

Package ‘Exact’

October 22, 2016

Version 1.7

Date 2016-10-21

Title Unconditional Exact Test

Author Peter Calhoun [aut, cre]

Maintainer Peter Calhoun <calhoun.peter@gmail.com>

Description

Performs unconditional exact tests and power calculations for 2x2 contingency tables. Unconditional exact tests are often more powerful than conditional exact tests and asymptotic tests.

License GPL-2

Imports graphics, stats

NeedsCompilation no

Repository CRAN

Date/Publication 2016-10-22 12:17:16

R topics documented:

Exact	1
exact.test	3
power.exact.test	6

Index	9
--------------	----------

Exact	<i>Unconditional Exact Tests for 2x2 Tables</i>
-------	---

Description

This package implements the `exact.test` function to perform unconditional exact tests. This package also includes the `power.exact.test` function to calculate the power to detect a significant difference using unconditional exact tests.

Details

Package: Exact
Type: Package
Version: 1.7
Date: 2016-10-21
License: GPL-2

Unconditional exact tests can be performed to test the independence of rows and columns in a 2x2 table. Unconditional tests (such as Barnard's and Boschloo's exact tests) are more powerful alternatives than conditional tests (such as Fisher's exact test). P-values can be computed for 2x2 tables under the binomial and multinomial models using various test statistics to find the 'as or more extreme' tables. The multinomial model assumes only the total sample size is known in advance (common in cross-sectional studies), the binomial model assumes only the row or column margins (but not both) are fixed (most common situation; occurs in case-control studies), and the hypergeometric model assumes both row and column margins are fixed (very unlikely situation; only design where Fisher's test should be performed).

The details of the test statistics are given in the `exact.test` function description. Suissa and Shuster suggested using a Z-pooled statistic, which they found to be uniformly more powerful than Fisher's test for balanced designs. Boschloo recommended using the p-value for Fisher's test as the test statistic. This method became known as Boschloo's test, and it is always uniformly more powerful than Fisher's test. Mato and Andres suggested using Barnard's CSM test. Additionally, Berger and Boos proposed considering only values of the nuisance parameter that are in a constructed confidence interval. The interval approach often yields more powerful tests. While there is still a disagreement on which test statistic to use, most researchers agree that Fisher's exact test should **not** be used to analyze 2x2 tables. All of these tests can be computed in this package.

Note

Throughout the years I have received help while creating this package. Special thanks goes to Philo Calhoun, Tal Galili, Kamil Erguler, Roger Berger, Karl Hufthammer, and the R community.

Author(s)

Peter Calhoun [aut, cre]

Maintainer: Peter Calhoun <calhoun.peter@gmail.com>

References

- Barnard, G.A. (1945) A new test for 2x2 tables. *Nature*, **156**, 177
- Barnard, G.A. (1947) Significance tests for 2x2 tables. *Biometrika*, **34**, 123-138
- Berger, R. and Boos D. (1994) P values maximized over a confidence set for the nuisance parameter. *Journal of the American Statistical Association*, **89**, 1012-1016
- Berger, R. (1994) Power comparison of exact unconditional tests for comparing two binomial proportions. *Institute of Statistics Mimeo Series No. 2266*
- Berger, R. (1996) More powerful tests from confidence interval p values. *American Statistician*, **50**, 314-318

Boschloo, R. D. (1970), Raised Conditional Level of Significance for the 2x2-table when Testing the Equality of Two Probabilities. *Statistica Neerlandica*, **24**, 1-35

Cardillo, G. (2009) MyBarnard: a very compact routine for Barnard's exact test on 2x2 matrix. <http://www.mathworks.com/matlabcentral/fileexchange/25760-mybarnard>

Mato, S. and Andres, M. (1997), Simplifying the calculation of the P-value for Barnard's test and its derivatives. *Statistics and Computing*, **7**, 137-143

