Supporting Appendix: Fisher's Exact Test for Count Data

Description

Performs Fisher's exact test for testing the null of independence of rows and columns in a contingency table with fixed marginals.

Usage

fisher.test(x, y = NULL, workspace = 200000, hybrid = FALSE,
or = 1, alternative = "two.sided," conf. level = 0.95)

Arguments

x	either a two-dimensional contingency table in matrix form, or a factor object.
y	a factor object; ignored if x is a matrix.
workspace	an integer specifying the size of the workspace used in the network algorithm.
hybrid	a logical indicating whether the exact probabilities (default) or a hybrid approximation thereof should be computed. In the hybrid case, asymptotic chi-squared probabilities are only used provided that the "Cochran" conditions are satisfied.
or	the hypothesized odds ratio. Only used in the 2 by 2 case.
alternative	indicates the alternative hypothesis and must be one of "two.sided," "greater," or "less." You can specify just the initial letter. Only used in the 2 by 2 case.
conf.level	confidence level for the returned confidence interval. Only used in the 2 by 2 case.

Details

If x is a matrix, it is taken as a two-dimensional contingency table, and hence its entries should be nonnegative integers. Otherwise, both x and y must be vectors of the same length. Incomplete cases are removed, the vectors are coerced into factor objects, and the contingency table is computed from these.

In the one-sided 2 by 2 cases, *P* values are obtained directly by using the hypergeometric distribution. Otherwise, computations are based on a C version of the FORTRAN subroutine FEXACT (1-2). The FORTRAN code can be accessed at www.netlib.org/toms/643.

In the 2 by 2 case, the null of conditional independence is equivalent to the hypothesis that the odds ratio equals one. Exact inference can be based on observing that in general, given all marginal totals fixed, the first element of the contingency table has a noncentral hypergeometric distribution with noncentrality parameter given by the odds ratio (3).

Value

A list with class “htest” containing the following components:

p.value	the *P* value of the test.
conf.int	a confidence interval for the odds ratio. Only present in the 2 by 2 case.
estimate	an estimate of the odds ratio. Note that the conditional maximum likelihood estimate (MLE) rather than the unconditional MLE (the sample odds ratio) is used. Only present in the 2 by 2 case.
null.value	the odds ratio under the null, or. Only present in the 2 by 2 case.
alternative	a character string describing the alternative hypothesis.
method	the character string “Fisher's Exact Test for Count Data.”
data.name	a character string giving the names of the data.

1. Mehta, C. R. & Patel, N. R. (1986) *ACM Transactions Mathematical Software* **12,** 154-161.

2. Clarkson, D. B., Fan, Y. & Joe, H. (1993) *ACM Transactions Mathematical Software* **19**, 484-488.

3. Fisher, R. A. (1935) *J. R. Stat. Soc. A* **98**, 39–54.