Package 'chisquare'

January 8, 2024

Title Chi-Square and G-Square Test of Independence, Power and Residual
Analysis, Measures of Categorical Association

Version 0.9

Description Provides the facility to perform the chi-square and G-square test of independence, calculates the power of the traditional chi-square test, compute permutation and Monte Carlo p-value, and provides measures of association such as Phi, odds ratio with 95 percent CI and p-value, adjusted contingency coefficient, Cramer's V and 95 percent CI, bias-corrected Cramer's V, W, Cohen's w, Goodman-Kruskal's lambda, gamma and its p-value, and tau, Cohen's k and its 95 percent CI. It also calculates standardized, moment-corrected standardized, and adjusted standardized residuals, and their significance. Different outputs are returned in nicely formatted tables.

Depends R (>= 4.0.0)

Imports graphics (>= 4.2.0), gt (>= 0.3.1), stats (>= 4.2.0)

License GPL (>= 2)

Encoding UTF-8

LazyData true

RoxygenNote 7.2.3

NeedsCompilation no

Author Gianmarco Alberti [aut, cre]

Maintainer Gianmarco Alberti <gianmarcoalberti@gmail.com>

Repository CRAN

Date/Publication 2024-01-08 15:00:02 UTC

R topics documented:

chisquare	*R function for Chi-square, (N-1) Chi-square, and G-Square test of independence, power calculation, measures of association, and standardized/moment-corrected standardized/adjusted standardized residuals, visualisation of odds ratio in 2xk tables (where k >= 2)*

Description

The function performs the chi-square test (both in its original format and in the N-1 version) and the G-square test of independence on the input contingency table. It also calculates the power of the traditional chi-square test and various measures of categorical association, returns standardized, moment-corrected standardized, and adjusted standardized residuals (with indication of their significance), and calculates relative and absolute contributions to the chi-square. The p value associated to the chi-square statistic is also calculated via both a permutation- and a Monte Carlo-based method. The 95 percent confidence interval around those p values is also calculated. Nicely-formatted output tables are rendered. Optionally, in 2xk tables (where k >= 2), a plot of the odds ratios can be rendered.

Visit this LINK to access the package's vignette.

Usage

```
chisquare(
  data,
  B = 1000,
  plot.or = FALSE,
  reference.level = 1,
  row.level = 1,
  or.alpha = 0.05,
  power.alpha = 0.05,
  adj.alpha = FALSE,
  format = "short",
  graph = FALSE,
  oneplot = TRUE,
  tfs = 13
)
```

Arguments

data	Dataframe containing the input contingency table.
B	Number of simulated tables to be used to calculate the permutation- and the Monte Carlo-based p value (1000 by default).
plot.or	Takes TRUE or FALSE (default) if the user wants a plot of the odds ratios to be rendered (only for 2xk tables, where k >= 2).

reference.level

  The index of the column reference level for odds ratio calculations (default: 1). The user must select the column level to serve as the reference level (only for 2xk tables, where k >= 2).

row.level The index of the row category to be used in odds ratio calculations (1 or 2; default: 1). The user must select the row level to which the calculation of the odds ratios make reference (only for 2xk tables, where k >= 2).

or.alpha The significance level used for the odds ratios' confidence intervals (default: 0.05).

power.alpha The significance level used for the calculation of the power of the traditional chi-square test (default: 0.05).

adj.alpha Takes TRUE or FALSE (default) if the user wants or does not want the significance level of the residuals (standardised, adjusted standardised, and moment-corrected) to be corrected using the Sidak's adjustment method (see Details).

format Takes *short* (default) if the dataset is a dataframe storing a contingency table; if the input dataset is a dataframe storing two columns that list the levels of the two categorical variables, *long* will preliminarily cross-tabulate the levels of the categorical variable in the 1st column against the levels of the variable stored in the 2nd column.

graph Takes TRUE or FALSE (default) if the user wants or does not want to plot the permutation and Monte Carlo distribution of the chi-square statistic accross the number of simulated tables set by the B parameter.

oneplot Takes TRUE (default) or FALSE if the user wants or does not want to render of the permutation and Monte Carlo distribution in the same plot.

tfs Numerical value to set the size of the font used in the main body of the various output tables (13 by default).

