

BFC3241

INVESTMENTS

Study Guide · Expanded Edition

Plain-English explanations · worked examples · common traps
Topics 1–11

Monash University

Topic 0 — Prerequisite Knowledge (Recap)

A quick recap of the finance fundamentals this unit assumes you already know.

1. Why we discount cash flows

A dollar today is worth more than a dollar in the future, so cash flows arriving at different times can't be compared directly. Discounting converts every future cash flow back to its value today, bringing them all to a common point so they can be added and compared.

2. Present value, future value and compounding

Growing a sum forward is compounding; bringing it back is discounting. When interest compounds more often than yearly, use the rate per period and the number of periods (e.g. 6% p.a. compounding semi-annually = 3% over 20 half-years for 10 years).

$$FV = PV(1 + r)^n \quad PV = \frac{FV}{(1 + r)^n}$$

where PV = present value (value today); FV = future value; r = interest rate per period; n = number of periods.

The effective annual rate (EAR) converts a quoted rate that compounds m times a year into a single annual figure, so rates with different compounding can be compared:

$$EAR = \left(1 + \frac{APR}{m}\right)^m - 1$$

where EAR = effective annual rate; APR = quoted (nominal) annual rate; m = number of compounding periods per year.

Examples: \$200,000 at 6% p.a. compounding semi-annually for 10 years → $200,000(1.03)^{20} = \$361,222.25$. A rate of 8% p.a. compounding monthly → $EAR = (1 + 0.08/12)^{12} - 1 \approx 8.30\%$.

3. Valuing a stream of cash flows

The present value of a stream is simply the sum of the present values of its individual cash flows. This holds because of the Law of One Price — if the bundle were priced differently from its parts, arbitrageurs would buy the cheaper and sell the dearer until the prices converged. Common patterns:

- Annuity — equal cash flows at the end of each period for n periods (e.g. fixed mortgage payments).
- Annuity due — the same, but paid at the start of each period.
- Perpetuity — equal cash flows that continue forever.
- Growing perpetuity — cash flows that grow at a constant rate g forever.

$$PV_{\text{perpetuity}} = \frac{C}{r} \quad PV_{\text{growing}} = \frac{C}{r - g}$$

$$PV_{\text{annuity}} = \frac{C}{r} \left[1 - \frac{1}{(1 + r)^n} \right]$$

where C = the constant periodic cash flow; r = discount rate per period; g = constant growth rate of the cash flow; n = number of periods.

4. Rate of return

An investment's return is the income plus the change in value over the period, divided by the amount invested. For a share it splits into a capital gains yield and a dividend yield:

$$r = \frac{P_1 - P_0 + D_1}{P_0} = \underbrace{\frac{P_1 - P_0}{P_0}}_{\text{capital gains yield}} + \underbrace{\frac{D_1}{P_0}}_{\text{dividend yield}}$$

where P_0 = price at the start; P_1 = price at the end; D_1 = dividend (or income) received over the period.

5. The weighted average cost of capital (WACC)

The WACC is a weighted average of the costs of the different sources of finance a firm uses (debt and equity), each weighted by its share of total firm value. It is often used as the project discount rate because a single, consistent rate avoids inconsistencies between analysts and limits the gaming of rates to favour pet projects — assuming the project's risk is similar to the firm's overall risk.

6. Portfolio weights

A portfolio weight is the value invested in an asset divided by the total portfolio value (the weights sum to 1):

$$w_i = \frac{\text{value invested in asset } i}{\text{total portfolio value}}$$

where w_i = weight of asset i in the portfolio; the numerator is the dollar value held in asset i .

Example: 135 shares of A at \$47 and 105 of B at \$41 give a \$10,650 portfolio, so $w_A = 6,345/10,650 = 0.5958$ and $w_B = 4,305/10,650 = 0.4042$.

7. The risk–return relationship

Investors demand higher expected returns to bear higher risk. Faced with a choice, they prefer the highest return for a given level of risk, or the lowest risk for a given level of return.

