

Paper 430-2011

MCPOWER: a Flexible Macro Suite for Generating Monte Carlo Power Estimates for Linear, Logistic, and Poisson Regression Models

Ken Kleinman, Harvard Medical School and Harvard Pilgrim Health Care Institute,
Boston, MA

ABSTRACT

I describe and demonstrate a macro suite which estimates power for linear, logistic, and Poisson regression using Monte Carlo methods. In this approach one repeatedly simulates data sets under the alternative and estimates the power as the proportion of data sets in which the null can be rejected. The macro allows a wide variety of predictor distributions and produces raw estimates as well as smoothed estimates which draw strength across the assessed sample sizes. It optionally generates plots and printouts including an interpolated estimate of the number of subjects necessary to achieve a target power.

INTRODUCTION

In statistics we spend a lot of time and theoretical energy ensuring that the size of the test is correct. The size, denoted α , is the probability that we make an error in rejecting the null hypothesis. In other words, if the truth is that there is no effect, α is the probability we reject it anyway.

The other kind of mistake we can make is if we conclude there is no effect when in fact there is one. The probability of this is denoted β . The power of a test, the probability we correctly conclude there is a difference when there is one, is $1 - \beta$. In general, power has received less analytic attention than size. However, the power is crucial in planning a study: without it we might waste time conducting a study with a very small chance of finding the effect we're curious about, or waste money collecting data on too many people.

Some of the most commonly used statistical methods are generalized linear models, which include linear, logistic, and Poisson (or count) regression. For linear regression, analytic power is simple: simple formulae show the power to detect slopes measured in standard deviation units of the outcome per standard deviation of the predictor. (This applies for dichotomous predictors as well.) These analytic results are implemented in PROC POWER and PROC GLMPower. For logistic and Poisson regression models, no general analytic results exist. Thus, for example, PROC POWER uses approximations to provide power for logistic regression. No SAS® procedures provide power for Poisson regression.

An alternative approach to analytic power calculation is to use Monte Carlo methods. Briefly, the notion here is to first observe that the notion of power requires specifying a particular alternative hypothesis. This is the alternative, which is true, making the null hypothesis of no effect be false. The Monte Carlo approach to power calculation is to simulate data under this alternative and determine whether the null hypothesis is rejected or not. Repeating this process many times provides increasingly precise estimates of the probability of rejecting the null under that alternative. And the probability the null is rejected is the power, if the alternative is true. Neither approximations nor complex mathematics tailored to the setting are required.

In this paper I describe the SAS macro suite MCPOWER which implements Monte Carlo power calculations for a spectrum of models. The macro uses the generality available in the RAND function to allow the user to choose between normal, Bernoulli, and Poisson outcomes for linear, logistic, and Poisson regression, respectively. It also allows a wide assortment of covariate distributions. It provides options to allow a range of alternatives and sample sizes to be generated in a single macro call. The output includes raw estimates of power as well as a smoothed estimate which draws information across a series of alternatives. Optionally, a plot is generated using PROC GPLOT or PROC SGPanel, and the sample size associated with 80% and 90% power for each alternative is estimated via interpolation, if possible. There is also a menu-based access to the macro suite through a %WINDOW macro.

MACRO DESCRIPTIONS AND FUNCTION

The MPOWER suite includes two macros. The %MANYPOWER macro is called by the user. It interprets the submitted parameters to identify each specific alternative hypothesis implied by the macro call. This may include a range of sample sizes and effect sizes (slopes, odds ratios, or risk ratios). It also collects details about the outcome and predictor distribution, the null hypothesis, and the number of times to simulate the data sets, as well as the desired output.

The %GLMPOWER macro is called by %MANYPOWER. It generates data sets under a single alternative, i.e., with a given sample size and a single effect size or linear slope across the predictor, and with a specified outcome distribution and predictor distribution. It then fits the model to each data set and determines whether the specified null hypothesis was rejected or not.

The definition of %MANYPOWER is shown below; the macros themselves are available for download from <http://www.kenkleinman.net/home/index.php/sas-and-r-code/sas-macros.html>.

```
%macro manypower(alt1=.,alt2=.,alt3=.,
                 n1=.,n2=.,n3=.,n4=.,n5=.,n6=.,n7=.,n8=.,n9=.,n10=.,
                 outcomedist = "NORMAL", null=0, scaley=1,
                 xdist="NORMAL", xloc=., xscale =.,
                 nreps=,plotrows=2, descending = desc, intercept = 0,
                 plotstyle=graphic, outputpower=YES, outputnsubs=YES, test=FALSE);
```

The macro options are described below, along with defaults, explanations, and any limitations.

