CFA Institute
Investment Foundations

A CLEAR UNDERSTANDING OF THE INDUSTRY

IS CFA INSTITUTE INVESTMENT FOUNDATIONS RIGHT FOR YOU?

Investment Foundations is a certificate program designed to give you a clear understanding of the investment management industry. Whether you're just starting a career in financial services or want to learn essential industry concepts, Investment Foundations offers an accessible solution for breaking through the complexities of the global investment industry and raising your professional profile.

THE BIG PICTURE

Investment Foundations is a comprehensive global education certificate program that provides a clear understanding of investment industry essentials. The certificate program is designed for all professional disciplines outside of investment roles, including IT, operations, accounting, administration, and marketing. There is no education or experience requirement and the exam can be taken at your convenience at available test centers around the world.

WHAT YOU LEARN

The certificate program covers the essentials of the investment management industry:

Module 1: Industry Overview

Module 2: Ethics and Regulation

Module 3: Inputs and Tools

Module 4: Investment Instruments

Module 5: Industry Structure

Module 6: Serving Client Needs

Module 7: Industry Controls

HOW WILL YOU BENEFIT

Clarity
Benefit from having a common understanding of industry structure and terminology, regardless of your job function or geographic location.

Collaboration
Work more effectively with global colleagues by understanding industry functions, building stronger relationships and raising your professional competence.

Confidence
Gain the knowledge to identify issues and the confidence to speak up. Get a better sense of your role and how you connect with the complex global industry at large.

"[CFA Institute Investment Foundations] is very relevant to the current market and I can study on the move. The program cleared up a lot of concepts for me and now I am much more comfortable speaking with clients about what is happening in the market."

MITALI BHANDARE
MORNINGSTAR,
INVESTMENT FOUNDATIONS CERTIFICATE HOLDER

"The main benefit of [CFA Institute Investment Foundations] was an ability to see a bigger picture of the finance industry and the role of our business within it."

ALEXANDER TARASOV
CITCO FUND SERVICES,
INVESTMENT FOUNDATIONS CERTIFICATE HOLDER

HOW TO FIND OUT MORE

Go to: cfa.is/InvFound

www.cfainstitute.org

CHAPTER 8

QUANTITATIVE CONCEPTS

by Michael J. Buckle, PhD, James Seaton, PhD, and Stephen Thomas, PhD

LEARNING OUTCOMES

After completing this chapter, you should be able to do the following:

- **a** Define the concept of interest;
- **b** Compare simple and compound interest;
- **c** Define present value, future value, and discount rate;
- **d** Describe how time and discount rate affect present and future values;
- **e** Explain the relevance of net present value in valuing financial investments;
- **f** Describe applications of time value of money;
- **g** Explain uses of mean, median, and mode, which are measures of frequency or central tendency;
- **h** Explain uses of range, percentile, standard deviation, and variance, which are measures of dispersion;
- **i** Describe and interpret the characteristics of a normal distribution;
- **j** Describe and interpret correlation.

INTRODUCTION

1

Knowledge of quantitative (mathematically based) concepts is extremely important to understanding the world of finance and investing. Quantitative concepts play a role in financial decisions, such as saving and borrowing, and also form the foundation for valuing investment opportunities and assessing their risks. The time value of money and descriptive statistics are two important quantitative concepts. They are not directly related to each other, but we combine them in this chapter because they are key quantitative concepts used in finance and investment.

The time value of money is useful in many walks of life: it helps savers to know how long it will take them to afford a certain item and how much they will have to put aside each week or month, it helps investors to assess whether an investment should provide a satisfactory return, and it helps companies to determine whether the profit from investing will exceed the cost.

Statistics are also used in a wide range of business and personal contexts. As you attempt to assess the large amount of personal and work-related data that are part of our everyday lives, you will probably realise that an efficient summary and description of data is helpful to make sense of it. Most people, for instance, look at summaries of weather information to make decisions about how to dress and whether to carry an umbrella or bring rain gear. Summary statistics help you understand and use information in making decisions, including financial decisions. For example, summary information about a company's or market's performance can help in investment decisions.

In short, quantitative concepts are fundamental to the investment industry. For anyone working in the industry, familiarity with the concepts described in this chapter is critical. **As always, you are *not* responsible for calculations, but the presentation of formulae and illustrative calculations may enhance your understanding.**

TIME VALUE OF MONEY

2

Valuing cash flows, which occur over different periods, is an important issue in finance. You may be concerned with how much money you will have in the future (the future value) as a result of saving or investing over time. You may want to know how much you should save in a certain amount of time to accumulate a specified amount in the future. You may want to know what your expected return is on an investment with specified cash flows at different points in time. These types of problems occur every day in investments (e.g., in buying a bond), personal finance (e.g., in arranging an automobile loan or a mortgage), and corporate finance (e.g., in evaluating whether to build a factory). These problems are known as "time value of money" problems because their solutions reflect the principle that the timing of a cash flow affects the cash flow's value.

2.1 Interest

Borrowing and lending are transactions with cash flow consequences. Someone who needs money borrows it from someone who does not need it in the present (a saver) and is willing to lend it. In the present, the borrower has money and the lender has given up money. In the future, the borrower will give up money to pay back the lender; the lender will receive money as repayment from the borrower in the form of **interest**, as shown below. The lender will also receive back the money lent to the borrower. The money originally borrowed, which interest is calculated on, is called the **principal**. Interest can be defined as payment for the use of borrowed money.

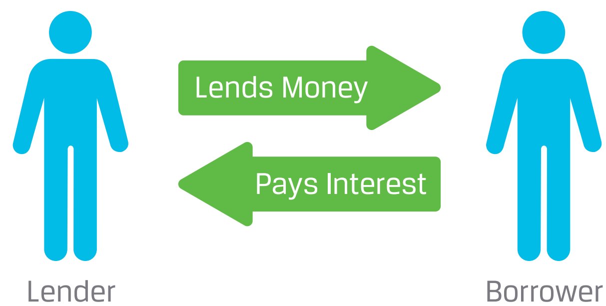

Interest is all about timing: someone needs money now while someone else is willing and able to give up money now, but at a price. The borrower pays a price for not being able to wait to have money and to compensate the lender for giving up potential current consumption or other investment opportunities; that price is interest. Interest is paid by a borrower and earned by the lender to compensate the lender for opportunity cost and risk. **Opportunity cost**, in general, is the value of alternative opportunities that have been given up by the lender, including lending to others, investing elsewhere, or simply spending the money. Opportunity cost can also be seen as compensation for deferring consumption. Lending delays consumption by the term of the loan (the time over which the loan is repaid). The longer the consumption is deferred, the more compensation (higher interest) the lender will demand.

The lender also bears risks, such as the risk of not getting the money back if the borrower defaults (fails to make a promised payment). The riskier the borrower or the less certain the borrower's ability to repay the loan, the higher the level of interest demanded by the lender. Another risk is that as a result of inflation (an increase in prices of goods and services), the money received may not be worth as much as expected. In other words, a lender's purchasing power may decline even if the money is repaid as promised. The greater the expected inflation, the higher the level of interest demanded by the lender.

From the borrower's perspective, interest is the cost of having access to money that they would not otherwise have. An interest rate is determined by two factors: opportunity cost and risk. Even if a loan is viewed as riskless (zero likelihood of default), there still has to be compensation for the lender's opportunity cost and for expected inflation. Exhibit 1 shows examples of borrowers and lenders.

Exhibit 1 Examples of Borrowers and Lenders

Borrowers and lenders can be people, companies, financial institutions, and so on. Here are some examples of borrowers and lenders that you may be familiar with.

If you get money from the bank and pay interest, ***you are the borrower*** and ***the bank is the lender.***

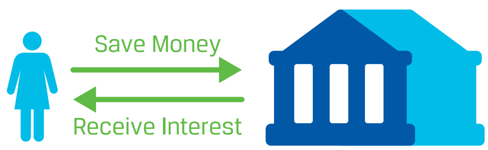

If you put money into a savings account and earn interest, ***you are a saver*** who lends your savings to the bank and ***the bank is the borrower.***

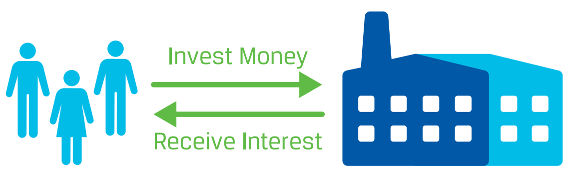

If people invest in a company and earn interest by buying bonds, ***they are the lenders*** and ***the company is the borrower***.

2.1.1 *Simple Interest*

A **simple interest rate** is the cost to the borrower or the rate of return to the lender, per period, on the original principal (the amount borrowed). Conventionally, interest rates are stated as annual rates, so the period is assumed to be one year unless stated otherwise. The cost or return is stated as a percentage rate of the original principal so the rates can then be compared, regardless of the amount of principal they apply to. For example, a loan with a 5% interest rate is more expensive to the borrower than a loan with a 3% interest rate. Similarly, a loan with a 5% interest rate provides a higher promised return to the lender than a loan with a 3% interest rate.

The actual amount of interest earned or paid depends on the simple interest rate, the amount of principal lent or borrowed, and the number of periods over which it is lent or borrowed. We can show this mathematically as follows:

$$\text{Simple interest} = \text{Simple interest rate} \times \text{Principal} \times \text{Number of periods}$$

If you put money in a bank account and the bank offers a simple interest rate of 10% per annum (or annually), then for every £100 you put in, you (as a lender to the bank) will receive £10 in the course of the year (assume at year end to simplify calculations):

$$\text{Interest} = 0.10 \times £100 \times 1 = £10$$

If your money is left in the bank for two years, the interest paid will be £20:

$$\text{Interest} = 0.10 \times £100 \times 2 = £20$$

Simple interest is not reinvested and is applied only to the original principal, as shown in Exhibit 2.