8. Diversification

Holding two or more assets whose values do not always move together reduces risk, because their movements partly cancel out. This is the foundation for portfolio theory later in the unit.

9. The three forms of market efficiency

- Weak form — prices reflect all information in past prices, so technical analysis cannot earn abnormal returns.
- Semi-strong form — prices reflect all publicly available information, so fundamental analysis of public data cannot; only private information could.
- Strong form — prices reflect all information, public and private, so not even insiders can earn abnormal returns.

Contents

Topic 0 — Prerequisite Knowledge (Recap)	2
Topic 1 — How Markets Work	5
Topic 2 — Portfolio Theory	11
Topic 3 — Asset Pricing Models	16
Topic 4 — Equity Analysis I: Valuation	20
Topic 5 — Equity Analysis II: The Big Picture	26
Topic 6 — Market Efficiency	29
Topic 7 — Behavioural Finance	34
Topic 8 — Investment Vehicles I: Passive Investments	38
Topic 9 — Investment Vehicles II: Active Investments	43
Topic 10 — Portfolio Management: Performance Evaluation	47
Topic 11 — CFA Code of Ethics and Professional Standards	52

Topic 1 — How Markets Work

What investing is really about

Before any stock-picking, the biggest decision an investor makes is asset allocation — how to split money across the broad asset classes (cash, bonds, shares). This matters more than which individual securities you choose: study after study finds that the mix of asset classes explains the large majority of a portfolio’s return over time, far more than individual security selection. The intuition is simple — deciding “70% shares, 30% bonds” shapes your outcome much more than deciding which shares fill that 70%. When you set that mix you are balancing three things: your return requirement, your risk tolerance, and your time horizon (a 25-year-old saving for retirement can ride out crashes a 64-year-old cannot).

Primary vs secondary markets

A primary market transaction is when a security is created and sold for the first time — the classic case is an IPO. The key feature: the money goes to the company, because the company is the one issuing new shares to raise capital.

A secondary market transaction is when existing securities change hands between investors — what happens on the ASX every day. Here the money flows between investors, not to the company. The firm is not involved; you are buying someone else’s existing shares.

A clean way to hold it: the primary market is buying a new car from the manufacturer (money to the maker); the secondary market is buying that car used from its owner (money to the owner, not the maker).

How you place a trade: order types

When you trade, you choose how the order executes. The two core types trade off price certainty against execution certainty:

- **Market order** — “fill me now at whatever the best available price is.” You are guaranteed to trade, but not at a guaranteed price. Good when getting in or out matters more than a few cents.
- **Limit order** — “only buy at or below \$X” (or “only sell at or above \$X”). You are guaranteed your price or better, but not guaranteed to trade at all — if the market never reaches your limit, nothing happens.

Two more are conditional orders built to manage a position you already hold:

- **Stop-loss order** — automatically sells once the price falls to a set level, capping your downside. If you buy at \$100 and set a stop at \$90, you have decided \$10 is the most you will stomach losing.
- **Take-profit order** — automatically sells once the price rises to a set level, locking in a gain before it can evaporate.

