Chapter 755

Assurance for Tests for the Difference Between Two Poisson Rates

Introduction

This procedure calculates the assurance of one- or two-sided tests of the difference between two independent Poisson event rates. The Poisson probability law gives the probability distribution of the number of events occurring in a specified interval of time or space. The Poisson distribution is characterized by a single parameter which is the mean number of occurrences during the specified interval. The Poisson distribution is often used to fit count data, such as the number of legions on a subject's legs or the number of episodes during a year.

The calculation is based on a user-specified prior distribution of the effect size parameters. This procedure may also be used to determine the needed sample size to obtain a specified assurance. The methods for assurance calculation in this procedure are based on O'Hagan, Stevens, and Campbell (2005).

Assurance

The assurance of a design is the expected value of the power with respect to one or more prior distributions of the design parameters. Assurance is also referred to as *Bayesian assurance*, *expected power*, *average power*, *statistical assurance*, *hybrid classical-Bayesian procedure,* or *probability of success*.

The power of a design is the probability of rejecting the null hypothesis, conditional on a given set of design attributes, such as the test statistic, the significance level, the sample size, and the effect size to be detected. As the effect size parameters are typically unknown quantities, the stated power may be very different from the true power if the specified parameter values are inaccurate.

While power is conditional on individual design parameter values, and is highly sensitive to those values, assurance is the average power across a presumed prior distribution of the effect size parameters. Thus, assurance adds a Bayesian element to the frequentist framework, resulting in a hybrid approach to the probability of trial or study success. It should be noted that when it comes time to perform the statistical test on the resulting data, these methods for calculating assurance assume that the traditional (frequentist) tests will be used.

The next section describes some of the ways in which the prior distributions for effect size parameters may be determined.

Elicitation

In order to calculate assurance, a suitable prior distribution for the effect size parameters must be determined. This process is called the *elicitation* of the prior distribution.

The elicitation may be as simple as choosing a distribution that seems plausible for the parameter(s) of interest, or as complex as combining the informed advice of several experts based on experience in the field, available pilot data, or previous studies. The accuracy of the assurance value depends on the accuracy

of the elicited prior distribution. The assumption (or hope) is that an informed prior distribution will produce a more accurate estimate of the probability of trial success than a single value estimate. Because clinical trials and other studies are often costly, many institutions now routinely require an elicitation step.

Two reference texts that focus on elicitation are O'Hagan, Buck, Daneshkhah, Eiser, Garthwaite, Jenkinson, Oakley, and Rakow (2006) and Dias, Morton, and Quigley (2018).

Test Procedure

These results follow Mathews (2010). Assume that all subjects in each group are observed for a fixed time period (which we set as one without loss of generality) and the number of events (outcomes or defects) is recorded. The following table presents the various terms used.

Group	1	2
Sample Size	N_1	N_2
Number of Events	X_1	X_2
Event Rate	λ_1	λ_2
Distribution of X	Poisson(λ_1)	Poisson(λ_2)

Mathews (2010) proposed two test statistics that can be used to test statistical hypotheses about the rate difference. The first is based is the *large-sample z-test* of the hypotheses H0: $\lambda_1 = \lambda_2$ versus Ha: $\lambda_1 \neq \lambda_2$.

$$z_{LS} = \frac{\hat{\lambda}_2 - \hat{\lambda}_1}{\sqrt{\dfrac{\hat{\lambda}_1}{N_1} + \dfrac{\hat{\lambda}_2}{N_2}}}$$

where $\hat{\lambda}_i = X_i/N_i$. This test is appropriate when $X_i > 30$.

The power of the one-sided level α hypothesis test is given by

$$Power = \Phi\left[\frac{\lambda_2 - \lambda_1}{\sqrt{\dfrac{\lambda_1}{N_1} + \dfrac{\lambda_2}{N_2}}} - z_{1-\alpha}\right]$$

where $z_x = \Phi(x)$ is the standard normal distribution function.

The second test is the *square root transform* version of the first. The square root transform helps the test statistic approach normality for smaller sample sizes. It is given by

$$z_{SR} = \frac{\sqrt{\hat{\lambda}_2} - \sqrt{\hat{\lambda}_1}}{\dfrac{1}{2}\sqrt{\dfrac{1}{N_1} + \dfrac{1}{N_2}}}$$

The power of the one-sided level α hypothesis test is given by

$$Power = \Phi\left[\frac{\sqrt{\lambda_2} - \sqrt{\lambda_1}}{\frac{1}{2}\sqrt{\frac{1}{N_1} + \frac{1}{N_2}}} - z_{1-\alpha}\right]$$

Assurance Calculation

This assurance computation described here is based on O'Hagan, Stevens, and Campbell (2005).

Let $P'(H|\lambda_1, \lambda_2)$ be the power function described above where H is the event that null hypothesis is rejected conditional on the parameter values. The specification of λ_1 and λ_2 is critical to the power calculation, but the actual values are seldom known. Assurance is defined as the expected power where the expectation is with respect to a joint prior distribution for the parameters λ_1 and λ_2. Hence, the definition of assurance is

$$Assurance = E_{\lambda_1, \lambda_2}\left(P'(H|\lambda_1, \lambda_2)\right) = \iint P'(H|\lambda_1, \lambda_2)f(\lambda_1, \lambda_2)d\lambda_1\, d\lambda_2$$

where $f(\lambda_1, \lambda_2)$ is the joint prior distribution of λ_1 and λ_2.

In **PASS**, the joint prior distribution can be specified as either a discrete approximation to the joint prior distribution, or as individual prior distributions, one for each parameter.

Specifying a Joint Prior Distribution

If the joint prior distribution is to be specified directly, the distribution is specified in **PASS** using a discrete approximation to the function $f(\lambda_1, \lambda_2)$. This provides flexibility in specifying the joint prior distribution. In the two-parameter case, three columns are entered on the spreadsheet: two for the parameters and a third for the probability. Each row gives a value for each parameter and the corresponding parameter-combination probability. The accuracy of the distribution approximation is controlled by the number of points (spreadsheet rows) that are used.

