

Sample Size Computations and Power Analysis with the SAS[®] System

John M. Castelloe, SAS Institute Inc., Cary, NC

Abstract

Statistical power analysis characterizes the ability of a study to detect a meaningful effect size—for example, the difference between two population means. It also determines the sample size required to provide a desired power for an effect of scientific interest. Proper planning reduces the risk of conducting a study that will not produce useful results and determines the most sensitive design for the resources available. Power analysis is now integral to the health and behavioral sciences, and its use is steadily increasing wherever empirical studies are performed.

SAS Institute is working to implement power analysis for common situations such as *t*-tests, ANOVA, comparison of binomial proportions, equivalence testing, survival analysis, contingency tables and linear models, and eventually for a wide range of models and designs. An effective graphical user interface reveals the contribution to power of factors such as effect size, sample size, inherent variability, type I error rate, and choice of design and analysis. This presentation demonstrates issues involved in power analysis, summarizes the current state of methodology and software, and outlines future directions.

Introduction

Suppose you have performed a small study and are disappointed to find that the results are unexpectedly insignificant. Where did you go wrong? You may need to do a larger study to detect the effect you suspect, but how much larger?

Alternatively, suppose you have performed a large study and found a *hugely* significant effect. In follow-up studies, can you make more efficient use of resources by using smaller sample sizes?

Power analysis can optimize the resource usage and design of a study, improving chances of conclusive results with maximum efficiency. Power analysis is

most effective when performed at the study planning stage, and as such it encourages early collaboration between researcher and statistician. It also focuses attention on effect sizes and variability in the underlying scientific process, concepts that both researcher and statistician should consider carefully at this stage. Muller and Benignus (1992) and O'Brien and Muller (1993) provide cogent discussions of these and related concepts. These references also provide a good general introduction to power analysis.

There are many factors involved in a power analysis, such as the research objective, design, data analysis method, power, sample size, type I error, variability, and effect size. By performing a power analysis, you can learn about the relationships between these factors, optimizing those that are under your control and exploring the implications of those that are fixed.

For the purposes of statistical testing, the research objective is usually to use a feasible sample of data to assess a given hypothesis, H_1 , that some effect exists in a much larger population of potential data. If the sample data lead you to conclude that H_1 is true, but the opposite is really the case—that is, if the (null) hypothesis H_0 is true that there really is *no* effect—this is called a *type I error*. The probability of a type I error is usually designated “alpha” or α , and statistical tests are designed to ensure that α is suitably small (for example, less than 0.05). But it is also important to control the probability β of making the opposite (*type II*) error—that is, concluding H_0 , that there is no effect, when there really is one. The probability $1 - \beta$ of correctly rejecting H_0 when it is false is traditionally called the *power* of the test. (Note, however, that another more technical definition of power is the probability of rejecting H_0 for any given set of circumstances, even those corresponding to H_0 being true.)

Power analysis is often problematic in practice, being performed infrequently or improperly. There are several reasons for this: it is technically complicated, usually under-represented in statistical curricula, and often not performed early enough to be effective (that is, in the study planning stage). Good software tools

for power analysis can alleviate these difficulties and help you to benefit from these techniques.

Some Power Analysis Scenarios

There are several different varieties of power analysis. Here are a few simple scenarios:

- A statistician in a manufacturing company is reviewing a proposed experiment designed to assess the effect of various operating conditions on the quality of a product. He would like to conduct a power analysis to see if the planned number of replications and experimental units will be sufficient to detect the postulated effects.
- An advertising executive is interested in studying several alternative marketing strategies, with the aim of deciding how many and which strategies to implement. She would like to get a ballpark idea of how many mailings are necessary to detect differences in response rates.
- A study performed by a behavioral scientist (without a prior power analysis) did not detect a significant difference between success rates of two alternative therapies. He is considering repeating the study, but would first like to assess the power of the first study to detect the minimal effect size in which he is interested. A finding of low power would encourage him to repeat the study with a larger sample size or more efficient design.

