

1: Power and Sample Size Calculators | HyLown

Sample size calculation Example Consider a population with proportion p . Let X be the number of successes in a random sample of size n with model $X \sim \text{Binomial}(n;p)$.

Underlying event rate in the population Standard deviation in the population. Some more factors that can be considered while calculating the final sample size include the expected drop-out rate, an unequal allocation ratio, and the objective and design of the study. To put it in different words, we are willing to accept the detection of a difference 5 out of times when actually no difference exists i. POWER Sometimes, and exactly conversely, we may commit another type of error where we fail to detect a difference when actually there is a difference. This is called the Type II error that detects a false negative difference, as against the one mentioned above where we detect a false positive difference when no difference actually exists or the Type I error. We must decide what is the false negative rate we are willing to accept to make our study adequately powered to accept or reject our null hypothesis accurately. The power of a study increases as the chances of committing a Type II error decrease. If the average weight loss following one diet program is 20 kg and following another is 10 kg, the absolute effect size would be 10 kg. In statistics, the difference between the value of the variable in the control group and that in the test drug group is known as effect size. This difference can be expressed as the absolute difference or the relative difference, e . We can estimate the effect size based on previously reported or preclinical studies. It is important to note that if the effect size is large between the study groups then the sample size required for the study is less and if the effect size between the study groups is small, the sample size required is large. In the case of observational studies, for example, if we want to find an association between smoking and lung cancer, since earlier studies have shown that there is a large effect size, a smaller sample would be needed to prove this effect. This unlike the level of significance and power is not selected by convention. Rather, it is estimated from previously reported studies. Sometimes it so happens that after a trial is initiated, the overall event rate proves to be unexpectedly low and the sample size may have to be adjusted, with all statistical precautions. While calculating the sample size an investigator needs to anticipate the variation in the measures that are being studied. It is easy to understand why we would require a smaller sample if the population is more homogenous and therefore has a smaller variance or standard deviation. Suppose we are studying the effect of an intervention on the weight and consider a population with weights ranging from 45 to kg. Naturally the standard deviation in this group will be great and we would need a larger sample size to detect a difference between interventions, else the difference between the two groups would be masked by the inherent difference between them because of the variance. If on the other hand, we were to take a sample from a population with weights between 80 and kg we would naturally get a tighter and more homogenous group, thus reducing the standard deviation and therefore the sample size. The sample size is calculated using the following formula:

2: Power and Sample Size Determination

Statistical power is a fundamental consideration when designing research experiments. It goes hand-in-hand with sample size. The formulas that our calculators use come from clinical trials, epidemiology, pharmacology, earth sciences, psychology, survey sampling basically every scientific discipline.

Sample Size Determining Sample Size: This simple question is a never-ending quandary for researchers who use statistically based calculations to answer different questions. A larger sample group can yield more accurate study results – but excessive responses can be pricey. Learn How to Determine Sample Size

Consequential research requires an understanding of the statistics that drive the range of sample size decisions you need to make. A simple equation will help you put the migraine pills away and sample confidently knowing that there is a high probability that your survey is statistically accurate with the correct sample size.

Sample Size Variables Based on Target Population Before you can calculate a sample size, you need to determine a few things about the target population and the sample you need: **Population Size** – How many total people fit your demographic? For instance, if you want to know about mothers living in the US, your population size would be the total number of mothers living in the US. Not all populations sizes need to be this large. Even if your population size is small, just know who fits into your demographics. It is common for the population to be unknown or approximated between two educated guesses. **Margin of Error Confidence Interval** – No sample will be perfect, so you must decide how much error to allow. The confidence interval determines how much higher or lower than the population mean you are willing to let your sample mean fall. For example, it will look something like this: **Standard of Deviation** – How much variance do you expect in your responses? **Calculating Sample Size** Okay, now that we have these values defined, we can calculate our needed sample size. This can be done using an online sample size calculator or with paper and pencil. Your confidence level corresponds to a Z-score. This is a constant value needed for this equation. Here are the z-scores for the most common confidence levels:

3: PowerSampleSize < Main < Vanderbilt Biostatistics Wiki

The sample size computations depend on the level of significance, α , the desired power of the test (equivalent to $1 - \beta^2$), the variability of the outcome, and the effect size. The effect size is the difference in the parameter of interest that represents a clinically meaningful difference.

