

Understanding statistical terms: 1

An increasing number of statistical terms appear in journal articles and other medical information. A working knowledge of these is essential in assessing clinical evidence. With this in mind, we are producing a series of explanatory articles covering various statistical terms and their uses. This, the first article in the series, will focus on some of the most common terms used in randomised controlled trials.¹

Describing a population or sample

Normal distribution

The normal distribution is the pattern for the distribution of a set of data which follows a bell-shaped curve. The x-axis represents the variable of interest (e.g. height or weight) and the y-axis shows how many people have a certain value of the variable. The curve peaks at the middle variable (the mean value; see below) and decreases symmetrically on either side. The shape of this distribution is shown below:




Figure 1. The normal distribution

Skewed distribution

If a distribution is skewed, it is not symmetrical but one of the two 'tails' that lie on either side of the mean value is longer than the other.

Mean

This is a way of describing 'the average'; it is derived by adding all the values and dividing by the total number of values. The mean of a population is written as ' μ ' (Greek letter mu) and the mean of a sample is written as ' \bar{x} ' ('x-bar').

Median

The median is an 'average' that may be most useful when values are not distributed evenly (e.g. in a skewed distribution); it is derived by listing the population or sample values in size order and identifying the middle one of the list. This means that there will be equal numbers of values that are higher and lower than the median value.

Mode

The mode is the value appearing most frequently in a distribution (i.e. the single most 'popular' value); it is another way of describing an 'average' value. It can be particularly useful when the data is skewed.

Range

The range is the difference between the smallest and the largest value in a population or sample distribution. It is therefore a measure of how 'spread-out' the values are. However, it only takes into account two of the values and therefore may, if taken in isolation, give a misleading impression of the nature of the population distribution. This could occur if it is wrongly assumed that the two values come from a normal distribution when, in reality, the distribution is markedly skewed.

Interquartile range

The interquartile range (IQR) is the numerical difference between the values that are positioned a quarter of the way up from the bottom of the distribution and a quarter of the way down from the top of the distribution. It is usually used when the values are not distributed evenly and when the median value is reported. The two values that define the IQR encompass the middle 50% of all the values in the distribution. Therefore, IQR gives an indication of how spread out the values are.

Standard deviation

The standard deviation (SD) is another way of describing how 'spread-out' the values are in a distribution, but unlike the range, it takes into account all the values rather than only the two ends of the distribution. If many data values are close to the mean, then the standard deviation is small; if many data values are far from the mean, then the standard deviation is large. If all the data values are identical, then the standard deviation is zero, indicating that there is no deviation at all from that one value. The standard deviation of the whole population is sometimes written as ' σ ' (Greek letter lower case sigma) and the standard deviation of a sample as ' s '.

The standard deviation of a population is calculated a formula (see box) that involves subtracting the mean value (μ) from each data value (x); squaring each of these differences; summing all these squared differences (the Greek upper case sigma, Σ , means 'sum all the values of the expression that follows it'); dividing this sum by $(n-1)$ where n is the number of data values; and taking the square root of the answer.

Formula for calculating the standard deviation of a population

$$\sigma = \sqrt{\frac{\sum [x-\mu]^2}{n}}$$

Where:
 σ = standard deviation
 Σ = sum of all the values of the expression that follows it
 $(x - \mu)$ = individual data value subtracting the mean value of the population
 n = number of data values

One potential use of the standard deviation is to convert a normal distribution into the so-called 'standard normal distribution' (a normal distribution with a mean of 0 and a standard deviation of 1). Normal distributions can be transformed to standard normal distributions by using a formula (see box).

Formula for calculating the standard normal distribution

Where:

$$z = \frac{x - \mu}{\sigma}$$

z = value in the standard normal distribution (also known as 'z-score')

x = data value from the original normal distribution

μ = mean of the original normal distribution

σ = standard deviation of the original normal distribution

In a standard normal distribution, a positive z-score represents the number of standard deviations the value lies above the mean, whereas a negative z-score represents the number of standard deviations the value lies below the mean.

Incidence

Incidence refers to the number of new cases of an event or phenomenon occurring in a specified time period, and is often expressed as a proportion of the relevant population. For example, if 10 new cases of angina are diagnosed each year in a population of 100,000 people, the incidence is 10 per 100,000 per year.