Details

The function produces the following **measures of categorical associations**:

- Phi (with indication of the magnitude of the effect size; only for 2x2 tables)

- Phi corrected (with indication of the magnitude of the effect size; only for 2x2 tables)

- Phi signed (with indication of the magnitude of the effect size; only for 2x2 tables)

- Yule's Q (only for 2x2 tables, includes p-value)

- Odds ratio (only for 2x2 tables, includes 95perc confidence interval, p value, and indication of the magnitude of the effect size)

- Adjusted contingency coefficient C (with indication of the magnitude of the effect size)

- Cramer's V (with 95perc confidence interval; includes indication of the magnitude of the effect size)

- Bias-corrected Cramer's V (with indication of the magnitude of the effect size)

- Cohen's w (with indication of the magnitude of the effect size)

- W coefficient (includes 95perc confidence interval and magnitude of the effect size)

- Goodman-Kruskal's lambda (both asymmetric and symmetric)

- Corrected version of lambda (both asymmetric and symmetric)

- Goodman-Kruskal's tau (asymmetric) and gamma (with p-value)

- Cohen's k (with 95perc confidence interval)

Indication of the magnitude of the association as indicated by the coefficients
The function provides indication of the mangitude of the association (effect size) for the Phi, Phi corrected, Phi signed, Cadj, Cramer's V, Cramer's V bias-corrected, Cohen's w, W, and for the Odds Ratio.

With the exception of the latter (for which see further down), the effect size for the other measures of association is based on Cohen 1988.

Phi, Phi corrected, Phi signed, and w are assessed against the well-known Cohen's classification scheme's thresholds (small 0.1, medium 0.3, large 0.5). For input cross-tabs larger than 2x2, the Cadj, V, V bias-corrected, and W coefficients are assessed against thresholds that depend on the table's df, which (as per Cohen 1988) correspond to the smaller between the rows and columns number, minus 1. On the basis of the table's df, the three thresholds are calculated as follows:

small effect: 0.100 / sqrt(min(nr,nc)-1)
medium effect: 0.300 / sqrt(min(nr,nc)-1)
large effect: 0.500 / sqrt(min(nr,nc)-1)

where nr and nc are the number of rows and number of columns respectively, and min(nr,nc)-1 corresponds to the table's df. Essentially, the thresholds for a small, medium, and large effect are computed by dividing the Cohen's thresholds for a 2x2 table (df=1) by the square root of the input table's df.

Consider a V value of (say) 0.35; its effect size interpretation changes based on the table's dimension:

for a 2x2 table, 0.35 corresponds to a "medium" effect;
for a 3x3 table, 0.35 still corresponds to a "medium" effect;
for a 4x4 table, 0.35 corresponds to a "large" effect.

The examples illustrate that for the same (say) V value, the interpreted effect size can shift from "medium" in a smaller table to "large" in a larger table. In simpler terms, the threshold for determining a "large" effect, for instance, becomes more accessible to reach as the table's size increases.

It is crucial to be aware of this as it highlights that the same coefficient value can imply different magnitudes of effect depending on the table's size

See: Cohen 1988; Sheskin 2011.

Power of the Traditional Chi-Square Test
The function calculates the power of the traditional chi-square test, which is the probability of correctly rejecting the null hypothesis when it is false. The power is determined by the observed

chi-square statistic, the sample size, and the degrees of freedom, without explicitly calculating an effect size, following the method described by Oyeyemi et al. 2010.

The degrees of freedom are calculated as (number of rows - 1) * (number of columns - 1). The alpha level is set by default at 0.05 and can be customized using the power.alpha parameter. The power is then estimated using the non-centrality parameter based on the observed chi-square statistic.

The calculation involves determining the critical chi-squared value based on the alpha level and degrees of freedom, and then computing the probability that the chi-squared distribution with the given degrees of freedom exceeds this critical value.

The resulting power value indicates how likely the test is to detect an effect if one exists. A power value close to 1 suggests a high probability of detecting a true effect, while a lower value indicates a higher risk of a Type II error. Typically, a power value of 0.8 or higher is considered robust in most research contexts.