One can select a power and determine an appropriate sample size beforehand or do power analysis afterwards. However, power analysis is beyond the scope of this course and predetermining sample size is best. **Sample Size Importance** An appropriate sample size is crucial to any well-planned research investigation. Although crucial, the simple question of sample size has no definite answer due to the many factors involved. We expect large samples to give more reliable results and small samples to often leave the null hypothesis unchallenged. Large samples may be justified and appropriate when the difference sought is small and the population variance large. Established statistical procedures help ensure appropriate sample sizes so that we reject the null hypothesis not only because of statistical significance, but also because of practical importance. These procedures must consider the size of the type I and type II errors as well as the population variance and the size of the effect. The probability of committing a type I error is the same as our level of significance, commonly, 0. Ideally both types of error are minimized. Alpha is generally established before-hand: The larger alpha values result in a smaller probability of committing a type II error which thus increases the power. Power is the area under the distribution of sampling means centered on which is beyond the critical value for the distribution of sampling means centered on. The basic factors which affect power are the directional nature of the alternative hypothesis number of tails ; the level of significance alpha ; n sample size ; and the effect size ES. We will consider each in turn. Suppose we change the example above from a one-tailed to a two-tailed test. There are now two regions to consider, one above 1. Most of the area from the sampling distribution centered on comes from above. One-tailed tests generally have more power. Since a larger value for alpha corresponds with a small confidence level, we need to be clear we are referred strictly to the magnitude of alpha and not the increased confidence we might associate with a smaller value! Increasing sample size increases power. For comparison we will summarize our results:

4: Sample size calculation

Overview of Power and Sample www.enganchecubano.com Calculators. We have 30 calculators. This is a quick-start guide.

This will be either that the means differ two sided or they differ in a particular direction one sided. The default is two sided. Decide the significance level you plan to use. Decide what power you want i . In fundamental studies where we may only be interested in large effects a Type II error may not have such serious consequences. Obtain an estimate of the noise, σ . This has to come from a previous study, the literature or a pilot study. Estimate the signal effect size that might interest you. How large a difference between the two means would be of scientific or clinical interest? If the difference is only small, you are probably not particularly interested in it. If it is large, then you certainly want to be able to detect it. The signal is the cutoff between these two alternatives. If the response is larger, then there will be an even greater chance of detecting it. What if there are more than two groups? It is technically possible to do a power analysis for an analysis of variance with several treatment groups. The problem is to specify an effect size of clinical or scientific importance when there are three or more groups. One alternative is to power the experiment assuming a t-test on the two groups likely to be most extreme such as the control and top dose assuming there are such groups. This would mean that if the response is stronger than expected, then differences between the control and an intermediate group would become statistically significant. A screen shot of such a calculation for an experiment with five treatment groups with an effect size of 1. This would require 25 animals. Note that large numbers are needed in some cases. A web site that will do the calculations Click the arrow below for a pdf paper giving more details on power analysis. Although there is probably sufficient information given in the table above and the example below for you to estimate your required sample size, you can click below for a web site which will do the calculations for you. It can be downloaded from this web site An example comparing two means A vet wants to compare the effect on blood pressure of two anesthetics for dogs under clinical conditions. He has published some preliminary data. The dogs were unsexed healthy animals weighing 3 kg. Mean systolic blood pressure was mm Hg with a standard deviation of 36mm, the noise Assume: A difference in blood pressure of 20 mmHg the signal or more would be of clinical importance a clinical not a statistical decision. A significance level of 0.05. Note that great accuracy is not needed as there are uncertainties in the estimates of the standard deviation and the effect size of clinical importance. However there are many statistical software packages will do the calculations. The output below is done using the R statistical package for this set of data. Note that the sample size needs to be rounded up to a whole number. Sixty-eight dogs per group in total is a lot of dogs and using such animals would be time-consuming. An alternative In the same journal an investigator was working with male Beagles weighing kg. These had a mean BP of mm Hg. Assume a 20mm difference between groups would be of clinical importance as before. The table below summarises the situation. This poses a problem. And is there ever any case for using genetically heterogeneous animals if all it does is increase noise and reduce the power of the experiment, leading to false negative results? Alternative approaches It would make no sense to go ahead and do the experiment simply using the heterogeneous dogs. But there are some obvious alternatives. If each dog could be given both anaesthetics say in random order on different days, then it would be possible to use small numbers of even quite heterogeneous dogs, assuming that there are no important breed differences in response. Technically, this would be a randomised block design discussed later 2. As far as possible there should be equal numbers in each group. This would indicate whether the two anesthetics differ over-all and whether breed differences need to be taken into account when choosing one of these anesthetics. If lots of characters are being measured it may not be clear which one is the most important There may be no estimate of the standard deviation if the character has not previously been measured In fundamental research it may be impossible to specify an effect size likely to be of scientific importance A power analysis is difficult with complex experiments involving many treatment groups and possible interactions. This depends on the law of diminishing returns. It needs an estimate of E: There may be a case for E being higher if it leads to a more balanced design, the likely cost of a

