

SAS/STAT[®] 12.1 User's Guide

The FREQ Procedure

(Chapter)



This document is an individual chapter from *SAS/STAT® 12.1 User's Guide*.

The correct bibliographic citation for the complete manual is as follows: SAS Institute Inc. 2012. *SAS/STAT® 12.1 User's Guide*. Cary, NC: SAS Institute Inc.

Copyright © 2012, SAS Institute Inc., Cary, NC, USA

All rights reserved. Produced in the United States of America.

For a Web download or e-book: Your use of this publication shall be governed by the terms established by the vendor at the time you acquire this publication.

The scanning, uploading, and distribution of this book via the Internet or any other means without the permission of the publisher is illegal and punishable by law. Please purchase only authorized electronic editions and do not participate in or encourage electronic piracy of copyrighted materials. Your support of others' rights is appreciated.

U.S. Government Restricted Rights Notice: Use, duplication, or disclosure of this software and related documentation by the U.S. government is subject to the Agreement with SAS Institute and the restrictions set forth in FAR 52.227-19, Commercial Computer Software-Restricted Rights (June 1987).

SAS Institute Inc., SAS Campus Drive, Cary, North Carolina 27513.

Electronic book 1, August 2012

SAS® Publishing provides a complete selection of books and electronic products to help customers use SAS software to its fullest potential. For more information about our e-books, e-learning products, CDs, and hard-copy books, visit the SAS Publishing Web site at support.sas.com/publishing or call 1-800-727-3228.

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 registered trademarks or trademarks of their respective companies.

Chapter 38

The FREQ Procedure

Contents

Overview: FREQ Procedure	2472
Getting Started: FREQ Procedure	2474
Frequency Tables and Statistics	2474
Agreement Study	2481
Syntax: FREQ Procedure	2484
PROC FREQ Statement	2484
BY Statement	2486
EXACT Statement	2487
OUTPUT Statement	2494
TABLES Statement	2505
TEST Statement	2540
WEIGHT Statement	2543
Details: FREQ Procedure	2544
Inputting Frequency Counts	2544
Grouping with Formats	2544
Missing Values	2545
In-Database Computation	2548
Statistical Computations	2549
Definitions and Notation	2549
Chi-Square Tests and Statistics	2550
Measures of Association	2556
Binomial Proportion	2565
Risks and Risk Differences	2572
Odds Ratio and Relative Risks for 2 x 2 Tables	2582
Cochran-Armitage Test for Trend	2585
Jonckheere-Terpstra Test	2586
Tests and Measures of Agreement	2588
Cochran-Mantel-Haenszel Statistics	2592
Gail-Simon Test for Qualitative Interactions	2601
Exact Statistics	2601
Computational Resources	2606
Output Data Sets	2606
Displayed Output	2609
ODS Table Names	2617
ODS Graphics	2621
Examples: FREQ Procedure	2622

Example 38.1: Output Data Set of Frequencies	2622
Example 38.2: Frequency Dot Plots	2625
Example 38.3: Chi-Square Goodness-of-Fit Tests	2628
Example 38.4: Binomial Proportions	2632
Example 38.5: Analysis of a 2x2 Contingency Table	2635
Example 38.6: Output Data Set of Chi-Square Statistics	2638
Example 38.7: Cochran-Mantel-Haenszel Statistics	2640
Example 38.8: Cochran-Armitage Trend Test	2642
Example 38.9: Friedman's Chi-Square Test	2646
Example 38.10: Cochran's Q Test	2647
References	2650

Overview: FREQ Procedure

The FREQ procedure produces one-way to n -way frequency and contingency (crosstabulation) tables. For two-way tables, PROC FREQ computes tests and measures of association. For n -way tables, PROC FREQ provides stratified analysis by computing statistics across, as well as within, strata.

For one-way frequency tables, PROC FREQ computes goodness-of-fit tests for equal proportions or specified null proportions. For one-way tables, PROC FREQ also provides confidence limits and tests for binomial proportions, including tests for noninferiority and equivalence.

For contingency tables, PROC FREQ can compute various statistics to examine the relationships between two classification variables. For some pairs of variables, you might want to examine the existence or strength of any association between the variables. To determine if an association exists, chi-square tests are computed. To estimate the strength of an association, PROC FREQ computes measures of association that tend to be close to zero when there is no association and close to the maximum (or minimum) value when there is perfect association. The statistics for contingency tables include the following:

- chi-square tests and measures
- measures of association
- risks (binomial proportions) and risk differences for 2×2 tables
- odds ratios and relative risks for 2×2 tables
- tests for trend
- tests and measures of agreement
- Cochran-Mantel-Haenszel statistics

PROC FREQ computes asymptotic standard errors, confidence intervals, and tests for measures of association and measures of agreement. Exact p -values and confidence intervals are available for many test statistics and measures. PROC FREQ also performs analyses that adjust for any stratification variables by computing statistics across, as well as within, strata for n -way tables. These statistics include Cochran-Mantel-Haenszel statistics and measures of agreement.

In choosing measures of association to use in analyzing a two-way table, you should consider the study design (which indicates whether the row and column variables are dependent or independent), the measurement scale of the variables (nominal, ordinal, or interval), the type of association that each measure is designed to detect, and any assumptions required for valid interpretation of a measure. You should exercise care in selecting measures that are appropriate for your data.

Similar comments apply to the choice and interpretation of test statistics. For example, the Mantel-Haenszel chi-square statistic requires an ordinal scale for both variables and is designed to detect a linear association. The Pearson chi-square, on the other hand, is appropriate for all variables and can detect any kind of association, but it is less powerful for detecting a linear association because its power is dispersed over a greater number of degrees of freedom (except for 2×2 tables).

For more information about selecting the appropriate statistical analyses, see Agresti (2007) and Stokes, Davis, and Koch (2012).

Several SAS procedures produce frequency counts; only PROC FREQ computes chi-square tests for one-way to n -way tables and measures of association and agreement for contingency tables. Other procedures to consider for counting include the TABULATE and UNIVARIATE procedures. When you want to produce contingency tables and tests of association for sample survey data, use PROC SURVEYFREQ. See Chapter 14, “[Introduction to Survey Procedures](#),” for more information. When you want to fit models to categorical data, use a procedure such as CATMOD, GENMOD, GLIMMIX, LOGISTIC, PROBIT, or SURVEYLOGISTIC. See Chapter 8, “[Introduction to Categorical Data Analysis Procedures](#),” for more information.

PROC FREQ uses the Output Delivery System (ODS), a SAS subsystem that provides capabilities for displaying and controlling the output from SAS procedures. ODS enables you to convert any of the output from PROC FREQ into a SAS data set. See the section “[ODS Table Names](#)” on page 2617 for more information.

PROC FREQ uses ODS Graphics to create graphs as part of its output. For general information about ODS Graphics, see Chapter 21, “[Statistical Graphics Using ODS](#).” For specific information about the statistical graphics available with the FREQ procedure, see the **PLOTS=** option in the TABLES statement and the section “[ODS Graphics](#)” on page 2621.

Getting Started: FREQ Procedure

Frequency Tables and Statistics

The FREQ procedure provides easy access to statistics for testing for association in a crosstabulation table.

In this example, high school students applied for courses in a summer enrichment program; these courses included journalism, art history, statistics, graphic arts, and computer programming. The students accepted were randomly assigned to classes with and without internships in local companies. Table 38.1 contains counts of the students who enrolled in the summer program by gender and whether they were assigned an internship slot.

Table 38.1 Summer Enrichment Data

Gender	Internship	Enrollment		
		Yes	No	Total
boys	yes	35	29	64
boys	no	14	27	41
girls	yes	32	10	42
girls	no	53	23	76

The SAS data set SummerSchool is created by inputting the summer enrichment data as cell count data, or providing the frequency count for each combination of variable values. The following DATA step statements create the SAS data set SummerSchool:

```
data SummerSchool;
  input Gender $ Internship $ Enrollment $ Count @@;
  datalines;
boys yes yes 35   boys yes no 29
boys no yes 14   boys no no 27
girls yes yes 32   girls yes no 10
girls no yes 53   girls no no 23
;
```

The variable Gender takes the values ‘boys’ or ‘girls,’ the variable Internship takes the values ‘yes’ and ‘no,’ and the variable Enrollment takes the values ‘yes’ and ‘no.’ The variable Count contains the number of students that correspond to each combination of data values. The double at sign (@@) indicates that more than one observation is included on a single data line. In this DATA step, two observations are included on each line.

Researchers are interested in whether there is an association between internship status and summer program enrollment. The Pearson chi-square statistic is an appropriate statistic to assess the association in the corresponding 2×2 table. The following PROC FREQ statements specify this analysis.

You specify the table for which you want to compute statistics with the TABLES statement. You specify the statistics you want to compute with options after a slash (/) in the TABLES statement.

```
proc freq data=SummerSchool order=data;
  tables Internship*Enrollment / chisq;
  weight Count;
run;
```

The ORDER= option controls the order in which variable values are displayed in the rows and columns of the table. By default, the values are arranged according to the alphanumeric order of their unformatted values. If you specify ORDER=DATA, the data are displayed in the same order as they occur in the input data set. Here, because 'yes' appears before 'no' in the data, 'yes' appears first in any table. Other options for controlling order include ORDER=FORMATTED, which orders according to the formatted values, and ORDER=FREQUENCY, which orders by descending frequency count.

In the TABLES statement, Internship*Enrollment specifies a table where the rows are internship status and the columns are program enrollment. The CHISQ option requests chi-square statistics for assessing association between these two variables. Because the input data are in cell count form, the WEIGHT statement is required. The WEIGHT statement names the variable Count, which provides the frequency of each combination of data values.

Figure 38.1 presents the crosstabulation of Internship and Enrollment. In each cell, the values printed under the cell count are the table percentage, row percentage, and column percentage, respectively. For example, in the first cell, 63.21 percent of the students offered courses with internships accepted them and 36.79 percent did not.

Figure 38.1 Crosstabulation Table

The FREQ Procedure				
Table of Internship by Enrollment				
Internship	Enrollment			
Frequency				
Percent				
Row Pct				
Col Pct	yes	no	Total	
-----+-----+-----+				
yes	67	39	106	
	30.04	17.49	47.53	
	63.21	36.79		
	50.00	43.82		
-----+-----+-----+				
no	67	50	117	
	30.04	22.42	52.47	
	57.26	42.74		
	50.00	56.18		
-----+-----+-----+				
Total	134	89	223	
	60.09	39.91	100.00	

Figure 38.2 displays the statistics produced by the CHISQ option. The Pearson chi-square statistic is labeled 'Chi-Square' and has a value of 0.8189 with 1 degree of freedom. The associated p -value is 0.3655, which means that there is no significant evidence of an association between internship status and program enrollment. The other chi-square statistics have similar values and are asymptotically equivalent. The other statistics (phi coefficient, contingency coefficient, and Cramér's V) are measures of association derived from the Pearson chi-square. For Fisher's exact test, the two-sided p -value is 0.4122, which also shows no association between internship status and program enrollment.

Figure 38.2 Statistics Produced with the CHISQ Option

Statistic	DF	Value	Prob
Chi-Square	1	0.8189	0.3655
Likelihood Ratio Chi-Square	1	0.8202	0.3651
Continuity Adj. Chi-Square	1	0.5899	0.4425
Mantel-Haenszel Chi-Square	1	0.8153	0.3666
Phi Coefficient		0.0606	
Contingency Coefficient		0.0605	
Cramer's V		0.0606	
Fisher's Exact Test			
Cell (1,1) Frequency (F)		67	
Left-sided Pr <= F		0.8513	
Right-sided Pr >= F		0.2213	
Table Probability (P)		0.0726	
Two-sided Pr <= P		0.4122	

The analysis, so far, has ignored gender. However, it might be of interest to ask whether program enrollment is associated with internship status after adjusting for gender. You can address this question by doing an analysis of a set of tables (in this case, by analyzing the set consisting of one for boys and one for girls). The Cochran-Mantel-Haenszel (CMH) statistic is appropriate for this situation: it addresses whether rows and columns are associated after controlling for the stratification variable. In this case, you would be stratifying by gender.

The PROC FREQ statements for this analysis are very similar to those for the first analysis, except that there is a third variable, Gender, in the TABLES statement. When you cross more than two variables, the two rightmost variables construct the rows and columns of the table, respectively, and the leftmost variables determine the stratification.

The following PROC FREQ statements also request frequency plots for the crosstabulation tables. PROC FREQ produces these plots by using ODS Graphics to create graphs as part of the procedure output. ODS Graphics must be enabled before producing plots. The PLOTS(ONLY)=FREQPLOT option requests frequency plots. The TWOWAY=CLUSTER *plot-option* specifies a cluster layout for the two-way frequency plots.

```
ods graphics on;
proc freq data=SummerSchool;
  tables Gender*Internship*Enrollment /
    chisq cmh plots(only)=freqplot(twoway=cluster);
  weight Count;
run;
ods graphics off;
```


This execution of PROC FREQ first produces two individual crosstabulation tables of Internship by Enrollment: one for boys and one for girls. Frequency plots and chi-square statistics are produced for each individual table. [Figure 38.3](#), [Figure 38.4](#), and [Figure 38.5](#) show the results for boys. Note that the chi-square statistic for boys is significant at the $\alpha = 0.05$ level of significance. Boys offered a course with an internship are more likely to enroll than boys who are not.

[Figure 38.4](#) displays the frequency plot of Internship by Enrollment for boys. By default, frequency plots are displayed as bar charts. You can use PLOTS= options to request dot plots instead of bar charts, to change the orientation of the bars from vertical to horizontal, and to change the scale from frequencies to percents. You can also use PLOTS= options to specify other two-way layouts (stacked, vertical groups, or horizontal groups) and to change the primary grouping from column levels to row levels.

[Figure 38.6](#), [Figure 38.7](#), and [Figure 38.8](#) display the crosstabulation table, frequency plot, and chi-square statistics for girls. You can see that there is no evidence of association between internship offers and program enrollment for girls.

Figure 38.3 Crosstabulation Table for Boys

The FREQ Procedure				
Table 1 of Internship by Enrollment Controlling for Gender=boys				
Internship	Enrollment			
Frequency				
Percent				
Row Pct				
Col Pct	no	yes	Total	
-----+-----+-----+				
no	27	14	41	
	25.71	13.33	39.05	
	65.85	34.15		
	48.21	28.57		
-----+-----+-----+				
yes	29	35	64	
	27.62	33.33	60.95	
	45.31	54.69		
	51.79	71.43		
-----+-----+-----+				
Total	56	49	105	
	53.33	46.67	100.00	


Figure 38.4 Frequency Plot for Boys**Figure 38.5** Chi-Square Statistics for Boys

Statistic	DF	Value	Prob
Chi-Square	1	4.2366	0.0396
Likelihood Ratio Chi-Square	1	4.2903	0.0383
Continuity Adj. Chi-Square	1	3.4515	0.0632
Mantel-Haenszel Chi-Square	1	4.1963	0.0405
Phi Coefficient		0.2009	
Contingency Coefficient		0.1969	
Cramer's V		0.2009	

Fisher's Exact Test		
Cell (1,1) Frequency (F)	27	
Left-sided Pr <= F	0.9885	
Right-sided Pr >= F	0.0311	
Table Probability (P)	0.0196	
Two-sided Pr <= P	0.0467	

Figure 38.6 Crosstabulation Table for Girls

Table 2 of Internship by Enrollment Controlling for Gender=girls				
Internship		Enrollment		
Frequency				
Percent				
Row Pct				
Col Pct	no	yes		Total
-----+-----+-----+				
no	23	53		76
	19.49	44.92		64.41
	30.26	69.74		
	69.70	62.35		
-----+-----+-----+				
yes	10	32		42
	8.47	27.12		35.59
	23.81	76.19		
	30.30	37.65		
-----+-----+-----+				
Total	33	85		118
	27.97	72.03		100.00

Figure 38.7 Frequency Plot for Girls**Figure 38.8** Chi-Square Statistics for Girls

Statistic	DF	Value	Prob
Chi-Square	1	0.5593	0.4546
Likelihood Ratio Chi-Square	1	0.5681	0.4510
Continuity Adj. Chi-Square	1	0.2848	0.5936
Mantel-Haenszel Chi-Square	1	0.5545	0.4565
Phi Coefficient		0.0688	
Contingency Coefficient		0.0687	
Cramer's V		0.0688	

Fisher's Exact Test		
Cell (1,1) Frequency (F)	23	
Left-sided Pr <= F	0.8317	
Right-sided Pr >= F	0.2994	
Table Probability (P)	0.1311	
Two-sided Pr <= P	0.5245	

These individual table results demonstrate the occasional problems with combining information into one table and not accounting for information in other variables such as Gender. Figure 38.9 contains the CMH results. There are three summary (CMH) statistics; which one you use depends on whether your rows and/or columns have an order in $r \times c$ tables. However, in the case of 2×2 tables, ordering does not matter and all three statistics take the same value. The CMH statistic follows the chi-square distribution under the hypothesis of no association, and here, it takes the value 4.0186 with 1 degree of freedom. The associated p -value is 0.0450, which indicates a significant association at the $\alpha = 0.05$ level.

Thus, when you adjust for the effect of gender in these data, there is an association between internship and program enrollment. But, if you ignore gender, no association is found. Note that the CMH option also produces other statistics, including estimates and confidence limits for relative risk and odds ratios for 2×2 tables and the Breslow-Day Test. These results are not displayed here.

Figure 38.9 Test for the Hypothesis of No Association

Cochran-Mantel-Haenszel Statistics (Based on Table Scores)				
Statistic	Alternative Hypothesis	DF	Value	Prob
1	Nonzero Correlation	1	4.0186	0.0450
2	Row Mean Scores Differ	1	4.0186	0.0450
3	General Association	1	4.0186	0.0450

Agreement Study

Medical researchers are interested in evaluating the efficacy of a new treatment for a skin condition. Dermatologists from participating clinics were trained to conduct the study and to evaluate the condition. After the training, two dermatologists examined patients with the skin condition from a pilot study and rated the same patients. The possible evaluations are terrible, poor, marginal, and clear. Table 38.2 contains the data.

Table 38.2 Skin Condition Data

Dermatologist 1	Dermatologist 2			
	Terrible	Poor	Marginal	Clear
Terrible	10	4	1	0
Poor	5	10	12	2
Marginal	2	4	12	5
Clear	0	2	6	13

The following DATA step statements create the SAS dataset `SkinCondition`. The dermatologists' evaluations of the patients are contained in the variables `Derm1` and `Derm2`; the variable `Count` is the number of patients given a particular pair of ratings.

```
data SkinCondition;
    input Derm1 $ Derm2 $ Count;
    datalines;
    terrible terrible 10
    terrible      poor 4
    terrible      marginal 1
    terrible      clear 0
    poor          terrible 5
    poor          poor 10
    poor          marginal 12
    poor          clear 2
    marginal      terrible 2
    marginal      poor 4
    marginal      marginal 12
    marginal      clear 5
    clear         terrible 0
    clear         poor 2
    clear         marginal 6
    clear         clear 13
    ;
```

The following PROC FREQ statements request an agreement analysis of the skin condition data. In order to evaluate the agreement of the diagnoses (a possible contribution to measurement error in the study), the *kappa coefficient* is computed.

The TABLES statement requests a crosstabulation of the variables `Derm1` and `Derm2`. The AGREE option in the TABLES statement requests the kappa coefficient, together with its standard error and confidence limits. The KAPPA option in the TEST statement requests a test for the null hypothesis that kappa equals zero, or that the agreement is purely by chance. The NOPRINT option in the TABLES statement suppresses the display of the two-way table. The PLOTS= option requests an agreement plot for the two dermatologists. ODS Graphics must be enabled before producing plots.

```
ods graphics on;
proc freq data=SkinCondition order=data;
    tables Derm1*Derm2 /
        agree noprint plots=agreeplot;
    test kappa;
    weight Count;
run;
ods graphics off;
```

Figure 38.10 and Figure 38.11 show the results. The kappa coefficient has the value 0.3449, which indicates some agreement between the dermatologists, and the hypothesis test confirms that you can reject the null hypothesis of no agreement. This conclusion is further supported by the confidence interval of (0.2030, 0.4868), which suggests that the true kappa is greater than zero. The AGREE option also produces Bowker's test for symmetry and the weighted kappa coefficient, but that output is not shown here. Figure 38.11 displays the agreement plot for the ratings of the two dermatologists.

Figure 38.10 Agreement Study





Figure 38.11 Agreement Plot



Syntax: FREQ Procedure

The following statements are available in the FREQ procedure:

```
PROC FREQ < options > ;
  BY variables ;
  EXACT statistic-options < / computation-options > ;
  OUTPUT < OUT=SAS-data-set > output-options ;
  TABLES requests < / options > ;
  TEST options ;
  WEIGHT variable < / option > ;
```

The PROC FREQ statement is the only required statement for the FREQ procedure. If you specify the following statements, PROC FREQ produces a one-way frequency table for each variable in the most recently created data set.

```
proc freq;
run;
```

The rest of this section gives detailed syntax information for the BY, EXACT, OUTPUT, TABLES, TEST, and WEIGHT statements in alphabetical order after the description of the PROC FREQ statement. Table 38.3 summarizes the basic function of each PROC FREQ statement.

Table 38.3 Summary of PROC FREQ Statements

Statement	Description
BY	Provides separate analyses for each BY group
EXACT	Requests exact tests
OUTPUT	Requests an output data set
TABLES	Specifies tables and requests analyses
TEST	Requests tests for measures of association and agreement
WEIGHT	Identifies a weight variable

PROC FREQ Statement

```
PROC FREQ < options > ;
```

The PROC FREQ statement invokes the FREQ procedure. Optionally, it also identifies the input data set. By default, the procedure uses the most recently created SAS data set.

Table 38.4 lists the *options* available in the PROC FREQ statement. Descriptions of the *options* follow in alphabetical order.

Table 38.4 PROC FREQ Statement Options

Option	Description
COMPRESS	Begins the next one-way table on the current page
DATA=	Names the input data set
FORMCHAR=	Specifies the outline and cell divider characters for crosstabulation tables
NLEVELS	Displays the number of levels for all TABLES variables
NOPRINT	Suppresses all displayed output
ORDER=	Specifies the order for reporting variable values
PAGE	Displays one table per page

You can specify the following *options* in the PROC FREQ statement.

COMPRESS

begins display of the next one-way frequency table on the same page as the preceding one-way table if there is enough space to begin the table. By default, the next one-way table begins on the current page only if the entire table fits on that page. The COMPRESS option is not valid with the [PAGE](#) option.

DATA=SAS-data-set

names the SAS data set to be analyzed by PROC FREQ. If you omit the DATA= option, the procedure uses the most recently created SAS data set.

FORMCHAR(1,2,7)= 'formchar-string'

defines the characters to be used for constructing the outlines and dividers for the cells of crosstabulation table displays. The *formchar-string* should be three characters long. The characters are used to draw the vertical separators (1), the horizontal separators (2), and the vertical-horizontal intersections (7). If you do not specify the FORMCHAR= option, PROC FREQ uses FORMCHAR(1,2,7)= '|-+' by default. [Table 38.5](#) summarizes the formatting characters used by PROC FREQ.

Table 38.5 Formatting Characters Used by PROC FREQ

Position	Default	Used to Draw
1		Vertical separators
2	-	Horizontal separators
7	+	Intersections of vertical and horizontal separators

The FORMCHAR= option can specify 20 different SAS formatting characters used to display output; however, PROC FREQ uses only the first, second, and seventh formatting characters. Therefore, the proper specification for PROC FREQ is FORMCHAR(1,2,7)= '*formchar-string*'.

Specifying all blanks for *formchar-string* produces crosstabulation tables with no outlines or dividers—for example, FORMCHAR(1,2,7)= ' '. You can use any character in *formchar-string*, including hexadecimal characters. If you use hexadecimal characters, you must put an *x* after the closing quote. For information about which hexadecimal codes to use for which characters, see the documentation for your hardware.

See the CALENDAR, PLOT, and TABULATE procedures in the *Base SAS Procedures Guide* for more information about form characters.

NLEVELS

displays the “Number of Variable Levels” table, which provides the number of levels for each variable named in the TABLES statements. See the section “[Number of Variable Levels Table](#)” on page 2609 for details. PROC FREQ determines the variable levels from the formatted variable values, as described in the section “[Grouping with Formats](#)” on page 2544.

NOPRINT

suppresses the display of all output. You can use the NOPRINT option when you only want to create an output data set. See the section “[Output Data Sets](#)” on page 2606 for information about the output data sets produced by PROC FREQ. Note that the NOPRINT option temporarily disables the Output Delivery System (ODS). For more information, see Chapter 20, “[Using the Output Delivery System](#).”

NOTE: A [NOPRINT](#) option is also available in the [TABLES](#) statement. It suppresses display of the crosstabulation tables but allows display of the requested statistics.

ORDER=DATA | FORMATTED | FREQ | INTERNAL

specifies the order of the variable levels in the frequency and crosstabulation tables, which you request in the [TABLES](#) statement.

The ORDER= option can take the following values:

Value of ORDER=	Levels Ordered By
DATA	Order of appearance in the input data set
FORMATTED	External formatted value, except for numeric variables with no explicit format, which are sorted by their unformatted (internal) value
FREQ	Descending frequency count; levels with the most observations come first in the order
INTERNAL	Unformatted value

By default, ORDER=INTERNAL. The FORMATTED and INTERNAL orders are machine-dependent. The ORDER= option does not apply to missing values, which are always ordered first.

For more information about sort order, see the chapter on the SORT procedure in the *Base SAS Procedures Guide* and the discussion of BY-group processing in *SAS Language Reference: Concepts*.

PAGE

displays only one table per page. Otherwise, PROC FREQ displays multiple tables per page as space permits. The PAGE option is not valid with the [COMPRESS](#) option.

BY Statement

BY *variables* ;

You can specify a BY statement with PROC FREQ to obtain separate analyses of observations in groups that are defined by the BY variables. When a BY statement appears, the procedure expects the input data set to be sorted in order of the BY variables. If you specify more than one BY statement, only the last one specified is used.

If your input data set is not sorted in ascending order, use one of the following alternatives:

- Sort the data by using the SORT procedure with a similar BY statement.
- Specify the NOTSORTED or DESCENDING option in the BY statement for the FREQ procedure. The NOTSORTED option does not mean that the data are unsorted but rather that the data are arranged in groups (according to values of the BY variables) and that these groups are not necessarily in alphabetical or increasing numeric order.
- Create an index on the BY variables by using the DATASETS procedure (in Base SAS software).

For more information about BY-group processing, see the discussion in *SAS Language Reference: Concepts*. For more information about the DATASETS procedure, see the discussion in the *Base SAS Procedures Guide*.

EXACT Statement

EXACT *statistic-options* < / *computation-options* > ;

The EXACT statement requests exact tests and confidence limits for selected statistics. The *statistic-options* identify which statistics to compute, and the *computation-options* specify options for computing exact statistics. See the section “[Exact Statistics](#)” on page 2601 for details.

NOTE: PROC FREQ computes exact tests by using fast and efficient algorithms that are superior to direct enumeration. Exact tests are appropriate when a data set is small, sparse, skewed, or heavily tied. For some large problems, computation of exact tests might require a considerable amount of time and memory. Consider using asymptotic tests for such problems. Alternatively, when asymptotic methods might not be sufficient for such large problems, consider using Monte Carlo estimation of exact *p*-values. You can request Monte Carlo estimation by specifying the [MC computation-option](#) in the EXACT statement. See the section “[Computational Resources](#)” on page 2604 for more information.

Statistic Options

The *statistic-options* specify which exact tests and confidence limits to compute. [Table 38.6](#) lists the available *statistic-options* and the exact statistics that are computed. Descriptions of the *statistic-options* follow the table in alphabetical order.

For one-way tables, exact *p*-values are available for binomial proportion tests, the chi-square goodness-of-fit test, and the likelihood ratio chi-square test. Exact (Clopper-Pearson) confidence limits are available for the binomial proportion.

For two-way tables, exact *p*-values are available for the following tests: Pearson chi-square test, likelihood ratio chi-square test, Mantel-Haenszel chi-square test, Fisher’s exact test, Jonckheere-Terpstra test, and Cochran-Armitage test for trend. Exact *p*-values are also available for tests of the following statistics: Pearson correlation coefficient, Spearman correlation coefficient, Kendall’s tau-*b*, Stuart’s tau-*c*, Somers’ $D(C|R)$, Somers’ $D(R|C)$, simple kappa coefficient, and weighted kappa coefficient.

For 2×2 tables, PROC FREQ provides the exact McNemar’s test, exact confidence limits for the odds ratio, and Barnard’s unconditional exact test for the risk (proportion) difference. PROC FREQ also provides exact

unconditional confidence limits for the risk difference and for the relative risk (ratio of proportions). For stratified 2×2 tables, PROC FREQ provides Zelen's exact test for equal odds ratios, exact confidence limits for the common odds ratio, and an exact test for the common odds ratio.

Most of the *statistic-option* names listed in Table 38.6 are identical to the corresponding option names in the TABLES and OUTPUT statements. You can request exact computations for groups of statistics by using *statistic-options* that are identical to the TABLES statement options CHISQ, MEASURES, and AGREE. For example, when you specify the CHISQ *statistic-option* in the EXACT statement, PROC FREQ computes exact p -values for the Pearson chi-square, likelihood ratio chi-square, and Mantel-Haenszel chi-square tests for two-way tables. You can request an exact test for an individual statistic by specifying the corresponding *statistic-option* from the list in Table 38.6.

Using the EXACT Statement with the TABLES Statement

You must use a TABLES statement with the EXACT statement. If you use only one TABLES statement, you do not need to specify the same options in both the TABLES and EXACT statements; when you specify a *statistic-option* in the EXACT statement, PROC FREQ automatically invokes the corresponding TABLES statement option. However, when you use an EXACT statement with multiple TABLES statements, you must specify options in the TABLES statements to request statistics. PROC FREQ then provides exact tests or confidence limits for those statistics that you also specify in the EXACT statement.

Table 38.6 EXACT Statement Statistic Options

Statistic Option	Exact Statistics
AGREE	McNemar's test (for 2×2 tables), simple kappa test, weighted kappa test
BARNARD	Barnard's test (for 2×2 tables)
BINOMIAL	Binomial proportion tests for one-way tables
CHISQ	Chi-square goodness-of-fit test for one-way tables; Pearson chi-square, likelihood ratio chi-square, and Mantel-Haenszel chi-square tests for two-way tables
COMOR	Confidence limits for the common odds ratio, common odds ratio test (for $h \times 2 \times 2$ tables)
EQOR ZELN	Zelen's test for equal odds ratios (for $h \times 2 \times 2$ tables)
FISHER	Fisher's exact test
JT	Jonckheere-Terpstra test
KAPPA	Test for the simple kappa coefficient
KENTB	Test for Kendall's tau- b
LRCHI	Likelihood ratio chi-square test (one-way or two-way tables)
MCNEM	McNemar's test (for 2×2 tables)
MEASURES	Tests for the Pearson correlation and Spearman correlation, confidence limits for the odds ratio (for 2×2 tables)
MHCHI	Mantel-Haenszel chi-square test
OR	Confidence limits for the odds ratio (for 2×2 tables)
PCHI	Pearson chi-square test
PCORR	Test for the Pearson correlation coefficient
RELRISK	Confidence limits for the relative risk (for 2×2 tables)
RISKDIFF	Confidence limits for the proportion difference (for 2×2 tables)
SCORR	Test for the Spearman correlation coefficient

Table 38.6 *continued*

Statistic Option	Exact Statistics
SMDCR	Test for Somers' $D(C R)$
SMDRC	Test for Somers' $D(R C)$
STUTC	Test for Stuart's tau- c
TREND	Cochran-Armitage test for trend
WTKAP	Test for the weighted kappa coefficient

You can specify the following *statistic-options* in the EXACT statement.

AGREE

requests the exact McNemar's test and exact tests for the simple and weighted kappa coefficients. See the sections "[Tests and Measures of Agreement](#)" on page 2588 and "[Exact Statistics](#)" on page 2601 for details.

The **AGREE** option in the TABLES statement provides the asymptotic McNemar's test and the kappa estimates, standard errors, and confidence limits. The **AGREE** option in the TEST statement provides asymptotic tests for the kappa coefficients.

Kappa coefficients are defined only for square tables, where the number of rows equals the number of columns. McNemar's test is available for 2×2 tables.

BARNARD

requests Barnard's exact unconditional test for the risk (proportion) difference for 2×2 tables. See the section "[Barnard's Unconditional Exact Test](#)" on page 2581 for details. The **RISKDIFF** option in the TABLES statement provides risk difference estimates, confidence limits, and asymptotic tests. See the section "[Risks and Risk Differences](#)" on page 2572 for more information.

BINOMIAL

requests exact tests for the binomial proportion (for one-way tables). See the section "[Binomial Tests](#)" on page 2568 for details. The **BINOMIAL** option in the TABLES statement provides confidence limits and asymptotic tests for the binomial proportion. See the section "[Binomial Proportion](#)" on page 2565 for more information.

CHISQ

requests exact tests for the Pearson chi-square, likelihood ratio chi-square, and Mantel-Haenszel chi-square for two-way tables. See the section "[Chi-Square Tests and Statistics](#)" on page 2550 for details. The **CHISQ** option in the TABLES statement provides asymptotic tests for these statistics.

For one-way tables, the **CHISQ** option requests the exact chi-square goodness-of-fit test. See the section "[Chi-Square Test for One-Way Tables](#)" on page 2551 for details.

COMOR

requests an exact test and exact confidence limits for the common odds ratio for multiway 2×2 tables. See the section "[Exact Confidence Limits for the Common Odds Ratio](#)" on page 2599 for details. The **CMH** option in the TABLES statement provides Mantel-Haenszel and logit estimates and asymptotic confidence limits for the common odds ratio.

EQOR

requests Zelen's exact test for equal odds ratios, which is available for multiway 2×2 tables. See the section "[Zelen's Exact Test for Equal Odds Ratios](#)" on page 2598 for details. The **CMH** option in the **TABLES** statement provides the asymptotic Breslow-Day test for homogeneity of odds ratios.

FISHER

request Fisher's exact test. See the sections "[Fisher's Exact Test](#)" on page 2554 and "[Exact Statistics](#)" on page 2601 for details. For 2×2 tables, the **CHISQ** option in the **TABLES** statement provides Fisher's exact test. For general $R \times C$ tables, Fisher's exact test is also known as the Freeman-Halton test.

JT

requests the exact Jonckheere-Terpstra test. See the sections "[Jonckheere-Terpstra Test](#)" on page 2586 and "[Exact Statistics](#)" on page 2601 for details. The **JT** option in the **TABLES** statement provides the asymptotic Jonckheere-Terpstra test.

KAPPA

requests an exact test for the simple kappa coefficient. See the sections "[Simple Kappa Coefficient](#)" on page 2589 and "[Exact Statistics](#)" on page 2601 for details. The **AGREE** option in the **TABLES** statement provides the simple kappa estimate, standard error, and confidence limits. The **KAPPA** option in the **TEST** statement provides the asymptotic test for the simple kappa coefficient.

Kappa coefficients are defined only for square tables, where the number of rows equals the number of columns. PROC FREQ does not compute kappa coefficients for tables that are not square.

KENTB

requests an exact test for Kendall's tau-*b*. See the sections "[Kendall's Tau-b](#)" on page 2558 and "[Exact Statistics](#)" on page 2601 for details. The **MEASURES** option in the **TABLES** statement provides the Kendall's tau-*b* estimate and standard error. The **KENTB** option in the **TEST** statement provides an asymptotic test for Kendall's tau-*b*.

LRCHI

requests an exact test for the likelihood ratio chi-square for two-way tables. See the sections "[Likelihood Ratio Chi-Square Test](#)" on page 2553 and "[Exact Statistics](#)" on page 2601 for details. For one-way tables, the **LRCHI** option requests an exact likelihood ratio goodness-of-fit test. See the section "[Likelihood Ratio Chi-Square Test for One-Way Tables](#)" on page 2552 for details.

The **CHISQ** option in the **TABLES** statement provides asymptotic likelihood ratio chi-square tests.

MCNEM

requests the exact McNemar's test. See the sections "[McNemar's Test](#)" on page 2588 and "[Exact Statistics](#)" on page 2601 for details. The **AGREE** option in the **TABLES** statement provides the asymptotic McNemar's test.

MEASURES

requests exact tests for the Pearson and Spearman correlations. See the sections "[Pearson Correlation Coefficient](#)" on page 2559, "[Spearman Rank Correlation Coefficient](#)" on page 2560, and "[Exact Statistics](#)" on page 2601 for details. The **PCORR** and **SCORR** options in the **TEST** statement provide asymptotic tests for the Pearson and Spearman correlations, respectively.

The **MEASURES** option also requests exact confidence limits for the odds ratio for 2×2 tables. See the sections "[Odds Ratio \(Case-Control Studies\)](#)" on page 2582 and "[Exact Confidence Limits for the Odds Ratio](#)" on page 2583 for details. The **MEASURES** and **RELRISK** options in the **TABLES** statement provide asymptotic confidence limits for the odds ratio.

MHCHI

requests an exact test for the Mantel-Haenszel chi-square. See the sections “[Mantel-Haenszel Chi-Square Test](#)” on page 2553 and “[Exact Statistics](#)” on page 2601 for details. The **CHISQ** option in the **TABLES** statement provides the asymptotic Mantel-Haenszel chi-square test.

OR

requests exact confidence limits for the odds ratio for 2×2 tables. See the sections “[Odds Ratio \(Case-Control Studies\)](#)” on page 2582 and “[Exact Confidence Limits for the Odds Ratio](#)” on page 2583 for details. The **MEASURES** and **RELRISK** options in the **TABLES** statement provide asymptotic confidence limits for the odds ratio.

PCHI

requests an exact test for the Pearson chi-square for two-way tables. See the sections “[Pearson Chi-Square Test for Two-Way Tables](#)” on page 2551 and “[Exact Statistics](#)” on page 2601 for details. For one-way tables, the **PCHI** option requests an exact chi-square goodness-of-fit test. See the section “[Chi-Square Test for One-Way Tables](#)” on page 2551 for details. The **CHISQ** option in the **TABLES** statement provides asymptotic chi-square tests.

PCORR

requests an exact test for the Pearson correlation coefficient. See the sections “[Pearson Correlation Coefficient](#)” on page 2559 and “[Exact Statistics](#)” on page 2601 for details. The **MEASURES** option in the **TABLES** statement provides the Pearson correlation estimate and standard error. The **PCORR** option in the **TEST** statement provides an asymptotic test for the Pearson correlation.

RELRISK <(options)>

requests exact unconditional confidence limits for the relative risk for 2×2 tables. **PROC FREQ** computes the confidence limits by inverting two separate one-sided exact tests (Santner and Snell 1980). By default, this computation uses the unstandardized relative risk as the test statistic. If you specify the **RELRISK(METHOD=SCORE)** option, the computation uses the score statistic (Chan and Zhang 1999). See the section “[Exact Unconditional Confidence Limits for the Relative Risk](#)” on page 2584 for more information.

You can set the confidence level by specifying the **ALPHA=** option in the **TABLES** statement. The default of **ALPHA=0.5** produces 95% confidence limits.

The **RELRISK** option in the **TABLES** statement provides asymptotic confidence limits for the relative risk. See the section “[Relative Risks \(Cohort Studies\)](#)” on page 2583 for more information.

You can specify the following *options* in parentheses after the **RELRISK** *statistic-option*:

COLUMN=1 | 2 | BOTH

specifies the 2×2 table column for which to compute the relative risk. The default is **COLUMN=1**, which provides exact confidence limits for the column 1 relative risk. If you specify **COLUMN=BOTH**, **PROC FREQ** provides exact confidence limits for both column 1 and column 2 relative risks.

METHOD=SCORE

requests exact unconditional confidence limits that are based on the score statistic (Chan and Zhang 1999). See the section “[Exact Unconditional Confidence Limits for the Relative Risk](#)” on page 2584 for more information. If you do not specify **METHOD=SCORE**, by default the exact confidence limit computations are based on the unstandardized relative risk.

RISKDIFF <(options)>

requests exact unconditional confidence limits for the risk difference for 2×2 tables. PROC FREQ computes the confidence limits by inverting two separate one-sided exact tests (Santner and Snell 1980). By default, this computation uses the unstandardized risk difference as the test statistic. If you specify the **RISKDIFF(METHOD=SCORE)** option, the computation uses the score statistic (Chan and Zhang 1999). See the section “Exact Unconditional Confidence Limits for the Risk Difference” on page 2580 for more information.

You can set the confidence level by specifying the **ALPHA=** option in the **TABLES** statement. The default of **ALPHA=0.5** produces 95% confidence limits.

The **RISKDIFF** option in the **TABLES** statement provides asymptotic confidence limits for the risk difference, including Wald, Newcombe, and Miettinen-Nurminen confidence limits. See the section “Risk Difference Confidence Limits” on page 2574 for more information.

You can specify the following *options* in parentheses after the **RISKDIFF** *statistic-option*:

COLUMN=1 | 2 | BOTH

specifies the 2×2 table column for which to compute the risk difference. The default is **COLUMN=BOTH**, which provides exact confidence limits for both column 1 and column 2 risk differences.

METHOD=SCORE

requests exact unconditional confidence limits that are based on the score statistic (Chan and Zhang 1999). See the section “Exact Unconditional Confidence Limits for the Risk Difference” on page 2580 for more information. If you do not specify **METHOD=SCORE**, by default the exact confidence limit computations are based on the unstandardized risk difference.

SCORR

requests an exact test for the Spearman correlation coefficient. See the sections “Spearman Rank Correlation Coefficient” on page 2560 and “Exact Statistics” on page 2601 for details. The **MEASURES** option in the **TABLES** statement provides the Spearman correlation estimate and standard error. The **SCORR** option in the **TEST** statement provides an asymptotic test for the Spearman correlation.

SMDCR

requests an exact test for Somers’ $D(C|R)$. See the sections “Somers’ D ” on page 2559 and “Exact Statistics” on page 2601 for details. The **MEASURES** option in the **TABLES** statement provides Somers’ $D(C|R)$ estimate and the standard error. The **SMDCR** option in the **TEST** statement provides an asymptotic test for Somers’ $D(C|R)$.

SMDRC

requests an exact test for Somers’ $D(R|C)$. See the sections “Somers’ D ” on page 2559 and “Exact Statistics” on page 2601 for details. The **MEASURES** option in the **TABLES** statement provides Somers’ $D(R|C)$ estimate and the standard error. The **SMDRC** option in the **TEST** statement requests an asymptotic test for Somers’ $D(R|C)$.

STUTC

requests an exact test for Stuart’s tau- c . See the sections “Stuart’s Tau- c ” on page 2558 and “Exact Statistics” on page 2601 for details. The **MEASURES** option in the **TABLES** statement provides Stuart’s tau- c estimate and the standard error. The **STUTC** option in the **TEST** statement provides an asymptotic test for Stuart’s tau- c .

TREND

requests the exact Cochran-Armitage test for trend. See the sections “[Cochran-Armitage Test for Trend](#)” on page 2585 and “[Exact Statistics](#)” on page 2601 for details. The **TREND** option in the TABLES statement provides the asymptotic Cochran-Armitage trend test. The trend test is available for tables of dimensions $2 \times C$ or $R \times 2$.

WTKAP

requests an exact test for the weighted kappa coefficient. See the sections “[Weighted Kappa Coefficient](#)” on page 2590 and “[Exact Statistics](#)” on page 2601 for details. The **AGREE** option in the TABLES statement provides the weighted kappa estimate, standard error, and confidence limits. The **WTKAP** option in the TEST statement provides the asymptotic test for the weighted kappa coefficient.

Kappa coefficients are defined only for square tables, where the number of rows equals the number of columns. PROC FREQ does not compute kappa coefficients for tables that are not square. For 2×2 tables, the weighted kappa coefficient equals the simple kappa coefficient, and PROC FREQ does not present separate results for the weighted kappa coefficient.

Computation Options

The *computation-options* specify options for computing exact statistics. You can specify the following *computation-options* in the EXACT statement after a slash (/).

ALPHA= α

specifies the level of the confidence limits for Monte Carlo p -value estimates. The value of α must be between 0 and 1, and the default is 0.01. A confidence level of α produces $100(1 - \alpha)\%$ confidence limits. The default of ALPHA=.01 produces 99% confidence limits for the Monte Carlo estimates.

The ALPHA= option invokes the **MC** option.

MAXTIME=value

specifies the maximum clock time (in seconds) that PROC FREQ can use to compute an exact p -value. If the procedure does not complete the computation within the specified time, the computation terminates. The value of MAXTIME= must be a positive number. The MAXTIME= option is valid for Monte Carlo estimation of exact p -values as well as for direct exact p -value computation. See the section “[Computational Resources](#)” on page 2604 for more information.

MC

6 requests Monte Carlo estimation of exact p -values instead of direct exact p -value computation. Monte Carlo estimation can be useful for large problems that require a considerable amount of time and memory for exact computations but for which asymptotic approximations might not be sufficient. See the section “[Monte Carlo Estimation](#)” on page 2604 for more information.

The MC option is available for all EXACT *statistic-options* except the BINOMIAL option and the following options that apply only to 2×2 or $h \times 2 \times 2$ tables: BARNARD, COMOR, EQOR, MCNEM, OR, RELRISK, and RISKDIFF. PROC FREQ computes only exact tests or confidence limits for these statistics.

The ALPHA=, N=, and SEED= options also invoke the MC option.

N=*n*

specifies the number of samples for Monte Carlo estimation. The value of *n* must be a positive integer, and the default is 10,000. Larger values of *n* produce more precise estimates of exact *p*-values. Because larger values of *n* generate more samples, the computation time increases.

The N= option invokes the [MC](#) option.

POINT

requests exact point probabilities for the test statistics.

The POINT option is available for all EXACT statement *statistic-options* except BARNARD, OR, RELRISK, and RISKDIFF. The POINT option is not available with the [MC](#) option.

SEED=*number*

specifies the initial seed for random number generation for Monte Carlo estimation. The value of the SEED= option must be an integer. If you do not specify the SEED= option or if the SEED= value is negative or zero, PROC FREQ uses the time of day from the computer's clock to obtain the initial seed.

The SEED= option invokes the [MC](#) option.

OUTPUT Statement

OUTPUT < **OUT=** *SAS-data-set* > *output-options* ;

The OUTPUT statement creates a SAS data set that contains statistics that are computed by PROC FREQ. [Table 38.7](#) lists the statistics that can be stored in the output data set. You identify which statistics to include by specifying *output-options*.

You must use a [TABLES](#) statement with the OUTPUT statement. The OUTPUT statement stores statistics for only one table request. If you use multiple TABLES statements, the contents of the output data set correspond to the last TABLES statement. If you use multiple table requests in a single TABLES statement, the contents of the output data set correspond to the last table request. Only one OUTPUT statement is allowed in a single invocation of the procedure.

For a one-way or two-way table, the output data set contains one observation that stores the requested statistics for the table. For a multiway table, the output data set contains an observation for each two-way table (stratum) of the multiway crosstabulation. If you request summary statistics for the multiway table, the output data set also contains an observation that stores the across-strata summary statistics. If you use a [BY](#) statement, the output data set contains an observation or set of observations for each BY group. For more information about the contents of the output data set, see the section “[Contents of the OUTPUT Statement Output Data Set](#)” on page 2608.

The output data set that is created by the OUTPUT statement is not the same as the output data set that is created by the [OUT=](#) option in the [TABLES](#) statement. The OUTPUT statement creates a data set that contains statistics (such as the Pearson chi-square and its *p*-value), and the [OUT=](#) option in the TABLES statement creates a data set that contains frequency table counts and percentages. See the section “[Output Data Sets](#)” on page 2606 for more information.

As an alternative to the OUTPUT statement, you can use the Output Delivery System (ODS) to store statistics that PROC FREQ computes. ODS can create a SAS data set from any table that PROC FREQ produces. See the section “[ODS Table Names](#)” on page 2617 for more information.

You can specify the following *options* in the OUTPUT statement:

OUT=SAS-data-set

specifies the name of the output data set. When you use an OUTPUT statement but do not use the OUT= option, PROC FREQ creates a data set and names it by using the *DATA**n* convention.

output-options

specify the statistics to include in the output data set. Table 38.7 lists the *output-options* that are available in the OUTPUT statement, together with the TABLES statement options that are required to produce the statistics. Descriptions of the *output-options* follow the table in alphabetical order.

You can specify *output-options* to request individual statistics, or you can request groups of statistics by using *output-options* that are identical to the group options in the TABLES statement (for example, the CHISQ, MEASURES, CMH, AGREE, and ALL options).

When you specify an *output-option*, the output data set includes statistics from the corresponding analysis. In addition to the estimate or test statistic, the output data set includes associated values such as standard errors, confidence limits, *p*-values, and degrees of freedom. See the section “Contents of the OUTPUT Statement Output Data Set” on page 2608 for details.

To store a statistic in the output data set, you must also request computation of that statistic with the appropriate TABLES, EXACT, or TEST statement option. For example, the PCHI *output-option* includes the Pearson chi-square in the output data set. You must also request computation of the Pearson chi-square by specifying the CHISQ option in the TABLES statement. Or, if you use only one TABLES statement, you can request computation of the Pearson chi-square by specifying the PCHI or CHISQ option in the EXACT statement. Table 38.7 lists the TABLES statement options that are required to produce the OUTPUT data set statistics.

Table 38.7 OUTPUT Statement Output Options

Output Option	Output Data Set Statistics	Required TABLES Statement Option
AGREE	McNemar's test (2×2 tables), Bowker's test, simple and weighted kappas; for multiple strata, overall simple and weighted kappas, tests for equal kappas, and Cochran's <i>Q</i> ($h \times 2 \times 2$ tables)	AGREE
AJCHI	Continuity-adjusted chi-square (2×2 tables)	CHISQ
ALL	CHISQ, MEASURES, and CMH statistics; N (number of nonmissing observations)	ALL
BDCHI BD	Breslow-Day test ($h \times 2 \times 2$ tables)	CMH, CMH1, or CMH2
BINOMIAL	Binomial statistics (one-way tables)	BINOMIAL
CHISQ	For one-way tables, goodness-of-fit test; for two-way tables, Pearson, likelihood ratio, continuity-adjusted, and Mantel-Haenszel chi-squares, Fisher's exact test (2×2 tables), phi and contingency coefficients, Cramér's <i>V</i>	CHISQ
CMH	Cochran-Mantel-Haenszel (CMH) correlation, row mean scores (ANOVA), and general association statistics; for 2×2 tables, logit and Mantel-Haenszel common odds ratios and relative risks, Breslow-Day test	CMH

Table 38.7 continued

Output Option	Output Data Set Statistics	Required TABLES Statement Option
CMH1	CMH statistics, except row mean scores (ANOVA) and general association statistics	CMH or CMH1
CMH2	CMH statistics, except general association statistic	CMH or CMH2
CMHCOR	CMH correlation statistic	CMH, CMH1, or CMH2
CMHGA	CMH general association statistic	CMH
CMHRMS	CMH row mean scores (ANOVA) statistic	CMH or CMH2
COCHQ	Cochran's Q ($h \times 2 \times 2$ tables)	AGREE
CONTGY	Contingency coefficient	CHISQ
CRAMV	Cramér's V	CHISQ
EQKAP	Test for equal simple kappas	AGREE
EQOR ZELEN	Zelen's test for equal odds ratios ($h \times 2 \times 2$ tables)	CMH and EXACT EQOR
EQWKP	Test for equal weighted kappas	AGREE
FISHER	Fisher's exact test	CHISQ or FISHER ¹
GAILSIMON GS	Gail-Simon test	CMH(GAILSIMON)
GAMMA	Gamma	MEASURES
JT	Jonckheere-Terpstra test	JT
KAPPA	Simple kappa coefficient	AGREE
KENTB TAUB	Kendall's tau- b	MEASURES
LAMCR	Lambda asymmetric ($C R$)	MEASURES
LAMDAS	Lambda symmetric	MEASURES
LAMRC	Lambda asymmetric ($R C$)	MEASURES
LGOR	Logit common odds ratio	CMH, CMH1, or CMH2
LGRRC1	Logit common relative risk, column 1	CMH, CMH1, or CMH2
LGRRC2	Logit common relative risk, column 2	CMH, CMH1, or CMH2
LRCHI	Likelihood ratio chi-square	CHISQ
MCNEM	McNemar's test (2×2 tables)	AGREE
MEASURES	Gamma, Kendall's tau- b , Stuart's tau- c , Somers' $D(C R)$ and $D(R C)$, Pearson and Spearman correlations, lambda asymmetric ($C R$) and ($R C$), lambda symmetric, uncertainty coefficients ($C R$) and ($R C$), symmetric uncertainty coefficient; odds ratio and relative risks (2×2 tables)	MEASURES
MHCHI	Mantel-Haenszel chi-square	CHISQ
MHOR COMOR	Mantel-Haenszel common odds ratio	CMH, CMH1, or CMH2
MHRRC1	Mantel-Haenszel common relative risk, column 1	CMH, CMH1, or CMH2
MHRRC2	Mantel-Haenszel common relative risk, column 2	CMH, CMH1, or CMH2
N	Number of nonmissing observations	
NMISS	Number of missing observations	
OR	Odds ratio (2×2 tables)	MEASURES or RELRISK

¹CHISQ computes Fisher's exact test for 2×2 tables. Use the FISHER option to compute Fisher's exact test for general $r \times c$ tables.

