Glossary

1 : M matched case–control study. See *matched study*.

2 × 2 contingency table. A 2 × 2 contingency table is used to describe the association between a binary independent variable and a binary response variable of interest.

2 × 2 × K contingency table. See *stratified 2 × 2 tables*.

acceptance region. In hypothesis testing, an acceptance region is a set of sample values for which the null hypothesis cannot be rejected or can be accepted. It is the complement of the rejection region.

accrual period or **recruitment period** or **accrual**. The accrual period (or recruitment period) is the period during which subjects are being enrolled (recruited) into a study. Also see *follow-up period*.

actual alpha, **actual significance level**. This is an attained or observed significance level.

actual confidence-interval width. This is the CI width that is computed using the rounded-up sample size when the population standard deviation is known.

actual power. This is power corresponding to the actual sample size.

actual probability of confidence-interval width. `ciwidth` will calculate the required sample size for a specified probability of CI width, and if it is fractional, will round it up to report an integer. The actual probability of CI width is calculated using the rounded sample-size estimates.

actual sample size. For a two-sample study, when specifying one of the sample sizes and a sample-size ratio that result in noninteger sample sizes, `power` and `ciwidth` will round down the noninteger sample sizes to the nearest integers and use these integers for computations. The actual sample size is the rounded-down sample size.

actual sample-size ratio. When specifying a sample-size ratio that results in noninteger sample sizes, `power` and `ciwidth` will round down the input sample sizes and round up the computed sample sizes to the nearest integers. The actual sample-size ratio is computed using the rounded sample sizes.

administrative censoring. Administrative censoring is the right-censoring that occurs when the study observation period ends. All subjects complete the course of the study and are known to have experienced one of two outcomes at the end of the study: survival or failure. This type of censoring should not be confused with withdrawal and loss to follow-up. Also see *censored, uncensored, left-censored, and right-censored*.

allocation ratio. This ratio n_2/n_1 represents the number of subjects in the comparison, experimental group relative to the number of subjects in the reference, control group. Also see [PSS-4] **Unbalanced designs**.

alpha. Alpha, α, denotes the significance level.

alternative hypothesis. In hypothesis testing, the alternative hypothesis represents the counterpoint to which the null hypothesis is compared. When the parameter being tested is a scalar, the alternative hypothesis can be either one sided or two sided.

alternative value, **alternative parameter**. This value of the parameter of interest under the alternative hypothesis is fixed by the investigator in a power and sample-size analysis. For example, alternative mean value and alternative mean refer to a value of the mean parameter under the alternative hypothesis.

analysis of variance, **ANOVA**. This is a class of statistical models that studies differences between means from multiple populations by partitioning the variance of the continuous outcome into independent sources of variation due to effects of interest and random variation. The test statistic is then formed as a ratio of the expected variation due to the effects of interest to the expected random variation. Also see *one-way ANOVA*, *two-way ANOVA*, *one-way repeated-measures ANOVA*, and *two-way repeated-measures ANOVA*.

balanced design. A balanced design represents an experiment in which the numbers of treated and untreated subjects are equal. For many types of two-sample hypothesis tests, the power of the test is maximized with balanced designs. For both PrSS and PSS analyses, balanced designs tend to require fewer subjects than their corresponding unbalanced designs.

beta. Beta, β, denotes the probability of committing a type II error, namely, failing to reject the null hypothesis even though it is false.

between-subjects design. This is an experiment that has only between-subjects factors. See [PSS-2] **power oneway** and [PSS-2] **power twoway**.

between-subjects factor. This is a factor for which each subject receives only one of the levels.

binomial test. A binomial test is a test for which the exact sampling distribution of the test statistic is binomial; see [R] **bitest**. Also see [PSS-2] **power oneproportion**.

bisection method. This method finds a root x of a function $f(x)$ such that $f(x) = 0$ by repeatedly subdividing an interval on which $f(x)$ is defined until the change in successive root estimates is within the requested tolerance and function $f(\cdot)$ evaluated at the current estimate is sufficiently close to zero.

case–control study. An observational study that retrospectively compares characteristics of subjects with a certain problem (cases) with characteristics of subjects without the problem (controls). For example, to study association between smoking and lung cancer, investigators will sample subjects with and without lung cancer and record their smoking status. Case–control studies are often used to study rare diseases.

CCT. See *controlled clinical trial study*.

cell means. These are means of the outcome of interest within cells formed by the cross-classification of the two factors. See [PSS-2] **power twoway** and [PSS-2] **power repeated**.

cell-means model. A cell-means model is an ANOVA model formulated in terms of cell means.