Type II error is high, the procedures are very mild or it is an in-vitro experiment with no ethical implications E is the number of degrees of freedom in an analysis of variance ANOVA. It is based on the need to obtain an adequate estimate of the standard deviation. The plot above right shows the amount of information in a sample of data as a function of E . However, if experimental units are inexpensive such as tissue culture dishes then Suppose you decide to do an experiment with four treatment groups a control and three dose levels and eight animals per group. So this is unnecessarily large. But it is crude compared with the power analysis. Use a power analysis where possible. Use the resource equation when a power analysis is not possible.

5: - Power and Sample Size Determination for Testing a Population Mean | STAT

To calculate power or sample size for these tests, you need to determine the minimum difference (effect) that you consider to be meaningful. Then, you can determine the power or the sample size you need to be able to reflect the null hypothesis when the true value differs from the hypothesized value by this minimum difference.

Acknowledgements PS is an interactive program for performing power and sample size calculations that may be downloaded for free. It can be used for studies with dichotomous, continuous, or survival response measures. The alternative hypothesis of interest may be specified either in terms of differing response rates, means, or survival times, or in terms of relative risks or odds ratios. Studies with dichotomous or continuous outcomes may involve either a matched or independent study design. The program can determine the sample size needed to detect a specified alternative hypothesis with the required power, the power with which a specific alternative hypothesis can be detected with a given sample size, or the specific alternative hypotheses that can be detected with a given power and sample size. The PS program can produce graphs to explore the relationships between power, sample size and detectable alternative hypotheses. It is often helpful to hold one of these variables constant and plot the other two against each other. The program can generate graphs of sample size versus power for a specific alternative hypothesis, sample size versus detectable alternative hypotheses for a specified power, or power versus detectable alternative hypotheses for a specified sample size. Linear or logarithmic axes may be used for either axes. Multiple curves can be plotted on a single graphic. We have also installed the program on Linux and Macintosh computers using a program called Wine that facilitates running Windows software on other operating systems. To avoid problems with the installation process, it is helpful if the target folder is empty. A file called pssetup3. Click the Overview button for an introduction to the program and instruction on its use. PS is a self-documented program with extensive interactive help. We know of a bug in one of the third-party tools that we use that causes the program to malfunction when the Windows language is set to something other than English. We are working to replace this tool with one of our own and we apologize for any inconvenience that this problem causes. Macintosh and Linux In the past we have tried to supply installation packages that allowed the installation of the PS program in one step. For a number of reasons, this has been a troublesome approach. There are many versions of the OSX and Linux distributions in use. It is difficult to make an installer that works correctly on all of these. PlayOnMac and PlayOnLinux are free packages that simplify the installation and use of Windows software on these other platforms. The process for installing PS involves 3 steps: Download