Exhibit 2 Simple Interest of 10% on £100 Original Principal

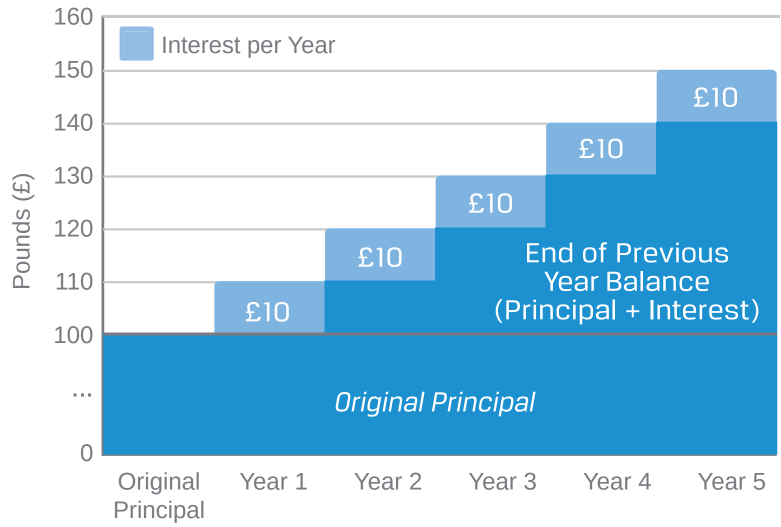

If the interest earned is added to the original principal, the relationship between the original principal and its future value with simple interest can be described as follows:

Future value = Original principal × [1 + (Simple interest rate × Number of periods)]

To extend our deposit example: £100 × [1 + (0.10 × 2)] = £100 × (1.20) = £120. The value at the end of two years is £120.

2.1.2 *Compound Interest*

Interest compounds when it is added to the original principal. **Compound interest** is often referred to as "interest on interest". As opposed to simple interest, interest is assumed to be reinvested so future interest is earned on principal and reinvested interest, not just on the original principal.

If a deposit of £100 is made and earns 10% and the money is reinvested (remains on deposit), then additional interest is earned in the course of the second year on the £10 of interest earned in the first year. The interest is being compounded. Total interest after two years will now be £21; £10 (= £100 × 0.10) for the first year, plus £11 (= £110 × 0.10) for the second year. The second year's interest is calculated on the original £100 principal plus the first year's interest of £10. As shown in Exhibit 3, the total interest after two years is £21 rather than £20 as in the case of simple interest shown in Exhibit 2.

Exhibit 3 Compound Interest of 10% on £100 Original Principal

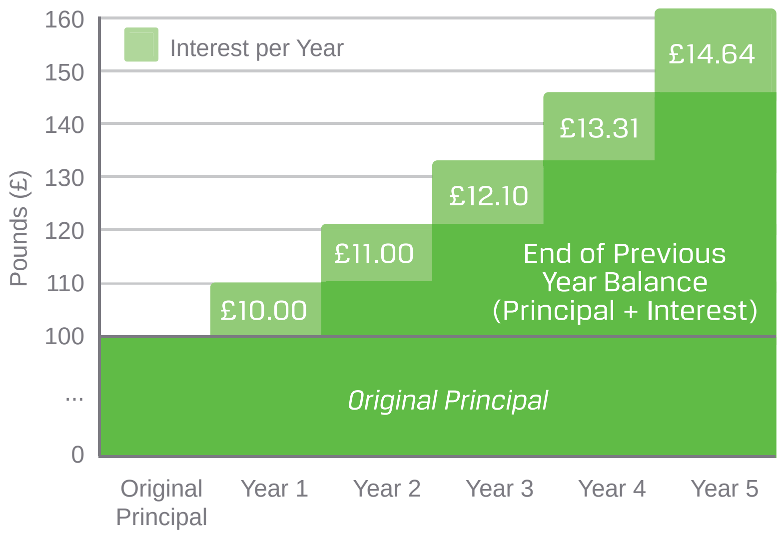

The relationship between the original principal and its future value when interest is compounded can be described as follows:

$$\text{Future value} = \text{Original principal} \times (1 + \text{Simple interest rate})^{\text{Number of periods}}$$

In the deposit example, £100 × $(1 + 0.10)^2$ = £100 × $(1.10)^2$ = £121. With compounding, the value at the end of two years is £121.

2.1.3 *Comparing Simple Interest and Compound Interest*

Compound interest is extremely powerful for savers; reinvesting the interest earned on investments is a way of growing savings. Somebody who invests £100 at 10% for two years will end up with £1 more by reinvesting the interest (£121) than with simple interest (£120). This amount may not look very impressive, but over a longer time period, say 20 years, £100 invested at 10% for 20 years becomes £300 with simple interest {£100 × [1 + (0.10 × 20)] = £100 × 3 = £300} but £673 with compound interest [£100 × $(1 + 0.10)^{20}$ = £100 × $(1.10)^{20}$ = £673]. This concept is illustrated in Exhibit 4.

Exhibit 4 Effects on Savings of Simple and Compound Interest

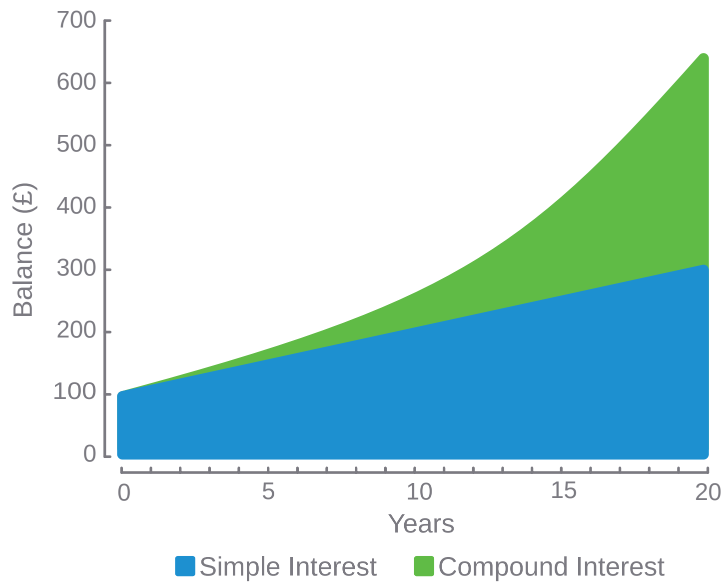

2.1.4 *Annual Percentage Rate and Effective Annual Rate*

Unless stated otherwise, interest rates are stated as annual rates. The rate quoted is often the **annual percentage rate** (APR), which is a simple interest rate that does not involve compounding. Another widely used rate is the **effective annual rate** (EAR). This rate involves annualising, through compounding, a rate that is paid more than once a year—usually monthly, quarterly, or semi-annually. The following equation shows how to determine the EAR given the APR.

$$\text{EAR} = \left[\left(1 + \frac{\text{APR}}{\text{Number of periods per year}}\right)^{\text{Number of periods per year}}\right] - 1$$

Example 1 shows a few types of financial products and their simple interest rates (APRs) and their compound rates (EARs).

EXAMPLE 1. SIMPLE AND COMPOUND INTEREST RATES

A credit card charges interest at an APR of 15.24%, compounded daily. A bank pays 0.2% monthly on the average amount on deposit over the month. A loan is made with a 6.0% annual rate, compounded quarterly. The following table shows what the expected annual rate is for each of these situations. The rate is higher than the APR because of compounding.

	Simple Interest Rate or APR	Compound Interest Rate or EAR
Credit card	15.24%	$16.46\% = \left[\left(1 + \frac{0.1524}{365}\right)^{365}\right] - 1$
Bank deposit	2.4% (= 0.2% × 12)	$2.43\% = \left[\left(1 + \frac{0.024}{12}\right)^{12}\right] - 1$
Loan	6.0%	$6.14\% = \left[\left(1 + \frac{0.06}{4}\right)^{4}\right] - 1$

As can be seen in Example 1, in general, whenever an interest rate compounds more often than annually, the EAR is greater than the APR. In other words, more frequent compounding leads to a higher EAR.

2.2 Present Value and Future Value

Two basic time value of money problems are finding the value of a set of cash flows now (present value) and the value as of a point of time in the future (future value).

2.2.1 *Present Value and Future Value*

If you are offered £1 today or £1 in a year's time, which would you choose? Most people say £1 today because it gives them the choice of whether to spend or invest the money today and avoid the risk of never getting it at all. The £1 to be received in the future is worth less than £1 received today. The £1 to be received in the future is today worth £1 minus the opportunity cost and the risk of being without it for one year. The present value is obtained by discounting the future cash flow by the interest rate. The rate of interest in this context can be called the **discount rate**.

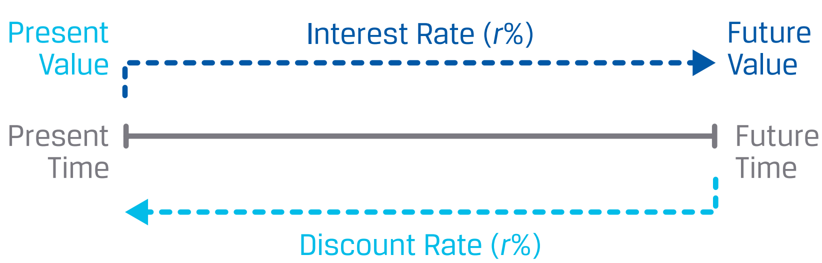

Time affects the value of money because delay creates opportunity costs and risk. If you earn a return of r% for waiting one year, £1 × (1 + r%) is the **future value** after one year of £1 invested today. Put another way, £1 is the **present value** of £1 × (1 + r%) received in a year's time.