An example of entering a joint prior distribution is included at the end of the chapter.

Specifying Individual Prior Distributions

Ciarleglio, Arendt, and Peduzzi (2016) suggest that more flexibility is available if the joint prior distribution is separated into two independent univariate distributions as follows

$$f(\lambda_1, \lambda_2) = f_1(\lambda_1)f_2(\lambda_2)$$

where $f_1(\lambda_1)$ is the prior distribution of λ_1 and $f_2(\lambda_2)$ is the prior distribution of λ_2. This method is also available in **PASS**. In this case, the definition of assurance becomes

$$Assurance = E_{\lambda_1, \lambda_2}\left(P'(H|\lambda_1, \lambda_2)\right) = \iint P'(H|\lambda_1, \lambda_2)f_1(\lambda_1)f_2(\lambda_2)d\lambda_1\, d\lambda_2$$

Using this definition, the assurance can be calculated using numerical integration. There are a variety of pre-programmed, univariate prior distributions available in **PASS**.

Fixed Values (No Prior) and Custom Values

For any given parameter, **PASS** also provides the option of entering a single fixed value for the prior distribution, or a series of values and corresponding probabilities (using the spreadsheet), rather than one of the pre-programmed distributions.

Numerical Integration in PASS (and Notes on Computation Speed)

When the prior distribution is specified as independent univariate distributions, **PASS** uses a numerical integration algorithm to compute the assurance value as follows:

The domain of each prior distribution is divided into M intervals. Since many of the available prior distributions are unbounded on one (e.g., Gamma) or both (e.g., Normal) ends, an approximation is made to make the domain finite. This is accomplished by truncating the distribution to a domain between the two quantiles: $q_{0.001}$ and $q_{0.999}$.

The value of M controls the accuracy and speed of the algorithm. If only one parameter is to be given a prior distribution, then a value of M between 50 and 100 usually gives an accurate result in a timely manner. However, if two parameters are given priors, the number of iterations needed increases from M to M^2. For example, if M is 100, 10000 iterations are needed. Reducing M from 100 to 50 reduces the number of iterations from 10000 to 2500.

The algorithm runtime increases when searching for sample size rather than solving for assurance, as a search algorithm is employed in this case. When solving for sample size, we recommend reducing M to 20 or less while exploring various scenarios, and then increasing M to 50 or more for a final, more accurate, result.

List of Available Univariate Prior Distributions

This section lists the univariate prior distributions that may be used for any of the applicable parameters when the Prior Entry Method is set to Individual.

No Prior

If 'No Prior' is chosen for a parameter, the parameter is assumed to take on a single, fixed value with probability one.

Beta (Shape 1, Shape 2, a, c)

A random variable *X* that follows the beta distribution is defined on a finite interval [a, c]. Two shape parameters (*α* and *β*) control the shape of this distribution. Two location parameters *a* and *c* give the minimum and maximum of *X*.

The probability density function of the beta distribution is

$$f(x|\alpha, \beta, a, c) = \frac{\left(\frac{x - a}{c - a}\right)^{\alpha-1} \left(\frac{c - x}{c - a}\right)^{\beta-1}}{(c - a)B(\alpha, \beta)}$$

where $B(\alpha, \beta) = \Gamma(\alpha)\,\Gamma(\beta)/\,\Gamma(\alpha + \beta)$ and $\Gamma(z)$ is the gamma function.

The mean of *X* is

$$\mu_X = \frac{\alpha c + \beta a}{\alpha + \beta}$$

Various distribution shapes are controlled by the values of *α* and *β*. These include

Symmetric and Unimodal

$\alpha = \beta > 1$

U Shaped

$\alpha = \beta < 1$

Bimodal

$\alpha, \beta < 1$

Uniform

$\alpha = \beta = 1$

Parabolic

$\alpha = \beta = 2$

Bell-Shaped

$\alpha = \beta > 2$

Gamma (Shape, Scale)

A random variable X that follows the gamma distribution is defined on the interval $(0, \infty)$. A shape parameter, κ, and a scale parameter, θ, control the distribution.

The probability density function of the gamma distribution is

$$f(x|\kappa, \theta) = \frac{x^{\kappa-1} e^{-\frac{x}{\theta}}}{\theta^{\kappa} \Gamma(\kappa)}$$

where $\Gamma(z)$ is the gamma function.

The mean of X is

$$\mu_X = \frac{\kappa}{\theta}$$

A truncated version of the distribution is constructed by dividing the density by $1 - \text{Prob}(Min \leq X \leq Max)$ where Min and Max are two truncation bounds.

Inverse-Gamma (Shape, Scale)

A random variable X that follows the inverse-gamma distribution is defined on the interval $(0, \infty)$. If $Y \sim$ gamma, then $X = 1 / Y \sim$ inverse-gamma. A shape parameter, α, and a scale parameter, β, control the distribution.

The probability density function of the inverse-gamma distribution is

$$f(x|\alpha, \beta) = \frac{\beta^{\alpha} x^{\alpha-1} e^{-\frac{\beta}{x}}}{\Gamma(\alpha)}$$

where $\Gamma(z)$ is the gamma function.

The mean of X is

$$\mu_X = \frac{\beta}{\alpha - 1} \text{ for } \alpha > 1$$

A truncated version of the distribution is constructed by dividing the density by $1 - \text{Prob}(Min \leq X \leq Max)$ where Min and Max are two truncation bounds.

Logistic (Location, Scale)

A random variable X that follows the logistic distribution is defined on the interval $(-\infty, \infty)$. A location parameter, μ, and a scale parameter, s, control the distribution.

The probability density function of the logistic distribution is

$$f(x|\mu, s) = \frac{e^{-\frac{x-\mu}{s}}}{s\left(1 + e^{-\frac{x-\mu}{s}}\right)^2}$$

The mean of X is

$$\mu_X = \mu$$

A truncated version of the distribution is constructed by dividing the density by $1 - \text{Prob}(Min \leq X \leq Max)$ where *Min* and *Max* are two truncation bounds.