the PS installer pssetup3. PlayOnMac and PlayOnLinux inspect the software to be installed and attempt to also install other packages and tools that are necessary for the PS program to run correctly. There are a number Wine implementations and Wine front ends available. Our recommendation is not the only approach. Click here for PlayOnMac instructions. If you have any questions or comments about our software please send email to dale. It will be appreciated. The method of Schlesselman is used for studies with independent case and control groups that will be analyzed using an uncorrected chi-squared test; the method of Casagrande et al. The alternative hypotheses may be specified in terms of odds ratios or exposure prevalence rates. The method of Dupont is used for studies with paired or matched cases and controls. The alternative hypotheses are specified in terms of odds ratios. Multiple 2 X 2 tables -- Mantel-Haenszel Test: The method of Wittes and Wallenstein is used. Assume that each 2 X 2 table consists of cases and controls selected from a different stratum that is defined by one or more confounding variables. The odds ratio for disease in exposed subjects compared to unexposed subjects is assumed to be equal within all strata. The alternative hypotheses are specified in terms of this odds ratio. The methods of Schlesselman , Casagrande et al. The alternative hypotheses may be specified in terms of relative risks or outcome probabilities. Linear Regression 1 Treatment -- Testing the slope of a simple linear regression line: The method of Dupont and Plummer is used to design studies in which we wish to detect a regression slope of a given magnitude. The values of the independent x variable of the regression line may either be specified by the investigator or determined observationally when the study is performed. In the latter case, the investigator must estimate the standard deviation of the independent variable s. Linear Regression 2

Treatments -- Comparing the slopes and intercepts of two independent linear regressions: The approach of Dupont and Plummer is used to design studies in which we wish to determine whether the slopes or intercepts of two independent regression lines differ by a given amount. The values of the independent x variables of the regression lines may either be specified by the investigator or determined observationally when the study is performed. In the latter case, the investigator must estimate the standard deviations of the independent variables. Survival Studies -- Evaluating independent cohorts using the log-rank test: The approach of Schoenfeld and Richter is used. The ratio of the number of control subjects per experimental subject in the cohorts being compared may be specified by the user. The alternative hypotheses are specified in terms of the hazard ratio for control subjects relative to experimental subjects or the median survival times for control and experimental subjects. The approach of Dupont and Plummer is used for paired and independent samples. The ratio of the number of control subjects per experimental subject may be specified by the user. This method produces results that are in close agreement with those of Pearson and Hartley Wittes J, Wallenstein S: Visual Components Sybase Inc. Journal of the National Cancer Institute, 22, Statistical Modeling for Biomedical Researchers, 2nd Edition. We are grateful to Gordon R. Bernard for his support and to Yuwei Zhu for her assistance in editing this program. This web page and the PS:

6: Power (statistics) - Wikipedia

The paradigmatic problem of power calculation is the estimation of a parameter $\hat{\mu}$ (for example, a regression coefficient) such as would arise in estimating a difference or treatment effect, with the sample size determining the standard.