Suggestion of a suitable chi-square testing method
The first rendered table includes a suggestion for the applicable chi-squared test method, derived from an internal analysis of the input contingency table. The decision logic used is as follows:

For 2x2 Tables:
- if the grand total is equal to or larger than 5 times the number of cells, the traditional Chi-Square test is suggested. Permutation or Monte Carlo methods can also be considered.

- if the grand total is smaller than 5 times the number of cells, the minimum expected count is checked:
(A) if it is equal to or larger than 1, the (N-1)/N adjusted Chi-Square test is suggested, with an option for Permutation or Monte Carlo methods.
(B) if it is less than 1, the Permutation or Monte Carlo method is recommended.

For Larger than 2x2 Tables:
- the logic is similar to that for 2x2 tables, with the same criteria for suggesting the traditional Chi-Square test, the (N-1)/N adjusted test, or the Permutation or Monte Carlo methods.

The rationale of a threshold for the applicability of the traditional chi-square test corresponding to 5 times the number of cells is based on the following.

Literature indicates that the traditional chi-squared test's validity is not as fragile as once thought, especially when considering the average expected frequency across all cells in the cross-tab, rather than the minimum expected value in any single cell. An average expected frequency of at least 5 across all cells of the input table should be sufficient for maintaining the chi-square test's reliability at the 0.05 significance level.

As a consequence, a table's grand total equal to or larger than 5 times the number of cells should ensure the applicability of the traditional chi-square test (at alpha 0.05).

See: Roscoe-Byars 1971; Greenwood-Nikulin 1996; Zar 2014.

For the rationale of the use of the (N-1)/N adjusted version of the chi-square test, and for the permutation and Monte Carlo method, see below.

Chi-square statistics adjusted using the (N-1)/N adjustment
The adjustment is done by multiplying the chi-square statistics by (N-1)/N, where N is the table grand total (sample size). The p-value of the corrected statistic is calculated the regular way (i.e., using the same degrees of freedom as in the traditional test). The correction seems particularly relevant for tables where N is smaller than 20 and where the expected frequencies are equal or larger than 1. The corrected chi-square test proves more conservative when the sample size is small. As N increases, the term (N-1)/N approaches 1, making the adjusted chi-square value virtually equivalent to the unadjusted value.

See: Upton 1982; Rhoades-Overall1982; Campbel 2007; Richardson 2011.

Permutation-based and Monte Carlo p-value for the chi-square statistic
The p-value of the observed chi-square statistic is also calculated on the basis of both a permutation-based and a Monte Carlo approach. In the first case, the dataset is permuted *B* times (1000 by default), whereas in the second method *B* establishes the number of random tables generated under the null hypothesis of independence (1000 by default).

As for the permutation method, the function does the following internally:
(1) Converts the input dataset to long format and expands to individual observations;
(2) Calculates the observed chi-squared statistic;
(3) Randomly shuffles (B times) the labels of the levels of one variable, and recalculates chi-squared statistic for each shuffled dataset; (4) Computes the p-value based on the distribution of permuted statistics (see below).

For the rationale of the permutation-based approach, see for instance Agresti et al 2022.

For the rationale of the Monte Carlo approach, see for instance the description in Beh-Lombardo 2014: 62-64.

Both simulated p-values are calculated as follows:

$sum(chistat.simulated >= chisq.stat)/B$, where

chistat.simulated is a vector storing the B chi-squared statistics generated under the Null Hypothesis, and
chisq.stat is the observed chi-squared statistic.

Both distributions can be optionally plotted setting the graph parameter to TRUE.

Confidence interval around the permutation-based and Monte Carlo p-value
The function calculates the 95 percent Confidence Interval around the simulated p-values. The Wald CI quantifies the uncertainty around the simulated p-value estimate. For a 95 percent CI, the

standard z-value of 1.96 is used. The standard error for the estimated p-value is computed as the square root of (estimated p-value * (1 - estimated p-value) / number of simulations-1).

The lower and upper bounds of the CI are then calculated as follows:
Lower Confidence Interval = estimated p-value - (z-value * standard error)
Upper Confidence Interval = estimated p-value + (z-value * standard error)

Finally, the lower and upper CIs are clipped to lie within 0 and 1.