A saver may want to know how much money is needed today to produce a certain sum in the future given the rate of interest, *r*. In the example in Exhibit 3, today's value is £100 and the interest rate is 10%, so the future value after two years is £100 × (1 + 0.10)^2 = £121. The present value—the equivalent value today—of £121 in two years, given that the annual interest rate is 10%, is £100.

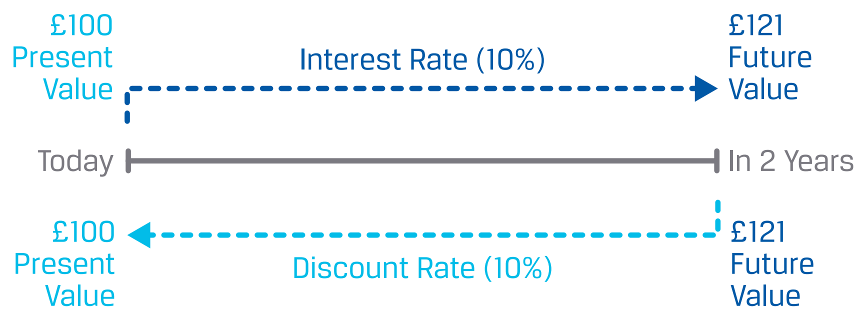

Before you can calculate present or future values, you must know the appropriate interest or discount rates to use. The rate will usually depend on the overall level of interest rates in the economy, the opportunity cost, and the riskiness of the investments under consideration. The following equations generalise the calculation of future and present values:

$$\text{Future value} = \text{Present value} \times (1 + \text{Interest rate})^{\text{Number of periods}}$$

$$\text{Present value} = \frac{\text{Future value}}{(1 + \text{Discount rate})^{\text{Number of periods}}}$$

Note that the interest and discount rates are the same percentage rates, but the terminology varies based on context. Calculating present values allows investors and analysts to translate cash flows of different amounts and at different points in the future into sums in the present that can be compared with each other. Likewise, the cash flows can be translated into the values they would be equivalent to at a common future point.

Example 2 compares two investments with the same initial outflow (investment) but with different future cash inflows at different points in time.

EXAMPLE 2. COMPARING INVESTMENTS

1 You are choosing between two investments of equal risk. You believe that given the risk, the appropriate discount rate to use is 9%. Your initial investment (outflow) for each is £500. One investment is expected to pay out £1,000 three years from now; the other investment is expected to pay out £1,350 five years from now. To choose between the two investments, you must compare the value of each investment at the same point in time.

Present value of £1,000 in three years discounted at 9%

$$= \frac{£1{,}000}{(1.09)^3} = \frac{£1{,}000}{1.295} = £772.18$$

Present value of £1,350 in five years discounted at 9%

$$= \frac{£1{,}350}{(1.09)^5} = \frac{£1{,}350}{1.5386} = £877.41$$

As you can see, the investment with a payout of £1,350 five years from now is worth more in present value terms, so it is the better investment.

2 You are choosing between the same two investments but you have reassessed their risks. You now consider the five-year investment to be more risky than the first and estimate that a 15% return is required to justify making this investment.

Present value of £1,350 in five years discounted at 15%

$$= \frac{£1{,}350}{(1.15)^5} = \frac{£1{,}350}{2.0114} = £671.19$$

The investment paying £1,000 in three years (discounted at 9%) is, in this case, preferable to the investment paying £1,350 in five years (discounted at 15%) in present value terms. Its present value of £771.28 is higher than the present value of £671.19 on the five-year investment.

Example 2 shows three elements that must be considered when comparing investments:

- the cash flows each investment will generate in the future,
- the timing of these cash flows, and
- the risk associated with each investment, which is reflected in the discount rate.

Present value considers the joint effect of these three elements and provides an effective way of comparing investments with different risks that have different future cash flows at different points in time.

2.2.2 *Net Present Value*

Present value is appropriate for comparing investments when the initial outflow for each investment is the same, as in Example 2. But investments may not have the same initial cash outflow, and outflows may occur at times other than time zero (the time of the initial outflow). The **net present value** (NPV) of an investment is the present value of future cash flows or returns minus the present value of the cost of the investment (which often, but not always, occurs only in the initial period). Using NPV rather than present value to evaluate investments is especially important when the investments have different initial costs. Example 3 below illustrates this.

EXAMPLE 3. COMPARING INVESTMENTS USING NET PRESENT VALUE

The NPV of the investment in Example 2 that is paying £1,350 in five years (discounted at 15%) if it initially cost £500 is:

£671.19 – £500.00 = £171.19

The NPV of the investment paying £1,000 in three years discounted at 9% if it initially cost £700 is:

£772.18 – £700 = £72.18.

This amount is less than £171.19, making the investment paying £1,350 in five years discounted at 15% worth more in present value terms. This conclusion differs from that reached when present value only was used.

If costs were to occur at times different from time zero, then they would also be discounted back to time zero for the purposes of comparison and calculation of the NPV. If the NPV is zero or greater, the investment is earning at least the discount rate. An NPV of less than zero indicates that the investment should not be made.

Calculating the NPV allows an investor to compare different investments using their projected cash flows and costs. The concepts of present value and net present value have widespread applications in the valuation of financial assets and products. For example, equities may pay dividends and/or be sold in the future, bonds may pay interest and principal in the future, and insurance may lead to future payouts.

Estimating values by using cash flows is also important to companies considering a range of investment opportunities. For example, should the sales team be supplied with tablets or laptops, or should the company open a new office in Asia or carry on visiting from the company's European headquarters? In order to choose, decision makers estimate the expected future cash flows of the alternatives available. The decision makers then discount the estimated cash flows by an appropriate discount rate that reflects the riskiness of these cash flows. They work out the discounted cash flows for each opportunity to estimate the value of the cash flows at the current time (the present value) and to arrive at the net present value. They then compare the net present values of all the opportunities and choose the opportunity or combination of opportunities with the largest positive net present value.

2.2.3 *Application of the Time Value of Money*

The time value of money concept can help to solve many common financial problems. If you save in a deposit account, it can tell you by how much your money will grow over a given number of years. Time value of money problems can involve both positive cash flows (inflows or savings) and negative cash flows (outflows or withdrawals). Example 4 illustrates, with two different sets of facts, how cash inflows and outflows affect future value.

EXAMPLE 4. FUTURE VALUE

1 You place £1,000 on deposit at an annual interest rate of 10% and make regular contributions of £250 at the end of each of the next two years. How much do you have in your account at the end of two years?

The initial £1,000 becomes £1,000 × $(1 + 0.10)^2$	=	£1,210
The first annual £250 payment becomes £250 × (1+ 0.10)	=	£275
The second annual payment is received at the end and earns no interest	=	£250
The total future value	=	£1,735

2 You place £1,000 on deposit and withdraw £250 at the end of the first year. The balance on deposit at the beginning of the year earns an annual interest rate of 10%. How much do you have in your account at the end of two years?

At the end of the first year, you have £1,000 × (1 + 0.10)	=	£1,100
You withdraw £250 and begin the second year with an amount	=	£850
At the end of the second year, you have £850 × (1 + 0.10)	=	£935

Time value of money can also help determine the value of a financial instrument. It can help you work out the value of an annuity or how long it will take to pay off the mortgage on your home.

2.2.3.1 Present Value and the Valuation of Financial Instruments People invest in financial products and instruments because they expect to get future benefits in the form of future cash flows. These cash flows can be in the form of income, such as dividends and interest, from the repayment of an amount lent, or from selling the financial product or instrument to someone else. An investor is exchanging a sum of money today for future cash flows, and some of these cash flows are more uncertain than others. The value (amount exchanged) today of a financial product should equal the value of its expected future cash flows. This concept is shown in Example 5.

EXAMPLE 5. VALUE OF A LOAN

Consider the example of a simple loan that was made three years ago. Two years from today, the loan will mature and the borrower should repay the principal value of the loan, which is £100. The investor who buys (or owns) this loan should also receive from the borrower two annual interest payments at the originally promised interest rate of 8%. The interest payments will be £8 (= 8% × £100), with the first interest payment received a year from now and the second two years from now.

How much would an investor pay today to secure these two years of cash flow if the appropriate discount rate is 10% (i.e. $r = 0.10$)? Note that the rate used for discounting the future cash flows should reflect the risk of the investment and interest rates in the market. In practice, it is unlikely that the discount rate will be equal to the loan's originally promised interest rate because the risk of the investment and interest rates in the market may change over time.

The first year's interest payment is worth $\frac{£8}{1.10^1} = £7.27$.

The second year's interest payment is worth $\frac{£8}{1.10^2} = £6.61$.

The repayment of the loan's principal value in two years is worth $\frac{£100}{1.10^2} = £82.64$.

So today, the cash flows returned by the loan are worth £7.27 + £6.61 + £82.64 = £96.52. So this loan is worth £96.52 to the investor. In other words, if the original lender wanted to sell this loan, an investor would pay £96.52.

Through the understanding of present value and knowing how to calculate it, investors can assess whether the price of a financial instrument trading in the marketplace is priced cheaply, priced fairly, or overpriced.

2.2.3.2 Time Value of Money and Regular Payments Many kinds of financial arrangements involve regular payments over time. For example, most consumer loans, including mortgages, involve regular periodic payments to pay off the loan. Each period, some of the payment covers the interest on the loan and the rest of the payment pays off some of the principal (the loaned amount). A pension savings scheme or pension plan may also involve regular contributions.