Background[edit] Statistical tests use data from samples to assess, or make inferences about, a statistical population. In the concrete setting of a two-sample comparison, the goal is to assess whether the mean values of some attribute obtained for individuals in two sub-populations differ. For example, to test the null hypothesis that the mean scores of men and women on a test do not differ, samples of men and women are drawn, the test is administered to them, and the mean score of one group is compared to that of the other group using a statistical test such as the two-sample z-test. The power of the test is the probability that the test will find a statistically significant difference between men and women, as a function of the size of the true difference between those two populations. Factors influencing power[edit] Statistical power may depend on a number of factors. Some factors may be particular to a specific testing situation, but at a minimum, power nearly always depends on the following three factors: The most commonly used criteria are probabilities of 0. If the criterion is 0. One easy way to increase the power of a test is to carry out a less conservative test by using a larger significance criterion, for example 0. This increases the chance of rejecting the null hypothesis i. But it also increases the risk of obtaining a statistically significant result i. The magnitude of the effect of interest in the population can be quantified in terms of an effect size δ , where there is greater power to detect larger effects. An effect size can be a direct value of the quantity of interest, or it can be a standardized measure that also accounts for the variability in the population. If constructed appropriately, a standardized effect size, along with the sample size, will completely determine the power. An unstandardized direct effect size will rarely be sufficient to determine the power, as it does not contain information about the variability in the measurements. The sample size determines the amount of sampling error inherent in a test result. Other things being equal, effects are harder to detect in smaller samples. Increasing sample size is often the easiest way to boost the statistical power of a test. How increased sample size translates to higher power is a measure of the efficiency of the test—for example, the sample size required for a given power. Consequently, power can often be improved by reducing the measurement error in the data. The design of an experiment or observational study often influences the power. For example, in a two-sample testing situation with a given total sample size n , it is optimal to have equal numbers of observations from the two populations being compared as long as the variances in the two populations are the same. In regression analysis and analysis of variance, there are extensive theories and practical strategies for improving the power based on optimally setting the values of the independent variables in the model. However, there will be times when this 4-to-1 weighting is inappropriate. In medicine, for example, tests are often designed in such a way that no false negatives Type II errors will be produced. But this inevitably raises the risk of obtaining a false positive a Type I error. In many contexts, the issue is less about determining if there is or is not a difference but rather with getting a more refined estimate of the population effect size. For example, if we were expecting a population correlation between intelligence and job performance of around 0. However, in doing this study we are probably more interested in knowing whether the correlation is 0. In this context we would need a much larger sample size in order to reduce the confidence interval of our estimate to a range that is acceptable for our purposes. Techniques similar to those employed in a traditional power analysis can be used to determine the sample size required for the width of a confidence interval to be less than a given value. Many statistical analyses involve the estimation of several unknown quantities. In simple cases, all but one of these quantities are nuisance parameters. In this setting, the only relevant power pertains to the single quantity that will undergo formal statistical inference. In some settings, particularly if the goals are more "exploratory", there may be a number of quantities of interest in the analysis. For example, in a multiple regression analysis we may include several covariates of potential interest. In situations such as this where several hypotheses are under consideration, it is common that the powers associated with the different hypotheses differ. For instance, in multiple regression analysis, the power for detecting an effect of a given size is related to the variance of the

covariate. Since different covariates will have different variances, their powers will differ as well. Any statistical analysis involving multiple hypotheses is subject to inflation of the type I error rate if appropriate measures are not taken. Such measures typically involve applying a higher threshold of stringency to reject a hypothesis in order to compensate for the multiple comparisons being made α . In this situation, the power analysis should reflect the multiple testing approach to be used. Thus, for example, a given study may be well powered to detect a certain effect size when only one test is to be made, but the same effect size may have much lower power if several tests are to be performed. It is also important to consider the statistical power of a hypothesis test when interpreting its results. A hypothesis test may fail to reject the null, for example, if a true difference exists between two populations being compared by a t-test but the effect is small and the sample size is too small to distinguish the effect from random chance. Post hoc analysis Power analysis can either be done before a priori or prospective power analysis or after post hoc or retrospective power analysis data are collected. A priori power analysis is conducted prior to the research study, and is typically used in estimating sufficient sample sizes to achieve adequate power. Post-hoc analysis of "observed power" is conducted after a study has been completed, and uses the obtained sample size and effect size to determine what the power was in the study, assuming the effect size in the sample is equal to the effect size in the population. Whereas the utility of prospective power analysis in experimental design is universally accepted, post hoc power analysis is fundamentally flawed. In particular, it has been shown that post-hoc "observed power" is a one-to-one function of the p-value attained. In frequentist statistics, an underpowered study is unlikely to allow one to choose between hypotheses at the desired significance level. In Bayesian statistics, hypothesis testing of the type used in classical power analysis is not done. In the Bayesian framework, one updates his or her prior beliefs using the data obtained in a given study. In principle, a study that would be deemed underpowered from the perspective of hypothesis testing could still be used in such an updating process. A study with low power is unlikely to lead to a large change in beliefs. Example[edit] The following is an example that shows how to compute power for a randomized experiment: Suppose the goal of an experiment is to study the effect of a treatment on some quantity, and compare research subjects by measuring the quantity before and after the treatment, analyzing the data using a paired t-test.