Perhaps the most basic distinction in power analysis is that between *prospective* and *retrospective* analyses. In the examples above, the first two are prospective, while the third is retrospective. A prospective power analysis looks ahead to a future study, while a retrospective power analysis attempts to characterize a completed study. Sometimes the distinction is a bit fuzzy: for example, a retrospective analysis for a recently completed study can become a prospective analysis if it leads to the planning of a new study to address the same research objectives with improved resource allocation.

Although a retrospective analysis is the most convenient kind of power analysis to perform, it is often uninformative or misleading, especially when power is computed for the observed effect size. See the section “Effective Power Analysis” for more details.

Power analyses can also be characterized by the factor(s) of primary interest. For example, you might want to estimate power, determine required sample size, or assess detectable effect sizes. Sometimes

the research goal involves the largest acceptable confidence interval width instead of the significance of a hypothesis test; in this case, rather than considering the criterion of power, you would focus on the probability of achieving the desired confidence interval precision. There are also Bayesian approaches to sample size determination for estimating parameters or maximizing utility functions.

Example: Prospective Analysis for a Clinical Trial

The purpose of this example is to introduce some of the issues involved in power analysis and to demonstrate the use of some simple SAS[®] software tools for addressing them. Let’s say you are a clinical researcher wanting to compare the effect of two drugs, A and B, on systolic blood pressure (SBP). You have enough resources to recruit 25 subjects for each drug. Will this be enough to ensure a reasonable chance of establishing a significant result if the mean SBPs of patients on each drug really differ? In other words, will your study have good power? The answer depends on many factors:

- How big is the underlying effect size that you are trying to detect? That is, what is the population difference in mean SBP between patients using drug A and patients using drug B? Of course, this is unknown; that is why you are doing the study! But you can make an educated guess or set a goal for the detectable effect size. Then the power analysis determines the chance of detecting this conjectured effect size. For example, suppose you have some results from a previous study involving drug A, and you believe that the mean SBP for drug B differs by about 10% from the mean SBP for drug A. If the mean SBP for drug A is 120, you thus posit an effect size of 12.
- What is the inherent variability in SBP? Suppose previous studies involving drug A have shown the standard deviation of SBP to be between 11 and 15, and that the standard deviations are expected to be roughly the same for the two drugs. You want to consider this range of variability in your power analysis.
- What data analysis method and level of type I error should you use? You decide to use the simple approach of a two-sample *t*-test (assuming equal variances) with $\alpha = 0.05$. To be conservative you use a two-sided test, although you suspect the mean SBP for drug B is higher.

With these specifications, the power can be computed

using the noncentral F distribution. The following SAS statements compute this power for the standard deviation of 15:

```
data twosample;
  Mu1=120; Mu2=132; StDev=15;
  N1=25; N2=25; Alpha=0.05;
  NCP = (Mu2-Mu1)**2/((StDev**2)*
    (1/N1 + 1/N2));
  CriticalValue = FINV(1-Alpha, 1,
    N1+N2-2, 0);
  Power = SDF('F', CriticalValue,
    1, N1+N2-2, NCP);
proc print data=twosample;
run;
```

The noncentrality parameter NCP is calculated from the conjectured means μ_1 and μ_2 , sample sizes N_1 and N_2 , and common standard deviation StDev . The critical value of the test statistic is then computed, and the power is the probability of a noncentral-F random variable with noncentrality parameter NCP, one numerator degree of freedom, and $N_1 + N_2 - 2$ denominator degrees of freedom exceeding this critical value. This probability is computed using the DATA step function SDF, which calculates survival distribution function values. In general, $\text{SDF} = 1 - \text{CDF}$; the SDF form is more accurate for values in the upper tail. The CDF and SDF functions, introduced in Release 6.11 and 6.12 of the SAS System, respectively, are documented in SAS Institute Inc. (1999b). Their use is recommended for applications requiring their enhanced numerical accuracy.

