

Data analysis: hypothesis testing



This item contains selected online content. It is for use alongside, not as a replacement for the module website, which is the primary study format and contains activities and resources that cannot be replicated in the printed versions.

About this free course

This free course is an adapted extract from the Open University course B126 *Business data analytics and decision making* - www.open.ac.uk/courses/modules/b126.

This version of the content may include video, images and interactive content that may not be optimised for your device.

You can experience this free course as it was originally designed on OpenLearn, the home of free learning from The Open University –

There you'll also be able to track your progress via your activity record, which you can use to demonstrate your learning.

First published 2023.

Unless otherwise stated, copyright © 2023 The Open University, all rights reserved.

Intellectual property

Unless otherwise stated, this resource is released under the terms of the Creative Commons Licence v4.0 <http://creativecommons.org/licenses/by-nc-sa/4.0/deed.en>. Within that The Open University interprets this licence in the following way:

www.open.edu/openlearn/about-openlearn/frequently-asked-questions-on-openlearn. Copyright and rights falling outside the terms of the Creative Commons Licence are retained or controlled by The Open University. Please read the full text before using any of the content.

We believe the primary barrier to accessing high-quality educational experiences is cost, which is why we aim to publish as much free content as possible under an open licence. If it proves difficult to release content under our preferred Creative Commons licence (e.g. because we can't afford or gain the clearances or find suitable alternatives), we will still release the materials for free under a personal end-user licence.

This is because the learning experience will always be the same high quality offering and that should always be seen as positive – even if at times the licensing is different to Creative Commons.

When using the content you must attribute us (The Open University) (the OU) and any identified author in accordance with the terms of the Creative Commons Licence.

The Acknowledgements section is used to list, amongst other things, third party (Proprietary), licensed content which is not subject to Creative Commons licensing. Proprietary content must be used (retained) intact and in context to the content at all times.

The Acknowledgements section is also used to bring to your attention any other Special Restrictions which may apply to the content. For example there may be times when the Creative Commons Non-Commercial Sharealike licence does not apply to any of the content even if owned by us (The Open University). In these instances, unless stated otherwise, the content may be used for personal and non-commercial use.

We have also identified as Proprietary other material included in the content which is not subject to Creative Commons Licence. These are OU logos, trading names and may extend to certain photographic and video images and sound recordings and any other material as may be brought to your attention.

Unauthorised use of any of the content may constitute a breach of the terms and conditions and/or intellectual property laws.

We reserve the right to alter, amend or bring to an end any terms and conditions provided here without notice.

All rights falling outside the terms of the Creative Commons licence are retained or controlled by The Open University.

Head of Intellectual Property, The Open University

Contents

Introduction	4
Learning outcomes	5
1 Two types of hypotheses	6
1.1 Formulating null and alternative hypotheses	7
1.2 Population mean (μ)	8
2 Testing with data	10
3 Alpha (α) levels	12
4 One-tailed vs two-tailed test	14
4.1 The normal distribution	14
4.2 Two-tailed tests	15
4.3 One-sided tests	17
4.4 Check your understanding	19
5 Mean and z-score ranges	20

Introduction

In this course, you will explore the process of testing hypotheses and making inferences about data. This will typically be a collection of individuals in a population in the statistical sense (people, companies, countries, etc.), and the quantity of interest will be a parameter describing the population.

You will learn about descriptive statistics, alpha (α) and confidence intervals as well as the distinction between one-tailed and two-tailed tests, and the concept of 'p-value'.

Furthermore, you will gain insight into how to test for differences in means and proportions. Ultimately, you will gain knowledge about the concept of 'tests for statistical significance'.



"Surprisingly, the soil samples provide new evidence for the 'giant ball of cheese' hypothesis."

Figure 1 Data sheds light on hypotheses

This course requires Microsoft Excel in some activities therefore this course should be completed on a desktop or laptop rather than a mobile device. If you do not have access to Microsoft Excel there are various other free options – such as Google Sheets, Apple Numbers or LibreOffice.

This OpenLearn course is an adapted extract from the Open University course

[B126 Business data analytics and decision making](#).

Learning outcomes

After studying this course, you should be able to:

- understand the principle of hypothesis testing
- understand the idea of alpha in hypothesis testing
- differentiate between one-tailed and two-tailed tests
- understand hypothesis testing of means and proportions
- report the exact p-value of a test.

1 Two types of hypotheses

During the process of data-driven decision making, managers typically follow four key steps:

1. Formulating a hypothesis
2. Identifying and obtaining the appropriate data to test the hypothesis
3. Executing the test
4. Making a decision based on the results.

In this course, your focus will be on gaining a comprehensive understanding of what a hypothesis is and how to effectively formulate one.