Most consumer loans result in a final balance of money equal to zero. That is, the loan is paid off. Two time value of money applications that require the final balance of money to be zero are annuities and mortgages.

An **annuity** involves the initial payment of an amount, usually to an insurance company, in exchange for a fixed number of future payments of a certain amount. Each period, the insurance company makes payments to the annuity holder; these payments are equivalent to the annuity holder making withdrawals. These withdrawals can be viewed as negative cash flows because they reduce the annuity balance. The initial payment to the insurer is called the value of the annuity and the final value is equal to zero.

A repayment or amortising mortgage involves a loan and a series of fixed payments. The initial amount of the loan is referred to as the principal. Although the payment amounts are fixed, the portion of each payment that is interest is based on the remaining principal at the beginning of each period. As some of the principal is repaid each period, the amount of interest decreases over time, and thus the amount of principal repaid increases with each successive payment until the value of the principal is reduced to zero. At this point, the loan is said to mature.

Example 6 illustrates the reduction of an annuity to zero over time and the reduction of a mortgage to zero over time. To simplify the examples, the assumption is that the annuity and the mortgage each mature in five years and entail a single withdrawal or payment each of the five years.

EXAMPLE 6. ANNUITY AND MORTGAGE

1 A retired French man pays an insurance company €10,000 in exchange for a promise by the insurance company to pay him €2,375 at the end of each of the next four years and €2,370 at the end of the fifth year. The insurance company is in effect paying him 6.0% interest on the annuity balance.

Year	Annuity Balance at Beginning of Year	Balance at End of Year before Withdrawal	Withdrawal (Payment by Insurance Company)
1	€10,000	€10,600 (= 10,000 × 1.06)	€2,375
2	€8,225	€8,719 (= 8,225 × 1.06)	€2,375
3	€6,344	€6,725 (= 6,344 × 1.06)	€2,375
4	€4,350	€4,611 (= 4,350 × 1.06)	€2,375
5	€2,236	€2,370 (= 2,236 × 1.06)	€2,370
6	€0		

2 You borrow £60,000 to buy a small cottage in the country. The interest rate on the mortgage is 4.60%. Your payment at the end of each year will be £13,706.

Year	Mortgage Outstanding at Beginning of Year	Total Mortgage Payment	Interest Paid	Principal Reduced
1	£60,000	£13,706	£2,760 (= 60,000 × 0.046)	£10,946
2	£49,054	£13,706	£2,257 (= 49,054 × 0.046)	£11,449
3	£37,605	£13,706	£1,730 (= 37,605 × 0.046)	£11,976
4	£25,630	£13,706	£1,179 (= 25,630 × 0.046)	£12,527
5	£13,103	£13,706	£603 (= 13,103 × 0.046)	£13,103
6	£0			

As you can see in Example 6, both the annuity and mortgage balances decline to zero over time.

3 DESCRIPTIVE STATISTICS

As the name suggests, descriptive statistics are used to describe data. Often, you are confronted by data that you need to organise in order to understand it. For example, you get the feeling that the drive home from work is getting slower and you are thinking of changing your route. How could you assess whether the journey really is getting slower? Suppose you calculated and compared the average daily commute time each month over a year. The first question you need to address is, what is meant by average? There are a number of different ways to calculate averages that are described in Section 3.1, each of which has advantages and disadvantages.

In general, descriptive statistics are numbers that summarise essential features of a data set. A data set relates to a particular variable—the time it takes to drive home from work in our example. The data set includes several observations—that is, observed values for the variable. For example, if you keep track of your daily commute time for a year, you will end up with approximately 250 observations. The **distribution** of a variable is the values a variable can take and the number of observations associated with each of these values.

We will discuss two types of descriptive statistics: those that describe the central tendency of a data set (e.g., the average or mean) and those that describe the dispersion or spread of the data (e.g., the standard deviation). In addition to knowing whether the drive to work is getting slower (by comparing monthly averages), you might also want to find a way to measure how much variation there is between journey times from one day to another (by using standard deviation).

Similar needs to summarise data arise in business. For example, when comparing the time taken to process two types of trades, a sample of the times required to process each trade would need to be collected. The average time it takes to process each type of trade could be calculated and the average times could then be compared. Descriptive statistics efficiently summarise the information from large quantities of data for the purpose of making comparisons. Descriptive statistics may also help in predicting future values and understanding risk. For example, if there was little variation in the times taken to process a trade, then presumably you would be confident that you had a good idea of the average time it takes to process a trade and comfortable with that as an estimate of how long it will take to process future trades. But if the time taken to process trades was highly variable, you would have less confidence in how long it would take on average to process future trades.

3.1 Measures of Frequency and Average

The purpose of measuring the frequency of outcomes or "central tendency" is to describe a group of individual data scores with a single measurement. The value used to describe the group will be the single value considered to be most representative of all the individual scores.

Measures of central tendency are useful for making comparisons between groups of individuals or between sets of figures. Such measures reduce a large number of measurements to a single figure. For instance, the mean or average temperature in country X in July from 1961 to 2011 is calculated to be 16.1°C. Over the same period in September, the average temperature is 13.6°C. Because it is a long time series, you can reasonably conclude that it is usually warmer in July than September in country X.

Common measures of central tendency are

- arithmetic mean,
- geometric mean,
- median, and
- mode.

The appropriate measure for a given data set depends on the features of the data and the purpose of your calculation. These measures are examined in the following sections.

3.1.1 *Arithmetic Mean*

The **arithmetic mean** is the most commonly used measure of central tendency and is familiar to most people. It is usually shortened to just "mean" or "average". To calculate the mean, you add all the numbers in the data set together and divide by the number of observations (items in the data set). The arithmetic mean assumes that each observation is equally probable (likely to occur). If each observation is not equally probable, you can get a weighted mean by multiplying the value of each observation by its probability and then summing these values. The sum of the probabilities always equals 1.

Exhibit 5 shows the annual returns earned on an investment over a 10-year period. The information contained in Exhibit 5 will be used in examples throughout this section.

Exhibit 5 Ten Years of Annual Returns

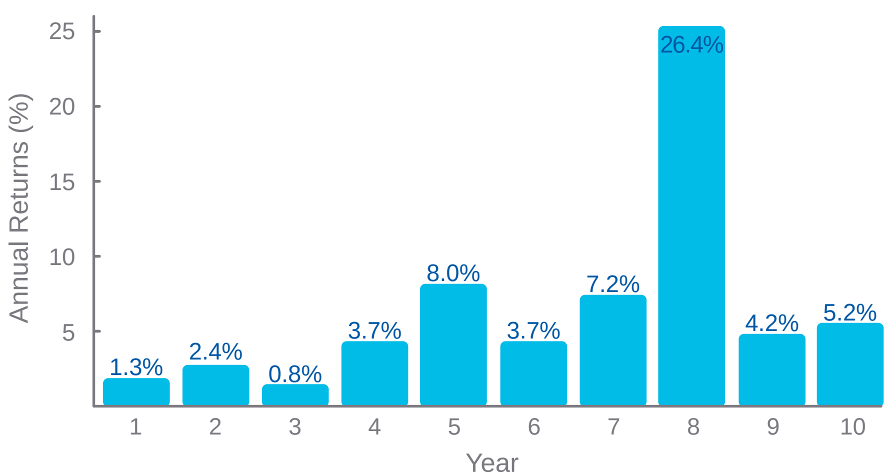

Example 7 shows the calculation of the arithmetic mean.

EXAMPLE 7. ARITHMETIC MEAN

An investment earns the returns shown in Exhibit 5 over a 10-year period.

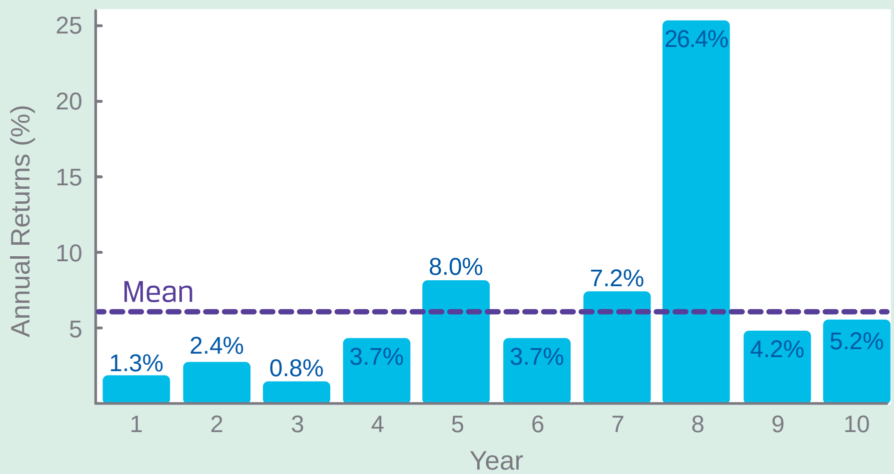

$$\frac{(1.3 + 2.4 + 0.8 + 3.7 + 8.0 + 3.7 + 7.2 + 26.4 + 4.2 + 5.2)}{10} = \textbf{6.3\% Mean}$$

The arithmetic mean return or average annual return over the 10-year period is 6.3%. The weighted mean return (shown in the following equation) is the same as the arithmetic return because the probability assigned to each return is the same: 10% or 0.1.

Weighted mean annual return

$$= (0.1 \times 1.3) + (0.1 \times 2.4) + (0.1 \times 0.8) + (0.1 \times 3.7) + (0.1 \times 8.0) + (0.1 \times 3.7) + (0.1 \times 7.2) + (0.1 \times 26.4) + (0.1 \times 4.2) + (0.1 \times 5.2)$$

$$= 6.3\%$$

The arithmetic mean annual return is 6.3%.