The resulting power is about 79%. If you would really like a power of 85% or more when the standard deviation is 15, then you will need more subjects. How many? One way to investigate required sample size is to construct a power curve, as shown in Figure 1. This curve was generated using the Sample Size task in the Analyst Application. Note that a sample size of 30 for each group would be sufficient to achieve 85% power.

Now suppose that a colleague brings to your attention the possibility of using a simple AB/BA cross-over design. Half of the subjects would get 6 weeks on drug A, a 4-week washout period, and then 6 weeks on drug B; the other half would follow the same pattern but with drug order switched. Assuming there are no period or carry-over effects, you can use a paired t -test to assess the difference between the two drugs. Each pair consists of the SBP for a patient while using drug A and the SBP for that same patient while using drug B. Suppose previous studies have shown that there is correlation of roughly $\rho = 0.8$ between pairs of SBP measurements for each subject.

What would the power for the study be if you use this




Figure 1. Power Curve for Two-Sample t -test

cross-over design with 25 subjects? You simply need to calculate the standard deviation of a pair difference, which is given by

$$\sigma_{\Delta} = \sqrt{\sigma_1^2 + \sigma_2^2 - 2\rho\sigma_1\sigma_2}$$

where σ_1 and σ_2 are the standard deviations for the two drug types (assumed to be equal in this case). The resulting values are $\sigma_{\Delta} = 6.96$ when $\sigma_1 = \sigma_2 = 11$, and $\sigma_{\Delta} = 9.49$ when $\sigma_1 = \sigma_2 = 15$. The following SAS statements compute the power for the larger standard deviation:

```
data paired;
  Mu1=120; Mu2=132; StDev1=15;
  StDev2=15; Corr=0.8; N=25;
  Alpha=0.05;
  StDevDiff = sqrt(StDev1**2 +
    StDev2**2 -
    2*Corr*StDev1*StDev2);
  NCP = (Mu2-Mu1)**2 /
    (StDevDiff**2/N);
  CriticalValue = FINV(1-Alpha, 1,
    N-1, 0);
  Power = SDF('F', CriticalValue,
    1, N-1, NCP);
proc print data=paired;
run;
```

The resulting power is over 99% with 25 subjects. A power curve generated using the Analyst Application, displayed in Figure 2, reveals that 85% power

is achieved with only 8 subjects in the cross-over design.

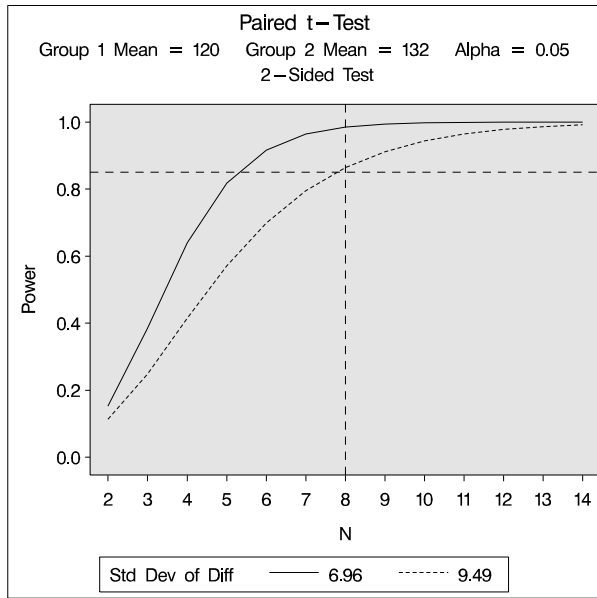


Figure 2. Power Curve for Paired *t*-test.

This example points out the need for the following power analysis tools:

- direct computation of the required sample size for 85% power
- determination of the effect size that can be detected with 85% power
- automatic generation of presentation-ready graphs, tables, and narrative reports
- assistants for computation of required input parameters

The Analyst Application computes basic sample sizes and constructs power curves; the other capabilities are planned for new software currently under development.