Lognormal (Mean, SD)

A random variable X that follows the lognormal distribution is defined on the interval $(0, \infty)$. A location parameter, $\mu_{\log(X)}$, and a scale parameter, $\sigma_{\log(X)}$, control the distribution. If Z ~ standard normal, then $X = e^{\mu + \sigma Z} \sim$ lognormal. Note that $\mu_{\log(X)} = E(\log(X))$ and $\sigma_{\log(X)} = Standard\ Deviation(\log(X))$.

The probability density function of the lognormal distribution is

$$f(x|\mu, \sigma) = \frac{e^{-\frac{1}{2}\left(\frac{\log x - \mu}{\sigma}\right)^2}}{x\sigma\sqrt{2\pi}}$$

The mean of X is

$$\mu_X = e^{\mu + \frac{\sigma^2}{2}}$$

A truncated version of the distribution is constructed by dividing the density by $1 - \text{Prob}(Min \leq X \leq Max)$ where *Min* and *Max* are two truncation bounds.

LogT (Mean, SD)

A random variable X that follows the logT distribution is defined on the interval $(0, \infty)$. A location parameter, $\mu_{\log(X)}$, a scale parameter, $\sigma_{\log(X)}$, and a shape parameter, ν, control the distribution. Note that ν is referred to as the *degrees of freedom*.

If t ~ Student's t, then $X = e^{\mu + \sigma t} \sim$ logT.

The probability density function of the logT distribution is

$$f(x|\mu, \sigma, \nu) = \frac{\Gamma\left(\frac{\nu + 1}{2}\right)}{x\Gamma\left(\frac{\nu}{2}\right)\sigma\sqrt{\nu\pi}}\left(1 + \frac{1}{\nu}\left(\frac{\log x - \mu}{\sigma}\right)^2\right)^{\left(\frac{-\nu - 1}{2}\right)}$$

The mean of X is not defined.

A truncated version of the distribution is constructed by dividing the density by $1 - \text{Prob}(Min \leq X \leq Max)$ where *Min* and *Max* are two truncation bounds.

Normal (Mean, SD)

A random variable X that follows the normal distribution is defined on the interval $(-\infty, \infty)$. A location parameter, μ, and a scale parameter, σ, control the distribution.

The probability density function of the normal distribution is

$$f(x|\mu, \sigma) = \frac{e^{-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2}}{\sigma\sqrt{2\pi}}$$

The mean of X is

$$\mu_X = \mu$$

A truncated version of the distribution is constructed by dividing the density by $1 - \text{Prob}(Min \leq X \leq Max)$ where *Min* and *Max* are two truncation bounds.

T (Mean, SD, DF)

A random variable X that follows Student's t distribution is defined on the interval $(-\infty, \infty)$. A location parameter, μ, a scale parameter, σ, and a shape parameter, ν, control the distribution. Note that ν is referred to as the *degrees of freedom* or *DF*.

The probability density function of the Student's t distribution is

$$f(x|\mu, \sigma, \nu) = \frac{\Gamma\left(\frac{\nu + 1}{2}\right)}{\Gamma\left(\frac{\nu}{2}\right)\sigma\sqrt{\nu\pi}}\left(1 + \frac{1}{\nu}\left(\frac{x - \mu}{\sigma}\right)^2\right)^{\left(\frac{-\nu - 1}{2}\right)}$$

The mean of X is μ if $\nu > 1$.

A truncated version of the distribution is constructed by dividing the density by $1 - \text{Prob}(Min \leq X \leq Max)$ where *Min* and *Max* are two truncation bounds.

Triangle (Mode, Min, Max)

Let a = minimum, b = maximum, and c = mode. A random variable X that follows a triangle distribution is defined on the interval (a, b).

The probability density function of the triangle distribution is

$$f(x|a, b, c) = \begin{cases} \dfrac{2(x-a)}{(b-a)(c-a)} & \text{for } a \leq x < c \\[2mm] \dfrac{2}{b-a} & \text{for } x = c \\[2mm] \dfrac{2(b-x)}{(b-a)(b-c)} & \text{for } c < x \leq b \end{cases}$$

The mean of X is

$$\frac{a+b+c}{3}$$

Uniform (Min, Max)

Let a = minimum and b = maximum. A random variable X that follows a uniform distribution is defined on the interval $[a, b]$.

The probability density function of the uniform distribution is

$$f(x|a, b) = \begin{cases} \dfrac{1}{b-a} & \text{for } a \leq x \leq b \end{cases}$$

The mean of X is

$$\frac{a+b}{2}$$

Weibull (Shape, Scale)

A random variable X that follows the Weibull distribution is defined on the interval $(0, \infty)$. A shape parameter, κ, and a scale parameter, λ, control the distribution.

The probability density function of the Weibull distribution is

$$f(x|\kappa, \lambda) = \frac{\kappa}{\lambda} \left(\frac{x}{\lambda}\right)^{\kappa-1} e^{-\left(\frac{x}{\lambda}\right)^{\kappa}}$$

The mean of X is

$$\mu_X = \kappa\Gamma\left(1 + \frac{1}{\kappa}\right)$$

A truncated version of the distribution is constructed by dividing the density by $1 - \text{Prob}(Min \leq X \leq Max)$ where Min and Max are two truncation bounds.

Custom (Values and Probabilities in Spreadsheet)

This custom prior distribution is represented by a set of user-specified points and associated probabilities, entered in two columns of the spreadsheet. The points make up the entire set of values that are used for this parameter in the calculation of assurance. The associated probabilities should sum to one. Note that custom values and probabilities can be used to approximate any continuous distribution.

For example, a prior distribution of X might be

X_i	P_i
10	0.2
20	0.2
30	0.3
40	0.2
50	0.1

In this example, the mean of X is

$$\mu_X = \sum_{i=1}^{5} X_i P_i$$

Example 1 – Assurance Over a Range of Sample Sizes

Researchers wish to compare two drugs to determine whether there is a meaningful difference in the event rate. A sample size analysis determined that 578 subjects per group were needed to achieve 90% power at a 0.05 level of significance when the event rates of the control and treatment groups were 1.0 and 1.2, respectively. (See Example 1 of the *Tests for the Difference Between Two Poisson Rates* procedure).