The mean has one main disadvantage: it is particularly susceptible to the influence of **outliers**. These are values that are unusual compared with the rest of the data set by being especially small or large in numerical value. The arithmetic mean is not very representative of the whole set of observations when there are outliers. Example 8 shows the effect of excluding an outlier from the calculation of the arithmetic mean.

EXAMPLE 8. EFFECT OF OUTLIER ON ARITHMETIC MEAN

In the case of the annual returns in Exhibit 5, there is one value—26.4%—that is much larger than the others. If this value is included, the mean is 6.3%, but excluding this value reduces the mean to 4.1%.

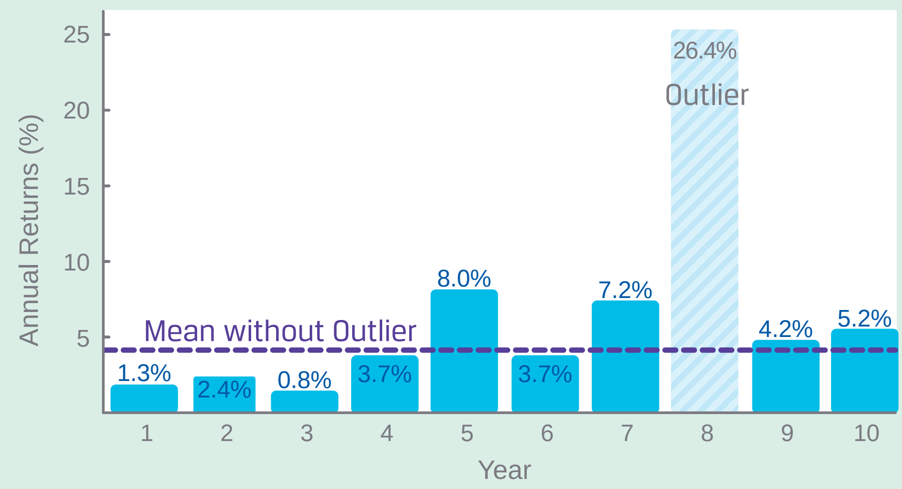

$$\frac{(1.3 + 2.4 + 0.8 + 3.7 + 8.0 + 3.7 + 7.2 + 4.2 + 5.2)}{9} = \textbf{4.1\% Mean}$$

The arithmetic mean excluding the outlier is 4.1%.

Including the outlier, the mean is dragged in the direction of the outlier. When there are one or more outliers in a set of data in one direction, the data are said to be **skewed** in that direction. In Example 7, ordering data so larger numbers are to the right of smaller numbers, 26.4% lies to the right of the other data. Thus, the data are said to be right skewed (or positively skewed). Other measures of central tendency may better accommodate outliers.

3.1.2 *Geometric Mean*

An alternative average to the arithmetic mean is the **geometric average** or **geometric mean**. Applied to investment returns, the geometric mean return is the average return assuming that returns are compounding. To illustrate how the geometric mean is calculated, let us start with the example of a three-year investment that returns 8% the first year, 3% the second year, and 7% the third year.

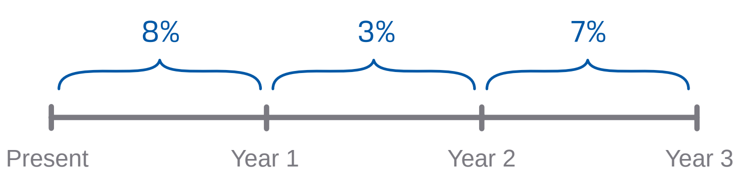

$$[(1 + 8\%) \times (1 + 3\%) \times (1 + 7\%)]^{1/3} - 1 \approx \mathbf{6.0\%}$$

The first step to calculate the geometric mean return is to multiply 1 plus each annual return and add them together, which gives you the amount you would have accumulated at the end of the three years per currency unit of investment: [(1 + 8%) × (1 + 3%) × (1 + 7%) ≈ 1.1903]. This value of 1.1903 reflects three years of investment, but the geometric mean return should capture an average rate of return for each of the three years. So, the second step requires moving from three years to one by raising the accumulation to the power of "one over the number of periods held," three in this particular case; this calculation can also be described as taking "the number of periods held" root of the value ($1.1903^{1/3} \approx 1.060$). This value of 1.060 includes both the original investment and the average yearly return on the investment each year (1 plus the geometric mean return). The last step is, therefore, to subtract 1 from this value to arrive at the return that would have to be earned on average each year to get to the total accumulation over the three years ($1.060 - 1 \approx 0.060$ or 6.0%). The geometric mean return is 6.0%, which in this case is the same as the arithmetic mean return. Geometric mean is frequently the preferred measure for the investment industry.

The following formula is used to arrive at the geometric mean return:

$$\text{Geometric mean return} = \left[(1 + r_1) \times \ldots (1 + r_t)\right]^{1/t} - 1$$

where

r_i = the return in period i expressed using decimals

t = the number of periods

Example 9 shows the calculation of the geometric mean return for the investment of Exhibit 5.

EXAMPLE 9. GEOMETRIC MEAN RETURN

If 1 currency unit was invested, you would have 1.8 currency units at the end of the 10 years.

Total accumulation after 10 years

$= [(1 + 1.3\%) \times (1 + 2.4\%) \times (1 + 0.8\%) \times (1 + 3.7\%) \times (1 + 8.0\%) \times (1 + 3.7\%) \times (1 + 7.2\%) \times (1 + 26.4\%) \times (1 + 4.2\%) \times (1 + 5.2\%)]$

$= [(1.013) \times (1.024) \times (1.008) \times (1.037) \times (1.08) \times (1.037) \times (1.072) \times (1.264) \times (1.042) \times (1.052)]$

$= 1.8$

Average accumulation per year = 10th root of 1.8 = $(1.8)^{1/10} = 1.061$

Geometric mean annual return = $1.061 - 1 = 0.061 = 6.1\%$

This can also be done as one calculation:

Geometric mean annual return

$= \{[(1 + 1.3\%) \times (1 + 2.4\%) \times (1 + 0.8\%) \times (1 + 3.7\%) \times (1 + 8.0\%) \times (1 + 3.7\%) \times (1 + 7.2\%) \times (1 + 26.4\%) \times (1 + 4.2\%) \times (1 + 5.2\%)]^{(1/10)}\} - 1$

$= 6.1\%$

The geometric mean annual return is 6.1%. One currency unit invested for 10 years and earning 6.1% per year would accumulate to approximately 1.8 units.

An important aspect to notice is that the geometric mean is lower than the arithmetic mean even though the annual returns over the 10-year holding period are identical. This result is because the returns are compounded when calculating the geometric mean return. Recall that compounding will result in a higher value over time, so a lower rate of return is required to reach the same amount. In fact, if the same set of numbers is used to calculate both means, the geometric mean return is never greater than the arithmetic mean return and is normally lower.

3.1.3 *Median*

If you put data in ascending order of size from the smallest to the largest, the **median** is the middle value. If there is an even number of items in a data set, then you average the two middle observations. Hence, in many cases (i.e., when the sample size is odd or when the two middle-ranked items of an even-numbered data set are the same) the median will be a number that actually occurs in the data set. Example 10 shows the calculation of the median for the investment of Exhibit 5.

EXAMPLE 10. MEDIAN

When the returns are ordered from low to high, the median value is the arithmetic mean of the fifth and sixth ordered observations.

Annual Returns Ordered Low to High

0.8% 1.3% 2.4% 3.7% $\boxed{\textbf{3.7\%}\ \ \textbf{4.2\%}}$ 5.2% 7.2% 8.0% 26.4%

$$\frac{(3.7 + 4.2)}{2} \approx \textbf{4.0\% Median}$$

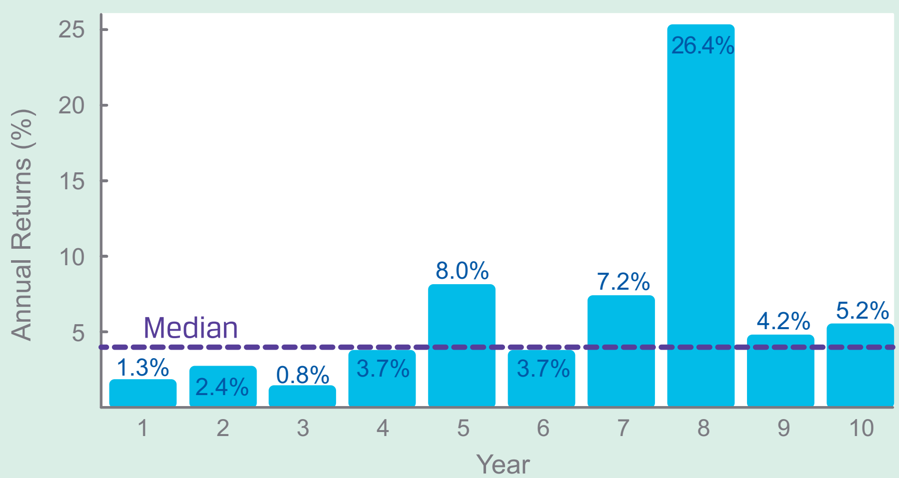

The median investment return over the 10-year period is 4.0%.

An advantage of the median over the mean is that it is not sensitive to outliers. In the case of the annual returns shown in Exhibit 5, the median of close to 4.0% is more representative of the data's central tendency. This 4.0% median return is close to the 4.1% arithmetic mean return when the outlier is excluded. The median is usually a better measure of central tendency than the mean when the data are skewed.