The implemented procedure aligns with the one described at this link: https://blogs.sas.com/content/iml/2015/10/28/simulatio exact-tables.html

Moment-corrected standardized residuals
The moment-corrected standardized residuals are calculated as follows:

$stand.res/(sqrt((nr-1)*(nc-1)/(nr*nc)))$, where

stand.res is each cell's standardized residual, *nr* and *nc* are the number of rows and columns respectively.

See Garcia-Perez-Nunez-Anton 2003: 827.

Adjusted standardized residuals
The adjusted standardized residuals are calculated as follows:

$stand.res[i,j]/sqrt((1-sr[i]/n)*(1-sc[j]/n))$, where

stand.res is the standardized residual for cell *ij*, *sr* is the row sum for row *i*, *sc* is the column sum for column *j*, and *n* is the table grand total. The *adjusted standardized residuals* should be used in place of the standardised residuals since the latter are not truly standarised because they have a nonunit variance. The standardised residuals therefore underestimate the divergence between the observed and the expected counts. The adjusted standardized residuals (and the moment-corrected ones) correct that deficiency.

For more info see: Haberman 1973.

Significance of the residuals
The significance of the residuals (standardized, moment-corrected standardized, and adjusted standardized) is assessed using alpha 0.05 or, optionally (by setting the parameter adj.alpha to TRUE), using an adjusted alpha calculated using the Sidak's method:

$alpha.adj = 1-(1-0.05)^{(}1/(nr*nc))$, where

nr and *nc* are the number of rows and columns in the table respectively. The adjusted alpha is then converted into a critical two-tailed z value.

See: Beasley-Schumacker 1995: 86, 89.

Cells' relative contribution (in percent) to the chi-square statistic
The cells' relative contribution (in percent) to the chi-square statistic is calculated as:

$chisq.values/chisq.stat * 100$, where

chisq.values and *chisq.stat* are the chi-square value in each individual cell of the table and the value of the chi-square statistic, respectively. The *average contribution* is calculated as $100/(nr * nc)$, where *nr* and *nc* are the number of rows and columns in the table respectively.

Cells' absolute contribution (in percent) to the chi-square statistic
The cells' absolute contribution (in percent) to the chi-square statistic is calculated as:

$chisq.values/n * 100$, where

chisq.values and *n* are the chi-square value in each individual cell of the table and the table's grant total, respectively. The *average contribution* is calculated as sum of all the absolute contributions divided by the number of cells in the table.

For both the relative and absolute contributions to the chi-square, see: Beasley-Schumacker 1995: 90.

Phi corrected
To further refine Phi, a corrected version has been introduced. It accounts for the fact that the original coefficient (1) might not reach its maximum value of 1 even when there is a perfect association between the variables, and (2) it is not directly comparable across tables with different marginals. To calculate Phi-corrected, one first computes Phi-max, which represents the maximum possible value of Phi under the given marginal totals. Phi-corrected is equal to Phi/Phi-max.

For more details see: Cureton 1959; Liu 1980; Davenport et al. 1991; Rash et al. 2011.

95perc confidence interval around Cramer's V
The calculation of the 95perc confidence interval around Cramer's V is based on Smithson 2003: 39-41, and builds on the R code made available by the author on the web (http://www.michaelsmithson.online/stats/CIstuff/CI.htr

Bias-corrected Cramer's V
The bias-corrected Cramer's V is based on Bergsma 2013: 323–328.

W coefficient
It addresses some limitations of Cramer's V. When the marginal probabilities are unevenly distributed, V may overstate the strength of the association, proving pretty high even when the overall association is weak. W is based on the distance between observed and expected frequencies. It uses the squared distance to adjust for the unevenness of the marginal distributions in the table. The indication of the magnitude of the association is based on Cohen 1988 (see above). Unlike Kvalseth

2018a, the calculation of the 95 percent confidence interval is based on a bootstrap approach (employing 10k resampled tables, and the 2.5th and 97.5th percentiles of the bootstrap distribution).

For more details see: Kvalseth 2018a.