The term hypothesis refers to an explanation proposed for a phenomenon or idea (e.g. 'the sky is blue'). In order for a hypothesis to be considered scientific, it must be tested scientifically. In simple terms, a hypothesis is a 'guess' that **can be tested**.



Figure 2 iPhone

For example, say when conducting market research, a respondent states that the selling price of Apple iPhones was too high (i.e. it was too expensive). This is a supposition but cannot be tested. This is because you do not know what is meant by the term 'expensive'. However, if the respondent mentions that the Apple iPhone costs over £500, this can be considered as expensive. In this case, the respondent's supposition becomes a hypothesis. In order to test the hypothesis, market researchers need to collect some data (typically by surveying a small group of people – a sample) to study the population. The concept of population refers to the complete collection of individuals or objects that are of interest in a given context. This can include, for example, all small- and medium-sized businesses within a particular country, or the entirety of online shoppers.

The concepts of null hypothesis and alternative hypothesis are fundamental in business decision making.

The null hypothesis is an algebraic statement that expresses the currently accepted value for a parameter in the population. It is commonly used as the default position when testing a hypothesis, and the researcher presumes it to be true unless there is sufficient evidence to the contrary. The null hypothesis is usually represented by ' H_0 '. For instance, if it is believed that the average annual salary in the UK is approximately £26,000, it can be considered as the null hypothesis, which can be formally stated as:

H_0 : The average annual salary in the UK is £26,000.

On the other hand, the alternative hypothesis is a statement or assumption that challenges the null hypothesis. It is the opposite of the null hypothesis and is what the

researcher aims to demonstrate to be true. The alternative hypothesis is denoted as ' H_a ' or ' H_1 '.

An alternative hypothesis is needed because it provides an alternative explanation or theory to the null hypothesis. You will look at more examples in the next section.

1.1 Formulating null and alternative hypotheses

In the world of scientific inquiry, you often begin with a null hypothesis (H_0), which expresses the currently accepted value for a parameter in the population. The alternative hypothesis (H_a), on the other hand, is the opposite of the null hypothesis and challenges the currently accepted value.

To illustrate this concept of null and alternative hypotheses, you will look at some well-known stories and examples.

In ancient and medieval times, the widely held belief was that all planets orbited around the Earth, as the Earth was considered the centre of the universe. This idea can be considered the null hypothesis, as it represents the currently accepted value for a parameter in the population. Thus, it can be written as:

H_0 : All planets orbit around the Earth.




Figure 3 Solar system

In the world of business and finance, the idea that paper money must be backed by gold (the gold standard) was also a commonly held belief for a long time. This belief can be considered a null hypothesis. However, following the Great Depression, people began to question this belief and broke the link between banknotes and gold. This alternative hypothesis challenged the gold standard, and it eventually became widely accepted that the value of paper money is not necessarily equal to a fixed amount of gold. Thus, H_0 and H_a statements can be written as:

H_0 : The value of paper money is equal to a fixed amount of gold.

H_a : The value of paper money is not equal to a fixed amount of gold.

In modern times, people generally place their trust in the value of banknotes issued by central banks or monetary authorities, which are backed by a strong government. This belief can be considered a null hypothesis. However, digital currency, such as Bitcoin, has emerged as an alternative to traditional paper money. Bitcoin is not backed by any central bank or monetary authority, and transactions involving Bitcoin are verified by network nodes using cryptography and recorded in a blockchain. This alternative hypothesis challenges the belief that the value of paper money is solely based on people's trust in central banks or monetary authorities. Thus, H_0 and H_a statements can be written as:

H_0 : The value of paper money is equal to people's trust in central banks or monetary authorities.

H_a : The value of paper money is not equal to people's trust in central banks or monetary authorities.

In conclusion, the alternative hypothesis always challenges the idea expressed in the null hypothesis. By testing the null hypothesis against the alternative hypothesis, you can determine which idea is more supported by the available data. The alternative hypothesis is often referred to as a 'research hypothesis' because it initiates the motivation and opportunities for further research.

Let's return to the first example given in Section 1. If you see that your friends and relatives make more or less than £26,000 annually on average, perhaps you should question the widely accepted proposition of £26,000 as the average annual salary in the UK. This will enable you to develop an alternative hypothesis:

H_a : Average annual salary in the UK is not equal to £26,000.

The following activity will test your knowledge of null and alternative hypotheses.

Activity 1 Null hypothesis versus alternative hypothesis



Allow approximately 10 minutes to complete this activity

Read the following statements. Can you develop a null hypothesis and an alternative hypothesis?