3.1.4 *Mode*

The **mode** is the most frequently occurring value in a data set. Example 11 shows how the mode is determined for the investment of Exhibit 5.

EXAMPLE 11. MODE

Looking at Exhibit 5, we see that one value occurs twice, 3.7%. This value is the mode of the data.

Annual Returns Ordered Low to High

0.8% 1.3% 2.4% **3.7% 3.7%** 4.2% 5.2% 7.2% 8.0% 26.4%

3.7% Mode

The mode can be used as a measure of central tendency for data that have been sorted into categories or groups. For example, if all the employees in a company were asked what form of transportation they used to get to work each day, it would be possible to group the answers into categories, such as car, bus, train, bicycle, and walking. The category with the highest number would be the mode.

A problem with the mode is that it is often not unique, in which case there is no mode. If there are two or more values that share the same frequency of occurrence, there is no agreed method to choose the representative value. The mode may also be difficult to compute if the data are continuous. **Continuous data** are data that can take on an infinite number of values between whole numbers—for example, weights of people. One person may weigh 62.435 kilos and another 62.346 kilos. By contrast, **discrete data** show observations only as distinct values—for example, the number of people employed at different companies. The number of people employed will be a whole number. For continuous data, it is less likely that any observation will occur more frequently than once, so the mode is generally not used for identifying central tendency for continuous data.

Another problem with the mode is that the most frequently occurring observation may be far away from the rest of the observations and does not meaningfully represent them.

3.2 Measures of Dispersion

Whereas measures of central tendency are used to estimate representative or central values of a set of data, measures of dispersion are important for describing the spread of the data or its variation around a central value. Two data sets may have the same mean or median but completely different levels of variability, or vice versa. A description of a data set should include both a measure of central tendency, such as the mean, and a measure of dispersion. Suppose two companies both have an average annual salary of $50,000, but in one company most salaries are clustered close to the average, whereas in the second they are spread out with many people earning very little and some earning a lot. It would be useful to have a measure of dispersion that can help identify such differences between data sets.

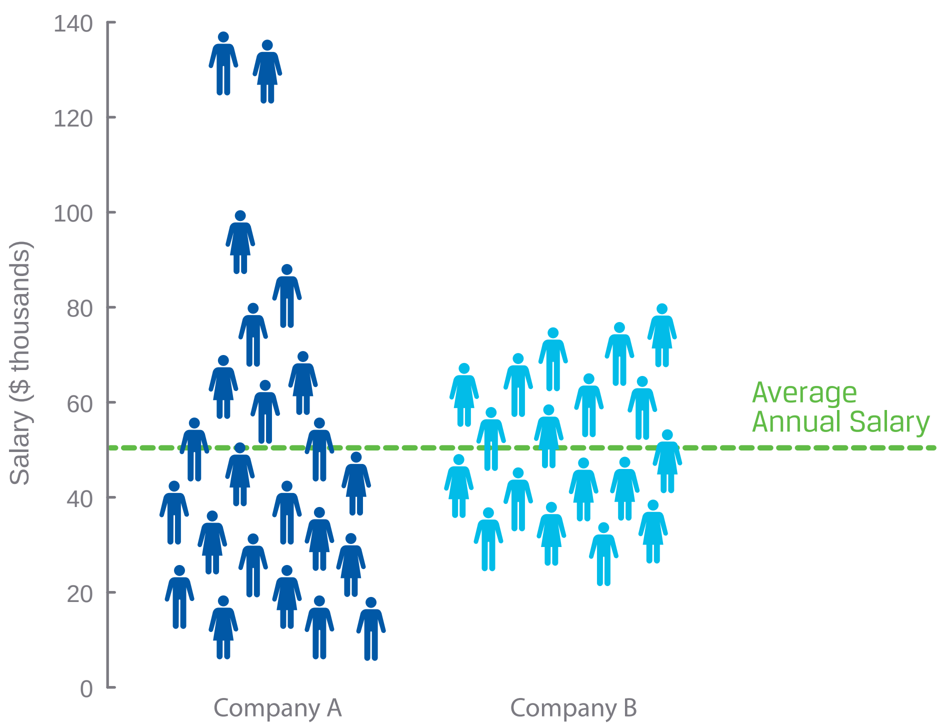

Another reason why measures of dispersion are important in finance is that investment risk is often measured using some measure of variability. When investors are considering investing in a security, they are interested in the likely (expected) return on that investment as well as in the risk that the return could differ from the expected return (its variability). A risk-averse investor considering two investments that have similar expected returns but very different measures of variability (risk) around those expected returns, typically prefers the security with the lower variability.

Two common measures of dispersion of a data set are the range and the standard deviation.

3.2.1 *Range*

The **range** is the difference between the highest and lowest values in a data set. It is the easiest measure of dispersion to calculate and understand, but it is very sensitive to outliers. Example 12 explains the calculation of the range of returns for the investment of Exhibit 5.

EXAMPLE 12. RANGE

In Exhibit 5 we see that the highest annual return is 26.4% and the lowest annual return is 0.8%.

Annual Returns Ordered Low to High

$\boxed{0.8\%}$ 1.3% 2.4% 3.7% 3.7% 4.2% 5.2% 7.2% 8.0% $\boxed{26.4\%}$

26.4% – 0.8% = 25.6% Range

If the extreme value at the upper end of the range is excluded, the next highest value, 8.0%, is used to estimate the range, and the range is reduced significantly.

Annual Returns Ordered Low to High

$\boxed{0.8\%}$ 1.3% 2.4% 3.7% 3.7% 4.2% 5.2% 7.2% $\boxed{8.0\%}$ ~~26.4%~~

8.0% – 0.8% = 7.2% Range

Clearly, the range is affected by extreme values and, if there are outliers, it says little about the distribution of the data between those extremes.

If there are a large number of observations ranked in order of size, the range can be divided into 100 equal-sized intervals. The dividing points are termed percentiles. The 50th percentile is the median and divides the observations so that 50% are higher and 50% are lower than the median. The 20th percentile is the value below which 20% of observations in the series fall. So, the dispersion of the observations can be described in terms of percentiles. Observations can be divided into other equal-sized intervals. Commonly used intervals are quartiles (the observations are divided into four equal-sized intervals) and deciles (the observations are divided into 10 equal-sized intervals)

3.2.2 *Standard Deviation*

A commonly used measure of dispersion is the **standard deviation**. It measures the variability or volatility of a data set around the average value (the arithmetic mean) of that data set. Although, as mentioned before, you are not responsible for any calculations, you may find it helpful to look at the formula for how standard deviation is calculated.

$$\text{Standard deviation} = \sqrt{\frac{[X_1 - E(X)]^2 + [X_2 - E(X)]^2 + \ldots + [X_n - E(X)]^2}{n}}$$

where

X_i = observation i (one of n possible outcomes for X)

n = number of observations of X

$E(X)$ = the mean (average) value of X or the expected value of X

$[X_i - E(X)]$ = difference between value of observation X_i and the mean value of X

The differences between the observed values of X and the mean value of X capture the variability of X. These differences are squared and summed. Note that because the differences are squared, what matters is the size of the difference not the sign of the difference. The sum is then divided by the number of observations. Finally, the square root of this value is taken to get the standard deviation.

The value before the square root is taken is known as the **variance**, which is another measure of dispersion. The standard deviation is the square root of the variance. The standard deviation and the variance capture the same thing—how far away from the mean the observations are. The advantage of the standard deviation is that it is expressed in the same unit as the mean. For example, if the mean is expressed as minutes of journey time, the standard deviation will also be expressed as minutes, whereas the variance will be expressed as minutes squared, making the standard deviation an easier measure to use and compare with the mean.

To illustrate the calculation of the standard deviation, let us return to the example of a three-year investment that returns 8% or 0.08 the first year, 3% or 0.03 the second year, and 7% or 0.07 the third year. The arithmetic mean return is 6% or 0.06. The standard deviation is approximately 2.16%.

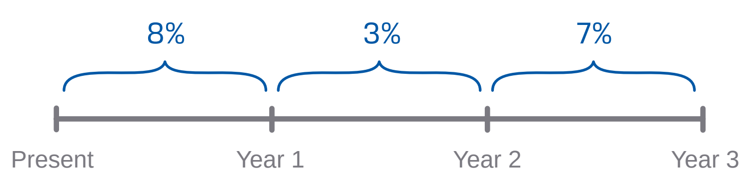

$$\text{Standard deviation} = \sqrt{\frac{(0.08 - 0.06)^2 + (0.03 - 0.06)^2 + (0.07 - 0.06)^2}{3}}$$

$$= \sqrt{\frac{(0.02)^2 + (-0.03)^2 + (0.01)^2}{3}}$$

$$= \sqrt{\frac{(0.0004) + (0.0009) + (0.0001)}{3}}$$

$$= \sqrt{\frac{(0.0014)}{3}} = 0.0216 = 2.16\%$$

Example 13 shows the calculation of the standard deviation for the investment in Exhibit 5.

EXAMPLE 13. STANDARD DEVIATION

The arithmetic mean annual return, as calculated in Example 7, is 6.3%.

Standard deviation

= square root of {[(0.013 – 0.063)^2 + (0.024 – 0.063)^2 + (0.008 – 0.063)^2 + (0.037 – 0.063)^2 + (0.08 – 0.063)^2 + (0.037 – 0.063)^2 + (0.072 – 0.063)^2 + (0.264 – 0.063)^2 + (0.042 – 0.063)^2 + (0.052 – 0.063)^2] ÷ 10}

= square root of [(0.0025 + 0.0015 + 0.0030 + 0.0007 + 0.0003 + 0.0007 + 0.0001 + .0404 + 0.0004 + 0.0001) ÷ 10]

= square root of 0.00497

= 0.0705, rounded to the nearest 10th percent = 7.1% (this value is used in Example 14).