χ^2 test. This test for which either an asymptotic sampling distribution or a sampling distribution of a test statistic is χ^2. See [PSS-2] **power onevariance** and [PSS-2] **power twoproportions**.

CI. See *confidence interval*.

CI precision. See *confidence-interval precision*.

CI precision graph. See *confidence-interval precision curve*.

CI width. See *confidence-interval width*.

clinical trial. A clinical trials is an experiment testing a medical treatment or procedure on human subjects.

clinically meaningful difference, **clinically meaningful effect**, **clinically significant difference**. Clinically meaningful difference represents the magnitude of an effect of interest that is of clinical importance. What is meant by "clinically meaningful" may vary from study to study. In clinical trials,

for example, if no prior knowledge is available about the performance of the considered clinical procedure, a standardized effect size (adjusted for standard deviation) between 0.25 and 0.5 may be considered of clinical importance.

cluster randomized design, **CRD**, **cluster randomized trial**, **CRT**, **group randomized trial**, **GRT**. Cluster randomized design is a type of randomized design in which groups of subjects or clusters are sampled instead of individual subjects. A cluster is the randomization unit, and an individual within a cluster is the analysis unit. Observations within a cluster tend to be correlated. The strength of the correlation is measured by the intraclass correlation. Also see *individual-level design*.

cluster size. The number of subjects in a group or cluster in a cluster randomized design. If cluster sizes vary between clusters, the coefficient of variation for cluster sizes is used for power and sample-size determination.

Cochran–Armitage test. The Cochran–Armitage test is a test for a linear trend in a probability of response in a $J \times 2$ contingency table. The test statistic has an asymptotic χ^2 distribution under the null hypothesis. See [PSS-2] **power trend**.

Cochran–Mantel–Haenszel test. See *Mantel–Haenszel test*.

coefficient of variation, **CV**. Coefficient of variation measures the spread or the variability of the observations relative to the mean.

cohort study. Typically an observational study, a cohort study may also be an experimental study in which a cohort, a group of subjects who have similar characteristics, is followed over time and evaluated at the end of the study. For example, cohorts of vaccinated and unvaccinated subjects are followed over time to study the effectiveness of influenza vaccines.

columns in graph. Think of `power, graph()` and `ciwidth, graph()` as graphing the columns of `power, table` and `ciwidth, table`, respectively. One of the columns will be placed on the x axis, another will be placed on the y axis, and, if you have more columns with varying values, separate plots will be created for each. Similarly, we use the terms "column symbol", "column name", and "column label" to refer to symbols, names, and labels that appear in tables when tabular output is requested.

common odds ratio. A measure of association in stratified 2×2 tables. It can be viewed as a weighted aggregate of stratum-specific odds ratios.

comparison value. See *alternative value*.

compound symmetry. A covariance matrix has a compound-symmetry structure if all the variances are equal and all the covariances are equal. This is a special case of the sphericity assumption.

concordant pairs. In a 2×2 contingency table, a concordant pair is a pair of observations that are both either successes or failures. Also see *discordant pairs* and *Introduction* under *Remarks and examples* in [PSS-2] **power pairedproportions**.

confidence bounds. See *confidence limits*.

confidence coefficient. See *confidence level*.

confidence interval. A confidence interval provides an interval estimate for a parameter of interest. It is constructed such that, in a repeated independent sampling, the proportion of confidence intervals containing the true parameter value will be larger than or equal to the specified confidence level, $1 - \alpha$. A confidence interval can also be viewed as a range of plausible values that cannot be rejected by the corresponding hypothesis test at a given significance level α. See *Confidence intervals* in [PSS-3] **Intro (ciwidth)**. Also see *one-sided confidence interval* and *two-sided confidence interval*.

confidence level. The confidence level sets the degree of certainty with which the CIs, constructed from repeated independent sampling, will be guaranteed to contain the true parameter value. For example, when specifying a confidence level of 95, the CI is guaranteed to contain the true parameter value 95% of the time. The confidence level is equal to $1 - \alpha$, where α is the significance level.

confidence limits. The confidence limits are the upper and lower limits of the confidence interval. For two-sided CIs, both confidence limits are finite. For one-sided CIs, one confidence limit is finite and the other is infinite. An upper one-sided CI has a lower confidence limit equal to negative infinity, whereas a lower one-sided CI has an upper confidence limit equal to infinity. See *Confidence intervals* in [PSS-3] **Intro (ciwidth)**.