Table 38.7 continued

Output Option	Output Data Set Statistics	Required TABLES Statement Option
PCHI	Chi-square goodness-of-fit test (one-way tables), Pearson chi-square (two-way tables)	CHISQ
PCORR	Pearson correlation coefficient	MEASURES
PHI	Phi coefficient	CHISQ
PLCORR	Polychoric correlation coefficient	PLCORR
RDIF1	Column 1 risk difference (row 1 – row 2)	RISKDIFF
RDIF2	Column 2 risk difference (row 1 – row 2)	RISKDIFF
RELRISK	Odds ratio and relative risks (2×2 tables)	MEASURES or RELRISK
RISKDIFF	Risks and risk differences (2×2 tables)	RISKDIFF
RISKDIFF1	Risks and risk difference, column 1	RISKDIFF
RISKDIFF2	Risks and risk difference, column 2	RISKDIFF
RRC1 RELRISK1	Relative risk, column 1	MEASURES or RELRISK
RRC2 RELRISK2	Relative risk, column 2	MEASURES or RELRISK
RSK1 RISK1	Column 1 overall risk	RISKDIFF
RSK11 RISK11	Column 1 risk for row 1	RISKDIFF
RSK12 RISK12	Column 2 risk for row 1	RISKDIFF
RSK2 RISK2	Column 2 overall risk	RISKDIFF
RSK21 RISK21	Column 1 risk for row 2	RISKDIFF
RSK22 RISK22	Column 2 risk for row 2	RISKDIFF
SCORR	Spearman correlation coefficient	MEASURES
SMDCR	Somers' $D(C R)$	MEASURES
SMDRC	Somers' $D(R C)$	MEASURES
STUTC TAUC	Stuart's tau- c	MEASURES
TREND	Cochran-Armitage test for trend	TREND
TSYMM BOWKER	Bowker's test of symmetry	AGREE
U	Symmetric uncertainty coefficient	MEASURES
UCR	Uncertainty coefficient ($C R$)	MEASURES
URC	Uncertainty coefficient ($R C$)	MEASURES
WTKAP	Weighted kappa coefficient	AGREE

You can specify the following *output-options* in the OUTPUT statement.

AGREE

includes the following tests and measures of agreement in the output data set: McNemar's test (for 2×2 tables), Bowker's test of symmetry, the simple kappa coefficient, and the weighted kappa coefficient. For multiway tables, the AGREE option also includes the following statistics in the output data set: overall simple and weighted kappa coefficients, tests for equal simple and weighted kappa coefficients, and Cochran's Q test.

The AGREE option in the TABLES statement requests computation of tests and measures of agreement. See the section “Tests and Measures of Agreement” on page 2588 for details about these statistics.

AGREE statistics are computed only for square tables, where the number of rows equals the number of columns. PROC FREQ provides Bowker's test of symmetry and weighted kappa coefficients only

for tables larger than 2×2 . (For 2×2 tables, Bowker's test is identical to McNemar's test, and the weighted kappa coefficient equals the simple kappa coefficient.) Cochran's Q is available for multiway 2×2 tables.

AJCHI

includes the continuity-adjusted chi-square in the output data set. The continuity-adjusted chi-square is available for 2×2 tables and is provided by the **CHISQ** option in the **TABLES** statement. See the section “Continuity-Adjusted Chi-Square Test” on page 2553 for details.

ALL

includes all statistics that are requested by the **CHISQ**, **MEASURES**, and **CMH** *output-options* in the output data set. **ALL** also includes the number of nonmissing observations, which you can request individually by specifying the **N** *output-option*.

BDCHI | BD

includes the Breslow-Day test in the output data set. The Breslow-Day test for homogeneity of odds ratios is computed for multiway 2×2 tables and is provided by the **CMH**, **CMH1**, and **CMH2** options in the **TABLES** statement. See the section “Breslow-Day Test for Homogeneity of the Odds Ratios” on page 2598 for details.

BINOMIAL

includes the binomial proportion estimate, confidence limits, and tests in the output data set. The **BINOMIAL** option in the **TABLES** statement requests computation of binomial statistics, which are available for one-way tables. See the section “Binomial Proportion” on page 2565 for details.

CHISQ

includes the following chi-square tests and measures in the output data set for two-way tables: Pearson chi-square, likelihood ratio chi-square, Mantel-Haenszel chi-square, phi coefficient, contingency coefficient, and Cramér's V . For 2×2 tables, **CHISQ** also includes Fisher's exact test and the continuity-adjusted chi-square in the output data set. See the section “Chi-Square Tests and Statistics” on page 2550 for details. For one-way tables, **CHISQ** includes the chi-square goodness-of-fit test in the output data set. See the section “Chi-Square Test for One-Way Tables” on page 2551 for details. The **CHISQ** option in the **TABLES** statement requests computation of these statistics.

If you specify the **CHISQ(WARN=OUTPUT)** option in the **TABLES** statement, the **CHISQ** option also includes the variable **WARN_PCHI** in the output data set. This variable indicates the validity warning for the asymptotic Pearson chi-square test.

CMH

includes the following Cochran-Mantel-Haenszel statistics in the output data set: correlation, row mean scores (ANOVA), and general association. For 2×2 tables, the **CMH** option also includes the Mantel-Haenszel and logit estimates of the common odds ratio and relative risks. For multiway (stratified) 2×2 tables, the **CMH** option includes the Breslow-Day test for homogeneity of odds ratios. The **CMH** option in the **TABLES** statement requests computation of these statistics. See the section “Cochran-Mantel-Haenszel Statistics” on page 2592 for details.

If you specify the **CMH(MF)** option in the **TABLES** statement, the **CMH** option includes the Mantel-Fleiss analysis in the output data set. The variables **MF_CMH** and **WARN_CMH** contain the Mantel-Fleiss criterion and the warning indicator, respectively.

CMH1

includes the [CMH](#) statistics in the output data set, with the exception of the row mean scores (ANOVA) statistic and the general association statistic. The [CMH1](#) option in the [TABLES](#) statement requests computation of these statistics. See the section “[Cochran-Mantel-Haenszel Statistics](#)” on page 2592 for details.

CMH2

includes the [CMH](#) statistics in the output data set, with the exception of the general association statistic. The [CMH2](#) option in the [TABLES](#) statement requests computation of these statistics. See the section “[Cochran-Mantel-Haenszel Statistics](#)” on page 2592 for details.

CMHCOR

includes the Cochran-Mantel-Haenszel correlation statistic in the output data set. The [CMH](#) option in the [TABLES](#) statement requests computation of this statistic. See the section “[Correlation Statistic](#)” on page 2594 for details.

CMHGA

includes the Cochran-Mantel-Haenszel general association statistic in the output data set. The [CMH](#) option in the [TABLES](#) statement requests computation of this statistic. See the section “[General Association Statistic](#)” on page 2595 for details.

CMHRMS

includes the Cochran-Mantel-Haenszel row mean scores (ANOVA) statistic in the output data set. The [CMH](#) option in the [TABLES](#) statement requests computation of this statistic. See the section “[ANOVA \(Row Mean Scores\) Statistic](#)” on page 2594 for details.

COCHQ

includes Cochran’s Q test in the output data set. The [AGREE](#) option in the [TABLES](#) statement requests computation of this test, which is available for multiway 2×2 tables. See the section “[Cochran’s \$Q\$ Test](#)” on page 2592 for details.

CONTGY

includes the contingency coefficient in the output data set. The [CHISQ](#) option in the [TABLES](#) statement requests computation of the contingency coefficient. See the section “[Contingency Coefficient](#)” on page 2555 for details.

CRAMV

includes Cramér’s V in the output data set. The [CHISQ](#) option in the [TABLES](#) statement requests computation of Cramér’s V . See the section “[Cramér’s \$V\$ ” on page 2555 for details.](#)

EQKAP

includes the test for equal simple kappa coefficients in the output data set. The [AGREE](#) option in the [TABLES](#) statement requests computation of this test, which is available for multiway, square ($h \times r \times r$) tables. See the section “[Tests for Equal Kappa Coefficients](#)” on page 2591 for details.

EQOR | ZELEN

includes Zelen’s exact test for equal odds ratios in the output data set. The [EQOR](#) option in the [EXACT](#) statement requests computation of this test, which is available for multiway 2×2 tables. See the section “[Zelen’s Exact Test for Equal Odds Ratios](#)” on page 2598 for details.

EQWKP

includes the test for equal weighted kappa coefficients in the output data set. The [AGREE](#) option in the [TABLES](#) statement requests computation of this test. The test for equal weighted kappas is available for multiway, square ($h \times r \times r$) tables where $r > 2$. See the section “[Tests for Equal Kappa Coefficients](#)” on page 2591 for details.

FISHER

includes Fisher’s exact test in the output data set. For 2×2 tables, the [CHISQ](#) option in the [TABLES](#) statement provides Fisher’s exact test. For tables larger than 2×2 , the [FISHER](#) option in the [EXACT](#) statement provides Fisher’s exact test. See the section “[Fisher’s Exact Test](#)” on page 2554 for details.

GAILSIMON | GS

includes the Gail-Simon test for qualitative interaction in the output data set. The [CMH\(GAILSIMON\)](#) option in the [TABLES](#) statement requests computation of this test. See the section “[Gail-Simon Test for Qualitative Interactions](#)” on page 2601 for details.

GAMMA

includes the gamma statistic in the output data set. The [MEASURES](#) option in the [TABLES](#) statement requests computation of the gamma statistic. See the section “[Gamma](#)” on page 2557 for details.

JT

includes the Jonckheere-Terpstra test in the output data set. The [JT](#) option in the [TABLES](#) statement requests the Jonckheere-Terpstra test. See the section “[Jonckheere-Terpstra Test](#)” on page 2586 for details.

KAPPA

includes the simple kappa coefficient in the output data set. The [AGREE](#) option in the [TABLES](#) statement requests computation of kappa, which is available for square tables (where the number of rows equals the number of columns). For multiway square tables, the [KAPPA](#) option also includes the overall kappa coefficient in the output data set. See the sections “[Simple Kappa Coefficient](#)” on page 2589 and “[Overall Kappa Coefficient](#)” on page 2591 for details.

KENTB | TAUB

includes Kendall’s tau-*b* in the output data set. The [MEASURES](#) option in the [TABLES](#) statement requests computation of Kendall’s tau-*b*. See the section “[Kendall’s Tau-b](#)” on page 2558 for details.

LAMCR

includes the asymmetric lambda $\lambda(C|R)$ in the output data set. The [MEASURES](#) option in the [TABLES](#) statement requests computation of lambda. See the section “[Lambda \(Asymmetric\)](#)” on page 2563 for details.

LAMDAS

includes the symmetric lambda in the output data set. The [MEASURES](#) option in the [TABLES](#) statement requests computation of lambda. See the section “[Lambda \(Symmetric\)](#)” on page 2563 for details.

LAMRC

includes the asymmetric lambda $\lambda(R|C)$ in the output data set. The [MEASURES](#) option in the [TABLES](#) statement requests computation of lambda. See the section “[Lambda \(Asymmetric\)](#)” on page 2563 for details.

LGOR

includes the logit estimate of the common odds ratio in the output data set. The **CMH** option in the **TABLES** statement requests computation of this statistic, which is available for 2×2 tables. See the section “[Adjusted Odds Ratio and Relative Risk Estimates](#)” on page 2596 for details.

LGRRC1

includes the logit estimate of the common relative risk (column 1) in the output data set. The **CMH** option in the **TABLES** statement requests computation of this statistic, which is available for 2×2 tables. See the section “[Adjusted Odds Ratio and Relative Risk Estimates](#)” on page 2596 for details.

LGRRC2

includes the logit estimate of the common relative risk (column 2) in the output data set. The **CMH** option in the **TABLES** statement requests computation of this statistic, which is available for 2×2 tables. See the section “[Adjusted Odds Ratio and Relative Risk Estimates](#)” on page 2596 for details.

LRCHI

includes the likelihood ratio chi-square in the output data set. The **CHISQ** option in the **TABLES** statement requests computation of the likelihood ratio chi-square. See the section “[Likelihood Ratio Chi-Square Test](#)” on page 2553 for details.

MCNEM

includes McNemar’s test (for 2×2 tables) in the output data set. The **AGREE** option in the **TABLES** statement requests computation of McNemar’s test. See the section “[McNemar’s Test](#)” on page 2588 for details.

MEASURES

includes the following measures of association in the output data set: gamma, Kendall’s tau-*b*, Stuart’s tau-*c*, Somers’ $D(C|R)$, Somers’ $D(R|C)$, Pearson and Spearman correlation coefficients, lambda (symmetric and asymmetric), and uncertainty coefficients (symmetric and asymmetric). For 2×2 tables, the **MEASURES** option also includes the odds ratio, column 1 relative risk, and column 2 relative risk. The **MEASURES** option in the **TABLES** statement requests computation of these statistics. See the section “[Measures of Association](#)” on page 2556 for details.

MHCHI

includes the Mantel-Haenszel chi-square in the output data set. The **CHISQ** option in the **TABLES** statement requests computation of the Mantel-Haenszel chi-square. See the section “[Mantel-Haenszel Chi-Square Test](#)” on page 2553 for details.

MHOR | COMOR

includes the Mantel-Haenszel estimate of the common odds ratio in the output data set. The **CMH** option in the **TABLES** statement requests computation of this statistic, which is available for 2×2 tables. See the section “[Adjusted Odds Ratio and Relative Risk Estimates](#)” on page 2596 for details.

MHRC1

includes the Mantel-Haenszel estimate of the common relative risk (column 1) in the output data set. The **CMH** option in the **TABLES** statement requests computation of this statistic, which is available for 2×2 tables. See the section “[Adjusted Odds Ratio and Relative Risk Estimates](#)” on page 2596 for details.

MHRRC2

includes the Mantel-Haenszel estimate of the common relative risk (column 2) in the output data set. The **CMH** option in the **TABLES** statement requests computation of this statistic, which is available for 2×2 tables. See the section “[Adjusted Odds Ratio and Relative Risk Estimates](#)” on page 2596 for details.

N

includes the number of nonmissing observations in the output data set.

NMISS

includes the number of missing observations in the output data set. See the section “[Missing Values](#)” on page 2545 for details about how PROC FREQ handles missing values.

OR

includes the odds ratio (for 2×2 tables) in the output data set. The **MEASURES** and **REL RISK** options in the **TABLES** statement request this statistic. See the section “[Odds Ratio and Relative Risks for 2 x 2 Tables](#)” on page 2582 for details.

PCHI

includes the Pearson chi-square in the output data set for two-way tables. See the section “[Pearson Chi-Square Test for Two-Way Tables](#)” on page 2551 for details. For one-way tables, the **PCHI** option includes the chi-square goodness-of-fit test in the output data set. See the section “[Chi-Square Test for One-Way Tables](#)” on page 2551 for details. The **CHISQ** option in the **TABLES** statement requests computation of these statistics.

If you specify the **CHISQ(WARN=OUTPUT)** option in the **TABLES** statement, the **PCHI** option also includes the variable **WARN_PCHI** in the output data set. This variable indicates the validity warning for the asymptotic Pearson chi-square test.

PCORR

includes the Pearson correlation coefficient in the output data set. The **MEASURES** option in the **TABLES** statement requests computation of the Pearson correlation. See the section “[Pearson Correlation Coefficient](#)” on page 2559 for details.

PHI

includes the phi coefficient in the output data set. The **CHISQ** option in the **TABLES** statement requests computation of the phi coefficient. See the section “[Phi Coefficient](#)” on page 2555 for details.

PLCORR

includes the polychoric correlation coefficient in the output data set. For 2×2 tables, this statistic is known as the tetrachoric correlation coefficient. The **PLCORR** option in the **TABLES** statement requests computation of the polychoric correlation. See the section “[Polychoric Correlation](#)” on page 2562 for details.

RDIF1

includes the column 1 risk difference (row 1 – row 2) in the output data set. The **RISKDIFF** option in the **TABLES** statement requests computation of risks and risk differences, which are available for 2×2 tables. See the section “[Risks and Risk Differences](#)” on page 2572 for details.

RDIF2

includes the column 2 risk difference (row 1 – row 2) in the output data set. The **RISKDIFF** option in the **TABLES** statement requests computation of risks and risk differences, which are available for 2×2 tables. See the section “[Risks and Risk Differences](#)” on page 2572 for details.

RELRIK

includes the column 1 and column 2 relative risks (for 2×2 tables) in the output data set. The **MEASURES** and **RELRIK** options in the **TABLES** statement request these statistics. See the section “[Odds Ratio and Relative Risks for 2 x 2 Tables](#)” on page 2582 for details.

RISKDIFF

includes risks (binomial proportions) and risk differences for 2×2 tables in the output data set. These statistics include the row 1 risk, row 2 risk, total (overall) risk, and risk difference (row 1 – row 2) for column 1 and column 2. The **RISKDIFF** option in the **TABLES** statement requests computation of these statistics. See the section “[Risks and Risk Differences](#)” on page 2572 for details.

RISKDIFF1

includes column 1 risks (binomial proportions) and risk differences for 2×2 tables in the output data set. These statistics include the row 1 risk, row 2 risk, total (overall) risk, and risk difference (row 1 – row 2). The **RISKDIFF** option in the **TABLES** statement requests computation of these statistics. See the section “[Risks and Risk Differences](#)” on page 2572 for details.

RISKDIFF2

includes column 2 risks (binomial proportions) and risk differences for 2×2 tables in the output data set. These statistics include the row 1 risk, row 2 risk, total (overall) risk, and risk difference (row 1 – row 2). The **RISKDIFF** option in the **TABLES** statement requests computation of these statistics. See the section “[Risks and Risk Differences](#)” on page 2572 for details.

RRC1 | RELRIK1

includes the column 1 relative risk in the output data set. The **MEASURES** and **RELRIK** options in the **TABLES** statement request relative risks, which are available for 2×2 tables. See the section “[Odds Ratio and Relative Risks for 2 x 2 Tables](#)” on page 2582 for details.

RRC2 | RELRIK2

includes the column 2 relative risk in the output data set. The **MEASURES** and **RELRIK** options in the **TABLES** statement request relative risks, which are available for 2×2 tables. See the section “[Odds Ratio and Relative Risks for 2 x 2 Tables](#)” on page 2582 for details.

RSK1 | RISK1

includes the overall column 1 risk in the output data set. The **RISKDIFF** option in the **TABLES** statement requests computation of risks and risk differences, which are available for 2×2 tables. See the section “[Risks and Risk Differences](#)” on page 2572 for details.

RSK11 | RISK11

includes the column 1 risk for row 1 in the output data set. The **RISKDIFF** option in the **TABLES** statement requests computation of risks and risk differences, which are available for 2×2 tables. See the section “[Risks and Risk Differences](#)” on page 2572 for details.

RSK12 | RISK12

includes the column 2 risk for row 1 in the output data set. The **RISKDIFF** option in the **TABLES** statement requests computation of risks and risk differences, which are available for 2×2 tables. See the section “[Risks and Risk Differences](#)” on page 2572 for details.

RSK2 | RISK2

includes the overall column 2 risk in the output data set. The **RISKDIFF** option in the **TABLES** statement requests computation of risks and risk differences. See the section “[Risks and Risk Differences](#)” on page 2572 for details.

RSK21 | RISK21

includes the column 1 risk for row 2 in the output data set. The **RISKDIFF** option in the **TABLES** statement requests computation of risks and risk differences, which are available for 2×2 tables. See the section “[Risks and Risk Differences](#)” on page 2572 for details.

RSK22 | RISK22

includes the column 2 risk for row 2 in the output data set. The **RISKDIFF** option in the **TABLES** statement requests computation of risks and risk differences, which are available for 2×2 tables. See the section “[Risks and Risk Differences](#)” on page 2572 for details.

SCORR

includes the Spearman correlation coefficient in the output data set. The **MEASURES** option in the **TABLES** statement requests computation of the Spearman correlation. See the section “[Spearman Rank Correlation Coefficient](#)” on page 2560 for details.

SMDCR

includes Somers’ $D(C|R)$ in the output data set. The **MEASURES** option in the **TABLES** statement requests computation of Somers’ D . See the section “[Somers’ \$D\$](#) ” on page 2559 for details.

SMDRC

includes Somers’ $D(R|C)$ in the output data set. The **MEASURES** option in the **TABLES** statement requests computation of Somers’ D . See the section “[Somers’ \$D\$](#) ” on page 2559 for details.

STUTC | TAUC

includes Stuart’s tau- c in the output data set. The **MEASURES** option in the **TABLES** statement requests computation of tau- c . See the section “[Stuart’s Tau- \$c\$](#) ” on page 2558 for details.

TREND

includes the Cochran-Armitage test for trend in the output data set. The **TREND** option in the **TABLES** statement requests computation of the trend test. This test is available for tables of dimension $2 \times C$ or $R \times 2$. See the section “[Cochran-Armitage Test for Trend](#)” on page 2585 for details.

TSYMM | BOWKER

includes Bowker’s test of symmetry in the output data set. The **AGREE** option in the **TABLES** statement requests computation of Bowker’s test. See the section “[Bowker’s Test of Symmetry](#)” on page 2588 for details.

U

includes the uncertainty coefficient (symmetric) in the output data set. The **MEASURES** option in the **TABLES** statement requests computation of the uncertainty coefficient. See the section “[Uncertainty Coefficient \(Symmetric\)](#)” on page 2565 for details.

UCR

includes the asymmetric uncertainty coefficient $U(C|R)$ in the output data set. The **MEASURES** option in the **TABLES** statement requests computation of the uncertainty coefficient. See the section “[Uncertainty Coefficients \(Asymmetric\)](#)” on page 2564 for details.

URC

includes the asymmetric uncertainty coefficient $U(R|C)$ in the output data set. The **MEASURES** option in the **TABLES** statement requests computation of the uncertainty coefficient. See the section “[Uncertainty Coefficients \(Asymmetric\)](#)” on page 2564 for details.

WTKAP

includes the weighted kappa coefficient in the output data set. The **AGREE** option in the **TABLES** statement requests computation of weighted kappa, which is available for square tables larger than 2×2 . For multiway tables, the **WTKAP** option also includes the overall weighted kappa coefficient in the output data set. See the sections “[Weighted Kappa Coefficient](#)” on page 2590 and “[Overall Kappa Coefficient](#)” on page 2591 for details.

TABLES Statement

TABLES *requests* *</ options>* ;

The **TABLES** statement requests one-way to n -way frequency and crosstabulation tables and statistics for those tables.

If you omit the **TABLES** statement, PROC FREQ generates one-way frequency tables for all data set variables that are not listed in the other statements.

The following argument is required in the **TABLES** statement.

requests

specify the frequency and crosstabulation tables to produce. A request is composed of one variable name or several variable names separated by asterisks. To request a one-way frequency table, use a single variable. To request a two-way crosstabulation table, use an asterisk between two variables. To request a multiway table (an n -way table, where $n > 2$), separate the desired variables with asterisks. The unique values of these variables form the rows, columns, and strata of the table. You can include up to 50 variables in a single multiway table request.

For two-way to multiway tables, the values of the last variable form the crosstabulation table columns, while the values of the next-to-last variable form the rows. Each level (or combination of levels) of the other variables forms one stratum. PROC FREQ produces a separate crosstabulation table for each stratum. For example, a specification of $A*B*C*D$ in a **TABLES** statement produces k tables, where k is the number of different combinations of values for A and B. Each table lists the values for C down the side and the values for D across the top.

You can use multiple TABLES statements in the PROC FREQ step. PROC FREQ builds all the table requests in one pass of the data, so that there is essentially no loss of efficiency. You can also specify any number of table requests in a single TABLES statement. To specify multiple table requests quickly, use a grouping syntax by placing parentheses around several variables and joining other variables or variable combinations. For example, the statements shown in Table 38.8 illustrate grouping syntax.

Table 38.8 Grouping Syntax

TABLES Request	Equivalent to
A*(B C)	A*B A*C
(A B)*(C D)	A*C B*C A*D B*D
(A B C)*D	A*D B*D C*D
A -- C	A B C
(A -- C)*D	A*D B*D C*D

The TABLES statement variables are one or more variables from the DATA= input data set. These variables can be either character or numeric, but the procedure treats them as categorical variables. PROC FREQ uses the formatted values of the TABLES variable to determine the categorical variable levels. So if you assign a format to a variable with a FORMAT statement, PROC FREQ formats the values before dividing observations into the levels of a frequency or crosstabulation table. See the FORMAT procedure in the *Base SAS Procedures Guide* and the FORMAT statement and SAS formats in *SAS Formats and Informats: Reference*.

If you use PROC FORMAT to create a user-written format that combines missing and nonmissing values into one category, PROC FREQ treats the entire category of formatted values as missing. See the discussion in the section “[Grouping with Formats](#)” on page 2544 for more information.

By default, the frequency or crosstabulation table lists the values of both character and numeric variables in ascending order based on internal (unformatted) variable values. You can change the order of the values in the table by specifying the [ORDER=](#) option in the PROC FREQ statement. To list the values in ascending order by formatted value, use ORDER=FORMATTED.

Without Options

If you request a one-way frequency table for a variable without specifying options, PROC FREQ produces frequencies, cumulative frequencies, percentages of the total frequency, and cumulative percentages for each value of the variable. If you request a two-way or an *n*-way crosstabulation table without specifying options, PROC FREQ produces crosstabulation tables that include cell frequencies, cell percentages of the total frequency, cell percentages of row frequencies, and cell percentages of column frequencies. The procedure excludes observations with missing values from the table but displays the total frequency of missing observations below each table.

Options

Table 38.9 lists the *options* available in the TABLES statement. Descriptions of the *options* follow in alphabetical order.

Table 38.9 TABLES Statement Options

Option	Description
Control Statistical Analysis	
AGREE	Requests tests and measures of classification agreement
ALL	Requests tests and measures of association produced by the CHISQ , MEASURES , and CMH options
ALPHA=	Sets the confidence level for confidence limits
BINOMIAL	Requests binomial proportions, confidence limits, and tests for one-way tables
CHISQ	Requests chi-square tests and measures based on chi-square
CL	Requests confidence limits for MEASURES statistics
CMH	Requests all Cochran-Mantel-Haenszel statistics
CMH1	Requests CMH correlation statistic, adjusted odds ratios, and adjusted relative risks
CMH2	Requests CMH correlation and row mean scores (ANOVA) statistics, adjusted odds ratios, and adjusted relative risks
CONVERGE=	Specifies convergence criterion for polychoric correlation
FISHER	Requests Fisher's exact test for tables larger than 2×2
GAILSIMON	Requests Gail-Simon test for qualitative interactions
JT	Requests Jonckheere-Terpstra test
MAXITER=	Specifies maximum number of iterations for polychoric correlation
MEASURES	Requests measures of association
MISSING	Treats missing values as nonmissing
PLCORR	Requests polychoric correlation
RELRISK	Requests odds ratios and relative risks for 2×2 tables
RISKDIFF	Requests risks and risk differences for 2×2 tables
SCORES=	Specifies the type of row and column scores
TREND	Requests Cochran-Armitage test for trend
Control Additional Table Information	
CELLCHI2	Displays cell contributions to the Pearson chi-square statistic
CUMCOL	Displays cumulative column percentages
DEVIATION	Displays deviations of cell frequencies from expected values
EXPECTED	Displays expected cell frequencies
MISSPRINT	Displays missing value frequencies
SPARSE	Includes all possible combinations of variable levels in the LIST table and OUT= data set
STDRES	Displays standardized residuals in the CROSSLIST table
TOTPCT	Displays percentages of total frequency for n -way tables ($n > 2$)
Control Displayed Output	
CONTENTS=	Specifies the contents label for crosstabulation tables
CROSSLIST	Displays crosstabulation tables in ODS column format
FORMAT=	Formats the frequencies in crosstabulation tables
LIST	Displays two-way to n -way tables in list format
MAXLEVELS=	Specifies maximum number of levels to display in one-way tables
NOCOL	Suppresses display of column percentages

Table 38.9 *continued*

Option	Description
NOCUM	Suppresses display of cumulative frequencies and percentages
NOFREQ	Suppresses display of frequencies
NOPERCENT	Suppresses display of percentages
NOPRINT	Suppresses display of crosstabulation tables but displays statistics
NOROW	Suppresses display of row percentages
NOSPARE	Suppresses zero frequency levels in the CROSSLIST table, LIST table, and OUT= data set
NOWARN	Suppresses log warning message for the chi-square test
PRINTKWT	Displays kappa coefficient weights
SCOROUT	Displays row and column scores
Produce Statistical Graphics	
PLOTS=	Requests plots from ODS Graphics
Create an Output Data Set	
OUT=	Names an output data set to contain frequency counts
OUTCUM	Includes cumulative frequencies and percentages in the output data set for one-way tables
OUTEXPECT	Includes expected frequencies in the output data set
OUTPCT	Includes row, column, and two-way table percentages in the output data set

You can specify the following *options* in a TABLES statement.

AGREE <(WT=FC)>

requests tests and measures of classification agreement for square tables. The AGREE option provides McNemar's test for 2×2 tables and Bowker's test of symmetry for square tables with more than two response categories. The AGREE option also produces the simple kappa coefficient, the weighted kappa coefficient, their asymptotic standard errors, and their confidence limits. When there are multiple strata, the AGREE option provides overall simple and weighted kappas as well as tests for equal kappas among strata. When there are multiple strata and two response categories, PROC FREQ computes Cochran's Q test. See the section "[Tests and Measures of Agreement](#)" on page 2588 for details about these statistics.

If you specify the WT=FC option in parentheses after the AGREE option, PROC FREQ uses Fleiss-Cohen weights to compute the weighted kappa coefficient. By default, PROC FREQ uses Cicchetti-Allison weights. See the section "[Weighted Kappa Coefficient](#)" on page 2590 for details. You can specify the [PRINTKWT](#) option to display the kappa coefficient weights.

AGREE statistics are computed only for square tables, where the number of rows equals the number of columns. If your table is not square due to observations with zero weights, you can specify the [ZEROS](#) option in the WEIGHT statement to include these observations. For more details, see the section "[Tables with Zero Rows and Columns](#)" on page 2592.

You can use the [TEST](#) statement to request asymptotic tests for the simple and weighted kappa coefficients. You can request exact p -values for the simple and weighted kappa coefficient tests, as well

as for McNemar's test, by specifying the corresponding options in the **EXACT** statement. See the section “[Exact Statistics](#)” on page 2601 for more information.

ALL

requests all tests and measures that are produced by the **CHISQ**, **MEASURES**, and **CMH** options. You can control the number of CMH statistics to compute by specifying the **CMH1** or **CMH2** option.

ALPHA= α

specifies the level of confidence limits. The value of α must be between 0 and 1, and the default is 0.05. A confidence level of α produces $100(1 - \alpha)\%$ confidence limits. The default of ALPHA=0.05 produces 95% confidence limits.

ALPHA= applies to confidence limits that are requested by TABLES statement options. There is a separate **ALPHA=** option in the EXACT statement that sets the level of confidence limits for Monte Carlo estimates of exact p -values, which you request by specifying the **MC** option in the EXACT statement.

BINOMIAL <(binomial-options)>

requests the binomial proportion for one-way tables. When you specify the BINOMIAL option, by default PROC FREQ also provides the asymptotic standard error, asymptotic Wald and exact (Clopper-Pearson) confidence limits, and the asymptotic equality test for the binomial proportion.

You can specify *binomial-options* in parentheses after the BINOMIAL option. The **LEVEL=** *binomial-option* identifies the variable level for which to compute the proportion. If you do not specify LEVEL=, PROC FREQ computes the proportion for the first level that appears in the output. The **P=** *binomial-option* specifies the null proportion for the binomial tests. If you do not specify P=, PROC FREQ uses P=0.5 by default.

You can also specify *binomial-options* to request additional tests and confidence limits for the binomial proportion. The **EQUIV**, **NONINF**, and **SUP** *binomial-options* request tests of equivalence, noninferiority, and superiority, respectively. The **CL=** *binomial-option* requests confidence limits for the binomial proportion. [Table 38.10](#) summarizes the *binomial-options*.

Available confidence limits for the binomial proportion include Agresti-Coull, exact (Clopper-Pearson), Jeffreys, Wald, and Wilson (score) confidence limits. You can specify more than one type of binomial confidence limits in the same analysis. If you do not specify any confidence limit requests, PROC FREQ computes Wald asymptotic confidence limits and exact (Clopper-Pearson) confidence limits by default. The **ALPHA=** option determines the confidence level, and the default of ALPHA=0.05 produces 95% confidence limits for the binomial proportion.

As part of the noninferiority, superiority, and equivalence analyses, PROC FREQ provides test-based confidence limits that have a confidence coefficient of $100(1 - 2\alpha)\%$ (Schuirmann 1999). The **ALPHA=** option determines the confidence level, and the default of ALPHA=0.05 produces 90% confidence limits. See the sections “[Noninferiority Test](#)” on page 2569 and “[Equivalence Test](#)” on page 2570 for details.

To request exact tests for the binomial proportion, specify the **BINOMIAL** option in the **EXACT** statement. PROC FREQ then computes exact p -values for all binomial tests that you request with *binomial-options*, which can include tests of noninferiority, superiority, and equivalence, in addition to the test of equality.

See the section “[Binomial Proportion](#)” on page 2565 for details.

Table 38.10 BINOMIAL Options

Option	Description
LEVEL=	Specifies the variable level
P=	Specifies the null proportion
CORRECT	Requests continuity correction
Request Confidence Limits	
CL=AGRESTICOULL AC	Requests Agresti-Coull confidence limits
CL=ALL	Requests all confidence limits
CL=EXACT CLOPPERPEARSON	Requests Clopper-Pearson confidence limits
CL=JEFFREYS J	Requests Jeffreys confidence limits
CL=WALD	Requests Wald confidence limits
CL=WILSON W	Requests Wilson (score) confidence limits
Request Tests	
EQUIV EQUIVALENCE	Requests an equivalence test
NONINF NONINFERIORITY	Requests a noninferiority test
SUP SUPERIORITY	Requests a superiority test
MARGIN=	Specifies the test margin
VAR=SAMPLE NULL	Specifies the test variance

You can specify the following *binomial-options* in parentheses after the BINOMIAL option:

CL=type | (types)

requests confidence limits for the binomial proportion. You can specify one or more *types* of confidence limits. When you specify only one *type*, you can omit the parentheses around the request.

PROC FREQ displays the confidence limits in the “Binomial Confidence Limits” table. The ALPHA= option determines the level of the confidence limits that the CL= *binomial-option* provides. The default of ALPHA=0.05 produces 95% confidence limits for the binomial proportion. This differs from the test-based confidence limits that are provided with the equivalence, noninferiority, and superiority tests (EQUIV, NONINF, and SUP), which have a confidence coefficient of $100(1-2\alpha)\%$ (Schuirmann 1999). See the sections “Equivalence Test” on page 2570 and “Noninferiority Test” on page 2569 for details.

You can specify the CL= *binomial-option* with or without requests for binomial tests. The confidence limits that CL= produces do not depend on the tests that you request and do not use the value of the test margin (which you specify by using the MARGIN= *binomial-option*).

The following *types* of binomial confidence limits are available:

AGRESTICOULL | AC

requests Agresti-Coull confidence limits for the binomial proportion. See the section “Agresti-Coull Confidence Limits” on page 2566 for details.

ALL

requests all available types of confidence limits for the binomial proportion. These include the following: Agresti-Coull, exact (Clopper-Pearson), Jeffreys, Wald, and Wilson (score) confidence limits.

EXACT | CLOPPERPEARSON

requests exact (Clopper-Pearson) confidence limits for the binomial proportion. See the section “[Exact \(Clopper-Pearson\) Confidence Limits](#)” on page 2567 for details. If you do not request any binomial confidence limits by specifying the `CL= binomial-option`, PROC FREQ produces Wald and exact (Clopper-Pearson) confidence limits by default. To request exact tests for the binomial proportion, specify the `BINOMIAL` option in the **EXACT** statement.

JEFFREYS | J

requests Jeffreys confidence limits for the binomial proportion. See the section “[Jeffreys Confidence Limits](#)” on page 2566 for details.

WALD < (CORRECT) >

requests Wald confidence limits for the binomial proportion. See the section “[Wald Confidence Limits](#)” on page 2566 for details. If you do not request any binomial confidence limits by specifying the `CL= binomial-option`, PROC FREQ produces Wald and exact (Clopper-Pearson) confidence limits by default.

If you specify `CL=WALD(CORRECT)`, the Wald confidence limits include a continuity correction. If you specify the `CORRECT binomial-option`, the Wald confidence limits and the Wald tests include continuity corrections.

WILSON | W | SCORE < (CORRECT) >

requests Wilson confidence limits for the binomial proportion. These are also known as *score* confidence limits. See the section “[Wilson \(Score\) Confidence Limits](#)” on page 2567 for details.

If you specify `CL=WILSON(CORRECT)` or the `CORRECT binomial-option`, the Wilson confidence limits include a continuity correction.

CORRECT

includes a continuity correction in the Wald confidence limits, Wald tests, and Wilson confidence limits.

You can request continuity corrections individually for Wald or Wilson confidence limits by specifying the `CL=WALD(CORRECT)` or `CL=WILSON(CORRECT) binomial-option`, respectively.

EQUIV | EQUIVALENCE

requests a test of equivalence for the binomial proportion. See the section “[Equivalence Test](#)” on page 2570 for details. You can specify the equivalence test margins, the null proportion, and the variance type with the `MARGIN=`, `P=`, and `VAR= binomial-options`, respectively. To request an exact equivalence test, specify the `BINOMIAL` option in the **EXACT** statement.

LEVEL=level-number | 'level-value'

specifies the variable level for the binomial proportion. By default, PROC FREQ computes the proportion of observations for the first variable level that appears in the output. To request a different level, use `LEVEL=level-number` or `LEVEL='level-value'`, where *level-number* is the variable level's number or order in the output, and *level-value* is the formatted value of the variable level. The value of *level-number* must be a positive integer. You must enclose *level-value* in single quotes.

MARGIN=*value* | (*lower,upper*)

specifies the margin for the noninferiority, superiority, and equivalence tests, which you request with the **NONINF**, **SUP**, and **EQUIV** *binomial-options*, respectively. If you do not specify **MARGIN=**, PROC FREQ uses a margin of 0.2 by default.

For noninferiority and superiority tests, specify a single *value* for the **MARGIN=** option. The **MARGIN=** *value* must be a positive number. You can specify *value* as a number between 0 and 1. Or you can specify *value* in percentage form as a number between 1 and 100, and PROC FREQ converts that number to a proportion. The procedure treats the value 1 as 1%.

For noninferiority and superiority tests, the test limits must be between 0 and 1. The limits are determined by the null proportion value (which you can specify with the **P=** *binomial-option*) and by the margin value. The noninferiority limit equals the null proportion minus the margin. By default, the null proportion equals 0.5 and the margin equals 0.2, which gives a noninferiority limit of 0.3. The superiority limit equals the null proportion plus the margin, which is 0.7 by default.

For an equivalence test, you can specify a single **MARGIN=** *value*, or you can specify both *lower* and *upper* values. If you specify a single **MARGIN=** *value*, it must be a positive number, as described previously. If you specify a single **MARGIN=** *value* for an equivalence test, PROC FREQ uses $-value$ as the lower margin and *value* as the upper margin for the test. If you specify both *lower* and *upper* values for an equivalence test, you can specify them in proportion form as numbers between -1 or 1. Or you can specify them in percentage form as numbers between -100 and 100, and PROC FREQ converts the numbers to proportions. The value of *lower* must be less than the value of *upper*.

The equivalence limits must be between 0 and 1. The equivalence limits are determined by the null proportion value (which you can specify with the **P=** *binomial-option*) and by the margin values. The lower equivalence limit equals the null proportion plus the lower margin. By default, the null proportion equals 0.5 and the lower margin equals -0.2 , which gives a lower equivalence limit of 0.3. The upper equivalence limit equals the null proportion plus the upper margin, which is 0.7 by default.

See the sections “**Noninferiority Test**” on page 2569 and “**Equivalence Test**” on page 2570 for details.

NONINF | **NONINFERIORITY**

requests a test of noninferiority for the binomial proportion. See the section “**Noninferiority Test**” on page 2569 for details. You can specify the noninferiority test margin, the null proportion, and the variance type with the **MARGIN=**, **P=**, and **VAR=** *binomial-options*, respectively. To request an exact noninferiority test, specify the **BINOMIAL** option in the **EXACT** statement.

P=*value*

specifies the null hypothesis proportion for the binomial tests. If you omit the **P=** option, PROC FREQ uses 0.5 as the null proportion. The null proportion *value* must be a positive number. You can specify *value* as a number between 0 and 1. Or you can specify *value* in percentage form as a number between 1 and 100, and PROC FREQ converts that number to a proportion. The procedure treats the value 1 as 1%.

SUP | SUPERIORITY

requests a test of superiority for the binomial proportion. See the section “[Superiority Test](#)” on page 2570 for details. You can specify the superiority test margin, the null proportion, and the variance type with the **MARGIN=**, **P=**, and **VAR=** *binomial-options*, respectively. To request an exact superiority test, specify the **BINOMIAL** option in the **EXACT** statement.

VAR=SAMPLE | NULL

specifies the type of variance estimate to use in the Wald tests of noninferiority, superiority, and equivalence. The default is **VAR=SAMPLE**, which estimates the variance from the sample proportion. **VAR=NULL** uses a test-based variance that is computed from the null hypothesis proportion (which is specified by the **P=** *binomial-option*). See the sections “[Noninferiority Test](#)” on page 2569 and “[Equivalence Test](#)” on page 2570 for details.

CELLCHI2

displays each crosstabulation table cell’s contribution to the total Pearson chi-square statistic. The cell contribution is computed as

$$(frequency - expected)^2 / expected$$

where *frequency* is the table cell frequency or count and *expected* is the expected cell frequency, which is computed under the null hypothesis that the row and column variables are independent. See the section “[Pearson Chi-Square Test for Two-Way Tables](#)” on page 2551 for details.

The **CELLCHI2** option has no effect for one-way tables or for tables that are displayed with the **LIST** option.

CHISQ <(chisq-options)>

requests chi-square tests of homogeneity or independence and measures of association that are based on the chi-square statistic. For two-way tables, the chi-square tests include the Pearson chi-square, likelihood ratio chi-square, and Mantel-Haenszel chi-square tests. The chi-square measures include the phi coefficient, contingency coefficient, and Cramér’s *V*. For 2×2 tables, the **CHISQ** option also provides Fisher’s exact test and the continuity-adjusted chi-square test. See the section “[Chi-Square Tests and Statistics](#)” on page 2550 for details.

For one-way tables, the **CHISQ** option provides the Pearson chi-square goodness-of-fit test. You can also request the likelihood ratio goodness-of-fit test for one-way tables by specifying the **LRCHISQ** *chisq-option* in parentheses after the **CHISQ** option. By default, the one-way chi-square tests are based on the null hypothesis of equal proportions. Alternatively, you can provide null hypothesis proportions or frequencies by specifying the **TESTP=** or **TESTF=** *chisq-option*, respectively. See the section “[Chi-Square Test for One-Way Tables](#)” on page 2551 for more information.

To request Fisher’s exact test for tables larger than 2×2 , specify the **FISHER** option in the **EXACT** statement. Exact *p*-values are also available for the Pearson, likelihood ratio, and Mantel-Haenszel chi-square tests. See the description of the **EXACT** statement for more information.

You can specify the following *chisq-options* in parentheses after the **CHISQ** option:

DF=df

specifies the degrees of freedom for the chi-square tests. The value of *df* must not be zero. If the value of *df* is positive, PROC FREQ uses *df* as the degrees of freedom for the chi-square tests. If the value of *df* is negative, PROC FREQ uses *df* to adjust the default degrees of freedom for the chi-square tests.

By default for one-way tables, the value of *df* is $(n - 1)$, where n is the number of variable levels in the table. By default for two-way tables, the value of *df* is $(r - 1)(c - 1)$, where r is the number of rows in the table and c is the number of columns. See the sections “[Chi-Square Test for One-Way Tables](#)” on page 2551 and “[Chi-Square Tests and Statistics](#)” on page 2550 for more information.

If you specify a negative value of *df*, PROC FREQ adjusts the default degrees of freedom by adding the (negative) value of *df* to the default value to produce the adjusted degrees of freedom. The adjusted degrees of freedom must be positive.

The `DF= chisq-option` specifies or adjusts the degrees of freedom for the following chi-square tests: the Pearson and likelihood ratio goodness-of-fit tests for one-way tables; and the Pearson, likelihood ratio, and Mantel-Haenszel chi-square tests for two-way tables.

LRCHISQ

requests the likelihood ratio goodness-of-fit test for one-way tables. See the section “[Likelihood Ratio Chi-Square Test for One-Way Tables](#)” on page 2552 for more information.

By default, this test is based on the null hypothesis of equal proportions. You can provide null hypothesis proportions or frequencies by specifying the `TESTP=` or `TESTF= chisq-option`, respectively. You can request an exact likelihood ratio goodness-of-fit test by specifying the `LRCHI` option in the EXACT statement.

TESTF=(values) | SAS-data-set

specifies null hypothesis frequencies for the one-way chi-square goodness-of-fit tests. See the section “[Chi-Square Test for One-Way Tables](#)” on page 2551 for details. You can list the null frequencies as *values* in parentheses after `TESTF=`. Or you can provide the null frequencies in a secondary input data set by specifying `TESTF=SAS-data-set`. The `TESTF=SAS-data-set` cannot be the same data set that you specify in the `DATA=` option. You can specify only one `TESTF=` or `TESTP=` data set in a single invocation of the procedure.

If you list the null frequencies as *values*, you can separate the *values* with blanks or commas. The *values* must be positive numbers. The number of *values* must equal the number of variable levels in the one-way table. The sum of the *values* must equal the total frequency for the one-way table. Order the *values* to match the order in which the corresponding variable levels appear in the one-way frequency table.

If you provide the null frequencies in a secondary input data set (`TESTF=SAS-data-set`), the variable that contains the null frequencies should be named `_TESTF_`, `TestFrequency`, or `Frequency`. The null frequencies must be positive numbers. The number of frequencies must equal the number of levels in the one-way frequency table, and the sum of the frequencies must equal the total frequency for the one-way table. Order the null frequencies in the data set to match the order in which the corresponding variable levels appear in the one-way frequency table.

TESTP=(values) | SAS-data-set

specifies null hypothesis proportions for the one-way chi-square goodness-of-fit tests. See the section “[Chi-Square Test for One-Way Tables](#)” on page 2551 for details. You can list the null proportions as *values* in parentheses after `TESTP=`. Or you can provide the null proportions in a secondary input data set by specifying `TESTP=SAS-data-set`. The `TESTP=SAS-data-set` cannot be the same data set that you specify in the `DATA=` option. You can specify only one `TESTF=` or `TESTP=` data set in a single invocation of the procedure.

If you list the null proportions as *values*, you can separate the *values* with blanks or commas. The *values* must be positive numbers. The number of *values* must equal the number of variable levels in the one-way table. Order the *values* to match the order in which the corresponding variable levels appear in the one-way frequency table. You can specify *values* in probability form as numbers between 0 and 1, where the proportions sum to 1. Or you can specify *values* in percentage form as numbers between 0 and 100, where the percentages sum to 100.

If you provide the null proportions in a secondary input data set (TESTP=SAS-data-set), the variable that contains the null proportions should be named _TESTP_, TestPercent, or Percent. The null proportions must be positive numbers. The number of proportions must equal the number of levels in the one-way frequency table. You can provide the proportions in probability form as numbers between 0 and 1, where the proportions sum to 1. Or you can provide the proportions in percentage form as numbers between 0 and 100, where the percentages sum to 100. Order the null proportions in the data set to match the order in which the corresponding variable levels appear in the one-way frequency table.

WARN=*value* | (*values*)

controls the warning message for the validity of the asymptotic Pearson chi-square test. By default, PROC FREQ displays a warning message when more than 20% of the table cells have expected frequencies that are less than 5. If you specify the **NOPRINT** option in the **PROC FREQ** statement, the procedure displays the warning in the log; otherwise, the procedure displays the warning as a footnote in the chi-square table. You can use the WARN= option to suppress the warning and to include a warning indicator in the output data set.

The WARN= option can take one or more of the following values. If you specify more than one value, enclose the values in parentheses after WARN=. For example, **warn = (output noprint)**.

Value of WARN=	Description
OUTPUT	Adds a warning indicator variable to the output data set
NOLOG	Suppresses the chi-square warning message in the log
NOPRINT	Suppresses the chi-square warning message in the display
NONE	Suppresses the chi-square warning message entirely

If you specify the WARN=OUTPUT option, the ODS output data set ChiSq contains a variable named Warning that equals 1 for the Pearson chi-square observation when more than 20% of the table cells have expected frequencies that are less than 5 and equals 0 otherwise. If you specify WARN=OUTPUT and also specify the CHISQ option in the **OUTPUT** statement, the statistics output data set contains a variable named WARN_PCHI that indicates the warning.

The WARN=NOLOG option has the same effect as the **NOWARN** option in the TABLES statement.

CL

requests confidence limits for the **MEASURES** statistics. If you omit the MEASURES option, the CL option invokes MEASURES. You can set the level of the confidence limits by using the **ALPHA=** option. The default of ALPHA=0.5 produces 95% confidence limits. See the sections “**Measures of Association**” on page 2556 and “**Confidence Limits**” on page 2556 for more information.

CMH <(cmh-options)>

requests Cochran-Mantel-Haenszel statistics, which test for association between the row and column variables after adjusting for the remaining variables in a multiway table. The Cochran-Mantel-Haenszel statistics include the nonzero correlation statistic, the row mean scores (ANOVA) statistic, and the general association statistic. In addition, for 2×2 tables, the CMH option provides the adjusted Mantel-Haenszel and logit estimates of the odds ratio and relative risks, together with their confidence limits. For stratified 2×2 tables, the CMH option provides the Breslow-Day test for homogeneity of odds ratios. (To request Tarone's adjustment for the Breslow-Day test, specify the **BDT** *cmh-option*.) See the section "[Cochran-Mantel-Haenszel Statistics](#)" on page 2592 for details.

You can use the **CMH1** or **CMH2** option to control the number of CMH statistics that PROC FREQ computes.

For stratified 2×2 tables, you can request Zelen's exact test for equal odds ratios by specifying the **EQOR** option in the **EXACT** statement. See the section "[Zelen's Exact Test for Equal Odds Ratios](#)" on page 2598 for details. You can request exact confidence limits for the common odds ratio by specifying the **COMOR** option in the **EXACT** statement. This option also provides a common odds ratio test. See the section "[Exact Confidence Limits for the Common Odds Ratio](#)" on page 2599 for details.

You can specify the following *cmh-options* in parentheses after the CMH option. These *cmh-options*, which apply to stratified 2×2 tables, are also available with the **CMH1** or **CMH2** option.

BDT

requests Tarone's adjustment in the Breslow-Day test for homogeneity of odds ratios. See the section "[Breslow-Day Test for Homogeneity of the Odds Ratios](#)" on page 2598 for details.

GAILSIMON | **GS** <(COLUMN=1 | 2)>

requests the Gail-Simon test for qualitative interaction, which applies to stratified 2×2 tables. See the section "[Gail-Simon Test for Qualitative Interactions](#)" on page 2601 for details.

The **COLUMN=** option specifies the column of the risk differences to use in computing the Gail-Simon test. By default, PROC FREQ uses column 1 risk differences. If you specify **COLUMN=2**, PROC FREQ uses column 2 risk differences.

The **GAILSIMON** *cmh-option* has the same effect as the **GAILSIMON** option in the **TABLES** statement.

MANTELFLEISS | **MF**

requests the Mantel-Fleiss criterion for the Mantel-Haenszel statistic for stratified 2×2 tables. See the section "[Mantel-Fleiss Criterion](#)" on page 2595 for details.

CMH1 <(cmh-options)>

requests the Cochran-Mantel-Haenszel correlation statistic. This option does not provide the CMH row mean scores (ANOVA) statistic or the general association statistic, which are provided by the **CMH** option. For tables larger than 2×2 , the CMH1 option requires less memory than the CMH option, which can require an enormous amount of memory for large tables.

For 2×2 tables, the CMH1 option also provides the adjusted Mantel-Haenszel and logit estimates of the odds ratio and relative risks, together with their confidence limits. For stratified 2×2 tables, the CMH1 option provides the Breslow-Day test for homogeneity of odds ratios.

The *cmh-options* for CMH1 are the same as the *cmh-options* that are available with the CMH option. See the description of the **CMH** option for details.

CMH2 <(*cmh-options*)>

requests the Cochran-Mantel-Haenszel correlation statistic and the row mean scores (ANOVA) statistic. This option does not provide the CMH general association statistic, which is provided by the [CMH](#) option. For tables larger than 2×2 , the CMH2 option requires less memory than the CMH option, which can require an enormous amount of memory for large tables.

For 2×2 tables, the CMH1 option also provides the adjusted Mantel-Haenszel and logit estimates of the odds ratio and relative risks, together with their confidence limits. For stratified 2×2 tables, the CMH1 option provides the Breslow-Day test for homogeneity of odds ratios.