7: Determining Sample Size: Find the # of response you need | Qualtrics

Power and sample size estimations are used by researchers to determine how many subjects are needed to answer the research question (or null hypothesis). An example is the case of thrombolysis in acute myocardial infarction (AMI).

Boston University School of Public Health Introduction A critically important aspect of any study is determining the appropriate sample size to answer the research question. This module will focus on formulas that can be used to estimate the sample size needed to produce a confidence interval estimate with a specified margin of error precision or to ensure that a test of hypothesis has a high probability of detecting a meaningful difference in the parameter. Studies should be designed to include a sufficient number of participants to adequately address the research question. Studies that have either an inadequate number of participants or an excessively large number of participants are both wasteful in terms of participant and investigator time, resources to conduct the assessments, analytic efforts and so on. These situations can also be viewed as unethical as participants may have been put at risk as part of a study that was unable to answer an important question. Studies that are much larger than they need to be to answer the research questions are also wasteful. The formulas presented here generate estimates of the necessary sample size n required based on statistical criteria. However, in many studies, the sample size is determined by financial or logistical constraints. For example, suppose a study is proposed to evaluate a new screening test for Down Syndrome. Suppose that the screening test is based on analysis of a blood sample taken from women early in pregnancy. In order to evaluate the properties of the screening test t . The amniocentesis is included as the gold standard and the plan is to compare the results of the screening test to the results of the amniocentesis. These financial constraints alone might substantially limit the number of women that can be enrolled. Just as it is important to consider both statistical and clinical significance when interpreting results of a statistical analysis, it is also important to weigh both statistical and logistical issues in determining the sample size for a study. Learning Objectives After completing this module, the student will be able to: Provide examples demonstrating how the margin of error, effect size and variability of the outcome affect sample size computations. Compute the sample size required to estimate population parameters with precision. Interpret statistical power in tests of hypothesis. Compute the sample size required to ensure high power when hypothesis testing. Issues in Estimating Sample Size for Confidence Intervals Estimates The module on confidence intervals provided methods for estimating confidence intervals for various parameters t . Confidence intervals for every parameter take the following general form: It involves a value from the t distribution, as opposed to one from the standard normal distribution, to reflect the desired level of confidence. When performing sample size computations, we use the large sample formula shown here. The resultant sample size might be small, and in the analysis stage, the appropriate confidence interval formula must be used. For example, suppose we want to estimate the mean weight of female college students. The margin of error is so wide that the confidence interval is uninformative. In order to determine the sample size needed, the investigator must specify the desired margin of error. It is important to note that this is not a statistical issue, but a clinical or a practical one. For example, suppose we want to estimate the mean birth weight of infants born to mothers who smoke cigarettes during pregnancy. Birth weights in infants clearly have a much more restricted range than weights of female college students. Therefore, we would probably want to generate a confidence interval for the mean birth weight that has a margin of error not exceeding 1 or 2 pounds. Our goal is to determine the sample size, n , that ensures that the margin of error, "E," does not exceed a specified value. We can take the formula above and, with some algebra, solve for n : First, multiply both sides of the equation by the square root of n . Then cancel out the square root of n from the numerator and denominator on the right side of the equation since any number divided by itself is equal to 1. Now divide both sides by "E" and cancel out "E" from the numerator and denominator on the left side. Finally, square both sides of the equation to get: This formula generates the sample size, n , required to ensure that the margin of error, E, does not exceed a specified value. When we use the sample size formula above or one of the other formulas that we will present in the sections that follow, we are planning a study to estimate the unknown mean of a particular outcome variable in a population. It is