confidence-interval half-width. The half-width of a confidence interval is equal to one half of the confidence-interval width, $w/2$, and is also known as the margin of error. The CI half-width is used as a measure of precision for a symmetric CI.

confidence-interval precision. The precision of a confidence interval is typically measured by its width. A larger width means a lower degree of precision and leads to a wider CI. A smaller width means a higher degree of precision and leads to a narrower CI.

confidence-interval precision curve. A confidence-interval precision curve is a graph of the estimated CI width as a function of some other study parameter, such as sample size or probability of CI width. The CI width is plotted on the y axis, and the sample size or other parameter is plotted on the x axis.

confidence-interval precision determination. This pertains to the computation of confidence-interval width given sample size, probability of CI width, and other study parameters.

confidence-interval width. For two-sided CIs, the width is defined as the difference between the upper and lower limits. For an upper one-sided CI, the width is the difference between the upper confidence limit and the point estimate. For a lower one-sided CI, the width is the difference between the lower confidence limit and the point estimate.

contrasts. Contrasts refers to a linear combination of cell means such that the sum of contrast coefficients is zero.

control covariates. See *reduced model*.

control group. A control group comprises subjects that are randomly assigned to a group where they receive no treatment or receives a standard treatment. In hypothesis testing, this is usually a reference group. Also see *experimental group*.

controlled clinical trial study. This is an experimental study in which treatments are assigned to two or more groups of subjects without the randomization.

CRD. See *cluster randomized design*.

critical region. See *rejection region*.

critical value. In hypothesis testing, a critical value is a boundary of the rejection region.

cross-sectional study. This type of observational study measures various population characteristics at one point in time or over a short period of time. For example, a study of the prevalence of breast cancer in the population is a cross-sectional study.

CRT. See *cluster randomized design*.

CV. See *coefficient of variation*.

delta. Delta, δ, in the context of power and sample-size calculations, denotes the effect size.

directional test. See *one-sided test*.

discordant pairs. In a 2×2 contingency table, discordant pairs are the success–failure or failure–success pairs of observations. Also see *concordant pairs* and *Introduction* under *Remarks and examples* in [PSS-2] **power pairedproportions**.

discordant proportion. This is a proportion of discordant pairs or discordant sets. Also see *Introduction* under *Remarks and examples* in [PSS-2] **power pairedproportions** as well as *Introduction* under *Remarks and examples* in [PSS-2] **power mcc**.

discordant sets. In a matched study with multiple controls matched to a case, discordant sets are the sets in which there is any success–failure or failure–success match between the case and any matched control. Also see *Introduction* under *Remarks and examples* in [PSS-2] **power mcc**.

dropout. Dropout is the withdrawal of subjects before the end of a study and leads to incomplete or missing data.

effect size. The effect size is the size of the clinically significant difference between the treatments being compared, typically expressed as a quantity that is independent of the unit of measure. For example, in a one-sample mean test, the effect size is a standardized difference between the mean and its reference value. In other cases, the effect size may be measured as an odds ratio or a risk ratio. See [PSS-2] **Intro (power)** to learn more about the relationship between effect size and the power of a test.

effect-size curve. The effect-size curve is a graph of the estimated effect size or target parameter as a function of some other study parameter such as the sample size. The effect size or target parameter is plotted on the y axis, and the sample size or other parameter is plotted on the x axis.

effect-size determination. This pertains to the computation of an effect size or a target parameter given power, sample size, and other study parameters.

equal-allocation design. See *balanced design*.

exact test. An exact test is one for which the probability of observing the data under the null hypothesis is calculated directly, often by enumeration. Exact tests do not rely on any asymptotic approximations and are therefore widely used with small datasets. See [PSS-2] **power oneproportion** and [PSS-2] **power twoproportions**.

experimental group. An experimental group is a group of subjects that receives a treatment or procedure of interest defined in a controlled experiment. In hypothesis testing, this is usually a comparison group. Also see *control group*.

experimental study. In an experimental study, as opposed to an observational study, the assignment of subjects to treatments is controlled by investigators. For example, a study that compares a new treatment with a standard treatment by assigning each treatment to a group of subjects is an experimental study.

exponential test. The exponential test is the parametric test comparing the hazard rates, λ_1 and λ_2, (or log hazards) from two independent exponential (constant only) regression models with the null hypothesis H_0: $\lambda_2 - \lambda_1 = 0$ [or H_0: $\ln(\lambda_2) - \ln(\lambda_1) = \ln(\lambda_2/\lambda_1) = 0$].

exposure odds ratio. An odds ratio of exposure in cases relative to controls in a case–control study.