The *cmh-options* for CMH2 are the same as the *cmh-options* that are available with the CMH option. See the description of the [CMH](#) option for details.

CONTENTS='string'

specifies the label to use for crosstabulation tables in the contents file, the Results window, and the trace record. For information about output presentation, see the *SAS Output Delivery System: User's Guide*.

If you omit the CONTENTS= option, the contents label for crosstabulation tables is “Cross-Tabular Freq Table” by default.

Note that contents labels for all crosstabulation tables that are produced by a single TABLES statement use the same text. To specify different contents labels for different crosstabulation tables, request the tables in separate TABLES statements and use the CONTENTS= option in each TABLES statement.

To remove the crosstabulation table entry from the contents file, you can specify a null label with CONTENTS=“.”

The CONTENTS= option affects only contents labels for crosstabulation tables. It does not affect contents labels for other PROC FREQ tables.

To specify the contents label for any PROC FREQ table, you can use PROC TEMPLATE to create a customized table definition. The CONTENTS_LABEL attribute in the DEFINE TABLE statement of PROC TEMPLATE specifies the contents label for the table. See the chapter “The TEMPLATE Procedure” in the *SAS Output Delivery System: User's Guide* for more information.

CONVERGE=value

specifies the convergence criterion for computing the polychoric correlation, which you request by specifying the [PLCORR](#) option. The CONVERGE= *value* must be a positive number. By default, CONVERGE=0.0001. Iterative computation of the polychoric correlation stops when the convergence measure falls below *value* or when the number of iterations exceeds the [MAXITER=](#) value, whichever happens first. For parameter values that are less than 0.01, the procedure evaluates convergence by using the absolute difference instead of the relative difference. See the section “[Polychoric Correlation](#)” on page 2562 for details.

CROSSLIST <STDRES>

displays crosstabulation tables in ODS column format instead of the default crosstabulation cell format. In the CROSSLIST table display, the rows correspond to the crosstabulation table cells, and the columns correspond to descriptive statistics such as Frequency and Percent. The CROSSLIST table displays the same information as the default crosstabulation table, but it uses an ODS column format instead of the table cell format. See the section “[Multiway Tables](#)” on page 2611 for details about the contents of the CROSSLIST table.

You can control the contents of a CROSSLIST table by specifying the same options available for the default crosstabulation table. These include the [NOFREQ](#), [NOPERCENT](#), [NOROW](#), and [NOCOL](#) options. You can request additional information in a CROSSLIST table by specifying the [CELLCHI2](#),

DEVIATION, EXPECTED, MISSPRINT, and TOTPCT options. You can also display standardized residuals in a CROSSLIST table by specifying the CROSSLIST(STDRES) option (which is not available for the default crosstabulation table).

The FORMAT= and CUMCOL options have no effect on CROSSLIST tables. You cannot specify both the LIST option and the CROSSLIST option in the same TABLES statement.

For CROSSLIST tables, you can use the NOSPARSE option to suppress display of variable levels that have zero frequencies. By default, PROC FREQ displays all levels of the column variable within each level of the row variable, including any levels that have zero frequencies. By default for multiway tables displayed as CROSSLIST tables, the procedure displays all levels of the row variable for each stratum of the table, including any row levels with zero frequencies in the stratum.

You can specify the following *option* in parentheses after the CROSSLIST option:

STDRES

displays the standardized residuals of the table cells in the CROSSLIST table. The standardized residual is the ratio of (*frequency* – *expected*) to its standard error, where *frequency* is the table cell frequency and *expected* is the expected table cell frequency, which is computed under the null hypothesis that the row and column variables are independent. See the section “[Standardized Residuals](#)” on page 2552 for more information.

CUMCOL

displays the cumulative column percentages in the cells of the crosstabulation table. The CUMCOL option does not apply to crosstabulation tables produced with the LIST or CROSSLIST option.

DEVIATION

displays the deviation of the frequency from the expected frequency for each cell of the crosstabulation table. See the section “[Pearson Chi-Square Test for Two-Way Tables](#)” on page 2551 for details. The DEVIATION option does not apply to crosstabulation tables produced with the LIST option.

EXPECTED

displays the expected cell frequencies under the hypothesis of independence (or homogeneity) for crosstabulation tables. See the section “[Pearson Chi-Square Test for Two-Way Tables](#)” on page 2551 for details. The EXPECTED option does not apply to tables produced with the LIST option.

FISHER | EXACT

requests Fisher’s exact test for tables that are larger than 2×2 . (For 2×2 tables, the CHISQ option provides Fisher’s exact test.) This test is also known as the Freeman-Halton test. See the sections “[Fisher’s Exact Test](#)” on page 2554 and “[Exact Statistics](#)” on page 2601 for more information.

If you omit the CHISQ option in the TABLES statement, the FISHER option invokes CHISQ. You can also request Fisher’s exact test by specifying the FISHER option in the EXACT statement.

NOTE: PROC FREQ computes exact tests by using fast and efficient algorithms that are superior to direct enumeration. Exact tests are appropriate when a data set is small, sparse, skewed, or heavily tied. For some large problems, computation of exact tests might require a substantial amount of time and memory. Consider using asymptotic tests for such problems. Alternatively, when asymptotic methods might not be sufficient for such large problems, consider using Monte Carlo estimation of exact *p*-values. You can request Monte Carlo estimation by specifying the MC *computation-option* in the EXACT statement. See the section “[Computational Resources](#)” on page 2604 for more information.

FORMAT=*format-name*

specifies a format for the following crosstabulation table cell values: frequency, expected frequency, and deviation. PROC FREQ also uses the specified format to display the row and column total frequencies and the overall total frequency in crosstabulation tables.

You can specify any standard SAS numeric format or a numeric format defined with the FORMAT procedure. The format length must not exceed 24. If you omit the FORMAT= option, by default PROC FREQ uses the BEST6. format to display frequencies less than 1E6, and the BEST7. format otherwise.

The FORMAT= option applies only to crosstabulation tables displayed in the default format. It does not apply to crosstabulation tables produced with the [LIST](#) or [CROSSLIST](#) option.

To change display formats in any FREQ table, you can use PROC TEMPLATE. See the chapter “The TEMPLATE Procedure” in the *SAS Output Delivery System: User’s Guide* for more information.

GAILSIMON | GS <(COLUMN=1 | 2)>

requests the Gail-Simon test for qualitative interaction, which applies to stratified 2×2 tables. See the section “[Gail-Simon Test for Qualitative Interactions](#)” on page 2601 for details.

The COLUMN= option specifies the column of the risk differences to use in computing the Gail-Simon test. By default, PROC FREQ uses column 1 risk differences. If you specify COLUMN=2, PROC FREQ uses column 2 risk differences.

JT

requests the Jonckheere-Terpstra test. See the section “[Jonckheere-Terpstra Test](#)” on page 2586 for details. To request exact p -values for the Jonckheere-Terpstra test, specify the JT option in the EXACT statement. See the section “[Exact Statistics](#)” on page 2601 for more information.

LIST

displays two-way to n -way crosstabulation tables in a list format instead of the default crosstabulation cell format. The LIST option displays the entire multiway table in one table, instead of displaying a separate two-way table for each stratum. See the section “[Multiway Tables](#)” on page 2611 for details.

The LIST option is not available when you also specify statistical options. You must use the standard crosstabulation table display or the [CROSSLIST](#) display when you request statistical tests or measures.

MAXITER=*number*

specifies the maximum number of iterations for computing the polychoric correlation, which you request by specifying the [PLCORR](#) option. The value of *number* must be a positive integer. By default MAXITER=20. Iterative computation of the polychoric correlation stops when the number of iterations exceeds the MAXITER= value or when the convergence measure falls below the [CONVERGE=](#) value, whichever happens first. See the section “[Polychoric Correlation](#)” on page 2562 for details.

MAXLEVELS=*n*

specifies the maximum number of variable levels to display in one-way frequency tables. The value of n must be a positive integer. PROC FREQ displays the first n variable levels, matching the order in which the levels appear in the one-way frequency table. (The [ORDER=](#) option controls the order of the variable levels. By default, ORDER=INTERNAL, which orders the variable levels by unformatted value.)

The MAXLEVELS= option also applies to one-way frequency plots, which you can request by specifying the **PLOTS=FREQPLOT** option when ODS Graphics is enabled.

If you specify the **MISSPRINT** option to display missing levels in the frequency table, the MAXLEVELS= option displays the first n nonmissing levels.

The MAXLEVELS= option does not apply to the **OUT=** output data set, which includes all variable levels. The MAXLEVELS= option does not affect the computation of percentages, statistics, or tests for the one-way table; these values are based on the complete table.

MEASURES

requests measures of association. The MEASURES option provides the following statistics and their asymptotic standard errors: gamma, Kendall's tau- b , Stuart's tau- c , Somers' $D(C|R)$, Somers' $D(R|C)$, Pearson and Spearman correlation coefficients, lambda (symmetric and asymmetric), and uncertainty coefficients (symmetric and asymmetric). To request confidence limits for these measures of association, you can specify the **CL** option.

For 2×2 tables, the MEASURES option also provides the odds ratio, column 1 relative risk, column 2 relative risk, and the corresponding confidence limits. Alternatively, you can request the odds ratio and relative risks separately (without the other measures of association), by specifying the **REL RISK** option.

See the section “[Measures of Association](#)” on page 2556 for details.

You can use the **TEST** statement to request asymptotic tests for the following measures of association: gamma, Kendall's tau- b , Stuart's tau- c , Somers' $D(C|R)$, Somers' $D(R|C)$, and Pearson and Spearman correlation coefficients. You can use the **EXACT** statement to request exact confidence limits for the odds ratio, exact unconditional confidence limits for the relative risks, and exact tests for the following measures of association: Kendall's tau- b , Stuart's tau- c , Somers' $D(C|R)$ and $D(R|C)$, and Pearson and Spearman correlation coefficients. See the descriptions of the **TEST** and **EXACT** statements for more information.

MISSING

treats missing values as a valid nonmissing level for all TABLES variables. The MISSING option displays the missing levels in frequency and crosstabulation tables and includes them in all calculations of percentages, tests, and measures.

By default, if you do not specify the MISSING or **MISSPRINT** option, an observation is excluded from a table if it has a missing value for any of the variables in the TABLES request. When PROC FREQ excludes observations with missing values, it displays the total frequency of missing observations below the table. See the section “[Missing Values](#)” on page 2545 for more information.

MISSPRINT

displays missing value frequencies in frequency and crosstabulation tables but does not include the missing value frequencies in any computations of percentages, tests, or measures.

By default, if you do not specify the **MISSING** or **MISSPRINT** option, an observation is excluded from a table if it has a missing value for any of the variables in the TABLES request. When PROC FREQ excludes observations with missing values, it displays the total frequency of missing observations below the table. See the section “[Missing Values](#)” on page 2545 for more information.

NOCOL

suppresses the display of column percentages in crosstabulation table cells.

NOCUM

suppresses the display of cumulative frequencies and percentages in one-way frequency tables. The NOCUM option also suppresses the display of cumulative frequencies and percentages in crosstabulation tables in list format, which you request with the [LIST](#) option.

NOFREQ

suppresses the display of cell frequencies in crosstabulation tables. The NOFREQ option also suppresses row total frequencies. This option has no effect for one-way tables or for crosstabulation tables in list format, which you request with the [LIST](#) option.

NOPERCENT

suppresses the display of overall percentages in crosstabulation tables. These percentages include the cell percentages of the total (two-way) table frequency, as well as the row and column percentages of the total table frequency. To suppress the display of cell percentages of row or column totals, use the [NOROW](#) or [NOCOL](#) option, respectively.

For one-way frequency tables and crosstabulation tables in list format, the NOPERCENT option suppresses the display of percentages and cumulative percentages.

NOPRINT

suppresses the display of frequency and crosstabulation tables but displays all requested tests and statistics. To suppress the display of all output, including tests and statistics, use the [NOPRINT](#) option in the PROC FREQ statement.

NOROW

suppresses the display of row percentages in crosstabulation table cells.

NOSPARE

suppresses the display of cells with a zero frequency count in [LIST](#) output and omits them from the [OUT=](#) data set. The NOSPARE option applies when you specify the [ZEROS](#) option in the [WEIGHT](#) statement to include observations with zero weights. By default, the ZEROS option invokes the [SPARSE](#) option, which displays table cells with a zero frequency count in the LIST output and includes them in the OUT= data set. See the description of the [ZEROS](#) option for more information.

The NOSPARE option also suppresses the display of variable levels with zero frequency in [CROSSLIST](#) tables. By default for CROSSLIST tables, PROC FREQ displays all levels of the column variable within each level of the row variable, including any column variable levels with zero frequency for that row. For multiway tables displayed with the CROSSLIST option, the procedure displays all levels of the row variable for each stratum of the table by default, including any row variable levels with zero frequency for the stratum.

NOWARN

suppresses the log warning message for the validity of the asymptotic Pearson chi-square test. By default, PROC FREQ provides a validity warning for the asymptotic Pearson chi-square test when more than 20 cells have expected frequencies that are less than 5. This warning message appears in the log if you specify the NOPRINT option in the PROC FREQ statement,

The NOWARN option is equivalent to the [CHISQ\(WARN=NOLOG\)](#) option. You can also use the [CHISQ\(WARN=\)](#) option to suppress the warning message in the display and to request a warning variable in the chi-square ODS output data set or in the OUTPUT data set.

OUT=SAS-data-set

names an output data set that contains frequency or crosstabulation table counts and percentages. If more than one table request appears in the TABLES statement, the contents of the OUT= data set correspond to the last table request in the TABLES statement. The OUT= data set variable COUNT contains the frequencies and the variable PERCENT contains the percentages. See the section “[Output Data Sets](#)” on page 2606 for details. You can specify the following options to include additional information in the OUT= data set: [OUTCUM](#), [OUTEXPECT](#), and [OUTPCT](#).

OUTCUM

includes cumulative frequencies and cumulative percentages in the [OUT=](#) data set for one-way tables. The variable CUM_FREQ contains the cumulative frequencies, and the variable CUM_PCT contains the cumulative percentages. See the section “[Output Data Sets](#)” on page 2606 for details. The OUTCUM option has no effect for two-way or multiway tables.

OUTEXPECT

includes expected cell frequencies in the [OUT=](#) data set for crosstabulation tables. The variable EXPECTED contains the expected cell frequencies. See the section “[Output Data Sets](#)” on page 2606 for details. The EXPECTED option has no effect for one-way tables.

OUTPCT

includes the following additional variables in the [OUT=](#) data set for crosstabulation tables:

PCT_COL	percentage of column frequency
PCT_ROW	percentage of row frequency
PCT_TABL	percentage of stratum (two-way table) frequency, for n -way tables where $n > 2$

See the section “[Output Data Sets](#)” on page 2606 for details. The OUTPCT option has no effect for one-way tables.

PLCORR

requests the polychoric correlation coefficient. For 2×2 tables, this statistic is more commonly known as the tetrachoric correlation coefficient, and it is labeled as such in the displayed output. See the section “[Polychoric Correlation](#)” on page 2562 for details.

The PLCORR option provides the polychoric correlation coefficient and its asymptotic standard error. If you also specify the [CL](#) option, PROC FREQ provides confidence limits for the polychoric correlation. You can specify the [PLCORR](#) option in the TEST statement, to request Wald and likelihood ratio tests for the polychoric correlation. If you omit the [MEASURES](#) option, the PLCORR option invokes the MEASURES option.

The [CONVERGE=](#) and [MAXITER=](#) options control the iterative computation of the polychoric correlation coefficient. The CONVERGE= option specifies the convergence criterion and the MAXITER= option specifies the maximum number of iterations. By default CONVERGE=0.0001 and MAXITER=20. Iterative computation of the polychoric correlation stops when the number of iterations exceeds the MAXITER= value or when the convergence measures falls below the CONVERGE= value, whichever happens first.

PLOTS < (*global-plot-options*) > < = *plot-request* < (*plot-options*) > >

PLOTS < (*global-plot-options*) >

< = (*plot-request* < (*plot-options*) > < ... *plot-request* < (*plot-options*) > >) >

controls the plots that are produced through ODS Graphics. *Plot-requests* identify the plots, and *plot-options* control the appearance and content of the plots. You can specify *plot-options* in parentheses after a *plot-request*. A *global-plot-option* applies to all plots for which it is available unless it is altered by a specific *plot-option*. You can specify *global-plot-options* in parentheses after the PLOTS option.

When you specify only one *plot-request*, you can omit the parentheses around the request. For example:

```
plots=all
plots=freqplot
plots=(freqplot oddsrationplot)
plots(only)=(cumfreqplot deviationplot)
```

ODS Graphics must be enabled before plots can be requested. For example:

```
ods graphics on;
proc freq;
  tables treatment*response / chisq plots=freqplot;
  weight wt;
run;
ods graphics off;
```

For more information about enabling and disabling ODS Graphics, see the section “[Enabling and Disabling ODS Graphics](#)” on page 600 in Chapter 21, “[Statistical Graphics Using ODS](#).”

If ODS Graphics is enabled but you do not specify the PLOTS= option, PROC FREQ produces all plots that are associated with the analyses that you request with the exception of the frequency, cumulative frequency, and mosaic plots. To produce a frequency plot or cumulative frequency plot when ODS Graphics is enabled, you must specify the [FREQPLOT](#) or [CUMFREQPLOT](#) *plot-request*, respectively, in the PLOTS= option, or you must specify the [PLOTS=ALL](#) option. To produce a mosaic plot when ODS Graphics is enabled, you must specify the [MOSAICPLOT](#) *plot-request* in the PLOTS= option, or you must specify the [PLOTS=ALL](#) option.

PROC FREQ produces the remaining plots (listed in [Table 38.11](#)) by default when you request the corresponding TABLES statement options. You can suppress default plots and request specific plots by using the [PLOTS\(ONLY\)=](#) option; [PLOTS\(ONLY\)=\(*plot-requests*\)](#) produces only the plots that are specified as *plot-requests*. You can suppress all plots by specifying the [PLOTS=NONE](#) option. The PLOTS= option has no effect when you specify the [NOPRINT](#) option in the [PROC FREQ](#) statement.

Plot Requests

[Table 38.11](#) lists the available *plot-requests* together with their required TABLES statement options. Descriptions of the *plot-requests* follow the table in alphabetical order.

Table 38.11 Plot Requests

Plot Request	Description	Required TABLES Statement Option
AGREEPLOT	Agreement plot	AGREE ($r \times r$ table)
ALL	All plots	None
CUMFREQPLOT	Cumulative frequency plot	One-way table request
DEVIATIONPLOT	Deviation plot	CHISQ (one-way table)
FREQPLOT	Frequency plot	Any table request
KAPPAPLOT	Kappa plot	AGREE ($h \times r \times r$ table)
MOSAICPLOT	Mosaic plot	Two-way or multiway table request
NONE	No plots	None
ODDSRATIOPLOT	Odds ratio plot	MEASURES or RELRISK ($h \times 2 \times 2$ table)
RELRIKSPLOT	Relative risk plot	MEASURES or RELRISK ($h \times 2 \times 2$ table)
RISKDIFFPLOT	Risk difference plot	RISKDIFF ($h \times 2 \times 2$ table)
WTKAPPAPLOT	Weighted kappa plot	AGREE ($h \times r \times r$ table, $r > 2$)

The following *plot-requests* are available with the PLOTS= option:

AGREEPLOT < (*plot-options*) >

requests an agreement plot (Bangdiwala and Bryan 1987). An agreement plot displays the strength of agreement in a two-way table, where the row and column variables represent two independent ratings of n subjects. For information about agreement plots, see Bangdiwala (1988); Bangdiwala et al. (2008), and Friendly (2000, Section 3.7.2).

To produce an agreement plot, you must also specify the **AGREE** option in the TABLES statement. Agreement statistics and plots are available for two-way square tables, where the number of rows equals the number of columns.

Table 38.12 lists the *plot-options* that are available for agreement plots. See the subsection “Plot Options” for descriptions of the *plot-options*.

Table 38.12 Plot Options for AGREEPLOT

Plot Option	Description	Values
LEGEND=	Legend	YES* or NO
PARTIAL=	Partial agreement	YES* or NO
SHOWSCALE=	Frequency scale	YES* or NO
STATS	Statistics	None

*Default

ALL

requests all plots that are associated with the specified analyses. If you specify the PLOTS=ALL option, PROC FREQ produces the frequency, cumulative frequency, and mosaic plots that are associated with the tables that you request. (PROC FREQ does not produce these plots by default when ODS Graphics is enabled.)

CUMFREQPLOT <(plot-options)>

requests a plot of cumulative frequencies. Cumulative frequency plots are available for one-way frequency tables.

To produce a cumulative frequency plot, you must specify the CUMFREQPLOT *plot-request* in the PLOTS= option, or you must specify the **PLOTS=ALL** option. PROC FREQ does not produce cumulative frequency plots by default when ODS Graphics is enabled.

Table 38.13 lists the *plot-options* that are available for cumulative frequency plots. See the subsection “Plot Options” for descriptions of the *plot-options*.

Table 38.13 Plot Options for CUMFREQPLOT

Plot Option	Description	Values
ORIENT=	Orientation	VERTICAL* or HORIZONTAL
SCALE=	Scale	FREQ* or PERCENT
TYPE=	Type	BARCHART* or DOTPLOT

*Default

DEVIATIONPLOT <(plot-options)>

requests a plot of relative deviations from expected frequencies. Deviation plots are available for chi-square analysis of one-way frequency tables. To produce a deviation plot, you must also specify the **CHISQ** option in the TABLES statement for a one-way frequency table.

Table 38.14 lists the *plot-options* that are available for deviation plots. See the subsection “Plot Options” for descriptions of the *plot-options*.

Table 38.14 Plot Options for DEVIATIONPLOT

Plot Option	Description	Values
NOSTAT	No statistic	None
ORIENT=	Orientation	VERTICAL* or HORIZONTAL
TYPE=	Type	BARCHART* or DOTPLOT

*Default

FREQPLOT <(plot-options)>

requests a frequency plot. Frequency plots are available for frequency and crosstabulation tables. For multiway crosstabulation tables, PROC FREQ provides a two-way frequency plot for each stratum (two-way table).

To produce a frequency plot, you must specify the FREQPLOT *plot-request* in the PLOTS= option, or you must specify the **PLOTS=ALL** option. PROC FREQ does not produce frequency plots by default when ODS Graphics is enabled.

By default, PROC FREQ displays frequency plots as bar charts. You can specify the **TYPE=DOTPLOT** *plot-option* to display frequency plots as dot plots. You can plot percentages instead of frequencies by specifying the **SCALE=PERCENT** *plot-option*. There are four frequency plot layouts available, which you can request by specifying the **TWOWAY=** *plot-option*. See the subsection “Plot Options” for more information.

By default, the primary grouping of graph cells in a two-way layout is by column variable. Row variable levels are then displayed within column variable levels. You can specify the **GROUPBY=ROW** *plot-option* to group first by row variable.

Table 38.15 lists the *plot-options* that are available for frequency plots. See the subsection “Plot Options” for descriptions of the *plot-options*.

The following *plot-options* are available for all frequency plots: **ORIENT=**, **SCALE=**, and **TYPE=**. The following *plot-options* are available for frequency plots of two-way (and multi-way) tables: **GROUPBY=**, **NPANELPOS=**, and **TWOWAY=**. The **NPANELPOS=** *plot-option* is not available with the **TWOWAY=CLUSTER** or **TWOWAY=STACKED** layout, which is always displayed in a single panel.

Table 38.15 Plot Options for FREQPLOT

Plot Option	Description	Values
GROUPBY=**	Primary group	COLUMN* or ROW
NPANELPOS=**	Sections per panel	Number (4*)
ORIENT=	Orientation	VERTICAL* or HORIZONTAL
SCALE=	Scale	FREQ* or PERCENT
TWOWAY=**	Two-way layout	CLUSTER, GROUPHORIZONTAL, GROUPVERTICAL*, or STACKED
TYPE=	Type	BARCHART* or DOTPLOT

*Default

**For two-way tables

KAPPAPLOT < (*plot-options*) >

requests a plot of kappa statistics with confidence limits. Kappa plots are available for multiway square tables and display the kappa statistic (with confidence limits) for each two-way table (stratum). Kappa plots also display the overall kappa statistic unless you specify the **COMMON=NO** *plot-option*. To produce a kappa plot, you must specify the **AGREE** option in the **TABLES** statement to compute kappa statistics.

Table 38.16 lists the *plot-options* that are available for kappa plots. See the subsection “Plot Options” for descriptions of the *plot-options*.

Table 38.16 Plot Options for KAPPAPLOT and WTKAPPAPLOT

Plot Option	Description	Values
CLDISPLAY=	Error bar type	SERIF*, LINE, or BAR
COMMON=	Overall kappa	YES* or NO
NPANELPOS=	Statistics per graphic	Number
ORDER=	Order of two-way levels	ASCENDING or DESCENDING
RANGE=	Range to display	Values or CLIP
STATS	Statistics	None

*Default

MOSAICPLOT < (*plot-option*) >

requests a mosaic plot. Mosaic plots are available for crosstabulation tables. For multiway tables, PROC FREQ provides a mosaic plot for each two-way table (stratum).

To produce a mosaic plot, you must specify the MOSAICPLOT *plot-request* in the PLOTS= option, or you must specify the **PLOTS=ALL** option. PROC FREQ does not produce mosaic plots by default when ODS Graphics is enabled.

Mosaic plots display tiles that correspond to the crosstabulation table cells. The areas of the tiles are proportional to the frequencies of the table cells. The column variable is displayed on the X axis, and the tile widths are proportional to the relative frequencies of the column variable levels. The row variable is displayed on the Y axis, and the tile heights are proportional to the relative frequencies of the row levels within column levels. The colors of the tiles correspond to the row variable levels. See Friendly (2000) for more information.

You can specify the following *plot-option* in parentheses after the MOSAICPLOT *plot-request*:

SQUARE

produces a square mosaic plot, where the height of the Y axis equals the width of the X axis. In a square mosaic plot, the scale of the relative frequencies is the same on both axes. By default, PROC FREQ produces a rectangular mosaic plot.

NONE

suppresses all plots.

ODDSRATIOPLOT < (*plot-options*) >

requests a plot of odds ratios with confidence limits. Odds ratio plots are available for multiway 2×2 tables and display the odds ratio (with confidence limits) for each 2×2 table (stratum). To produce an odds ratio plot, you must also specify the **MEASURES** or **RELRISK** option in the TABLES statement to compute odds ratios.

If you specify the **CMH** option in the TABLES statement to compute the common odds ratio for the multiway table, by default the plot displays the common odds ratio. The **COMMON=NO** *plot-option* suppresses display of the common odds ratio. If you do not specify the **CMH** option, the common odds ratio is not available to be displayed in the plot.

By default, odds ratio plots display asymptotic confidence limits. You can specify the **CL=EXACT** *plot-option* to display exact confidence limits for the odds ratios. You must also request computation of exact confidence limits by specifying the **OR** option in the EXACT statement. To display the common odds ratio with exact confidence limits, you must also specify the **COMOR** option in the EXACT statement.

Table 38.17 lists the *plot-options* that are available for odds ratio plots. See the subsection “Plot Options” for descriptions of the *plot-options*.

Table 38.17 Plot Options for ODDSRATIOPLOT, RELRISKPLOT, and RISKDIFFPLOT

Plot Option	Description	Values
CL=	Confidence limit type	Type
CLDISPLAY=	Error bar type	SERIF*, LINE, or BAR
COLUMN=**	Risk column	1* or 2
LOGBASE=***	Axis scale	2, E, or 10
NPANELPOS=	Statistics per graphic	Number (All*)
ORDER=	Order of two-way levels	ASCENDING or DESCENDING
RANGE=	Range to display	Values or CLIP
STATS	Statistics	None

*Default

** Available for RELRISKPLOT and RISKDIFFPLOT

*** Available for ODDSRATIOPLOT and RELRISKPLOT

RELRISKPLOT <(plot-options)>

requests a plot of relative risks with confidence limits. Relative risk plots are available for multiway 2×2 tables and display the relative risk (with confidence limits) for each 2×2 table (stratum). To produce a relative risk plot, you must also specify the **MEASURES** or **RELRISK** option in the TABLES statement to compute relative risks.

If you specify the **CMH** option in the TABLES statement to compute the common relative risk for the multiway table, by default the plot displays the common relative risk. The **COMMON=NO** *plot-option* suppresses display of the common relative risk. If you do not specify the **CMH** option, the common relative risk is not available to be displayed in the plot.

By default, relative risk plots display asymptotic confidence limits. You can specify the **CL=EXACT** *plot-option* to display exact confidence limits for the relative risks. You must also request computation of exact confidence limits by specifying the **RELRISK** option in the EXACT statement. When you specify the **CL=EXACT** *plot-option*, the relative risk plot does not display the common relative risk.

Table 38.17 lists the *plot-options* that are available for relative risk plots. See the subsection “Plot Options” for descriptions of the *plot-options*.

RISKDIFFPLOT <(plot-options)>

requests a plot of risk (proportion) differences with confidence limits. Risk difference plots are available for multiway 2×2 tables and display the risk difference (with confidence limits) for each 2×2 table (stratum). To produce a risk difference plot, you must also specify the **RISKDIFF** option in the TABLES statement to compute risk differences.

By default, risk difference plots display asymptotic Wald confidence limits. You can specify the **CL=EXACT** *plot-option* to display exact confidence limits for the risk differences. You must also request computation of exact confidence limits by specifying the **RISKDIFF** option in the EXACT statement.

In addition to Wald and exact confidence limits, you can display the following confidence limit types in a risk difference plot: Agresti-Caffo, Hauck-Anderson, Miettinen-Nurminen, and Newcombe. See the descriptions of the **CL=EXACT** *plot-option* and the **RISKDIFF** option for more information.

Table 38.17 lists the *plot-options* that are available for risk difference plots. See the subsection “Plot Options” for descriptions of the *plot-options*.

WTKAPPAPLOT <(plot-options)>

requests a plot of weighted kappa statistics with confidence limits. Weighted kappa plots are available for multiway square tables and display the weighted kappa statistic (with confidence limits) for each two-way table (stratum). Weighted kappa plots also display the overall weighted kappa statistic unless you specify the **COMMON=NO** *plot-option*.

To produce a weighted kappa plot, you must also specify the **AGREE** option in the TABLES statement to compute weighted kappa statistics. Simple kappa and weighted kappa statistics are equal for 2×2 tables; therefore, PROC FREQ displays weighted kappa statistics and plots only for tables that are larger than 2×2 .

Table 38.16 lists the *plot-options* that are available for weighted kappa plots. See the subsection “Plot Options” for descriptions of the *plot-options*.

Global Plot Options

A *global-plot-option* applies to all plots for which the option is available unless it is altered by an individual *plot-option*. You can specify *global-plot-options* in parentheses after the PLOTS option. For example:

```
plots(order=ascending stats)=(riskdiffplot oddsratioplot)
plots(only)=freqplot
```

The following *plot-options* are available as *global-plot-options*: **CLDISPLAY=**, **COLUMN=**, **COMMON=**, **EXACT**, **LOGBASE=**, **NPANELPOS=**, **ORDER=**, **ORIENT=**, **RANGE=**, **SCALE=**, **STATS**, and **TYPE=**. See the subsection “Plot Options” for description of these *plot-options*.

In addition to these *plot-options*, you can specify the following *global-plot-option* in parentheses after the PLOTS option:

ONLY

suppresses the default plots and requests only the plots that are specified as *plot-requests*.

Plot Options

You can specify the following *plot-options* in parentheses after a *plot-request*:

CL=type

specifies the type of confidence limits to display. The **CL=** *plot-option* is available for the following plots: **ODDSRATIOPLOT**, **RELRIKSPLOT**, and **RISKDIFFPLOT**.

For odds ratio and relative risk plots, the available confidence limit types are asymptotic and exact, which you can request by specifying **CL=ASYMPTOTIC** and **CL=EXACT**, respectively. By default, these plots display asymptotic confidence limits. When you specify **CL=EXACT** to display exact confidence limits, you must also request computation of exact confidence limits by specifying the corresponding option (**OR** or **RELRIKSPLOT**) in the EXACT statement. See the sections “Exact Confidence Limits for the Odds Ratio” on page 2583 and “Exact Unconditional Confidence Limits for the Relative Risk” on page 2584 for more information.

For risk difference plots, the available confidence limit types include the following: **CL=AC** (Agresti-Caffo), **CL=EXACT** (exact unconditional), **CL=HA** (Hauck-Anderson), **CL=MN**

(Miettinen-Nurminen), CL=NEWCOMBE (Newcombe), and CL=WALD (Wald). See the description of the **CL= riskdiff-option** and see the section “Risk Difference Confidence Limits” on page 2574 for more information. If you specify CL=EXACT to display exact confidence limits in the plot, you must also request computation of exact confidence limits by specifying the **RISKDIFF** option in the EXACT statement.

CLDISPLAY=SERIF | LINE | BAR < width >

controls the appearance of the confidence limit error bars. The CLDISPLAY= *plot-option* is available for the following plots: **KAPPAPLOT**, **ODDSRATIO**PLOT, **RELRISK**PLOT, **RISKDIFF**PLOT, and **WTKAPP**PLOT.

The default value is CLDISPLAY=SERIF, which displays the confidence limits as lines with serifs. CLDISPLAY=LINE displays the confidence limits as plain lines without serifs.

CLDISPLAY=BAR displays the confidence limits as bars. By default, the width of the bars equals the size of the marker for the estimate. You can control the width of the bars and the size of the marker by specifying the value of *width* as a percentage of the distance between bars, $0 < width \leq 1$. The bar might disappear when the value of *width* is very small.

COLUMN=1 | 2

specifies the 2×2 table column to use to compute the risk (proportion). The COLUMN= *plot-option* is available for the relative risk plot (**RELRISK**PLOT) and the risk difference plot (**RISKDIFF**PLOT). If you specify COLUMN=1, the plot displays the column 1 relative risks or the column 1 risk difference. Similarly, if you specify COLUMN=2, the plot displays the column 2 relative risks or risk difference. The default is COLUMN=1.

COMMON=YES | NO

controls the display of the common (overall) statistic in plots that display stratum (two-way table) statistics for a multiway table. The COMMON= *plot-option* is available for the following plots: **KAPPAPLOT**, **ODDSRATIO**PLOT, **RELRISK**PLOT, and **WTKAPP**PLOT.

By default, COMMON=YES and the plots display the common value when it is available. When you specify the **AGREE** option in the TABLES statement for a multiway table, PROC FREQ computes overall kappa statistics in addition to kappa statistics for each stratum. When you specify the **MEASURES** or **RELRISK** option in the TABLES statement, PROC FREQ computes only stratum-level odds ratios and relative risks; to compute common odds ratios and relative risks, you must also specify the **CMH** option in the TABLES statement.

COMMON=YES is not available for relative risk plots when you specify the CL=EXACT *plot-option* to display exact confidence limits.

EXACT

requests display of exact confidence limits instead of asymptotic confidence limits. The EXACT *plot-option* is available for the odds ratio plot (**ODDSRATIO**PLOT), relative risk plot (**RELRISK**PLOT), and risk difference plot (**RISKDIFF**PLOT). The EXACT *plot-option* is equivalent to the **CL=EXACT** *plot-option*.

When you specify the EXACT *plot-option*, you must also request computation of exact confidence limits by specifying the appropriate *statistic-option* in the **EXACT** statement.

GROUPBY=COLUMN | ROW

specifies the primary grouping for two-way frequency plots. The default is **GROUPBY=COLUMN**, which groups graph cells first by column variable and displays row variable levels within column variable levels. You can specify **GROUPBY=ROW** to group first by row variable. In two-way and multiway table requests, the column variable is the last variable specified and forms the columns of the crosstabulation table. The row variable is the next-to-last variable specified and forms the rows of the table.

By default for a bar chart that is displayed in the **TWOWAY=STACKED** layout, bars correspond to the column variable levels, and row levels are displayed (stacked) within each column bar. By default for a bar chart that is displayed in the **TWOWAY=CLUSTER** layout, bars are first grouped by column variable levels, and row levels are displayed as adjacent bars within each column-level group. You can reverse the default row and column variable grouping by specifying **GROUPBY=ROW**.

LOGBASE=2 | E | 10

applies to the odds ratio plot (**ODDSRATIOPLOT**) and the relative risk plot (**RELRIKSPLOT**). **LOGBASE=** displays the odds ratio or relative risk axis on the log scale that you specify.

LEGEND=YES | NO

applies to the agreement plot (**AGREEPLOT**). **LEGEND=NO** suppresses the legend that identifies the areas of exact and partial agreement. The default is **LEGEND=YES**.

NOSTAT

applies to the deviation plot (**DEVIATIONPLOT**). **NOSTAT** suppresses the chi-square *p*-value that deviation plot displays by default.

NPANELPOS=*n*

applies to the following plots: **FREQPLOT** (for two-way and multiway tables), **KAPPAPLOT**, **ODDSRATIOPLOT**, **RELRIKSPLOT**, **RISKDIFFPLOT**, and **WTKAPPAPLOT**.

For kappa, odds ratio, relative risk, risk difference, and weighted kappa plots, **NPANELPOS=** divides the plot into multiple panels that display at most $|n|$ statistics per panel. If *n* is positive, the number of statistics per panel is balanced; if *n* is negative, the number of statistics per panel is not balanced. For example, suppose you want to display 21 odds ratios. **NPANELPOS=20** displays two panels, the first with 11 odds ratios and the second with 10 odds ratios; **NPANELPOS=-20** displays 20 odds ratios in the first panel but only one in the second panel. By default, *n* = 0 and all statistics are displayed in a single panel.

For two-way frequency plots, the **NPANELPOS= plot-option** divides the plot into multiple panels that display at most $|n|$ sections per panel. Sections correspond to row or column variable levels, depending on the type of plot and the grouping.

The **NPANELPOS= plot-option** applies to two-way plots that are displayed in the **TWOWAY=GROUPVERTICAL** or **TWOWAY=GROUPHORIZONTAL** layout. The **NPANELPOS= plot-option** does not apply to the **TWOWAY=CLUSTER** and **TWOWAY=STACKED** layouts, which are always displayed in a single panel.

If *n* is positive, the number of sections per panel is balanced; if *n* is negative, the number of sections per panel is not balanced. By default for frequency plots, *n* = 4 and each panel includes at most four sections.

ORDER=ASCENDING | DESCENDING

displays the statistics in sorted order. By default, the statistics are displayed in the order in which the two-way table strata appear in the multiway table display. The ORDER= *plot-option* applies to the following plots: [KAPPAPLOT](#), [ODDSRATIOPLOT](#), [RELRISKPLOT](#), [RISKDIFFPLOT](#), and [WTKAPPAPLOT](#).

ORIENT=HORIZONTAL | VERTICAL

controls the orientation of the plot. The ORIENT= *plot-option* applies to the following plots: [CUMFREQPLOT](#), [DEVIATIONPLOT](#), and [FREQPLOT](#).

ORIENT=HORIZONTAL places the variable levels on the Y axis and the frequencies, percentages, or statistic values on the X axis. ORIENT=VERTICAL places the variable levels on the X axis. The default orientation is ORIENT=VERTICAL for bar charts ([TYPE=BARCHART](#)) and ORIENT=HORIZONTAL for dot plots ([TYPE=DOTPLOT](#)).

PARTIAL=YES | NO

controls the display of partial agreement in the agreement plot ([AGREEPLOT](#)). PARTIAL=NO suppresses the display of partial agreement. When you specify PARTIAL=NO, the agreement plot displays only exact agreement. Exact agreement includes the diagonal cells of the square table, where the row and column variable levels are the same. Partial agreement includes the adjacent off-diagonal table cells, where the row and column values are within one level of exact agreement. The default is PARTIAL=YES.

RANGE=(*< min>* *< ,max>*) | CLIP

specifies the range of values to display. The RANGE= *plot-option* applies to the following plots: [KAPPAPLOT](#), [ODDSRATIOPLOT](#), [RELRISKPLOT](#), [RISKDIFFPLOT](#), and [WTKAPPAPLOT](#). If you specify RANGE=CLIP, the confidence limits are clipped and the display range is determined by the minimum and maximum values of the statistics. By default, the display range includes all confidence limits.

SCALE=FREQ | LOG | PERCENT | SQRT

specifies the scale of the frequencies to display. The SCALE= *plot-option* applies to the frequency plot ([FREQPLOT](#)) and the cumulative frequency plot ([CUMFREQPLOT](#)).

The default is SCALE=FREQ, which displays unscaled frequencies. SCALE=LOG displays log (base 10) frequencies. SCALE=PERCENT displays percentages (relative frequencies). SCALE=SQRT displays square roots of the frequencies, which produces a plot known as a *rootogram*.

SHOWSCALE=YES | NO

controls the display of the cumulative frequency scale on the right side of the agreement plot ([AGREEPLOT](#)). SHOWSCALE=NO suppresses the display of the scale. The default is SHOWSCALE=YES.

SQUARE

requests a square mosaic plot ([MOSAICPLOT](#)), where the height of the Y axis equals the width of the X axis. In a square mosaic plot, the scale of the relative frequencies is the same on both axes. By default, the mosaic plot area is rectangular.

STATS

displays statistic values in the plot. For the following plots, the *STATS plot-option* displays the statistics and their confidence limits on the right side of the plot: [KAPPAPLOT](#), [ODDSRATIOPLLOT](#), [RELRISKPLLOT](#), [RISKDIFFPLLOT](#), and [WTKAPPAPLOT](#).

For the agreement plot ([AGREEPLLOT](#)), *STATS* displays the values of the kappa statistic, the weighted kappa statistic, and the B_n measure (Bangdiwala and Bryan 1987).

If you do not request the *STATS plot-option*, these plots do not display the statistic values.

TWOWAY=CLUSTER | GROUPTHORIZONTAL | GROUPTHVERTICAL | STACKED

specifies the layout for two-way frequency plots.

All *TWOWAY=* layouts are available for bar charts (*TYPE=BARCHART*). All *TWOWAY=* layouts except *TWOWAY=CLUSTER* are available for dot plots (*TYPE=DOTPLOT*). The *ORIENT=* and *GROUPBY=* *plot-options* are available for all *TWOWAY=* layouts.

The default two-way layout is *TWOWAY=GROUPTHVERTICAL*, which produces a grouped plot that has a vertical common baseline. By default for bar charts (*TYPE=BARCHART*, *ORIENT=VERTICAL*), the X axis displays column variable levels, and the Y axis displays frequencies. The plot includes a vertical (Y-axis) block for each row variable level. The relative positions of the graph cells in this plot layout are the same as the relative positions of the table cells in the crosstabulation table. You can reverse the default row and column grouping by specifying the *GROUPBY=ROW* *plot-option*.

The *TWOWAY=GROUPTHORIZONTAL* layout produces a grouped plot that has a horizontal common baseline. By default (*GROUPBY=COLUMN*), the plot displays a block on the X axis for each column variable level. Within each column-level block, the plot displays row variable levels.

The *TWOWAY=STACKED* layout produces stacked displays of frequencies. By default (*GROUPBY=COLUMN*) in a stacked bar chart, the bars correspond to column variable levels, and row levels are stacked within each column level. By default in a stacked dot plot, the dotted lines correspond to column levels, and cell frequencies are plotted as data dots on the corresponding column line. The dot color identifies the row level.

The *TWOWAY=CLUSTER* layout, which is available only for bar charts, displays groups of adjacent bars. By default, the primary grouping is by column variable level, and row levels are displayed within each column level.

You can reverse the default row and column grouping in any layout by specifying the *GROUPBY=ROW* *plot-option*. The default is *GROUPBY=COLUMN*, which groups first by column variable.

TYPE=BARCHART | DOTPLOT

specifies the plot type for frequency (*FREQPLOT*), cumulative frequency (*CUMFREQPLOT*), and deviation plots (*DEVIATIONPLOT*). *TYPE=BARCHART* produces a bar chart and *TYPE=DOTPLOT* produces a dot plot. The default is *TYPE=BARCHART*.

PRINTKWT

displays the weights that PROC FREQ uses to compute the weighted kappa coefficient. You must also specify the **AGREE** option to request the weighted kappa coefficient. You can specify (WT=FC) with the **AGREE** option to request Fleiss-Cohen weights. By default, PROC FREQ uses Cicchetti-Allison weights to compute the weighted kappa coefficient. See the section “[Weighted Kappa Coefficient](#)” on page 2590 for details.

RELRISK | OR

requests relative risk measures and their confidence limits for 2×2 tables. These measures include the odds ratio and the column 1 and 2 relative risks. See the section “[Odds Ratio and Relative Risks for \$2 \times 2\$ Tables](#)” on page 2582 for details.

You can also obtain the RELRISK measures by specifying the **MEASURES** option, which produces other measures of association in addition to the relative risks.

You can request exact confidence limits for the odds ratio by specifying the **OR** option in the **EXACT** statement. You can request exact unconditional confidence limits for the relative risks by specifying the **RELRISK** option in the **EXACT** statement. See the sections “[Exact Confidence Limits for the Odds Ratio](#)” on page 2583 and “[Exact Unconditional Confidence Limits for the Relative Risk](#)” on page 2584 for more information.

RISKDIFF <(riskdiff-options)>

requests risks (binomial proportions) and risk differences for 2×2 tables. When you specify the **RISKDIFF** option, PROC FREQ provides the row 1 risk, row 2 risk, total (overall) risk, and risk difference (row 1 – row 2), together with their asymptotic standard errors and Wald confidence limits. PROC FREQ also provides exact (Clopper-Pearson) confidence limits for the row 1, row 2, and total risks by default. You can request exact unconditional confidence limits for the risk difference by specifying the **RISKDIFF** option in the **EXACT** statement. See the section “[Risks and Risk Differences](#)” on page 2572 for details. PROC FREQ displays these results in the column 1 and column 2 “Risk Estimates” tables.

You can specify *riskdiff-options* in parentheses after the **RISKDIFF** option to request tests and additional confidence limits for the risk difference. [Table 38.18](#) summarizes the *riskdiff-options*.

The **EQUIV**, **NONINF**, and **SUP** *riskdiff-options* request tests of equivalence, noninferiority, and superiority, respectively, for the risk difference. Available test methods include Farrington-Manning (score), Hauck-Anderson, and Newcombe (hybrid-score), in addition to the Wald test.

As part of the noninferiority, superiority, and equivalence analyses, PROC FREQ provides test-based confidence limits that have a confidence coefficient of $100(1 - 2\alpha)\%$ (Schuirmann 1999). The **ALPHA=** option determines the confidence level, and the default of **ALPHA=0.05** produces 90% confidence limits. See the sections “[Noninferiority Tests](#)” on page 2577 and “[Equivalence Tests](#)” on page 2580 for details.

The **CL=** *riskdiff-option* requests confidence limits for the risk difference. Available confidence limit types include Agresti-Caffo, exact unconditional, Hauck-Anderson, Miettinen-Nurminen, Newcombe, and Wald. Continuity-corrected Newcombe and Wald confidence limits are also available. You can request more than one type of confidence limits in the same analysis. PROC FREQ displays the confidence limits in the “Proportion (Risk) Difference Confidence Limits” table.

The **ALPHA=** option determines the level of the confidence limits that the **CL= riskdiff-option** provides. The default of **ALPHA=0.05** produces 95% confidence limits for the risk difference.

The **CL=EXACT riskdiff-option** displays exact unconditional confidence limits in the “Proportion (Risk) Difference Confidence Limits” table. When you use **CL=EXACT**, you must also request computation of the exact confidence limits by specifying the **RISKDIFF** option in the **EXACT** statement.

Table 38.18 RISKDIFF (Proportion Difference) Options

Option	Description
COLUMN=1 2	Specifies the risk column
CORRECT	Requests continuity correction
NORISKS	Suppresses default risk tables
Request Confidence Limits	
CL=AC	Requests Agresti-Caffo confidence limits
CL=EXACT	Displays exact confidence limits
CL=HA	Requests Hauck-Anderson confidence limits
CL=MN	Requests Miettinen-Nurminen confidence limits
CL=NEWCOMBE	Requests Newcombe confidence limits
CL=WALD	Requests Wald confidence limits
Request Tests	
EQUAL	Requests an equality test
EQUIV EQUIVALENCE	Requests an equivalence test
NONINF NONINFERIORITY	Requests a noninferiority test
SUP SUPERIORITY	Requests a superiority test
MARGIN=	Specifies the test margin
METHOD=	Specifies the test method
VAR=SAMPLE NULL	Specifies the test variance

You can specify the following *riskdiff-options* in parentheses after the **RISKDIFF** option:

CL=type | (types)

requests confidence limits for the risk difference. You can specify one or more *types* of confidence limits. When you specify only one *type*, you can omit the parentheses around the request. PROC FREQ displays the confidence limits in the “Proportion (Risk) Difference Confidence Limits” table.

The **ALPHA=** option determines the confidence level, and the default of **ALPHA=0.05** produces 95% confidence limits for the risk difference. This differs from the test-based confidence limits that are provided with the equivalence, noninferiority, and superiority tests (**EQUIV**, **NONINF**, and **SUP**), which have a confidence coefficient of $100(1 - 2\alpha)\%$ (Schuirmann 1999).

You can specify **CL=** with or without requests for risk difference tests. The confidence limits produced by **CL=** do not depend on the tests that you request and do not use the value of the test margin (which you specify by using **MARGIN= riskdiff-option**).

You can control the risk column for the confidence limits by specifying the **COLUMN= riskdiff-option**. If you do not specify **COLUMN=**, PROC FREQ provides confidence limits for the column 1 risk difference by default.

The following *types* of confidence limits are available:

AC | AGRESTICAFFO

requests Agresti-Caffo confidence limits for the risk difference. See the subsection “Agresti-Caffo Confidence Limits” in the section “Risk Difference Confidence Limits” on page 2574 for details.

EXACT

displays exact unconditional confidence limits for the risk difference in the “Proportion (Risk) Difference Confidence Limits” table. You must also request computation of the exact confidence limits by specifying the **RISKDIFF** option in the **EXACT** statement.

PROC FREQ computes the confidence limits by inverting two separate one-sided exact tests (tail method), where the tests are based on the unstandardized risk difference by default. If you specify the **RISKDIFF(METHOD=SCORE)** option in the **EXACT** statement, the tests are based on the score statistic. See the **RISKDIFF** option in the **EXACT** statement and the section “Exact Unconditional Confidence Limits for the Risk Difference” on page 2580 for more information.

By default, PROC FREQ also displays these exact confidence limits in the “Risk Estimates” table. You can suppress this table by specifying the **NORISKS riskdiff-option**.

HA

requests Hauck-Anderson confidence limits for the risk difference. See the subsection “Hauck-Anderson Confidence Limits” in the section “Risk Difference Confidence Limits” on page 2574 for details.

MN < (CORRECT=NO | MEE) >

requests Miettinen-Nurminen confidence limits for the risk difference. See the subsection “Miettinen-Nurminen Confidence Limits” in the section “Risk Difference Confidence Limits” on page 2574 for details. By default, the Miettinen-Nurminen confidence limits include a bias correction factor (Miettinen and Nurminen 1985; Newcombe and Nurminen 2011). If you specify **CL=MN(CORRECT=NO)**, PROC FREQ provides the uncorrected form of the confidence limits (Mee 1984).

NEWCOMBE | SCORE | WILSON <(CORRECT)>

requests Newcombe hybrid-score confidence limits for the risk difference. If you specify **CL=NEWCOMBE(CORRECT)** or the **CORRECT riskdiff-option**, the Newcombe confidence limits include a continuity correction. See the subsection “Newcombe Confidence Limits” in the section “Risk Difference Confidence Limits” on page 2574 for details.

WALD < (CORRECT) >

requests Wald confidence limits for the risk difference. If you specify **CL=WALD(CORRECT)** or the **CORRECT riskdiff-option**, the Wald confidence limits include a continuity correction. See the subsection “Wald Confidence Limits” in the section “Risk Difference Confidence Limits” on page 2574 for details.

COLUMN=1 | 2 | BOTH

specifies the table column for which to compute the risk difference tests (**EQUAL**, **EQUIV**, **NONINF**, and **SUP**) and the risk difference confidence limits (which are requested by the **CL= riskdiff-option**).

If you do not specify `COLUMN=`, PROC FREQ provides the risk difference tests and confidence limits for column 1 by default. The `COLUMN=` option has no effect on the “Risk Estimates” table, which is produced for both column 1 and column 2. You can suppress the “Risk Estimates” table by specifying the [NORISKS riskdiff-option](#).

CORRECT

includes a continuity correction in the Wald confidence limits, Wald tests, and Newcombe confidence limits. See the section “[Risks and Risk Differences](#)” on page 2572 for details.

EQUAL

requests a test of the null hypothesis that the risk difference equals zero. PROC FREQ provides an asymptotic Wald test of equality. If you specify the [CORRECT riskdiff-option](#), the Wald test includes a continuity correction. If you specify the [VAR=NULL riskdiff-option](#), the test uses the null (test-based) variance instead of the sample-based variance. See the section “[Equality Test](#)” on page 2576 for details.

EQUIV | EQUIVALENCE

requests a test of equivalence for the risk difference. See the section “[Equivalence Tests](#)” on page 2580 for details. You can specify the equivalence test margins with the [MARGIN= riskdiff-option](#) and the test method with the [METHOD= riskdiff-option](#). PROC FREQ uses `METHOD=WALD` by default.