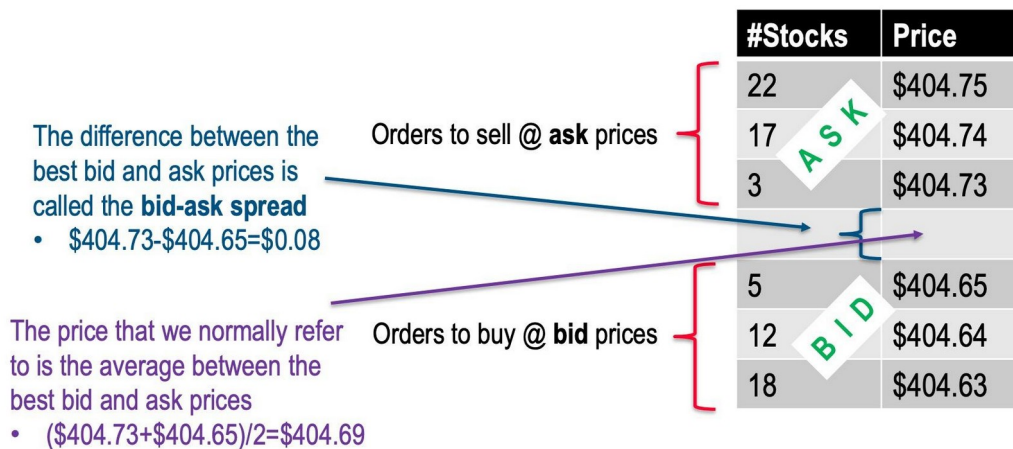
Example

You want to buy a stock currently at \$404.70. A market order fills immediately at the best ask. A limit order at \$404.65 only fills if a seller drops to your price — if the stock keeps climbing, your order just sits there unfilled.

The cost of trading

Trading is never free. Costs come in three forms, and two of them are hidden:

- **Commission (brokerage)** — the explicit, visible fee your broker charges.
- **Bid-ask spread** — an implicit cost. At any moment there is a higher price you can buy at (the ask) and a lower price you can sell at (the bid). The gap between them is a cost you pay just by transacting — you effectively lose the spread the instant you buy then sell.



- **Price impact** — for large orders, the act of trading moves the price against you. If only a few shares are available at the best price, your next share has to reach up to the next, worse price. Big investors manage this by slicing large orders into smaller chunks over time, or by using dark pools — private venues where large blocks trade without broadcasting the order to the whole market.

Trading on margin (borrowing to buy)

Buying on margin means borrowing money from your broker to buy more shares than your own cash would allow — it is leverage. Two terms define it:

- **Initial margin** — the share of the purchase you fund with your own money at the outset (the rest is the broker's loan).
- **Maintenance margin** — the minimum equity percentage you must keep in the account afterward. If your equity falls below it, you get a margin call: top up cash or the broker sells your position to recover the loan.

The reason margin matters is that it amplifies both gains and losses. You are controlling a bigger position with the same cash, so a given percentage move in the stock translates into a larger percentage move in your equity.

Example

You buy 1,000 shares at \$70 (total \$70,000) with 50% initial margin, so you put in \$35,000 and borrow \$35,000. Maintenance margin is 40%.

If the price falls to \$60, the stock is worth \$60,000, the loan is unchanged at \$35,000, so your equity is \$25,000. Your margin is now $25,000 / 60,000 = 41.67\%$ — still above the 40% floor, so no margin call yet.

At what price does the call hit? Solve for the price P where $\text{equity} / \text{value} = 40\%$: $(1,000P - 35,000) / 1,000P = 0.40$, which gives $P = \$58.33$. Below that, the broker calls.

Assets		Liabilities & Equity	
Value of stock	\$70,000	Loan	\$35,000
		Equity	\$35,000
Total	\$70,000	Total	\$70,000

New share price \$60

Assets		Liabilities & Equity	
Value of stock	\$60,000	Loan	\$35,000
		Equity	\$25,000
Total	\$60,000	Total	\$60,000

A margin call can feed a market crash: falling prices force margin sales, those sales push prices down further, triggering yet more calls — a self-reinforcing spiral.

Short selling (profiting when prices fall)

Short selling flips the usual order around. You borrow shares, sell them now at today's price, and aim to buy them back later at a lower price to return to the lender — pocketing the difference. You are betting the price falls.

The asymmetry is the thing to understand: when you buy a stock normally, the most you can lose is 100% (it goes to zero) but your upside is unlimited. Shorting is the mirror image — your maximum gain is capped (the price can only fall to zero) while your potential loss is unlimited, because a stock can rise without limit. That is why shorting is risky.