F test. An F test is a test for which a sampling distribution of a test statistic is an F distribution. See [PSS-2] **power twovariances**.

factor, **factor variables**. This is a categorical explanatory variable with any number of levels.

finite population correction. When sampling is performed without replacement from a finite population, a finite population correction is applied to the standard error of the estimator to reduce sampling variance.

Fisher–Irwin's exact test. See *Fisher's exact test*.

Fisher's exact test. Fisher's exact test is an exact small-sample test of independence between rows and columns in a 2×2 contingency table. Conditional on the marginal totals, the test statistic has a hypergeometric distribution under the null hypothesis. See [PSS-2] **power twoproportions** and [R] **tabulate twoway**.

Fisher's z test. This is a z test comparing one or two correlations. See [PSS-2] **power onecorrelation** and [PSS-2] **power twocorrelations**. Also see *Fisher's z transformation*.

Fisher's z transformation. Fisher's z transformation applies an inverse hyperbolic tangent transformation to the sample correlation coefficient. This transformation is useful for testing hypothesis concerning Pearson's correlation coefficient. The exact sampling distribution of the correlation coefficient is complicated, while the transformed statistic is approximately standard normal.

fixed effects. Fixed effects represent all levels of the factor that are of interest.

follow-up period or **follow-up**. The (minimum) follow-up period is the period after the last subject entered the study until the end of the study. The follow-up defines the phase of a study during which subjects are under observation and no new subjects enter the study. If T is the total duration of a study, and r is the accrual period of the study, then follow-up period f is equal to $T - r$. Also see *accrual period*.

follow-up study. See *cohort study*.

fractional sample size. Fractional (noninteger) sample sizes occur when specifying an odd number for the total sample size in studies with an equal-allocation design. They may also occur when specifying noninteger sample-size ratios.

full model. In the regression context, a full model is a regression model that includes all covariates of interest. Also see *reduced model*.

Greenhouse–Geisser correction. See *nonsphericity correction*.

group randomized trial, **GRT**. See *cluster randomized design*.

hypothesis. A hypothesis is a statement about a population parameter of interest.

hypothesis testing, **hypothesis test**. This method of inference evaluates the validity of a hypothesis based on a sample from the population. See *Hypothesis testing* under *Remarks and examples* in [PSS-2] **Intro (power)**.

hypothesized value. See *null value*.

individual-level design. Individual-level design is a classical randomized design in which individual subjects or observations are sampled; thus they represent both units of randomization and units of analysis. In contrast, see *cluster randomized design*.

interaction effects. Interaction effects measure the dependence of the effects of one factor on the levels of the other factor. Mathematically, they can be defined as the differences among treatment means that are left after main effects are removed from these differences.

intraclass correlation. Intraclass correlation measures the dependence of observations in the same group or cluster.

J $\times$ 2 contingency table. A $J \times 2$ contingency table is used to describe the association between an ordinal independent variable with J levels and a binary response variable of interest.

Lagrange multiplier test. See *score test*.

likelihood-ratio test. The likelihood-ratio (LR) test is one of the three classical testing procedures used to compare the fit of two models, one of which, the constrained model, is nested within the full (unconstrained) model. Under the null hypothesis, the constrained model fits the data as well as the full model. The LR test requires one to determine the maximal value of the log-likelihood function for both the constrained and the full models. See [PSS-2] **power twoproportions** and [R] **lrtest**.

loss to follow-up. Subjects are lost to follow-up if they do not complete the course of the study for reasons unrelated to the event of interest. For example, loss to follow-up occurs if subjects move to a different area or decide to no longer participate in a study. Loss to follow-up should not be confused with administrative censoring. If subjects are lost to follow-up, the information about the outcome these subjects would have experienced at the end of the study, had they completed the study, is unavailable. Also see *withdrawal*, *administrative censoring*, and *follow-up period or follow-up*.

lower one-sided confidence interval. A lower one-sided confidence interval contains a range of values that are greater than or equal to the lower confidence limit ll. The confidence interval is defined by a finite lower confidence limit and an upper confidence limit of infinity: $[ll, \infty)$.

lower one-sided test, **lower one-tailed test**. A lower one-sided test is a one-sided test of a scalar parameter in which the alternative hypothesis is lower one sided, meaning that the alternative hypothesis states that the parameter is less than the value conjectured under the null hypothesis. Also see *One-sided test versus two-sided test* under *Remarks and examples* in [PSS-2] **Intro (power)**.

main effects. These are average, additive effects that are associated with each level of each factor. For example, the main effect of level j of a factor is the difference between the mean of all observations on the outcome of interest at level j and the grand mean.