To complete their sample size study, the researchers want to run an assurance analysis for a range of group sample sizes from 300 to 600. An elicitation exercise determined that the prior distribution of the $\lambda 1$ should be normal with mean 1.0 and standard deviation 0.03. The elicitation also concluded that prior distribution of $\lambda 2$ should be normal with mean 1.2 and standard deviation 0.05.

Setup

If the procedure window is not already open, use the PASS Home window to open it. The parameters for this example are listed below and are stored in the **Example 1** settings file. To load these settings to the procedure window, click **Open Example Settings File** in the Help Center or File menu.

Design Tab

Solve For ....**Assurance**
Prior Entry Method.......................................**Individual (Enter a prior distribution for each applicable parameter)**
Alternative Hypothesis ....**Two-Sided (H1: λ1 ≠ λ2)**
Test Statistic ....**Large Samples**
Alpha...**0.05**
Group Allocation ....**Equal (N1 = N2)**
Sample Size Per Group **300 400 500 600**
Prior Distribution of λ1..................................**Normal (Mean, SD)**
Mean...**1**
SD..**0.03**
Truncation Boundaries..................................**None**
Prior Distribution of λ2..................................**Normal (Mean, SD)**
Mean...**1.2**
SD..**0.05**
Truncation Boundaries..................................**None**

Options Tab

Number of Computation Points for each........**50**
  Prior Distribution
Maximum N1 in Sample Size Search **5000**

Output

Click the Calculate button to perform the calculations and generate the following output.

Numeric Reports

Numeric Results
_____

Solve For: Assurance
Hypotheses: H0: λ1 = λ2   vs.   H1: λ1 ≠ λ2
Test Statistic: Large-Sample
Prior Type: Independent Univariate Distributions

Prior Distributions
λ1: Normal (Mean = 1, SD = 0.03).
λ2: Normal (Mean = 1.2, SD = 0.05).
_____

| | | Sample Size | | | Expected Event Rate Group 1 | Expected Event Rate Group 2 | |
Assurance*	Power‡	N1	N2	N	E(λ1)	E(λ2)	Alpha
0.62222	0.64638	300	300	600	1	1.2	0.05
0.72002	0.76939	400	400	800	1	1.2	0.05
0.78781	0.85432	500	500	1000	1	1.2	0.05
0.83552	0.91035	600	600	1200	1	1.2	0.05
_____

* The number of points used for computation of the prior(s) was 50.
‡ Power was calculated using λ1 = E(λ1) = 1 and λ2 = E(λ2) = 1.2.

Assurance   The expected power where the expectation is with respect to the prior distribution(s).
Power   The power calculated using the parameter values shown in the footnote. Note that these parameter values may
                    be different from those shown in the report.
N1   The number of subjects in group 1.
N2   The number of subjects in group 2.
N   The total sample size. N = N1 + N2.
E(λ1)   The expected value over its prior distribution of the group 1 mean event rate.
E(λ2)   The expected value over its prior distribution of the group 2 mean event rate.
Alpha   The probability of rejecting a true null hypothesis.

Summary Statements
_____

A parallel two-group design will be used to test whether the Group 2 (treatment) Poisson event rate is different from
the Group 1 (control) Poisson event rate (H0: λ1 = λ2 versus H1: λ1 ≠ λ2). The comparison will be made using a
two-sided, large-sample Z-test of the rate difference, with a Type I error rate (α) of 0.05. The prior distribution used
for the Group 1 event rate is Normal (Mean = 1, SD = 0.03). The prior distribution used for the Group 2 event rate
is Normal (Mean = 1.2, SD = 0.05). With sample sizes of 300 for Group 1 (control) and 300 for Group 2 (treatment),
the assurance (average power) is 0.62222.
_____

Dropout-Inflated Sample Size

	Sample Size			Dropout-Inflated Enrollment Sample Size			Expected Number of Dropouts		
Dropout Rate	N1	N2	N	N1'	N2'	N'	D1	D2	D
20%	300	300	600	375	375	750	75	75	150
20%	400	400	800	500	500	1000	100	100	200
20%	500	500	1000	625	625	1250	125	125	250
20%	600	600	1200	750	750	1500	150	150	300

Dropout Rate	The percentage of subjects (or items) that are expected to be lost at random during the course of the study and for whom no response data will be collected (i.e., will be treated as "missing"). Abbreviated as DR.
N1, N2, and N	The evaluable sample sizes at which power is computed (as entered by the user). If N1 and N2 subjects are evaluated out of the N1' and N2' subjects that are enrolled in the study, the design will achieve the stated power.
N1', N2', and N'	The number of subjects that should be enrolled in the study in order to obtain N1, N2, and N evaluable subjects, based on the assumed dropout rate. N1' and N2' are calculated by inflating N1 and N2 using the formulas N1' = N1 / (1 - DR) and N2' = N2 / (1 - DR), with N1' and N2' always rounded up. (See Julious, S.A. (2010) pages 52-53, or Chow, S.C., Shao, J., Wang, H., and Lokhnygina, Y. (2018) pages 32-33.)
D1, D2, and D	The expected number of dropouts. D1 = N1' - N1, D2 = N2' - N2, and D = D1 + D2.

Dropout Summary Statements

Anticipating a 20% dropout rate, 375 subjects should be enrolled in Group 1, and 375 in Group 2, to obtain final group sample sizes of 300 and 300, respectively.

References

O'Hagan, A., Stevens, J.W., and Campbell, M.J. 2005. 'Assurance in clinical trial design'. Pharmaceutical Statistics, Volume 4, Pages 187-201.

Ciarleglio, M.M., Arendt, C.D., and Peduzzi, P.N. 2016. 'Selection of the effect size for sample size determination for a continuous response in a superiority clinical trial using a hybrid classical and Bayesian procedure'. Clinical Trials, Volume 13(3), pages 275-285.

Dias, L.C., Morton, A., and Quigley, J. 2018. Elicitation, The Science and Art of Structuring Judgement. Springer.