'It is believed that a high-end coffee machine produces a cup of caffè latte with an average of 1 cm of foam. The hotel employee claims that after the machine has been repaired, it is no longer able to produce a cup of caffè latte with 1cm foam.'

Provide your answer...

.....

Discussion

H_0 : a coffee machine makes a cup of caffè latte with 1cm foam on average.

H_a : a coffee machine cannot make a cup of caffè latte with 1 cm foam on average.

If you have developed the hypotheses H_0 and H_a as mentioned in the discussion to Activity 1, you have shown that you are familiar with the structure of different types of hypotheses. However, in the next section you will explore the concept of hypothesis formulation further.

1.2 Population mean (μ)

As you have seen so far in Section 1, the formulation of H_0 is based on widely accepted beliefs, which can also be interpreted as the population mean (μ), representing the average value in the population (where μ is the symbol for a population mean).

To illustrate this point, consider the example of the average UK salary, which is 'widely believed' to be £26,000 per year. In this case, μ is equivalent to £26,000, which can be expressed as:

$H_0: \mu = £26,000$

As H_0 and H_a are always opposite (i.e. since the purpose of H_a is to challenge the belief of H_0), you can also express H_a as:

$H_a: \mu \neq £26,000$.

If you revisit the caffè latte foam example from Activity 1, H_0 and H_a can also be expressed as:

$H_0: \mu = 1\text{cm foam}$ (representing the population mean for the height of foam in a caffè latte)

$H_a: \mu \neq 1\text{cm foam}$ (indicating a departure from the widely accepted belief about the population mean).

In summary, the process of hypothesis formulation is a critical step in business decision making. It involves setting up the null hypothesis (H_0), which represents the widely accepted belief or the status quo (or μ), and the alternative hypothesis (H_a), which challenges the belief and suggests the possibility of a significant difference. A clear understanding of the concept of hypothesis formulation is crucial for accurate business decision making.

Having discussed how hypotheses are formulated, you will now turn your attention to how to test them. In the following sections, you will look at test statistics and alpha levels, which play a critical role in determining the statistical significance of a hypothesis test.

2 Testing with data

Once hypotheses have been developed, the next step is to find existing data or design a study (to generate some primary data) to test them. The results of the testing will have two possible outcomes.

1. Reject null hypothesis (H_0) – this also means that the alternative hypothesis (H_a) is accepted.
2. Fail to reject null hypothesis (H_0) – this means that you accept H_0 and reject alternative hypothesis (H_a).

You can test hypotheses empirically using test statistics that involve calculating sample data and using the results to decide whether you need to reject the null hypothesis or fail to reject the null hypothesis.

A 'test statistic' is a number calculated by a statistical test, which provides information about how much the relationship between variables in the test differs from the null hypothesis. There are different types of test statistics and the choice of test statistic depends on the type of hypothesis being tested and the nature of the data. For example, if the data is continuous and normally distributed, the t-test or z-test may be used as a test statistic. If the data is categorical, the chi-squared test may be used. Some of these will be discussed later in this course.

To begin, you will first obtain an overview of the general guidelines for using the test statistic to determine whether to reject or fail to reject the null hypothesis. Consider the coffee example from Activity 1. In order to test the null hypothesis ($H_0: \mu = 1\text{cm foam}$), you might sample 60 cups of caffè latte made by the coffee machine – i.e. select 60 cups of caffè latte at random (random sampling) from all the cups of caffè latte that the coffee machine made throughout the day. Then, you would measure the height of the foam to obtain the average height and calculate the test statistic.

You may be wondering why here the null hypothesis and not the alternative hypothesis ($H_0: \mu \neq 1\text{cm foam}$) is being exclusively discussed. In the realm of hypothesis testing, you commence by assuming that the null hypothesis is true and then employ sample data to determine whether or not to reject it. If you reject the null hypothesis, you can infer that the alternative hypothesis is true. As you may recall, the null and alternative hypotheses are complementary and mutually exclusive. The null hypothesis denotes the prevailing belief, whereas the alternative hypothesis represents what you aim to demonstrate. Therefore, rejecting the null hypothesis implies that you have evidence to support the alternative hypothesis. It is noteworthy that you cannot directly prove the alternative hypothesis. Instead, you can only reject or fail to reject the null hypothesis. The decision to reject or fail to reject the null hypothesis is grounded in the test statistics.



Figure 4 Caffè latte

For the study:

- Researcher 1 samples 60 cups of caffè latte and gets the average foam height equal to 1.1 cm.
- Researcher 2 samples 60 cups of caffè latte and gets the average foam height equal to 1.5 cm.
- Researcher 3 samples 60 cups of caffè latte and gets the average foam height equal to 2.6 cm.