Mantel–Haenszel test. The Mantel–Haenszel test evaluates whether the overall degree of association in stratified 2×2 tables is significant assuming that the exposure effect is the same across strata. See [PSS-2] **power cmh**.

margin of error. See *confidence-interval half-width*.

marginal homogeneity. Marginal homogeneity refers to the equality of one or more row marginal proportions with the corresponding column proportions. Also see *Introduction* under *Remarks and examples* in [PSS-2] **power pairedproportions**.

marginal proportion. This represents a ratio of the number of observations in a row or column of a contingency table relative to the total number of observations. Also see *Introduction* under *Remarks and examples* in [PSS-2] **power pairedproportions**.

matched study. In a matched study, an observation from one group is matched to one or more observations from another group with respect to one or more characteristics of interest. When multiple matches occur, the study design is $1 : M$, where M is the number of matches. Also see *paired data*, also known as $1 : 1$ matched data.

McNemar's test. McNemar's test is a test used to compare two dependent binary populations. The null hypothesis is formulated in the context of a 2×2 contingency table as a hypothesis of marginal homogeneity. See [PSS-2] **power pairedproportions** and the `mcc` command in [R] **Epitab**.

MDES. See *minimum detectable effect size*.

mean contrasts. See *contrasts*.

minimum detectable effect size. The minimum detectable effect size is the smallest effect size that can be detected by hypothesis testing for a given power and sample size.

minimum detectable value. The minimum detectable value represents the smallest amount or concentration of a substance that can be reliably measured.

mixed design. A mixed design is an experiment that has at least one between-subjects factor and one within-subject factor. See [PSS-2] **power repeated**.

multiple partial correlation. In the regression context, multiple partial correlation is the measure of association between the dependent variable and one or more independent variables of interest, while controlling for the effect of other variables in the model.

negative effect size. In power and sample-size analysis, we obtain a negative effect size when the postulated value of the parameter under the alternative hypothesis is less than the hypothesized value of the parameter under the null hypothesis. Also see *positive effect size*.

nominal alpha, **nominal significance level**. This is a desired or requested significance level.

noncentrality parameter. In power and sample-size analysis, a noncentrality parameter is the expected value of the test statistic under the alternative hypothesis.

nondirectional test. See *two-sided test*.

nonsphericity correction. This is a correction used for the degrees of freedom of a regular F test in a repeated-measures ANOVA to compensate for the lack of sphericity of the repeated-measures covariance matrix.

null hypothesis. In hypothesis testing, the null hypothesis typically represents the conjecture that one is attempting to disprove. Often the null hypothesis is that a treatment has no effect or that a statistic is equal across populations.

null value, **null parameter**. This value of the parameter of interest under the null hypothesis is fixed by the investigator in a power and sample-size analysis. For example, null mean value and null mean refer to the value of the mean parameter under the null hypothesis.

number of clusters. The number of independent sampling units, groups or clusters, in a cluster randomized design.

observational study. In an observational study, as opposed to an experimental study, the assignment of subjects to treatments happens naturally and is thus beyond the control of investigators. Investigators can only observe subjects and measure their characteristics. For example, a study that evaluates the effect of exposure of children to household pesticides is an observational study.

observed level of significance. See *p-value*.

odds and odds ratio. The odds in favor of an event are $\text{Odds} = p/(1-p)$, where p is the probability of the event. Thus if $p = 0.2$, the odds are 0.25, and if $p = 0.8$, the odds are 4.

The log of the odds is $\ln(\text{Odds}) = \text{logit}(p) = \ln\{p/(1-p)\}$, and logistic regression models, for instance, fit $\ln(\text{Odds})$ as a linear function of the covariates.

The odds ratio is a ratio of two odds: $\text{Odds}_2/\text{Odds}_1$. The individual odds that appear in the ratio are usually for an experimental group and a control group or for two different demographic groups.

one-sample test. A one-sample test compares a parameter of interest from one sample with a reference value. For example, a one-sample mean test compares a mean of the sample with a reference value.

one-sided confidence interval. See *upper one-sided confidence interval* and *lower one-sided confidence interval*.

one-sided test, one-tailed test. A one-sided test is a hypothesis test of a scalar parameter in which the alternative hypothesis is one sided, meaning that the alternative hypothesis states that the parameter is either less than or greater than the value conjectured under the null hypothesis but not both. Also see *One-sided test versus two-sided test* under *Remarks and examples* in [PSS-2] **Intro (power)**.

one-way ANOVA, one-way analysis of variance. A one-way ANOVA model has a single factor. Also see [PSS-2] **power oneway**.

