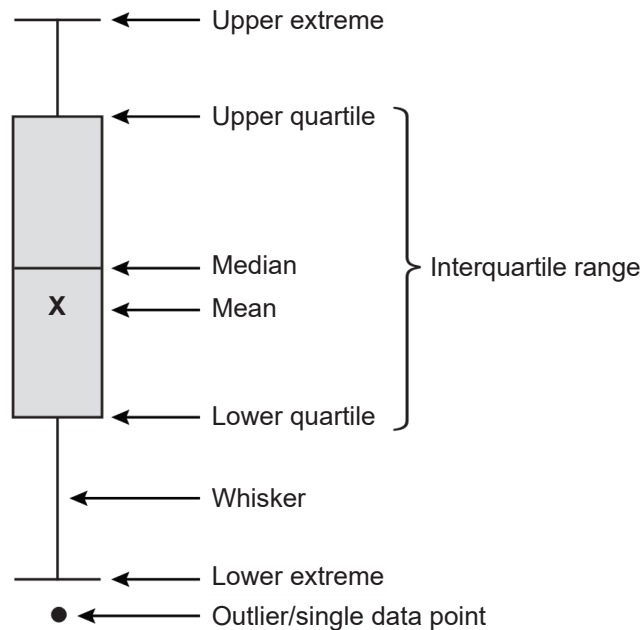
4%, 5%, 10%, 10%, 13%, 17%, 19%, 23%

Now there are eight observations. The first quartile = $(8 + 1) \times 25/100 = 2.25$.

This means that, once the data set has been rearranged in ascending order, the first quartile is the second observation plus 0.25 times the difference between the second and third observations. Therefore, one-fourth of the observed returns are below 6.25% [= $5\% + 0.25(10\% - 5\%)$].

Quantiles may be visualized by a box-and-whisker plot, which shows whiskers at the highest and lowest values and boxes for the middle quantile ranges. Exhibit 1 shows a box-and-whisker plot for a quartile:

Exhibit 1 Box and whisker plot



Quantiles in Investment Practice

In finance, quantiles are used to rank performance. Managers are often evaluated based on the percentile of their performance relative to their peers, so a manager ranking in the 95% percentile (ie, 95% of all returns below that manager's return) is considered superior to one ranked in the 80% percentile.

Quantiles are also used for investment research. Dividing data into quantiles based on criteria such as sales or market share allows for evaluation of that criterion on metrics such as return on assets.

Measures of Dispersion



LOS: Calculate, interpret, and evaluate measures of dispersion to address an investment problem.

Dispersion describes variability of the data series around the central tendency and is often described as risk in the context of average returns.

The Range

The range is the difference between a data set's maximum and minimum values. Although it cannot describe the shape of the distribution, the range is the simplest dispersion measure. If extreme outliers exist, the range will not be representative of the observations. For instance, if the average of 20 annual returns is 3%, and for 19 returns, the lowest is 2.9% and the highest is 3.1%, a return of 10% will result in a far greater range than if only the 19 "representative" returns were used.

The Mean Absolute Deviation

Given the limitations of the range, another possibility is to use the arithmetic mean of a data set's deviations around the set's mean. However, the differences between each observation and the mean sum to zero, so taking the average of those deviations also equals zero. The mean absolute deviation (MAD), however, uses the absolute value of these differences:

$$\text{MAD} = \frac{|(X_1 - \bar{X})| + |(X_2 - \bar{X})| + \dots + |(X_i - \bar{X})|}{n}$$

Although difficult to manipulate mathematically, the MAD does use all observations in the data set, whereas the range does not.



Example 3 Calculating mean absolute deviation

Problem

Calculate the mean absolute deviation of the following data set that contains average returns earned by a portfolio manager over the last 6 years:

4%, 1%, 8%, 10%, 3%, 16%

Solution

$$\text{The mean} = \frac{4\% + 1\% + 8\% + 10\% + 3\% + 16\%}{6} = 7\%$$

$$\begin{aligned} \text{MAD} &= \frac{|0.04 - 0.07| + |0.01 - 0.07| + |0.08 - 0.07| + |0.10 - 0.07| + |0.03 - 0.07| + |0.16 - 0.07|}{6} \\ &= \frac{(0.03 + 0.06 + 0.01 + 0.03 + 0.04 + 0.09)}{6} = \frac{0.26}{6} = 0.0433 = 4.3\% \end{aligned}$$

On average, the observations in the data set deviate 4.3% from the mean return of 7%.

Sample Variance and Sample Standard Deviation

Variance is defined as the average of the squared deviations around the mean; squaring the deviations takes care of the problem of differences from mean canceling each other out. The downside is that the unit values of the original data are not preserved. Standard deviation is the positive square root of the variance and is stated in units of the original variable.