Mehrotra, D., Chan, I., Berger, R. (2003), A Cautionary Note on Exact Unconditional Inference for a Difference Between Two Independent Binomial Proportions. *Biometrics*, **59**, 441-450

Ruxton, G. D. and Neuhauser M (2010), Good practice in testing for an association in contingency tables. *Behavioral Ecology and Sociobiology*, **64**, 1505-1513

Suissa, S. and Shuster, J. J. (1985), Exact Unconditional Sample Sizes for the 2x2 Binomial Trial, *Journal of the Royal Statistical Society, Ser. A*, **148**, 317-327

exact.test

Unconditional exact tests for 2x2 tables

Description

Calculates Barnard's or Boschloo's unconditional exact test for binomial or multinomial models

Usage

```
exact.test(data, alternative = c("two.sided", "less", "greater"), npNumbers = 100,
           beta = 0.001, interval = FALSE,
           method = c("z-pooled", "z-unpooled", "boschloo", "santner and snell",
                     "csm", "csm approximate", "csm modified"),
           model = c("Binomial", "Multinomial"), cond.row = TRUE, to.plot = TRUE,
           ref.pvalue = TRUE)
```

Arguments

data	A two dimensional contingency table in matrix form
alternative	Indicates the alternative hypothesis: must be either "less", "two.sided", or "greater"
npNumbers	Number: The number of nuisance parameters considered
beta	Number: Confidence level for constructing the interval of nuisance parameters considered. Only used if interval=TRUE
interval	Logical: Indicates if a confidence interval on the nuisance parameter should be computed
method	Indicates the method for finding tables as or more extreme than the observed table: must be either "Z-pooled", "Z-unpooled", "Santner and Snell", "Boschloo", "CSM", "CSM approximate", or "CSM modified". CSM tests cannot be calculated for multinomial models
model	The model being used: must be either "Binomial" or "Multinomial"

cond.row	Logical: Indicates if row margins are fixed in the binomial models. Only used if model="Binomial"
to.plot	Logical: Indicates if plot of p-value vs. nuisance parameter should be generated. Only used if model="Binomial"
ref.pvalue	Logical: Indicates if p-value should be refined by maximizing the p-value function after the nuisance parameter is selected. Only used if model="Binomial"

Details

Unconditional exact tests can be used for binomial or multinomial models. The binomial model assumes the row or column margins (but not both) are known in advance, while the multinomial model assumes only the total sample size is known beforehand. Conditional tests have both row and column margins fixed. The null hypothesis is that the rows and columns are independent. Under the binomial model, the user will need to input which margin is fixed (default is rows).

See the following formulas in the Referene Manual: <https://CRAN.R-project.org/package=Exact>.

Let X denote a generic 2x2 table with fixed sample sizes n_1 and n_2 , X_0 denote the observed table, and $T(X)$ represent the test statistic function. The null hypothesis can be written as $p_1 = p_2 \equiv p$. The p-value function with rows fixed is the product of two independent binomials:

$$P(X|p) = \sup_{0 \leq p \leq 1} \sum_{T(X) \geq T(X_0)} \binom{n_1}{x_{11}} \binom{n_2}{x_{21}} p^{x_{11}+x_{21}} (1-p)^{x_{12}+x_{13}}$$

The multinomial model is similar except the summand has a multinomial distribution with two nuisance parameters.

There are several possible test statistics to determine the 'as or more extreme' tables seen in the index of summation. The method variable lets the user choose the test statistic being used. A brief description for each test statistic is given below (see References for more details):

Let $\hat{p}_1 = x_{11}/n_1$, $\hat{p}_2 = x_{21}/n_2$, and $\hat{p} = (x_{11} + x_{21})/(n_1 + n_2)$.