one-way repeated-measures ANOVA. A one-way repeated-measures ANOVA model has a single within-subject factor. Also see [PSS-2] **power repeated**.

paired data. Paired data consist of pairs of observations that share some characteristics of interest. For example, measurements on twins, pretest and posttest measurements, before and after measurements, repeated measurements on the same individual. Paired data are correlated and thus must be analyzed by using a paired test. See [PSS-3] **ciwidth pairedmeans** for PrSS analysis for a paired-means-difference CI.

paired observations. See *paired data*.

paired test. A paired test is used to test whether the parameters of interest of two paired populations are equal. The test takes into account the dependence between measurements. For this reason, paired tests are usually more powerful than their two-sample counterparts. For example, a paired-means or paired-difference test is used to test whether the means of two paired (correlated) populations are equal.

partial correlation. Partial correlation is the measure of association between two continuous variables, while controlling for the effect of other variables.

Pearson's correlation. Pearson's correlation ρ, also known as the product-moment correlation, measures the degree of association between two variables. Pearson's correlation equals the variables' covariance divided by their respective standard deviations, and ranges between -1 and 1. Zero indicates no correlation between the two variables.

population parameter. See *target parameter*.

positive effect size. In power and sample-size analysis, we obtain a positive effect size when the postulated value of the parameter under the alternative hypothesis is greater than the hypothesized value of the parameter under the null hypothesis. Also see *negative effect size*.

postulated value. See *alternative value*.

power. The power of a test is the probability of correctly rejecting the null hypothesis when it is false. It is often denoted as $1-\beta$ in the statistical literature, where β is the type II error probability. Commonly used values for power are 80% and 90%. See [PSS-2] **Intro (power)** for more details about power.

power and sample-size analysis. Power and sample-size analysis investigates the optimal allocation of study resources to increase the likelihood of the successful achievement of a study objective. The focus of power and sample-size analysis is on studies that use hypothesis testing for inference. Power and sample-size analysis provides an estimate of the sample size required to achieve the desired power of a test in a future study. See [PSS-2] **Intro (power)**. Also see *precision and sample-size analysis*.

power curve. A power curve is a graph of the estimated power as a function of some other study parameter such as the sample size. The power is plotted on the y axis, and the sample size or other parameter is plotted on the x axis. See [PSS-2] **power, graph**.

power determination. This pertains to the computation of a power given sample size, effect size, and other study parameters.

power function. The power functions is a function of the population parameter θ, defined as the probability that the observed sample belongs to the rejection region of a test for given θ. See *Hypothesis testing* under *Remarks and examples* in [PSS-2] **Intro (power)**.

power graph. See *power curve*.

precision and sample-size analysis. Just like power and sample-size analysis, precision and sample-size analysis investigates the optimal allocation of study resources to increase the likelihood of the successful achievement of a study objective. The focus of precision and sample-size analysis is on studies that use confidence intervals for inference. Precision and sample-size analysis provides an estimate of the sample size required to achieve the desired precision of a confidence interval in a future study. See [PSS-3] **Intro (ciwidth)**.

precision of a confidence interval. See *confidence-interval precision*.

probability of a type I error. This is the probability of committing a type I error of incorrectly rejecting the null hypothesis. Also see *significance level*.

probability of a type II error. This is the probability of committing a type II error of incorrectly accepting the null hypothesis. Common values for the probability of a type II error are 0.1 and 0.2 or, equivalently, 10% and 20%. Also see *beta* and *power*.

probability of confidence-interval width. The probability of CI width is the probability that the width of a CI in a future study will be no greater than a prespecified value.

probability of confidence-interval width determination. This pertains to the computation of the probability of CI width given CI width, sample size, and other study parameters.

prospective study. In a prospective study, the population or cohort is classified according to specific risk factors, such that the outcome of interest, typically various manifestations of a disease, can be observed over time and tied in to the initial classification. Also see *retrospective study*.

PrSS analysis. See *precision and sample-size analysis*.

PSS analysis. See *power and sample-size analysis*.

PSS Control Panel. The PSS Control Panel is a point-and-click graphical user interface for power and sample-size analysis. See [PSS-2] **GUI (power)**.

p-value. P-value is a probability of obtaining a test statistic as extreme or more extreme as the one observed in a sample assuming the null hypothesis is true.

R^2. See *coefficient of determination*.

random effects. Random effects represent a random sample of levels from all possible levels, and the interest lies in all possible levels.

randomized controlled trial. In this experimental study, treatments are randomly assigned to two or more groups of subjects.