The question now is how to make your decision. Using only the data collected by Researcher 1, the average height of the foam is 1.1 cm. It is not exactly equal to 1 cm, but it is very close. In this case, you may say that you cannot reject the null hypothesis. However, if you look at the data collected by Researcher 3 (the average height of the foam is equal to 2.6 cm), this is far beyond the commonly accepted idea of 1 cm foam for a caffè latte. Therefore, you can reject the null hypothesis and accept the alternative hypothesis ($\mu \neq 1$ cm foam). Considering the data collected by Researchers 1 and 3, it is easy to decide whether to reject the null hypothesis or not. However, based on the data collected by Researcher 2, it is extremely difficult to make a decision. Although the average foam height of 1.5 cm is not far away from 1 cm, does it go far enough to be considered sufficient?

In order to answer this question, you would need to introduce the concept of 'statistical significance', which you will look at in more detail in the next section.

3 Alpha (α) levels

The level of statistical significance is the threshold at which you decide whether to reject the null hypothesis (to make a decision). A result is statistically significant when the situation described in the null hypothesis is highly unlikely to have occurred.

By using the concept of statistical significance, you can have a concrete way of examining the claim concerning the null hypothesis, using the data collected, to make a clear decision on when to reject the null hypothesis and when to leave it. In this way, you do not have to guess whether the test statistic is too high or too low.

How can you determine the appropriate level of statistical significance to use? In order to answer this question, you must introduce the concept of 'level of confidence' (C) – how confident are you in your decision to reject the null hypothesis? C is related to the concept of 'confidence intervals' that you may have seen elsewhere. C is usually set at 95% when establishing a confidence interval. Having determined C , you can calculate the level of statistical significance that you want to use as $1 - C$.

If you choose the level of confidence (C) at 95%, then the level of statistical significance that you set up for your decision is equal to $1 - 95\% = 5\%$.

The level of statistical significance is referred to as 'alpha' or ' α '.

So, returning back to the question on caffè latte foam, is 1.5 cm average foam height, which was only slightly different from the expected value of 1 cm, significant enough to reject or fail to reject the null hypothesis (H_0)? The researchers had to consider the alpha level to determine the appropriate course of action. The alpha level is a crucial factor in hypothesis testing and is used to determine the threshold for statistical significance. In other words, it helps to determine whether the results obtained are due to chance or if they represent a genuine difference between the expected and observed values.

Calculating the alpha level can be tricky, but it is a necessary step in hypothesis testing (you will discuss how to calculate the alpha level in later sections). Once the alpha level has been determined, it can be used to decide whether to accept or reject the null hypothesis based on the findings of Researcher 2. The decision of whether to reject or fail to reject H_0 based on the findings of Researcher 2 will ultimately depend on the alpha level chosen, so the researchers will have to carefully consider this factor before drawing any conclusions from the study.

You will be tested on your understanding of how to transfer levels of confidence to alpha levels in the following activity.

Activity 2 Level of alpha



Allow approximately 10 minutes to complete this activity

Imagine you are working as a marketing analyst for a new product launch, and you want to make sure that your market research is accurate and reliable. You need to determine the appropriate alpha level for your survey results, based on the desired confidence level.

Can you determine α ?

Table 1 Level of alpha

Level of Confidence (C)	α
90%	<input type="text" value="Provide your answer..."/>
95%	<input type="text" value="Provide your answer..."/>
99%	<input type="text" value="Provide your answer..."/>

Discussion**Table 1 Level of alpha
(completed)**

Level of Confidence (C)	α
90%	10%
95%	5%
99%	1%

α is calculated as $1 - C$.

$$1 - 90\% = 10\%$$

$$1 - 95\% = 5\%$$

$$1 - 99\% = 1\%$$

There is an inverse relationship between confidence levels and alpha levels, as increasing the confidence level leads to a decrease in the alpha level, and vice versa. It is important to carefully consider the appropriate alpha level and confidence level when conducting hypothesis testing to ensure that the appropriate level of risk and certainty are balanced.

4 One-tailed vs two-tailed test

To gain a deeper understanding of how to conduct a hypothesis test, this section will delve into the concepts of one-tailed and two-tailed tests. These tests are vital tools in statistical hypothesis testing, and the decision of which test to employ depends on the research question and hypothesis under examination. It is crucial to give careful thought to the suitable type of test to ensure that the hypothesis is thoroughly tested and precise conclusions are derived from the data. This section will elaborate on this topic in greater detail.

To commence, complete the following activity pertaining to the formulation of null and alternative hypotheses. This exercise may be somewhat challenging, but it serves as an excellent introduction to upcoming discussions – don't be concerned if you find it difficult!

Activity 3 Hypotheses setting

 Allow approximately 10 minutes to complete this activity

Read the following statements and then develop a null hypothesis and an alternative hypothesis.