Z-unpooled (or Wald):

$$Z_u(x_{11}, x_{21}) = \frac{\hat{p}_2 - \hat{p}_1}{\sqrt{\frac{\hat{p}_1(1-\hat{p}_1)}{n_1} + \frac{\hat{p}_2(1-\hat{p}_2)}{n_2}}}$$

Z-pooled (or Score):

$$Z_p(x_{11}, x_{21}) = \frac{\hat{p}_2 - \hat{p}_1}{\sqrt{\frac{\hat{p}(1-\hat{p})}{n_1} + \frac{\hat{p}(1-\hat{p})}{n_2}}}$$

Santner and Snell:

$$D(x_{11}, x_{21}) = \hat{p}_2 - \hat{p}_1$$

Boschloo:

Uses the p-value from Fisher's exact test as the test statistic.

CSM:

Starts with the most extreme table and adds other 'as or more extreme' tables one step at a time by maximizing the summand of the p-value function. This approach can be computationally intensive.

CSM modified:

Starts with all tables that must be more extreme and adds other 'as or more extreme' tables one step at a time by maximizing the summand of the p-value function. This approach can be computationally intensive.

CSM approximate:

Maximizes the summand of the p-value function for each possible table. Thus, the test statistic is the p-value function without the summation. This approach is less computationally intensive than the CSM test because the maximization is not repeated at each step.

The supremum of the common success probability is taken over all values between 0 and 1. Another approach, proposed by Berger and Boos, is to take the supremum over a Clopper-Pearson confidence interval. This approach adds a small penalty to the p-value to ensure a level- α test, but eliminates unlikely probabilities from inflating the p-value. The p-value function becomes:

$$P(X|p) = \left(\sup_{p \in C_\beta} \sum_{T(X) \geq T(X_0)} \binom{n_1}{x_{11}} \binom{n_2}{x_{21}} p^{x_{11}+x_{21}} (1-p)^{x_{12}+x_{13}} \right) + \beta$$

where C_β is the $100(1 - \beta)\%$ confidence interval of p

There are many ways to define the two-sided p-value; this code uses the `fisher.test` approach by summing the probabilities for both sides of the table.

Value

A list with class "htest" containing the following components:

<code>p.value</code>	The computed p-value
<code>test.statistic</code>	The observed test statistic
<code>estimate</code>	An estimate of the parameter tested
<code>alternative</code>	A character string describing the alternative hypothesis
<code>model</code>	A character string describing the model design ("Binomial" or "Multinomial")
<code>method</code>	A character string describing the method to determine 'as or more extreme' tables
<code>np</code>	The nuisance parameter that maximizes the p-value. For multinomial models, both nuisance parameters are given
<code>np.range</code>	The range of nuisance parameters considered. For multinomial models, both nuisance parameter ranges are given
<code>data.name</code>	A character string giving the names of the data

Warning

Multinomial models and CSM tests may take a very long time, even for sample sizes less than 100.

Note

CSM test and multinomial models are much more computationally intensive. I have also spent a greater amount of time making the computations for the binomial models more efficient; future work will be devoted to improving the multinomial models. Boschloo's test also takes longer due to calculating Fisher's p-value for every possible table; however, a created function that calculates Fisher's test efficiently is utilized. Increasing the number of nuisance parameters considered and refining the p-value will increase the computation time.

Author(s)

Peter Calhoun

References

This code was influenced by the FORTRAN program located at <http://www4.stat.ncsu.edu/~boos/exact/>

See Also

fisher.test and **exact2x2**

Examples

```
data <- matrix(c(7, 8, 12, 3), 2, 2, byrow=TRUE)
exact.test(data, alternative="less", to.plot=TRUE)
exact.test(data, alternative="two.sided", interval=TRUE, beta=0.001, npNumbers=100,
            method="Z-pooled", to.plot=FALSE)
exact.test(data, alternative="two.sided", interval=TRUE, beta=0.001, npNumbers=100,
            method="Boschloo", to.plot=FALSE)

#Example from Barnard's (1947) appendix:
data <- matrix(c(4, 0, 3, 7), 2, 2,
              dimnames=list(c("Box 1", "Box 2"), c("Defective", "Not Defective")))
exact.test(data, method="CSM", alternative="two.sided")

data <- matrix(c(6, 8, 4, 3), 2, 2, byrow=TRUE)
exact.test(data, model="Multinomial", alternative="less", method="Z-pooled")
```

power.exact.test

Power calculation for unconditional exact test

Description

Calculates the power of the design for known sample sizes and true probabilities.