RCT. See *randomized controlled trial*.

recruitment period. See *accrual period*.

reduced model. In the regression context, a reduced model is a regression model that contains only a subset of covariates from the corresponding full model. These covariates are referred to as "control covariates". The covariates that are not in the reduced model are referred to as "tested covariates".

reference value. See *null value*.

rejection region. In hypothesis testing, a rejection region is a set of sample values for which the null hypothesis can be rejected.

relative risk. See *risk ratio*.

retrospective study. In a retrospective study, a group with a disease of interest is compared with a group without the disease, and information is gathered in a retrospective way about the exposure in each group to various risk factors that might be associated with the disease. Also see *prospective study*.

risk difference. A risk difference is defined as the probability of an event occurring when a risk factor is increased by one unit minus the probability of the event occurring without the increase in the risk factor.

When the risk factor is binary, the risk difference is the probability of the outcome when the risk factor is present minus the probability when the risk factor is not present.

When one compares two populations, a risk difference is defined as a difference between the probabilities of an event in the two groups. It is typically a difference between the probability in the comparison group or experimental group and the probability in the reference group or control group.

risk factor. A risk factor is a variable that is associated with an increased or decreased probability of an outcome.

risk ratio. A risk ratio, also called a relative risk, measures the increase in the likelihood of an event occurring when a risk factor is increased by one unit. It is the ratio of the probability of the event when the risk factor is increased by one unit over the probability without that increase.

When the risk factor is binary, the risk ratio is the ratio of the probability of the event when the risk factor occurs over the probability when the risk factor does not occur.

When one compares two populations, a risk ratio is defined as a ratio of the probabilities of an event in the two groups. It is typically a ratio of the probability in the comparison group or experimental group to the probability in the reference group or control group.

sample size. This is the number of subjects in a sample. See [PSS-2] **Intro (power)** to learn more about the relationship between sample size and the power of a test.

sample-size curve. A sample-size curve is a graph of the estimated sample size as a function of some other study parameter such as power or CI width. The sample size is plotted on the y axis, and the power or other parameter is plotted on the x axis.

sample-size determination. This pertains to the computation of a sample size given either power and effect size or CI width and probability of CI width and any other study parameters. In a cluster randomized design, sample-size determination consists of determining the number of clusters given the cluster size or the cluster size given the number of clusters.

sample-size ratio. The ratio of the experimental-group sample size relative to the control-group sample size, n_2/n_1.

Satterthwaite's t test. Satterthwaite's t test is a modification of the two-sample t test to account for unequal variances in the two populations. See *Methods and formulas* in [PSS-2] **power twomeans** for details.

score test. A score test, also known as a Lagrange multiplier test, is one of the three classical testing procedures used to compare the fit of two models, one of which, the constrained model, is nested within the full (unconstrained) model. The null hypothesis is that the constrained model fits the data as well as the full model. The score test only requires one to fit the constrained model. See [PSS-2] **power oneproportion** and [R] **prtest**.

sensitivity analysis. Sensitivity analysis investigates the effect of varying study parameters on power, CI precision, probability of CI width, sample size, and other components of a study. The true values of study parameters are usually unknown, and analyses of power, precision, and sample size use best guesses for these values. It is therefore important to evaluate the sensitivity of the computed power, CI precision, or sample size in response to changes in study parameters. See [PSS-2] **power, table**, [PSS-2] **power, graph**, [PSS-3] **ciwidth, table**, and [PSS-3] **ciwidth, graph** for details.

sign test. A sign test is used to test the null hypothesis that the median of a distribution is equal to some reference value. A sign test is carried out as a test of binomial proportion with a reference value of 0.5. See [PSS-2] **power oneproportion** and [R] **bitest**.

significance level. In hypothesis testing, the significance level α is an upper bound for a probability of a type I error. See [PSS-2] **Intro (power)** to learn more about the relationship between significance level and the power of a test.

size of test. See *significance level*.

sphericity assumption. All differences between levels of the within-subject factor within-subject factor have the same variances.

stratified 2 × 2 tables. Stratified 2×2 tables describe the association between a binary independent variable and a binary response variable of interest. The analysis is stratified by a nominal (categorical) variable with K levels.

symmetry. In a 2×2 contingency table, symmetry refers to the equality of the off-diagonal elements. For a 2×2 table, a test of marginal homogeneity reduces to a test of symmetry.

t test. A t test is a test for which the sampling distribution of the test statistic is a Student's t distribution.

A one-sample t test is used to test whether the mean of a population is equal to a specified value when the variance must also be estimated. The test statistic follows Student's t distribution with $N - 1$ degrees of freedom, where N is the sample size.