'It is believed that OU students need to set aside no longer than, on average, 15 hours to study an entire session of an OU course. However, a researcher believes that OU students spend longer studying an entire session of an OU course.'

Provide your answer...

Discussion

H_0 : OU students spend, on average, no more than 15 hours studying an entire session of OU course.

H_a : OU students spend, on average, more than 15 hours studying an entire session of OU course.

They can also be written as:

$H_0: \mu \leq 15$ hours studies

$H_a: \mu > 15$ hours studies

μ is a symbol for a population mean. Remember, H_0 and H_a are always opposites.

Did you identify any differences between the hypotheses you developed in Activity 1 and Activity 3? The set of hypotheses in Activity 1 has an equal (=) or not equal (\neq) supposition (sign) in the statement. However, in Activity 3, the set of hypotheses has less than or equal to (\leq) and greater than ($>$) supposition (sign) in the statement. This creates different conditions that lead to acceptance or rejection of the null hypothesis.

4.1 The normal distribution

Here, you will look at the concept of normal distribution and the bell-shaped curve. The peak point (the top of the bell) represents the most probable occurrences, while other possible occurrences are distributed symmetrically around the peak point, creating a downward-sloping curve on either side of the peak point.




Figure 5 The peak point is average

In order to test hypotheses, you need to calculate the test statistic and compare it with the value in the bell curve. This will be done by using the concept of 'normal distribution'.

A normal distribution is a probability distribution that is symmetric about the mean, indicating that data near the mean are more likely to occur than data far from it. In graph form, a normal distribution appears as a bell curve. The values in the x-axis of the normal distribution graph represent the z-scores. The test statistic that you wish to use to test the set of hypotheses is the **z-score**. A z-score is used to measure how far the observation (sample mean) is from the 0 value of the bell curve (population mean). In statistics, this distance is measured by standard deviation. Therefore, when the z-score is equal to 2, the observation is 2 standard deviations away from the value 0 in the normal distribution curve.




Figure 6 Normal distribution – bell curve

4.2 Two-tailed tests

Hypotheses that have an equal ($=$) or not equal (\neq) supposition (sign) in the statement are called **non-directional hypotheses**. In non-directional hypotheses, the researcher is interested in whether there is a statistically significant difference or relationship between two or more variables, but does not have any specific expectation about which group or variable will be higher or lower. For example, a non-directional hypothesis might be: 'There is a difference in the preference for brand X between male and female consumers.' In this hypothesis, the researcher is interested in whether there is a statistically significant

difference in the preference for brand X between male and female consumers, but does not have a specific prediction about which gender will have a higher preference. The researcher may conduct a survey or experiment to collect data on the brand preference of male and female consumers and then use statistical analysis to determine whether there is a significant difference between the two groups.

Non-directional hypotheses are also known as two-tailed hypotheses. The term 'two-tailed' comes from the fact that the statistical test used to evaluate the hypothesis is based on the assumption that the difference or relationship could occur in either direction, resulting in two 'tails' in the probability distribution. Using the coffee foam example (from Activity 1), you have the following set of hypotheses:

$H_0: \mu = 1\text{cm foam}$

$H_a: \mu \neq 1\text{cm foam}$

In this case, the researcher can reject the null hypothesis for the mean value that is either 'much higher' or 'much lower' than 1 cm foam. This is called a **two-tailed test** because the rejection region includes outcomes from both the upper and lower tails of the sample distribution when determining a decision rule. To give an illustration, if you set alpha level (α) equal to 0.05, that would give you a 95% confidence level. Then, you would reject the null hypothesis for obtained values of $z < -1.96$ and $z > 1.96$ (you will look at how to calculate z-scores later in the course).

This can be plotted on a graph as shown in Figure 7.




Figure 7 Two-tailed test

In a two-tailed hypothesis test, the null hypothesis assumes that there is no significant difference or relationship between the two groups or variables, and the alternative hypothesis suggests that there is a significant difference or relationship, but does not specify the direction of the difference or relationship.

When performing a two-tailed test, you need to determine the level of significance, which is denoted by alpha (α). The value of alpha, in this case, is 0.05. To perform a two-tailed test at a significance level of 0.05, you need to divide alpha by 2, giving a significance level of 0.025 for each distribution tail ($0.05/2 = 0.025$). This is done because the two-tailed test is looking for significance in either tail of the distribution. If the calculated test statistic falls in the rejection region of either tail of the distribution, then the null hypothesis is rejected and the alternative hypothesis is accepted. In this case, the researcher can conclude that there is a significant difference or relationship between the two groups or variables.