Mathews, Paul. 2010. Sample Size Calculations: Practical Methods for Engineers and Scientists. Mathews Malnar and Bailey. Fairport Harbor, OH. www.mmbstatical.com

Smith, P.G. and Morrow, R.H. 1996. Field Trials of Health Intervention in Developing Countries: A Toolbox. Macmillan Education. Oxford, England.

Campbell, M.J. and Walters, S.J. 2014. How to Design, Analyse and Report Cluster Randomised Trials in Medicine and Health Related Research. John Wiley. New York.

This report shows the assurance values obtained by the various sample sizes.

Plots Section

Plots
_____
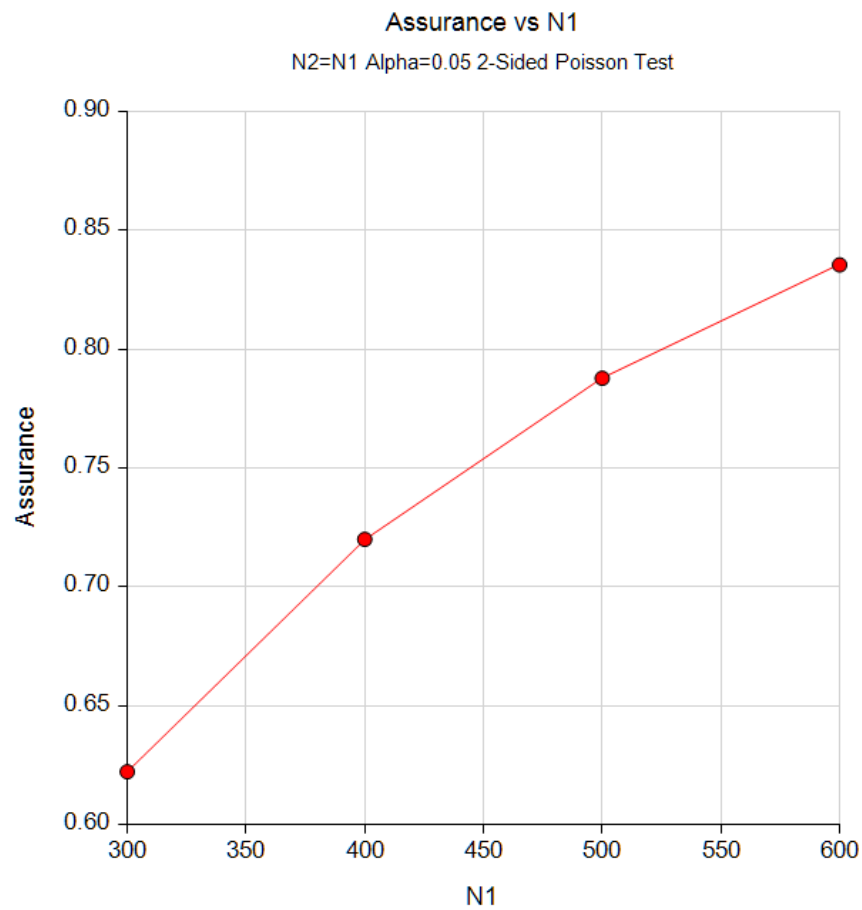
This plot shows the relationship between the assurance and sample size. Note the diminishing impact on assurance of each increase in the number of subjects.

Comparison Plots Section

Comparison Plots
_____

This plot compares the assurance and power across values of sample size. Note that assurance does not increase nearly as fast as does power.

Example 2 – Validation using Hand Computation

We could not find a validation example in the literature, so we have developed a validation example of our own.

Suppose a one-sided, large-sample z-test is used in which N1 = N2 = 500 and the significance level is 0.025.
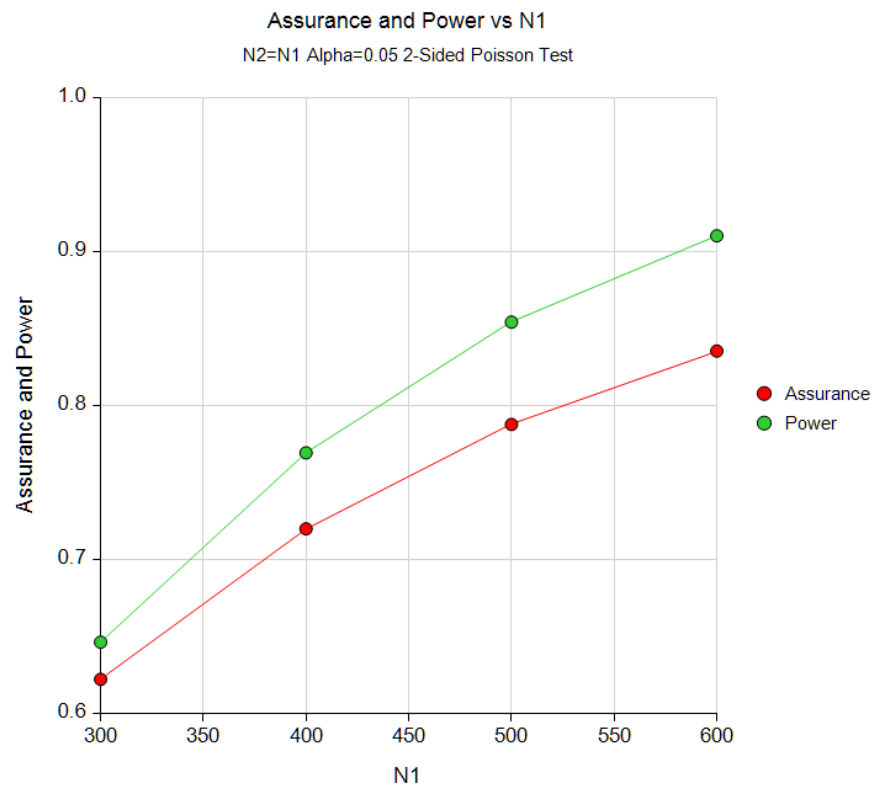
The prior distribution of λ1 is approximated by the following table.