Corrected Goodman-Kruskal's lambda
The corrected Goodman-Kruskal's lambda adeptly addresses skewed or unbalanced marginal probabilities which create problems to the traditional lambda. By emphasizing categories with higher probabilities through a process of squaring maximum probabilities and normalizing with marginal probabilities, this refined coefficient addresses inherent limitations of lambda.

For more details see: Kvalseth 2018b.

Odds Ratio
The odds ratio is calculated for 2x2 tables. In case of zeros along any of the table's diagonal, the *Haldane-Anscombe* correction is applied. It consists in adding 0.5 to every cell of the table before calculating the odds ratio. For tables of size 2xk (where k >= 2), pairwise odds ratios can be plotted (along with their confidence interval) by setting the or.alpha parameter to TRUE. The mentioned correction is also applied to the calculation of those pairwise odds ratios (for more information on the plot, see further below).

For the Haldane-Anscombe correction see, for instance, Fleiss-Levin-Paik 2003: 102-103.

Odds Ratio effect size magnitude
The magnitude of the associaiton indicated by the odds ratio is based on the thresholds (and corresponding reciprocal) suggested by Chen et al 2010:

- OR < 1.68 - Very small
- 1.68 <= OR < 3.47 - Small
- 3.47 <= OR < 6.71 - Medium
- OR >= 6.71 - Large

Odd Ratios plot
For 2xk table, where k >= 2:
by setting the plor.or parameter to TRUE, a plot showing the odds ratios and their 95percent confidence interval will be rendered. The confidence level can be modified via the or.alpha parameter. The odds ratios are calculated for the column levels, and one of them is to be selected by the user as a reference for comparison via the reference.level parameter (set to 1 by default). Also, the user may want to select the row category to which the calculation of the odds ratios makes reference (using the row.level parameter, which is set to 1 by default). If any of the pairwisely-generated 2x2 tables on which the odds ratio is calculated features zeros along any of the diagonal, the *Haldane-Anscombe* correction is applied (see above).

To better understand the rationale of plotting the odds ratios, consider the following example, which uses on the famous Titanic data:

Create a 2x3 contingency table:
```
mytable <- matrix(c(123, 158, 528, 200, 119, 181), nrow = 2, byrow = TRUE)
colnames(mytable) <- c("1st", "2nd", "3rd")
rownames(mytable) <- c("Died", "Survived")
```

Now, we perform the test and visualise the odds ratios:
```
chisquare(mytable, plot.or=TRUE, reference.level=1, row.level=1)
```

In the rendered plot, we can see the odds ratios and confidence intervals for the second and third column level (i.e., 2nd class and 3rd class) because the first column level has been selected as reference level. The odds ratios are calculated making reference to the first row category (i.e., *Died*). From the plot, we can see that, compared to the 1st class, passengers on the 2nd class have 2.16 times larger odds of dying; passengers on the 3rd class have 4.74 times larger odds of dying compared to the 1st class.

Note that if we set the `row.level` parameter to 2, we make reference to the second row category, i.e. *Survived*:
```
chisquare(mytable, plot.or=TRUE, reference.level=1, row.level=2)
```

In the plot, we can see that passengers in the 2nd class have 0.46 times the odds of surviving of passengers in the 1st class, while passengers from the 3rd class have 0.21 times the odds of surviving of those travelling in the 1st class.

Other measures of categorical association
For the other measures of categorical association provided by the function, see for example Sheskin 2011: 1415-1427.

Additional notes on calculations:

- the **Phi** coefficient is based on the chi-square statistic as per Sheskin 2011's equation 16.21, whereas the **Phi signed** is after Sheskin's equation 16.20;

- the **2-sided p value of Yule's Q** is calculated following Sheskin 2011's equation 16.24;

- **Cohen's w** is calculated as $V * sqrt(min(nr, nc) - 1)$, where *V* is Cramer's V, and *nr* and *nc* are the number of rows and columns respectively; see Sheskin 2011: 679;

- the **2-tailed p value** of **Goodman-Kruskal's gamma** is based on the associated z-score calculated as per Sheskin 2011's equation 32.2;

- the **symmetric** version of **Goodman-Kruskal's lambda** is calculated as per Reynolds 1984: 55-57;

- **Goodman-Kruskal's tau** is calculated as per Reynolds 1984: 57-60;

- **Cohen's k** is calculated as per Sheskin 2011: 688-689 (equation 16.30).