unlikely that we would know the standard deviation of that variable. In sample size computations, investigators often use a value for the standard deviation from a previous study or a study done in a different, but comparable, population. The sample size computation is not an application of statistical inference and therefore it is reasonable to use an appropriate estimate for the standard deviation. The estimate can be derived from a different study that was reported in the literature; some investigators perform a small pilot study to estimate the standard deviation. A pilot study usually involves a small number of participants. Regardless of how the estimate of the variability of the outcome is derived, it should always be conservative. The formula produces the minimum sample size to ensure that the margin of error in a confidence interval will not exceed E . In planning studies, investigators should also consider attrition or loss to follow-up. The formula above gives the number of participants needed with complete data to ensure that the margin of error in the confidence interval does not exceed E . We will illustrate how attrition is addressed in planning studies through examples in the following sections.

Sample Size for One Sample, Continuous Outcome

In studies where the plan is to estimate the mean of a continuous outcome variable in a single population, the formula for determining sample size is given below: An investigator wants to estimate the mean systolic blood pressure in children with congenital heart disease who are between the ages of 3 and 5. How many children should be enrolled in the study? The standard deviation of systolic blood pressure is unknown, but the investigators conduct a literature search and find that the standard deviation of systolic blood pressures in children with other cardiac defects is between 15 and 20. To estimate the sample size, we consider the larger standard deviation in order to obtain the most conservative largest sample size. We always round up; the sample size formulas always generate the minimum number of subjects needed to ensure the specified precision. Because the estimates of the standard deviation were derived from studies of children with other cardiac defects, it would be advisable to use the larger standard deviation and plan for a study with 62 children. Selecting the smaller sample size could potentially produce a confidence interval estimate with a larger margin of error.

An investigator wants to estimate the mean birth weight of infants born full term approximately 40 weeks gestation to mothers who are 19 years of age and under. The mean birth weight of infants born full-term to mothers 20 years of age and older is 3,400 grams with a standard deviation of 400 grams. Try to work through the calculation before you look at the answer. The equation to determine the sample size for determining p seems to require knowledge of p , but this is obviously a circular argument, because if we knew the proportion of successes in the population, then a study would not be necessary! What we really need is an approximate value of p or an anticipated value. The range of p is 0 to 1, and therefore the range of $p(1-p)$ is 0 to 1.

An investigator wants to estimate the proportion of freshmen at his University who currently smoke cigarettes. Because we have no information on the proportion of freshmen who smoke, we use 0.25. If the investigator believes that this is a reasonable estimate of prevalence 2 years later, it can be used to plan the next study.

An investigator wants to estimate the prevalence of breast cancer among women who are between 40 and 45 years of age living in Boston. How many women must be involved in the study to ensure that the estimate is precise? National data suggest that 1 in 100 women are diagnosed with breast cancer by age 40. This translates to a proportion of 0.01. The sample size is computed as follows: This is a situation where investigators might decide that a sample of this size is not feasible. Suppose that the investigators thought a sample of size 5, would be reasonable from a practical point of view. Recall that the confidence interval formula to estimate prevalence is: Assuming that the prevalence of breast cancer in the sample will be close to that based on national data, we would expect the margin of error to be approximately equal to the following: The investigators must decide if this would be sufficiently precise to answer the research question. Note that the above is based on the assumption that the prevalence of breast cancer in Boston is similar to that reported nationally. This may or may not be a reasonable assumption. In fact, it is the objective of the current study to estimate the prevalence in Boston.