Example

Stock at \$60, you short 100 shares, so you receive \$6,000 now. Initial margin 50%. MMR 30%

Assets		Liabilities & Equity	
Cash	\$6,000	Short position	\$6,000
T-bills	\$3,000	Equity	\$3,000
Total	\$9,000	Total	\$9,000

New Share price \$65

Assets		Liabilities & Equity	
Cash	\$6,000	Short position	\$6,500
T-bills	\$3,000	Equity	\$2,500
Total	\$9,000	Total	\$9,000

New margin = $2500/6500 = 38.46\%$

Margin has dropped from 50% to 38.46% > 30% MMR

c. How far must the share price rise to receive a margin call

$$0.3 = (9000 - x)/x$$

$$X = 6923 = \$69.23$$

Measuring return

The basic building block is the holding-period return (HPR) — the total return over the time you held an asset:

$$\text{HPR} = \frac{\text{ending price} - \text{beginning price} + \text{income}}{\text{beginning price}}$$

where income = any dividends or interest received while the asset was held.

It has two parts: the capital gain (price change) and the income yield (dividends or interest). To compare investments held for different lengths of time, you annualise — and the correct way compounds the return rather than just multiplying, because returns earn returns.

Two ways to average a return series: the arithmetic average (simple mean — best estimate of a single future period) and the geometric average (compounded — the true growth rate of \$1 over the whole horizon):

$$\text{AAR} = \frac{1}{n} \sum_{t=1}^n \text{HPR}_t$$

where HPR_t = holding-period return in period t ; n = number of periods. Arithmetic (simple) average.

$$\text{GAR} = \left[\prod_{t=1}^n (1 + \text{HPR}_t) \right]^{1/n} - 1$$

where HPR_t = holding-period return in period t ; n = number of periods. Geometric (compounded) average.

Expected return and risk

Looking forward, we do not know which outcome will occur, so we work with probabilities. The expected return is the probability-weighted average across possible scenarios:

$$E(r) = \sum_{s=1}^S p(s) r(s)$$

where $p(s)$ = probability of scenario s ; $r(s)$ = return in scenario s ; the sum runs over all S possible scenarios.

Risk is how much actual outcomes scatter around that expected value. We measure it with variance — the probability-weighted average of squared deviations from the mean — and then take its square root to get the standard deviation, which lands back in the same units as returns (%). Standard deviation is the workhorse risk measure: loosely, it is the “typical distance” an outcome lands from the average.

$$\sigma^2 = \sum_{s=1}^S p(s) [r(s) - E(r)]^2$$

where $p(s)$ = probability of scenario s ; $r(s)$ = return in scenario s ; $E(r)$ = expected return.

Example

A stock will return +20% with probability 0.5 and -5% with probability 0.5. Expected return = $0.5(20\%) + 0.5(-5\%) = 7.5\%$. The outcomes sit well away from 7.5%, so the standard deviation (the risk) is high relative to a stock whose outcomes cluster near its mean.

Portfolios: weights and return

A portfolio weight is simply the fraction of your total money in a given asset: weight = value in that asset / total portfolio value (the weights add to 100%). A portfolio’s expected return is just the weighted average of the expected returns of its holdings:

$$E(r_{\text{portfolio}}) = w_1 E(r_1) + w_2 E(r_2) + \dots$$

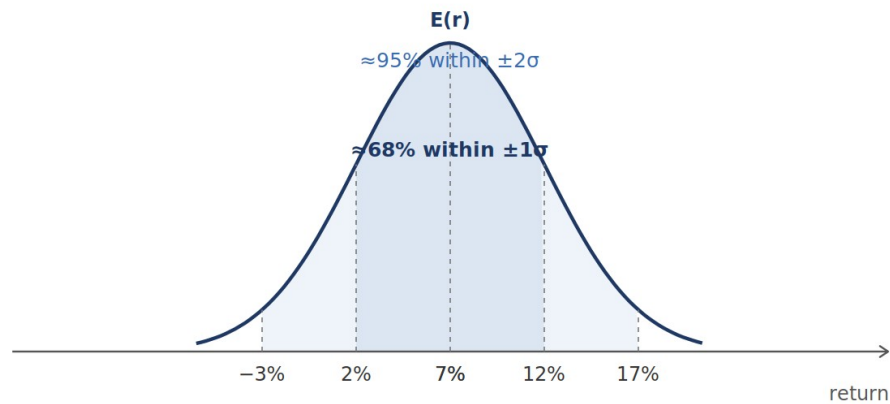
where w_i = fraction of the portfolio held in asset i (weights sum to 1); $E(r_i)$ = expected return of asset i .