Value

The function produces **optional charts** (distribution of the permuted chi-square statistic and a plot of the odds ratios between a reference column level and the other ones, the latter only for 2xk tables

where k >= 2), and a number of **output tables** that are nicely formatted with the help of the *gt* package. The output tables are listed below:

- Input contingency table (with some essential analytical results annotated at the bottom)
- Expected frequencies
- Cells' chi-square value
- Cells' relative contribution (in percent) to the chi-square statistic (cells in RED feature a larger-than-average contribution)
- Cells' absolute contribution (in percent) to the chi-square statistic (colour same as above)
- Standardized residuals (RED for large significant residuals, BLUE for small significant residuals)
- Moment-corrected standardized residuals (colour same as above)
- Adjusted standardized residuals (colour same as above)
- Table of output statistics, p values, and association measures

Also, the function returns a **list containing the following elements**:

- **input.table**:
    - *crosstab*: input contingency table.
- **chi.sq.related.results**:
    - *exp.freq*: table of expected frequencies.
    - *smallest.exp.freq*: smallest expected frequency.
    - *avrg.exp.freq*: average expected frequency.
    - *chisq.values*: cells' chi-square value.
    - *chisq.relat.contrib*: cells' relative contribution (in percent) to the chi-square statistic.
    - *chisq.abs.contrib*: cells' absolute contribution (in percent) to the chi-square statistic.
    - *chisq.statistic*: observed chi-square value.
    - *chisq.p.value*: p value of the chi-square statistic.
    - *chi.sq.power*: power of the traditional chi-square test.
    - *chisq.adj*: chi-square statistic adjusted using the (N-1)/N correction.
    - *chisq.adj.p.value*: p value of the adjusted chi-square statistic.
    - *chisq.p.value.perm*: permutation-based p value, based on B permuted tables.
    - *chisq.p.value.perm CI lower boundary*: lower boundary of the 95 percent CI around the permutation-based p value.
    - *chisq.p.value.perm CI upper boundary*: upper boundary of the 95 percent CI around the permutation-based p value.
    - *chisq.p.value.MC*: Monte Carlo p value, based on B random tables.
    - *chisq.p.value.MC CI lower boundary*: lower boundary of the 95 percent CI around the Monte Carlo p value.
    - *chisq.p.value.MC CI upper boundary*: upper boundary of the 95 percent CI around the Monte Carlo p value.
- **G.square**:
    - *Gsq.statistic*: observed G-square value.

  – *Gsq.p.value*: p value of the G-square statistic.
- **residuals**:
  – *stand.resid*: table of chi-square standardized residuals.
  – *mom.corr.stand.resid*: table of moment-corrected standardized residuals.
  – *adj.stand.resid*: table of adjusted standardized residuals.
- **chi.sq.based.assoc.measures**:
  – *Phi*: Phi coefficient (only for 2x2 tables).
  – *Phi corr*: corrected Phi coefficient (only for 2x2 tables).
  – *Phi signed*: signed Phi coefficient (only for 2x2 tables).
  – *Cadj*: adjusted contingency coefficient C.
  – *Cramer's V*: Cramer's V coefficient.
  – *Cramer's V CI lower boundary*: lower boundary of the 95perc CI.
  – *Cramer's V CI upper boundary*: upper boundary of the 95perc CI.
  – *Cramer's Vbc*: bias-corrected Cramer's V coefficient.
  – *w*: Cohen's w.
  – *W*: W coefficient.
  – *W CI lower boundary*: lower boundary of the 95perc CI.
  – *W CI upper boundary*: upper boundary of the 95perc CI.
- **non.chi.sq.based.assoc.measures**:
  – *Yule's Q*: Q coefficient (only for 2x2 tables).
  – *Yule's Q p.value*: 2-tailed p value of Yule's Q.
  – *Odds ratio*: odds ratio (only for 2x2 tables).
  – *Odds ratio CI lower boundary*: lower boundary of the 95perc CI.
  – *Odds ratio CI upper boundary*: upper boundary of the 95perc CI.
  – *Odds ratio p.value*: p value of the odds ratio.
  – *lambda (rows dep.)*: Goodman-Kruskal's lambda coefficient (considering the rows being the dependent variable).
  – *lambda (cols dep.)*: Goodman-Kruskal's lambda coefficient (considering the columns being the dependent variable).
  – *lambda (symmetric)*: Goodman-Kruskal's symmetric lambda coefficient.
  – *lambda corrected (rows dep.)*: corrected version of the lambda coefficient (considering the rows being the dependent variable).
  – *lambda corrected (cols dep.)*: corrected version of the lambda coefficient (considering the columns being the dependent variable).
  – *lambda corrected (symmetric)*: corrected version of the symmetric lambda coefficient.
  – *tau (rows dep.)*: Goodman-Kruskal's tau coefficient (considering the rows being the dependent variable).
  – *tau (cols dep.)*: Goodman-Kruskal's tau coefficient (considering the columns being the dependent variable).
  – *gamma*: Goodman-Kruskal's gamma coefficient.
  – *gamma.p.value*: 2-sided p value for the Goodman-Kruskal's gamma coefficient.
  – *k*: Cohen'k.