Prevalence

Prevalence refers to the proportion of a population that has a particular condition or characteristic at a specific point in time (point prevalence), or during a certain time period (period prevalence). For example, if 1 in 100 people in a population study have diabetes mellitus at a particular time, the point prevalence of the condition is 1%. If 1 in 50 people are recognised as having diabetes mellitus during a 3 month period (i.e. both previously known and newly diagnosed cases), the period prevalence of this condition is 2%.

Terms relating to clinical trial design

Null hypothesis

A null hypothesis supposes that the treatments being compared in a study are equally effective, for example, a new drug is no better than the old one or that a new drug is no better than placebo. If the results of the study allow rejection of this hypothesis (with a specific probability), this suggests that there is true a difference between the two drugs.

Sample size

The sample size is the number of people that are included in a study. Ideally, the appropriate sample size for a study should be chosen at the design stage to ensure that the study is big enough to have a realistic chance of finding evidence to support or reject the null hypothesis.

Power

The power of a study is a measure of how likely it is to be able to find a certain size of difference (deemed clinically important) between the groups being compared, assuming such a difference exists. In general, the larger the study, the greater the power. A study with too few people in it is underpowered, so it is unlikely to be able to provide convincing evidence on whether or not there is a real difference between the treatments being compared. Power also depends on how large a difference is expected between groups. For example, if only a small extra benefit is expected for one drug over another, more people will be required to achieve the same power to detect the difference than if a very large difference is expected between the treatments. Study power is usually set at 80%, which means that there is a 20% likelihood of missing a real difference between the two groups being compared.

Confounding factors

A confounding factor is something that accounts for a particular outcome in a study but is also related independently to another factor that has no such causal influence. Focusing exclusively on this second, non-causal factor could then give the incorrect impression that it is directly responsible for the outcome, because the effect of the confounding factor (the real explanation) is either not known or has been overlooked. For example, if all the men in a study were given one drug and all the women given the other, the results might be wholly or partly due to the differences between the groups (gender) rather than the effect of the treatment each group receives. If so, gender is said to be a confounding factor or confounder. Randomisation (see below) aims to distribute potential confounding factors roughly equally between the groups before the treatment is given (e.g. roughly the same percentage of males in each treatment group). Assuming randomisation is successful, any difference between the outcomes of the groups at the end of the study can be reasonably attributed to the difference between the treatments, rather than the confounding factor. Confounders can be overcome by stratified randomisation or minimisation (see below).

Bias

Bias is any distorting influence that makes the results of a study an inaccurate reflection of the true situation. For example, if all the new patients attending a clinic are given a new drug, and their results are compared with everyone attending the clinic the previous year (new and longer-standing patients), differences in the outcome could be due to the inclusion of patients with chronic disease in only one of the groups, (so-called 'selection bias'), rather than a real effect of the new drug.

Blinding

If the people measuring the outcome in a trial know which treatment participants had, this may influence how they perform such assessment. For example, in a trial of a new antihypertensive drug, the assessors might round down the blood pressure of the patients they know have been on

this treatment, and round up the measurements of those who have not had this therapy. 'Blinding' (also known as 'masking') of outcome assessors in order to maximise the apparent benefit of this new treatment (i.e. making sure they do not know which treatment the patient has had) helps to eliminate this source of bias. Similarly, if patients know which treatment they have had, this might influence how they report outcomes. For example, a patient might be more inclined to report a benefit if they know they have had an active treatment, and to report no difference if they know they have had a placebo. 'Blinding' the patient to their treatment allocation can help in eliminating such biases.

If both the patient and the person measuring the outcome do not know the patient's treatment allocation, the study is said to be double-blind. If only one person (patient or outcome assessor) does not know but the other person does, the study is said to be single-blind. This may occur, for example, if the patient was treated with intravenously infused or oral antibacterial therapy (the patient would know if they had had a drip or not, but this information could be concealed from the outcome assessor to reduce the likelihood of biased assessment). In other circumstances, the person administering a treatment or assessing its outcome might know whether the patient has been given a real or a sham treatment (say, real or sham laser therapy); the patient could be blinded to the treatment to help avoid biasing the results.

Randomisation

Randomisation is a method of allocating patients to different treatments in a trial by chance rather than by specific patient characteristics. Done correctly, randomisation should ensure that, on average, the patients in each group have similar distributions of those characteristics that could affect the outcome (e.g. age, gender, other potential confounding factors); this is more likely the larger the sample size.