Assuming that the population follows a normal distribution, the tail located below the critical value of $z = -1.96$ (in a later section, you will discuss how this value was determined) and the tail above the critical value of $z = +1.96$ each represent a proportion of 0.025. These tails are referred to as the lower and upper tails, respectively, and they

correspond to the extreme values of the distribution that are far from the central part of the bell curve. These critical values are used in a two-tailed hypothesis test to determine whether to reject or fail to reject the null hypothesis. The null hypothesis represents the default assumption that there is no significant difference between the observed data and what would be expected under a specific condition.

If the calculated test statistic falls within the critical values, then the null hypothesis cannot be rejected at the 0.05 level of significance. However, if the calculated test statistic falls outside the critical values (orange-coloured areas in Figure 7), then the null hypothesis can be rejected in favour of the alternative hypothesis, suggesting that there is evidence of a significant difference between the observed data and what would be expected under the specified condition.

4.3 One-sided tests

As well as non-directional hypotheses, you will also encounter hypotheses that have a less than or equal to (\leq) and greater than ($>$) supposition (sign) in the statement (as you saw in Activity 3). This is called a **directional** hypothesis. A directional hypothesis is a type of research hypothesis that aims to predict the direction of the relationship or difference between two variables. Essentially, it specifies the anticipated outcome of a study prior to the collection of data.

For example, a directional hypothesis might propose that a marketing campaign will increase product sales, predicting the direction of the relationship (i.e. the marketing campaign will lead to an *increase* in product sales). In contrast, a non-directional hypothesis simply states that there is a relationship between two variables without specifying the direction of that relationship, such as: 'There is a relationship between the marketing campaign and product sales.'

Directional hypotheses are often preferred in scientific research because they provide a more precise and focused prediction than non-directional hypotheses. In business management, a directional hypothesis can also be a useful tool. For example, a company may use a directional hypothesis to design a study that examines the effectiveness of a marketing campaign in enhancing sales. This approach provides a clearer understanding of the impact of the campaign and enables the company to make more informed decisions about future marketing strategies.

A one-tailed test is a statistical test employed to evaluate a directional hypothesis, which predicts the direction of the difference or association between two variables. Its objective is to ascertain if the data supports the anticipated direction.

To illustrate, consider the hypotheses from Activity 3:

H_0 : $\mu \leq 15$ hours of studies

H_a : $\mu > 15$ hours of studies

The null hypothesis (H_0) posits that the population mean (μ) is less than or equal to 15 hours of studies, while the alternative hypothesis (H_a) predicts that the population mean is greater than 15 hours of studies.

To conduct a one-tailed test, a critical value must be established to determine whether the null hypothesis should be rejected or retained. Typically, a significance level (α) is set for this purpose. For instance, assuming $\alpha = 0.05$, the z-score for a one-tailed test with $\alpha = 0.05$ in a normal distribution is 1.645. Consequently, the null hypothesis would be rejected if the z-score exceeds 1.645. In other words, only the upper tail region of the distribution is rejected for a one-tailed test. Additionally, you employ distinct z-scores since, in contrast to a two-tailed test, the alpha level does not need to be divided by two. In a normal distribution, the area in the tail above $z = +1.645$ represents 0.5 of the distribution. This

portion of the distribution is significantly remote from the centre of the bell curve at 0. Consequently, the null hypothesis would be rejected if the z-score exceeds 1.645 (as depicted in Figure 8).




Figure 8 One-tailed test – right (upper) tail

In summary, a one-tailed test is used to assess a directional hypothesis in which the direction of the difference or association between two variables is predicted. The critical value for a one-tailed test is determined by the selected significance level (α), and the test is conducted to ascertain whether the data supports the predicted direction.

In addition, the one-tailed test is not limited to a single direction (greater than) but can also be employed in the opposite direction (less than). An example can be used to illustrate this type of hypothesis testing. Consider a situation where the management team believes that the average amount spent by customers during their visits to a department store is £65. However, the service manager observes that customers spend less than that amount during their visits. In this case, you can formulate the following set of hypotheses:

$$H_0: \mu \geq £65$$

$$H_a: \mu < £65$$

To test this directional hypothesis, a one-tailed test must be conducted. The alternative hypothesis states that the specific value of μ will be lower than the value specified in the hypothesis. Therefore, you must reject the region in the lower tail of the normal distribution. More specifically, the rejection region of the one-tailed test at alpha levels equals 0.05. The lower tail of the normal distribution has a z-score lower than -1.645. Any hypothesis in this region will be rejected. The graph in Figure 9 illustrates this.




Figure 9 One-tailed test – left (lower) tail

In conclusion, the one-tailed test is not restricted to a specific direction and can be used in either direction, depending on the research question and the hypothesis being tested. The

test is used to determine if the data supports a directional hypothesis, and a critical value is established based on the significance level chosen for the test.