For populations (ie, all n observations are included), the variance formula uses n in the denominator. For samples, the variance formula uses $n - 1$ observations in the denominator. The formula for sample variance is:

$$s^2 = \frac{(X_1 - \bar{X})^2 + (X_2 - \bar{X})^2 + \dots + (X_i - \bar{X})^2}{n - 1}$$

And the sample deviation is:

$$s = \sqrt{s^2}$$



Example 4 Calculating sample variance and sample standard deviation

Problem

Calculate the variance and standard deviation for five golfers whose scores are 67, 71, 72, 75, and 68.

Solution

$$\text{The mean score} = \frac{67 + 71 + 72 + 75 + 68}{5} = 70.6$$

$$s^2 = \frac{((67 - 70.6)^2 + (71 - 70.6)^2 + (72 - 70.6)^2 + (75 - 70.6)^2 + (68 - 70.6)^2)}{5 - 1}$$

$$s^2 = \frac{41.2}{4} = 10.3$$

$$s = \sqrt{10.3} = 3.21$$

The average for the golfers is a score of 70.6 with a standard deviation of 3.21.

Compare this to the standard deviation if the number of golfers represented the entire population:

$$\sigma^2 = \frac{41.2}{5} = 8.24, \quad \sigma = 2.87$$

Downside Deviation and Coefficient of Variation

Target downside deviation, also known as target semi-deviation, measures the dispersion of observations below the target rather than around the average. First, identify observations less than the target, and then determine the average dispersion below the target:

$$s_{\text{target}} = \left(\frac{(X_1 - B)^2 + (X_2 - B)^2 + \dots + (X_i - B)^2}{n - 1} \right)^{1/2}$$

$$X_{1,2,\dots,i} = \text{All observations that are below } B$$

$$B = \text{Target}$$

This resembles the formula for standard deviation, but the numerator is the sum of the squared differences between each observation and the target, instead of each observation and the mean. This better focuses on downside risk than the standard deviation, since, typically, B is less than the mean \bar{X} .



Example 5 Calculating downside deviation

Problem

An asset manager's returns for the last 8 years are shown below:

Year	Return (%)
20X0	15
20X1	-9
20X2	4
20X3	4
20X4	8
20X5	-2
20X6	0
20X7	3

Calculate the downside deviation if the target return is 3%

Solution

The returns for 20X0, 20X2, 20X3, 20X4, and 20X6 are greater than or equal to 3%, so downside deviation would not apply for those years. Instead, downside deviation applies only to 20X1, 20X5, and 20X7:

$$s^2 = \frac{(-9 - 3)^2 + (-2 - 3)^2 + (0 - 3)^2}{8 - 1}$$

$$s^2 = \frac{(-12)^2 + (-5)^2 + (-3)^2}{7} = \frac{178}{7} \approx 25.43$$

$$s = (25.43)^{1/2} \approx 5.04\%$$

Comparing the downside deviation to the standard deviation of the sample:

$$s \approx 7.06\% \text{ (calculations not shown)}$$

The downside deviation is smaller than standard deviation, since standard deviation captures *all* deviations from the mean. The downside deviation captures only the deviations between an observation and a specified target return.

Coefficient of Variation

Relative dispersion describes dispersion relative to a reference value or benchmark. One measure of relative dispersion is the coefficient of variation, which identifies risk per unit of return:

$$CV = \frac{s}{\bar{X}}$$

For returns, the CV measures the standard deviation of returns per unit of average return. The statistic is meaningless, however, if average return is negative.

Measures of a Shape of a Distribution



LOS: Interpret and evaluate measures of skewness and kurtosis to address an investment problem.


The mean and standard deviation measures may not completely describe investment returns, especially whether large deviations are likely to be positive or negative. In fact, non-normal distributions are unlikely to have a useful standard deviation measure and other measures will provide more useful information.

In a standard normal distribution, the mean, median, and mode are the same.

Skewness

A skewed distribution is nonsymmetrical and can be said to be either positively or negatively skewed:

Exhibit 2 Positive, negative, and normal skew



A positively skewed distribution has frequent small losses and a few extremely large gains. Investors tend to prefer positively skewed distributions because the average gain is greater than the median of all gains.

In contrast, a negatively skewed distribution has frequent small gains and a few extremely large losses.

The degree to which a distribution is skewed is called its *skewness* or *relative skewness*. When the number of observations is large (ie, $n > 100$), the **sample skewness** is calculated as:

$$\text{Skewness} = \left(\frac{1}{n} \right) \times \frac{(X_1 - \bar{X})^3 + (X_2 - \bar{X})^3 + \dots + (X_i - \bar{X})^3}{s^3}$$

$X_{1,2,\dots,i}$ = Observations

n = Number of observations

s = Standard deviation

Kurtosis

Kurtosis refers to the proportion of observations in the tails compared to a normal distribution. The kurtosis for a normal distribution is 3, so the formula below calculates excess kurtosis (ie, the difference from 3):

$$\text{Kurtosis} = \left[\left(\frac{1}{n} \right) \times \frac{(X_1 - \bar{X})^4 + (X_2 - \bar{X})^4 + \dots + (X_i - \bar{X})^4}{s^4} \right] - 3$$

$X_{1,2,\dots,i}$ = Observations

n = Number of observations

s = Standard deviation

Exhibit 3 displays a graph of different types of kurtosis and their properties:

Exhibit 3 Kurtosis

Kurtosis indicates whether a distribution is more peaked relative to a normal distribution. A normal distribution has kurtosis of 3.