Sample Sizes for Two Independent Samples, Continuous Outcome

In studies where the plan is to estimate the difference in means between two independent populations, the formula for determining the sample sizes required in each comparison group is given below: Recall from the module on confidence intervals that, when we generated a confidence interval estimate for the difference in means, we used S_p , the pooled estimate of the common standard deviation, as a measure of variability in the outcome based on

pooling the data, where S_p is computed as follows: If data are available on variability of the outcome in each comparison group, then S_p can be computed and used in the sample size formula. However, it is more often the case that data on the variability of the outcome are available from only one group, often the untreated one. When planning a clinical trial to investigate a new drug or procedure, data are often available from other trials that involved a placebo or an active control group. The standard deviation of the outcome variable measured in patients assigned to the placebo, control or unexposed group can be used to plan a future trial, as illustrated below. Note that the formula for the sample size generates sample size estimates for samples of equal size. If a study is planned where different numbers of patients will be assigned or different numbers of patients will comprise the comparison groups, then alternative formulas can be used. An investigator wants to plan a clinical trial to evaluate the efficacy of a new drug designed to increase HDL cholesterol the "good" cholesterol. The plan is to enroll participants and to randomly assign them to receive either the new drug or a placebo. HDL cholesterol will be measured in each participant after 12 weeks on the assigned treatment. The investigator would like the margin of error to be no more than 3 units. How many patients should be recruited into the study? The sample sizes are computed as follows: To plan this study, we can use data from the Framingham Heart Study. In participants who attended the seventh examination of the Offspring Study and were not on treatment for high cholesterol, the standard deviation of HDL cholesterol is 12. We will use this value and the other inputs to compute the sample sizes as follows: Again, these sample sizes refer to the numbers of participants with complete data. In order to ensure that the total sample size of 1000 is available at 12 weeks, the investigator needs to recruit more participants to allow for attrition. An investigator wants to compare two diet programs in children who are obese.

8: Minitab's Power and Sample Size Tools - Minitab

This page contains links to JavaScript based forms for simple power/sample size calculations. An extensive list of alternative and more comprehensive resources is available at UCSF Biostatistics: Power and Sample Size Programs.

9: Power and sample size

The program can determine the sample size needed to detect a specified alternative hypothesis with the required power, the power with which a specific alternative hypothesis can be detected with a given sample size, or the specific alternative hypotheses that can be detected with a given power and sample size.

Know when to quit Canadian press style guide The rules of life richard templar Voyages of Discovery in the Arctic and Antarctic Seas, and Round the World The liberal tradition in American thought Goethes knowledge of English literature. The Ark in the Garden 12th chemistry notes file Introduction: Revisiting the intellectual transformation of nineteenth-century France The new atheism and the assault on theism The T. R. Connection The dawning of a new century: 2000-2007. All aboard the goodnight train Preparation for the NCLEX-RN examination : transitional issues for the foreign-educated nurse The girl on the train kickass Snow Bound, Among the Hills, Songs of Labor and Other Poems Cryptography and network security principles and practices 7th edition Traffic, Environment, Politics (Europaische Hochschulschriften. Reihe Xxxi, Politikwissenschaft, Bd. 386. Selected games (of William Steinitz Marine control technology Motorola mh230r manual espaÃ±ol 1994 cadillac seville service manual Cancer, the outlaw cell Gleim Cpa Examination Review Problems and Solutions 10ed (CPA Examination Review) Personal computing (1977) Barbara Delinskys Mixed Prepack Touch typing Made Simple Coal Mining and Ground Water Resources U.S. Changing beads desk diva Sail question paper for mechanical 2013 Integrative health a holistic approach for health professionals The archaeology of medieval bookbinding The third act : George W. Bush and the Gulf Still Rhymin on the Range Complete Preparation for the PCAT Advances In Urology Volume 7 (ADVANCES IN UROLOGY) Discretionary managerial behavior Next to normal score Screen printing with a drawing fluid resist The Complete Poetry of Edgar Allan Poe