MARGIN=value | (lower,upper)

specifies the margin for the noninferiority, superiority, and equivalence tests, which you request with the [NONINF](#), [SUP](#), and [EQUIV riskdiff-options](#), respectively. If you do not specify `MARGIN=`, PROC FREQ uses a margin of 0.2 by default.

For noninferiority and superiority tests, specify a single *value* for `MARGIN=`. The `MARGIN=value` must be a positive number. You can specify *value* as a number between 0 and 1. Or you can specify *value* in percentage form as a number between 1 and 100, and PROC FREQ converts that number to a proportion. The procedure treats the value 1 as 1%.

For an equivalence test, you can specify a single `MARGIN= value`, or you can specify both *lower* and *upper* values. If you specify a single `MARGIN= value`, it must be a positive number, as described previously. If you specify a single `MARGIN= value` for an equivalence test, PROC FREQ uses $-value$ as the lower margin and *value* as the upper margin for the test. If you specify both *lower* and *upper* values for an equivalence test, you can specify them in proportion form as numbers between -1 or 1 . Or you can specify them in percentage form as numbers between -100 and 100 , and PROC FREQ converts the numbers to proportions. The value of *lower* must be less than the value of *upper*.

METHOD=method

specifies the method for the noninferiority, superiority, and equivalence analyses, which you request with the [NONINF](#), [SUP](#), and [EQUIV riskdiff-options](#), respectively. If you do not specify the `METHOD= riskdiff-option`, PROC FREQ uses `METHOD=WALD` by default.

The following *methods* are available:

FM

requests Farrington-Manning (score) tests and test-based confidence limits for the equivalence, noninferiority, and superiority analyses. See the subsection “[Farrington-Manning Test](#)” in the section “[Noninferiority Tests](#)” on page 2577 for details.

HA

requests Hauck-Anderson tests and confidence limits for the equivalence, noninferiority, and superiority analyses. See the subsection “[Hauck-Anderson Test](#)” in the section “[Noninferiority Tests](#)” on page 2577 for details.

NEWCOMBE | SCORE | WILSON

requests Newcombe (hybrid-score) confidence limits for the equivalence, noninferiority, and superiority analyses. If you specify the [CORRECT riskdiff-option](#), the Newcombe confidence limits include a continuity correction. See the subsection “[Newcombe Noninferiority Analysis](#)” in the section “[Noninferiority Tests](#)” on page 2577 for details.

WALD

requests Wald tests and confidence limits for the equivalence, noninferiority, and superiority analyses. If you specify the [CORRECT riskdiff-option](#), the Wald confidence limits include a continuity correction. If you specify the [VAR=NULL riskdiff-option](#), the tests and confidence limits use the null (test-based) variance instead of the sample-based variance. See the subsection “[Wald Test](#)” in the section “[Noninferiority Tests](#)” on page 2577 for details.

NONINF | NONINFERIORITY

requests a test of noninferiority for the risk difference. See the section “[Noninferiority Tests](#)” on page 2577 for details. You can specify the test margin with the [MARGIN= riskdiff-option](#) and the test method with the [METHOD= riskdiff-option](#). PROC FREQ uses METHOD=WALD by default.

NORISKS

suppresses display of the “Risk Estimates” tables, which the RISKDIFF option produces by default for column 1 and column 2. The “Risk Estimates” tables contain the risks and risk differences, together with their asymptotic standard errors, Wald confidence limits, and exact confidence limits.

SUP | SUPERIORITY

requests a test of superiority for the binomial proportion. See the section “[Superiority Test](#)” on page 2579 for details. You can specify the test margin with the [MARGIN= riskdiff-option](#) and the test method with the [METHOD= riskdiff-option](#). PROC FREQ uses METHOD=WALD by default.

VAR=SAMPLE | NULL

specifies the type of variance estimate to use in the Wald tests of noninferiority, superiority, equivalence, and equality. The default is VAR=SAMPLE, which uses the sample-based variance. VAR=NULL uses a test-based variance that is computed from the null hypothesis risk difference value. See the sections “[Equality Test](#)” on page 2576 and “[Noninferiority Tests](#)” on page 2577 for details.

SCORES=type

specifies the type of row and column scores that PROC FREQ uses to compute the following statistics: Mantel-Haenszel chi-square, Pearson correlation, Cochran-Armitage test for trend, weighted kappa coefficient, and Cochran-Mantel-Haenszel statistics. The value of *type* can be one of the following:

- MODRIDIT
- RANK
- RIDIT
- TABLE

See the section “[Scores](#)” on page 2549 for descriptions of these score types.

If you do not specify the `SCORES=` option, PROC FREQ uses `SCORES=TABLE` by default. For character variables, the row and column TABLE scores are the row and column numbers. That is, the TABLE score is 1 for row 1, 2 for row 2, and so on. For numeric variables, the row and column TABLE scores equal the variable values. See the section “[Scores](#)” on page 2549 for details. Using MODRIDIT, RANK, or RIDIT scores yields nonparametric analyses.

You can use the `SCOROUT` option to display the row and column scores.

SCOROUT

displays the row and column scores that PROC FREQ uses to compute score-based tests and statistics. You can specify the score type with the `SCORES=` option. See the section “[Scores](#)” on page 2549 for details.

The scores are computed and displayed only when PROC FREQ computes statistics for two-way tables. You can use ODS to store the scores in an output data set. See the section “[ODS Table Names](#)” on page 2617 for more information.

SPARSE

reports all possible combinations of the variable values for an n -way table when $n > 1$, even if a combination does not occur in the data. The SPARSE option applies only to crosstabulation tables displayed in LIST format and to the `OUT=` output data set. If you do not use the `LIST` or `OUT=` option, the SPARSE option has no effect.

When you specify the SPARSE and LIST options, PROC FREQ displays all combinations of variable values in the table listing, including those with a frequency count of zero. By default, without the SPARSE option, PROC FREQ does not display zero-frequency levels in LIST output. When you use the SPARSE and `OUT=` options, PROC FREQ includes empty crosstabulation table cells in the output data set. By default, PROC FREQ does not include zero-frequency table cells in the output data set.

See the section “[Missing Values](#)” on page 2545 for more information.

TOTPCT

displays the percentage of the total multiway table frequency in crosstabulation tables for n -way tables, where $n > 2$. By default, PROC FREQ displays the percentage of the individual two-way table frequency but does not display the percentage of the total frequency for multiway crosstabulation tables. See the section “[Multiway Tables](#)” on page 2611 for more information.

The percentage of total multiway table frequency is displayed by default when you specify the `LIST` option. It is also provided by default in the PERCENT variable in the `OUT=` output data set.

TREND

requests the Cochran-Armitage test for trend. The table must be $2 \times C$ or $R \times 2$ to compute the trend test. See the section “[Cochran-Armitage Test for Trend](#)” on page 2585 for details. To request exact p -values for the trend test, specify the `TREND` option in the `EXACT` statement. See the section “[Exact Statistics](#)” on page 2601 for more information.

TEST Statement

TEST *test-options* ;

The TEST statement requests asymptotic tests for measures of association and measures of agreement. The *test-options* identify which tests to compute. Table 38.19 lists the available *test-options*, together with their corresponding TABLES statement options. Descriptions of the *test-options* follow the table in alphabetical order.

For each measure of association or agreement that you request in the TEST statement, PROC FREQ provides an asymptotic test that the measure equals zero. The procedure displays the asymptotic standard error under the null hypothesis, the test statistic, and the one-sided and two-sided *p*-values. Additionally, PROC FREQ reports the confidence limits for the measure. The ALPHA= option in the TABLES statement determines the confidence level, which by default equals 0.05 and provides 95% confidence limits. See the sections “Asymptotic Tests” on page 2556 and “Confidence Limits” on page 2556 for details. Also see the sections “Measures of Association” on page 2556 and “Tests and Measures of Agreement” on page 2588 for information about the individual measures.

You can also request exact tests for selected measures of association and agreement by using the EXACT statement. See the section “Exact Statistics” on page 2601 for more information.

Using the TEST Statement with the TABLES Statement

You must use a TABLES statement with the TEST statement. If you use only one TABLES statement, you do not need to specify the same options in both the TABLES and TEST statements; when you specify an option in the TEST statement, PROC FREQ automatically invokes the corresponding TABLES statement option. However, when you use the TEST statement with multiple TABLES statements, you must specify options in the TABLES statements to request statistics. PROC FREQ then provides asymptotic tests for those statistics that you also specify in the TEST statement.

Table 38.19 TEST Statement Options

Test Option	Asymptotic Tests	Required TABLES Statement Option
AGREE	Simple and weighted kappa coefficients	AGREE
GAMMA	Gamma	ALL or MEASURES
KAPPA	Simple kappa coefficient	AGREE
KENTB	Kendall's tau- <i>b</i>	ALL or MEASURES
MEASURES	Gamma, Kendall's tau- <i>b</i> , Stuart's tau- <i>c</i> , Somers' $D(C R)$, Somers' $D(R C)$, Pearson and Spearman correlations	ALL or MEASURES
PCORR	Pearson correlation coefficient	ALL or MEASURES
PLCORR	Polychoric correlation	PLCORR
SCORR	Spearman correlation coefficient	ALL or MEASURES
SMDCR	Somers' $D(C R)$	ALL or MEASURES
SMDRC	Somers' $D(R C)$	ALL or MEASURES
STUTC	Stuart's tau- <i>c</i>	ALL or MEASURES
WTKAP	Weighted kappa coefficient	AGREE

You can specify the following *test-options* in the TEST statement.

AGREE

requests asymptotic tests for the simple kappa coefficient and the weighted kappa coefficient. See the sections “[Simple Kappa Coefficient](#)” on page 2589 and “[Weighted Kappa Coefficient](#)” on page 2590 for details.

The **AGREE** option in the TABLES statement provides estimates, standard errors, and confidence limits for kappa coefficients. You can request exact tests for kappa coefficients by using the **EXACT** statement.

Kappa coefficients are defined only for square tables, where the number of rows equals the number of columns. Kappa coefficients are not computed for tables that are not square. For 2×2 tables, the weighted kappa coefficient is identical to the simple kappa coefficient, and PROC FREQ presents only the simple kappa coefficient.

GAMMA

requests an asymptotic test for the gamma statistic. See the section “[Gamma](#)” on page 2557 for details. The **MEASURES** option in the TABLES statement provides the gamma statistic and its asymptotic standard error.

KAPPA

requests an asymptotic test for the simple kappa coefficient. See the section “[Simple Kappa Coefficient](#)” on page 2589 for details.

The **AGREE** option in the TABLES statement provides the kappa statistic, its standard error, and its confidence limits. You can request an exact test for the simple kappa coefficient by specifying the **KAPPA** option in the EXACT statement.

Kappa coefficients are defined only for square tables, where the number of rows equals the number of columns. PROC FREQ does not compute kappa coefficients for tables that are not square.

KENTB | TAUB

requests an asymptotic test for Kendall’s tau-*b*. See the section “[Kendall’s Tau-b](#)” on page 2558 for details.

The **MEASURES** option in the TABLES statement provides Kendall’s tau-*b* and its standard error. You can request an exact test for Kendall’s tau-*b* by specifying the **KENTB** option in the EXACT statement.

MEASURES

requests asymptotic tests for the following measures of association: gamma, Kendall’s tau-*b*, Pearson correlation coefficient, Somers’ $D(C|R)$, Somers’ $D(R|C)$, Spearman correlation coefficient, and Stuart’s tau-*c*. See the section “[Measures of Association](#)” on page 2556 for details.

The **MEASURES** option in the TABLES statement provides measures of association and their asymptotic standard errors. You can request exact tests for selected measures by using the **EXACT** statement.

PCORR

requests an asymptotic test for the Pearson correlation coefficient. See the section “[Pearson Correlation Coefficient](#)” on page 2559 for details.

The **MEASURES** option in the TABLES statement provides the Pearson correlation and its standard error. You can request an exact test for the Pearson correlation by specifying the **PCORR** option in the EXACT statement.

PLCORR

requests Wald and likelihood ratio tests for the polychoric correlation coefficient. See the section “[Polychoric Correlation](#)” on page 2562 for details.

The **PLCORR** option in the TABLES statement provides the polychoric correlation and its standard error.

SCORR

requests an asymptotic test for the Spearman correlation coefficient. See the section “[Spearman Rank Correlation Coefficient](#)” on page 2560 for details.

The **MEASURES** option in the TABLES statement provides the Spearman correlation and its standard error. You can request an exact test for the Spearman correlation by specifying the **SCORR** option in the EXACT statement.

SMDCR

requests an asymptotic test for Somers’ $D(C|R)$. See the section “[Somers’ D](#)” on page 2559 for details.

The **MEASURES** option in the TABLES statement provides Somers’ $D(C|R)$ and its standard error. You can request an exact test for Somers’ $D(C|R)$ by specifying the **SMDCR** option in the EXACT statement.

SMDRC

requests an asymptotic test for Somers’ $D(R|C)$. See the section “[Somers’ D](#)” on page 2559 for details.

The **MEASURES** option in the TABLES statement provides Somers’ $D(R|C)$ and its standard error. You can request an exact test for Somers’ $D(R|C)$ by specifying the **SMDRC** option in the EXACT statement.

STUTC | TAUC

requests an asymptotic test for Stuart’s tau- c . See the section “[Stuart’s Tau-c](#)” on page 2558 for details.

The **MEASURES** option in the TABLES statement provides Stuart’s tau- c and its standard error. You can request an exact test for Stuart’s tau- c by specifying the **STUTC** option in the EXACT statement.

WTKAP

requests an asymptotic test for the weighted kappa coefficient. See the section “[Weighted Kappa Coefficient](#)” on page 2590 for details.

The **AGREE** option in the TABLES statement provides the weighted kappa statistic, its standard error, and its confidence limits. You can request an exact test for weighted kappa by specifying the **WTKAP** option in the EXACT statement.

Kappa coefficients are defined only for square tables, where the number of rows equals the number of columns. PROC FREQ does not compute kappa coefficients for tables that are not square. For 2×2 tables, the weighted kappa coefficient is identical to the simple kappa coefficient, and PROC FREQ presents only the simple kappa coefficient.

WEIGHT Statement

WEIGHT *variable* </ *option* > ;

The WEIGHT statement names a numeric variable that provides a weight for each observation in the input data set. The WEIGHT statement is most commonly used to input cell count data. See the section “[Inputting Frequency Counts](#)” on page 2544 for more information. If you use a WEIGHT statement, PROC FREQ assumes that an observation represents n observations, where n is the value of *variable*. The value of the WEIGHT variable is not required to be an integer.

If the value of the WEIGHT variable is missing, PROC FREQ does not use that observation in the analysis. If the value of the WEIGHT variable is zero, PROC FREQ ignores the observation unless you specify the [ZEROS](#) option, which includes observations with zero weights. If you do not specify a WEIGHT statement, PROC FREQ assigns a weight of one to each observation. The sum of the WEIGHT variable values represents the total number of observations.

If any value of the WEIGHT variable is negative, PROC FREQ displays the frequencies computed from the weighted values but does not compute percentages and statistics. If you create an output data set by using the [OUT=](#) option in the TABLES statement, PROC FREQ assigns missing values to the PERCENT variable. PROC FREQ also assigns missing values to the variables that the OUTEXPECT and OUTPCT options provide. If any value of the WEIGHT variable is negative, you cannot create an output data set by using the [OUTPUT](#) statement because statistics are not computed when there are negative weights.

You can specify the following *option* in the WEIGHT statement:

ZEROS

includes observations with zero weight values. By default, PROC FREQ ignores observations with zero weights.

If you specify the ZEROS option, frequency and crosstabulation tables display any levels corresponding to observations with zero weights. Without the ZEROS option, PROC FREQ does not process observations with zero weights, and so does not display levels that contain only observations with zero weights.

With the ZEROS option, PROC FREQ includes levels with zero weights in the chi-square goodness-of-fit test for one-way tables. Also, PROC FREQ includes any levels with zero weights in binomial computations for one-way tables. This makes it possible to compute binomial tests and estimates when the specified level contains no observations with positive weights.

For two-way tables, the ZEROS option enables computation of kappa statistics when there are levels that contain no observations with positive weight. For more information, see the section “[Tables with Zero Rows and Columns](#)” on page 2592.

Note that even with the ZEROS option, PROC FREQ does not compute the [CHISQ](#) or [MEASURES](#) statistics for two-way tables when the table has a zero row or zero column because most of these statistics are undefined in this case.

The ZEROS option invokes the [SPARSE](#) option in the TABLES statement, which includes table cells with a zero frequency count in the [LIST](#) output and in the [OUT=](#) data set. By default, without the SPARSE option, PROC FREQ does not include zero frequency cells in the LIST output or in the [OUT=](#) data set. If you specify the ZEROS option in the WEIGHT statement but do not want the SPARSE option, you can specify the [NOSPARSE](#) option in the TABLES statement.

Details: FREQ Procedure

Inputting Frequency Counts

PROC FREQ can use either raw data or cell count data to produce frequency and crosstabulation tables. *Raw data*, also known as case-record data, report the data as one record for each subject or sample member. *Cell count data* report the data as a table, listing all possible combinations of data values along with the frequency counts. This way of presenting data often appears in published results.

The following DATA step statements store raw data in a SAS data set:

```
data Raw;
  input Subject $ R C @@;
  datalines;
01 1 1  02 1 1  03 1 1  04 1 1  05 1 1
06 1 2  07 1 2  08 1 2  09 2 1  10 2 1
11 2 1  12 2 1  13 2 2  14 2 2  14 2 2
;
```

You can store the same data as cell counts by using the following DATA step statements:

```
data CellCounts;
  input R C Count @@;
  datalines;
1 1 5   1 2 3
2 1 4   2 2 3
;
```

The variable *R* contains the values for the rows, and the variable *C* contains the values for the columns. The variable *Count* contains the cell count for each row and column combination.

Both the *Raw* data set and the *CellCounts* data set produce identical frequency counts, two-way tables, and statistics. When using the *CellCounts* data set, you must include a *WEIGHT* statement to specify that the variable *Count* contains cell counts. For example, the following PROC FREQ statements create a two-way crosstabulation table by using the *CellCounts* data set:

```
proc freq data=CellCounts;
  tables R*C;
  weight Count;
run;
```

Grouping with Formats

PROC FREQ groups a variable's values according to its formatted values. If you assign a format to a variable with a *FORMAT* statement, PROC FREQ formats the variable values before dividing observations into the levels of a frequency or crosstabulation table.

For example, suppose that variable X has the values 1.1, 1.4, 1.7, 2.1, and 2.3. Each of these values appears as a level in the frequency table. If you decide to round each value to a single digit, include the following statement in the PROC FREQ step:

```
format X 1.;
```

Now the table lists the frequency count for formatted level 1 as two and for formatted level 2 as three.

PROC FREQ treats formatted character variables in the same way. The formatted values are used to group the observations into the levels of a frequency table or crosstabulation table. PROC FREQ uses the entire value of a character format to classify an observation.

You can also use the FORMAT statement to assign formats that were created with the FORMAT procedure to the variables. User-written formats determine the number of levels for a variable and provide labels for a table. If you use the same data with different formats, then you can produce frequency counts and statistics for different classifications of the variable values.

When you use PROC FORMAT to create a user-written format that combines missing and nonmissing values into one category, PROC FREQ treats the entire category of formatted values as missing. For example, a questionnaire codes 1 as yes, 2 as no, and 8 as a no answer. The following PROC FORMAT statements create a user-written format:

```
proc format;
  value Questfmt 1   ='Yes'
                2   ='No'
                8,. ='Missing';
run;
```

When you use a FORMAT statement to assign Questfmt. to a variable, the variable's frequency table no longer includes a frequency count for the response of 8. You must use the MISSING or MISSPRINT option in the TABLES statement to list the frequency for no answer. The frequency count for this level includes observations with either a value of 8 or a missing value (.).

The frequency or crosstabulation table lists the values of both character and numeric variables in ascending order based on internal (unformatted) variable values unless you change the order with the ORDER= option. To list the values in ascending order by formatted values, use ORDER=FORMATTED in the PROC FREQ statement.

For more information about the FORMAT statement, see *SAS Formats and Informats: Reference*.

Missing Values

When the value of the WEIGHT variable is missing, PROC FREQ does not include that observation in the analysis.

PROC FREQ treats missing BY variable values like any other BY variable value. The missing values form a separate BY group.

If an observation has a missing value for a variable in a TABLES request, by default PROC FREQ does not include that observation in the frequency or crosstabulation table. Also by default, PROC FREQ does not include observations with missing values in the computation of percentages and statistics. The procedure displays the number of missing observations below each table.

PROC FREQ also reports the number of missing values in output data sets. The TABLES statement OUT= data set includes an observation that contains the missing value frequency. The NMISS option in the OUTPUT statement provides an output data set variable that contains the missing value frequency.

The following options change the way in which PROC FREQ handles missing values of TABLES variables:

- MISSPRINT** displays missing value frequencies in frequency or crosstabulation tables but does not include them in computations of percentages or statistics.
- MISSING** treats missing values as a valid nonmissing level for all TABLES variables. Displays missing levels in frequency and crosstabulation tables and includes them in computations of percentages and statistics.

This example shows the three ways that PROC FREQ can handle missing values of TABLES variables. The following DATA step statements create a data set with a missing value for the variable A:

```
data one;
  input A Freq;
  datalines;
1 2
2 2
. 2
;
```

The following PROC FREQ statements request a one-way frequency table for the variable A. The first request does not specify a missing value option. The second request specifies the MISSPRINT option in the TABLES statement. The third request specifies the MISSING option in the TABLES statement.

```
proc freq data=one;
  tables A;
  weight Freq;
  title 'Default';
run;
proc freq data=one;
  tables A / missprint;
  weight Freq;
  title 'MISSPRINT Option';
run;
proc freq data=one;
  tables A / missing;
  weight Freq;
  title 'MISSING Option';
run;
```

Figure 38.12 displays the frequency tables produced by this example. The first table shows PROC FREQ's default behavior for handling missing values. The observation with a missing value of the TABLES variable A is not included in the table, and the frequency of missing values is displayed below the table. The second table, for which the MISSPRINT option is specified, displays the missing observation but does not include its frequency when computing the total frequency and percentages. The third table shows that PROC FREQ treats the missing level as a valid nonmissing level when the MISSING option is specified. The table displays the missing level, and PROC FREQ includes this level when computing frequencies and percentages.

Figure 38.12 Missing Values in Frequency Tables

Default				
The FREQ Procedure				
A	Frequency	Percent	Cumulative Frequency	Cumulative Percent
1	2	50.00	2	50.00
2	2	50.00	4	100.00
Frequency Missing = 2				
MISSPRINT Option				
The FREQ Procedure				
A	Frequency	Percent	Cumulative Frequency	Cumulative Percent
.	2	.	.	.
1	2	50.00	2	50.00
2	2	50.00	4	100.00
Frequency Missing = 2				
MISSING Option				
The FREQ Procedure				
A	Frequency	Percent	Cumulative Frequency	Cumulative Percent
.	2	33.33	2	33.33
1	2	33.33	4	66.67
2	2	33.33	6	100.00

When a combination of variable values for a two-way table is missing, PROC FREQ assigns zero to the frequency count for the table cell. By default, PROC FREQ does not display missing combinations in LIST format. Also, PROC FREQ does not include missing combinations in the OUT= output data set by default. To include missing combinations, you can specify the SPARSE option with the LIST or OUT= option in the TABLES statement.

In-Database Computation

The FREQ procedure can use in-database computation to construct frequency and crosstabulation tables when the **DATA=** input data set is stored as a table in a supported database management system (DBMS). Supported databases include Teradata, DB2 under UNIX, and Oracle. In-database computation can provide the advantages of faster processing and reduced data transfer between the database and SAS software. For information about in-database computation, see the section “In-Database Procedures” in *SAS/ACCESS for Relational Databases: Reference*.

PROC FREQ performs in-database computation by using SQL implicit pass-through. The procedure generates SQL queries that are based on the tables that you request in the **TABLES** statement. The database executes these SQL queries to construct initial summary tables, which are then transmitted to PROC FREQ. The procedure uses this summary information to perform the remaining analyses and tasks in the usual way (out of the database). So instead of transferring the entire data set over the network between the database and SAS software, the in-database method transfers only the summary tables. This can substantially reduce processing time when the dimensions of the summary tables (in terms of rows and columns) are much smaller than the dimensions of the entire database table (in terms of individual observations). Additionally, in-database summarization uses efficient parallel processing, which can also provide performance advantages.

In-database computation is controlled by the SQLGENERATION option, which you can specify in either a LIBNAME statement or an OPTIONS statement. See the section “In-Database Procedures” in *SAS/ACCESS for Relational Databases: Reference* for details about the SQLGENERATION option and other options that affect in-database computation. By default, PROC FREQ uses in-database computation when possible. There are no FREQ procedure options that control in-database computation.

PROC FREQ uses formatted values to group observations into the levels of frequency and crosstabulation tables. See the section “**Grouping with Formats**” on page 2544 for more information. If formats are available in the database, then in-database summarization uses the formats. If formats are not available in the database, then in-database summarization is based on the raw data values, and PROC FREQ performs the final, formatted classification (out of the database). For more information, see the section “Deploying and Using SAS Formats in Teradata” in *SAS/ACCESS for Relational Databases: Reference*.

The order of observations is not inherently defined for DBMS tables. The following options relate to the order of observations and therefore should not be specified for PROC FREQ in-database computation:

- If you specify the **FIRSTOBS=** or **OBS=** data set option, PROC FREQ does not perform in-database computation.
- If you specify the **NOTSORTED** option in the **BY** statement, PROC FREQ in-database computation ignores it and uses the default **ASCENDING** order for **BY** variables.
- If you specify the **ORDER=DATA** option for input data in a DBMS table, PROC FREQ computation might produce different results for separate runs of the same analysis. In addition to determining the order of variable levels in crosstabulation table displays, the **ORDER=** option can also affect the values of many of the test statistics and measures that PROC FREQ computes.

Statistical Computations

Definitions and Notation

A two-way table represents the crosstabulation of row variable X and column variable Y . Let the table row values or levels be denoted by $X_i, i = 1, 2, \dots, R$, and the column values by $Y_j, j = 1, 2, \dots, C$. Let n_{ij} denote the frequency of the table cell in the i th row and j th column and define the following notation:

$$n_{i.} = \sum_j n_{ij} \quad (\text{row totals})$$

$$n_{.j} = \sum_i n_{ij} \quad (\text{column totals})$$

$$n = \sum_i \sum_j n_{ij} \quad (\text{overall total})$$

$$p_{ij} = n_{ij}/n \quad (\text{cell percentages})$$

$$p_{i.} = n_{i.}/n \quad (\text{row percentages of total})$$

$$p_{.j} = n_{.j}/n \quad (\text{column percentages of total})$$

$$R_i = \text{score for row } i$$

$$C_j = \text{score for column } j$$

$$\bar{R} = \sum_i n_{i.} R_i / n \quad (\text{average row score})$$

$$\bar{C} = \sum_j n_{.j} C_j / n \quad (\text{average column score})$$

$$A_{ij} = \sum_{k>i} \sum_{l>j} n_{kl} + \sum_{k<i} \sum_{l<j} n_{kl}$$

$$D_{ij} = \sum_{k>i} \sum_{l<j} n_{kl} + \sum_{k<i} \sum_{l>j} n_{kl}$$

$$P = \sum_i \sum_j n_{ij} A_{ij} \quad (\text{twice the number of concordances})$$

$$Q = \sum_i \sum_j n_{ij} D_{ij} \quad (\text{twice the number of discordances})$$

Scores

PROC FREQ uses scores of the variable values to compute the Mantel-Haenszel chi-square, Pearson correlation, Cochran-Armitage test for trend, weighted kappa coefficient, and Cochran-Mantel-Haenszel statistics. The SCORES= option in the TABLES statement specifies the score type that PROC FREQ uses. The available score types are TABLE, RANK, RIDIT, and MODRIDIT scores. The default score type is TABLE. Using MODRIDIT, RANK, or RIDIT scores yields nonparametric analyses.

For numeric variables, table scores are the values of the row and column levels. If the row or column variable is formatted, then the table score is the internal numeric value corresponding to that level. If two or more numeric values are classified into the same formatted level, then the internal numeric value for that

level is the smallest of these values. For character variables, table scores are defined as the row numbers and column numbers (that is, 1 for the first row, 2 for the second row, and so on).

Rank scores, which you request with the SCORES=RANK option, are defined as

$$R_i^1 = \sum_{k < i} n_{k.} + (n_{i.} + 1)/2 \quad i = 1, 2, \dots, R$$

$$C_j^1 = \sum_{l < j} n_{.l} + (n_{.j} + 1)/2 \quad j = 1, 2, \dots, C$$

where R_i^1 is the rank score of row i , and C_j^1 is the rank score of column j . Note that rank scores yield midranks for tied values.

Ridit scores, which you request with the SCORES=RIDIT option, are defined as rank scores standardized by the sample size (Bross 1958; Mack and Skillings 1980). Ridit scores are derived from the rank scores as

$$R_i^2 = R_i^1/n \quad i = 1, 2, \dots, R$$

$$C_j^2 = C_j^1/n \quad j = 1, 2, \dots, C$$

Modified ridit scores (SCORES=MODRIDIT) represent the expected values of the order statistics of the uniform distribution on (0,1) (van Elteren 1960; Lehmann 2006). Modified ridit scores are derived from rank scores as

$$R_i^3 = R_i^1/(n + 1) \quad i = 1, 2, \dots, R$$

$$C_j^3 = C_j^1/(n + 1) \quad j = 1, 2, \dots, C$$

Chi-Square Tests and Statistics

The CHISQ option provides chi-square tests of homogeneity or independence and measures of association that are based on the chi-square statistic. When you specify the CHISQ option in the TABLES statement, PROC FREQ computes the following chi-square tests for each two-way table: Pearson chi-square, likelihood ratio chi-square, and Mantel-Haenszel chi-square tests. PROC FREQ provides the following measures of association that are based on the Pearson chi-square statistic: phi coefficient, contingency coefficient, and Cramér's V . For 2×2 tables, the CHISQ option also provides Fisher's exact test and the continuity-adjusted chi-square statistic. You can request Fisher's exact test for general $R \times C$ tables by specifying the FISHER option in the TABLES or EXACT statement.

If you specify the CHISQ option for one-way tables, PROC FREQ provides a one-way Pearson chi-square goodness-of-fit test. If you specify the CHISQ(LRCHISQ) option for one-way tables, PROC FREQ also provides a one-way likelihood ratio chi-square test. The other tests and statistics that the CHISQ option produces are available only for two-way tables.

For two-way tables, the null hypothesis for the chi-square tests is no association between the row variable and the column variable. When the sample size n is large, the test statistics have asymptotic chi-square distributions under the null hypothesis. When the sample size is not large, or when the data set is sparse or heavily tied, exact tests might be more appropriate than asymptotic tests. PROC FREQ provides exact p -values for the Pearson chi-square, likelihood ratio chi-square, and Mantel-Haenszel chi-square tests, in addition to Fisher's exact test. For one-way tables, PROC FREQ provides exact p -values for the Pearson and likelihood ratio chi-square goodness-of-fit tests. You can request these exact tests by specifying the

corresponding options in the EXACT statement. See the section “Exact Statistics” on page 2601 for more information.

The Mantel-Haenszel chi-square statistic is appropriate only when both variables lie on an ordinal scale. The other chi-square tests and statistics in this section are appropriate for either nominal or ordinal variables. The following sections give the formulas that PROC FREQ uses to compute the chi-square tests and statistics. For more information about these statistics, see Agresti (2007) and Stokes, Davis, and Koch (2012), and the other references cited.

Chi-Square Test for One-Way Tables

For one-way frequency tables, the CHISQ option in the TABLES statement provides a chi-square goodness-of-fit test. Let C denote the number of classes, or levels, in the one-way table. Let f_i denote the frequency of class i (or the number of observations in class i) for $i = 1, 2, \dots, C$. Then PROC FREQ computes the one-way chi-square statistic as

$$Q_P = \sum_{i=1}^C (f_i - e_i)^2 / e_i$$

where e_i is the expected frequency for class i under the null hypothesis.

In the test for equal proportions, which is the default for the CHISQ option, the null hypothesis specifies equal proportions of the total sample size for each class. Under this null hypothesis, the expected frequency for each class equals the total sample size divided by the number of classes,

$$e_i = n / C \quad \text{for } i = 1, 2, \dots, C$$

In the test for specified frequencies, which PROC FREQ computes when you input null hypothesis frequencies by using the TESTF= option, the expected frequencies are the TESTF= values that you specify. In the test for specified proportions, which PROC FREQ computes when you input null hypothesis proportions by using the TESTP= option, the expected frequencies are determined from the specified TESTP= proportions p_i as

$$e_i = p_i \times n \quad \text{for } i = 1, 2, \dots, C$$

Under the null hypothesis (of equal proportions, specified frequencies, or specified proportions), Q_P has an asymptotic chi-square distribution with $C-1$ degrees of freedom.

In addition to the asymptotic test, you can request an exact one-way chi-square test by specifying the CHISQ option in the EXACT statement. See the section “Exact Statistics” on page 2601 for more information.

Pearson Chi-Square Test for Two-Way Tables

The Pearson chi-square for two-way tables involves the differences between the observed and expected frequencies, where the expected frequencies are computed under the null hypothesis of independence. The Pearson chi-square statistic is computed as

$$Q_P = \sum_i \sum_j (n_{ij} - e_{ij})^2 / e_{ij}$$

where n_{ij} is the observed frequency in table cell (i, j) and e_{ij} is the expected frequency for table cell (i, j) . The expected frequency is computed under the null hypothesis that the row and column variables are independent,

$$e_{ij} = (n_{i.} n_{.j}) / n$$

When the row and column variables are independent, Q_P has an asymptotic chi-square distribution with $(R-1)(C-1)$ degrees of freedom. For large values of Q_P , this test rejects the null hypothesis in favor of the alternative hypothesis of general association.

In addition to the asymptotic test, you can request an exact Pearson chi-square test by specifying the PCHI or CHISQ option in the EXACT statement. See the section “Exact Statistics” on page 2601 for more information.

For 2×2 tables, the Pearson chi-square is also appropriate for testing the equality of two binomial proportions. For $R \times 2$ and $2 \times C$ tables, the Pearson chi-square tests the homogeneity of proportions. See Fienberg (1980) for details.

Standardized Residuals

When you specify the CROSSLIST(STDRES) option in the TABLES statement for two-way or multiway tables, PROC FREQ displays the standardized residuals in the CROSSLIST table.

The standardized residual of a crosstabulation table cell is the ratio of (*frequency* – *expected*) to its standard error, where *frequency* is the table cell frequency and *expected* is the estimated expected cell frequency. The expected frequency is computed under the null hypothesis that the row and column variables are independent. See the section “Pearson Chi-Square Test for Two-Way Tables” on page 2551 for more information.

PROC FREQ computes the standardized residual of table cell (i, j) as

$$(n_{ij} - e_{ij}) / \sqrt{e_{ij}(1 - p_{i.})(1 - p_{.j})}$$

where n_{ij} is the observed frequency of table cell (i, j) , e_{ij} is the expected frequency of the table cell, $p_{i.}$ is the proportion in row i ($n_{i.}/n$), and $p_{.j}$ is the proportion in column j ($n_{.j}/n$). The expected frequency of table cell (i, j) is computed as

$$e_{ij} = (n_{i.} n_{.j}) / n$$

Under the null hypothesis of independence, each standardized residual has an asymptotic standard normal distribution. See section 2.4.5 of Agresti (2007) for more information.

Likelihood Ratio Chi-Square Test for One-Way Tables

For one-way frequency tables, the CHISQ(LRCHISQ) option in the TABLES statement provides a likelihood ratio chi-square goodness-of-fit test. By default, the likelihood ratio test is based on the null hypothesis of equal proportions in the C classes (levels) of the one-way table. If you specify null hypothesis proportions or frequencies by using the CHISQ(TESTP=) or CHISQ(TESTF=) option, respectively, the likelihood ratio test is based on the null hypothesis values that you specify.

PROC FREQ computes the one-way likelihood ratio test as

$$G^2 = 2 \sum_{i=1}^C f_i \ln(f_i / e_i)$$

where f_i is the observed frequency of class i , and e_i is the expected frequency of class i under the null hypothesis.

For the null hypothesis of equal proportions, the expected frequency of each class equals the total sample size divided by the number of classes,

$$e_i = n / C \quad \text{for } i = 1, 2, \dots, C$$

If you provide null hypothesis frequencies by specifying the CHISQ(TESTF=) option in the TABLES statement, the expected frequencies are the TESTF= values that you specify. If you provide null hypothesis proportions by specifying the CHISQ(ATESTP=) option in the TABLES statement, PROC FREQ computes the expected frequencies as

$$e_i = p_i \times n \quad \text{for } i = 1, 2, \dots, C$$

where the proportions p_i are the AATESTP= values that you specify.

Under the null hypothesis (of equal proportions, specified frequencies, or specified proportions), the likelihood ratio statistic G^2 has an asymptotic chi-square distribution with $C-1$ degrees of freedom.

In addition to the asymptotic test, you can request an exact one-way likelihood ratio chi-square test by specifying the LRCHISQ option in the EXACT statement. See the section “Exact Statistics” on page 2601 for more information.

Likelihood Ratio Chi-Square Test

The likelihood ratio chi-square involves the ratios between the observed and expected frequencies. The likelihood ratio chi-square statistic is computed as

$$G^2 = 2 \sum_i \sum_j n_{ij} \ln(n_{ij}/e_{ij})$$

where n_{ij} is the observed frequency in table cell (i, j) and e_{ij} is the expected frequency for table cell (i, j) .

When the row and column variables are independent, G^2 has an asymptotic chi-square distribution with $(R-1)(C-1)$ degrees of freedom.

In addition to the asymptotic test, you can request an exact likelihood ratio chi-square test by specifying the LRCHI or CHISQ option in the EXACT statement. See the section “Exact Statistics” on page 2601 for more information.

Continuity-Adjusted Chi-Square Test

The continuity-adjusted chi-square for 2×2 tables is similar to the Pearson chi-square, but it is adjusted for the continuity of the chi-square distribution. The continuity-adjusted chi-square is most useful for small sample sizes. The use of the continuity adjustment is somewhat controversial; this chi-square test is more conservative (and more like Fisher’s exact test) when the sample size is small. As the sample size increases, the continuity-adjusted chi-square becomes more like the Pearson chi-square.

The continuity-adjusted chi-square statistic is computed as

$$Q_C = \sum_i \sum_j (\max(0, |n_{ij} - e_{ij}| - 0.5))^2 / e_{ij}$$

Under the null hypothesis of independence, Q_C has an asymptotic chi-square distribution with $(R-1)(C-1)$ degrees of freedom.

Mantel-Haenszel Chi-Square Test

The Mantel-Haenszel chi-square statistic tests the alternative hypothesis that there is a linear association between the row variable and the column variable. Both variables must lie on an ordinal scale. The Mantel-Haenszel chi-square statistic is computed as

$$Q_{MH} = (n-1)r^2$$

where r is the Pearson correlation between the row variable and the column variable. For a description of the Pearson correlation, see the “[Pearson Correlation Coefficient](#)” on page 2559. The Pearson correlation and thus the Mantel-Haenszel chi-square statistic use the scores that you specify in the SCORES= option in the TABLES statement. See Mantel and Haenszel (1959) and Landis, Heyman, and Koch (1978) for more information.

Under the null hypothesis of no association, Q_{MH} has an asymptotic chi-square distribution with one degree of freedom.

In addition to the asymptotic test, you can request an exact Mantel-Haenszel chi-square test by specifying the MHCHI or CHISQ option in the EXACT statement. See the section “[Exact Statistics](#)” on page 2601 for more information.

Fisher's Exact Test

Fisher's exact test is another test of association between the row and column variables. This test assumes that the row and column totals are fixed, and then uses the hypergeometric distribution to compute probabilities of possible tables conditional on the observed row and column totals. Fisher's exact test does not depend on any large-sample distribution assumptions, and so it is appropriate even for small sample sizes and for sparse tables.

2 × 2 Tables For 2 × 2 tables, PROC FREQ gives the following information for Fisher's exact test: table probability, two-sided p -value, left-sided p -value, and right-sided p -value. The table probability equals the hypergeometric probability of the observed table, and is in fact the value of the test statistic for Fisher's exact test.

Where p is the hypergeometric probability of a specific table with the observed row and column totals, Fisher's exact p -values are computed by summing probabilities p over defined sets of tables,

$$\text{Prob} = \sum_A p$$

The two-sided p -value is the sum of all possible table probabilities (conditional on the observed row and column totals) that are less than or equal to the observed table probability. For the two-sided p -value, the set A includes all possible tables with hypergeometric probabilities less than or equal to the probability of the observed table. A small two-sided p -value supports the alternative hypothesis of association between the row and column variables.

For 2 × 2 tables, one-sided p -values for Fisher's exact test are defined in terms of the frequency of the cell in the first row and first column of the table, the (1,1) cell. Denoting the observed (1,1) cell frequency by n_{11} , the left-sided p -value for Fisher's exact test is the probability that the (1,1) cell frequency is less than or equal to n_{11} . For the left-sided p -value, the set A includes those tables with a (1,1) cell frequency less than or equal to n_{11} . A small left-sided p -value supports the alternative hypothesis that the probability of an observation being in the first cell is actually less than expected under the null hypothesis of independent row and column variables.

Similarly, for a right-sided alternative hypothesis, A is the set of tables where the frequency of the (1,1) cell is greater than or equal to that in the observed table. A small right-sided p -value supports the alternative that the probability of the first cell is actually greater than that expected under the null hypothesis.

Because the (1,1) cell frequency completely determines the 2 × 2 table when the marginal row and column sums are fixed, these one-sided alternatives can be stated equivalently in terms of other cell probabilities or

ratios of cell probabilities. The left-sided alternative is equivalent to an odds ratio less than 1, where the odds ratio equals $(n_{11}n_{22}/n_{12}n_{21})$. Additionally, the left-sided alternative is equivalent to the column 1 risk for row 1 being less than the column 1 risk for row 2, $p_{1|1} < p_{1|2}$. Similarly, the right-sided alternative is equivalent to the column 1 risk for row 1 being greater than the column 1 risk for row 2, $p_{1|1} > p_{1|2}$. See Agresti (2007) for details.

R × C Tables Fisher's exact test was extended to general $R \times C$ tables by Freeman and Halton (1951), and this test is also known as the Freeman-Halton test. For $R \times C$ tables, the two-sided p -value definition is the same as for 2×2 tables. The set A contains all tables with p less than or equal to the probability of the observed table. A small p -value supports the alternative hypothesis of association between the row and column variables. For $R \times C$ tables, Fisher's exact test is inherently two-sided. The alternative hypothesis is defined only in terms of general, and not linear, association. Therefore, Fisher's exact test does not have right-sided or left-sided p -values for general $R \times C$ tables.

For $R \times C$ tables, PROC FREQ computes Fisher's exact test by using the network algorithm of Mehta and Patel (1983), which provides a faster and more efficient solution than direct enumeration. See the section "Exact Statistics" on page 2601 for more details.

Phi Coefficient

The phi coefficient is a measure of association derived from the Pearson chi-square. The range of the phi coefficient is $-1 \leq \phi \leq 1$ for 2×2 tables. For tables larger than 2×2 , the range is $0 \leq \phi \leq \min(\sqrt{R-1}, \sqrt{C-1})$ (Liebetrau 1983). The phi coefficient is computed as

$$\phi = (n_{11}n_{22} - n_{12}n_{21}) / \sqrt{n_{1.}n_{2.}n_{.1}n_{.2}} \quad \text{for } 2 \times 2 \text{ tables}$$

$$\phi = \sqrt{Q_P/n} \quad \text{otherwise}$$

See Fleiss, Levin, and Paik (2003, pp. 98–99) for more information.

Contingency Coefficient

The contingency coefficient is a measure of association derived from the Pearson chi-square. The range of the contingency coefficient is $0 \leq P \leq \sqrt{(m-1)/m}$, where $m = \min(R, C)$ (Liebetrau 1983). The contingency coefficient is computed as

$$P = \sqrt{Q_P / (Q_P + n)}$$

See Kendall and Stuart (1979, pp. 587–588) for more information.

Cramér's V

Cramér's V is a measure of association derived from the Pearson chi-square. It is designed so that the attainable upper bound is always 1. The range of Cramér's V is $-1 \leq V \leq 1$ for 2×2 tables; for tables larger than 2×2 , the range is $0 \leq V \leq 1$. Cramér's V is computed as

$$V = \phi \quad \text{for } 2 \times 2 \text{ tables}$$

$$V = \sqrt{\frac{Q_P/n}{\min(R-1, C-1)}} \quad \text{otherwise}$$

See Kendall and Stuart (1979, p. 588) for more information.

Measures of Association

When you specify the MEASURES option in the TABLES statement, PROC FREQ computes several statistics that describe the association between the row and column variables of the contingency table. The following are measures of ordinal association that consider whether the column variable *Y* tends to increase as the row variable *X* increases: gamma, Kendall's tau-*b*, Stuart's tau-*c*, and Somers' *D*. These measures are appropriate for ordinal variables, and they classify pairs of observations as *concordant* or *discordant*. A pair is concordant if the observation with the larger value of *X* also has the larger value of *Y*. A pair is discordant if the observation with the larger value of *X* has the smaller value of *Y*. See Agresti (2007) and the other references cited for the individual measures of association.

The Pearson correlation coefficient and the Spearman rank correlation coefficient are also appropriate for ordinal variables. The Pearson correlation describes the strength of the linear association between the row and column variables, and it is computed by using the row and column scores specified by the SCORES= option in the TABLES statement. The Spearman correlation is computed with rank scores. The polychoric correlation (requested by the PLCORR option) also requires ordinal variables and assumes that the variables have an underlying bivariate normal distribution. The following measures of association do not require ordinal variables and are appropriate for nominal variables: lambda asymmetric, lambda symmetric, and the uncertainty coefficients.

PROC FREQ computes estimates of the measures according to the formulas given in the following sections. For each measure, PROC FREQ computes an asymptotic standard error (ASE), which is the square root of the asymptotic variance denoted by Var in the following sections.

Confidence Limits

If you specify the CL option in the TABLES statement, PROC FREQ computes asymptotic confidence limits for all MEASURES statistics. The confidence coefficient is determined according to the value of the ALPHA= option, which, by default, equals 0.05 and produces 95% confidence limits.

The confidence limits are computed as

$$\text{Est} \pm (z_{\alpha/2} \times \text{ASE})$$

where Est is the estimate of the measure, $z_{\alpha/2}$ is the $100(1 - \alpha/2)$ percentile of the standard normal distribution, and ASE is the asymptotic standard error of the estimate.

Asymptotic Tests

For each measure that you specify in the TEST statement, PROC FREQ computes an asymptotic test of the null hypothesis that the measure equals zero. Asymptotic tests are available for the following measures of association: gamma, Kendall's tau-*b*, Stuart's tau-*c*, Somers' $D(C|R)$, Somers' $D(R|C)$, the Pearson correlation coefficient, and the Spearman rank correlation coefficient. To compute an asymptotic test, PROC FREQ uses a standardized test statistic *z*, which has an asymptotic standard normal distribution under the null hypothesis. The test statistic is computed as

$$z = \text{Est} / \sqrt{\text{Var}_0(\text{Est})}$$

where Est is the estimate of the measure and $\text{Var}_0(\text{Est})$ is the variance of the estimate under the null hypothesis. Formulas for $\text{Var}_0(\text{Est})$ for the individual measures of association are given in the following sections.

Note that the ratio of Est to $\sqrt{\text{Var}_0(\text{Est})}$ is the same for the following measures: gamma, Kendall's tau-*b*, Stuart's tau-*c*, Somers' $D(C|R)$, and Somers' $D(R|C)$. Therefore, the tests for these measures are identical. For example, the *p*-values for the test of H_0 : gamma = 0 equal the *p*-values for the test of H_0 : tau-*b* = 0.

PROC FREQ computes one-sided and two-sided *p*-values for each of these tests. When the test statistic *z* is greater than its null hypothesis expected value of zero, PROC FREQ displays the right-sided *p*-value, which is the probability of a larger value of the statistic occurring under the null hypothesis. A small right-sided *p*-value supports the alternative hypothesis that the true value of the measure is greater than zero. When the test statistic is less than or equal to zero, PROC FREQ displays the left-sided *p*-value, which is the probability of a smaller value of the statistic occurring under the null hypothesis. A small left-sided *p*-value supports the alternative hypothesis that the true value of the measure is less than zero. The one-sided *p*-value P_1 can be expressed as

$$P_1 = \begin{cases} \text{Prob}(Z > z) & \text{if } z > 0 \\ \text{Prob}(Z < z) & \text{if } z \leq 0 \end{cases}$$

where *Z* has a standard normal distribution. The two-sided *p*-value P_2 is computed as

$$P_2 = \text{Prob}(|Z| > |z|)$$

Exact Tests

Exact tests are available for the following measures of association: Kendall's tau-*b*, Stuart's tau-*c*, Somers' $D(C|R)$ and $(R|C)$, the Pearson correlation coefficient, and the Spearman rank correlation coefficient. If you request an exact test for a measure of association in the EXACT statement, PROC FREQ computes the exact test of the hypothesis that the measure equals zero. See the section "Exact Statistics" on page 2601 for details.

Gamma

The gamma (Γ) statistic is based only on the number of concordant and discordant pairs of observations. It ignores tied pairs (that is, pairs of observations that have equal values of *X* or equal values of *Y*). Gamma is appropriate only when both variables lie on an ordinal scale. The range of gamma is $-1 \leq \Gamma \leq 1$. If the row and column variables are independent, then gamma tends to be close to zero. Gamma is computed as

$$G = (P - Q) / (P + Q)$$

and the asymptotic variance is

$$\text{Var}(G) = \frac{16}{(P + Q)^4} \sum_i \sum_j n_{ij} (QA_{ij} - PD_{ij})^2$$

For 2×2 tables, gamma is equivalent to Yule's *Q*. See Goodman and Kruskal (1979) and Agresti (2002) for more information.

The variance under the null hypothesis that gamma equals zero is computed as

$$\text{Var}_0(G) = \frac{4}{(P + Q)^2} \left(\sum_i \sum_j n_{ij} (A_{ij} - D_{ij})^2 - (P - Q)^2/n \right)$$

See Brown and Benedetti (1977) for details.

Kendall's Tau-b

Kendall's tau- b (τ_b) is similar to gamma except that tau- b uses a correction for ties. Tau- b is appropriate only when both variables lie on an ordinal scale. The range of tau- b is $-1 \leq \tau_b \leq 1$. Kendall's tau- b is computed as

$$t_b = (P - Q) / \sqrt{w_r w_c}$$

and the asymptotic variance is

$$\text{Var}(t_b) = \frac{1}{w^4} \left(\sum_i \sum_j n_{ij} (2w d_{ij} + t_b v_{ij})^2 - n^3 t_b^2 (w_r + w_c)^2 \right)$$

where

$$\begin{aligned} w &= \sqrt{w_r w_c} \\ w_r &= n^2 - \sum_i n_{i.}^2 \\ w_c &= n^2 - \sum_j n_{.j}^2 \\ d_{ij} &= A_{ij} - D_{ij} \\ v_{ij} &= n_{i.} w_c + n_{.j} w_r \end{aligned}$$

See Kendall (1955) for more information.

The variance under the null hypothesis that tau- b equals zero is computed as

$$\text{Var}_0(t_b) = \frac{4}{w_r w_c} \left(\sum_i \sum_j n_{ij} (A_{ij} - D_{ij})^2 - (P - Q)^2 / n \right)$$

See Brown and Benedetti (1977) for details.

PROC FREQ also provides an exact test for the Kendall's tau- b . You can request this test by specifying the KENTB option in the EXACT statement. See the section “[Exact Statistics](#)” on page 2601 for more information.

Stuart's Tau-c

Stuart's tau- c (τ_c) makes an adjustment for table size in addition to a correction for ties. Tau- c is appropriate only when both variables lie on an ordinal scale. The range of tau- c is $-1 \leq \tau_c \leq 1$. Stuart's tau- c is computed as

$$t_c = m(P - Q) / n^2(m - 1)$$

and the asymptotic variance is

$$\text{Var}(t_c) = \frac{4m^2}{(m - 1)^2 n^4} \left(\sum_i \sum_j n_{ij} d_{ij}^2 - (P - Q)^2 / n \right)$$

where $m = \min(R, C)$ and $d_{ij} = A_{ij} - D_{ij}$. The variance under the null hypothesis that tau- c equals zero is the same as the asymptotic variance Var,

$$\text{Var}_0(t_c) = \text{Var}(t_c)$$

See Brown and Benedetti (1977) for details.

PROC FREQ also provides an exact test for the Stuart's tau- c . You can request this test by specifying the STUTC option in the EXACT statement. See the section “[Exact Statistics](#)” on page 2601 for more information.