λ1	Prob
0.98	0.3
1.00	0.4
1.02	0.3

The prior distribution of λ2 is approximated by the following table.

λ2	Prob
1.12	0.2
1.20	0.6
1.28	0.2

The *Tests for the Difference Between Two Poisson Rates* procedure is used to compute the power for each of the nine combinations of λ1 and λ2. The results of these calculations are shown next.

Numeric Results

Solve For: Power
Alternative Hypothesis: One-Sided
Test Statistic: Large-Sample Z-Test
Groups: 1 = Control, 2 = Treatment

| | Sample Size | | | Event Rate | | | | |
| | | | | | | Difference | Ratio | |
Power	N1	N2	N	λ1	λ2	Diff	RR	Alpha
0.57937	500	500	1000	0.98	1.12	0.14	1.14286	0.025
0.91494	500	500	1000	0.98	1.20	0.22	1.22449	0.025
0.99383	500	500	1000	0.98	1.28	0.30	1.30612	0.025
0.45340	500	500	1000	1.00	1.12	0.12	1.12000	0.025
0.85432	500	500	1000	1.00	1.20	0.20	1.20000	0.025
0.98561	500	500	1000	1.00	1.28	0.28	1.28000	0.025
0.33308	500	500	1000	1.02	1.12	0.10	1.09804	0.025
0.77077	500	500	1000	1.02	1.20	0.18	1.17647	0.025
0.96950	500	500	1000	1.02	1.28	0.26	1.25490	0.025

The assurance calculation is made by summing the quantities $\left[(power_{i,j})(p(\lambda 1_i))(p(\lambda 2_j))\right]$ as follows

$$Assurance = (0.57937 \times 0.3 \times 0.2) + (0.91494 \times 0.3 \times 0.6) + \cdots + (0.96950 \times 0.3 \times 0.2) = 0.79613.$$

To run this example, the spreadsheet will need to be loaded with the following four columns in which the first two are for λ1 and the second two are for λ2.

C1	C2	C3	C4
0.98	0.3	1.12	0.2
1.00	0.4	1.20	0.6
1.02	0.3	1.28	0.2

Setup

If the procedure window is not already open, use the PASS Home window to open it. The parameters for this example are listed below and are stored in the **Example 2** settings file. To load these settings to the procedure window, click **Open Example Settings File** in the Help Center or File menu.

Design Tab

Solve For ....**Assurance**
Prior Entry Method...**Individual (Enter a prior distribution for each applicable parameter)**
Alternative Hypothesis ....**One-Sided (H1: λ1 < λ2)**
Test Statistic ....**Large Samples**
Alpha...**0.025**
Group Allocation ....**Equal (N1 = N2)**
Sample Size Per Group **500**
Prior Distribution of λ1...................................**Custom (Values and Probabilities in Spreadsheet)**
Column of Values **C1**
Column of Pr(Values).....................................**C2**
Prior Distribution of λ2...................................**Custom (Values and Probabilities in Spreadsheet)**
Column of Values **C3**
Column of Pr(Values).....................................**C4**

Options Tab

Number of Computation Points for each........**50**
  Prior Distribution
Maximum N1 in Sample Size Search **5000**

Input Spreadsheet Data

Row	C1	C2	C3	C4
1	0.98	0.3	1.12	0.2
2	1.00	0.4	1.20	0.6
3	1.02	0.3	1.28	0.2

Output

Click the Calculate button to perform the calculations and generate the following output.

Numeric Results

Solve For: Assurance
Hypotheses: H0: λ1 ≥ λ2   vs.   H1: λ1 < λ2
Test Statistic: Large-Sample
Prior Type: Independent Univariate Distributions

Prior Distributions
λ1:   Point List (Values = C1, Probs = C2).
  C1:   0.98 1 1.02
  C2:   0.3 0.4 0.3
λ2:   Point List (Values = C3, Probs = C4).
  C3:   1.12 1.2 1.28
  C4:   0.2 0.6 0.2

		Sample Size			Expected Event Rate Group 1	Expected Event Rate Group 2	
Assurance	Power‡	N1	N2	N	E(λ1)	E(λ2)	Alpha
0.79613	0.85432	500	500	1000	1	1.2	0.025

‡ Power was calculated using λ1 = E(λ1) = 1 and λ2 = E(λ2) = 1.2.

PASS has also calculated the assurance as 0.79613 which validates the procedure.

Example 3 – Finding the Sample Size Needed to Achieve a Specified Assurance

Continuing with Example 1, the researchers want to investigate the sample sizes necessary to achieve assurances of 0.4, 0.5, 0.6, 0.7, and 0.8.

Setup

If the procedure window is not already open, use the PASS Home window to open it. The parameters for this example are listed below and are stored in the **Example 3** settings file. To load these settings to the procedure window, click **Open Example Settings File** in the Help Center or File menu.

Design Tab

Solve For ....**Sample Size**
Prior Entry Method.......................................**Individual (Enter a prior distribution for each applicable parameter)**
Alternative Hypothesis ....**Two-Sided (H1: λ1 ≠ λ2)**
Test Statistic ....**Large Samples**
Assurance...**0.4 0.5 0.6 0.7 0.8**
Alpha...**0.05**
Group Allocation ....**Equal (N1 = N2)**
Prior Distribution of λ1.................................**Normal (Mean, SD)**
Mean...**1**
SD..**0.03**
Truncation Boundaries.................................**None**
Prior Distribution of λ2.................................**Normal (Mean, SD)**
Mean...**1.2**
SD..**0.05**
Truncation Boundaries.................................**None**

Options Tab

Number of Computation Points for each........**50**
  Prior Distribution
Maximum N1 in Sample Size Search **5000**

Output

Click the Calculate button to perform the calculations and generate the following output.