The standard deviation is 7.1%.

Larger values of standard deviation relative to the mean indicate greater variation in a data set. Also, by using standard deviation, you can determine how likely it is that any given observation will occur based on its distance from the mean. Example 14 compares the returns of the investment shown in Exhibit 5 and the returns on another investment over the same period using mean and standard deviation.

EXAMPLE 14. COMPARISON OF INVESTMENTS

An investment earns the returns shown in Exhibit 5 over a 10-year period:

Number of observations = 10

Mean = 6.3%

Standard deviation = 7.1%

Another investment over the same time period has the following characteristics:

Number of observations = 10

Mean = 6.5%

Standard deviation = 2.6%

Based on mean and standard deviation, the second investment is better than the first investment. It has a higher mean return and less variability, which implies less risk, in its returns.

3.2.3 *Normal Distribution*

The arithmetic mean and standard deviation are two powerful ways of describing many distributions of data. A distribution is simply the set of values that a variable can take, showing their observed or theoretical frequency of occurrence. For example, consider the distribution of salaries earned by employees in two companies as shown in Exhibit 6. The observations in these distributions are grouped into different salary ranges.

Exhibit 6 Number of Employees in Various Salary Ranges

	Number of Employees	
Salary ($)	**Company X**	**Company Y**
15,000–20,000	5	1
20,001–25,000	8	1
25,001–30,000	20	3
30,001–35,000	30	8
35,001–40,000	22	10
40,001–45,000	12	15
45,001–50,000	6	20
50,001–55,000	2	9
55,001–60,000	1	7

Sometimes it is helpful to look at a picture of the distribution to understand it. The shape of the distribution has a bearing on how you interpret the summary measures of the distribution. This data can be shown pictorially using a **histogram**—a bar chart with bars that are proportional to the frequency of occurrence of each group of observations—as shown in Exhibits 7A and 7B.

Exhibit 7A Salaries of Employees at Company X

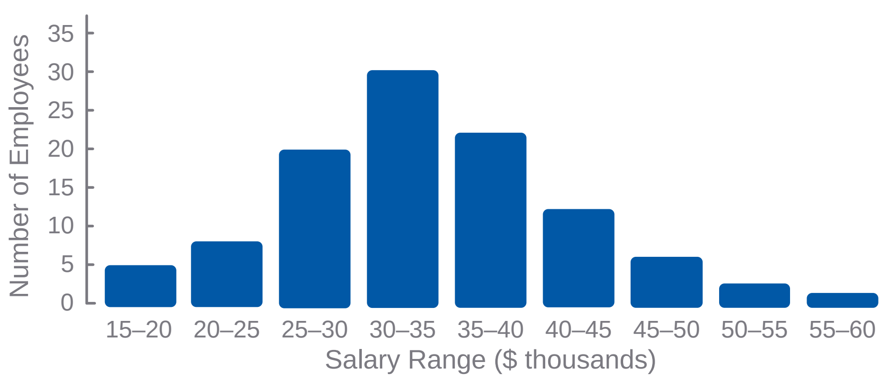

Exhibit 7B Salaries of Employees at Company Y

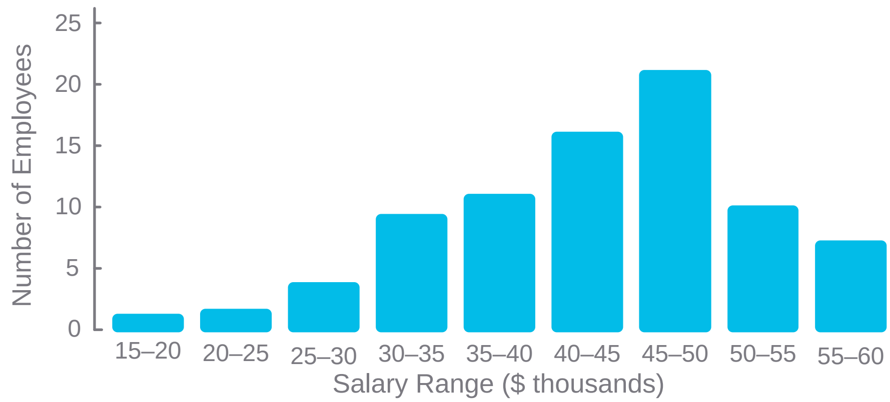

Note that the two distributions are not symmetrical. A symmetrical distribution would have observations falling off fairly evenly on either side of the centre of the range of salaries ($35,001–$40,000). Instead, in each of these distributions, the bulk of the observations are stacked towards one end of the range and tail off gradually towards the other end. The two distributions are different in that each is stacked towards a different end. Such distributions are considered skewed; the distribution for Company X is positively skewed (i.e., the majority of the observations are on the left and the skew or tail is on the right), whereas the distribution for Company Y is negatively skewed (left skewed).

Although the range of the observations is the same in each case, the mean for each is very different. Company X's mean is approximately $35,000, whereas Company Y's mean is approximately $44,000.

For a perfectly symmetrical distribution, such as a normal distribution (see Exhibit 8), the mean, median, and mode will be identical. If the distribution is skewed, these three measures of central tendency will differ. Looking again at Company X's salary data, for instance, we do not have enough detailed information to identify a mode. The mean is larger than the median because the mean is more affected by extreme values than the median. The distribution is skewed to the right, so the mean is dragged towards the extreme positive values. The reverse is true for distributions that are negatively skewed, such as in Company Y's salary data. In this case, the mean is smaller than the median because the mean is pulled left in the direction of the skew.

A **normal distribution** is represented in a graph by a bell curve; an example of a bell-shaped curve is shown in Exhibit 8. The shape of the curve is symmetrical with a single central peak at the mean of the data and the graph falling off evenly on either side of the mean; 50% of the distribution lies to the left of the mean, and 50% lies to the right of the mean. The shape of a normal distribution depends on the mean and the standard deviation. The mean of the distribution determines the location of the centre of the curve, and the standard deviation determines the height and width of the curve. When the standard deviation is large, the curve is short and wide; when the standard deviation is small, the curve is tall and narrow.

A normal distribution has special importance in statistics because many variables have the approximate shape of a normal distribution—for example, height, blood pressure, and lengths of objects produced by machines. This distribution is often useful as a description of data when there are a large number of observations.

A normal distribution is a distribution of a continuous random variable (i.e., a variable that can take on an infinite number of values). The vertical axis for the normal distribution is the probability or likelihood of occurrence. By contrast, on the histogram shown earlier, the vertical axis was frequency of occurrence. The mean (and median) is the centre of the distribution and has the highest probability of occurrence. Half of the observations lie on one side of the mean and half on the other. Approximately two-thirds of the observations are within one standard deviation of the mean, and 95% of observations are within two standard deviations of the mean. Exhibit 8 shows a normal distribution.

Exhibit 8 Standard Deviation (SD) and Normal Distribution

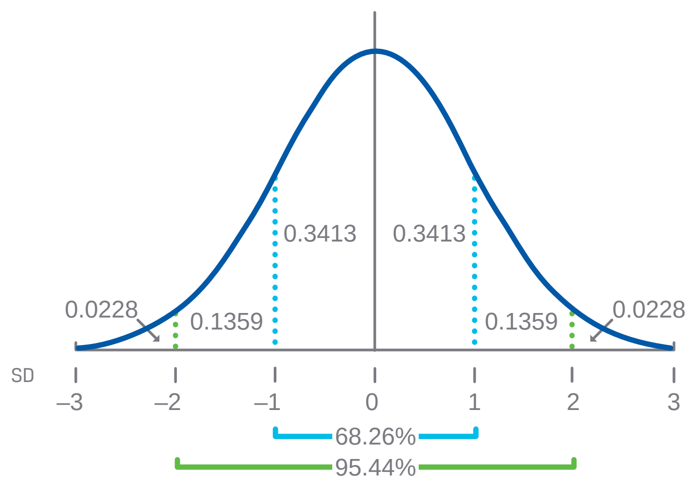

The total area under the curve or bell is 100% of the distribution. The area under the curve that is within one standard deviation of the mean is about 68% of all the observations. In other words, given a mean of 0 and a standard deviation of 1, about 68% of the observations fall between –1 and +1, and 32% of the observations are more than one standard deviation from the mean. The area under the curve that is within 2 standard deviations of the mean is about 95% of the observations. Given a mean of 0 and a standard deviation of 1, about 95% of the observations fall between –2 and +2, and 5% of the observations are more than two standard deviations from the mean. The area under the curve that is within three standard deviations of the mean represents about 99% of the observations. Given a mean of 0 and a standard deviation of 1, about 99% of the observations fall between –3 and +3, and less than 1% of the observations occur more than three standard deviations away from the mean.

The observations that are more than a specified number of standard deviations from the mean can be described as lying in the tails of the distribution. Assuming that returns on a portfolio of stocks are normally distributed, the chance of extreme losses (a return more than three standard deviations lower than the mean return) is relatively small. The chance of the return being in the left tail more than two standard deviations from the mean (which would be an extreme loss under typical circumstances) is just 2.5%. In other words, out of 200 days, 5 days are expected to have observations that are more than two standard deviations from the mean. But during the financial crisis of 2008, the losses that were incurred by some banks over several days in a row were 25 standard deviations below the mean.

To put this in perspective, if returns are normally distributed, a return that is 7.26 standard deviations below the mean would be expected to occur once every 13.7 billion years. That is approximately the age of the universe. The frequency of extreme events during the financial crisis of 2008 was, therefore, much higher than predicted by the normal distribution. This inconsistency is often referred to as the distribution having "fat tails", meaning that the probability of observing extreme outcomes is higher than that predicted by a normal distribution.