Distributions:

Parameter	Default	Purpose
outcomedist	"NORMAL"	defines the type of regression: "BERNOULLI" for logistic, "POISSON" for Poisson, "NORMAL" for linear
scaley	1	residual variance of the outcome, if outcome is normal; ignored otherwise
intercept	0	intercept for the model generating the outcome data
xdist	"NORMAL"	distribution of x, uses the rand function
xloc	.	first parameter for the x distribution
xscale	.	second parameter for the x distribution

Table 1. %MANYPOWER macro parameters about the predictor and outcome distributions

The intercept is key for logistic regression, for which it is the probability the outcome is 1 if the covariate $x = 0$, and Poisson regression, for which it is the mean of the outcome when the covariate $x = 0$. For linear regression it is irrelevant. Any distribution from the RAND function can be used for the predictor x , provided it uses 0, 1, or 2 parameters. This excludes only the hypergeometric distribution and the multinomial ("TABLE") with more than three categories. The first parameter is referred to as *xloc*, within the macro, while the second is *xscale*. If only one parameter is needed, *xscale* **must** have a value of "." (without the quotes); if no parameter is needed, the both *xloc* and *xscale* **must** be ".". For the special case of the "EXPONENTIAL" distribution, a mean can be submitted as the *xloc* and this will be implemented by the macro, though the RAND function generates exponentials with mean = 1.

Hypotheses:

Parameter	Default	Purpose
null	0	the comparator for the estimated parameters from the models, in terms of model parameters; see note below

alt1, alt2, alt3		truths under the alternative—slope (if the <i>outcomedist</i> is "NORMAL"), odds ratio ("BERNOULLI"), or risk ratio ("POISSON")
n1,n2, ... n10		number of subjects to simulate in each data set

Table 2. %MANYPOWER macro parameters about the null and alternative hypotheses

The *null* hypothesis parameter, unlike the *alternative* truths, is expressed as a model parameter, meaning the **log** odds ratio for logistic regression and the **log** risk ratio for Poisson regression. The default value of 0 implies no slope for linear regression and an odds ratio or risk ratio of 1 for logistic and Poisson regression respectively.

Output:

Parameter	Default	Purpose/notes
outputpower	YES	prints a table with the actual and smoothed estimated power for each alternative and n; any other value will suppress printing the table
outputnsubs	YES	prints a table with the linearly interpolated n (and confidence interval for n) required for power = .1, ... , .9, within the range of the estimated power, for each alternative; any other value will suppress the table
plotstyle	graphic	if plotstyle=sg, the macro produces a PROC SGPANEL plot with a panel for each alternative slope; if plotstyle=graphic, produces a PROC GPLOT plot with a curve for each alternative slope; if any other value, no plot is made
plotrows	2	number of rows if plotstyle = sg, otherwise ignored

Table 3. %MANYPOWER macro parameters about the output

In the output data set, named *probitpower*, the actual estimated power is called "power." Smoothed power is called "ppower." This data set is created regardless of whether *outputpower* = YES or not.

EXAMPLE CODE AND OUTPUT

Here I show the settings call for Poisson regression with: dichotomous predictor with $p = .5$; 100 data sets per alternative value vs. null of log rate ratio = 0; 1000, 1100, 1200, 1400, 1700, and 2000 subjects per simulation; alternative rate ratios of 1.15, 1.2, 1.25; PROC GPLOT output; and the Poisson outcome has a mean of 1 when $x = 0$. The results are shown in figures 1, 2, and 3.

```
%manypower(null = 0, scaley=1,xdist = "bernoulli", xloc=.5, xscale=.,
            nreps = 100, n1=1000,n2=1100,n3=1200,n4=1400,n5=1700,n6=2000,
            alt1 = 1.15, alt2 = 1.20, alt3 = 1.25,
            plotrows=2,plotstyle=graphic, outcomedist="poisson", intercept = 1);
```