Numeric Reports

Numeric Results
───
Solve For: Sample Size
Hypotheses: H0: λ1 = λ2   vs.   H1: λ1 ≠ λ2
Test Statistic: Large-Sample
Prior Type: Independent Univariate Distributions

Prior Distributions
λ1:   Normal (Mean = 1, SD = 0.03).
λ2:   Normal (Mean = 1.2, SD = 0.05).
───

| Assurance*			Sample Size			Expected Event Rate Group 1	Expected Event Rate Group 2	
Actual	Target	Power‡	N1	N2	N	E(λ1)	E(λ2)	Alpha
0.40149	0.4	0.39148	156	156	312	1	1.2	0.05
0.50067	0.5	0.50137	212	212	424	1	1.2	0.05
0.60043	0.6	0.61959	282	282	564	1	1.2	0.05
0.70062	0.7	0.74478	377	377	754	1	1.2	0.05
0.80028	0.8	0.86943	523	523	1046	1	1.2	0.05
───

*   The number of points used for computation of the prior(s) was 50.
‡   Power was calculated using λ1 = E(λ1) = 1 and λ2 = E(λ2) = 1.2.

This report shows the required sample size for each assurance target.

Plots Section

Plots
_____
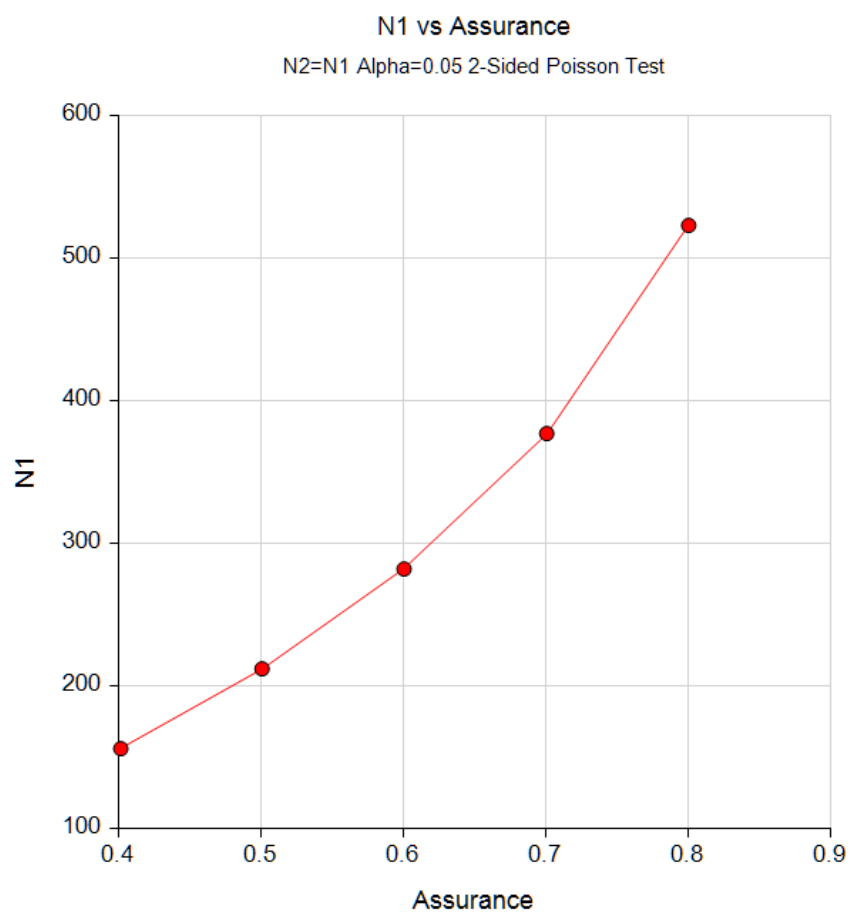
This plot shows the relationship between the sample size and assurance.

Example 4 – Joint Prior Distribution

The following example shows the complexity required to specify a joint distribution for two parameters.

Suppose a two-sided, unpooled-variance z-test will be used in which N1 = N2 = 2000 and the significance level is 0.05.

Further suppose that the joint prior distribution of the λ1 (control) and λ2 (treatment) is approximated by the following table. In a real study, the values in this table would be provided by an elicitation study.

Note that the program will rescale the probabilities so they sum to one.

λ1	λ2	Prob
0.32	0.34	0.05
0.36	0.34	0.10
0.44	0.34	0.25
0.34	0.35	0.20
0.37	0.35	0.25
0.45	0.35	0.40
0.34	0.36	0.50
0.38	0.36	0.55
0.46	0.36	0.70
0.35	0.37	0.50
0.39	0.37	0.55
0.47	0.37	0.70
0.36	0.38	0.20
0.40	0.38	0.25
0.48	0.38	0.40
0.37	0.39	0.05
0.41	0.39	0.10
0.49	0.39	0.25

To run this example, the spreadsheet will need to be loaded with the following three columns.

C1	C2	C3
0.32	0.34	0.05
0.36	0.34	0.10
0.44	0.34	0.25
0.34	0.35	0.20
0.37	0.35	0.25
0.45	0.35	0.40
0.34	0.36	0.50
0.38	0.36	0.55
0.46	0.36	0.70
0.35	0.37	0.50
0.39	0.37	0.55
0.47	0.37	0.70
0.36	0.38	0.20
0.40	0.38	0.25

0.48	0.38	0.40
0.37	0.39	0.05
0.41	0.39	0.10
0.49	0.39	0.25

Setup

If the procedure window is not already open, use the PASS Home window to open it. The parameters for this example are listed below and are stored in the **Example 4** settings file. To load these settings to the procedure window, click **Open Example Settings File** in the Help Center or File menu.

Design Tab

Solve For ....**Assurance**
Prior Entry Method.......................................**Combined (Enter parameter values and**
                                                                         probabilities on spreadsheet)
Alternative Hypothesis ....**Two-Sided (H1: λ1 ≠ λ2)**
Test Statistic ....**Large Samples**
Alpha..**0.05**
Group Allocation ....**Equal (N1 = N2)**
Sample Size Per Group **2000**
Column of λ1 Values.....................................**C1**
Column of λ2 Values.....................................**C2**
Column of Pr(Values)....................................**C3**

Options Tab

Number of Computation Points for each........**50**
  Prior Distribution
Maximum N1 in Sample Size Search **5000**

Input Spreadsheet Data

Row	C1	C2	C3
1	0.32	0.34	0.05
2	0.36	0.34	0.10
3	0.44	0.34	0.25
4	0.34	0.35	0.20
5	0.37	0.35	0.25
6	0.45	0.35	0.40
7	0.34	0.36	0.50
8	0.38	0.36	0.55
9	0.46	0.36	0.70
10	0.35	0.37	0.50
11	0.39	0.37	0.55
12	0.47	0.37	0.70
13	0.36	0.38	0.20
14	0.40	0.38	0.25
15	0.48	0.38	0.40
16	0.37	0.39	0.05
17	0.41	0.39	0.10
18	0.49	0.39	0.25

Output

Click the Calculate button to perform the calculations and generate the following output.