Somers' D

Somers' $D(C|R)$ and Somers' $D(R|C)$ are asymmetric modifications of tau- b . $C|R$ indicates that the row variable X is regarded as the independent variable and the column variable Y is regarded as dependent. Similarly, $R|C$ indicates that the column variable Y is regarded as the independent variable and the row variable X is regarded as dependent. Somers' D differs from tau- b in that it uses a correction only for pairs that are tied on the independent variable. Somers' D is appropriate only when both variables lie on an ordinal scale. The range of Somers' D is $-1 \leq D \leq 1$. Somers' $D(C|R)$ is computed as

$$D(C|R) = (P - Q) / w_r$$

and its asymptotic variance is

$$\text{Var}(D(C|R)) = \frac{4}{w_r^4} \sum_i \sum_j n_{ij} (w_r d_{ij} - (P - Q)(n - n_{i.}))^2$$

where $d_{ij} = A_{ij} - D_{ij}$ and

$$w_r = n^2 - \sum_i n_{i.}^2.$$

For more information, see Somers (1962); Goodman and Kruskal (1979); Liebetrau (1983).

The variance under the null hypothesis that $D(C|R)$ equals zero is computed as

$$\text{Var}_0(D(C|R)) = \frac{4}{w_r^2} \left(\sum_i \sum_j n_{ij} (A_{ij} - D_{ij})^2 - (P - Q)^2 / n \right)$$

See Brown and Benedetti (1977) for details.

Formulas for Somers' $D(R|C)$ are obtained by interchanging the indices.

PROC FREQ also provides exact tests for Somers' $D(C|R)$ and $(R|C)$. You can request these tests by specifying the SMDCR and SMDCR options in the EXACT statement. See the section “[Exact Statistics](#)” on page 2601 for more information.

Pearson Correlation Coefficient

The Pearson correlation coefficient (ρ) is computed by using the scores specified in the SCORES= option. This measure is appropriate only when both variables lie on an ordinal scale. The range of the Pearson correlation is $-1 \leq \rho \leq 1$. The Pearson correlation coefficient is computed as

$$r = v/w = ss_{rc} / \sqrt{ss_r ss_c}$$

and its asymptotic variance is

$$\text{Var}(r) = \frac{1}{w^4} \sum_i \sum_j n_{ij} \left(w(R_i - \bar{R})(C_j - \bar{C}) - \frac{b_{ij}v}{2w} \right)^2$$

where R_i and C_j are the row and column scores and

$$ss_r = \sum_i \sum_j n_{ij} (R_i - \bar{R})^2$$

$$ss_c = \sum_i \sum_j n_{ij} (C_j - \bar{C})^2$$

$$ss_{rc} = \sum_i \sum_j n_{ij} (R_i - \bar{R})(C_j - \bar{C})$$

$$b_{ij} = (R_i - \bar{R})^2 ss_c + (C_j - \bar{C})^2 ss_r$$

$$v = ss_{rc}$$

$$w = \sqrt{ss_r ss_c}$$

See Snedecor and Cochran (1989) for more information.

The SCORES= option in the TABLES statement determines the type of row and column scores used to compute the Pearson correlation (and other score-based statistics). The default is SCORES=TABLE. See the section “[Scores](#)” on page 2549 for details about the available score types and how they are computed.

The variance under the null hypothesis that the correlation equals zero is computed as

$$\text{Var}_0(r) = \left(\sum_i \sum_j n_{ij} (R_i - \bar{R})^2 (C_j - \bar{C})^2 - ss_{rc}^2 / n \right) / ss_r ss_c$$

Note that this expression for the variance is derived for multinomial sampling in a contingency table framework, and it differs from the form obtained under the assumption that both variables are continuous and normally distributed. See Brown and Benedetti (1977) for details.

PROC FREQ also provides an exact test for the Pearson correlation coefficient. You can request this test by specifying the PCORR option in the EXACT statement. See the section “[Exact Statistics](#)” on page 2601 for more information.

Spearman Rank Correlation Coefficient

The Spearman correlation coefficient (ρ_s) is computed by using rank scores, which are defined in the section “[Scores](#)” on page 2549. This measure is appropriate only when both variables lie on an ordinal scale. The range of the Spearman correlation is $-1 \leq \rho_s \leq 1$. The Spearman correlation coefficient is computed as

$$r_s = v / w$$

and its asymptotic variance is

$$\text{Var}(r_s) = \frac{1}{n^2 w^4} \sum_i \sum_j n_{ij} (z_{ij} - \bar{z})^2$$

where R_i^1 and C_j^1 are the row and column rank scores and

$$v = \sum_i \sum_j n_{ij} R(i) C(j)$$

$$w = \frac{1}{12} \sqrt{FG}$$

$$F = n^3 - \sum_i n_{i.}^3$$

$$G = n^3 - \sum_j n_{.j}^3$$

$$R(i) = R_i^1 - n/2$$

$$C(j) = C_j^1 - n/2$$

$$\bar{z} = \frac{1}{n} \sum_i \sum_j n_{ij} z_{ij}$$

$$z_{ij} = w v_{ij} - v w_{ij}$$

$$v_{ij} = n \left(R(i) C(j) + \frac{1}{2} \sum_l n_{il} C(l) + \frac{1}{2} \sum_k n_{kj} R(k) + \right. \\ \left. \sum_l \sum_{k>i} n_{kl} C(l) + \sum_k \sum_{l>j} n_{kl} R(k) \right)$$

$$w_{ij} = \frac{-n}{96w} (F n_{.j}^2 + G n_{i.}^2)$$

See Snedecor and Cochran (1989) for more information.

The variance under the null hypothesis that the correlation equals zero is computed as

$$\text{Var}_0(r_s) = \frac{1}{n^2 w^2} \sum_i \sum_j n_{ij} (v_{ij} - \bar{v})^2$$

where

$$\bar{v} = \sum_i \sum_j n_{ij} v_{ij} / n$$

Note that the asymptotic variance is derived for multinomial sampling in a contingency table framework, and it differs from the form obtained under the assumption that both variables are continuous and normally distributed. See Brown and Benedetti (1977) for details.

PROC FREQ also provides an exact test for the Spearman correlation coefficient. You can request this test by specifying the SCORR option in the EXACT statement. See the section “Exact Statistics” on page 2601 for more information.

Polychoric Correlation

When you specify the PLCORR option in the TABLES statement, PROC FREQ computes the polychoric correlation and its standard error. The polychoric correlation is based on the assumption that the two ordinal, categorical variables of the frequency table have an underlying bivariate normal distribution. The polychoric correlation coefficient is the maximum likelihood estimate of the product-moment correlation between the underlying normal variables. The range of the polychoric correlation is from -1 to 1 . For 2×2 tables, the polychoric correlation is also known as the tetrachoric correlation (and it is labeled as such in the displayed output). See Drasgow (1986) for an overview of polychoric correlation coefficient.

Olsson (1979) gives the likelihood equations and the asymptotic standard errors for estimating the polychoric correlation. The underlying continuous variables relate to the observed crosstabulation table through thresholds, which define a range of numeric values that correspond to each categorical (table) level. PROC FREQ uses Olsson’s maximum likelihood method for simultaneous estimation of the polychoric correlation and the thresholds. (Olsson also presents a two-step method that estimates the thresholds first.)

PROC FREQ iteratively solves the likelihood equations by using a Newton-Raphson algorithm. The initial estimates of the thresholds are computed from the inverse of the normal distribution function at the cumulative marginal proportions of the table. Iterative computation of the polychoric correlation stops when the convergence measure falls below the convergence criterion or when the maximum number of iterations is reached, whichever occurs first. For parameter values that are less than 0.01, the procedure evaluates convergence by using the absolute difference instead of the relative difference. The PLCORR(CONVERGE=) option specifies the convergence criterion, which is 0.0001 by default. The PLCORR(MAXITER=) option specifies the maximum number of iterations, which is 20 by default.

If you specify the CL option in the TABLES statement, PROC FREQ provides confidence limits for the polychoric correlation. The confidence limits are computed as

$$\hat{\rho} \pm (z_{\alpha/2} \times SE(\hat{\rho}))$$

where $\hat{\rho}$ is the estimate of the polychoric correlation, $z_{\alpha/2}$ is the $100(1 - \alpha/2)$ percentile of the standard normal distribution, and $SE(\hat{\rho})$ is the standard error of the polychoric correlation estimate.

If you specify the PLCORR option in the TEST statement, PROC FREQ provides Wald and likelihood ratio tests of the null hypothesis that the polychoric correlation equals 0. The Wald test statistic is computed as

$$z = \hat{\rho} / SE(\hat{\rho})$$

which has a standard normal distribution under the null hypothesis. PROC FREQ computes one-sided and two-sided p -values for the Wald test. When the test statistic z is greater than its null expected value of 0, PROC FREQ displays the right-sided p -value. When the test statistic is less than or equal to 0, PROC FREQ displays the left-sided p -value.

The likelihood ratio statistic for the polychoric correlation is computed as

$$G^2 = -2 \ln(L_0/L_1)$$

where L_0 is the value of the likelihood function (Olsson 1979) when the polychoric correlation is 0, and L_1 is the value of the likelihood function at the maximum (where all parameters are replaced by their maximum likelihood estimates). Under the null hypothesis, the likelihood ratio statistic has an asymptotic chi-square distribution with one degree of freedom.

Lambda (Asymmetric)

Asymmetric lambda, $\lambda(C|R)$, is interpreted as the probable improvement in predicting the column variable Y given knowledge of the row variable X . The range of asymmetric lambda is $0 \leq \lambda(C|R) \leq 1$. Asymmetric lambda $\lambda(C|R)$ is computed as

$$\lambda(C|R) = \frac{\sum_i r_i - r}{n - r}$$

and its asymptotic variance is

$$\text{Var}(\lambda(C|R)) = \frac{n - \sum_i r_i}{(n - r)^3} \left(\sum_i r_i + r - 2 \sum_i (r_i | l_i = l) \right)$$

where

$$r_i = \max_j(n_{ij})$$

$$r = \max_j(n_{.j})$$

$$c_j = \max_i(n_{ij})$$

$$c = \max_i(n_{i.})$$

The values of l_i and l are determined as follows. Denote by l_i the unique value of j such that $r_i = n_{ij}$, and let l be the unique value of j such that $r = n_{.j}$. Because of the uniqueness assumptions, ties in the frequencies or in the marginal totals must be broken in an arbitrary but consistent manner. In case of ties, l is defined as the smallest value of j such that $r = n_{.j}$.

For those columns containing a cell (i, j) for which $n_{ij} = r_i = c_j$, cs_j records the row in which c_j is assumed to occur. Initially cs_j is set equal to -1 for all j . Beginning with $i=1$, if there is at least one value j such that $n_{ij} = r_i = c_j$, and if $cs_j = -1$, then l_i is defined to be the smallest such value of j , and cs_j is set equal to i . Otherwise, if $n_{il} = r_i$, then l_i is defined to be equal to l . If neither condition is true, then l_i is taken to be the smallest value of j such that $n_{ij} = r_i$.

The formulas for lambda asymmetric $\lambda(R|C)$ can be obtained by interchanging the indices.

See Goodman and Kruskal (1979) for more information.

Lambda (Symmetric)

The nondirectional lambda is the average of the two asymmetric lambdas, $\lambda(C|R)$ and $\lambda(R|C)$. Its range is $0 \leq \lambda \leq 1$. Lambda symmetric is computed as

$$\lambda = \frac{\sum_i r_i + \sum_j c_j - r - c}{2n - r - c} = \frac{w - v}{w}$$

and its asymptotic variance is computed as

$$\text{Var}(\lambda) = \frac{1}{w^4} \left(wvy - 2w^2 \left(n - \sum_i \sum_j (n_{ij} \mid j = l_i, i = k_j) \right) - 2v^2(n - n_{kl}) \right)$$

where

$$r_i = \max_j(n_{ij})$$

$$r = \max_j(n_{.j})$$

$$c_j = \max_i(n_{ij})$$

$$c = \max_i(n_{i.})$$

$$w = 2n - r - c$$

$$v = 2n - \sum_i r_i - \sum_j c_j$$

$$x = \sum_i (r_i \mid l_i = l) + \sum_j (c_j \mid k_j = k) + r_k + c_l$$

$$y = 8n - w - v - 2x$$

The definitions of l_i and l are given in the previous section. The values k_j and k are defined in a similar way for lambda asymmetric ($R|C$).

See Goodman and Kruskal (1979) for more information.

Uncertainty Coefficients (Asymmetric)

The uncertainty coefficient $U(C|R)$ measures the proportion of uncertainty (entropy) in the column variable Y that is explained by the row variable X. Its range is $0 \leq U(C|R) \leq 1$. The uncertainty coefficient is computed as

$$U(C|R) = (H(X) + H(Y) - H(XY)) / H(Y) = v/w$$

and its asymptotic variance is

$$\text{Var}(U(C|R)) = \frac{1}{n^2 w^4} \sum_i \sum_j n_{ij} \left(H(Y) \ln \left(\frac{n_{ij}}{n_{i.}} \right) + (H(X) - H(XY)) \ln \left(\frac{n_{.j}}{n} \right) \right)^2$$

where

$$\begin{aligned}
 v &= H(X) + H(Y) - H(XY) \\
 w &= H(Y) \\
 H(X) &= -\sum_i \left(\frac{n_{i\cdot}}{n}\right) \ln\left(\frac{n_{i\cdot}}{n}\right) \\
 H(Y) &= -\sum_j \left(\frac{n_{\cdot j}}{n}\right) \ln\left(\frac{n_{\cdot j}}{n}\right) \\
 H(XY) &= -\sum_i \sum_j \left(\frac{n_{ij}}{n}\right) \ln\left(\frac{n_{ij}}{n}\right)
 \end{aligned}$$

The formulas for the uncertainty coefficient $U(R|C)$ can be obtained by interchanging the indices.

See Theil (1972, pp. 115–120) and Goodman and Kruskal (1979) for more information.

Uncertainty Coefficient (Symmetric)

The uncertainty coefficient U is the symmetric version of the two asymmetric uncertainty coefficients. Its range is $0 \leq U \leq 1$. The uncertainty coefficient is computed as

$$U = 2 (H(X) + H(Y) - H(XY)) / (H(X) + H(Y))$$

and its asymptotic variance is

$$\text{Var}(U) = 4 \sum_i \sum_j \frac{n_{ij} \left(H(XY) \ln\left(\frac{n_{i\cdot} n_{\cdot j}}{n^2}\right) - (H(X) + H(Y)) \ln\left(\frac{n_{ij}}{n}\right) \right)^2}{n^2 (H(X) + H(Y))^4}$$

where $H(X)$, $H(Y)$, and $H(XY)$ are defined in the previous section. See Goodman and Kruskal (1979) for more information.

Binomial Proportion

If you specify the BINOMIAL option in the TABLES statement, PROC FREQ computes the binomial proportion for one-way tables. By default, this is the proportion of observations in the first variable level that appears in the output. (You can use the LEVEL= option to specify a different level for the proportion.) The binomial proportion is computed as

$$\hat{p} = n_1 / n$$

where n_1 is the frequency of the first (or designated) level and n is the total frequency of the one-way table. The standard error of the binomial proportion is computed as

$$\text{se}(\hat{p}) = \sqrt{\hat{p} (1 - \hat{p}) / n}$$

Binomial Confidence Limits

By default, PROC FREQ provides Wald and exact (Clopper-Pearson) confidence limits for the binomial proportion. If you do not specify any confidence limit requests with the `CL= binomial-option`, PROC FREQ computes Wald asymptotic confidence limits. You can also request Agresti-Coull, Jeffreys, and Wilson (score) confidence limits for the binomial proportion. For details about these binomial confidence limits, including comparisons of their performance, see Brown, Cai, and DasGupta (2001); Agresti and Coull (1998); Newcombe (1998b).

Wald Confidence Limits Wald asymptotic confidence limits are based on the normal approximation to the binomial distribution. PROC FREQ computes the Wald confidence limits for the binomial proportion as

$$\hat{p} \pm (z_{\alpha/2} \times \text{se}(\hat{p}))$$

where $z_{\alpha/2}$ is the $100(1 - \alpha/2)$ percentile of the standard normal distribution. The confidence level α is determined by the `ALPHA=` option, which, by default, equals 0.05 and produces 95% confidence limits.

If you specify `CL=WALD(CORRECT)` or the `CORRECT binomial-option`, PROC FREQ includes a continuity correction of $1/2n$ in the Wald asymptotic confidence limits. The purpose of this correction is to adjust for the difference between the normal approximation and the discrete binomial distribution. See Fleiss, Levin, and Paik (2003) for more information. The continuity-corrected Wald confidence limits for the binomial proportion are computed as

$$\hat{p} \pm (z_{\alpha/2} \times \text{se}(\hat{p}) + (1/2n))$$

Agresti-Coull Confidence Limits If you specify the `CL=AGRESTICOULL binomial-option`, PROC FREQ computes Agresti-Coull confidence limits for the binomial proportion as

$$\tilde{p} \pm (z_{\alpha/2} \times \sqrt{\tilde{p}(1 - \tilde{p}) / \tilde{n}})$$

where

$$\begin{aligned}\tilde{n}_1 &= n_1 + z_{\alpha/2}^2/2 \\ \tilde{n} &= n + z_{\alpha/2}^2 \\ \tilde{p} &= \tilde{n}_1 / \tilde{n}\end{aligned}$$

The Agresti-Coull confidence interval has the same general form as the standard Wald interval but uses $\tilde{p}$ in place of $\hat{p}$. For $\alpha = 0.05$, the value of $z_{\alpha/2}$ is close to 2, and this interval is the “add 2 successes and 2 failures” adjusted Wald interval of Agresti and Coull (1998).

Jeffreys Confidence Limits If you specify the `CL=JEFFREYS binomial-option`, PROC FREQ computes Jeffreys confidence limits for the binomial proportion as

$$(\beta(\alpha/2, n_1 + 1/2, n - n_1 + 1/2), \beta(1 - \alpha/2, n_1 + 1/2, n - n_1 + 1/2))$$

where $\beta(\alpha, b, c)$ is the α percentile of the beta distribution with shape parameters b and c . The lower confidence limit is set to 0 when $n_1 = 0$, and the upper confidence limit is set to 1 when $n_1 = n$. This is an equal-tailed interval based on the noninformative Jeffreys prior for a binomial proportion. See Brown, Cai, and DasGupta (2001) for details. See Berger (1985) for information about using beta priors for inference on the binomial proportion.

Wilson (Score) Confidence Limits If you specify the CL=WILSON *binomial-option*, PROC FREQ computes Wilson confidence limits for the binomial proportion. These are also known as score confidence limits (Wilson 1927). The confidence limits are based on inverting the normal test that uses the null proportion in the variance (the score test). Wilson confidence limits are the roots of

$$|p - \hat{p}| = z_{\alpha/2} \sqrt{p(1-p)/n}$$

and are computed as

$$\left(\hat{p} + z_{\alpha/2}^2/2n \pm z_{\alpha/2} \sqrt{(\hat{p}(1-\hat{p}) + z_{\alpha/2}^2/4n)/n} \right) / \left(1 + z_{\alpha/2}^2/n \right)$$

If you specify CL=WILSON(CORRECT) or the CORRECT *binomial-option*, PROC FREQ provides continuity-corrected Wilson confidence limits, which are computed as the roots of

$$|p - \hat{p}| - 1/2n = z_{\alpha/2} \sqrt{p(1-p)/n}$$

The Wilson interval has been shown to have better performance than the Wald interval and the exact (Clopper-Pearson) interval. For more information, see Agresti and Coull (1998); Brown, Cai, and DasGupta (2001); Newcombe (1998b).

Exact (Clopper-Pearson) Confidence Limits Exact (Clopper-Pearson) confidence limits for the binomial proportion are constructed by inverting the equal-tailed test based on the binomial distribution. This method is attributed to Clopper and Pearson (1934). The exact confidence limits p_L and p_U satisfy the following equations, for $n_1 = 1, 2, \dots, n-1$:

$$\sum_{x=n_1}^n \binom{n}{x} p_L^x (1-p_L)^{n-x} = \alpha/2$$

$$\sum_{x=0}^{n_1} \binom{n}{x} p_U^x (1-p_U)^{n-x} = \alpha/2$$

The lower confidence limit equals 0 when $n_1 = 0$, and the upper confidence limit equals 1 when $n_1 = n$.

PROC FREQ computes the exact (Clopper-Pearson) confidence limits by using the F distribution as

$$p_L = \left(1 + \frac{n - n_1 + 1}{n_1 F(\alpha/2, 2n_1, 2(n - n_1 + 1))} \right)^{-1}$$

$$p_U = \left(1 + \frac{n - n_1}{(n_1 + 1) F(1 - \alpha/2, 2(n_1 + 1), 2(n - n_1))} \right)^{-1}$$

where $F(\alpha/2, b, c)$ is the $\alpha/2$ percentile of the F distribution with b and c degrees of freedom. See Leemis and Trivedi (1996) for a derivation of this expression. Also see Collett (1991) for more information about exact binomial confidence limits.

Because this is a discrete problem, the confidence coefficient (coverage probability) of the exact (Clopper-Pearson) interval is not exactly $(1 - \alpha)$ but is at least $(1 - \alpha)$. Thus, this confidence interval is conservative. Unless the sample size is large, the actual coverage probability can be much larger than the target value. For more information about the performance of these confidence limits, see Agresti and Coull (1998); Brown, Cai, and DasGupta (2001); Leemis and Trivedi (1996).

Binomial Tests

The BINOMIAL option provides an asymptotic equality test for the binomial proportion by default. You can also specify *binomial-options* to request tests of noninferiority, superiority, and equivalence for the binomial proportion. If you specify the BINOMIAL option in the EXACT statement, PROC FREQ also computes exact p -values for the tests that you request with the *binomial-options*.

Equality Test PROC FREQ computes an asymptotic test of the hypothesis that the binomial proportion equals p_0 , where you can specify the value of p_0 with the $P=$ *binomial-option*. If you do not specify a null value with $P=$, PROC FREQ uses $p_0 = 0.5$ by default. The binomial test statistic is computed as

$$z = (\hat{p} - p_0)/se$$

By default, the standard error is based on the null hypothesis proportion as

$$se = \sqrt{p_0(1 - p_0)/n}$$

If you specify the VAR=SAMPLE *binomial-option*, the standard error is computed from the sample proportion as

$$se = \sqrt{\hat{p}(1 - \hat{p})/n}$$

If you specify the CORRECT *binomial-option*, PROC FREQ includes a continuity correction in the asymptotic test statistic, towards adjusting for the difference between the normal approximation and the discrete binomial distribution. See Fleiss, Levin, and Paik (2003) for details. The continuity correction of $(1/2n)$ is subtracted from the numerator of the test statistic if $(\hat{p} - p_0)$ is positive; otherwise, the continuity correction is added to the numerator.

PROC FREQ computes one-sided and two-sided p -values for this test. When the test statistic z is greater than zero (its expected value under the null hypothesis), PROC FREQ computes the right-sided p -value, which is the probability of a larger value of the statistic occurring under the null hypothesis. A small right-sided p -value supports the alternative hypothesis that the true value of the proportion is greater than p_0 . When the test statistic is less than or equal to zero, PROC FREQ computes the left-sided p -value, which is the probability of a smaller value of the statistic occurring under the null hypothesis. A small left-sided p -value supports the alternative hypothesis that the true value of the proportion is less than p_0 . The one-sided p -value P_1 can be expressed as

$$P_1 = \begin{cases} \text{Prob}(Z > z) & \text{if } z > 0 \\ \text{Prob}(Z < z) & \text{if } z \leq 0 \end{cases}$$

where Z has a standard normal distribution. The two-sided p -value is computed as $P_2 = 2 \times P_1$.

If you specify the BINOMIAL option in the EXACT statement, PROC FREQ also computes an exact test of the null hypothesis $H_0: p = p_0$. To compute the exact test, PROC FREQ uses the binomial probability function,

$$\text{Prob}(X = x \mid p_0) = \binom{n}{x} p_0^x (1 - p_0)^{(n-x)} \quad \text{for } x = 0, 1, 2, \dots, n$$

where the variable X has a binomial distribution with parameters n and p_0 . To compute the left-sided p -value, $\text{Prob}(X \leq n_1)$, PROC FREQ sums the binomial probabilities over x from zero to n_1 . To compute

the right-sided p -value, $\text{Prob}(X \geq n_1)$, PROC FREQ sums the binomial probabilities over x from n_1 to n . The exact one-sided p -value is the minimum of the left-sided and right-sided p -values,

$$P_1 = \min(\text{Prob}(X \leq n_1 | p_0), \text{Prob}(X \geq n_1 | p_0))$$

and the exact two-sided p -value is computed as $P_2 = 2 \times P_1$.

Noninferiority Test If you specify the NONINF *binomial-option*, PROC FREQ provides a noninferiority test for the binomial proportion. The null hypothesis for the noninferiority test is

$$H_0: p - p_0 \leq -\delta$$

versus the alternative

$$H_a: p - p_0 > -\delta$$

where δ is the noninferiority margin and p_0 is the null proportion. Rejection of the null hypothesis indicates that the binomial proportion is not inferior to the null value. See Chow, Shao, and Wang (2003) for more information.

You can specify the value of δ with the MARGIN= *binomial-option*, and you can specify p_0 with the P= *binomial-option*. By default, $\delta = 0.2$ and $p_0 = 0.5$.

PROC FREQ provides an asymptotic Wald test for noninferiority. The test statistic is computed as

$$z = (\hat{p} - p_0^*) / \text{se}$$

where p_0^* is the noninferiority limit,

$$p_0^* = p_0 - \delta$$

By default, the standard error is computed from the sample proportion as

$$\text{se} = \sqrt{\hat{p}(1 - \hat{p})/n}$$

If you specify the VAR=NULL *binomial-option*, the standard error is based on the noninferiority limit (determined by the null proportion and the margin) as

$$\text{se} = \sqrt{p_0^*(1 - p_0^*)/n}$$

If you specify the CORRECT *binomial-option*, PROC FREQ includes a continuity correction in the asymptotic test statistic z . The continuity correction of $(1/2n)$ is subtracted from the numerator of the test statistic if $(\hat{p} - p_0^*)$ is positive; otherwise, the continuity correction is added to the numerator.

The p -value for the noninferiority test is

$$P_z = \text{Prob}(Z > z)$$

where Z has a standard normal distribution.

As part of the noninferiority analysis, PROC FREQ provides asymptotic Wald confidence limits for the binomial proportion. These confidence limits are computed as described in the section “[Wald Confidence Limits](#)” on page 2566 but use the same standard error (VAR=NULL or VAR=SAMPLE) as the noninferiority

test statistic z . The confidence coefficient is $100(1 - 2\alpha)\%$ (Schuirmann 1999). By default, if you do not specify the ALPHA= option, the noninferiority confidence limits are 90% confidence limits. You can compare the confidence limits to the noninferiority limit, $p_0^* = p_0 - \delta$.

If you specify the BINOMIAL option in the EXACT statement, PROC FREQ provides an exact noninferiority test for the binomial proportion. The exact p -value is computed by using the binomial probability function with parameters p_0^* and n ,

$$P_x = \sum_{k=n_1}^{k=n} \binom{n}{k} (p_0^*)^k (1 - p_0^*)^{(n-k)}$$

See Chow, Shao, and Wang (2003, p. 116) for details. If you request exact binomial statistics, PROC FREQ also includes exact (Clopper-Pearson) confidence limits for the binomial proportion in the equivalence analysis display. See the section “[Exact \(Clopper-Pearson\) Confidence Limits](#)” on page 2567 for details.

Superiority Test If you specify the SUP *binomial-option*, PROC FREQ provides a superiority test for the binomial proportion. The null hypothesis for the superiority test is

$$H_0: p - p_0 \leq \delta$$

versus the alternative

$$H_a: p - p_0 > \delta$$

where δ is the superiority margin and p_0 is the null proportion. Rejection of the null hypothesis indicates that the binomial proportion is superior to the null value. You can specify the value of δ with the MARGIN= *binomial-option*, and you can specify the value of p_0 with the P= *binomial-option*. By default, $\delta = 0.2$ and $p_0 = 0.5$.

The superiority analysis is identical to the noninferiority analysis but uses a positive value of the margin δ in the null hypothesis. The superiority limit equals $p_0 + \delta$. The superiority computations follow those in the section “[Noninferiority Test](#)” on page 2569 but replace $-\delta$ with δ . See Chow, Shao, and Wang (2003) for more information.

Equivalence Test If you specify the EQUIV *binomial-option*, PROC FREQ provides an equivalence test for the binomial proportion. The null hypothesis for the equivalence test is

$$H_0: p - p_0 \leq \delta_L \quad \text{or} \quad p - p_0 \geq \delta_U$$

versus the alternative

$$H_a: \delta_L < p - p_0 < \delta_U$$

where δ_L is the lower margin, δ_U is the upper margin, and p_0 is the null proportion. Rejection of the null hypothesis indicates that the binomial proportion is equivalent to the null value. See Chow, Shao, and Wang (2003) for more information.

You can specify the value of the margins δ_L and δ_U with the MARGIN= *binomial-option*. If you do not specify MARGIN=, PROC FREQ uses lower and upper margins of -0.2 and 0.2 by default. If you specify a single margin value δ , PROC FREQ uses lower and upper margins of $-\delta$ and δ . You can specify the null proportion p_0 with the P= *binomial-option*. By default, $p_0 = 0.5$.

PROC FREQ computes two one-sided tests (TOST) for equivalence analysis (Schuirmann 1987). The TOST approach includes a right-sided test for the lower margin and a left-sided test for the upper margin. The overall p -value is taken to be the larger of the two p -values from the lower and upper tests.

For the lower margin, the asymptotic Wald test statistic is computed as

$$z_L = (\hat{p} - p_L^*) / \text{se}$$

where the lower equivalence limit is

$$p_L^* = p_0 + \delta_L$$

By default, the standard error is computed from the sample proportion as

$$\text{se} = \sqrt{\hat{p}(1 - \hat{p})/n}$$

If you specify the VAR=NULL *binomial-option*, the standard error is based on the lower equivalence limit (determined by the null proportion and the lower margin) as

$$\text{se} = \sqrt{p_L^*(1 - p_L^*)/n}$$

If you specify the CORRECT *binomial-option*, PROC FREQ includes a continuity correction in the asymptotic test statistic z_L . The continuity correction of $(1/2n)$ is subtracted from the numerator of the test statistic $(\hat{p} - p_L^*)$ if the numerator is positive; otherwise, the continuity correction is added to the numerator.

The p -value for the lower margin test is

$$P_{z,L} = \text{Prob}(Z > z_L)$$

The asymptotic test for the upper margin is computed similarly. The Wald test statistic is

$$z_U = (\hat{p} - p_U^*) / \text{se}$$

where the upper equivalence limit is

$$p_U^* = p_0 + \delta_U$$

By default, the standard error is computed from the sample proportion. If you specify the VAR=NULL *binomial-option*, the standard error is based on the upper equivalence limit as

$$\text{se} = \sqrt{p_U^*(1 - p_U^*)/n}$$

If you specify the CORRECT *binomial-option*, PROC FREQ includes a continuity correction of $(1/2n)$ in the asymptotic test statistic z_U .

The p -value for the upper margin test is

$$P_{z,U} = \text{Prob}(Z < z_U)$$

Based on the two one-sided tests (TOST), the overall p -value for the test of equivalence equals the larger p -value from the lower and upper margin tests, which can be expressed as

$$P_z = \max(P_{z,L}, P_{z,U})$$

As part of the equivalence analysis, PROC FREQ provides asymptotic Wald confidence limits for the binomial proportion. These confidence limits are computed as described in the section “[Wald Confidence Limits](#)” on page 2566, but use the same standard error (VAR=NULL or VAR=SAMPLE) as the equivalence test statistics and have a confidence coefficient of $100(1 - 2\alpha)\%$ (Schuirmann 1999). By default, if you do not specify the ALPHA= option, the equivalence confidence limits are 90% limits. If you specify VAR=NULL, separate standard errors are computed for the lower and upper margin tests, each based on the null proportion and the corresponding (lower or upper) margin. The confidence limits are computed by using the maximum of these two standard errors. You can compare the confidence limits to the equivalence limits, $(p_0 + \delta_L, p_0 + \delta_U)$.

If you specify the BINOMIAL option in the EXACT statement, PROC FREQ also provides an exact equivalence test by using two one-sided exact tests (TOST). The procedure computes lower and upper margin exact tests by using the binomial probability function as described in the section “[Noninferiority Test](#)” on page 2569. The overall exact p -value for the equivalence test is taken to be the larger p -value from the lower and upper margin exact tests. If you request exact statistics, PROC FREQ also includes exact (Clopper-Pearson) confidence limits in the equivalence analysis display. The confidence coefficient is $100(1 - 2\alpha)\%$ (Schuirmann 1999). See the section “[Exact \(Clopper-Pearson\) Confidence Limits](#)” on page 2567 for details.

Risks and Risk Differences

The RISKDIFF option in the TABLES statement provides estimates of risks (binomial proportions) and risk differences for 2×2 tables. This analysis might be appropriate when comparing the proportion of some characteristic for two groups, where row 1 and row 2 correspond to the two groups, and the columns correspond to two possible characteristics or outcomes. For example, the row variable might be a treatment or dose, and the column variable might be the response. For more information, see Collett (1991); Fleiss, Levin, and Paik (2003); Stokes, Davis, and Koch (2012).

Let the frequencies of the 2×2 table be represented as follows.

	Column 1	Column 2	Total
Row 1	n_{11}	n_{12}	$n_{1\cdot}$
Row 2	n_{21}	n_{22}	$n_{2\cdot}$
Total	$n_{\cdot 1}$	$n_{\cdot 2}$	n

By default when you specify the RISKDIFF option, PROC FREQ provides estimates of the row 1 risk (proportion), the row 2 risk, the overall risk, and the risk difference for column 1 and for column 2 of the 2×2 table. The risk difference is defined as the row 1 risk minus the row 2 risk. The risks are binomial proportions of their rows (row 1, row 2, or overall), and the computation of their standard errors and Wald confidence limits follow the binomial proportion computations, which are described in the section “[Binomial Proportion](#)” on page 2565.

The column 1 risk for row 1 is the proportion of row 1 observations classified in column 1,

$$\hat{p}_1 = n_{11} / n_{1\cdot}$$

which estimates the conditional probability of the column 1 response, given the first level of the row variable. The column 1 risk for row 2 is the proportion of row 2 observations classified in column 1,

$$\hat{p}_2 = n_{21} / n_{2\cdot}$$

The overall column 1 risk is the proportion of all observations classified in column 1,

$$\hat{p} = n_{.1} / n$$

The column 1 risk difference compares the risks for the two rows, and it is computed as the column 1 risk for row 1 minus the column 1 risk for row 2,

$$\hat{d} = \hat{p}_1 - \hat{p}_2$$

The standard error of the column 1 risk for row i is computed as

$$se(\hat{p}_i) = \sqrt{\hat{p}_i (1 - \hat{p}_i) / n_{i.}}$$

The standard error of the overall column 1 risk is computed as

$$se(\hat{p}) = \sqrt{\hat{p} (1 - \hat{p}) / n}$$

Where the two rows represent independent binomial samples, the standard error of the column 1 risk difference is computed as

$$se(\hat{d}) = \sqrt{\hat{p}_1(1 - \hat{p}_1) + \hat{p}_2(1 - \hat{p}_2)}$$

The computations are similar for the column 2 risks and risk difference.

Confidence Limits

By default, the RISKDIFF option provides Wald asymptotic confidence limits for the risks (row 1, row 2, and overall) and the risk difference. By default, the RISKDIFF option also provides exact (Clopper-Pearson) confidence limits for the risks. You can suppress the display of this information by specifying the NORISKS *riskdiff-option*. You can specify *riskdiff-options* to request tests and other types of confidence limits for the risk difference. See the sections “[Risk Difference Confidence Limits](#)” on page 2574 and “[Risk Difference Tests](#)” on page 2576 for more information.

The risks are equivalent to the binomial proportions of their corresponding rows. This section describes the Wald confidence limits that are provided by default when you specify the RISKDIFF option. The BINOMIAL option provides additional confidence limit types and tests for risks (binomial proportions). See the sections “[Binomial Confidence Limits](#)” on page 2566 and “[Binomial Tests](#)” on page 2568 for details.

The Wald confidence limits are based on the normal approximation to the binomial distribution. PROC FREQ computes the Wald confidence limits for the risks and risk differences as

$$\text{Est} \pm (z_{\alpha/2} \times se(\text{Est}))$$

where Est is the estimate, $z_{\alpha/2}$ is the 100(1 - $\alpha/2$) percentile of the standard normal distribution, and $se(\text{Est})$ is the standard error of the estimate. The confidence level α is determined by the value of the ALPHA= option; the default of ALPHA=0.5 produces 95% confidence limits.

If you specify the CORRECT *riskdiff-option*, PROC FREQ includes continuity corrections in the Wald confidence limits for the risks and risk differences. The purpose of a continuity correction is to adjust for the difference between the normal approximation and the binomial distribution, which is discrete. See Fleiss, Levin, and Paik (2003) for more information. The continuity-corrected Wald confidence limits are computed as

$$\text{Est} \pm (z_{\alpha/2} \times se(\text{Est}) + cc)$$

where cc is the continuity correction. For the row 1 risk, $cc = (1/2n_{1.})$; for the row 2 risk, $cc = (1/2n_{2.})$; for the overall risk, $cc = (1/2n)$; and for the risk difference, $cc = ((1/n_{1.} + 1/n_{2.})/2)$. The column 1 and column 2 risks use the same continuity corrections.

By default when you specify the RISKDIFF option, PROC FREQ also provides exact (Clopper-Pearson) confidence limits for the column 1, column 2, and overall risks. These confidence limits are constructed by inverting the equal-tailed test that is based on the binomial distribution. See the section “[Exact \(Clopper-Pearson\) Confidence Limits](#)” on page 2567 for details.

Risk Difference Confidence Limits You can request additional confidence limits for the risk difference by specifying the `CL= riskdiff-option`. Available confidence limit types include Agresti-Caffo, exact unconditional, Hauck-Anderson, Miettinen-Nurminen, Newcombe (hybrid-score), and Wald confidence limits. Continuity-corrected Newcombe and Wald confidence limits are also available.

The confidence coefficient for the confidence limits produced by the `CL= riskdiff-option` is $100(1 - \alpha)\%$, where the value of α is determined by the `ALPHA=` option. The default of `ALPHA=0.05` produces 95% confidence limits. This differs from the test-based confidence limits that are provided with the equivalence, noninferiority, and superiority tests, which have a confidence coefficient of $100(1 - 2\alpha)\%$ (Schuirmann 1999). See the section “[Risk Difference Tests](#)” on page 2576 for details.

The section “[Exact Unconditional Confidence Limits for the Risk Difference](#)” on page 2580 describes the computation of the exact confidence limits. The confidence limits are constructed by inverting two separate one-sided exact tests (tail method). By default, the tests are based on the unstandardized risk difference. If you specify the `RISKDIFF(METHOD=SCORE)` option, the tests are based on the score statistic.

The following sections describe the computation of the Agresti-Coull, Hauck-Anderson, Miettinen-Nurminen, Newcombe (hybrid-score), and Wald confidence limits for the risk difference.

Agresti-Caffo Confidence Limits The Agresti-Caffo confidence limits for the risk difference are computed as

$$\tilde{d} \pm (z_{\alpha/2} \times \text{se}(\tilde{d}))$$

where $\tilde{d} = \tilde{p}_1 - \tilde{p}_2$, $\tilde{p}_i = (n_{i1} + 1)/(n_{i.} + 2)$,

$$\text{se}(\tilde{d}) = \sqrt{\tilde{p}_1(1 - \tilde{p}_1)/(n_{1.} + 2) + \tilde{p}_2(1 - \tilde{p}_2)/(n_{2.} + 2)}$$

and $z_{\alpha/2}$ is the $100(1 - \alpha/2)$ percentile of the standard normal distribution.

The Agresti-Caffo interval adjusts the Wald interval for the risk difference by adding a pseudo-observation of each type (success and failure) to each sample. See Agresti and Caffo (2000) and Agresti and Coull (1998) for more information.

Hauck-Anderson Confidence Limits The Hauck-Anderson confidence limits for the risk difference are computed as

$$\hat{d} \pm (cc + z_{\alpha/2} \times \text{se}(\hat{d}))$$

where $\hat{d} = \hat{p}_1 - \hat{p}_2$ and $z_{\alpha/2}$ is the $100(1 - \alpha/2)$ percentile of the standard normal distribution. The standard error is computed from the sample proportions as

$$\text{se}(\hat{d}) = \sqrt{\hat{p}_1(1 - \hat{p}_1)/(n_{1.} - 1) + \hat{p}_2(1 - \hat{p}_2)/(n_{2.} - 1)}$$

The Hauck-Anderson continuity correction cc is computed as

$$cc = 1 / (2 \min(n_{1.}, n_{.2}))$$

See Hauck and Anderson (1986) for more information. The subsection “[Hauck-Anderson Test](#)” in the section “[Noninferiority Tests](#)” on page 2577 describes the corresponding noninferiority test.

Miettinen-Nurminen Confidence Limits The Miettinen-Nurminen confidence limits for the risk difference (Miettinen and Nurminen 1985; Newcombe and Nurminen 2011) are based on the likelihood score test for the null hypothesis that the risk difference is zero. The test statistic is

$$T(\delta) = (\hat{d} - \delta) / \sqrt{\widetilde{\text{Var}}(\delta)}$$

where $\hat{d}$ is the observed value of the risk difference ($\hat{p}_1 - \hat{p}_2$),

$$\widetilde{\text{Var}}(\delta) = (n/(n-1)) (\tilde{p}_1(\delta)(1 - \tilde{p}_1(\delta))/n_1 + \tilde{p}_2(\delta)(1 - \tilde{p}_2(\delta))/n_2)$$

and $\tilde{p}_1(\delta)$ and $\tilde{p}_2(\delta)$ are the maximum likelihood estimates of the risks (proportions) when the risk difference is δ . PROC FREQ computes the maximum likelihood estimates $\tilde{p}_1(\delta)$ and $\tilde{p}_2(\delta)$ as described in the subsection “[Farrington-Manning Test](#)” in the section “[Noninferiority Tests](#)” on page 2577.

The $100(1 - \alpha)\%$ confidence interval for the risk difference consists of all values of δ for which the score test statistic $T(\delta)$ falls in the acceptance region,

$$\{\delta : T(\delta) < z_{\alpha/2}\}$$

PROC FREQ finds the confidence limits by iterative computation, which stops when the iteration increment falls below the convergence criterion or when the maximum number of iterations is reached, whichever occurs first. By default, the convergence criterion is 0.00000001 and the maximum number of iterations is 100.

By default, the Miettinen-Nurminen confidence limits include the bias correction factor $n/(n-1)$ in the computation of $\widetilde{\text{Var}}(\delta)$, as given by Miettinen and Nurminen (1985, p. 216). See Newcombe and Nurminen (2011) for more information. If you specify the `CL=MN(CORRECT=NO) riskdiff-option`, PROC FREQ does not include the bias correction factor in the computation of the confidence limits (Mee 1984). See also Agresti (2002, p. 77). The uncorrected confidence limits are labeled as “Miettinen-Nurminen-Mee” confidence limits in the displayed output.

Newcombe Confidence Limits Newcombe (hybrid-score) confidence limits for the risk difference are constructed from the Wilson score confidence limits for each of the two individual proportions. The confidence limits for the individual proportions are used in the standard error terms of the Wald confidence limits for the proportion difference. See Newcombe (1998a) and Barker et al. (2001) for more information.

Wilson score confidence limits for p_1 and p_2 are the roots of

$$|p_i - \hat{p}_i| = z_{\alpha/2} \sqrt{p_i(1 - p_i)/n_i}$$

for $i = 1, 2$. The confidence limits are computed as

$$\left(\hat{p}_i + z_{\alpha/2}^2/2n_i \pm z_{\alpha/2} \sqrt{(\hat{p}_i(1 - \hat{p}_i) + z_{\alpha/2}^2/4n_i)/n_i} \right) / \left(1 + z_{\alpha/2}^2/n_i \right)$$

See the section “[Wilson \(Score\) Confidence Limits](#)” on page 2567 for details.

Denote the lower and upper Wilson score confidence limits for p_1 as L_1 and U_1 , and denote the lower and upper confidence limits for p_2 as L_2 and U_2 . The Newcombe confidence limits for the proportion difference ($d = p_1 - p_2$) are computed as

$$d_L = (\hat{p}_1 - \hat{p}_2) - \sqrt{(\hat{p}_1 - L_1)^2 + (U_2 - \hat{p}_2)^2}$$

$$d_U = (\hat{p}_1 - \hat{p}_2) + \sqrt{(U_1 - \hat{p}_1)^2 + (\hat{p}_2 - L_2)^2}$$

If you specify the **CORRECT** *riskdiff-option*, PROC FREQ provides continuity-corrected Newcombe confidence limits. By including a continuity correction of $1/2n_{i\cdot}$, the Wilson score confidence limits for the individual proportions are computed as the roots of

$$|p_i - \hat{p}_i| - 1/2n_{i\cdot} = z_{\alpha/2} \sqrt{p_i(1 - p_i)/n_{i\cdot}}$$

The continuity-corrected confidence limits for the individual proportions are then used to compute the proportion difference confidence limits d_L and d_U .

Wald Confidence Limits The Wald confidence limits for the risk difference are computed as

$$\hat{d} \pm (z_{\alpha/2} \times \text{se}(\hat{d}))$$

where $\hat{d} = \hat{p}_1 - \hat{p}_2$, $z_{\alpha/2}$ is the $100(1 - \alpha/2)$ percentile of the standard normal distribution, and the standard error is computed from the sample proportions as

$$\text{se}(\hat{d}) = \sqrt{\hat{p}_1(1 - \hat{p}_1)/n_{1\cdot} + \hat{p}_2(1 - \hat{p}_2)/n_{2\cdot}}$$

If you specify the **CORRECT** *riskdiff-option*, the Wald confidence limits include a continuity correction cc ,

$$\hat{d} \pm (cc + z_{\alpha/2} \times \text{se}(\hat{d}))$$

where $cc = (1/n_{1\cdot} + 1/n_{2\cdot})/2$.

The subsection “**Wald Test**” in the section “**Noninferiority Tests**” on page 2577 describes the corresponding noninferiority test.

Risk Difference Tests

You can specify *riskdiff-options* to request tests of the risk (proportion) difference. You can request tests of equality, noninferiority, superiority, and equivalence for the risk difference. The test of equality is a standard Wald asymptotic test, available with or without a continuity correction. For noninferiority, superiority, and equivalence tests of the risk difference, the following test methods are provided: Wald (with and without continuity correction), Hauck-Anderson, Farrington-Manning (score), and Newcombe (with and without continuity correction). You can specify the test method with the **METHOD=** *riskdiff-option*. By default, PROC FREQ uses **METHOD=WALD**.

Equality Test If you specify the **EQUAL** *riskdiff-option*, PROC FREQ computes a test of equality, or a test of the null hypothesis that the risk difference equals zero. For the column 1 (or 2) risk difference, this test can be expressed as $H_0: d = 0$ versus the alternative $H_a: d \neq 0$, where $d = p_1 - p_2$ denotes the column 1 (or 2) risk difference. PROC FREQ provides a Wald asymptotic test of equality. The test statistic is computed as

$$z = \hat{d} / \text{se}(\hat{d})$$

By default, the standard error is computed from the sample proportions as

$$se(\hat{d}) = \sqrt{\hat{p}_1(1 - \hat{p}_1)/n_{1\cdot} + \hat{p}_2(1 - \hat{p}_2)/n_{2\cdot}}.$$

If you specify the VAR=NULL *riskdiff-option*, the standard error is based on the null hypothesis that the row 1 and row 2 risks are equal,

$$se(\hat{d}) = \sqrt{\hat{p}(1 - \hat{p}) \times (1/n_{1\cdot} + 1/n_{2\cdot})}$$

where $\hat{p} = n_{\cdot 1}/n$ estimates the overall column 1 risk.

If you specify the CORRECT *riskdiff-option*, PROC FREQ includes a continuity correction in the test statistic. If $\hat{d} > 0$, the continuity correction is subtracted from $\hat{d}$ in the numerator of the test statistic; otherwise, the continuity correction is added to the numerator. The value of the continuity correction is $(1/n_{1\cdot} + 1/n_{2\cdot})/2$.

PROC FREQ computes one-sided and two-sided p -values for this test. When the test statistic z is greater than 0, PROC FREQ displays the right-sided p -value, which is the probability of a larger value occurring under the null hypothesis. The one-sided p -value can be expressed as

$$P_1 = \begin{cases} \text{Prob}(Z > z) & \text{if } z > 0 \\ \text{Prob}(Z < z) & \text{if } z \leq 0 \end{cases}$$

where Z has a standard normal distribution. The two-sided p -value is computed as $P_2 = 2 \times P_1$.

Noninferiority Tests If you specify the NONINF *riskdiff-option*, PROC FREQ provides a noninferiority test for the risk difference, or the difference between two proportions. The null hypothesis for the noninferiority test is

$$H_0: p_1 - p_2 \leq -\delta$$

versus the alternative

$$H_a: p_1 - p_2 > -\delta$$

where δ is the noninferiority margin. Rejection of the null hypothesis indicates that the row 1 risk is not inferior to the row 2 risk. See Chow, Shao, and Wang (2003) for more information.

You can specify the value of δ with the MARGIN= *riskdiff-option*. By default, $\delta = 0.2$. You can specify the test method with the METHOD= *riskdiff-option*. The following methods are available for the risk difference noninferiority analysis: Wald (with and without continuity correction), Hauck-Anderson, Farrington-Manning (score), and Newcombe (with and without continuity correction). The Wald, Hauck-Anderson, and Farrington-Manning methods provide tests and corresponding test-based confidence limits; the Newcombe method provides only confidence limits. If you do not specify METHOD=, PROC FREQ uses the Wald test by default.

The confidence coefficient for the test-based confidence limits is $100(1 - 2\alpha)\%$ (Schuirmann 1999). By default, if you do not specify the ALPHA= option, these are 90% confidence limits. You can compare the confidence limits to the noninferiority limit, $-\delta$.

The following sections describe the noninferiority analysis methods for the risk difference.

Wald Test If you specify the METHOD=WALD *riskdiff-option*, PROC FREQ provides an asymptotic Wald test of noninferiority for the risk difference. This is also the default method. The Wald test statistic is computed as

$$z = (\hat{d} + \delta) / \text{se}(\hat{d})$$

where $\hat{d} = \hat{p}_1 - \hat{p}_2$ estimates the risk difference and δ is the noninferiority margin.

By default, the standard error for the Wald test is computed from the sample proportions as

$$\text{se}(\hat{d}) = \sqrt{\hat{p}_1(1 - \hat{p}_1)/n_{1\cdot} + \hat{p}_2(1 - \hat{p}_2)/n_{2\cdot}}$$

If you specify the VAR=NULL *riskdiff-option*, the standard error is based on the null hypothesis that the risk difference equals $-\delta$ (Dunnett and Gent 1977). The standard error is computed as

$$\text{se}(\hat{d}) = \sqrt{\tilde{p}(1 - \tilde{p})/n_{2\cdot} + (\tilde{p} - \delta)(1 - \tilde{p} + \delta)/n_{1\cdot}}$$

where

$$\tilde{p} = (n_{11} + n_{21} + \delta n_{1\cdot})/n$$

If you specify the CORRECT *riskdiff-option*, the test statistic includes a continuity correction. The continuity correction is subtracted from the numerator of the test statistic if the numerator is greater than zero; otherwise, the continuity correction is added to the numerator. The value of the continuity correction is $(1/n_{1\cdot} + 1/n_{2\cdot})/2$.

The p -value for the Wald noninferiority test is $P_z = \text{Prob}(Z > z)$, where Z has a standard normal distribution.

Hauck-Anderson Test If you specify the METHOD=HA *riskdiff-option*, PROC FREQ provides the Hauck-Anderson test for noninferiority. The Hauck-Anderson test statistic is computed as

$$z = (\hat{d} + \delta \pm cc) / \text{se}(\hat{d})$$

where $\hat{d} = \hat{p}_1 - \hat{p}_2$ and the standard error is computed from the sample proportions as

$$\text{se}(\hat{d}) = \sqrt{\hat{p}_1(1 - \hat{p}_1)/(n_{1\cdot} - 1) + \hat{p}_2(1 - \hat{p}_2)/(n_{2\cdot} - 1)}$$

The Hauck-Anderson continuity correction cc is computed as

$$cc = 1 / (2 \min(n_{1\cdot}, n_{2\cdot}))$$

The p -value for the Hauck-Anderson noninferiority test is $P_z = \text{Prob}(Z > z)$, where Z has a standard normal distribution. See Hauck and Anderson (1986) and Schuirmann (1999) for more information.

Farrington-Manning Test If you specify the METHOD=FM *riskdiff-option*, PROC FREQ provides the Farrington-Manning (score) test of noninferiority for the risk difference. The score test statistic is computed as

$$z = (\hat{d} + \delta) / \text{se}(\hat{d})$$

where $\hat{d} = \hat{p}_1 - \hat{p}_2$ and

$$\text{se}(\hat{d}) = \sqrt{\tilde{p}_1(1 - \tilde{p}_1)/n_{1\cdot} + \tilde{p}_2(1 - \tilde{p}_2)/n_{2\cdot}}$$

where $\tilde{p}_1$ and $\tilde{p}_2$ are the maximum likelihood estimators of p_1 and p_2 under the null hypothesis that the risk difference equals $-\delta$ (Miettinen and Nurminen 1985; Farrington and Manning 1990; Dann and Koch 2005). The p -value for the noninferiority test is then $P_z = \text{Prob}(Z > z)$, where Z has a standard normal distribution.