Stratified randomisation involves the process of separating the population into groups based on whether or not they have a potential confounder. These stratified groups are then randomised separately to the trial or treatments, so

ensuring that the resulting treatment groups each contain the same proportion of patients with the confounding factor. Consequently, any difference in results between these groups is likely to represent an effect of the treatment rather than of the confounder.

Methods of randomisation should be truly random (i.e. the sequence of allocation to treatments cannot be predicted in advance). For example, assigning patients to two groups alternately is not true randomisation as the group to which the next patient is to be allocated can be predicted from the previous allocation.

Minimisation

Minimisation is a method of allocating patients into treatment groups in a trial not randomly, but instead using means to reduce deliberately any imbalance between the groups on various characteristics that could be confounding factors. For example, if the next patient to be allocated in a study is female, and one treatment group already has more females, the patient would be allocated to the group with fewer females to make the groups more balanced on this criterion. This technique is especially useful if the sample size is small since, in these circumstances, using simple randomisation techniques could lead to differences between the groups on important confounding factors. It may also be used in large trials when there are many factors that need to be balanced.

[R=randomised controlled trial; M=meta-analysis]

1. Lang TA, Secic M. *How to report statistics in medicine: annotated guidelines for authors, editors, and reviewers*. Second edition. American College of Physicians: Philadelphia, 2006.

DOI: 10.1136/dtb.2009.01.0003



Editor: Ike Iheanacho BSc, MB BS.
Deputy Editor: Jenny Blythe MBChB, MRCP.
Associate Editors:
 Helen Barnett BPharm, MSc, LicAc, MBAcC; Jacoby Patterson MA Cantab, MB BChir, FFPH, MD; Andrea Tarr BPharm, MSc, MRPharmS.
Business Manager: Donna Walton BA, MPhil, DipM, ACIM.
Editorial Assistant: Michael Allen BSc.
Administrator: Ashling Connolly.
Contributing Editor: Dr Teck Khong MD, MRCP, St George's University Hospital, London.

Advisory Council: Martin Duerden B Med Sci, MRCP, Dip Ther, DPH, Conwy Local Health Board and Wales College of Medicine, Cardiff University; Wing May Kong PhD, MRCP, Imperial College & Central Middlesex Hospital, London; Prof. Hywel C. Williams BSc, MSc, FRCP, PhD, University Hospital, Nottingham.

Editorial Board: Prof. Ian Anderson MB BS, MD, MRCP, FRCPsych, University of Manchester & Manchester Mental Health & Social Care Trust; Peter Burrill B Pharm, Dip Pres Sci, MRPharmS, FFPMM, FCPP, Derbyshire County PCT; Paul Caldwell MB BS, MRCP, Goole Health Centre; Ernest H.S. Choy MD, MRCP, GKT School of Medicine, King's College London; Jo Congleton MA, MD, MRCP, Worthing General Hospital; Joanna Girling MA, MBBS, MRCP, FRCOG, West Middlesex University Hospital, London; Prof. Aroon Hingorani MA, PhD, FRCP, University College London; Sean Kelly MB, ChB, MD, FRCP, York Hospital; Tony Marson MB ChB, MD, FRCP, University of Liverpool & the Walton Centre; Mahdad Noursadeghi MB BS, MRCP, PhD, University College London; George Rylance MB.

Editorial correspondence: DTB, BMJ Group, BMA House, Tavistock Square, London WC1H 9JR, UK



© BMJ Group 2009. Material published in DTB may not be used in any form of advertising, sales or publicity, nor be reproduced without written permission. DTB is commissioned, researched and published by BMJ Group, BMA House, Tavistock Square London WC1H 9JR, UK. Printed by Hobbs the Printers Ltd, Southampton, Hampshire.

Email: dtb@bmjgroup.com Website: www.dtb.bmj.com
Changes of addresses and subscription/distribution enquiries:
 All changes of address and orders to DTB, Fulfillment Customer Service, BMJ Group, BMA House, Tavistock Square, London WC1H 9JR, UK. Tel: +44(0)20 7383 6270. Fax: +44(0)20 7383 6402.
 E-mail: support@bmjgroup.com.

Annual individual print plus online subscription £86. For other subscription rates, please email subscriptions@bmjgroup.com for more details. Retired doctors and pharmacists, students £43.
 Single issues are available—please contact subscriptions@bmjgroup.com for more information.