4.4 Check your understanding

In the following activity, you will be tested on your understanding of one-tailed and two-tailed tests with respect to a bell-shaped curve.

Activity 4 One-tailed or two-tailed test

 Allow approximately 10 minutes to complete this activity

Can you match the statements with the bell curves that indicate either a one-tailed or two-tailed test?

Interactive content is not available in this format.



In conclusion, the choice between a one-tailed test and a two-tailed test (as well as between a directional and non-directional hypothesis) depends on the research question and the nature of the relationship or difference being studied. Two-tailed tests and non-directional hypotheses are useful when the research question is exploratory, or the direction of the relationship is unknown, while one-tailed tests and directional hypotheses are useful when a specific direction of the relationship or difference is predicted and is of interest. Careful consideration and planning are necessary when choosing between these options to ensure that the hypothesis is properly tested and the correct conclusions are drawn from the data. Ultimately, the choice of the type of test and hypothesis should be guided by the research question, the theory behind the study, and the available data.

5 Mean and z-score ranges

In the previous section on hypothesis testing using the normal distribution, the z-score was frequently mentioned. This is because the z-score is the test statistic that is used to determine whether the null hypothesis should be accepted or rejected.

One way to gain an understanding of the calculated z-score and its alpha value is through the use of a z-score table. Using Excel functions, you can create a table that indicates z-scores and the corresponding area under the normal distribution curve. As you continue through the course, you will see how the z-score table can be used in hypothesis testing. For now, however, you will focus your attention on the use of Excel to create a table of z-scores, using the following steps.

Step 1: In an Excel spreadsheet, you can create a z-score table yourself that shows the z-score associated with any level of probability. To give an illustration, the screenshot below (Figure 10) shows a table with the values – 0.0; 0.1; 0.2; 0.3; 0.4; 0.5 – in a row and the values – 0.00; 0.01; 0.02; 0.03; 0.04; 0.05; 0.06; 0.07; 0.08; 0.09 – in columns. These are the first and second decimal digits of the probability whose z-score will be in the table. For example, the z-score in the first row (0.0) and second column (0.01) in the as-yet empty table refers to the probability $p=0.01$ (i.e., $0.0+0.01=0.01$). The table is really just a long list for probabilities from $p=0$ to $p=0.5$ formatted as a table.

Z Score	0.00	0.01	0.02	0.03	0.04	0.05	0.06	0.07	0.08	0.09
0.0										
0.1										
0.2										
0.3										
0.4										
0.5										

Figure 10 Creating a z-score table step 1

Step 2: To use the 'NORM.S.DIST' Excel formula, begin by typing '=NORM.S.DIST(' into the desired cell. The Excel software will prompt you to complete the formula's value entry by entering the appropriate values for 'z' and 'cumulative'.

B C D E F G H I J K L M										
			NORM.S.DIST(z, cumulative)							
Z Score	0.00	0.01	0.02	0.03	0.04	0.05	0.06	0.07	0.08	0.09
0.0	.DIST(
0.1										
0.2										
0.3										
0.4										
0.5										

Figure 11 Creating a z-score table step 2

Step 3: After initiating the NORM.S.DIST formula in the designated cell, you can assign a value to 'z' by adding the value in the row cell to the value in the column cell.

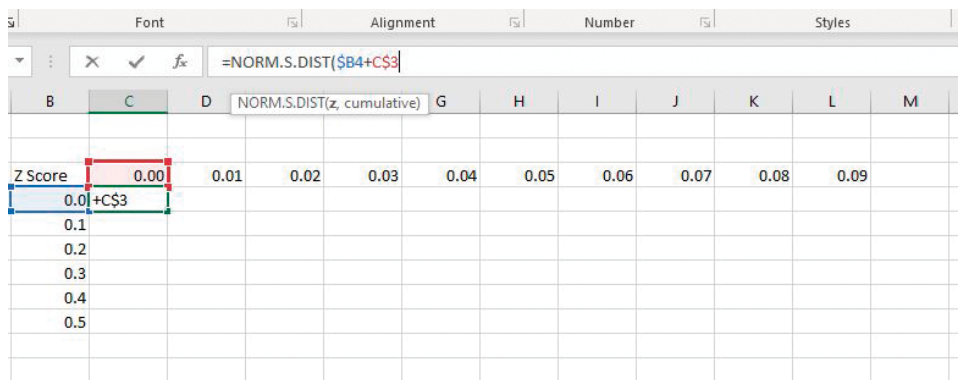


Figure 12 Creating a z-score table step 3

Step 4: To indicate the cumulative distribution function, set the 'cumulative' argument to 'TRUE'.