From Farrington and Manning (1990, p. 1453), the solution to the maximum likelihood equation is

$$\tilde{p}_1 = 2u \cos(w) - b/3a \quad \text{and} \quad \tilde{p}_2 = \tilde{p}_1 + \delta$$

where

$$\begin{aligned} w &= (\pi + \cos^{-1}(v/u^3))/3 \\ v &= b^3/(3a)^3 - bc/6a^2 + d/2a \\ u &= \text{sign}(v) \sqrt{b^2/(3a)^2 - c/3a} \\ a &= 1 + \theta \\ b &= -(1 + \theta + \hat{p}_1 + \theta \hat{p}_2 - \delta(\theta + 2)) \\ c &= \delta^2 - \delta(2\hat{p}_1 + \theta + 1) + \hat{p}_1 + \theta \hat{p}_2 \\ d &= \hat{p}_1 \delta(1 - \delta) \\ \theta &= n_{2\cdot}/n_{1\cdot}. \end{aligned}$$

Newcombe Noninferiority Analysis If you specify the `METHOD=NEWCOMBE riskdiff-option`, PROC FREQ provides a noninferiority analysis that is based on Newcombe hybrid-score confidence limits for the risk difference. The confidence coefficient for the confidence limits is $100(1 - 2\alpha)\%$ (Schuirmann 1999). By default, if you do not specify the `ALPHA=` option, these are 90% confidence limits. You can compare the confidence limits with the noninferiority limit, $-\delta$. If you specify the `CORRECT riskdiff-option`, the confidence limits includes a continuity correction. See the subsection “Newcombe Confidence Limits” in the section “[Risk Difference Confidence Limits](#)” on page 2574 for more information.

Superiority Test If you specify the `SUP riskdiff-option`, PROC FREQ provides a superiority test for the risk difference. The null hypothesis is

$$H_0: p_1 - p_2 \leq \delta$$

versus the alternative

$$H_a: p_1 - p_2 > \delta$$

where δ is the superiority margin. Rejection of the null hypothesis indicates that the row 1 proportion is superior to the row 2 proportion. You can specify the value of δ with the `MARGIN= riskdiff-option`. By default, $\delta = 0.2$.

The superiority analysis is identical to the noninferiority analysis but uses a positive value of the margin δ in the null hypothesis. The superiority computations follow those in the section “[Noninferiority Tests](#)” on page 2577 by replacing $-\delta$ by δ . See Chow, Shao, and Wang (2003) for more information.

Equivalence Tests If you specify the *EQUIV riskdiff-option*, PROC FREQ provides an equivalence test for the risk difference, or the difference between two proportions. The null hypothesis for the equivalence test is

$$H_0: p_1 - p_2 \leq -\delta_L \quad \text{or} \quad p_1 - p_2 \geq \delta_U$$

versus the alternative

$$H_a: \delta_L < p_1 - p_2 < \delta_U$$

where δ_L is the lower margin and δ_U is the upper margin. Rejection of the null hypothesis indicates that the two binomial proportions are equivalent. See Chow, Shao, and Wang (2003) for more information.

You can specify the value of the margins δ_L and δ_U with the *MARGIN= riskdiff-option*. If you do not specify *MARGIN=*, PROC FREQ uses lower and upper margins of -0.2 and 0.2 by default. If you specify a single margin value δ , PROC FREQ uses lower and upper margins of $-\delta$ and δ . You can specify the test method with the *METHOD= riskdiff-option*. The following methods are available for the risk difference equivalence analysis: Wald (with and without continuity correction), Hauck-Anderson, Farrington-Manning, and Newcombe (with and without continuity correction). The Wald, Hauck-Anderson, and Farrington-Manning methods provide tests and corresponding test-based confidence limits; the Newcombe method provides only confidence limits. If you do not specify *METHOD=*, PROC FREQ uses the Wald test by default.

PROC FREQ computes two one-sided tests (TOST) for equivalence analysis (Schuirmann 1987). The TOST approach includes a right-sided test for the lower margin δ_L and a left-sided test for the upper margin δ_U . The overall p -value is taken to be the larger of the two p -values from the lower and upper tests.

The section “**Noninferiority Tests**” on page 2577 gives details about the Wald, Hauck-Anderson, Farrington-Manning and Newcombe methods for the risk difference. The lower margin equivalence test statistic takes the same form as the noninferiority test statistic but uses the lower margin value δ_L in place of $-\delta$. The upper margin equivalence test statistic take the same form as the noninferiority test statistic but uses the upper margin value δ_U in place of $-\delta$.

The test-based confidence limits for the risk difference are computed according to the equivalence test method that you select. If you specify *METHOD=WALD* with *VAR=NULL*, or *METHOD=FM*, separate standard errors are computed for the lower and upper margin tests. In this case, the test-based confidence limits are computed by using the maximum of these two standard errors. The confidence limits have a confidence coefficient of $100(1 - 2\alpha)\%$ (Schuirmann 1999). By default, if you do not specify the *ALPHA=* option, these are 90% confidence limits. You can compare the confidence limits to the equivalence limits, (δ_L, δ_U) .

Exact Unconditional Confidence Limits for the Risk Difference

If you specify the *RISKDIFF* option in the *EXACT* statement, PROC FREQ provides exact unconditional confidence limits for the risk difference. PROC FREQ computes the confidence limits by inverting two separate one-sided tests (tail method), where the size of each test is at most $\alpha/2$ and the confidence coefficient is at least $(1 - \alpha)$. Exact conditional methods, described in the section “**Exact Statistics**” on page 2601, do not apply to the risk difference due to the presence of a nuisance parameter (Agresti 1992). The unconditional approach eliminates the nuisance parameter by maximizing the p -value over all possible values of the parameter (Santner and Snell 1980).

By default, PROC FREQ uses the unstandardized risk difference as the test statistic in the confidence limit computations. If you specify the *RISKDIFF(METHOD=SCORE)* option, the procedure uses the score

statistic (Chan and Zhang 1999). The score statistic is a less discrete statistic than the raw risk difference and produces less conservative confidence limits (Agresti and Min 2001). See also Santner et al. (2007) for comparisons of methods for computing exact confidence limits for the risk difference.

PROC FREQ computes the confidence limits as follows. The risk difference is defined as the difference between the row 1 and row 2 risks (proportions), $d = p_1 - p_2$, and n_1 and n_2 denote the row totals of the 2×2 table. The joint probability function for the table can be expressed in terms of the table cell frequencies, the risk difference, and the nuisance parameter p_2 as

$$f(n_{11}, n_{21}; n_1, n_2, d, p_2) = \binom{n_1}{n_{11}} (d + p_2)^{n_{11}} (1 - d - p_2)^{n_1 - n_{11}} \times \binom{n_2}{n_{21}} p_2^{n_{21}} (1 - p_2)^{n_2 - n_{21}}$$

The $100(1 - \alpha/2)\%$ confidence limits for the risk difference are computed as

$$\begin{aligned} d_L &= \sup (d_* : P_U(d_*) > \alpha/2) \\ d_U &= \inf (d_* : P_L(d_*) > \alpha/2) \end{aligned}$$

where

$$\begin{aligned} P_U(d_*) &= \sup_{p_2} \left(\sum_{A, T(a) \geq t_0} f(n_{11}, n_{21}; n_1, n_2, d_*, p_2) \right) \\ P_L(d_*) &= \sup_{p_2} \left(\sum_{A, T(a) \leq t_0} f(n_{11}, n_{21}; n_1, n_2, d_*, p_2) \right) \end{aligned}$$

The set A includes all 2×2 tables with row sums equal to n_1 and n_2 , and $T(a)$ denotes the value of the test statistic for table a in A . To compute $P_U(d_*)$, the sum includes probabilities of those tables for which $(T(a) \geq t_0)$, where t_0 is the value of the test statistic for the observed table. For a fixed value of d_* , $P_U(d_*)$ is taken to be the maximum sum over all possible values of p_2 .

By default, PROC FREQ uses the unstandardized risk difference as the test statistic T . If you specify the `RISKDIFF(METHOD=SCORE)` option, the procedure uses the risk difference score statistic as the test statistic (Chan and Zhang 1999). The computation of the risk difference score statistic is described in the subsection “[Farrington-Manning Test](#)” in the section “[Noninferiority Tests](#)” on page 2577. See Miettinen and Nurminen (1985) and Farrington and Manning (1990) for more information.

Barnard's Unconditional Exact Test

The BARNARD option in the EXACT statement provides an unconditional exact test for the risk (proportion) difference for 2×2 tables. The reference set for the unconditional exact test consists of all 2×2 tables that have the same row sums as the observed table (Barnard 1945, 1947, 1949). This differs from the reference set for exact conditional inference, which is restricted to the set of tables that have the same row sums and the same column sums as the observed table. See the sections “[Fisher's Exact Test](#)” on page 2554 and “[Exact Statistics](#)” on page 2601 for more information.

The test statistic is the standardized risk difference, which is computed as

$$T = d / \sqrt{p_{\cdot 1}(1 - p_{\cdot 1})(1/n_1 + 1/n_2)}$$

where the risk difference d is defined as the difference between the row 1 and row 2 risks (proportions), $d = (n_{11}/n_1 - n_{21}/n_2)$; n_1 and n_2 are the row 1 and row 2 totals, respectively; and $p_{.1}$ is the overall proportion in column 1, $(n_{11} + n_{21})/n$.

Under the null hypothesis that the risk difference is 0, the joint probability function for a table can be expressed in terms of the table cell frequencies, the row totals, and the unknown parameter π as

$$f(n_{11}, n_{21}; n_1, n_2, \pi) = \binom{n_1}{n_{11}} \binom{n_2}{n_{21}} \pi^{n_{11}+n_{21}} (1-\pi)^{n-n_{11}-n_{21}}$$

where π is the common value of the risk (proportion).

PROC FREQ sums the table probabilities over the reference set for those tables where the test statistic is greater than or equal to the observed value of the test statistic. This sum can be expressed as

$$\text{Prob}(\pi) = \sum_{A, T(a) \geq t_0} f(n_{11}, n_{21}; n_1, n_2, \pi)$$

where the set A contains all 2×2 tables with row sums equal to n_1 and n_2 , and $T(a)$ denotes the value of the test statistic for table a in A . The sum includes probabilities of those tables for which $(T(a) \geq t_0)$, where t_0 is the value of the test statistic for the observed table.

The sum $\text{Prob}(\pi)$ depends on the unknown value of π . To compute the exact p -value, PROC FREQ eliminates the nuisance parameter π by taking the maximum value of $\text{Prob}(\pi)$ over all possible values of π ,

$$\text{Prob} = \sup_{(0 \leq \pi \leq 1)} (\text{Prob}(\pi))$$

See Suissa and Shuster (1985) and Mehta and Senchaudhuri (2003).

Odds Ratio and Relative Risks for 2 x 2 Tables

Odds Ratio (Case-Control Studies)

The odds ratio is a useful measure of association for a variety of study designs. For a retrospective design called a *case-control study*, the odds ratio can be used to estimate the relative risk when the probability of positive response is small (Agresti 2002). In a case-control study, two independent samples are identified based on a binary (yes-no) response variable, and the conditional distribution of a binary explanatory variable is examined, within fixed levels of the response variable. See Stokes, Davis, and Koch (2012) and Agresti (2007).

The odds of a positive response (column 1) in row 1 is n_{11}/n_{12} . Similarly, the odds of a positive response in row 2 is n_{21}/n_{22} . The odds ratio is formed as the ratio of the row 1 odds to the row 2 odds. The odds ratio for a 2×2 table is defined as

$$OR = \frac{n_{11}/n_{12}}{n_{21}/n_{22}} = \frac{n_{11} n_{22}}{n_{12} n_{21}}$$

The odds ratio can be any nonnegative number. When the row and column variables are independent, the true value of the odds ratio equals 1. An odds ratio greater than 1 indicates that the odds of a positive response are higher in row 1 than in row 2. Values less than 1 indicate the odds of positive response are higher in row 2. The strength of association increases with the deviation from 1.

The transformation $G = (OR - 1)/(OR + 1)$ transforms the odds ratio to the range $(-1, 1)$ with $G = 0$ when $OR = 1$; $G = -1$ when $OR = 0$; and G approaches 1 as OR approaches infinity. G is the gamma statistic, which PROC FREQ computes when you specify the MEASURES option.

The asymptotic $100(1 - \alpha)\%$ confidence limits for the odds ratio are

$$(OR \times \exp(-z\sqrt{v}), OR \times \exp(z\sqrt{v}))$$

where

$$v = \text{Var}(\ln OR) = \frac{1}{n_{11}} + \frac{1}{n_{12}} + \frac{1}{n_{21}} + \frac{1}{n_{22}}$$

and z is the $100(1 - \alpha/2)$ percentile of the standard normal distribution. If any of the four cell frequencies are zero, the estimates are not computed.

Exact Confidence Limits for the Odds Ratio When you specify the OR option in the EXACT statement, PROC FREQ computes exact confidence limits for the odds ratio. Because this is a discrete problem, the confidence coefficient for the exact confidence interval is not exactly $(1 - \alpha)$ but is at least $(1 - \alpha)$. Thus, these confidence limits are conservative. See Agresti (1992) for more information.

PROC FREQ computes exact confidence limits for the odds ratio by using an algorithm based on Thomas (1971). See also Gart (1971). The following two equations are solved iteratively to determine the lower and upper confidence limits, ϕ_1 and ϕ_2 :

$$\sum_{i=n_{11}}^{n_{1\cdot}} \binom{n_{1\cdot}}{i} \binom{n_{2\cdot}}{n_{1\cdot}-i} \phi_1^i / \sum_{i=0}^{n_{1\cdot}} \binom{n_{1\cdot}}{i} \binom{n_{2\cdot}}{n_{1\cdot}-i} \phi_1^i = \alpha/2$$

$$\sum_{i=0}^{n_{11}} \binom{n_{1\cdot}}{i} \binom{n_{2\cdot}}{n_{1\cdot}-i} \phi_2^i / \sum_{i=0}^{n_{1\cdot}} \binom{n_{1\cdot}}{i} \binom{n_{2\cdot}}{n_{1\cdot}-i} \phi_2^i = \alpha/2$$

When the odds ratio equals zero, which occurs when either $n_{11} = 0$ or $n_{22} = 0$, PROC FREQ sets the lower exact confidence limit to zero and determines the upper limit with level α . Similarly, when the odds ratio equals infinity, which occurs when either $n_{12} = 0$ or $n_{21} = 0$, PROC FREQ sets the upper exact confidence limit to infinity and determines the lower limit with level α .

Relative Risks (Cohort Studies)

These measures of relative risk are useful in *cohort* (prospective) study designs, where two samples are identified based on the presence or absence of an explanatory factor. The two samples are observed in future time for the binary (yes-no) response variable under study. Relative risk measures are also useful in cross-sectional studies, where two variables are observed simultaneously. See Stokes, Davis, and Koch (2012) and Agresti (2007) for more information.

The column 1 relative risk is the ratio of the column 1 risk for row 1 to row 2. The column 1 risk for row 1 is the proportion of the row 1 observations classified in column 1,

$$p_1 = n_{11} / n_{1\cdot}$$

Similarly, the column 1 risk for row 2 is

$$p_2 = n_{21} / n_{2\cdot}$$

The column 1 relative risk is computed as

$$RR_1 = p_1 / p_2$$

A relative risk greater than 1 indicates that the probability of positive response is greater in row 1 than in row 2. Similarly, a relative risk less than 1 indicates that the probability of positive response is less in row 1 than in row 2. The strength of association increases with the deviation from 1.

Asymptotic $100(1 - \alpha)\%$ confidence limits for the column 1 relative risk are computed as

$$(RR_1 \times \exp(-z\sqrt{v}), RR_1 \times \exp(z\sqrt{v}))$$

where

$$v = \text{Var}(\ln RR_1) = ((1 - p_1)/n_{11}) + ((1 - p_2)/n_{21})$$

and z is the $100(1 - \alpha/2)$ percentile of the standard normal distribution. If either n_{11} or n_{21} is zero, the estimates are not computed.

PROC FREQ computes the column 2 relative risks in the same way.

Exact Unconditional Confidence Limits for the Relative Risk If you specify the RELRISK option in the EXACT statement, PROC FREQ provides exact unconditional confidence limits for the relative risk. PROC FREQ computes the confidence limits by inverting two separate one-sided tests (tail method), where the size of each test is at most $\alpha/2$ and the confidence coefficient is at least $(1 - \alpha)$. Exact conditional methods, described in the section “Exact Statistics” on page 2601, do not apply to the relative risk due to the presence of a nuisance parameter (Agresti 1992). The unconditional approach eliminates the nuisance parameter by maximizing the p -value over all possible values of the parameter (Santner and Snell 1980).

By default, PROC FREQ uses the unstandardized relative risk as the test statistic in the confidence limit computations. If you specify the RELRISK(METHOD=SCORE) option, the procedure uses the relative risk score statistic (Chan and Zhang 1999). The score statistic is a less discrete statistic than the raw relative risk and produces less conservative confidence limits (Agresti and Min 2001). See also Santner et al. (2007) for comparisons of methods for computing exact confidence limits.

See the section “Exact Unconditional Confidence Limits for the Risk Difference” on page 2580 for a description of the method that PROC FREQ uses to compute confidence limits for the relative risk. The test statistic for the relative risk computation is either the unstandardized relative risk (by default) or the relative risk score statistic (if you specify the RELRISK(METHOD=SCORE) option). PROC FREQ uses the following form of the unstandardized relative risk, which adds 0.05 to each frequency, to ensure that the statistic is defined when there are zero table cells (Gart and Nam 1988):

$$rr = \frac{(n_{11} + 0.5) / (n_{1.} + 0.5)}{(n_{21} + 0.5) / (n_{2.} + 0.5)}$$

If you specify the RELRISK(METHOD=SCORE) option, PROC FREQ uses the relative risk score statistic (Miettinen and Nurminen 1985; Farrington and Manning 1990). This test statistic is computed as

$$z = (\hat{p}_1 - R_0 \hat{p}_2) / \text{se}(rr)$$

where

$$se(rr) = \sqrt{\tilde{p}_1(1 - \tilde{p}_1)/n_{1\cdot} + R_0^2 \tilde{p}_2(1 - \tilde{p}_2)/n_{2\cdot}}$$

where $\tilde{p}_1$ and $\tilde{p}_2$ are the maximum likelihood estimators of p_1 and p_2 under the null hypothesis that the relative risk equals R_0 . From Farrington and Manning (1990, p. 1454), the maximum likelihood solution is

$$\tilde{p}_1 = (-b - \sqrt{b^2 - 4ac})/2a \quad \text{and} \quad \tilde{p}_2 = \tilde{p}_1/R_0$$

where

$$\begin{aligned} a &= 1 + \theta \\ b &= -(R_0(1 + \theta \hat{p}_2) + \theta + \hat{p}_1) \\ c &= R_0(\hat{p}_1 + \theta \hat{p}_2) \\ \theta &= n_{2\cdot}/n_{1\cdot}. \end{aligned}$$

Cochran-Armitage Test for Trend

The TREND option in the TABLES statement provides the Cochran-Armitage test for trend, which tests for trend in binomial proportions across levels of a single factor or covariate. This test is appropriate for a two-way table where one variable has two levels and the other variable is ordinal. The two-level variable represents the response, and the other variable represents an explanatory variable with ordered levels. When the two-way has two columns and R rows, PROC FREQ tests for trend across the R levels of the row variable, and the binomial proportion is computed as the proportion of observations in the first column. When the table has two rows and C columns, PROC FREQ tests for trend across the C levels of the column variable, and the binomial proportion is computed as the proportion of observations in the first row.

The trend test is based on the regression coefficient for the weighted linear regression of the binomial proportions on the scores of the explanatory variable levels. See Margolin (1988) and Agresti (2002) for details. If the table has two columns and R rows, the trend test statistic is computed as

$$T = \sum_{i=1}^R n_{i1}(R_i - \bar{R}) / \sqrt{p_{\cdot 1}(1 - p_{\cdot 1}) s^2}$$

where R_i is the score of row i , $\bar{R}$ is the average row score, and

$$s^2 = \sum_{i=1}^R n_{i\cdot}(R_i - \bar{R})^2$$

The SCORES= option in the TABLES statement determines the type of row scores used in computing the trend test (and other score-based statistics). The default is SCORES=TABLE. See the section “[Scores](#)” on page 2549 for details. For character variables, the table scores for the row variable are the row numbers (for example, 1 for the first row, 2 for the second row, and so on). For numeric variables, the table score for each row is the numeric value of the row level. When you perform the trend test, the explanatory variable might be numeric (for example, dose of a test substance), and the variable values might be appropriate scores. If

the explanatory variable has ordinal levels that are not numeric, you can assign meaningful scores to the variable levels. Sometimes equidistant scores, such as the table scores for a character variable, might be appropriate. For more information on choosing scores for the trend test, see Margolin (1988).

The null hypothesis for the Cochran-Armitage test is no trend, which means that the binomial proportion $p_{i1} = n_{i1}/n_i$ is the same for all levels of the explanatory variable. Under the null hypothesis, the trend statistic has an asymptotic standard normal distribution.

PROC FREQ computes one-sided and two-sided p -values for the trend test. When the test statistic is greater than its null hypothesis expected value of zero, PROC FREQ displays the right-sided p -value, which is the probability of a larger value of the statistic occurring under the null hypothesis. A small right-sided p -value supports the alternative hypothesis of increasing trend in proportions from row 1 to row R . When the test statistic is less than or equal to zero, PROC FREQ displays the left-sided p -value. A small left-sided p -value supports the alternative of decreasing trend.

The one-sided p -value for the trend test is computed as

$$P_1 = \begin{cases} \text{Prob}(Z > T) & \text{if } T > 0 \\ \text{Prob}(Z < T) & \text{if } T \leq 0 \end{cases}$$

where Z has a standard normal distribution. The two-sided p -value is computed as

$$P_2 = \text{Prob}(|Z| > |T|)$$

PROC FREQ also provides exact p -values for the Cochran-Armitage trend test. You can request the exact test by specifying the TREND option in the EXACT statement. See the section “Exact Statistics” on page 2601 for more information.

Jonckheere-Terpstra Test

The JT option in the TABLES statement provides the Jonckheere-Terpstra test, which is a nonparametric test for ordered differences among classes. It tests the null hypothesis that the distribution of the response variable does not differ among classes. It is designed to detect alternatives of ordered class differences, which can be expressed as $\tau_1 \leq \tau_2 \leq \dots \leq \tau_R$ (or $\tau_1 \geq \tau_2 \geq \dots \geq \tau_R$), with at least one of the inequalities being strict, where τ_i denotes the effect of class i . For such ordered alternatives, the Jonckheere-Terpstra test can be preferable to tests of more general class difference alternatives, such as the Kruskal–Wallis test (produced by the WILCOXON option in the NPAR1WAY procedure). See Pirie (1983) and Hollander and Wolfe (1999) for more information about the Jonckheere-Terpstra test.

The Jonckheere-Terpstra test is appropriate for a two-way table in which an ordinal column variable represents the response. The row variable, which can be nominal or ordinal, represents the classification variable. The levels of the row variable should be ordered according to the ordering you want the test to detect. The order of variable levels is determined by the ORDER= option in the PROC FREQ statement. The default is ORDER=INTERNAL, which orders by unformatted values. If you specify ORDER=DATA, PROC FREQ orders values according to their order in the input data set. For more information about how to order variable levels, see the ORDER= option.

The Jonckheere-Terpstra test statistic is computed by first forming $R(R-1)/2$ Mann-Whitney counts $M_{i,i'}$, where $i < i'$, for pairs of rows in the contingency table,

$$\begin{aligned} M_{i,i'} &= \{ \text{number of times } X_{i,j} < X_{i',j'}, \quad j = 1, \dots, n_i; \quad j' = 1, \dots, n_{i'} \} \\ &+ \frac{1}{2} \{ \text{number of times } X_{i,j} = X_{i',j'}, \quad j = 1, \dots, n_i; \quad j' = 1, \dots, n_{i'} \} \end{aligned}$$

where $X_{i,j}$ is response j in row i . The Jonckheere-Terpstra test statistic is computed as

$$J = \sum_{1 \leq i < i' \leq R} \sum M_{i,i'}$$

This test rejects the null hypothesis of no difference among classes for large values of J . Asymptotic p -values for the Jonckheere-Terpstra test are obtained by using the normal approximation for the distribution of the standardized test statistic. The standardized test statistic is computed as

$$J^* = (J - E_0(J)) / \sqrt{\text{Var}_0(J)}$$

where $E_0(J)$ and $\text{Var}_0(J)$ are the expected value and variance of the test statistic under the null hypothesis,

$$E_0(J) = \left(n^2 - \sum_i n_{i.}^2 \right) / 4$$

$$\text{Var}_0(J) = A/72 + B/(36n(n-1)(n-2)) + C/(8n(n-1))$$

where

$$A = n(n-1)(2n+5) - \sum_i n_{i.}(n_{i.}-1)(2n_{i.}+5) - \sum_j n_{.j}(n_{.j}-1)(2n_{.j}+5)$$

$$B = \left(\sum_i n_{i.}(n_{i.}-1)(n_{i.}-2) \right) \left(\sum_j n_{.j}(n_{.j}-1)(n_{.j}-2) \right)$$

$$C = \left(\sum_i n_{i.}(n_{i.}-1) \right) \left(\sum_j n_{.j}(n_{.j}-1) \right)$$

PROC FREQ computes one-sided and two-sided p -values for the Jonckheere-Terpstra test. When the standardized test statistic is greater than its null hypothesis expected value of zero, PROC FREQ displays the right-sided p -value, which is the probability of a larger value of the statistic occurring under the null hypothesis. A small right-sided p -value supports the alternative hypothesis of increasing order from row 1 to row R . When the standardized test statistic is less than or equal to zero, PROC FREQ displays the left-sided p -value. A small left-sided p -value supports the alternative of decreasing order from row 1 to row R .

The one-sided p -value for the Jonckheere-Terpstra test, P_1 , is computed as

$$P_1 = \begin{cases} \text{Prob}(Z > J^*) & \text{if } J^* > 0 \\ \text{Prob}(Z < J^*) & \text{if } J^* \leq 0 \end{cases}$$

where Z has a standard normal distribution. The two-sided p -value, P_2 , is computed as

$$P_2 = \text{Prob}(|Z| > |J^*|)$$

PROC FREQ also provides exact p -values for the Jonckheere-Terpstra test. You can request the exact test by specifying the JT option in the EXACT statement. See the section “[Exact Statistics](#)” on page 2601 for more information.

Tests and Measures of Agreement

When you specify the AGREE option in the TABLES statement, PROC FREQ computes tests and measures of agreement for square tables (that is, for tables where the number of rows equals the number of columns). For two-way tables, these tests and measures include McNemar's test for 2×2 tables, Bowker's test of symmetry, the simple kappa coefficient, and the weighted kappa coefficient. For multiple strata (n -way tables, where $n > 2$), PROC FREQ also computes the overall simple kappa coefficient and the overall weighted kappa coefficient, as well as tests for equal kappas (simple and weighted) among strata. Cochran's Q is computed for multiway tables when each variable has two levels, that is, for $h \times 2 \times 2$ tables.

PROC FREQ computes the kappa coefficients (simple and weighted), their asymptotic standard errors, and their confidence limits when you specify the AGREE option in the TABLES statement. If you also specify the KAPPA option in the TEST statement, then PROC FREQ computes the asymptotic test of the hypothesis that simple kappa equals zero. Similarly, if you specify the WTKAP option in the TEST statement, PROC FREQ computes the asymptotic test for weighted kappa.

In addition to the asymptotic tests described in this section, PROC FREQ provides exact p -values for McNemar's test, the simple kappa coefficient test, and the weighted kappa coefficient test. You can request these exact tests by specifying the corresponding options in the EXACT statement. See the section "[Exact Statistics](#)" on page 2601 for more information.

The following sections provide the formulas that PROC FREQ uses to compute the AGREE statistics. For information about the use and interpretation of these statistics, see Agresti (2002, 2007); Fleiss, Levin, and Paik (2003), and the other references cited for each statistic.

McNemar's Test

PROC FREQ computes McNemar's test for 2×2 tables when you specify the AGREE option. McNemar's test is appropriate when you are analyzing data from matched pairs of subjects with a dichotomous (yes-no) response. It tests the null hypothesis of marginal homogeneity, or $p_{1\cdot} = p_{\cdot 1}$. McNemar's test is computed as

$$Q_M = (n_{12} - n_{21})^2 / (n_{12} + n_{21})$$

Under the null hypothesis, Q_M has an asymptotic chi-square distribution with one degree of freedom. See McNemar (1947), as well as the general references cited in the preceding section. In addition to the asymptotic test, PROC FREQ also computes the exact p -value for McNemar's test when you specify the MCNEM option in the EXACT statement.

Bowker's Test of Symmetry

For Bowker's test of symmetry, the null hypothesis is that the cell proportions are symmetric, or that $p_{ij} = p_{ji}$ for all pairs of table cells. For 2×2 tables, Bowker's test is identical to McNemar's test, and so PROC FREQ provides Bowker's test for square tables larger than 2×2 .

Bowker's test of symmetry is computed as

$$Q_B = \sum_{i < j} \sum (n_{ij} - n_{ji})^2 / (n_{ij} + n_{ji})$$

For large samples, Q_B has an asymptotic chi-square distribution with $R(R-1)/2$ degrees of freedom under the null hypothesis of symmetry. See Bowker (1948) for details.

Simple Kappa Coefficient

The simple kappa coefficient, introduced by Cohen (1960), is a measure of interrater agreement. PROC FREQ computes the simple kappa coefficient as

$$\hat{\kappa} = (P_o - P_e) / (1 - P_e)$$

where $P_o = \sum_i p_{ii}$ and $P_e = \sum_i p_{i\cdot} p_{\cdot i}$. If the two response variables are viewed as two independent ratings of the n subjects, the kappa coefficient equals +1 when there is complete agreement of the raters. When the observed agreement exceeds chance agreement, kappa is positive, with its magnitude reflecting the strength of agreement. Although this is unusual in practice, kappa is negative when the observed agreement is less than chance agreement. The minimum value of kappa is between -1 and 0, depending on the marginal proportions.

The asymptotic variance of the simple kappa coefficient is computed as

$$\text{Var}(\hat{\kappa}) = (A + B - C) / (1 - P_e)^2 n$$

where

$$A = \sum_i p_{ii} (1 - (p_{i\cdot} + p_{\cdot i})(1 - \hat{\kappa}))^2$$

$$B = (1 - \hat{\kappa})^2 \sum_{i \neq j} p_{ij} (p_{\cdot i} + p_{\cdot j})^2$$

$$C = (\hat{\kappa} - P_e(1 - \hat{\kappa}))^2$$

See Fleiss, Cohen, and Everitt (1969) for details.

PROC FREQ computes confidence limits for the simple kappa coefficient as

$$\hat{\kappa} \pm (z_{\alpha/2} \times \sqrt{\text{Var}(\hat{\kappa})})$$

where $z_{\alpha/2}$ is the 100(1 - $\alpha/2$) percentile of the standard normal distribution. The value of α is determined by the value of the ALPHA= option, which, by default, equals 0.05 and produces 95% confidence limits.

To compute an asymptotic test for the kappa coefficient, PROC FREQ uses the standardized test statistic $\hat{\kappa}^*$, which has an asymptotic standard normal distribution under the null hypothesis that kappa equals zero. The standardized test statistic is computed as

$$\hat{\kappa}^* = \hat{\kappa} / \sqrt{\text{Var}_0(\hat{\kappa})}$$

where $\text{Var}_0(\hat{\kappa})$ is the variance of the kappa coefficient under the null hypothesis,

$$\text{Var}_0(\hat{\kappa}) = \left(P_e + P_e^2 - \sum_i p_{i\cdot} p_{\cdot i} (p_{i\cdot} + p_{\cdot i}) \right) / (1 - P_e)^2 n$$

See Fleiss, Levin, and Paik (2003) for details.

PROC FREQ also provides an exact test for the simple kappa coefficient. You can request the exact test by specifying the KAPPA or AGREE option in the EXACT statement. See the section “Exact Statistics” on page 2601 for more information.

Weighted Kappa Coefficient

The weighted kappa coefficient is a generalization of the simple kappa coefficient that uses weights to quantify the relative difference between categories. For 2×2 tables, the weighted kappa coefficient equals the simple kappa coefficient. PROC FREQ displays the weighted kappa coefficient only for tables larger than 2×2 . PROC FREQ computes the kappa weights from the column scores, by using either Cicchetti-Allison weights or Fleiss-Cohen weights, both of which are described in the following section. The weights w_{ij} are constructed so that $0 \leq w_{ij} < 1$ for all $i \neq j$, $w_{ii} = 1$ for all i , and $w_{ij} = w_{ji}$. The weighted kappa coefficient is computed as

$$\hat{\kappa}_w = (P_{o(w)} - P_{e(w)}) / (1 - P_{e(w)})$$

where

$$P_{o(w)} = \sum_i \sum_j w_{ij} p_{ij}$$

$$P_{e(w)} = \sum_i \sum_j w_{ij} p_{i \cdot} p_{\cdot j}$$

The asymptotic variance of the weighted kappa coefficient is

$$\text{Var}(\hat{\kappa}_w) = \left(\sum_i \sum_j p_{ij} (w_{ij} - (\bar{w}_{i \cdot} + \bar{w}_{\cdot j})(1 - \hat{\kappa}_w))^2 - (\hat{\kappa}_w - P_{e(w)}(1 - \hat{\kappa}_w))^2 \right) / (1 - P_{e(w)})^2 n$$

where

$$\bar{w}_{i \cdot} = \sum_j p_{\cdot j} w_{ij}$$

$$\bar{w}_{\cdot j} = \sum_i p_{i \cdot} w_{ij}$$

See Fleiss, Cohen, and Everitt (1969) for details.

PROC FREQ computes confidence limits for the weighted kappa coefficient as

$$\hat{\kappa}_w \pm (z_{\alpha/2} \times \sqrt{\text{Var}(\hat{\kappa}_w)})$$

where $z_{\alpha/2}$ is the $100(1 - \alpha/2)$ percentile of the standard normal distribution. The value of α is determined by the value of the ALPHA= option, which, by default, equals 0.05 and produces 95% confidence limits.

To compute an asymptotic test for the weighted kappa coefficient, PROC FREQ uses the standardized test statistic $\hat{\kappa}_w^*$, which has an asymptotic standard normal distribution under the null hypothesis that weighted kappa equals zero. The standardized test statistic is computed as

$$\hat{\kappa}_w^* = \hat{\kappa}_w / \sqrt{\text{Var}_0(\hat{\kappa}_w)}$$

where $\text{Var}_0(\hat{\kappa}_w)$ is the variance of the weighted kappa coefficient under the null hypothesis,

$$\text{Var}_0(\hat{\kappa}_w) = \left(\sum_i \sum_j p_{i \cdot} p_{\cdot j} (w_{ij} - (\bar{w}_{i \cdot} + \bar{w}_{\cdot j}))^2 - P_{e(w)}^2 \right) / (1 - P_{e(w)})^2 n$$

See Fleiss, Levin, and Paik (2003) for details.

PROC FREQ also provides an exact test for the weighted kappa coefficient. You can request the exact test by specifying the WTKAPPA or AGREE option in the EXACT statement. See the section “[Exact Statistics](#)” on page 2601 for more information.

Weights PROC FREQ computes kappa coefficient weights by using the column scores and one of the two available weight types. The column scores are determined by the SCORES= option in the TABLES statement. The two available types of kappa weights are Cicchetti-Allison and Fleiss-Cohen weights. By default, PROC FREQ uses Cicchetti-Allison weights. If you specify (WT=FC) with the AGREE option, then PROC FREQ uses Fleiss-Cohen weights to compute the weighted kappa coefficient.

PROC FREQ computes Cicchetti-Allison kappa coefficient weights as

$$w_{ij} = 1 - \frac{|C_i - C_j|}{C_C - C_1}$$

where C_i is the score for column i and C is the number of categories or columns. See Cicchetti and Allison (1971) for details.

The SCORES= option in the TABLES statement determines the type of column scores used to compute the kappa weights (and other score-based statistics). The default is SCORES=TABLE. See the section “Scores” on page 2549 for details. For numeric variables, table scores are the values of the variable levels. You can assign numeric values to the levels in a way that reflects their level of similarity. For example, suppose you have four levels and order them according to similarity. If you assign them values of 0, 2, 4, and 10, the Cicchetti-Allison kappa weights take the following values: $w_{12} = 0.8$, $w_{13} = 0.6$, $w_{14} = 0$, $w_{23} = 0.8$, $w_{24} = 0.2$, and $w_{34} = 0.4$. Note that when there are only two categories (that is, $C = 2$), the weighted kappa coefficient is identical to the simple kappa coefficient.

If you specify (WT=FC) with the AGREE option in the TABLES statement, PROC FREQ computes Fleiss-Cohen kappa coefficient weights as

$$w_{ij} = 1 - \frac{(C_i - C_j)^2}{(C_C - C_1)^2}$$

See Fleiss and Cohen (1973) for details.

For the preceding example, the Fleiss-Cohen kappa weights are: $w_{12} = 0.96$, $w_{13} = 0.84$, $w_{14} = 0$, $w_{23} = 0.96$, $w_{24} = 0.36$, and $w_{34} = 0.64$.

Overall Kappa Coefficient

When there are multiple strata, PROC FREQ combines the stratum-level estimates of kappa into an overall estimate of the supposed common value of kappa. Assume there are q strata, indexed by $h = 1, 2, \dots, q$, and let $\text{Var}(\hat{\kappa}_h)$ denote the variance of $\hat{\kappa}_h$. The estimate of the overall kappa coefficient is computed as

$$\hat{\kappa}_T = \sum_{h=1}^q \frac{\hat{\kappa}_h}{\text{Var}(\hat{\kappa}_h)} / \sum_{h=1}^q \frac{1}{\text{Var}(\hat{\kappa}_h)}$$

See Fleiss, Levin, and Paik (2003) for details.

PROC FREQ computes an estimate of the overall weighted kappa in the same way.

Tests for Equal Kappa Coefficients

When there are multiple strata, the following chi-square statistic tests whether the stratum-level values of kappa are equal:

$$Q_K = \sum_{h=1}^q (\hat{\kappa}_h - \hat{\kappa}_T)^2 / \text{Var}(\hat{\kappa}_h)$$

Under the null hypothesis of equal kappas for the q strata, Q_K has an asymptotic chi-square distribution with $q-1$ degrees of freedom. See Fleiss, Levin, and Paik (2003) for more information. PROC FREQ computes a test for equal weighted kappa coefficients in the same way.

Cochran's Q Test

Cochran's Q is computed for multiway tables when each variable has two levels, that is, for $2 \times 2 \cdots \times 2$ tables. Cochran's Q statistic is used to test the homogeneity of the one-dimensional margins. Let m denote the number of variables and N denote the total number of subjects. Cochran's Q statistic is computed as

$$Q_C = m(m-1) \left(\sum_{j=1}^m T_j^2 - T^2 \right) / \left(mT - \sum_{k=1}^N S_k^2 \right)$$

where T_j is the number of positive responses for variable j , T is the total number of positive responses over all variables, and S_k is the number of positive responses for subject k . Under the null hypothesis, Cochran's Q has an asymptotic chi-square distribution with $m-1$ degrees of freedom. See Cochran (1950) for details. When there are only two binary response variables ($m=2$), Cochran's Q simplifies to McNemar's test. When there are more than two response categories, you can test for marginal homogeneity by using the repeated measures capabilities of the CATMOD procedure.

Tables with Zero Rows and Columns

The AGREE statistics are defined only for square tables, where the number of rows equals the number of columns. If the table is not square, PROC FREQ does not compute AGREE statistics. In the kappa statistic framework, where two independent raters assign ratings to each of n subjects, suppose one of the raters does not use all possible r rating levels. If the corresponding table has r rows but only $r-1$ columns, then the table is not square and PROC FREQ does not compute AGREE statistics. To create a square table in this situation, use the ZEROS option in the WEIGHT statement, which requests that PROC FREQ include observations with zero weights in the analysis. Include zero-weight observations in the input data set to represent any rating levels that are not used by a rater, so that the input data set has at least one observation for each possible rater and rating combination. The analysis then includes all rating levels, even when all levels are not actually assigned by both raters. The resulting table (of rater 1 by rater 2) is a square table, and AGREE statistics can be computed.

For more information, see the description of the ZEROS option. By default, PROC FREQ does not process observations that have zero weights, because these observations do not contribute to the total frequency count, and because any resulting zero-weight row or column causes many of the tests and measures of association to be undefined. However, kappa statistics are defined for tables with a zero-weight row or column, and the ZEROS option makes it possible to input zero-weight observations and construct the tables needed to compute kappas.

Cochran-Mantel-Haenszel Statistics

The CMH option in the TABLES statement gives a stratified statistical analysis of the relationship between the row and column variables after controlling for the strata variables in a multiway table. For example, for the table request $A*B*C*D$, the CMH option provides an analysis of the relationship between C and D , after controlling for A and B . The stratified analysis provides a way to adjust for the possible confounding effects of A and B without being forced to estimate parameters for them.

The CMH analysis produces Cochran-Mantel-Haenszel statistics, which include the correlation statistic, the ANOVA (row mean scores) statistic, and the general association statistic. For 2×2 tables, the CMH option

also provides Mantel-Haenszel and logit estimates of the common odds ratio and the common relative risks, as well as the Breslow-Day test for homogeneity of the odds ratios.

Exact statistics are also available for stratified 2×2 tables. If you specify the EQOR option in the EXACT statement, PROC FREQ provides Zelen's exact test for equal odds ratios. If you specify the COMOR option in the EXACT statement, PROC FREQ provides exact confidence limits for the common odds ratio and an exact test that the common odds ratio equals one.

Let the number of strata be denoted by q , indexing the strata by $h = 1, 2, \dots, q$. Each stratum contains a contingency table with X representing the row variable and Y representing the column variable. For table h , denote the cell frequency in row i and column j by n_{hij} , with corresponding row and column marginal totals denoted by $n_{hi\cdot}$ and $n_{h\cdot j}$, and the overall stratum total by n_h .

Because the formulas for the Cochran-Mantel-Haenszel statistics are more easily defined in terms of matrices, the following notation is used. Vectors are presumed to be column vectors unless they are transposed ($'$).

$$\begin{aligned} \mathbf{n}'_{hi} &= (n_{hi1}, n_{hi2}, \dots, n_{hiC}) & (1 \times C) \\ \mathbf{n}'_h &= (\mathbf{n}'_{h1}, \mathbf{n}'_{h2}, \dots, \mathbf{n}'_{hR}) & (1 \times RC) \\ p_{hi\cdot} &= n_{hi\cdot} / n_h & (1 \times 1) \\ p_{h\cdot j} &= n_{h\cdot j} / n_h & (1 \times 1) \\ \mathbf{P}'_{h*} &= (p_{h1\cdot}, p_{h2\cdot}, \dots, p_{hR\cdot}) & (1 \times R) \\ \mathbf{P}'_{h*} &= (p_{h\cdot 1}, p_{h\cdot 2}, \dots, p_{h\cdot C}) & (1 \times C) \end{aligned}$$

Assume that the strata are independent and that the marginal totals of each stratum are fixed. The null hypothesis, H_0 , is that there is no association between X and Y in any of the strata. The corresponding model is the multiple hypergeometric; this implies that, under H_0 , the expected value and covariance matrix of the frequencies are, respectively,

$$\begin{aligned} \mathbf{m}_h &= \mathbf{E}[\mathbf{n}_h | H_0] = n_h (\mathbf{P}_{h*} \otimes \mathbf{P}_{h*}) \\ \text{Var}[\mathbf{n}_h | H_0] &= c \left((\mathbf{D}_{\mathbf{P}_{h*}} - \mathbf{P}_{h*} \mathbf{P}'_{h*}) \otimes (\mathbf{D}_{\mathbf{P}_{h*}} - \mathbf{P}_{h*} \mathbf{P}'_{h*}) \right) \end{aligned}$$

where

$$c = n_h^2 / (n_h - 1)$$

and where $\otimes$ denotes Kronecker product multiplication and $\mathbf{D}_a$ is a diagonal matrix with the elements of $\mathbf{a}$ on the main diagonal.

The generalized CMH statistic (Landis, Heyman, and Koch 1978) is defined as

$$Q_{CMH} = \mathbf{G}' \mathbf{V}_G^{-1} \mathbf{G}$$

where

$$\begin{aligned} \mathbf{G} &= \sum_h \mathbf{B}_h (\mathbf{n}_h - \mathbf{m}_h) \\ \mathbf{V}_G &= \sum_h \mathbf{B}_h (\text{Var}[\mathbf{n}_h | H_0]) \mathbf{B}'_h \end{aligned}$$

and where

$$\mathbf{B}_h = \mathbf{C}_h \otimes \mathbf{R}_h$$

is a matrix of fixed constants based on column scores $\mathbf{C}_h$ and row scores $\mathbf{R}_h$. When the null hypothesis is true, the CMH statistic has an asymptotic chi-square distribution with degrees of freedom equal to the rank of $\mathbf{B}_h$. If $\mathbf{V}_G$ is found to be singular, PROC FREQ prints a message and sets the value of the CMH statistic to missing.

PROC FREQ computes three CMH statistics by using this formula for the generalized CMH statistic, with different row and column score definitions for each statistic. The CMH statistics that PROC FREQ computes are the correlation statistic, the ANOVA (row mean scores) statistic, and the general association statistic. These statistics test the null hypothesis of no association against different alternative hypotheses. The following sections describe the computation of these CMH statistics.

CAUTION: The CMH statistics have low power for detecting an association in which the patterns of association for some of the strata are in the opposite direction of the patterns displayed by other strata. Thus, a nonsignificant CMH statistic suggests either that there is no association or that no pattern of association has enough strength or consistency to dominate any other pattern.

Correlation Statistic

The correlation statistic, popularized by Mantel and Haenszel, has one degree of freedom and is known as the Mantel-Haenszel statistic (Mantel and Haenszel 1959; Mantel 1963).

The alternative hypothesis for the correlation statistic is that there is a linear association between X and Y in at least one stratum. If either X or Y does not lie on an ordinal (or interval) scale, then this statistic is not meaningful.

To compute the correlation statistic, PROC FREQ uses the formula for the generalized CMH statistic with the row and column scores determined by the SCORES= option in the TABLES statement. See the section “Scores” on page 2549 for more information about the available score types. The matrix of row scores $\mathbf{R}_h$ has dimension $1 \times R$, and the matrix of column scores $\mathbf{C}_h$ has dimension $1 \times C$.

When there is only one stratum, this CMH statistic reduces to $(n - 1)r^2$, where r is the Pearson correlation coefficient between X and Y . When nonparametric (RANK or RIDIT) scores are specified, the statistic reduces to $(n - 1)r_s^2$, where r_s is the Spearman rank correlation coefficient between X and Y . When there is more than one stratum, this CMH statistic becomes a stratum-adjusted correlation statistic.

ANOVA (Row Mean Scores) Statistic

The ANOVA statistic can be used only when the column variable Y lies on an ordinal (or interval) scale so that the mean score of Y is meaningful. For the ANOVA statistic, the mean score is computed for each row of the table, and the alternative hypothesis is that, for at least one stratum, the mean scores of the R rows are unequal. In other words, the statistic is sensitive to location differences among the R distributions of Y .

The matrix of column scores $\mathbf{C}_h$ has dimension $1 \times C$, and the column scores are determined by the SCORES= option.

The matrix of row scores $\mathbf{R}_h$ has dimension $(R - 1) \times R$ and is created internally by PROC FREQ as

$$\mathbf{R}_h = [\mathbf{I}_{R-1}, -\mathbf{J}_{R-1}]$$

where $\mathbf{I}_{R-1}$ is an identity matrix of rank $R - 1$ and $\mathbf{J}_{R-1}$ is an $(R - 1) \times 1$ vector of ones. This matrix has the effect of forming $R - 1$ independent contrasts of the R mean scores.

When there is only one stratum, this CMH statistic is essentially an analysis of variance (ANOVA) statistic in the sense that it is a function of the variance ratio F statistic that would be obtained from a one-way ANOVA on the dependent variable Y . If nonparametric scores are specified in this case, then the ANOVA statistic is a Kruskal-Wallis test.

If there is more than one stratum, then this CMH statistic corresponds to a stratum-adjusted ANOVA or Kruskal-Wallis test. In the special case where there is one subject per row and one subject per column in the contingency table of each stratum, this CMH statistic is identical to Friedman's chi-square. See [Example 38.9](#) for an illustration.

General Association Statistic

The alternative hypothesis for the general association statistic is that, for at least one stratum, there is some kind of association between X and Y . This statistic is always interpretable because it does not require an ordinal scale for either X or Y .

For the general association statistic, the matrix $\mathbf{R}_h$ is the same as the one used for the ANOVA statistic. The matrix $\mathbf{C}_h$ is defined similarly as

$$\mathbf{C}_h = [\mathbf{I}_{C-1}, -\mathbf{J}_{C-1}]$$

PROC FREQ generates both score matrices internally. When there is only one stratum, then the general association CMH statistic reduces to $Q_P(n-1)/n$, where Q_P is the Pearson chi-square statistic. When there is more than one stratum, then the CMH statistic becomes a stratum-adjusted Pearson chi-square statistic. Note that a similar adjustment can be made by summing the Pearson chi-squares across the strata. However, the latter statistic requires a large sample size in each stratum to support the resulting chi-square distribution with $q(R-1)(C-1)$ degrees of freedom. The CMH statistic requires only a large overall sample size because it has only $(R-1)(C-1)$ degrees of freedom.

See Cochran (1954); Mantel and Haenszel (1959); Mantel (1963); Birch (1965); Landis, Heyman, and Koch (1978).

Mantel-Fleiss Criterion

If you specify the CMH(MANTELFLISS) option in the TABLES statement, PROC FREQ computes the Mantel-Fleiss criterion for stratified 2×2 tables. The Mantel-Fleiss criterion can be used to assess the validity of the chi-square approximation for the distribution of the Mantel-Haenszel statistic for 2×2 tables. For details, see Mantel and Fleiss (1980); Mantel and Haenszel (1959); Stokes, Davis, and Koch (2012); Dmitrienko et al. (2005).

The Mantel-Fleiss criterion is computed as

$$MF = \min \left(\left[\sum_h m_{h11} - \sum_h (n_{h11})_L \right], \left[\sum_h (n_{h11})_U - \sum_h m_{h11} \right] \right)$$

where m_{h11} is the expected value of n_{h11} under the hypothesis of no association between the row and column variables in table h , $(n_{h11})_L$ is the minimum possible value of the table cell frequency, and $(n_{h11})_U$ is the maximum possible value,

$$m_{h11} = n_{h1\cdot} \cdot n_{h\cdot 1} / n_h$$

$$(n_{h11})_L = \max(0, n_{h1\cdot} - n_{h\cdot 2})$$

$$(n_{h11})_U = \min(n_{h\cdot 1}, n_{h1\cdot})$$

The Mantel-Fleiss guideline accepts the validity of the Mantel-Haenszel approximation when the value of the criterion is at least 5. When the criterion is less than 5, PROC FREQ displays a warning.

Adjusted Odds Ratio and Relative Risk Estimates

The CMH option provides adjusted odds ratio and relative risk estimates for stratified 2×2 tables. For each of these measures, PROC FREQ computes a Mantel-Haenszel estimate and a logit estimate. These estimates apply to n -way table requests in the TABLES statement, when the row and column variables both have two levels.

For example, for the table request A*B*C*D, if the row and column variables C and D both have two levels, PROC FREQ provides odds ratio and relative risk estimates, adjusting for the confounding variables A and B.

The choice of an appropriate measure depends on the study design. For case-control (retrospective) studies, the odds ratio is appropriate. For cohort (prospective) or cross-sectional studies, the relative risk is appropriate. See the section “Odds Ratio and Relative Risks for 2×2 Tables” on page 2582 for more information on these measures.

Throughout this section, z denotes the $100(1 - \alpha/2)$ percentile of the standard normal distribution.

Odds Ratio, Case-Control Studies PROC FREQ provides Mantel-Haenszel and logit estimates for the common odds ratio for stratified 2×2 tables.

Mantel-Haenszel Estimator The Mantel-Haenszel estimate of the common odds ratio is computed as

$$OR_{MH} = \left(\sum_h n_{h11} n_{h22} / n_h \right) / \left(\sum_h n_{h12} n_{h21} / n_h \right)$$

It is always computed unless the denominator is zero. See Mantel and Haenszel (1959) and Agresti (2002) for details.