A two-sample t test is used to test whether the means of two populations are equal when the variances of the populations must also be estimated. When the two populations' variances are unequal, a modification to the standard two-sample t test is used; see *Satterthwaite's t test*.

target parameter. In power and sample-size analysis, the target parameter is the parameter of interest or the parameter in the study about which hypothesis tests are conducted.

test statistic. In hypothesis testing, a test statistic is a function of the sample that does not depend on any unknown parameters.

tested covariates. See *reduced model*.

two-independent-samples test. See *two-sample test*.

two-sample paired test. See *paired test*.

two-sample test. A two-sample test is used to test whether the parameters of interest of the two independent populations are equal. For example, two-sample means test, two-sample variances, two-sample proportions test, two-sample correlations test.

two-sided confidence interval. A two-sided CI contains a plausible finite range of values for a parameter of interest. Two-sided CIs contain a finite upper limit for plausible values greater than the point estimate and a finite lower limit for plausible values less than the point estimate. See *Confidence intervals* in [PSS-3] **Intro (ciwidth)**.

two-sided test, **two-tailed test**. A two-sided test is a hypothesis test of a parameter in which the alternative hypothesis is the complement of the null hypothesis. In the context of a test of a scalar parameter, the alternative hypothesis states that the parameter is less than or greater than the value conjectured under the null hypothesis.

two-way ANOVA, **two-way analysis of variance**. A two-way ANOVA model contains two factors. Also see [PSS-2] **power twoway**.

two-way repeated-measures ANOVA, **two-factor** ANOVA. This is a repeated-measures ANOVA model with one within-subject factor and one between-subjects factor. The model can be additive (contain only main effects of the factors) or can contain main effects and an interaction between the two factors. Also see [PSS-2] **power repeated**.

type I error. The type I error of a test is the error of rejecting the null hypothesis when it is true; see [PSS-2] **Intro (power)** for more details.

type I error probability. See *probability of a type I error*.

type I study. A type I study is a study in which all subjects fail (or experience an event) by the end of the study; that is, no censoring of subjects occurs.

type II error. The type II error of a test is the error of not rejecting the null hypothesis when it is false; see [PSS-2] **Intro (power)** for more details.

type II error probability. See *probability of a type II error*.

type II study. A type II study is a study in which there are subjects who do not fail (or do not experience an event) by the end of the study. These subjects are known to be censored.

unbalanced design. An unbalanced design indicates an experiment in which the numbers of treated and untreated subjects differ. Also see [PSS-4] **Unbalanced designs**.

unequal-allocation design. See *unbalanced design*.

upper one-sided confidence interval. An upper one-sided confidence interval contains a range of values that are less than or equal to the upper confidence limit ul. The confidence interval is defined by a finite upper confidence limit and a lower confidence limit of negative infinity: $(-\infty, ul]$.

upper one-sided test, **upper one-tailed test**. An upper one-sided test is a one-sided test of a scalar parameter in which the alternative hypothesis is upper one sided, meaning that the alternative hypothesis states that the parameter is greater than the value conjectured under the null hypothesis. Also see *One-sided test versus two-sided test* under *Remarks and examples* in [PSS-2] **Intro (power)**.

Wald test. A Wald test is one of the three classical testing procedures used to compare the fit of two models, one of which, the constrained model, is nested within the full (unconstrained) model. Under the null hypothesis, the constrained model fits the data as well as the full model. The Wald test requires one to fit the full model but does not require one to fit the constrained model. Also see [PSS-2] **power oneproportion** and [R] **test**.

withdrawal. Withdrawal is the process under which subjects withdraw from a study for reasons unrelated to the event of interest. For example, withdrawal occurs if subjects move to a different area or decide to no longer participate in a study. Withdrawal should not be confused with administrative censoring. If subjects withdraw from the study, the information about the outcome those subjects would have experienced at the end of the study, had they completed the study, is unavailable. Also see *loss to follow-up* and *administrative censoring*.

within-subject design. This is an experiment that has at least one within-subject factor. See [PSS-2] **power repeated**.

within-subject factor. This is a factor for which each subject receives several of or all the levels.

z test. A z test is a test for which a potentially asymptotic sampling distribution of the test statistic is a normal distribution. For example, a one-sample z test of means is used to test whether the mean of a population is equal to a specified value when the variance is assumed to be known. The distribution of its test statistic is normal. See [PSS-2] **power onemean**, [PSS-2] **power twomeans**, and [PSS-2] **power pairedmeans**.