Risk, importantly, is not a simple weighted average — that is the whole point of Topic 2.

Why we lean on the normal distribution

Finance often assumes returns follow a normal distribution (the bell curve) because it is fully described by just two numbers: the mean (expected return) and the standard deviation (risk). That is enormously convenient — it is exactly why standard deviation works as the risk measure. Useful rule of thumb: about 68% of outcomes fall within one standard deviation of the mean, and about 95% within two. So a portfolio with a 7% expected return and 10% standard deviation lands between -3% and +17% roughly two years in three.

Distribution of returns



Common trap Real returns are not perfectly normal — they have “fat tails” (extreme crashes and rallies happen more often than the bell curve predicts). The normal assumption is a simplification, and exam questions sometimes probe whether you know its limits.

A note on short-selling mechanics

When you short-sell, the sale proceeds are not yours to spend — the broker keeps them as a pledge (collateral) against the borrowed shares. You earn a small interest rate on those held proceeds, but you also pay the broker a stock-borrow fee for lending you the shares, and you must post your own margin on top. If the price rises far enough, you face a margin call just as a leveraged long position does.

Topic 2 — Portfolio Theory

The big idea

Topic 1 measured the risk and return of single assets. Portfolio theory asks the next question: when you combine assets, what happens to risk and return together? The headline result is that risk does not simply average out the way return does — combine assets that do not move in lockstep and you can lower risk without giving up return. That “free lunch” is diversification, and the whole week builds on it.

Risky and risk-free assets

A risk-free asset (think a short-term government bill) has a known, certain return — zero standard deviation. A risky asset has an uncertain return with positive standard deviation. The first building block is a portfolio that mixes one risky asset (or portfolio) with the risk-free asset.

The capital allocation line (CAL)

If you split money between the risk-free asset and a risky portfolio, every possible mix plots on a straight line in return–risk space called the capital allocation line. Its intercept is the risk-free rate (0% in the risky asset) and it rises towards the risky portfolio. Because the risk-free asset has no risk and no co-movement, the line is straight: both expected return and standard deviation scale linearly with how much you put in the risky asset.

The slope of the CAL is the Sharpe ratio — extra return earned per unit of risk:

$$\text{Sharpe ratio} = \frac{E(r_p) - r_f}{\sigma_p}$$

where $E(r_p)$ = portfolio expected return; r_f = risk-free rate; σ_p = portfolio standard deviation (total risk).

A steeper CAL is strictly better: more return for the same risk. So the best risky portfolio to hold is the one that produces the steepest possible CAL.