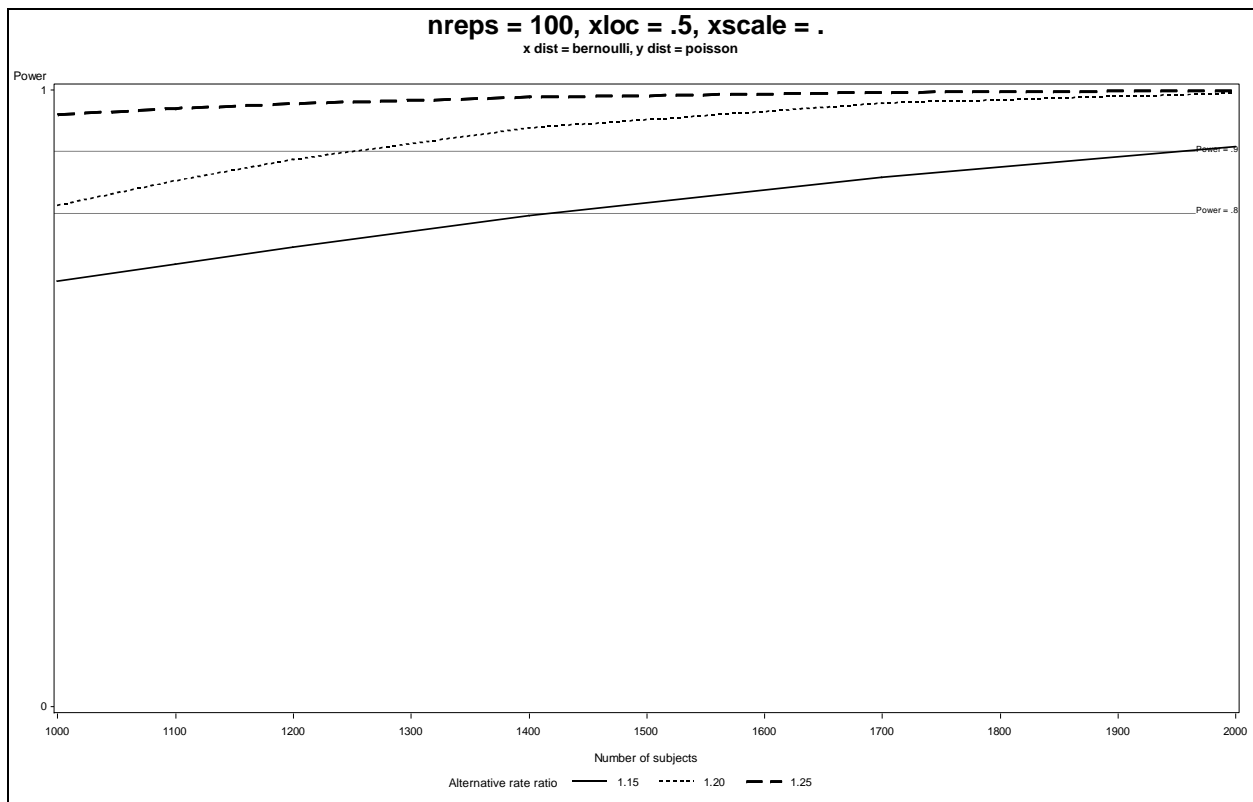


Figure 1. Plot results with a smoothed line for each alternative risk ratio

Obs	value	reps	nsubs	altslope	power	ppower	lowclpower	upclpower
1	69	100	1000	1.15	0.69	0.68855	0.68271	0.69435
2	74	100	1100	1.15	0.74	0.71723	0.71243	0.72199
3	69	100	1200	1.15	0.69	0.74457	0.74057	0.74855
4	84	100	1400	1.15	0.84	0.79482	0.79149	0.79813
5	84	100	1700	1.15	0.84	0.85812	0.85422	0.86195
6	91	100	2000	1.15	0.91	0.90668	0.90225	0.91095
7	81	100	1000	1.20	0.81	0.81192	0.80647	0.81727
8	85	100	1100	1.20	0.85	0.85239	0.84854	0.85617
9	88	100	1200	1.20	0.88	0.88657	0.88365	0.88943
10	95	100	1400	1.20	0.95	0.93724	0.93483	0.93957
11	98	100	1700	1.20	0.98	0.97816	0.97631	0.97989
12	99	100	2000	1.20	0.99	0.99383	0.99287	0.99468
13	97	100	1000	1.25	0.97	0.95992	0.95703	0.96265
14	95	100	1100	1.25	0.95	0.96975	0.96787	0.97154
15	100	100	1200	1.25	1.00	0.97749	0.97612	0.97880
16	97	100	1400	1.25	0.97	0.98806	0.98694	0.98910
17	100	100	1700	1.25	1.00	0.99587	0.99503	0.99658
18	100	100	2000	1.25	1.00	0.99875	0.99829	0.99909

Figure 2. Power results with estimated power (power) and smoothed power (ppower) with upper and lower 95% confidence intervals for each smoothed power.

Obs	altslope	desiredpower	nsubs Req	lowcInsubs Req	upcInsubs Req
1	1.15	0.7	1040	1021	1059
2	1.15	0.8	1425	1409	1441
3	1.15	0.9	1959	1933	1986
4	1.20	0.9	1254	1243	1264

Figure 3. Number of subjects required results with upper and lower confidence limits for desired power

CONCLUSION

The MCPower macro suite includes the common settings of linear, logistic, and Poisson regression. It allows a wide variety of covariate distributions to be specified. Data set, printed, and graphical output are available. PROC POWER and PROC GLMPower provide access to analytic results for linear models and approximate analytic results for logistic regression, as well as to ANOVA and multiple regression generally, survival analyses, table-based tests, correlations, and one-sample tests. In contrast, the MCPower macro suite provides empirical but non-approximate results for logistic regression with a single predictor, as well as Poisson regression. Its linear regression results are wholly replicable with analytic results including those available in PROC POWER, but its logistic and Poisson regression results cannot be replicated by PROC POWER. PROC GLMPower provides access to complex study designs, but is limited to normal outcomes. The MCPower macro suite cannot generate results equivalent to PROC GLMPower.

Future plans for the MCPower suite include the inclusion of negative binomial regression (an approach to count regression which allows more flexibility), the extension to interaction effects, and a separate extension for clustered data.

ACKNOWLEDGEMENTS

This work was supported by the Department of Population Medicine, Harvard Medical School and Harvard Pilgrim Health Care.

SAS and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. ® indicates USA registration.

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

CONTACT INFORMATION

Your comments and questions would be welcome and valued. Contact the author at:

Ken Kleinman
 Harvard Medical School and Harvard Pilgrim Health Care
 133 Brookline Ave.
 Boston, MA 02215
 ken.kleinman@gmail.com