Exhibit 9 gives examples of different types of bell-shaped distributions. How would you describe each curve? What does each tell you about the likelihood of extreme outcomes?

Exhibit 9 Bell-Shaped Distributions with Fat and Thin Tails

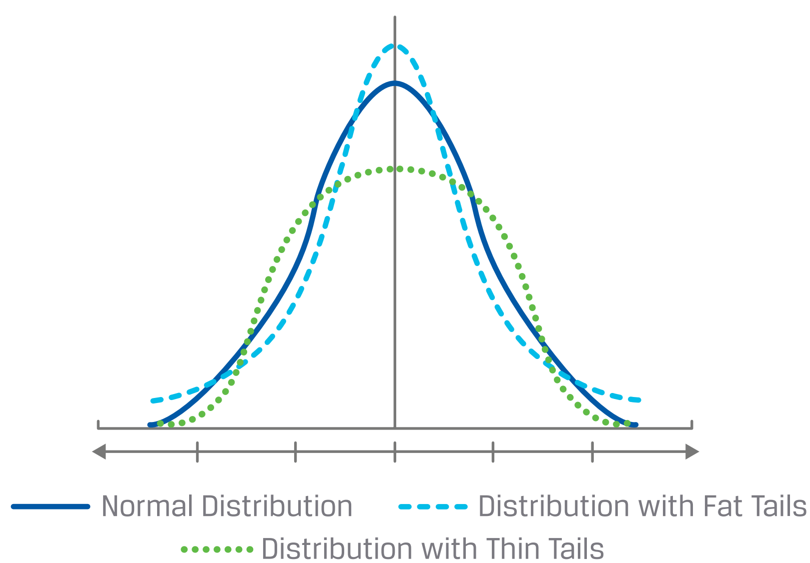

In Exhibit 9, the curve with the solid line represents the normal distribution. The curve with the dotted line is an example of distribution with thinner tails than the normal distribution, indicating a reduced probability of extreme outcomes. By contrast, the curve with the dashed line is an example of a distribution with fatter tails than the normal distribution, indicating increased likelihood of extreme outcomes.

3.3 Correlation

Another way of using and understanding data is identifying connections between data sets. The strength of a relationship between two variables, such as growth in gross domestic product (GDP) and stock market returns, can be measured by using **correlation**. Essentially, two variables are correlated when a change in one variable helps predict change in another variable.

When both variables change in the same direction, the variables are positively correlated. If we take the example of traders at an investment bank, salary and age are positively correlated if salaries increase as age increases. If the variables move in the opposite direction, then they are negatively correlated. For example, the size of a transaction and the fees expressed as a percentage of the transaction are negatively correlated if the larger the transaction, the smaller the associated fees. When there is no clear tendency for one variable to move in a particular direction (up or down) relative to changes in the other variable, then the variables are close to being uncorrelated. In practice, it is difficult to find two variables that have absolutely no relationship, even if just by chance.

Correlation is measured by the **correlation coefficient**, which has a scale of –1 to +1. When two variables move exactly in step with each other in the same direction—if one goes up, the other goes up in the same proportion—the variables are said to be perfectly positively correlated. In that case, the correlation coefficient is at its maximum of +1. When the two variables move exactly in step in opposite directions, they are perfectly negatively correlated and the correlation coefficient is –1. Variables with no relationship to each other will have a correlation coefficient close to 0.

Correlation measures both the direction of the relationship between two variables (negative or positive) and the strength of that relationship (the closer to +1 or –1, the stronger the relationship). In practice, it is unusual to find variables that are perfectly positively or perfectly negatively correlated. The stronger the relationship between two variables—the higher the degree of correlation—the more confidently one variable can be predicted given the other variable. For example, there may be a high correlation between stock market index returns and expected economic growth. In that case, if economic growth in the future is expected to be high then returns on the stock market index are likely to be high too.

It is important, however, to realise that correlation does not imply causation. For example, historically in the United States, stock market returns and snowfall are both higher in January, and from that you may assume a correlation. But obviously snowfall does not cause an increase in stock market returns, and an increase in stock market returns clearly does not cause snowfall. There may be situations in which a correlation implies some causal relationship. For example, a high correlation has been found between power production and job growth. It may follow that the more workers there are, the more power is consumed, but it does not necessarily follow that an increase in power generation will create jobs.

Correlation is important in investing because the rise or fall in value of a variable may help predict the rise or fall in value of another variable. It is also important because when two or more securities that are not perfectly correlated are combined together in a portfolio, there is normally a reduction in risk (measured by the portfolio's standard deviation of returns). As long as the returns on the securities do not have a correlation of +1 (that is, they are less than perfectly correlated), then the risk of the portfolio will be less than the weighted average of the risks of the securities in the portfolio because it is not likely that all the securities will perform poorly at the same time.

The practice of combining securities in a portfolio to reduce risk is known as **diversification**. An extreme example of an undiversified portfolio is one holding only one security. This approach is risky because it is not unusual for a single security to go down in value by a large amount in one year. It is much less common for a diversified portfolio of 20 different securities to go down by a large amount, even if they are selected at random. If the securities are selected from a variety of sectors, industries, company sizes, asset classes, and markets, it is even less likely. One caveat is that the benefits of diversification are much reduced in periods of financial crisis. In such periods, the correlation between returns on different securities (and different asset classes) tends to increase towards +1.

SUMMARY

The better your understanding of quantitative concepts, the easier it will be for you to make sense of the financial world. Knowledge of quantitative concepts, such as time value of money and descriptive statistics, is important to the understanding of many of the key products in the financial industry. Understanding the time value of money allows you to interpret cash flows and thus value them. Meanwhile, knowledge of statistical concepts will help in identifying the important information in a large amount of data, as well as in understanding what statistical measures reported by others mean. It is easy to misinterpret or be misled by statistics, such as mean and correlation, so an understanding of their uses and limitations is crucial.

- Interest is return earned by a lender that compensates for opportunity cost and risk. For the borrower, it is the cost of borrowing.
- The simple interest rate is the cost to the borrower or the rate of return to the lender, per period, on the original principal borrowed. A commonly quoted simple interest rate is the annual percentage rate (APR).
- Compound interest is the return to the lender or the cost to the borrower when interest is reinvested and added to the original principal.
- The effective annual rate (EAR) of interest is calculated by annualising a rate that is compounded more than once a year. The EAR is equal to or greater than the annual percentage rate.
- The present value of a future sum of money is found by discounting the future sum by an appropriate discount rate. (The present value of multiple cash flows is the sum of the present value of each cash flow.)
- Three elements must be considered when comparing investments: the cash flows each investment will generate in the future, the timing of these cash flows, and the risk associated with each investment. The discount rate reflects the riskiness of the cash flows.

- All else being equal (in other words, only one of the three elements differs):
  - the higher the cash flows, the higher the present and future values.
  - the earlier the cash flows, the higher the present and future values.
  - the lower the discount rate, the higher the present value.
  - the higher the interest rate, the higher the future value.
- The net present value is the present value of future cash flows net of the investment required to obtain them. It is useful when comparing alternatives that require different initial investments.
- Financial instruments can be valued as the present value of their expected future cash flows.
- An annuity involves an initial payment (outflow) in exchange for a fixed number of future receipts (inflows), each of an equal amount. Mortgages are amortising loans; the periodic payment is fixed, and in each period some of the payment covers the interest on the loan and the rest of the payment pays off some of the principal. Over time, the portion of the payment that reduces the principal increases.
- The role of descriptive statistics is to summarise the information given in large quantities of data for the purpose of making comparisons, predicting future values, and better understanding the data.
- The purpose of measures of frequency or central tendency is to describe a group of individual data scores with a single measurement. This measure is intended to be representative of the individual scores. Measures of central tendency include arithmetic mean, geometric mean, median, and mode. Different measures are appropriate for different types of data.
- The arithmetic mean is the most commonly used measure. It represents the sum of all the observations divided by the number of observations. It is an easy measure to understand but may not be a good representative measure when there are outliers.
- The geometric mean return is the average compounded return for each period—that is, the average return for each period assuming that returns are compounding. It is frequently the preferred measure of central tendency for returns in the investment industry.
- When observations are ranked in order of size, the median is the middle value. It is not sensitive to outliers and may be a more representative measure than the mean when data are skewed.
- The mode is the most frequently occurring value in a data set. A data set may have no identifiable unique mode. It may not be a meaningful representative measure of central tendency.
- Measures of dispersion are important for describing the spread of the data, or its variation around a central value. Two common measures of dispersion are range and standard deviation.

- Range is the difference between the highest and lowest values in a data set. It is easy to measure, but it is sensitive to outliers.
- Standard deviation measures the variability of a data set around the mean of the data set. It is in the same unit of measurement as the mean.
- A distribution is simply the values that a variable can take, showing its observed or theoretical frequency of occurrence.
- For a perfectly symmetrical distribution, the mean, median, and mode will be identical.
- A common symmetrical distribution is the normal distribution, a bell-shaped curve that is represented by its mean and standard deviation. In a normal distribution, 68% of all the observations lie within one standard deviation of the mean and about 95% of the observations are within two standard deviations.
- The strength of a relationship between two variables can be measured by using correlation.
- Correlation is measured by the correlation coefficient on a scale from –1 to +1. When two variables move exactly in tandem with each other, the variables are said to be perfectly positively correlated and the correlation coefficient is +1. When two variables move exactly in opposite directions, they are perfectly negatively correlated and the correlation coefficient is –1. Variables with no relationship to each other will have a correlation coefficient close to 0.
- It is important to realise that correlation does not imply causation.