Effective Power Analysis

Power analysis is most effective when performed as part of study planning. Several issues must be addressed. Muller and Benignus (1992, p. 216) list five primary considerations: opportunity costs, ethical trade-offs, the size of effect considered important, the uncertainty of parameter estimates, and the analyst's preference for amount of power. Effective power analysis software represents the relationship between all components clearly, revealing which quantities must be specified to compute the quantity of interest. It can provide a guiding structure for the collaboration

between researcher and statistician. A useful representation of the components focuses attention on the scientific issues and terminology rather than the mathematical details of computations. Other helpful features are intuitive organization of results and facilities for exporting results to formats that can be incorporated directly into reports.

There is some confusion in practice about the meaning of "the size of effect considered important." Just how should effect size be postulated? One alternative is to specify the effect size that represents minimal clinical significance; then the result of the power analysis reveals the chances of detecting a minimally meaningful effect size. Often this minimal effect size is so small that it requires excessive resources to detect. Another alternative is to make an educated guess of the true underlying effect size. Then the power analysis determines the chance of detecting the effect size that is believed to be true. The choice is ultimately determined by the research goals. Finally, you can specify a collection of possible values, perhaps spanning the range between minimally meaningful effects and larger surmised effects.

You can arrive at values for required quantities in a power analysis, such as effect sizes and measures of variability, in many different ways. For example, you can use pilot data, results of previous studies reported in literature, educated guesses derived from theory, or educated guesses derived from partial data (a small sample or even just quantiles). Effective software assists you in using pilot data and partial descriptions to obtain reasonable values.

Uncertainty is a fact of life in any power analysis, since at least some of the numbers used are best guesses of unknown values. The result of a power calculation, whether it be achieved power or required sample size or something else, serves only as a point estimate, conditional on the conjectured values of the other components. It is not feasible in general to quantify the variability involved in using educated guesses or undocumented results to specify these components. If observed data are used, relevant adjustments for variability in the data tend to be problematic in the sense of producing confidence intervals for power that are too wide for practical use. But there is a useful way for you to characterize the uncertainty in your power analysis, and also discover the extent to which statistical power is affected by each component. You can posit a reasonable range for each input component, vary each one within its range, and observe the variety of results in the form of tables or graphs. A good general power analysis package makes such a sensitivity analysis easy to perform.

You need to take special care in retrospective power

analysis. Thomas (1997) explains why the use of both the observed effect size and observed variability is generally uninformative. Other authors demonstrate a different viewpoint; for example, Taylor and Muller (1996) use the observed means and variances in a study, along with new sample sizes, to assess the power of a planned replicate study with more subjects. Bias is introduced not only by the use of observed statistics, but also by the tendency to conduct retrospective analyses more often for insignificant study results. Numerous new methods account for sources of bias and uncertainty in power estimates in certain cases. Refer to Muller and Pasour (1997), Taylor and Muller (1996), and O'Brien and Muller (1993) for results in linear models. These methods may be incorporated into future practice as they become more standardized and thorough in coverage.

Currently Available Methods

Power analysis methodology has been well-developed and standardized for some statistical methods, such as paired and pooled t -tests, fixed-effect ANOVA and regression models, binomial proportion comparisons, bioequivalence, correlation, and simple survival analysis models. However, for a surprising number of models and tests, if power analysis methods exist at all, they are approximate, sometimes unreliably so. Often, researchers are forced to do the analysis for a simplified version of the situation and hope that it extrapolates accurately to the one at hand.

For some statistical models and tests, power analysis calculations are exact in the sense of utilizing a mathematical formula that expresses power directly in terms of the other components. Such formulas typically involve either enumeration or noncentral versions of the distribution of the test statistic. In the absence of exact mathematical results, approximate formulas can sometimes be used. When neither exact power computations nor reasonable approximations are possible, simulation provides an increasingly viable alternative. You specify values for model parameters and use them to randomly generate a large number of hypothetical data sets. Applying the statistical test to each data set, you estimate power with the percentage of times the null hypothesis is rejected. While the simulation approach is computationally intensive, faster computing makes this less of an issue. A simulation-based power analysis is always a valid option, and, with a large number of data set replications, it can often be more accurate than approximations.