To compute confidence limits for the common odds ratio, PROC FREQ uses the Robins, Breslow, and Greenland (1986) variance estimate for $\ln(OR_{MH})$. The $100(1 - \alpha/2)\%$ confidence limits for the common odds ratio are

$$(OR_{MH} \times \exp(-z\hat{\sigma}), OR_{MH} \times \exp(z\hat{\sigma}))$$

where

$$\begin{aligned} \hat{\sigma}^2 &= \widehat{\text{Var}}(\ln(OR_{MH})) \\ &= \frac{\sum_h (n_{h11} + n_{h22})(n_{h11} n_{h22}) / n_h^2}{2 (\sum_h n_{h11} n_{h22} / n_h)^2} \\ &\quad + \frac{\sum_h [(n_{h11} + n_{h22})(n_{h12} n_{h21}) + (n_{h12} + n_{h21})(n_{h11} n_{h22})] / n_h^2}{2 (\sum_h n_{h11} n_{h22} / n_h) (\sum_h n_{h12} n_{h21} / n_h)} \\ &\quad + \frac{\sum_h (n_{h12} + n_{h21})(n_{h12} n_{h21}) / n_h^2}{2 (\sum_h n_{h12} n_{h21} / n_h)^2} \end{aligned}$$

Note that the Mantel-Haenszel odds ratio estimator is less sensitive to small n_h than the logit estimator.

Logit Estimator The adjusted logit estimate of the common odds ratio (Woolf 1955) is computed as

$$OR_L = \exp \left(\sum_h w_h \ln(OR_h) / \sum_h w_h \right)$$

and the corresponding $100(1 - \alpha)\%$ confidence limits are

$$\left(OR_L \times \exp \left(-z / \sqrt{\sum_h w_h} \right), OR_L \times \exp \left(z / \sqrt{\sum_h w_h} \right) \right)$$

where OR_h is the odds ratio for stratum h , and

$$w_h = 1/\text{Var}(\ln(OR_h))$$

If any table cell frequency in a stratum h is zero, PROC FREQ adds 0.5 to each cell of the stratum before computing OR_h and w_h (Haldane 1955) for the logit estimate. The procedure prints a warning when this occurs.

Relative Risks, Cohort Studies PROC FREQ provides Mantel-Haenszel and logit estimates of the common relative risks for stratified 2×2 tables.

Mantel-Haenszel Estimator The Mantel-Haenszel estimate of the common relative risk for column 1 is computed as

$$RR_{MH} = \left(\sum_h n_{h11} n_{h2\cdot} / n_h \right) / \left(\sum_h n_{h21} n_{h1\cdot} / n_h \right)$$

It is always computed unless the denominator is zero. See Mantel and Haenszel (1959) and Agresti (2002) for more information.

To compute confidence limits for the common relative risk, PROC FREQ uses the Greenland and Robins (1985) variance estimate for $\log(RR_{MH})$. The $100(1 - \alpha/2)\%$ confidence limits for the common relative risk are

$$(RR_{MH} \times \exp(-z\hat{\sigma}), RR_{MH} \times \exp(z\hat{\sigma}))$$

where

$$\hat{\sigma}^2 = \widehat{\text{Var}}(\ln(RR_{MH})) = \frac{\sum_h (n_{h1\cdot} n_{h2\cdot} n_{h\cdot 1} - n_{h11} n_{h21} n_h) / n_h^2}{(\sum_h n_{h11} n_{h2\cdot} / n_h) (\sum_h n_{h21} n_{h1\cdot} / n_h)}$$

Logit Estimator The adjusted logit estimate of the common relative risk for column 1 is computed as

$$RR_L = \exp \left(\sum_h w_h \ln(RR_h) / \sum_h w_h \right)$$

and the corresponding $100(1 - \alpha)\%$ confidence limits are

$$\left(RR_L \times \exp \left(-z / \sqrt{\sum_h w_h} \right), RR_L \times \exp \left(z / \sqrt{\sum_h w_h} \right) \right)$$

where RR_h is the column 1 relative risk estimate for stratum h and

$$w_h = 1 / \text{Var}(\ln(RR_h))$$

If n_{h11} or n_{h21} is zero, then PROC FREQ adds 0.5 to each cell of the stratum before computing RR_h and w_h for the logit estimate. The procedure prints a warning when this occurs. See Kleinbaum, Kupper, and Morgenstern (1982, Sections 17.4 and 17.5) for details.

Breslow-Day Test for Homogeneity of the Odds Ratios

When you specify the CMH option, PROC FREQ computes the Breslow-Day test for stratified 2×2 tables. It tests the null hypothesis that the odds ratios for the q strata are equal. When the null hypothesis is true, the statistic has approximately a chi-square distribution with $q-1$ degrees of freedom. See Breslow and Day (1980) and Agresti (2007) for more information.

The Breslow-Day statistic is computed as

$$Q_{BD} = \sum_h (n_{h11} - E(n_{h11} | OR_{MH}))^2 / \text{Var}(n_{h11} | OR_{MH})$$

where E and Var denote expected value and variance, respectively. The summation does not include any table with a zero row or column. If OR_{MH} equals zero or if it is undefined, then PROC FREQ does not compute the statistic and prints a warning message.

For the Breslow-Day test to be valid, the sample size should be relatively large in each stratum, and at least 80% of the expected cell counts should be greater than 5. Note that this is a stricter sample size requirement than the requirement for the Cochran-Mantel-Haenszel test for $q \times 2 \times 2$ tables, in that each stratum sample size (not just the overall sample size) must be relatively large. Even when the Breslow-Day test is valid, it might not be very powerful against certain alternatives, as discussed in Breslow and Day (1980).

If you specify the BDT option, PROC FREQ computes the Breslow-Day test with Tarone's adjustment, which subtracts an adjustment factor from Q_{BD} to make the resulting statistic asymptotically chi-square. The Breslow-Day-Tarone statistic is computed as

$$Q_{BDT} = Q_{BD} - \left(\sum_h (n_{h11} - E(n_{h11} | OR_{MH})) \right)^2 / \sum_h \text{Var}(n_{h11} | OR_{MH})$$

See Tarone (1985); Jones et al. (1989); Breslow (1996) for more information.

Zelen's Exact Test for Equal Odds Ratios

If you specify the EQOR option in the EXACT statement, PROC FREQ computes Zelen's exact test for equal odds ratios for stratified 2×2 tables. Zelen's test is an exact counterpart to the Breslow-Day asymptotic test for equal odds ratios. The reference set for Zelen's test includes all possible $q \times 2 \times 2$ tables with the same row, column, and stratum totals as the observed multiway table and with the same sum of cell (1,1) frequencies as the observed table. The test statistic is the probability of the observed $q \times 2 \times 2$ table conditional on the fixed margins, which is a product of hypergeometric probabilities.

The p -value for Zelen's test is the sum of all table probabilities that are less than or equal to the observed table probability, where the sum is computed over all tables in the reference set determined by the fixed margins and the observed sum of cell (1,1) frequencies. This test is similar to Fisher's exact test for two-way tables. For more information, see Zelen (1971); Hirji (2006); Agresti (1992). PROC FREQ computes Zelen's exact test by using the polynomial multiplication algorithm of Hirji et al. (1996).

Exact Confidence Limits for the Common Odds Ratio

If you specify the COMOR option in the EXACT statement, PROC FREQ computes exact confidence limits for the common odds ratio for stratified 2×2 tables. This computation assumes that the odds ratio is constant over all the 2×2 tables. Exact confidence limits are constructed from the distribution of $S = \sum_h n_{h11}$, conditional on the marginal totals of the 2×2 tables.

Because this is a discrete problem, the confidence coefficient for these exact confidence limits is not exactly $(1 - \alpha)$ but is at least $(1 - \alpha)$. Thus, these confidence limits are conservative. See Agresti (1992) for more information.

PROC FREQ computes exact confidence limits for the common odds ratio by using an algorithm based on Vollset, Hirji, and Elashoff (1991). See also Mehta, Patel, and Gray (1985).

Conditional on the marginal totals of 2×2 table h , let the random variable S_h denote the frequency of table cell (1,1). Given the row totals $n_{h1\cdot}$ and $n_{h2\cdot}$ and column totals $n_{h\cdot 1}$ and $n_{h\cdot 2}$, the lower and upper bounds for S_h are l_h and u_h ,

$$\begin{aligned} l_h &= \max(0, n_{h1\cdot} - n_{h\cdot 2}) \\ u_h &= \min(n_{h1\cdot}, n_{h\cdot 1}) \end{aligned}$$

Let C_{s_h} denote the hypergeometric coefficient,

$$C_{s_h} = \binom{n_{h\cdot 1}}{s_h} \binom{n_{h\cdot 2}}{n_{h1\cdot} - s_h}$$

and let ϕ denote the common odds ratio. Then the conditional distribution of S_h is

$$P(S_h = s_h \mid n_{1\cdot}, n_{\cdot 1}, n_{\cdot 2}) = C_{s_h} \phi^{s_h} / \sum_{x=l_h}^{x=u_h} C_x \phi^x$$

Summing over all the 2×2 tables, $S = \sum_h S_h$, and the lower and upper bounds of S are l and u ,

$$l = \sum_h l_h \quad \text{and} \quad u = \sum_h u_h$$

The conditional distribution of the sum S is

$$P(S = s \mid n_{h1\cdot}, n_{h\cdot 1}, n_{h\cdot 2}; h = 1, \dots, q) = C_s \phi^s / \sum_{x=l}^{x=u} C_x \phi^x$$

where

$$C_s = \sum_{s_1 + \dots + s_q = s} \left(\prod_h C_{s_h} \right)$$

Let s_0 denote the observed sum of cell (1,1) frequencies over the q tables. The following two equations are solved iteratively for lower and upper confidence limits for the common odds ratio, ϕ_1 and ϕ_2 :

$$\begin{aligned} \sum_{x=s_0}^{x=u} C_x \phi_1^x / \sum_{x=l}^{x=u} C_x \phi_1^x &= \alpha/2 \\ \sum_{x=l}^{x=s_0} C_x \phi_2^x / \sum_{x=l}^{x=u} C_x \phi_2^x &= \alpha/2 \end{aligned}$$

When the observed sum s_0 equals the lower bound l , PROC FREQ sets the lower confidence limit to zero and determines the upper limit with level α . Similarly, when the observed sum s_0 equals the upper bound u , PROC FREQ sets the upper confidence limit to infinity and determines the lower limit with level α .

When you specify the COMOR option in the EXACT statement, PROC FREQ also computes the exact test that the common odds ratio equals one. Setting $\phi = 1$, the conditional distribution of the sum S under the null hypothesis becomes

$$P_0(S = s \mid n_{h1\cdot}, n_{h\cdot 1}, n_{h\cdot 2}; h = 1, \dots, q) = C_s / \sum_{x=l}^{x=u} C_x$$

The point probability for this exact test is the probability of the observed sum s_0 under the null hypothesis, conditional on the marginals of the stratified 2×2 tables, and is denoted by $P_0(s_0)$. The expected value of S under the null hypothesis is

$$E_0(S) = \sum_{x=l}^{x=u} x C_x / \sum_{x=l}^{x=u} C_x$$

The one-sided exact p -value is computed from the conditional distribution as $P_0(S \geq s_0)$ or $P_0(S \leq s_0)$, depending on whether the observed sum s_0 is greater or less than $E_0(S)$,

$$P_1 = P_0(S \geq s_0) = \sum_{x=s_0}^{x=u} C_x / \sum_{x=l}^{x=u} C_x \quad \text{if } s_0 > E_0(S)$$

$$P_1 = P_0(S \leq s_0) = \sum_{x=l}^{x=s_0} C_x / \sum_{x=l}^{x=u} C_x \quad \text{if } s_0 \leq E_0(S)$$

PROC FREQ computes two-sided p -values for this test according to three different definitions. A two-sided p -value is computed as twice the one-sided p -value, setting the result equal to one if it exceeds one,

$$P_2^a = 2 \times P_1$$

Additionally, a two-sided p -value is computed as the sum of all probabilities less than or equal to the point probability of the observed sum s_0 , summing over all possible values of s , $l \leq s \leq u$,

$$P_2^b = \sum_{l \leq s \leq u: P_0(s) \leq P_0(s_0)} P_0(s)$$

Also, a two-sided p -value is computed as the sum of the one-sided p -value and the corresponding area in the opposite tail of the distribution, equidistant from the expected value,

$$P_2^c = P_0(|S - E_0(S)| \geq |s_0 - E_0(S)|)$$

Gail-Simon Test for Qualitative Interactions

The GAILSIMON option in the TABLES statement provides the Gail-Simon test for qualitative interaction for stratified 2×2 tables. For details, see Gail and Simon (1985); Silvapulle (2001); Dmitrienko et al. (2005).

The Gail-Simon test is based on the risk differences in stratified 2×2 tables, where the risk difference is defined as the row 1 risk (proportion in column 1) minus the row 2 risk. See the section “[Risks and Risk Differences](#)” on page 2572 for details. By default, the procedure uses column 1 risks to compute the Gail-Simon test. If you specify the GAILSIMON(COLUMN=2) option, the procedure uses column 2 risks.

PROC FREQ computes the Gail-Simon test statistics as described in Gail and Simon (1985),

$$Q- = \sum_h (d_h/s_h)^2 I(d_h > 0)$$

$$Q+ = \sum_h (d_h/s_h)^2 I(d_h < 0)$$

$$Q = \min(Q-, Q+)$$

where d_h is the risk difference in table h , s_h is the standard error of the risk difference, and $I(d_h > 0)$ equals 1 if $d_h > 0$ and 0 otherwise. Similarly, $I(d_h < 0)$ equals 1 if $d_h < 0$ and 0 otherwise. The q 2×2 tables (strata) are indexed by $h = 1, 2, \dots, q$.

The p -values for the Gail-Simon statistics are computed as

$$p(Q-) = \sum_h (1 - F_h(Q-)) B(h; n = q, p = 0.5)$$

$$p(Q+) = \sum_h (1 - F_h(Q+)) B(h; n = q, p = 0.5)$$

$$p(Q) = \sum_{h=1}^{q-1} (1 - F_h(Q)) B(h; n = (q - 1), p = 0.5)$$

where $F_h(\cdot)$ is the cumulative chi-square distribution function with h degrees of freedom and $B(h; n, p)$ is the binomial probability function with parameters n and p . The statistic Q tests the null hypothesis of no qualitative interaction. The statistic $Q-$ tests the null hypothesis of positive risk differences. A small p -value for $Q-$ indicates negative differences; similarly, a small p -value for $Q+$ indicates positive risk differences.

Exact Statistics

Exact statistics can be useful in situations where the asymptotic assumptions are not met, and so the asymptotic p -values are not close approximations for the true p -values. Standard asymptotic methods involve the assumption that the test statistic follows a particular distribution when the sample size is sufficiently large. When the sample size is not large, asymptotic results might not be valid, with the asymptotic p -values

differing perhaps substantially from the exact p -values. Asymptotic results might also be unreliable when the distribution of the data is sparse, skewed, or heavily tied. See Agresti (2007) and Bishop, Fienberg, and Holland (1975) for more information. Exact computations are based on the statistical theory of exact conditional inference for contingency tables, reviewed by Agresti (1992).

In addition to computation of exact p -values, PROC FREQ provides the option of estimating exact p -values by Monte Carlo simulation. This can be useful for problems that are so large that exact computations require a great amount of time and memory, but for which asymptotic approximations might not be sufficient.

Exact statistics are available for many PROC FREQ tests. For one-way tables, PROC FREQ provides exact p -values for the binomial proportion tests and the chi-square goodness-of-fit test. Exact (Clopper-Pearson) confidence limits are available for the binomial proportion. For two-way tables, PROC FREQ provides exact p -values for the following tests: Pearson chi-square test, likelihood ratio chi-square test, Mantel-Haenszel chi-square test, Fisher's exact test, Jonckheere-Terpstra test, and Cochran-Armitage test for trend. PROC FREQ also computes exact p -values for tests of the following statistics: Kendall's tau- b , Stuart's tau- c , Somers' $D(C|R)$, Somers' $D(R|C)$, Pearson correlation coefficient, Spearman correlation coefficient, simple kappa coefficient, and weighted kappa coefficient. For 2×2 tables, PROC FREQ provides McNemar's exact test and exact confidence limits for the odds ratio. PROC FREQ also provides exact unconditional confidence limits for the proportion (risk) difference and for the relative risk. For stratified 2×2 tables, PROC FREQ provides Zelen's exact test for equal odds ratios, exact confidence limits for the common odds ratio, and an exact test for the common odds ratio.

The following sections summarize the exact computational algorithms, define the exact p -values that PROC FREQ computes, discuss the computational resource requirements, and describe the Monte Carlo estimation option.

Computational Algorithms

PROC FREQ computes exact p -values for general $R \times C$ tables by using the network algorithm developed by Mehta and Patel (1983). This algorithm provides a substantial advantage over direct enumeration, which can be very time-consuming and feasible only for small problems. See Agresti (1992) for a review of algorithms for computation of exact p -values, and see Mehta, Patel, and Tsiatis (1984) and Mehta, Patel, and Senchaudhuri (1991) for information about the performance of the network algorithm.

The reference set for a given contingency table is the set of all contingency tables with the observed marginal row and column sums. Corresponding to this reference set, the network algorithm forms a directed acyclic network consisting of nodes in a number of stages. A path through the network corresponds to a distinct table in the reference set. The distances between nodes are defined so that the total distance of a path through the network is the corresponding value of the test statistic. At each node, the algorithm computes the shortest and longest path distances for all the paths that pass through that node. For statistics that can be expressed as a linear combination of cell frequencies multiplied by increasing row and column scores, PROC FREQ computes shortest and longest path distances by using the algorithm of Agresti, Mehta, and Patel (1990). For statistics of other forms, PROC FREQ computes an upper bound for the longest path and a lower bound for the shortest path by following the approach of Valz and Thompson (1994).

The longest and shortest path distances or bounds for a node are compared to the value of the test statistic to determine whether all paths through the node contribute to the p -value, none of the paths through the node contribute to the p -value, or neither of these situations occurs. If all paths through the node contribute, the p -value is incremented accordingly, and these paths are eliminated from further analysis. If no paths contribute, these paths are eliminated from the analysis. Otherwise, the algorithm continues, still processing this node and the associated paths. The algorithm finishes when all nodes have been accounted for.

In applying the network algorithm, PROC FREQ uses full numerical precision to represent all statistics, row and column scores, and other quantities involved in the computations. Although it is possible to use rounding to improve the speed and memory requirements of the algorithm, PROC FREQ does not do this because it can result in reduced accuracy of the p -values.

For one-way tables, PROC FREQ computes the exact chi-square goodness-of-fit test by the method of Radlow and Alf (1975). PROC FREQ generates all possible one-way tables with the observed total sample size and number of categories. For each possible table, PROC FREQ compares its chi-square value with the value for the observed table. If the table's chi-square value is greater than or equal to the observed chi-square, PROC FREQ increments the exact p -value by the probability of that table, which is calculated under the null hypothesis by using the multinomial frequency distribution. By default, the null hypothesis states that all categories have equal proportions. If you specify null hypothesis proportions or frequencies by using the TESTP= or TESTF= option in the TABLES statement, then PROC FREQ calculates the exact chi-square test based on that null hypothesis.

Other exact computations are described in sections about the individual statistics. See the section “[Binomial Proportion](#)” on page 2565 for details about how PROC FREQ computes exact confidence limits and tests for the binomial proportion. See the section “[Odds Ratio and Relative Risks for 2 x 2 Tables](#)” on page 2582 for information about computation of exact confidence limits for the odds ratio for 2×2 tables. Also, see the sections “[Exact Unconditional Confidence Limits for the Risk Difference](#)” on page 2580, “[Exact Confidence Limits for the Common Odds Ratio](#)” on page 2599, and “[Zelen's Exact Test for Equal Odds Ratios](#)” on page 2598.

Definition of p -Values

For several tests in PROC FREQ, the test statistic is nonnegative, and large values of the test statistic indicate a departure from the null hypothesis. Such nondirectional tests include the Pearson chi-square, the likelihood ratio chi-square, the Mantel-Haenszel chi-square, Fisher's exact test for tables larger than 2×2 , McNemar's test, and the one-way chi-square goodness-of-fit test. The exact p -value for a nondirectional test is the sum of probabilities for those tables having a test statistic greater than or equal to the value of the observed test statistic.

There are other tests where it might be appropriate to test against either a one-sided or a two-sided alternative hypothesis. For example, when you test the null hypothesis that the true parameter value equals 0 ($T = 0$), the alternative of interest might be one-sided ($T \leq 0$, or $T \geq 0$) or two-sided ($T \neq 0$). Such tests include the Pearson correlation coefficient, Spearman correlation coefficient, Jonckheere-Terpstra test, Cochran-Armitage test for trend, simple kappa coefficient, and weighted kappa coefficient. For these tests, PROC FREQ displays the right-sided p -value when the observed value of the test statistic is greater than its expected value. The right-sided p -value is the sum of probabilities for those tables for which the test statistic is greater than or equal to the observed test statistic. Otherwise, when the observed test statistic is less than or equal to the expected value, PROC FREQ displays the left-sided p -value. The left-sided p -value is the sum of probabilities for those tables for which the test statistic is less than or equal to the one observed. The one-sided p -value P_1 can be expressed as

$$P_1 = \begin{cases} \text{Prob}(\text{Test Statistic} \geq t) & \text{if } t > E_0(T) \\ \text{Prob}(\text{Test Statistic} \leq t) & \text{if } t \leq E_0(T) \end{cases}$$

where t is the observed value of the test statistic and $E_0(T)$ is the expected value of the test statistic under the null hypothesis. PROC FREQ computes the two-sided p -value as the sum of the one-sided p -value and the corresponding area in the opposite tail of the distribution of the statistic, equidistant from the expected

value. The two-sided p -value P_2 can be expressed as

$$P_2 = \text{Prob} (|\text{Test Statistic} - E_0(T)| \geq |t - E_0(T)|)$$

If you specify the POINT option in the EXACT statement, PROC FREQ also displays exact point probabilities for the test statistics. The exact point probability is the exact probability that the test statistic equals the observed value.

Computational Resources

PROC FREQ uses relatively fast and efficient algorithms for exact computations. These recently developed algorithms, together with improvements in computer power, now make it feasible to perform exact computations for data sets where previously only asymptotic methods could be applied. Nevertheless, there are still large problems that might require a prohibitive amount of time and memory for exact computations, depending on the speed and memory available on your computer. For large problems, consider whether exact methods are really needed or whether asymptotic methods might give results quite close to the exact results, while requiring much less computer time and memory. When asymptotic methods might not be sufficient for such large problems, consider using Monte Carlo estimation of exact p -values, as described in the section “[Monte Carlo Estimation](#)” on page 2604.

A formula does not exist that can predict in advance how much time and memory are needed to compute an exact p -value for a certain problem. The time and memory required depend on several factors, including which test is being performed, the total sample size, the number of rows and columns, and the specific arrangement of the observations into table cells. Generally, larger problems (in terms of total sample size, number of rows, and number of columns) tend to require more time and memory. Additionally, for a fixed total sample size, time and memory requirements tend to increase as the number of rows and columns increases, because this corresponds to an increase in the number of tables in the reference set. Also for a fixed sample size, time and memory requirements increase as the marginal row and column totals become more homogeneous. See Agresti, Mehta, and Patel (1990) and Gail and Mantel (1977) for more information.

At any time while PROC FREQ is computing exact p -values, you can terminate the computations by pressing the system interrupt key sequence (see the *SAS Companion* for your system) and choosing to stop computations. After you terminate exact computations, PROC FREQ completes all other remaining tasks. The procedure produces the requested output and reports missing values for any exact p -values that were not computed by the time of termination.

You can also use the MAXTIME= option in the EXACT statement to limit the amount of time PROC FREQ uses for exact computations. You specify a MAXTIME= value that is the maximum amount of clock time (in seconds) that PROC FREQ can use to compute an exact p -value. If PROC FREQ does not finish computing an exact p -value within that time, it terminates the computation and completes all other remaining tasks.

Monte Carlo Estimation

If you specify the option MC in the EXACT statement, PROC FREQ computes Monte Carlo estimates of the exact p -values instead of directly computing the exact p -values. Monte Carlo estimation can be useful for large problems that require a great amount of time and memory for exact computations but for which asymptotic approximations might not be sufficient. To describe the precision of each Monte Carlo estimate, PROC FREQ provides the asymptotic standard error and $100(1 - \alpha)\%$ confidence limits. The confidence level α is determined by the ALPHA= option in the EXACT statement, which, by default, equals 0.01 and produces 99% confidence limits. The N= n option in the EXACT statement specifies the number of samples that PROC FREQ uses for Monte Carlo estimation; the default is 10000 samples. You can specify a larger value for n to improve the precision of the Monte Carlo estimates. Because larger values of n generate more

samples, the computation time increases. Alternatively, you can specify a smaller value of n to reduce the computation time.

To compute a Monte Carlo estimate of an exact p -value, PROC FREQ generates a random sample of tables with the same total sample size, row totals, and column totals as the observed table. PROC FREQ uses the algorithm of Agresti, Wackerly, and Boyett (1979), which generates tables in proportion to their hypergeometric probabilities conditional on the marginal frequencies. For each sample table, PROC FREQ computes the value of the test statistic and compares it to the value for the observed table. When estimating a right-sided p -value, PROC FREQ counts all sample tables for which the test statistic is greater than or equal to the observed test statistic. Then the p -value estimate equals the number of these tables divided by the total number of tables sampled.

$$\begin{aligned}\hat{P}_{MC} &= M / N \\ M &= \text{number of samples with (Test Statistic} \geq t) \\ N &= \text{total number of samples} \\ t &= \text{observed Test Statistic}\end{aligned}$$

PROC FREQ computes left-sided and two-sided p -value estimates in a similar manner. For left-sided p -values, PROC FREQ evaluates whether the test statistic for each sampled table is less than or equal to the observed test statistic. For two-sided p -values, PROC FREQ examines the sample test statistics according to the expression for P_2 given in the section “[Definition of \$p\$ -Values](#)” on page 2603.

The variable M is a binomially distributed variable with N trials and success probability p . It follows that the asymptotic standard error of the Monte Carlo estimate is

$$se(\hat{P}_{MC}) = \sqrt{\hat{P}_{MC} (1 - \hat{P}_{MC}) / (N - 1)}$$

PROC FREQ constructs asymptotic confidence limits for the p -values according to

$$\hat{P}_{MC} \pm (z_{\alpha/2} \times se(\hat{P}_{MC}))$$

where $z_{\alpha/2}$ is the $100(1 - \alpha/2)$ percentile of the standard normal distribution and the confidence level α is determined by the ALPHA= option in the EXACT statement.

When the Monte Carlo estimate $\hat{P}_{MC}$ equals 0, PROC FREQ computes the confidence limits for the p -value as

$$(0, 1 - \alpha^{(1/N)})$$

When the Monte Carlo estimate $\hat{P}_{MC}$ equals 1, PROC FREQ computes the confidence limits as

$$(\alpha^{(1/N)}, 1)$$

Computational Resources

For each variable in a table request, PROC FREQ stores all of the levels in memory. If all variables are numeric and not formatted, this requires about 84 bytes for each variable level. When there are character variables or formatted numeric variables, the memory that is required depends on the formatted variable lengths, with longer formatted lengths requiring more memory. The number of levels for each variable is limited only by the largest integer that your operating environment can store.

For any single crosstabulation table requested, PROC FREQ builds the entire table in memory, regardless of whether the table has zero cell counts. Thus, if the numeric variables A, B, and C each have 10 levels, PROC FREQ requires 2520 bytes to store the variable levels for the table request A*B*C, as follows:

3 variables * 10 levels/variable * 84 bytes/level

In addition, PROC FREQ requires 8000 bytes to store the table cell frequencies

1000 cells * 8 bytes/cell

even though there might be only 10 observations.

When the variables have many levels or when there are many multiway tables, your computer might not have enough memory to construct the tables. If PROC FREQ runs out of memory while constructing tables, it stops collecting levels for the variable with the most levels and returns the memory that is used by that variable. The procedure then builds the tables that do not contain the disabled variables.

If there is not enough memory for your table request and if increasing the available memory is impractical, you can reduce the number of multiway tables or variable levels. If you are not using the CMH or AGREE option in the TABLES statement to compute statistics across strata, reduce the number of multiway tables by using PROC SORT to sort the data set by one or more of the variables or by using the DATA step to create an index for the variables. Then remove the sorted or indexed variables from the TABLES statement and include a BY statement that uses these variables. You can also reduce memory requirements by using a FORMAT statement in the PROC FREQ step to reduce the number of levels. Additionally, reducing the formatted variable lengths reduces the amount of memory that is needed to store the variable levels. For more information about using formats, see the section “[Grouping with Formats](#)” on page 2544.

Output Data Sets

PROC FREQ produces two types of output data sets that you can use with other statistical and reporting procedures. You can request these data sets as follows:

- Specify the OUT= option in a TABLES statement. This creates an output data set that contains frequency or crosstabulation table counts and percentages
- Specify an OUTPUT statement. This creates an output data set that contains statistics.

PROC FREQ does not display the output data sets. Use PROC PRINT, PROC REPORT, or any other SAS reporting tool to display an output data set.

In addition to these two output data sets, you can create a SAS data set from any piece of PROC FREQ output by using the Output Delivery System. See the section “[ODS Table Names](#)” on page 2617 for more information.

Contents of the TABLES Statement Output Data Set

The OUT= option in the TABLES statement creates an output data set that contains one observation for each combination of variable values (or table cell) in the last table request. By default, each observation contains the frequency and percentage for the table cell. When the input data set contains missing values, the output data set also contains an observation with the frequency of missing values. The output data set includes the following variables:

- BY variables
- table request variables, such as A, B, C, and D in the table request A*B*C*D
- COUNT, which contains the table cell frequency
- PERCENT, which contains the table cell percentage

If you specify the OUTEXPECT option in the TABLES statement for a two-way or multiway table, the output data set also includes expected frequencies. If you specify the OUTPCT option for a two-way or multiway table, the output data set also includes row, column, and table percentages. The additional variables are as follows:

- EXPECTED, which contains the expected frequency
- PCT_TABL, which contains the percentage of two-way table frequency, for n -way tables where $n > 2$
- PCT_ROW, which contains the percentage of row frequency
- PCT_COL, which contains the percentage of column frequency

If you specify the OUTCUM option in the TABLES statement for a one-way table, the output data set also includes cumulative frequencies and cumulative percentages. The additional variables are as follows:

- CUM_FREQ, which contains the cumulative frequency
- CUM_PCT, which contains the cumulative percentage

The OUTCUM option has no effect for two-way or multiway tables.

The following PROC FREQ statements create an output data set of frequencies and percentages:

```
proc freq;
  tables A A*B / out=D;
run;
```

The output data set D contains frequencies and percentages for the table of A by B, which is the last table request listed in the TABLES statement. If A has two levels (1 and 2), B has three levels (1,2, and 3), and no table cell count is zero or missing, then the output data set D includes six observations, one for each combination of A and B levels. The first observation corresponds to A=1 and B=1; the second observation corresponds to A=1 and B=2; and so on. The data set includes the variables COUNT and PERCENT. The value of COUNT is the number of observations with the given combination of A and B levels. The value of PERCENT is the percentage of the total number of observations with that A and B combination.

When PROC FREQ combines different variable values into the same formatted level, the output data set contains the smallest internal value for the formatted level. For example, suppose a variable X has the values 1.1, 1.4, 1.7, 2.1, and 2.3. When you submit the statement

```
format X 1.;
```

in a PROC FREQ step, the formatted levels listed in the frequency table for X are 1 and 2. If you create an output data set with the frequency counts, the internal values of the levels of X are 1.1 and 1.7. To report the internal values of X when you display the output data set, use a format of 3.1 for X.

Contents of the OUTPUT Statement Output Data Set

The **OUTPUT** statement creates a SAS data set that contains statistics computed by PROC FREQ. [Table 38.7](#) lists the statistics that can be stored in the output data set. You identify which statistics to include by specifying *output-options*. See the description of the OUTPUT statement for details.

If you specify multiple **TABLES** statements or multiple table requests in a single TABLES statement, the contents of the output data set correspond to the last table request.

For a one-way table or a two-way table, the output data set contains one observation that stores the requested statistics for the table. For a multiway table, the output data set contains an observation for each two-way table (stratum) of the multiway crosstabulation. If you request summary statistics for the multiway table, the output data set also contains an observation that stores the across-strata summary statistics. If you use a BY statement, the output data set contains an observation (for one-way or two-way tables) or set of observations (for multiway tables) for each BY group.

The OUTPUT data set can include the following variables:

- BY variables
- Variables that identify the stratum for multiway tables, such as A and B in the table request A*B*C*D
- Variables that contain the specified statistics

In addition to the specified estimate or test statistic, the output data set includes associated values such as standard errors, confidence limits, *p*-values, and degrees of freedom.

PROC FREQ constructs variable names for the statistics in the output data set by enclosing the *output-option* names in underscores. Variable names for the corresponding standard errors, confidence limits, *p*-values, and degrees of freedom are formed by combining the *output-option* names with prefixes that identify the associated values. [Table 38.20](#) lists the prefixes and their descriptions.

Table 38.20 Output Data Set Variable Name Prefixes

Prefix	Description
E_	Asymptotic standard error (ASE)
L_	Lower confidence limit
U_	Upper confidence limit
E0_	Null hypothesis ASE
Z_	Standardized value
DF_	Degrees of freedom
P_	<i>p</i> -value
P2_	Two-sided <i>p</i> -value
PL_	Left-sided <i>p</i> -value
PR_	Right-sided <i>p</i> -value
XP_	Exact <i>p</i> -value
XP2_	Exact two-sided <i>p</i> -value
XPL_	Exact left-sided <i>p</i> -value
XPR_	Exact right-sided <i>p</i> -value
XPT_	Exact point probability
XL_	Exact lower confidence limit
XU_	Exact upper confidence limit

For example, the **PCHI** *output-option* in the **OUTPUT** statement includes the Pearson chi-square test in the output data set. The variable names for the Pearson chi-square statistic, its degrees of freedom, and the corresponding *p*-value are **_PCHI_**, **DF_PCHI**, and **P_PCHI**, respectively. If the length of the prefix plus the *output-option* name exceeds eight characters, PROC FREQ truncates the variable name to eight characters.

Displayed Output

Number of Variable Levels Table

If you specify the **NLEVELS** option in the PROC FREQ statement, PROC FREQ displays the “Number of Variable Levels” table. This table provides the number of levels for all variables named in the TABLES statements. PROC FREQ determines the variable levels from the formatted variable values. See “[Grouping with Formats](#)” on page 2544 for details. The “Number of Variable Levels” table contains the following information:

- Variable name
- Levels, which is the total number of levels of the variable
- Number of Nonmissing Levels, if there are missing levels for any of the variables
- Number of Missing Levels, if there are missing levels for any of the variables

One-Way Frequency Tables

PROC FREQ displays one-way frequency tables for all one-way table requests in the **TABLES** statements, unless you specify the **NOPRINT** option in the PROC FREQ statement or the **NOPRINT** option in the TABLES statement. For a one-way table showing the frequency distribution of a single variable, PROC FREQ displays the name of the variable and its values. For each variable value or level, PROC FREQ displays the following information:

- Frequency count, which is the number of observations in the level
- Test Frequency count, if you specify the **CHISQ** and **TESTF=** options to request a chi-square goodness-of-fit test for specified frequencies
- Percent, which is the percentage of the total number of observations. (The **NOPERCENT** option suppresses this information.)
- Test Percent, if you specify the **CHISQ** and **TESTP=** options to request a chi-square goodness-of-fit test for specified percents. (The **NOPERCENT** option suppresses this information.)
- Cumulative Frequency count, which is the sum of the frequency counts for that level and all other levels listed above it in the table. The last cumulative frequency is the total number of nonmissing observations. (The **NOCUM** option suppresses this information.)
- Cumulative Percent, which is the percentage of the total number of observations in that level and in all other levels listed above it in the table. (The **NOCUM** or the **NOPERCENT** option suppresses this information.)

The one-way table also displays the Frequency Missing, which is the number of observations with missing values.

Statistics for One-Way Frequency Tables

For one-way tables, two statistical options are available in the **TABLES** statement. The **CHISQ** option provides a chi-square goodness-of-fit test, and the **BINOMIAL** option provides binomial proportion statistics and tests. PROC FREQ displays the following information, unless you specify the **NOPRINT** option in the PROC FREQ statement:

- If you specify the **CHISQ** option for a one-way table, PROC FREQ provides a chi-square goodness-of-fit test, displaying the Chi-Square statistic, the degrees of freedom (DF), and the probability value ($Pr > ChiSq$). If you specify the **CHISQ** option in the **EXACT** statement, PROC FREQ also displays the exact probability value for this test. If you specify the **POINT** option with the **CHISQ** option in the **EXACT** statement, PROC FREQ displays the exact point probability for the test statistic.
- If you specify the **BINOMIAL** option for a one-way table, PROC FREQ displays the estimate of the binomial Proportion, which is the proportion of observations in the first class listed in the one-way table. PROC FREQ also displays the asymptotic standard error (ASE) and the asymptotic (Wald) and exact (Clopper-Pearson) confidence limits by default. For the binomial proportion test, PROC FREQ displays the asymptotic standard error under the null hypothesis (ASE Under H0), the standardized test statistic (Z), and the one-sided and two-sided probability values.

If you specify the BINOMIAL option in the **EXACT** statement, PROC FREQ also displays the exact one-sided and two-sided probability values for this test. If you specify the **POINT** option with the BINOMIAL option in the EXACT statement, PROC FREQ displays the exact point probability for the test.

- If you request additional binomial confidence limits by specifying *binomial-options*, PROC FREQ provides a table that displays the lower and upper confidence limits for each type that you request. In addition to the Wald and exact (Clopper-Pearson) confidence limits, you can request Agresti-Coull, Jeffreys, and Wilson (score) confidence limits for the binomial proportion.
- If you request a binomial noninferiority or superiority test by specifying the NONINF or SUP *binomial-option*, PROC FREQ displays the following information: the binomial Proportion, the test ASE (under H0 or Sample), the test statistic Z, the probability value, the noninferiority or superiority limit, and the test confidence limits. If you specify the BINOMIAL option in the **EXACT** statement, PROC FREQ also provides the exact probability value for the test, and exact test confidence limits.
- If you request a binomial equivalence test by specifying the EQUIV *binomial-option*, PROC FREQ displays the binomial Proportion and the test ASE (under H0 or Sample). PROC FREQ displays two one-sided tests (TOST) for equivalence, which include test statistics (Z) and probability values for the Lower and Upper tests, together with the Overall probability value. PROC FREQ also displays the equivalence limits and the test-based confidence limits. If you specify the BINOMIAL option in the **EXACT** statement, PROC FREQ provides exact probability values for the TOST and exact test-based confidence limits.

Multiway Tables

PROC FREQ displays all multiway table requests in the **TABLES** statements, unless you specify the **NO-PRINT** option in the **PROC FREQ** statement or the **NOPRINT** option in the **TABLES** statement.

For two-way to multiway crosstabulation tables, the values of the last variable in the table request form the table columns. The values of the next-to-last variable form the rows. Each level (or combination of levels) of the other variables forms one stratum.

There are three ways to display multiway tables in PROC FREQ. By default, PROC FREQ displays multiway tables as separate two-way crosstabulation tables for each stratum of the multiway table. Also by default, PROC FREQ displays these two-way crosstabulation tables in table cell format. Alternatively, if you specify the **CROSSLIST** option, PROC FREQ displays the two-way crosstabulation tables in ODS column format. If you specify the **LIST** option, PROC FREQ displays multiway tables in list format, which presents the entire multiway crosstabulation in a single table.

Crosstabulation Tables

By default, PROC FREQ displays two-way crosstabulation tables in table cell format. The row variable values are listed down the side of the table, the column variable values are listed across the top of the table, and each row and column variable level combination forms a table cell.

Each cell of a crosstabulation table can contain the following information:

- Frequency, which is the number of observations in the table cell. (The **NOFREQ** option suppresses this information.)

- Expected frequency under the hypothesis of independence, if you specify the **EXPECTED** option
- Deviation of the cell frequency from the expected value, if you specify the **DEVIATION** option
- Cell Chi-Square, which is the cell's contribution to the total chi-square statistic, if you specify the **CELLCHI2** option
- Tot Pct, which is the cell's percentage of the total multiway table frequency, for n -way tables when $n > 2$, if you specify the **TOTPCT** option
- Percent, which is the cell's percentage of the total (two-way table) frequency. (The **NOPERCENT** option suppresses this information.)
- Row Pct, or the row percentage, which is the cell's percentage of the total frequency for its row. (The **NOROW** option suppresses this information.)
- Col Pct, or column percentage, which is the cell's percentage of the total frequency for its column. (The **NOCOL** option suppresses this information.)
- Cumulative Col%, or cumulative column percentage, if you specify the **CUMCOL** option

The table also displays the Frequency Missing, which is the number of observations with missing values.

CROSSLIST Tables

If you specify the **CROSSLIST** option, PROC FREQ displays two-way crosstabulation tables in ODS column format. The **CROSSLIST** column format is different from the default crosstabulation table cell format, but the **CROSSLIST** table provides the same information (frequencies, percentages, and other statistics) as the default crosstabulation table.

In the **CROSSLIST** table format, the rows of the display correspond to the crosstabulation table cells, and the columns of the display correspond to descriptive statistics such as frequencies and percentages. Each table cell is identified by the values of its **TABLES** row and column variable levels, with all column variable levels listed within each row variable level. The **CROSSLIST** table also provides row totals, column totals, and overall table totals.

For a crosstabulation table in **CROSSLIST** format, PROC FREQ displays the following information:

- the row variable name and values
- the column variable name and values
- Frequency, which is the number of observations in the table cell. (The **NOFREQ** option suppresses this information.)
- Expected cell frequency under the hypothesis of independence, if you specify the **EXPECTED** option
- Deviation of the cell frequency from the expected value, if you specify the **DEVIATION** option
- Cell Chi-Square, which is the cell's contribution to the total chi-square statistic, if you specify the **CELLCHI2** option
- Total Percent, which is the cell's percentage of the total multiway table frequency, for n -way tables when $n > 2$, if you specify the **TOTPCT** option

- Percent, which is the cell's percentage of the total (two-way table) frequency. (The **NOPERCENT** option suppresses this information.)
- Row Percent, which is the cell's percentage of the total frequency for its row. (The **NOROW** option suppresses this information.)
- Column Percent, the cell's percentage of the total frequency for its column. (The **NOCOL** option suppresses this information.)

The table also displays the Frequency Missing, which is the number of observations with missing values.

LIST Tables

If you specify the **LIST** option in the **TABLES** statement, **PROC FREQ** displays multiway tables in a list format rather than as crosstabulation tables. The **LIST** option displays the entire multiway table in one table, instead of displaying a separate two-way table for each stratum. The **LIST** option is not available when you also request statistical options. Unlike the default crosstabulation output, the **LIST** output does not display row percentages, column percentages, and optional information such as expected frequencies and cell chi-squares.

For a multiway table in list format, **PROC FREQ** displays the following information:

- the variable names and values
- Frequency, which is the number of observations in the level (with the indicated variable values)
- Percent, which is the level's percentage of the total number of observations. (The **NOPERCENT** option suppresses this information.)
- Cumulative Frequency, which is the accumulated frequency of the level and all other levels listed above it in the table. The last cumulative frequency in the table is the total number of nonmissing observations. (The **NOCUM** option suppresses this information.)
- Cumulative Percent, which is the accumulated percentage of the level and all other levels listed above it in the table. (The **NOCUM** or the **NOPERCENT** option suppresses this information.)

The table also displays the Frequency Missing, which is the number of observations with missing values.

Statistics for Multiway Tables

PROC FREQ computes statistical tests and measures for crosstabulation tables, depending on which statements and options you specify. You can suppress the display of these results by specifying the **NOPRINT** option in the **PROC FREQ** statement. With any of the following information, **PROC FREQ** also displays the Sample Size and the Frequency Missing.

- If you specify the **SCOROUT** option in the **TABLES** statement, **PROC FREQ** displays the Row Scores and Column Scores that it uses for statistical computations. The Row Scores table displays the row variable values and the Score corresponding to each value. The Column Scores table displays the column variable values and the corresponding Scores. **PROC FREQ** also identifies the score type used to compute the row and column scores. You can specify the score type with the **SCORES=** option in the **TABLES** statement.

- If you specify the **CHISQ** option, PROC FREQ displays the following statistics for each two-way table: Pearson Chi-Square, Likelihood Ratio Chi-Square, Continuity-Adjusted Chi-Square (for 2×2 tables), Mantel-Haenszel Chi-Square, the Phi Coefficient, the Contingency Coefficient, and Cramér's V. For each test statistic, PROC FREQ also displays the degrees of freedom (DF) and the probability value (Prob).
- If you specify the **CHISQ** option for 2×2 tables, PROC FREQ also displays Fisher's exact test. The test output includes the cell (1,1) frequency (F), the exact left-sided and right-sided probability values, the table probability (P), and the exact two-sided probability value.
- If you specify the **FISHER** option in the TABLES statement (or, equivalently, the **FISHER** option in the **EXACT** statement), PROC FREQ displays Fisher's exact test for tables larger than 2×2 . The test output includes the table probability (P) and the probability value. In addition, PROC FREQ displays the CHISQ output listed earlier, even if you do not also specify the CHISQ option.
- If you specify the **PCHI**, **LRCHI**, or **MHCHI** option in the **EXACT** statement, PROC FREQ displays the corresponding exact test: Pearson Chi-Square, Likelihood Ratio Chi-Square, or Mantel-Haenszel Chi-Square, respectively. The test output includes the test statistic, the degrees of freedom (DF), and the asymptotic and exact probability values. If you also specify the **POINT** option in the **EXACT** statement, PROC FREQ displays the point probability for each exact test requested. If you specify the **CHISQ** option in the **EXACT** statement, PROC FREQ displays exact probability values for all three of these chi-square tests.
- If you specify the **MEASURES** option, PROC FREQ displays the following statistics and their asymptotic standard errors (ASE) for each two-way table: Gamma, Kendall's Tau-b, Stuart's Tau-c, Somers' $D(C|R)$, Somers' $D(R|C)$, Pearson Correlation, Spearman Correlation, Lambda Asymmetric ($C|R$), Lambda Asymmetric ($R|C$), Lambda Symmetric, Uncertainty Coefficient ($C|R$), Uncertainty Coefficient ($R|C$), and Uncertainty Coefficient Symmetric. If you specify the **CL** option, PROC FREQ also displays confidence limits for these measures.
- If you specify the **PLCORR** option, PROC FREQ displays the tetrachoric correlation for 2×2 tables or the polychoric correlation for larger tables. In addition, PROC FREQ displays the MEASURES output listed earlier, even if you do not also specify the MEASURES option.
- If you specify the **GAMMA**, **KENTB**, **STUTC**, **SMDCR**, **SMDRC**, **PCORR**, or **SCORR** option in the **TEST** statement, PROC FREQ displays asymptotic tests for Gamma, Kendall's Tau-b, Stuart's Tau-c, Somers' $D(C|R)$, Somers' $D(R|C)$, the Pearson Correlation, or the Spearman Correlation, respectively. If you specify the **MEASURES** option in the **TEST** statement, PROC FREQ displays all these asymptotic tests. The test output includes the statistic, its asymptotic standard error (ASE), Confidence Limits, the ASE under the null hypothesis H_0 , the standardized test statistic (Z), and the one-sided and two-sided probability values.
- If you specify the **KENTB**, **STUTC**, **SMDCR**, **SMDRC**, **PCORR**, or **SCORR** option in the **EXACT** statement, PROC FREQ displays asymptotic and exact tests for the corresponding measure of association: Kendall's Tau-b, Stuart's Tau-c, Somers' $D(C|R)$, Somers' $D(R|C)$, the Pearson Correlation, or the Spearman correlation, respectively. The test output includes the correlation, its asymptotic standard error (ASE), Confidence Limits, the ASE under the null hypothesis H_0 , the standardized test statistic (Z), and the asymptotic and exact one-sided and two-sided probability values. If you also specify the **POINT** option in the **EXACT** statement, PROC FREQ displays the point probability for each exact test requested.

- If you specify the **RISKDIFF** option for 2×2 tables, PROC FREQ displays the Column 1 and Column 2 Risk Estimates. For each column, PROC FREQ displays the Row 1 Risk, Row 2 Risk, Total Risk, and Risk Difference, together with their asymptotic standard errors (ASE) and Asymptotic Confidence Limits. PROC FREQ also displays Exact Confidence Limits for the Row 1 Risk, Row 2 Risk, and Total Risk. If you specify the **RISKDIFF** option in the **EXACT** statement, PROC FREQ provides unconditional Exact Confidence Limits for the Risk Difference.
- If you specify the **RISKDIFF(CL=)** option for 2×2 tables, PROC FREQ displays the Proportion Difference Confidence Limits. For each confidence limit Type that you request (Exact, Farrington-Manning, Hauck-Anderson, Newcombe Score, or Wald), PROC FREQ displays the Lower and Upper Confidence Limits.
- If you request a noninferiority or superiority test for the proportion difference (**RISKDIFF**) by specifying the **NONINF** or **SUP riskdiff-option**, and if you specify **METHOD=HA** (Hauck-Anderson), **METHOD=FM** (Farrington-Manning), or **METHOD=WALD** (Wald), PROC FREQ displays the following information: the Proportion Difference, the test ASE (H0, Sample, Sample H-A, or FM, depending on the method you specify), the test statistic Z, the probability value, the Noninferiority or Superiority Limit, and the test-based Confidence Limits. If you specify **METHOD=NEWCOMBE** (Newcombe score), PROC FREQ displays the Proportion Difference, the Noninferiority or Superiority Limit, and the Newcombe Confidence Limits.
- If you request an equivalence test for the proportion difference (**RISKDIFF**) by specifying the **EQUIV riskdiff-option**, and if you specify **METHOD=HA** (Hauck-Anderson), **METHOD=FM** (Farrington-Manning), or **METHOD=WALD** (Wald), PROC FREQ displays the following information: the Proportion Difference and the test ASE (H0, Sample, Sample H-A, or FM, depending on the method you specify). PROC FREQ displays a two one-sided test (TOST) for equivalence, which includes test statistics (Z) and probability values for the Lower and Upper tests, together with the Overall probability value. PROC FREQ also displays the Equivalence Limits and the test-based Confidence Limits. If you specify **METHOD=NEWCOMBE** (Newcombe), PROC FREQ displays the Proportion Difference, the Equivalence Limits, and the score Confidence Limits.
- If you request an equality test for the proportion difference (**RISKDIFF**) by specifying the **EQUAL riskdiff-option**, PROC FREQ displays the following information: the Proportion Difference and the test ASE (H0 or Sample), the test statistic Z, the One-Sided probability value ($\Pr > Z$ or $\Pr < Z$), and the Two-Sided probability value, $\Pr > |Z|$.
- If you specify the **MEASURES** option or the **RELRISK** option for 2×2 tables, PROC FREQ displays Estimates of the Relative Risk for Case-Control and Cohort studies, together with their Confidence Limits. These measures are also known as the Odds Ratio and the Column 1 and 2 Relative Risks. If you specify the **OR** option in the **EXACT** statement, PROC FREQ also displays Exact Confidence Limits for the Odds Ratio. If you specify the **RELRISK** option in the **EXACT** statement, PROC FREQ displays unconditional Exact Confidence Limits for the Relative Risk.
- If you specify the **TREND** option, PROC FREQ displays the Cochran-Armitage Trend Test for tables that are $2 \times C$ or $R \times 2$. For this test, PROC FREQ gives the Statistic (Z) and the one-sided and two-sided probability values. If you specify the **TREND** option in the **EXACT** statement, PROC FREQ also displays the exact one-sided and two-sided probability values for this test. If you specify the **POINT** option with the **TREND** option in the **EXACT** statement, PROC FREQ displays the exact point probability for the test statistic.