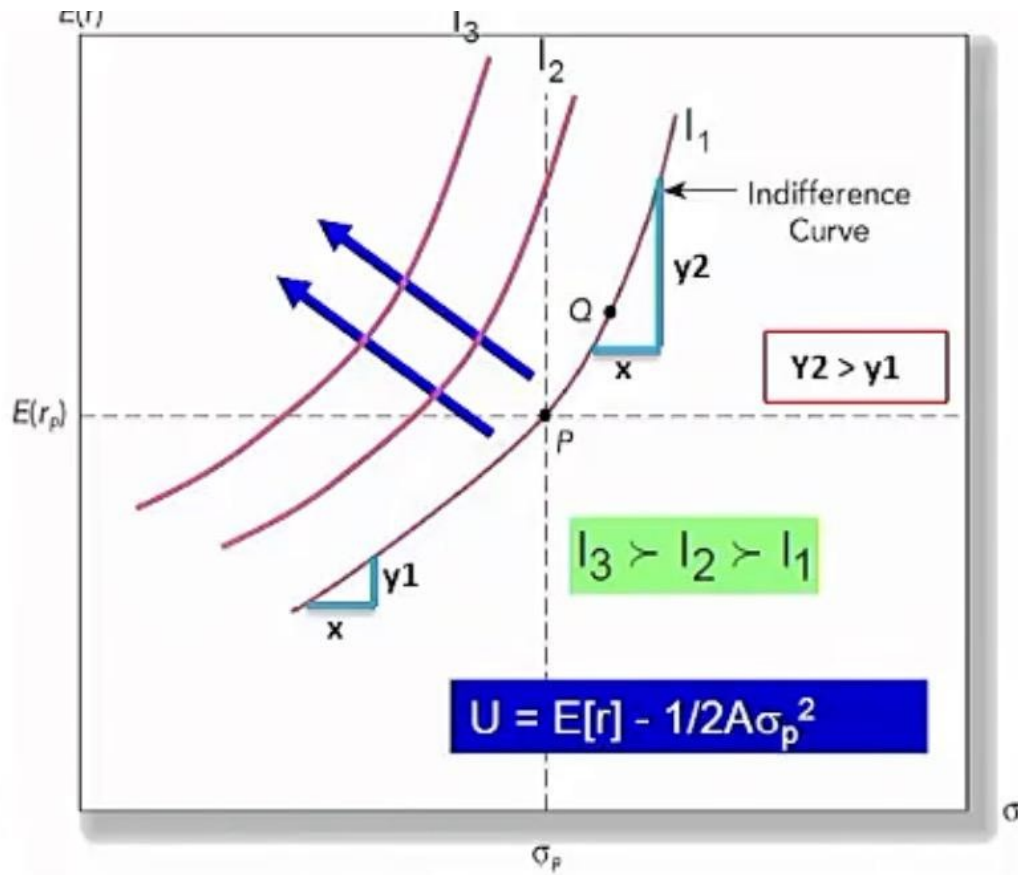
Risk aversion and where you sit on the line

The CAL is the menu; risk aversion decides your seat. We score a portfolio with a utility function:

$$U = E(r_c) - \frac{1}{2} A \sigma_c^2$$

where $E(r_c)$, σ_c = expected return and standard deviation of the complete portfolio; A = the investor’s risk-aversion coefficient.

A is the investor’s risk-aversion coefficient — the higher it is, the more heavily variance is penalised. A very risk-averse investor (high A) holds more of the risk-free asset and sits low on the CAL; a risk-tolerant investor sits higher, even borrowing to hold more than 100% of the risky portfolio. Graphically, you pick the point where your indifference curve is just tangent to the CAL.



Example

Risk-free rate 4%. The risky portfolio offers 12% with a standard deviation of 20%. Put 60% in the risky portfolio and 40% in the risk-free asset. Expected return = 0.4(4%) + 0.6(12%) = 8.8%. Standard deviation = 0.6 × 20% = 12% (only the risky part contributes). Every other split slides you up or down the same straight line.

Two risky assets: return, variance, and the role of correlation

Now drop the risk-free asset and combine two risky assets. Expected return is still the simple weighted average:

$$E(r_p) = w_1 E(r_1) + w_2 E(r_2)$$

where w_1, w_2 = weights in the two assets ($w_1 + w_2 = 1$); $E(r_1), E(r_2)$ = their expected returns.

But variance is not — and this is the crux of the whole subject:

$$\sigma_p^2 = w_1^2 \sigma_1^2 + w_2^2 \sigma_2^2 + 2 w_1 w_2 \rho_{12} \sigma_1 \sigma_2$$

where σ_1, σ_2 = standard deviations of the two assets; ρ_{12} = correlation between them; w_1, w_2 = weights.

The correlation coefficient itself is the covariance scaled by the two standard deviations, which keeps it between -1 and +1:

$$\rho_{12} = \frac{\text{Cov}(r_1, r_2)}{\sigma_1 \sigma_2}$$