Currently, you can perform a variety of power analyses with the SAS System:

- The Analyst Application implements power analyses for t -tests, confidence interval precision, equivalence of means, and one-way ANOVA, with capacity to solve for either power or sample size and generate power curves.
- UnifyPow, a SAS macro described in O'Brien (1998), covers a wide range of methods, including parametric and nonparametric methods for mean comparisons, proportions, logit analysis, log-linear models, regression, and correlation. Its primary strengths include the use of many exact power computations and support for unbalanced designs where relevant.
- A SAS/IML macro described in Muller, LaVange, Ramey, and Ramey (1992) performs power computations for general multivariate linear models with fixed effects. This macro is demonstrated in Timm and Mieczkowski (1997, pp. 253–254 and 523–550) and is available in the SAS Online Samples library.

Recent Research Developments

Power analysis methodology is an active area of research. A collection of short descriptions and primary references for some promising recent results are given below, categorized by statistical models and tests:

- *Fixed-effect linear models*: Methods for these are probably the most thoroughly developed. Power can be calculated exactly using noncentral-F and noncentral- t distributions for many special cases such as t -tests and ANOVA (with contrasts). Good power approximations are available for general multivariate linear models. O'Brien and Muller (1993) and O'Brien and Shieh (2000) present a thorough discussion.
- *Satterthwaite (unpooled) t -test*: DiSantostefano and Muller (1995) outline several alternative power approximations using noncentral-F and noncentral- t distributions, along with recommendations.
- *Mixed models*: Power analysis with random effects is an area of ongoing research, and the best current approach is simulation.

- *Comparison of two binomial proportions:* There are a number of statistical tests used in practice. The usual Pearson chi-square test (or the equivalent z-test) is appropriate for large samples. There are several available exact unconditional tests that perform well even for small sample sizes. Fisher's Exact test continues to be widely used. Power analysis methods have been developed for most of these tests, ranging from exact power computations to reasonably good power approximations. Refer to O'Brien and Muller (1993, section 8.5.3), Dozier and Muller (1993), Suissa and Shuster (1985), and Fleiss (1981, Chapter 3). For comparison of correlated proportions, refer to Selicato and Muller (1998).
- *Contingency tables:* Power analysis methods for relevant tests, such as contrasts among several proportions or independence in a RxC contingency table, generally involve the noncentral chi-square distribution and consist of exact computations or good approximations.
- *Equivalence of two means or proportions:* In these tests, the null hypothesis specifies a nonzero difference. Several approximate and exact power computation methods are available. For means, refer to Phillips (1990) for the additive model with normal data and Diletti et al. (1991) for the multiplicative model with lognormal data. For proportions, refer to Frick (1994).
- *Correlation coefficients:* In analyses involving one correlation coefficient, or comparison of two, power computations for statistical tests involving Fisher's z-transformation and normal approximations are simply based on the normal distribution. Power for tests of any correlation that arises in a fixed-effect general linear model can be computed using the methods described above for fixed-effect linear models. Gatsonis and Sampson (1989) outline power computations for tests of multiple or partial correlation in a model with all Gaussian predictors and a Gaussian response.
- *Nonparametric tests:* Power analysis methods for tests such as the Wilcoxon and Wilcoxon-Mann-Whitney are complicated, but involve good approximations. Refer to O'Brien (1998) and Noether (1987).
- *Generalized linear models:* There are currently no standardized approaches for power analysis in generalized linear models. Available approximate methods, described in Self, Mauritsen, and Ohara (1992), require discretization of

continuous predictors, and they have not been tested on a wide range of models. Simulation is a safe alternative for complicated models.