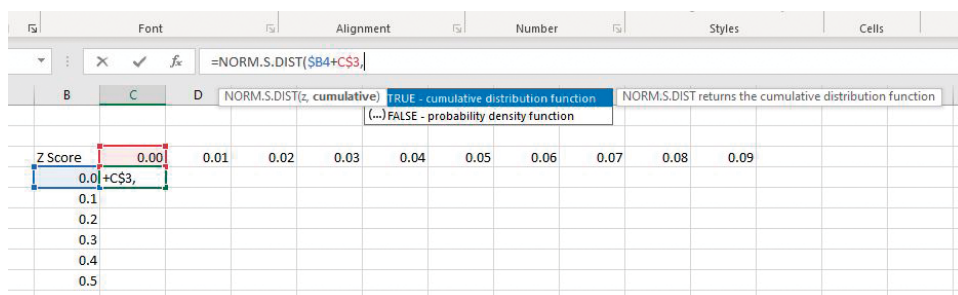


Figure 13 Creating a z-score table step 4

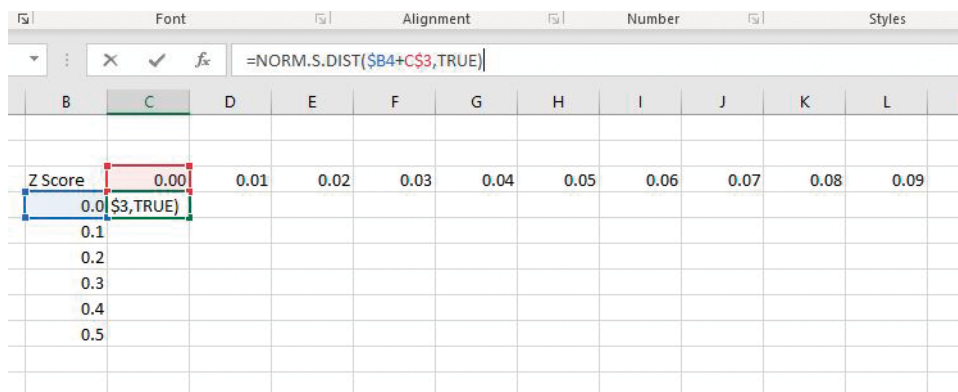


Figure 14 Creating a z-score table step 4

Step 5: After completing the formula with the appropriate values, press 'Enter' to calculate the result for the selected cell (Figure 15). To apply the Excel function to all cells within the table, click and drag the green box to cover the desired area (Figure 16).

Font Alignment Number Styles											
=NORM.S.DIST(\$B4+C\$3,TRUE)											
B	C	D	E	F	G	H	I	J	K	L	M
Z Score	0.00	0.01	0.02	0.03	0.04	0.05	0.06	0.07	0.08	0.09	
0.0	0.5000										
0.1											
0.2											
0.3											
0.4											
0.5											

Figure 15 Creating a z-score table step 5

Font Alignment Number Styles											
=NORM.S.DIST(\$B4+C\$3,TRUE)											
B	C	D	E	F	G	H	I	J	K	L	
Z Score	0.00	0.01	0.02	0.03	0.04	0.05	0.06	0.07	0.08	0.09	
0.0	0.5000	0.5040	0.5080	0.5120	0.5160	0.5199	0.5239	0.5279	0.5319	0.5359	
0.1	0.5398	0.5438	0.5478	0.5517	0.5557	0.5596	0.5636	0.5675	0.5714	0.5753	
0.2	0.5793	0.5832	0.5871	0.5910	0.5948	0.5987	0.6026	0.6064	0.6103	0.6141	
0.3	0.6179	0.6217	0.6255	0.6293	0.6331	0.6368	0.6406	0.6443	0.6480	0.6517	
0.4	0.6554	0.6591	0.6628	0.6664	0.6700	0.6736	0.6772	0.6808	0.6844	0.6879	
0.5	0.6915	0.6950	0.6985	0.7019	0.7054	0.7088	0.7123	0.7157	0.7190	0.7224	

Figure 16 Creating a z-score table step 5

In the following activity, you can practice using these Excel formulas to create a table of values of the normal distribution that correspond to a range of z-scores. These will be useful when testing hypotheses later.

Activity 5 Z-Score table

 Allow approximately 20 minutes to complete this activity

Using steps 1 to 5 described above, create a z-score table in Excel for z-scores ranging from -3 to 3.

Once you have completed these steps, reveal the discussion and compare your answers.

Discussion

The following table illustrates a z-score that ranges from -3 to 3. The following is what you should have accomplished as a result of this activity.