Numeric Results

Solve For: Assurance
Hypotheses: H0: λ1 = λ2   vs.   H1: λ1 ≠ λ2
Test Statistic: Large-Sample
Prior Type: Joint Multivariate Distribution

Prior Distribution
Point Lists
λ1:   C1:   0.32 0.36 0.44 0.34 0.37 0.45 0.34 0.38 0.46 0.35 0.39 0.47 0.36 0.4 0.48 0.37 0.41 0.49
λ2:   C2:   0.34 0.34 0.34 0.35 0.35 0.35 0.36 0.36 0.36 0.37 0.37 0.37 0.38 0.38 0.38 0.39 0.39 0.39
Prob: C3:   0.05 0.1 0.25 0.2 0.25 0.4 0.5 0.55 0.7 0.5 0.55 0.7 0.2 0.25 0.4 0.05 0.1 0.25

| | | | | | Event Rate | | |
| | | | | | Group 1 | Group 2 | |
Assurance	Power‡	N1	N2	N	E(λ1)	E(λ2)	Alpha
0.54566	0.65239	2000	2000	4000	0.41133	0.365	0.05

‡ Power was calculated using λ1 = E(λ1) = 0.41133 and λ2 = E(λ2) = 0.365.

PASS has calculated the assurance as 0.54566.

Example 5 – Joint Prior Validation

The problem given in Example 2 will be used to validate the joint prior distribution method. This will be done by running the independent-prior scenario used in that example through the joint-prior method and checking that the assurance values match.

In Example 2, the prior distributions of the λ1 and λ2 are

λ1	Prob
0.98	0.3
1.00	0.4
1.02	0.3

λ2	Prob
1.12	0.2
1.20	0.6
1.28	0.2

The joint prior distribution can be found by multiplying the three independent probabilities in each row. This results in the following discrete probability distribution.

λ1	λ2	P(λ1)	P(λ2)	Prob
0.98	1.12	0.3	0.2	0.06
0.98	1.20	0.3	0.6	0.18
0.98	1.28	0.3	0.2	0.06
1.00	1.12	0.4	0.2	0.08
1.00	1.20	0.4	0.6	0.24
1.00	1.28	0.4	0.2	0.08
1.02	1.12	0.3	0.2	0.06
1.02	1.20	0.3	0.6	0.18
1.02	1.28	0.3	0.2	0.06

To run this example, the spreadsheet is loaded with the following four columns.

C1	C2	C3	C4	C5
0.98	1.12	0.3	0.2	0.06
0.98	1.20	0.3	0.6	0.18
0.98	1.28	0.3	0.2	0.06
1.00	1.12	0.4	0.2	0.08
1.00	1.20	0.4	0.6	0.24
1.00	1.28	0.4	0.2	0.08
1.02	1.12	0.3	0.2	0.06
1.02	1.20	0.3	0.6	0.18
1.02	1.28	0.3	0.2	0.06

Setup

If the procedure window is not already open, use the PASS Home window to open it. The parameters for this example are listed below and are stored in the **Example 5** settings file. To load these settings to the procedure window, click **Open Example Settings File** in the Help Center or File menu.

Design Tab

Solve For ....**Assurance**
Prior Entry Method.......................................**Combined (Enter parameter values and**
                                                                                    probabilities on spreadsheet)
Alternative Hypothesis ....**One-Sided (H1: λ1 < λ2)**
Test Statistic ....**Large Samples**
Alpha..**0.025**
Group Allocation ....**Equal (N1 = N2)**
Sample Size Per Group **500**
Column of λ1 Values.....................................**C1**
Column of λ2 Values.....................................**C2**
Column of Pr(Values).....................................**C3**

Options Tab

Number of Computation Points for each........**50**
  Prior Distribution
Maximum N1 in Sample Size Search **5000**

Input Spreadsheet Data

Row	C1	C2	C3	C4	C5
1	0.98	1.12	0.3	0.2	0.06
2	0.98	1.20	0.3	0.6	0.18
3	0.98	1.28	0.3	0.2	0.06
4	1.00	1.12	0.4	0.2	0.08
5	1.00	1.20	0.4	0.6	0.24
6	1.00	1.28	0.4	0.2	0.08
7	1.02	1.12	0.3	0.2	0.06
8	1.02	1.20	0.3	0.6	0.18
9	1.02	1.28	0.3	0.2	0.06

Output

Click the Calculate button to perform the calculations and generate the following output.

Numeric Results

Solve For: Assurance
Hypotheses: H0: λ1 ≥ λ2   vs.   H1: λ1 < λ2
Test Statistic: Large-Sample
Prior Type: Joint Multivariate Distribution

Prior Distribution
Point Lists
λ1:   C1:   0.98 0.98 0.98 1 1 1 1.02 1.02 1.02
λ2:   C2:   1.12 1.2 1.28 1.12 1.2 1.28 1.12 1.2 1.28
Prob: C3:   0.3 0.3 0.3 0.4 0.4 0.4 0.3 0.3 0.3

		Sample Size			Expected Event Rate		
Assurance	Power‡	N1	N2	N	Group 1 E(λ1)	Group 2 E(λ2)	Alpha
0.76193	0.85432	500	500	1000	1	1.2	0.025

‡ Power was calculated using λ1 = E(λ1) = 1 and λ2 = E(λ2) = 1.2.

PASS has also calculated the assurance as 0.79613 which matches Example 2 and thus validates the procedure.