Compared to normal distribution:

Leptokurtic (kurtosis > 3)

- More peaked top, fatter tails
- Heavier concentration of smaller deviations around mean, larger deviations in tails

Platykurtic (kurtosis < 3)

- Less peaked top, thinner tails
- Lighter concentration of smaller deviations around mean, larger deviations in tails



Most equity return series are fat-tailed (leptokurtic), meaning that adjustments should be made to avoid having unexpectedly large gains and losses in a security or a portfolio. Distributions may have both skewness and kurtosis.


Correlation Between Two Variables



LOS: Interpret correlation between two variables to address an investment problem.

Scatterplots

Scatterplots are a means of visually observing the relationship between two variables. In finance, analysts often use scatterplots to chart the returns of two assets.



In addition to possible relationships (and patterns) between variables, scatterplots can provide other valuable information:

- Visual evidence of the range of observations (ie, highest and lowest)
- Obvious outliers

Covariance and Correlation

Covariance shows how two variables move together. The formula for sample covariance is:

$$\text{Cov}_{XY} = \sum_{i=1}^n \frac{(X_i - \bar{X}) \times (Y_i - \bar{Y})}{n - 3}$$

The covariance of X and Y may also be written as, but this leads to confusion with standard deviation. Using $n - 1$ in the denominator ensures that the sample variance is an unbiased estimator of the population variance. Note that this formula could also be used to find variance for a single variable, since the covariance of any variable with itself is its variance.

Positive covariance occurs when the paired variables X and Y move in the same direction; that is, the data pairs are both above their mean during the same periods. Negative covariance occurs when one of the paired variables is above its mean and the other is below its mean.

Positive covariance means that the two variables in general move together; if one goes up, so does the other. Negative covariance means that the movements of the two are inverse to each other. However, covariance can range from positive to negative infinity, so covariance gives information about the direction but not the strength of the relationship.

Correlation, a normalized version of covariance, is useful for conveying the strength of a relationship. The formula for correlation, ρ_{XY} , is:

$$\rho_{XY} = \frac{\text{Cov}_{XY}}{s_X \times s_Y}$$

where:

s = Standard deviation for each variable X and Y


Properties of Correlation

The main properties of correlation are:

- Correlations range from -1 to $+1$.
 - -1 indicates perfect *negative* correlation.
 - $+1$ indicates perfect *positive* correlation.
- If correlation = 0 , there is no linear relationship between X and Y .


Exhibit 5 demonstrates how scatterplots can visually offer evidence of correlation:

Exhibit 5 Scatterplots for different values of correlation coefficient



Correlation coefficient should not be used to represent co-movement for nonlinear relationships, which can be identified by plotting *X* variables and *Y* variables (ie, a scatter plot), as shown in Exhibit 6:

Exhibit 6 Variables with nonlinear association



©CFA Institute



Example 6 Correlation matrix

Correlations between pairs of various funds appear in the shaded areas of the following exhibit. The correlation of each fund’s return with itself appears in the appropriately matching row and column along the diagonal. Since the order of stating differences does not lead to a different result, the area below the diagonal of a correlation matrix would have the same information as the shaded area and is often left blank.

	Fund A	Fund B	Fund C
Fund A	1.0000	0.9436	-0.2454
Fund B		1.0000	-0.4586
Fund C			1.0000

The correlation between Fund A and Fund B is nearly +1, indicating a strongly positive relationship between Funds A and B. When Fund A’s return is above its mean, Fund B’s return is also very likely to be above its mean. Whereas Funds A and B move together, both funds have a negative correlation coefficient with Fund C. Fund B’s relationship is fairly pronounced, whereas Fund A’s relationship has only mildly negative correlation.

Limitations of Correlation Analysis

Although correlation analysis can be useful for identifying patterns or trends, there are some limitations to its usefulness:

- There must be a linear relationship between the variables for correlation to have a significance. Although Exhibit 6 demonstrates a strong nonlinear relationship between *X* and *Y*, there is no meaningful correlation between them.
- Correlation can be unreliable when there are outliers in one or both variables.

- Correlation does not imply causation. **Spurious correlation** refers to falsely assigning “cause and effect” to correlations:
 - that occur by chance within a given data set.
 - resulting from calculating each of two variables with a third variable.
 - that arise not from a direct relationship between two variables, but instead from their relationship with a third variable.