- *Logistic regression:* The approach of Self, Mauritsen, and Ohara (1992) can be used, with the necessary discretization of continuous predictors. Another power approximation is provided by Whittemore (1981), with a recommended restriction to small response probabilities. Simulation is the preferred method whenever one doubts the accuracy of these approximations in a given situation.
- *Log-linear models:* Power computations described in O'Brien (1986), Agresti (1990, pp. 241–244) and O'Brien (1998), involving exemplary data (a hypothetical data set specified to produce conjectured parameter values) and the noncentral chi-square distribution, provide good approximations that are similar to those in Self, Mauritsen, and Ohara (1992) but easier to calculate.
- *Survival data:* Power analysis for survival data is particularly complex. Current mainstream methodology covers straightforward situations, such as the comparison of two survival distributions assuming either exponential survival rates or proportional hazards: refer to Shuster (1990), Lachin (1981), Donner (1984), and Goldman and Hillman (1992). When factors such as accrual rates and dropout rates are involved, it is usually necessary to compute power via simulation.

Conclusion

Power analysis is a vital tool for study planning. The standard statistical testing paradigm implicitly assumes that type I errors (mistakenly concluding significance when there is no true effect) are more costly than type II errors (missing a truly significant result). This may be appropriate for your situation, or the relative costs of the two types of error may be reversed. For example, in screening experiments for drug development, it is often less damaging to carry a few false positives forward for follow-up testing than to miss potential leads. A power analysis can help you optimize your studies to achieve your desired balance between type I and type II errors. You can improve your chances of detecting effects that might otherwise have been ignored; you can save money and time, and perhaps minimize risks to subjects, with optimal designs and sample sizes.

While many basic sample size computations are currently available with the SAS System, there is room for

an integrated user interface that provides more analyses with more features. Developers are currently at work on such an application, and SUGI attendees can see the work in progress in the demo room.

Acknowledgments

I am grateful to Virginia Clark, Keith Muller, Ralph O'Brien, Bob Rodriguez, Maura Stokes, and Randy Tobias for valuable assistance in the preparation of this paper.