– *k CI lower boundary*: lower boundary of the 95perc CI.

– *k CI upper boundary*: upper boundary of the 95perc CI.

Note that the *p-values* returned in the above list are expressed in scientific notation, whereas the ones reported in the output table featuring the tests' result and measures of association are reported as broken down into classes (e.g., <0.05, or <0.01, etc), with the exception of the Monte Carlo p-value and its CI.

The **following examples**, which use in-built datasets, can be run to familiarise with the function:

-perform the test on the in-built 'social_class' dataset:
```
result <- chisquare(social_class)
```

-perform the test on a 2x2 subset of the 'diseases' dataset:
```
mytable <- diseases[3:4,1:2]
result <- chisquare(mytable)
```

-perform the test on a 2x2 subset of the 'safety' dataset:
```
mytable <- safety[c(4,1),c(1,6)]
result <- chisquare(mytable)
```

-build a toy dataset in 'long' format (gender vs. opinion about death sentence):
```
mytable <- data.frame(GENDER=c(rep("F", 360), rep("M", 340)),OPINION=c(rep("oppose",
235), rep("favour", 125), rep("oppose", 160), rep("favour", 180)))
```
-perform the test specifying that the input table is in 'long' format:
```
result <- chisquare(mytable, format="long")
```

References

Agresti, A., Franklin, C., & Klingenberg, B. (2022). Statistics: The Art and Science of Learning from Data, (5th ed.). Pearson Education.

Beh E.J., Lombardo R. 2014. Correspondence Analysis: Theory, Practice and New Strategies, Chichester, Wiley.

Beasley TM and Schumacker RE. 1995. Multiple Regression Approach to Analyzing Contingency Tables: Post Hoc and Planned Comparison Procedures. The Journal of Experimental Education, 64(1).

Bergsma, W. 2013. A bias correction for Cramér's V and Tschuprow's T. Journal of the Korean Statistical Society. 42 (3).

Campbell, I. (2007). Chi-squared and Fisher–Irwin tests of two-by-two tables with small sample recommendations. In Statistics in Medicine (Vol. 26, Issue 19, pp. 3661–3675).

Chen, H., Cohen, P., and Chen, S. (2010). How Big is a Big Odds Ratio? Interpreting the Magnitudes of Odds Ratios in Epidemiological Studies. In Communications in Statistics - Simulation and Computation (Vol. 39, Issue 4, pp. 860–864).

Cohen, J. 1988. Statistical power analysis for the behavioral sciences (2nd ed). Hillsdale, N.J: L. Erlbaum Associates.

Cureton, E. E. (1959). Note on phi/phimax. In Psychometrika (Vol. 24, Issue 1, pp. 89–91).

Davenport, E. C., Jr., & El-Sanhurry, N. A. (1991). Phi/Phimax: Review and Synthesis. In Educational and Psychological Measurement (Vol. 51, Issue 4, pp. 821–828).