Usage

```
power.exact.test(p1, p2, n1, n2, npNumbers = 100, alpha = 0.05,
  alternative = c("two.sided", "less", "greater"),
  interval = FALSE, beta = 0.001,
  method = c("z-pooled", "z-unpooled", "boschloo", "santner and snell",
    "csm", "csm approximate", "csm modified", "fisher"),
  ref.pvalue = TRUE, simulation = FALSE, nsim = 100)
```

Arguments

p1	The probability of success given in first group
p2	The probability of success given in second group
n1	The sample size in first group
n2	The sample size in second group
npNumbers	Number: The number of nuisance parameters considered
alpha	Significance level
alternative	Indicates the alternative hypothesis: must be either "less", "two.sided", or "greater"
interval	Logical: Indicates if a confidence interval on the nuisance parameter should be computed
beta	Number: Confidence level for constructing the interval of nuisance parameters considered. Only used if interval=TRUE
method	Indicates the method for finding tables as or more extreme than the observed table: must be either "Z-pooled", "Z-unpooled", "Santner and Snell", "Boschloo", "CSM", "CSM modified", or "CSM approximate"
ref.pvalue	Logical: Indicates if p-value should be refined by maximizing the p-value function after the nuisance parameter is selected
simulation	Logical: Indicates if the power calculation is exact or estimated by simulation
nsim	Number of simulations run. Only used if simulation=TRUE

Details

The power calculations are for binomial models. The design must know the fixed sample sizes in advance. There are $(n_1 + 1) \times (n_2 + 1)$ possible tables that could be produced. There are two ways to calculate the power: simulate the tables under two independent binomial distributions or consider all possible tables and calculate the exact power. The calculations can be done for any `exact.test` computation or using Fisher's exact test.

Value

The function returns the computed power.

Note

The code takes a very long time for the CSM test. Not refining the p-value often yields similar results and decreases the computation time.

Author(s)

Peter Calhoun

References

Berger, R. (1994) Power comparison of exact unconditional tests for comparing two binomial proportions. *Institute of Statistics Mimeo Series No.* 2266

Berger, R. (1996) More powerful tests from confidence interval p values. *American Statistician*, **50**, 314-318

Boschloo, R. D. (1970), Raised Conditional Level of Significance for the 2x2-table when Testing the Equality of Two Probabilities. *Statistica Neerlandica*, **24**, 1-35

See Also

statmod

Examples

```
power.exact.test(0.20, 0.80, 10, 20)
power.exact.test(0.20, 0.80, 10, 20, method="Fisher")
set.seed(218461)
power.exact.test(0.20, 0.80, 10, 20, interval=TRUE, method="Boschloo",
                 simulation=TRUE, nsim=100)
```


Index

*Topic **Barnard**

exact.test, 3
power.exact.test, 6

*Topic **Boschloo**

exact.test, 3
power.exact.test, 6

*Topic **Exact Tests**

Exact, 1
exact.test, 3
power.exact.test, 6

*Topic **Power**

power.exact.test, 6

*Topic **Unconditional**

exact.test, 3
power.exact.test, 6

binom.CI (Exact), 1
binomialCode (exact.test), 3
boschloo_TX (Exact), 1

csm_Tbls (Exact), 1
csmApprox_TX (Exact), 1
csmMod_Tbls (Exact), 1

Exact, 1
exact.test, 1, 2, 3, 7

fisher.2x2 (Exact), 1

maxPvalue (Exact), 1
multinomialCode (exact.test), 3

power.exact.test, 1, 6

santner_TX (Exact), 1

zpooled_TX (Exact), 1
zunpooled_TX (Exact), 1