References

- Agresti, A. (1990), *Categorical Data Analysis*, New York: John Wiley & Sons.
- Diletti, D., Hauschke, D., and Steinijans, V.W. (1991), "Sample Size Determination for Bioequivalence Assessment by Means of Confidence Intervals," *International Journal of Clinical Pharmacology, Therapy and Toxicology*, 29, 1–8.
- DiSantostefano, R.L. and Muller, K.E. (1995), "A Comparison of Power Approximations for Satterthwaite's Test," *Communications in Statistics – Simulation and Computation*, 24, 583–593.
- Donner, A. (1984), "Approaches to Sample Size Estimation in the Design of Clinical Trials—A Review," *Statistics in Medicine*, 3, 199–214.
- Dozier, W.G. and Muller, K.E. (1993), "Small-Sample Power of Uncorrected and Satterthwaite Corrected t Tests for Comparing Binomial Proportions," *Communications in Statistics – Simulation*, 22, 245–264.
- Fleiss, J.L. (1981), *Statistical Methods for Rates and Proportions*, New York: Wiley & Sons.
- Frick, H. (1994), "On Approximate and Exact Sample Size for Equivalence Tests for Binomial Proportions," *Biometrical Journal*, 36, 841–854.
- Gatsonis, C. and Sampson, A.R. (1989), "Multiple Correlation: Exact Power and Sample Size Calculations," *Psychological Bulletin*, 106, 516–524.
- Goldman, A.I. and Hillman, D.W. (1992), "Exemplary Data: Sample Size and Power in the Design of Event-Time Clinical Trials," *Controlled Clinical Trials*, 13, 256–271.
- Lachin, J.M. (1981), "Introduction to Sample Size Determination and Power Analysis for Clinical Trials," *Controlled Clinical Trials*, 2, 93–113.
- Muller, K.E. and Benignus, V.A. (1992), "Increasing Scientific Power with Statistical Power," *Neurotoxicology and Teratology*, 14, 211–219.
- Muller, K.E., LaVange, L.M., Ramey, S.L., and Ramey, C.T. (1992), "Power Calculations for General Linear Multivariate Models Including Repeated Measures Applications," *Journal of the American Statistical Association*, 87, 1209–1226.
- Muller, K.E. and Pasour, V.B. (1997), "Bias in Linear Model Power and Sample Size Due to Estimating Variance," *Communications in Statistics – Theory and Methods*, 26, 839–851.
- Noether, G.E. (1987), "Sample Size Determination for Some Common Nonparametric Tests," *Journal of the American Statistical Association*, 82, 645–647.
- O'Brien, R.G. (1986), "Using the SAS System to Perform Power Analyses for Log-linear Models," *Proceedings of the Eleventh Annual SAS Users Group International Conference*, Cary, NC: SAS Institute Inc., 778–782.
- O'Brien, R.G. (1998), "A Tour of UnifyPow: A SAS Module/Macro for Sample-Size Analysis," *Proceedings of the Twenty-Third Annual SAS Users Group International Conference*, Cary, NC: SAS Institute Inc., 1346–1355. Software and updates to this article can be found at www.bio.ri.ccf.org/UnifyPow.
- O'Brien, R.G. and Muller, K.E. (1993), "Unified Power Analysis for t-Tests Through Multivariate Hypotheses," In Edwards, L.K., ed. (1993), *Applied Analysis of Variance in Behavioral Science*, New York: Marcel Dekker, Chapter 8, 297–344.
- O'Brien, R.G. and Shieh, G. (2000), "Pragmatic, Unifying Algorithm Gives Power Probabilities for Common F Tests of the Multivariate General Linear Hypothesis." Manuscript downloadable in PDF form from www.bio.ri.ccf.org/UnifyPow.
- Phillips, K.F. (1990), "Power of the Two One-sided Tests Procedure in Bioequivalence," *Journal of Pharmacokinetics and Biopharmaceutics*, 18, 137–144.
- SAS Institute Inc. (1999a). *The Analyst Application, First Edition*, Cary, NC: SAS Institute Inc.
- SAS Institute Inc. (1999b). *SAS Language Reference: Dictionary, Version 7-1*, Cary, NC: SAS Institute Inc.
- Self, S.G., Mauritsen, R.H., and Ohara, J. (1992), "Power Calculations for Likelihood Ratio Tests in Generalized Linear Models," *Biometrics*, 48, 31–39.
- Selicato, G.R. and Muller, K.E. (1998), "Approximating Power of the Unconditional Test for Correlated

Binary Pairs," *Communications in Statistics –Simulation*, 27, 553–564.

Shuster, J.J. (1990), *Handbook of Sample Size Guidelines for Clinical Trials*, Boca Raton, FL: CRC Press.

Suissa, S. and Shuster, J.J. (1985), "Exact Unconditional Sample Sizes for the 2 X 2 Comparative Trial," *Journal of the Royal Statistical Society A*, 148, 317–327.

Taylor, D.J. and Muller, K.E. (1996), "Bias in Linear Model Power and Sample Size Calculation Due to Estimating Noncentrality," *Communications in Statistics –Theory and Methods*, 25, 1595–1610.

Thomas, L. (1997), "Retrospective Power Analysis," *Conservation Biology* 11, 276–280.

Timm, N.H. and Mieczkowski, T.A. (1997), *Univariate and Multivariate General Linear Models: Theory and Applications Using SAS® Software*, Cary, NC: SAS Institute Inc.

Whittemore, A.S. (1981), "Sample Size for Logistic Regression with Small Response Probability," *Journal of the American Statistical Association*, 76, 27–32.

SAS and SAS/IML software are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. ® indicates USA registration.

Other brand and product names are registered trademarks or trademarks of their respective companies.

Contact Information

John M. Castelloe, SAS Institute Inc., SAS Campus Drive, Cary, NC 27513. Phone (919) 531-5728, FAX (919) 677-4444, Email John.Castelloe@sas.com.

Version 1.2