Fleiss, J. L., Levin, B., & Paik, M. C. 2003. Statistical Methods for Rates and Proportions (3rd ed.). Wiley.

Garcia-Perez, MA, and Nunez-Anton, V. 2003. Cellwise Residual Analysis in Two-Way Contingency Tables. Educational and Psychological Measurement, 63(5).

Greenwood, P. E., & Nikulin, M. S. (1996). A guide to chi-squared testing. John Wiley & Sons.

Haberman, S. J. (1973). The Analysis of Residuals in Cross-Classified Tables. In Biometrics (Vol. 29, Issue 1, p. 205).

Kvålseth, T. O. (2018a). An alternative to Cramér's coefficient of association. In Communications in Statistics - Theory and Methods (Vol. 47, Issue 23, pp. 5662–5674).

Kvålseth, T. O. (2018b). Measuring association between nominal categorical variables: an alternative to the Goodman–Kruskal lambda. In Journal of Applied Statistics (Vol. 45, Issue 6, pp. 1118–1132).

Oyeyemi, G. M., Adewara, A. A., Adebola, F. B., & Salau, S. I. (2010). On the Estimation of Power and Sample Size in Test of Independence. In Asian Journal of Mathematics and Statistics (Vol. 3, Issue 3, pp. 139–146).

Rasch, D., Kubinger, K. D., & Yanagida, T. (2011). Statistics in Psychology Using R and SPSS. Wiley.

Reynolds, H. T. 1984. Analysis of Nominal Data (Quantitative Applications in the Social Sciences) (1st ed.). SAGE Publications.

Rhoades, H. M., & Overall, J. E. (1982). A sample size correction for Pearson chi-square in 2×2 contingency tables. In Psychological Bulletin (Vol. 91, Issue 2, pp. 418–423).

Richardson, J. T. E. (2011). The analysis of 2×2 contingency tables-Yet again. In Statistics in Medicine (Vol. 30, Issue 8, pp. 890–890).

Roscoe, J. T., & Byars, J. A. (1971). An Investigation of the Restraints with Respect to Sample Size Commonly Imposed on the Use of the Chi-Square Statistic. Journal of the American Statistical Association, 66(336), 755–759.

Sheskin, D. J. 2011. Handbook of Parametric and Nonparametric Statistical Procedures, Fifth Edition (5th ed.). Chapman and Hall/CRC.

Smithson M.J. 2003. Confidence Intervals, Quantitative Applications in the Social Sciences Series, No. 140. Thousand Oaks, CA: Sage.

Upton, G. J. G. (1982). A Comparison of Alternative Tests for the 2×2 Comparative Trial. In Journal of the Royal Statistical Society. Series A (General) (Vol. 145, Issue 1, p. 86).

Zar, J. H. (2014). Biostatistical analysis (5th ed.). Pearson New International Edition.

Examples

```
# Perform the test on the in-built 'social_class' dataset
result <- chisquare(social_class, B=99)


# Perform the test on a 2x2 subset
```

```
result <- chisquare(social_class[c(1:2), c(1:2)], B=99)
```

---

diseases	*Dataset: Cross-tabulation of quantity of tobacco smoked daily vs. cause of death*

Description

Cross-tabulation (15x4) of the amount of tobacco smoked on a daily basis (in gramms) against cause of death.

After: Velleman P F, Hoaglin D C, Applications, Basics, and Computing of Exploratory Data Analysis, Wadsworth Pub Co 1984 (Exhibit 8-1)

Usage

```
data(diseases)
```

Format

dataframe

---

safety	*Dataset: Cross-tabulation of people's feeling of safety vs. town size*

Description

Cross-tabulation (4x6).

Usage

```
data(safety)
```

Format

dataframe

social_class	*Dataset: Cross-tabulation of social class vs. diagnostic category for a sample of psychiatric patients*

Description

Cross-tabulation (3x4) after: Everitt B.S (1992), The Analysis of Contingency Tables, Chapman&Hall/CRC, second edition, table 3.13.

Usage

```
data(social_class)
```

Format

dataframe

Index