- If you specify the **JT** option, PROC FREQ displays the Jonckheere-Terpstra Test, showing the Statistic (JT), the standardized test statistic (Z), and the one-sided and two-sided probability values. If you specify the JT option in the **EXACT** statement, PROC FREQ also displays the exact one-sided and two-sided probability values for this test. If you specify the **POINT** option with the JT option in the EXACT statement, PROC FREQ displays the exact point probability for the test statistic.
- If you specify the **AGREE** option and the **PRINTKWT** option, PROC FREQ displays the Kappa Coefficient Weights for square tables larger than 2×2 .
- If you specify the **AGREE** option, for two-way tables PROC FREQ displays McNemar's Test and the Simple Kappa Coefficient for 2×2 tables. For square tables larger than 2×2 , PROC FREQ displays Bowker's Test of Symmetry, the Simple Kappa Coefficient, and the Weighted Kappa Coefficient. For McNemar's Test and Bowker's Test of Symmetry, PROC FREQ displays the Statistic (S), the degrees of freedom (DF), and the probability value ($\text{Pr} > S$). If you specify the **MCNEM** option in the **EXACT** statement, PROC FREQ also displays the exact probability value for McNemar's test. If you specify the **POINT** option with the **MCNEM** option in the EXACT statement, PROC FREQ displays the exact point probability for the test statistic. For the simple and weighted kappa coefficients, PROC FREQ displays the kappa values, asymptotic standard errors (ASE), and Confidence Limits.
- If you specify the **KAPPA** or **WTKAP** option in the **TEST** statement, PROC FREQ displays asymptotic tests for the simple kappa coefficient or the weighted kappa coefficient, respectively. If you specify the **AGREE** option in the TEST statement, PROC FREQ displays both these asymptotic tests. The test output includes the kappa coefficient, its asymptotic standard error (ASE), Confidence Limits, the ASE under the null hypothesis H_0 , the standardized test statistic (Z), and the one-sided and two-sided probability values.
- If you specify the **KAPPA** or **WTKAP** option in the **EXACT** statement, PROC FREQ displays asymptotic and exact tests for the simple kappa coefficient or the weighted kappa coefficient, respectively. The test output includes the kappa coefficient, its asymptotic standard error (ASE), Confidence Limits, the ASE under the null hypothesis H_0 , the standardized test statistic (Z), and the asymptotic and exact one-sided and two-sided probability values. If you specify the **POINT** option in the EXACT statement, PROC FREQ displays the point probability for each exact test requested.
- If you specify the **MC** option in the **EXACT** statement, PROC FREQ displays Monte Carlo estimates for all exact p -values requested by *statistic-options* in the EXACT statement. The Monte Carlo output includes the p -value Estimate, its Confidence Limits, the Number of Samples used to compute the Monte Carlo estimate, and the Initial Seed for random number generation.
- If you specify the **AGREE** option, for multiple strata PROC FREQ displays Overall Simple and Weighted Kappa Coefficients, with their asymptotic standard errors (ASE) and Confidence Limits. PROC FREQ also displays Tests for Equal Kappa Coefficients, giving the Chi-Squares, degrees of freedom (DF), and probability values ($\text{Pr} > \text{ChiSq}$) for the Simple Kappa and Weighted Kappa. For multiple strata of 2×2 tables, PROC FREQ displays Cochran's Q , giving the Statistic (Q), the degrees of freedom (DF), and the probability value ($\text{Pr} > Q$).
- If you specify the **CMH** option, PROC FREQ displays Cochran-Mantel-Haenszel Statistics for the following three alternative hypotheses: Nonzero Correlation, Row Mean Scores Differ (ANOVA Statistic), and General Association. For each of these statistics, PROC FREQ gives the degrees of freedom (DF) and the probability value (Prob). If you specify the **MANTELFLEISS** option, PROC FREQ displays the Mantel-Fleiss Criterion for 2×2 tables. For 2×2 tables, PROC FREQ also

displays Estimates of the Common Relative Risk for Case-Control and Cohort studies, together with their confidence limits. These include both Mantel-Haenszel and Logit stratum-adjusted estimates of the common Odds Ratio, Column 1 Relative Risk, and Column 2 Relative Risk. Also for 2×2 tables, PROC FREQ displays the Breslow-Day Test for Homogeneity of the Odds Ratios. For this test, PROC FREQ gives the Chi-Square, the degrees of freedom (DF), and the probability value ($\text{Pr} > \text{ChiSq}$).

- If you specify the **CMH** option in the TABLES statement and also specify the COMOR option in the **EXACT** statement, PROC FREQ displays exact confidence limits for the Common Odds Ratio for multiple strata of 2×2 tables. PROC FREQ also displays the Exact Test of H_0 : Common Odds Ratio = 1. The test output includes the Cell (1,1) Sum (S), Mean of S Under H_0 , One-sided $\text{Pr} \leq S$, and Point $\text{Pr} = S$. PROC FREQ also provides exact two-sided probability values for the test, computed according to the following three methods: $2 * \text{One-sided}$, Sum of probabilities $\leq \text{Point probability}$, and $\text{Pr} \geq |S - \text{Mean}|$.
- If you specify the **CMH** option in the TABLES statement and also specify the EQOR option in the **EXACT** statement, PROC FREQ computes Zelen's exact test for equal odds ratios for $h \times 2 \times 2$ tables. PROC FREQ displays Zelen's test along with the asymptotic Breslow-Day test produced by the CMH option. PROC FREQ displays the test statistic, Zelen's Exact Test (P), and the probability value, Exact $\text{Pr} \leq P$.
- If you specify the **GAILSIMON** option in the TABLES statement for a multiway 2×2 tables, PROC FREQ displays the Gail-Simon test for qualitative interactions. The display include the following statistics and their p -values: Q+ (Positive Risk Differences), Q- (Negative Risk Differences), and Q (Two-Sided).

ODS Table Names

PROC FREQ assigns a name to each table that it creates. You can use these names to refer to tables when you use the Output Delivery System (ODS) to select tables and create output data sets. For more information about ODS, see Chapter 20, "Using the Output Delivery System."

Table 38.21 lists the ODS table names together with their descriptions and the options required to produce the tables. Note that the ALL option in the TABLES statement invokes the CHISQ, MEASURES, and CMH options.

Table 38.21 ODS Tables Produced by PROC FREQ

ODS Table Name	Description	Statement	Option
BarnardsTest	Barnard's exact test	EXACT	BARNARD
BinomialCLs	Binomial confidence limits	TABLES	BINOMIAL(AC J W)
BinomialEquiv	Binomial equivalence analysis	TABLES	BINOMIAL(EQUIV)
BinomialEquivLimits	Binomial equivalence limits	TABLES	BINOMIAL(EQUIV)
BinomialEquivTest	Binomial equivalence test	TABLES	BINOMIAL(EQUIV)
BinomialNoninf	Binomial noninferiority test	TABLES	BINOMIAL(NONINF)
BinomialProp	Binomial proportion	TABLES	BINOMIAL
BinomialPropTest	Binomial proportion test	TABLES	BINOMIAL
BinomialSup	Binomial superiority test	TABLES	BINOMIAL(SUP)

Table 38.21 continued

ODS Table Name	Description	Statement	Option
BreslowDayTest	Breslow-Day test	TABLES	CMH ($h \times 2 \times 2$ table)
CMH	Cochran-Mantel-Haenszel statistics	TABLES	CMH
ChiSq	Chi-square tests	TABLES	CHISQ
CochransQ	Cochran's Q	TABLES	AGREE ($h \times 2 \times 2$ table)
ColScores	Column scores	TABLES	SCOROUT
CommonOdds- RatioCI	Exact confidence limits for the common odds ratio	EXACT	COMOR ($h \times 2 \times 2$ table)
CommonOdds- RatioTest	Common odds ratio exact test	EXACT	COMOR ($h \times 2 \times 2$ table)
CommonRelRisks	Common relative risks	TABLES	CMH ($h \times 2 \times 2$ table)
CrossList	Crosstabulation table in column format	TABLES	CROSSLIST (n -way table, $n > 1$)
CrossTabFreqs	Crosstabulation table	TABLES	(n -way table, $n > 1$)
EqualKappaTest	Test for equal simple kappas	TABLES	AGREE ($h \times 2 \times 2$ table)
EqualKappaTests	Tests for equal kappas	TABLES	AGREE ($h \times r \times r$, $r > 2$)
EqualOddsRatios	Tests for equal odds ratios	EXACT	EQOR ($h \times 2 \times 2$ table)
GailSimon	Gail-Simon test	TABLES	GAILSIMON ($h \times 2 \times 2$ table)
FishersExact	Fisher's exact test	EXACT or TABLES or TABLES	FISHER FISHER or EXACT CHISQ (2×2 table)
FishersExactMC	Monte Carlo estimates for Fisher's exact test	EXACT	FISHER / MC
Gamma	Gamma	TEST	GAMMA
GammaTest	Gamma test	TEST	GAMMA
JTTest	Jonckheere-Terpstra test	TABLES	JT
JTTestMC	Monte Carlo estimates for Jonckheere-Terpstra exact test	EXACT	JT / MC
KappaStatistics	Kappa statistics	TABLES	AGREE, no TEST or EXACT ($r \times r$ table, $r > 2$)
KappaWeights	Kappa weights	TABLES	AGREE and PRINTKWT
List	List format multiway table	TABLES	LIST
LRChiSq	Likelihood ratio chi-square exact test	EXACT	LRCHI
LRChiSqMC	Monte Carlo exact test for likelihood ratio chi-square	EXACT	LRCHI / MC

Table 38.21 *continued*

ODS Table Name	Description	Statement	Option
MantelFleiss	Mantel-Fleiss criterion	TABLES	CMH(MF) ($h \times 2 \times 2$ table)
McNemarsTest	McNemar's test	TABLES	AGREE (2×2 table)
Measures	Measures of association	TABLES	MEASURES
MHChiSq	Mantel-Haenszel chi-square exact test	EXACT	MHCHI
MHChiSqMC	Monte Carlo exact test for Mantel-Haenszel chi-square	EXACT	MHCHI / MC
NLevels	Number of variable levels	PROC	NLEVELS
OddsRatioCL	Exact confidence limits for the odds ratio	EXACT	OR (2×2 table)
OneWayChiSq	One-way chi-square test	TABLES	CHISQ (one-way table)
OneWayChiSqMC	Monte Carlo exact test for one-way chi-square	EXACT	CHISQ / MC (one-way table)
OneWayFreqs	One-way frequencies	PROC or TABLES	(no TABLES stmt) (one-way table)
OneWayLRChiSq	One-way likelihood ratio chi-square test	TABLES	CHISQ(LRCHISQ) (one-way table)
OverallKappa	Overall simple kappa	TABLES	AGREE ($h \times 2 \times 2$ table)
OverallKappas	Overall kappa coefficients	TABLES	AGREE ($h \times r \times r, r > 2$)
PdiffCLs	Proportion difference confidence limits	TABLES	RISKDIFF(CL=) (2×2 table)
PdiffEquiv	Equivalence analysis for the proportion difference	TABLES	RISKDIFF(EQUIV) (2×2 table)
PdiffEquivLimits	Equivalence limits for the proportion difference	TABLES	RISKDIFF(EQUIV) (2×2 table)
PdiffEquivTest	Equivalence test for the proportion difference	TABLES	RISKDIFF(EQUIV) (2×2 table)
PdiffNoninf	Noninferiority test for the proportion difference	TABLES	RISKDIFF(NONINF) (2×2 table)
PdiffSup	Superiority test for the proportion difference	TABLES	RISKDIFF(SUP) (2×2 table)
PdiffTest	Proportion difference test	TABLES	RISKDIFF(EQUAL) (2×2 table)
PearsonChiSq	Pearson chi-square exact test	EXACT	PCHI
PearsonChiSqMC	Monte Carlo exact test for Pearson chi-square	EXACT	PCHI / MC
PearsonCorr	Pearson correlation	TEST or EXACT	PCORR PCORR
PearsonCorrMC	Monte Carlo exact test for Pearson correlation	EXACT	PCORR / MC

Table 38.21 continued

ODS Table Name	Description	Statement	Option
PearsonCorrTest	Pearson correlation test	TEST or EXACT	PCORR
PlCorr	Polychoric correlation	TEST	PLCORR
PlCorrTest	Polychoric correlation test	TEST	PLCORR
RelativeRisks	Relative risk estimates	TABLES	REL RISK or MEASURES (2 × 2 table)
RelRisk1CL	Exact confidence limits for column 1 relative risk	EXACT	REL RISK (2 × 2 table)
RelRisk2CL	Exact confidence limits for column 2 relative risk	EXACT	REL RISK (2 × 2 table)
RiskDiffCol1	Column 1 risk estimates	TABLES	RISKDIFF (2 × 2 table)
RiskDiffCol2	Column 2 risk estimates	TABLES	RISKDIFF (2 × 2 table)
RowScores	Row scores	TABLES	SCOROUT
SimpleKappa	Simple kappa coefficient	TEST or EXACT	KAPPA
SimpleKappaMC	Monte Carlo exact test for simple kappa	EXACT	KAPPA / MC
SimpleKappaTest	Simple kappa test	TEST or EXACT	KAPPA
SomersDCR	Somers' $D(C R)$	TEST or EXACT	SMDCR
SomersDCRMC	Monte Carlo exact test for Somers' $D(C R)$	EXACT	SMDCR / MC
SomersDCRTest	Somers' $D(C R)$ test	TEST or EXACT	SMDCR
SomersDRC	Somers' $D(R C)$	TEST or EXACT	SMDRC
SomersDRCMC	Monte Carlo exact test for Somers' $D(R C)$	EXACT	SMDRC / MC
SomersDRCTest	Somers' $D(R C)$ test	TEST or EXACT	SMDRC
SpearmanCorr	Spearman correlation	TEST or EXACT	SCORR
SpearmanCorrMC	Monte Carlo exact test for Spearman correlation	EXACT	SCORR / MC
SpearmanCorrTest	Spearman correlation test	TEST or EXACT	SCORR
SymmetryTest	Test of symmetry	TABLES	AGREE
TauB	Kendall's tau- b	TEST or EXACT	KENTB
TauBMC	Monte Carlo exact test for Kendall's tau- b	EXACT	KENTB / MC

Table 38.21 *continued*

ODS Table Name	Description	Statement	Option
TauBTest	Kendall's tau- <i>b</i> test	TEST or EXACT	KENTB KENTB
TauC	Stuart's tau- <i>c</i>	TEST or EXACT	STUTC STUTC
TauCMC	Monte Carlo exact test for Stuart's tau- <i>c</i>	EXACT	STUTC / MC
TauCTest	Stuart's tau- <i>c</i> test	TEST or EXACT	STUTC STUTC
TrendTest	Cochran-Armitage trend test	TABLES	TREND
TrendTestMC	Monte Carlo exact test for trend	EXACT	TREND / MC
WeightedKappa	Weighted kappa	TEST or EXACT	WTKAP WTKAP
WeightedKappaMC	Monte Carlo exact test for weighted kappa	EXACT	WTKAP / MC
WeightedKappaTest	Weighted kappa test	TEST or EXACT	WTKAP WTKAP

ODS Graphics

Statistical procedures use ODS Graphics to create graphs as part of their output. ODS Graphics is described in detail in Chapter 21, “[Statistical Graphics Using ODS](#).”

Before you create graphs, ODS Graphics must be enabled (for example, by specifying the ODS GRAPHICS ON statement). For more information about enabling and disabling ODS Graphics, see the section “[Enabling and Disabling ODS Graphics](#)” on page 600 in Chapter 21, “[Statistical Graphics Using ODS](#).”

The overall appearance of graphs is controlled by ODS styles. Styles and other aspects of using ODS Graphics are discussed in the section “[A Primer on ODS Statistical Graphics](#)” on page 599 in Chapter 21, “[Statistical Graphics Using ODS](#).”

When ODS Graphics is enabled, you can request specific plots with the **PLOTS=** option in the **TABLES** statement. To produce a frequency plot or cumulative frequency plot, you must specify the **FREQPLOT** or **CUMFREQPLOT** *plot-request*, respectively, in the **PLOTS=** option. To produce a mosaic plot, you must specify the **MOSAICPLOT** *plot-request* in the **PLOTS=** option. You can also produce frequency, cumulative frequency, and mosaic plots by specifying the **PLOTS=ALL** option. By default, PROC FREQ produces all other plots that are associated with the analyses that you request in the **TABLES** statement. You can suppress the default plots and request specific plots by using the **PLOTS(ONLY)=** option. See the description of the **PLOTS=** option for details.

PROC FREQ assigns a name to each graph that it creates with ODS Graphics. You can use these names to refer to the graphs. [Table 38.22](#) lists the names of the graphs that PROC FREQ generates together with their descriptions, their **PLOTS=** options (*plot-requests*), and the **TABLES** statement options that are required to produce the graphs.

Table 38.22 Graphs Produced by PROC FREQ

ODS Graph Name	Description	PLOTS= Option	TABLES Statement Option
AgreePlot	Agreement plot	AGREEPLOT	AGREE ($r \times r$ table)
CumFreqPlot	Cumulative frequency plot	CUMFREQPLOT	One-way table request
DeviationPlot	Deviation plot	DEVIATIONPLOT	CHISQ (one-way table)
FreqPlot	Frequency plot	FREQPLOT	Any table request
KappaPlot	Kappa plot	KAPPAPLOT	AGREE ($h \times r \times r$ table)
MosaicPlot	Mosaic plot	MOSAICPLOT	Two-way or multiway table request
ORPlot	Odds ratio plot	ODDSRATIOPLOT	MEASURES or RELRISK ($h \times 2 \times 2$ table)
RelRiskPlot	Relative risk plot	RELRIKSPLOT	MEASURES or RELRISK ($h \times 2 \times 2$ table)
RiskDiffPlot	Risk difference plot	RISKDIFFPLOT	RISKDIFF ($h \times 2 \times 2$ table)
WtKappaPlot	Weighted kappa plot	WTKAPPAPLOT	AGREE ($h \times r \times r$ table, $r > 2$)

Examples: FREQ Procedure

Example 38.1: Output Data Set of Frequencies

The eye and hair color of children from two different regions of Europe are recorded in the data set Color. Instead of recording one observation per child, the data are recorded as cell counts, where the variable Count contains the number of children exhibiting each of the 15 eye and hair color combinations. The data set does not include missing combinations.

The following DATA step statements create the SAS data set Color:

```
data Color;
  input Region Eyes $ Hair $ Count @@;
  label Eyes  = 'Eye Color'
        Hair  = 'Hair Color'
        Region= 'Geographic Region';
  datalines;
1 blue fair 23 1 blue red 7 1 blue medium 24
1 blue dark 11 1 green fair 19 1 green red 7
1 green medium 18 1 green dark 14 1 brown fair 34
1 brown red 5 1 brown medium 41 1 brown dark 40
1 brown black 3 2 blue fair 46 2 blue red 21
2 blue medium 44 2 blue dark 40 2 blue black 6
2 green fair 50 2 green red 31 2 green medium 37
2 green dark 23 2 brown fair 56 2 brown red 42
2 brown medium 53 2 brown dark 54 2 brown black 13
;
```

The following PROC FREQ statements read the Color data set and create an output data set that contains the frequencies, percentages, and expected cell frequencies of the two-way table of Eyes by Hair. The TABLES

statement requests three tables: a frequency table for Eyes, a frequency table for Hair, and a crosstabulation table for Eyes by Hair. The OUT= option creates the FreqCount data set, which contains the crosstabulation table frequencies. The OUTEXPECT option outputs the expected table cell frequencies to FreqCount, and the SPARSE option includes zero cell frequencies in the output data set. The WEIGHT statement specifies that the variable Count contains the observation weights. These statements create [Output 38.1.1](#) through [Output 38.1.3](#).

```
proc freq data=Color;
  tables Eyes Hair Eyes*Hair / out=FreqCount outexpect sparse;
  weight Count;
  title 'Eye and Hair Color of European Children';
run;

proc print data=FreqCount noobs;
  title2 'Output Data Set from PROC FREQ';
run;
```

[Output 38.1.1](#) displays the two frequency tables produced by PROC FREQ: one showing the distribution of eye color, and one showing the distribution of hair color. By default, PROC FREQ lists the variables values in alphabetical order. The 'Eyes*Hair' specification produces a crosstabulation table, shown in [Output 38.1.2](#), with eye color defining the table rows and hair color defining the table columns. A zero cell frequency for green eyes and black hair indicates that this eye and hair color combination does not occur in the data.

The output data set FreqCount ([Output 38.1.3](#)) contains frequency counts and percentages for the last table requested in the TABLES statement, Eyes by Hair. Because the SPARSE option is specified, the data set includes the observation with a zero frequency. The variable Expected contains the expected frequencies, as requested by the OUTEXPECT option.

Output 38.1.1 Frequency Tables

Eye and Hair Color of European Children				
The FREQ Procedure				
Eye Color				
Eyes	Frequency	Percent	Cumulative Frequency	Cumulative Percent
blue	222	29.13	222	29.13
brown	341	44.75	563	73.88
green	199	26.12	762	100.00
Hair Color				
Hair	Frequency	Percent	Cumulative Frequency	Cumulative Percent
black	22	2.89	22	2.89
dark	182	23.88	204	26.77
fair	228	29.92	432	56.69
medium	217	28.48	649	85.17
red	113	14.83	762	100.00

Output 38.1.2 Crosstabulation Table

Table of Eyes by Hair						
Eyes(Eye Color)	Hair(Hair Color)					
Frequency						
Percent						
Row Pct						
Col Pct	black	dark	fair	medium	red	Total
blue	6	51	69	68	28	222
	0.79	6.69	9.06	8.92	3.67	29.13
	2.70	22.97	31.08	30.63	12.61	
	27.27	28.02	30.26	31.34	24.78	
brown	16	94	90	94	47	341
	2.10	12.34	11.81	12.34	6.17	44.75
	4.69	27.57	26.39	27.57	13.78	
	72.73	51.65	39.47	43.32	41.59	
green	0	37	69	55	38	199
	0.00	4.86	9.06	7.22	4.99	26.12
	0.00	18.59	34.67	27.64	19.10	
	0.00	20.33	30.26	25.35	33.63	
Total	22	182	228	217	113	762
	2.89	23.88	29.92	28.48	14.83	100.00

Output 38.1.3 Output Data Set of Frequencies

Eye and Hair Color of European Children				
Output Data Set from PROC FREQ				
Eyes	Hair	COUNT	EXPECTED	PERCENT
blue	black	6	6.409	0.7874
blue	dark	51	53.024	6.6929
blue	fair	69	66.425	9.0551
blue	medium	68	63.220	8.9239
blue	red	28	32.921	3.6745
brown	black	16	9.845	2.0997
brown	dark	94	81.446	12.3360
brown	fair	90	102.031	11.8110
brown	medium	94	97.109	12.3360
brown	red	47	50.568	6.1680
green	black	0	5.745	0.0000
green	dark	37	47.530	4.8556
green	fair	69	59.543	9.0551
green	medium	55	56.671	7.2178
green	red	38	29.510	4.9869

Example 38.2: Frequency Dot Plots

This example produces frequency dot plots for the children's eye and hair color data from [Example 38.1](#).

PROC FREQ produces plots by using ODS Graphics to create graphs as part of the procedure output. Frequency plots are available for any frequency or crosstabulation table request. You can display frequency plots as bar charts or dot plots. You can use *plot-options* to specify the orientation (vertical or horizontal), scale, and layout of the plots.


The following PROC FREQ statements request frequency tables and dot plots. The first TABLES statement requests a one-way frequency table of Hair and a crosstabulation table of Eyes by Hair. The PLOTS= option requests frequency plots for the tables, and the TYPE=DOTPLOT *plot-option* specifies dot plots. By default, frequency plots are produced as bar charts. ODS Graphics must be enabled before producing plots.


The second TABLES statement requests a crosstabulation table of Region by Hair and a frequency dot plot for this table. The SCALE=PERCENT *plot-option* plots percentages instead of frequency counts. SCALE=LOG and SCALE=SQRT *plot-options* are also available to plot log frequencies and square roots of frequencies, respectively.

The ORDER=FREQ option in the PROC FREQ statement orders the variable levels by frequency. This order applies to the frequency and crosstabulation table displays and also to the corresponding frequency plots.


```
ods graphics on;
proc freq data=Color order=freq;
  tables Hair Hair*Eyes / plots=freqplot(type=dotplot);
  tables Hair*Region / plots=freqplot(type=dotplot scale=percent);
  weight Count;
  title 'Eye and Hair Color of European Children';
run;
ods graphics off;
```

[Output 38.2.1](#), [Output 38.2.2](#), and [Output 38.2.3](#) display the dot plots produced by PROC FREQ. By default, the orientation of dot plots is horizontal, which places the variable levels on the Y axis. You can specify the ORIENT=VERTICAL *plot-option* to request a vertical orientation. For two-way plots, you can use the TWOWAY= *plot-option* to specify the plot layout. The default layout (shown in [Output 38.2.2](#) and [Output 38.2.3](#)) is GROUPVERTICAL. Two-way layouts STACKED and GROUPTHORIZONTAL are also available.

Output 38.2.1 One-Way Frequency Dot Plot

Output 38.2.2 Two-Way Frequency Dot Plot

Output 38.2.3 Two-Way Percent Dot Plot



Example 38.3: Chi-Square Goodness-of-Fit Tests

This example examines whether the children’s hair color (from [Example 38.1](#)) has a specified multinomial distribution for the two geographical regions. The hypothesized distribution of hair color is 30% fair, 12% red, 30% medium, 25% dark, and 3% black.

In order to test the hypothesis for each region, the data are first sorted by Region. Then the FREQ procedure uses a BY statement to produce a separate table for each BY group (Region). The option ORDER=DATA orders the variable values (hair color) in the frequency table by their order in the input data set. The TABLES statement requests a frequency table for hair color, and the option NOCUM suppresses the display of the cumulative frequencies and percentages.

The CHISQ option requests a chi-square goodness-of-fit test for the frequency table of Hair. The TESTP= option specifies the hypothesized (or test) percentages for the chi-square test; the number of percentages listed equals the number of table levels, and the percentages sum to 100%. The TESTP= percentages are listed in the same order as the corresponding variable levels appear in frequency table.

The PLOTS= option requests a deviation plot, which is associated with the CHISQ option and displays the relative deviations from the test frequencies. The TYPE=DOTPLOT *plot-option* requests a dot plot instead

of the default type, which is a bar chart. ODS Graphics must be enabled before producing plots. These statements produce [Output 38.3.1](#) through [Output 38.3.4](#).

```
proc sort data=Color;
  by Region;
run;

ods graphics on;
proc freq data=Color order=data;
  tables Hair / nocum chisq testp=(30 12 30 25 3)
           plots(only)=deviationplot(type=dotplot);
  weight Count;
  by Region;
  title 'Hair Color of European Children';
run;
ods graphics off;
```


Output 38.3.1 Frequency Table and Chi-Square Test for Region 1

Hair Color of European Children			
----- Geographic Region=1 -----			
The FREQ Procedure			
Hair Color			
Hair	Frequency	Percent	Test Percent
-----	-----	-----	-----
fair	76	30.89	30.00
red	19	7.72	12.00
medium	83	33.74	30.00
dark	65	26.42	25.00
black	3	1.22	3.00
----- Geographic Region=1 -----			
Chi-Square Test for Specified Proportions			
Chi-Square	7.7602		
DF	4		
Pr > ChiSq	0.1008		

[Output 38.3.1](#) shows the frequency table and chi-square test for Region 1. The frequency table lists the variable values (hair color) in the order in which they appear in the data set. The “Test Percent” column lists the hypothesized percentages for the chi-square test. Always check that you have ordered the TESTP= percentages to correctly match the order of the variable levels.

Output 38.3.2 shows the deviation plot for Region 1, which displays the relative deviations from the hypothesized values. The relative deviation for a level is the difference between the observed and hypothesized (test) percentage divided by the test percentage. You can suppress the chi-square p -value that is displayed by default in the deviation plot by specifying the NOSTATS *plot-option*.


Output 38.3.2 Deviation Plot for Region 1



Output 38.3.3 and Output 38.3.4 show the results for Region 2. PROC FREQ computes a chi-square statistic for each region. The chi-square statistic is significant at the 0.05 level for Region 2 ($p=0.0003$) but not for Region 1. This indicates a significant departure from the hypothesized percentages in Region 2.

Output 38.3.3 Frequency Table and Chi-Square Test for Region 2

Hair Color of European Children			
----- Geographic Region=2 -----			
The FREQ Procedure			
Hair Color			
Hair	Frequency	Percent	Test Percent
fair	152	29.46	30.00
red	94	18.22	12.00
medium	134	25.97	30.00
dark	117	22.67	25.00
black	19	3.68	3.00
----- Geographic Region=2 -----			
Chi-Square Test for Specified Proportions			
Chi-Square	21.3824		
DF	4		
Pr > ChiSq	0.0003		

Output 38.3.4 Deviation Plot for Region 2

Example 38.4: Binomial Proportions

In this example, PROC FREQ computes binomial proportions, confidence limits, and tests. The example uses the eye and hair color data from [Example 38.1](#). By default, PROC FREQ computes the binomial proportion as the proportion of observations in the first level of the one-way table. You can designate a different level by using the `LEVEL= binomial-option`.

The following PROC FREQ statements compute the proportion of children with brown eyes (from the data set in [Example 38.1](#)) and test the null hypothesis that the population proportion equals 50%. These statements also compute an equivalence for the proportion of children with fair hair.

The first TABLES statement requests a one-way frequency table for the variable Eyes. The BINOMIAL option requests the binomial proportion, confidence limits, and test. PROC FREQ computes the proportion with Eyes = 'brown', which is the first level displayed in the table. The AC, WILSON, and EXACT *binomial-options* request the following confidence limits types: Agresti-Coull, Wilson (score), and exact (Clopper-Pearson). By default, PROC FREQ provides Wald and exact (Clopper-Pearson) confidence limits for the binomial proportion. The BINOMIAL option also produces an asymptotic Wald test that the proportion equals 0.5. You can specify a different test proportion with the `P= binomial-option`. The ALPHA=0.1 option specifies that $\alpha = 10\%$, which produces 90% confidence limits.

The second TABLES statement requests a one-way frequency table for the variable Hair. The BINOMIAL option requests the proportion for the first level, Hair = 'fair'. The EQUIV *binomial-option* requests an equivalence test for the binomial proportion. The P=.28 option specifies 0.28 as the null hypothesis proportion, and the MARGIN=.1 option specifies 0.1 as the equivalence test margin.

```
proc freq data=Color order=freq;
  tables Eyes / binomial(ac wilson exact) alpha=.1;
  tables Hair / binomial(equiv p=.28 margin=.1);
  weight Count;
  title 'Hair and Eye Color of European Children';
run;
```

Output 38.4.1 displays the results for eye color, and Output 38.4.2 displays the results for hair color.

Output 38.4.1 Binomial Proportion for Eye Color

Hair and Eye Color of European Children				
The FREQ Procedure				
Eye Color				
Eyes	Frequency	Percent	Cumulative Frequency	Cumulative Percent
brown	341	44.75	341	44.75
blue	222	29.13	563	73.88
green	199	26.12	762	100.00
Binomial Proportion				
Eyes = brown				
Proportion		0.4475		
ASE		0.0180		
Type		90% Confidence Limits		
Wilson		0.4181	0.4773	
Agresti-Coull		0.4181	0.4773	
Clopper-Pearson (Exact)		0.4174	0.4779	
Test of H0: Proportion = 0.5				
ASE under H0		0.0181		
Z		-2.8981		
One-sided Pr < Z		0.0019		
Two-sided Pr > Z		0.0038		

The frequency table in Output 38.4.1 displays the values of Eyes in order of descending frequency count. PROC FREQ computes the proportion of children in the first level displayed in the frequency table, Eyes = 'brown'. Output 38.4.1 displays the binomial proportion confidence limits and test. The confidence

limits are 90% confidence limits. If you do not specify the ALPHA= option, PROC FREQ computes 95% confidence limits by default. Because the value of Z is less than zero, PROC FREQ displays the a left-sided p -value (0.0019). This small p -value supports the alternative hypothesis that the true value of the proportion of children with brown eyes is less than 50%.

Output 38.4.2 displays the equivalence test results produced by the second TABLES statement. The null hypothesis proportion is 0.28 and the equivalence margins are -0.1 and 0.1 , which yield equivalence limits of 0.18 and 0.38 . PROC FREQ provides two one-sided tests (TOST) for equivalence. The small p -value indicates rejection of the null hypothesis in favor of the alternative that the proportion is equivalent to the null value.

Output 38.4.2 Binomial Proportion for Hair Color

Hair Color				
Hair	Frequency	Percent	Cumulative Frequency	Cumulative Percent
fair	228	29.92	228	29.92
medium	217	28.48	445	58.40
dark	182	23.88	627	82.28
red	113	14.83	740	97.11
black	22	2.89	762	100.00

Equivalence Analysis				
H0: $P - p_0 \leq \text{Lower Margin}$ or $\geq \text{Upper Margin}$				
Ha: $\text{Lower Margin} < P - p_0 < \text{Upper Margin}$				
$p_0 = 0.28$ Lower Margin = -0.1 Upper Margin = 0.1				
Proportion		ASE (Sample)		
0.2992		0.0166		
Two One-Sided Tests (TOST)				
Test	Z	P-Value		
Lower Margin	7.1865	Pr > Z	<.0001	
Upper Margin	-4.8701	Pr < Z	<.0001	
Overall			<.0001	
Equivalence Limits		90% Confidence Limits		
0.1800	0.3800	0.2719	0.3265	

Example 38.5: Analysis of a 2x2 Contingency Table

This example computes chi-square tests and Fisher's exact test to compare the probability of coronary heart disease for two types of diet. It also estimates the relative risks and computes exact confidence limits for the odds ratio.

The data set `FatComp` contains hypothetical data for a case-control study of high fat diet and the risk of coronary heart disease. The data are recorded as cell counts, where the variable `Count` contains the frequencies for each exposure and response combination. The data set is sorted in descending order by the variables `Exposure` and `Response`, so that the first cell of the 2×2 table contains the frequency of positive exposure and positive response. The `FORMAT` procedure creates formats to identify the type of exposure and response with character values.

```
proc format;
    value ExpFmt 1='High Cholesterol Diet'
                0='Low Cholesterol Diet';
    value RspFmt 1='Yes'
                0='No';
run;

data FatComp;
    input Exposure Response Count;
    label Response='Heart Disease';
    datalines;
0 0 6
0 1 2
1 0 4
1 1 11
;

proc sort data=FatComp;
    by descending Exposure descending Response;
run;
```

In the following `PROC FREQ` statements, `ORDER=DATA` option orders the contingency table values by their order in the input data set. The `TABLES` statement requests a two-way table of `Exposure` by `Response`. The `CHISQ` option produces several chi-square tests, while the `RELRISK` option produces relative risk measures. The `EXACT` statement requests the exact Pearson chi-square test and exact confidence limits for the odds ratio.

```
proc freq data=FatComp order=data;
    format Exposure ExpFmt. Response RspFmt.;
    tables Exposure*Response / chisq relrisk;
    exact pchi or;
    weight Count;
    title 'Case-Control Study of High Fat/Cholesterol Diet';
run;
```

The contingency table in [Output 38.5.1](#) displays the variable values so that the first table cell contains the frequency for the first cell in the data set (the frequency of positive exposure and positive response).

Output 38.5.1 Contingency Table

Case-Control Study of High Fat/Cholesterol Diet				
The FREQ Procedure				
Table of Exposure by Response				
Exposure	Response(Heart Disease)			
Frequency				
Percent				
Row Pct				
Col Pct		Yes	No	Total
-----+-----+-----+				
High Cholesterol		11	4	15
Diet		47.83	17.39	65.22
		73.33	26.67	
		84.62	40.00	
-----+-----+-----+				
Low Cholesterol		2	6	8
Diet		8.70	26.09	34.78
		25.00	75.00	
		15.38	60.00	
-----+-----+-----+				
Total		13	10	23
		56.52	43.48	100.00

[Output 38.5.2](#) displays the chi-square statistics. Because the expected counts in some of the table cells are small, PROC FREQ gives a warning that the asymptotic chi-square tests might not be appropriate. In this case, the exact tests are appropriate. The alternative hypothesis for this analysis states that coronary heart disease is more likely to be associated with a high fat diet, so a one-sided test is desired. Fisher's exact right-sided test analyzes whether the probability of heart disease in the high fat group exceeds the probability of heart disease in the low fat group; because this p -value is small, the alternative hypothesis is supported.

The odds ratio, displayed in [Output 38.5.3](#), provides an estimate of the relative risk when an event is rare. This estimate indicates that the odds of heart disease is 8.25 times higher in the high fat diet group; however, the wide confidence limits indicate that this estimate has low precision.

Output 38.5.2 Chi-Square Statistics

Statistic	DF	Value	Prob
Chi-Square	1	4.9597	0.0259
Likelihood Ratio Chi-Square	1	5.0975	0.0240
Continuity Adj. Chi-Square	1	3.1879	0.0742
Mantel-Haenszel Chi-Square	1	4.7441	0.0294
Phi Coefficient		0.4644	
Contingency Coefficient		0.4212	
Cramer's V		0.4644	

WARNING: 50% of the cells have expected counts less than 5.
(Asymptotic) Chi-Square may not be a valid test.

Pearson Chi-Square Test

Chi-Square	4.9597
DF	1
Asymptotic Pr > ChiSq	0.0259
Exact Pr >= ChiSq	0.0393

Fisher's Exact Test

Cell (1,1) Frequency (F)	11
Left-sided Pr <= F	0.9967
Right-sided Pr >= F	0.0367
Table Probability (P)	0.0334
Two-sided Pr <= P	0.0393

Output 38.5.3 Relative Risk

Estimates of the Relative Risk (Row1/Row2)

Type of Study	Value	95% Confidence Limits	
Case-Control (Odds Ratio)	8.2500	1.1535	59.0029
Cohort (Col1 Risk)	2.9333	0.8502	10.1204
Cohort (Col2 Risk)	0.3556	0.1403	0.9009

Odds Ratio (Case-Control Study)

Odds Ratio	8.2500
Asymptotic Conf Limits	
95% Lower Conf Limit	1.1535
95% Upper Conf Limit	59.0029
Exact Conf Limits	
95% Lower Conf Limit	0.8677
95% Upper Conf Limit	105.5488

Example 38.6: Output Data Set of Chi-Square Statistics

This example uses the Color data from [Example 38.1](#) to output the Pearson chi-square and the likelihood ratio chi-square statistics to a SAS data set. The following PROC FREQ statements create a two-way table of eye color versus hair color.

```
proc freq data=Color order=data;
  tables Eyes*Hair / expected cellchi2 norow nocol chisq;
  output out=ChiSqData n nmiss pchi lrchi;
  weight Count;
  title 'Chi-Square Tests for 3 by 5 Table of Eye and Hair Color';
run;

proc print data=ChiSqData noobs;
  title1 'Chi-Square Statistics for Eye and Hair Color';
  title2 'Output Data Set from the FREQ Procedure';
run;
```

The EXPECTED option displays expected cell frequencies in the crosstabulation table, and the CELLCHI2 option displays the cell contribution to the overall chi-square. The NOROW and NOCOL options suppress the display of row and column percents in the crosstabulation table. The CHISQ option produces chi-square tests.

The OUTPUT statement creates the ChiSqData output data set and specifies the statistics to include. The N option requests the number of nonmissing observations, the NMISS option stores the number of missing observations, and the PCHI and LRCHI options request Pearson and likelihood ratio chi-square statistics, respectively, together with their degrees of freedom and p -values.

The preceding statements produce [Output 38.6.1](#) and [Output 38.6.2](#). The contingency table in [Output 38.6.1](#) displays eye and hair color in the order in which they appear in the Color data set. The Pearson chi-square statistic in [Output 38.6.2](#) provides evidence of an association between eye and hair color ($p=0.0073$). The cell chi-square values show that most of the association is due to more green-eyed children with fair or red hair and fewer with dark or black hair. The opposite occurs with the brown-eyed children.

[Output 38.6.3](#) displays the output data set created by the OUTPUT statement. It includes one observation that contains the sample size, the number of missing values, and the chi-square statistics and corresponding degrees of freedom and p -values as in [Output 38.6.2](#).

Output 38.6.1 Contingency Table

Chi-Square Tests for 3 by 5 Table of Eye and Hair Color						
The FREQ Procedure						
Table of Eyes by Hair						
Eyes(Eye Color)	Hair(Hair Color)					
Frequency						
Expected						
Cell Chi-Square						
Percent	fair	red	medium	dark	black	Total
blue	69	28	68	51	6	222
	66.425	32.921	63.22	53.024	6.4094	
	0.0998	0.7357	0.3613	0.0772	0.0262	
	9.06	3.67	8.92	6.69	0.79	29.13
green	69	38	55	37	0	199
	59.543	29.51	56.671	47.53	5.7454	
	1.5019	2.4422	0.0492	2.3329	5.7454	
	9.06	4.99	7.22	4.86	0.00	26.12
brown	90	47	94	94	16	341
	102.03	50.568	97.109	81.446	9.8451	
	1.4187	0.2518	0.0995	1.935	3.8478	
	11.81	6.17	12.34	12.34	2.10	44.75
Total	228	113	217	182	22	762
	29.92	14.83	28.48	23.88	2.89	100.00

Output 38.6.2 Chi-Square Statistics

Statistic	DF	Value	Prob
Chi-Square	8	20.9248	0.0073
Likelihood Ratio Chi-Square	8	25.9733	0.0011
Mantel-Haenszel Chi-Square	1	3.7838	0.0518
Phi Coefficient		0.1657	
Contingency Coefficient		0.1635	
Cramer's V		0.1172	

Output 38.6.3 Output Data Set

Chi-Square Statistics for Eye and Hair Color							
Output Data Set from the FREQ Procedure							
N	NMISS	_PCHI_	DF_PCHI	P_PCHI	_LRCHI_	DF_LRCHI	P_LRCHI
762	0	20.9248	8	.007349898	25.9733	8	.001061424

Example 38.7: Cochran-Mantel-Haenszel Statistics

The data set *Migraine* contains hypothetical data for a clinical trial of migraine treatment. Subjects of both genders receive either a new drug therapy or a placebo. Their response to treatment is coded as ‘Better’ or ‘Same’. The data are recorded as cell counts, and the number of subjects for each treatment and response combination is recorded in the variable *Count*.

```
data Migraine;
  input Gender $ Treatment $ Response $ Count @@;
  datalines;
female Active Better 16   female Active Same 11
female Placebo Better 5   female Placebo Same 20
male Active Better 12     male Active Same 16
male Placebo Better 7     male Placebo Same 19
;
```

The following PROC FREQ statements create a multiway table stratified by Gender, where Treatment forms the rows and Response forms the columns. The RELRISK option in the TABLES statement requests the odds ratio and relative risks for the two-way tables of Treatment by Response. The PLOTS= option requests a relative risk plot, which shows the relative risk and its confidence limits for each level of Gender and overall. The CMH option requests Cochran-Mantel-Haenszel statistics for the multiway table. For this stratified 2×2 table, the CMH option also produces estimates of the common relative risk and the Breslow-Day test for homogeneity of the odds ratios. The NOPRINT option suppresses the display of the crosstabulation tables.


```
ods graphics on;
proc freq data=Migraine;
  tables Gender*Treatment*Response /
         relrisk plots(only)=relriskplot(stats) cmh noprint;
  weight Count;
  title 'Clinical Trial for Treatment of Migraine Headaches';
run;
ods graphics off;
```

Output 38.7.1 through Output 38.7.4 show the results of the analysis. The relative risk plot (Output 38.7.1) displays the relative risks and confidence limits for the two levels of Gender and for the overall (common) relative risk. Output 38.7.2 displays the CMH statistics. For a stratified 2×2 table, the three CMH statistics test the same hypothesis. The significant p -value (0.004) indicates that the association between treatment and response remains strong after adjusting for gender.

The CMH option also produces a table of overall relative risks, as shown in Output 38.7.3. Because this is a prospective study, the relative risk estimate assesses the effectiveness of the new drug; the “Cohort (Coll Risk)” values are the appropriate estimates for the first column (the risk of improvement). The probability of migraine improvement with the new drug is just over two times the probability of improvement with the placebo.

The large p -value for the Breslow-Day test (0.2218) in Output 38.7.4 indicates no significant gender difference in the odds ratios.

Output 38.7.1 Relative Risk Plot



Output 38.7.2 Cochran-Mantel-Haenszel Statistics

Cochran-Mantel-Haenszel Statistics (Based on Table Scores)				
Statistic	Alternative Hypothesis	DF	Value	Prob
1	Nonzero Correlation	1	8.3052	0.0040
2	Row Mean Scores Differ	1	8.3052	0.0040
3	General Association	1	8.3052	0.0040

Output 38.7.3 CMH Option: Common Relative Risks

Estimates of the Common Relative Risk (Row1/Row2)				
Type of Study	Method	Value	95% Confidence Limits	
Case-Control (Odds Ratio)	Mantel-Haenszel	3.3132	1.4456	7.5934
	Logit	3.2941	1.4182	7.6515
Cohort (Col1 Risk)	Mantel-Haenszel	2.1636	1.2336	3.7948
	Logit	2.1059	1.1951	3.7108
Cohort (Col2 Risk)	Mantel-Haenszel	0.6420	0.4705	0.8761
	Logit	0.6613	0.4852	0.9013

Output 38.7.4 CMH Option: Breslow-Day Test

Breslow-Day Test for Homogeneity of the Odds Ratios	
Chi-Square	1.4929
DF	1
Pr > ChiSq	0.2218

Example 38.8: Cochran-Armitage Trend Test

The data set `Pain` contains hypothetical data for a clinical trial of a drug therapy to control pain. The clinical trial investigates whether adverse responses increase with larger drug doses. Subjects receive either a placebo or one of four drug doses. An adverse response is recorded as `Adverse='Yes'`; otherwise, it is recorded as `Adverse='No'`. The number of subjects for each drug dose and response combination is contained in the variable `Count`.

```
data pain;
  input Dose Adverse $ Count @@;
  datalines;
0 No 26    0 Yes  6
1 No 26    1 Yes  7
2 No 23    2 Yes  9
3 No 18    3 Yes 14
4 No  9    4 Yes 23
;
```

The following PROC FREQ statements provide a trend analysis. The TABLES statement requests a table of Adverse by Dose. The MEASURES option produces measures of association, and the CL option produces confidence limits for these measures. The TREND option tests for a trend across the ordinal values of the variable Dose with the Cochran-Armitage test. The PLOTS= option requests a mosaic plot of Adverse by Dose.


The EXACT statement produces exact p -values for this test, and the MAXTIME= option terminates the exact computations if they do not complete within 60 seconds. The TEST statement computes an asymptotic test for Somers' $D(R|C)$.

```
ods graphics on;
proc freq data=Pain;
  tables Adverse*Dose / trend measures cl
    plots=mosaicplot;
  test smdrc;
  exact trend / maxtime=60;
  weight Count;
  title 'Clinical Trial for Treatment of Pain';
run;
ods graphics off;
```

Output 38.8.1 through Output 38.8.4 display the results of the analysis. The “Col Pct” values in Output 38.8.1 show the expected increasing trend in the proportion of adverse effects with the increasing dosage (from 18.75% to 71.88%). The corresponding mosaic plot (Output 38.8.2) also shows this increasing trend.

Output 38.8.1 Contingency Table

Clinical Trial for Treatment of Pain						
The FREQ Procedure						
Table of Adverse by Dose						
Adverse	Dose					
Frequency						
Percent						
Row Pct						
Col Pct	0	1	2	3	4	Total
-----+-----+-----+-----+-----+-----+-----						
No	26	26	23	18	9	102
	16.15	16.15	14.29	11.18	5.59	63.35
	25.49	25.49	22.55	17.65	8.82	
	81.25	78.79	71.88	56.25	28.13	
-----+-----+-----+-----+-----+-----+-----						
Yes	6	7	9	14	23	59
	3.73	4.35	5.59	8.70	14.29	36.65
	10.17	11.86	15.25	23.73	38.98	
	18.75	21.21	28.13	43.75	71.88	
-----+-----+-----+-----+-----+-----+-----						
Total	32	33	32	32	32	161
	19.88	20.50	19.88	19.88	19.88	100.00

Output 38.8.2 Mosaic Plot

Output 38.8.3 displays the measures of association produced by the MEASURES option. Somers' $D(R|C)$ measures the association treating the row variable (Adverse) as the response and the column variable (Dose) as a predictor. Because the asymptotic 95% confidence limits do not contain zero, this indicates a strong positive association. Similarly, the Pearson and Spearman correlation coefficients show evidence of a strong positive association, as hypothesized.

The Cochran-Armitage test (**Output 38.8.4**) supports the trend hypothesis. The small left-sided p -values for the Cochran-Armitage test indicate that the probability of the Row 1 level (Adverse='No') decreases as Dose increases or, equivalently, that the probability of the Row 2 level (Adverse='Yes') increases as Dose increases. The two-sided p -value tests against either an increasing or decreasing alternative. This is an appropriate hypothesis when you want to determine whether the drug has progressive effects on the probability of adverse effects but the direction is unknown.

Output 38.8.3 Measures of Association

Statistic	Value	ASE	95% Confidence Limits	
Gamma	0.5313	0.0935	0.3480	0.7146
Kendall's Tau-b	0.3373	0.0642	0.2114	0.4631
Stuart's Tau-c	0.4111	0.0798	0.2547	0.5675
Somers' D C R	0.4427	0.0837	0.2786	0.6068
Somers' D R C	0.2569	0.0499	0.1592	0.3547
Pearson Correlation	0.3776	0.0714	0.2378	0.5175
Spearman Correlation	0.3771	0.0718	0.2363	0.5178
Lambda Asymmetric C R	0.1250	0.0662	0.0000	0.2547
Lambda Asymmetric R C	0.2373	0.0837	0.0732	0.4014
Lambda Symmetric	0.1604	0.0621	0.0388	0.2821
Uncertainty Coefficient C R	0.0515	0.0191	0.0140	0.0890
Uncertainty Coefficient R C	0.1261	0.0467	0.0346	0.2175
Uncertainty Coefficient Symmetric	0.0731	0.0271	0.0199	0.1262

Somers' D R C	
Somers' D R C	0.2569
ASE	0.0499
95% Lower Conf Limit	0.1592
95% Upper Conf Limit	0.3547

Test of H0: Somers' D R C = 0	
ASE under H0	0.0499
Z	5.1511
One-sided Pr > Z	<.0001
Two-sided Pr > Z	<.0001

Output 38.8.4 Trend Test

Cochran-Armitage Trend Test	

Statistic (Z)	-4.7918
Asymptotic Test	
One-sided Pr < Z	<.0001
Two-sided Pr > Z	<.0001
Exact Test	
One-sided Pr <= Z	7.237E-07
Two-sided Pr >= Z	1.324E-06

Example 38.9: Friedman's Chi-Square Test

Friedman's test is a nonparametric test for treatment differences in a randomized complete block design. Each block of the design might be a subject or a homogeneous group of subjects. If blocks are groups of subjects, the number of subjects in each block must equal the number of treatments. Treatments are randomly assigned to subjects within each block. If there is one subject per block, then the subjects are repeatedly measured once under each treatment. The order of treatments is randomized for each subject.

In this setting, Friedman's test is identical to the ANOVA (row means scores) CMH statistic when the analysis uses rank scores (SCORES=RANK). The three-way table uses subject (or subject group) as the stratifying variable, treatment as the row variable, and response as the column variable. PROC FREQ handles ties by assigning midranks to tied response values. If there are multiple subjects per treatment in each block, the ANOVA CMH statistic is a generalization of Friedman's test.

The data set Hypnosis contains data from a study investigating whether hypnosis has the same effect on skin potential (measured in millivolts) for four emotions (Lehmann 2006, p. 264). Eight subjects are asked to display fear, joy, sadness, and calmness under hypnosis. The data are recorded as one observation per subject for each emotion.

```
data Hypnosis;
  length Emotion $ 10;
  input Subject Emotion $ SkinResponse @@;
  datalines;
1 fear 23.1 1 joy 22.7 1 sadness 22.5 1 calmness 22.6
2 fear 57.6 2 joy 53.2 2 sadness 53.7 2 calmness 53.1
3 fear 10.5 3 joy 9.7 3 sadness 10.8 3 calmness 8.3
4 fear 23.6 4 joy 19.6 4 sadness 21.1 4 calmness 21.6
5 fear 11.9 5 joy 13.8 5 sadness 13.7 5 calmness 13.3
6 fear 54.6 6 joy 47.1 6 sadness 39.2 6 calmness 37.0
7 fear 21.0 7 joy 13.6 7 sadness 13.7 7 calmness 14.8
8 fear 20.3 8 joy 23.6 8 sadness 16.3 8 calmness 14.8
;
```

In the following PROC FREQ statements, the TABLES statement creates a three-way table stratified by Subject and a two-way table; the variables Emotion and SkinResponse form the rows and columns of each table. The CMH2 option produces the first two Cochran-Mantel-Haenszel statistics, the option SCORES=RANK specifies that rank scores are used to compute these statistics, and the NOPRINT option suppresses the contingency tables. These statements produce [Output 38.9.1](#) and [Output 38.9.2](#).

```
proc freq data=Hypnosis;
  tables Subject*Emotion*SkinResponse /
    cmh2 scores=rank noprint;
run;

proc freq data=Hypnosis;
  tables Emotion*SkinResponse /
    cmh2 scores=rank noprint;
run;
```

Because the CMH statistics in [Output 38.9.1](#) are based on rank scores, the Row Mean Scores Differ statistic is identical to Friedman's chi-square ($Q = 6.45$). The p -value of 0.0917 indicates that differences in skin potential response for different emotions are significant at the 10% level but not at the 5% level.

When you do not stratify by subject, the Row Mean Scores Differ CMH statistic is identical to a Kruskal-Wallis test and is not significant ($p = 0.9038$ in [Output 38.9.2](#)). Thus, adjusting for subject is critical to reducing the background variation due to subject differences.

Output 38.9.1 CMH Statistics: Stratifying by Subject

Clinical Trial for Treatment of Pain				
The FREQ Procedure				
Summary Statistics for Emotion by SkinResponse Controlling for Subject				
Cochran-Mantel-Haenszel Statistics (Based on Rank Scores)				
Statistic	Alternative Hypothesis	DF	Value	Prob
1	Nonzero Correlation	1	0.2400	0.6242
2	Row Mean Scores Differ	3	6.4500	0.0917

Output 38.9.2 CMH Statistics: No Stratification

Clinical Trial for Treatment of Pain				
The FREQ Procedure				
Summary Statistics for Emotion by SkinResponse				
Cochran-Mantel-Haenszel Statistics (Based on Rank Scores)				
Statistic	Alternative Hypothesis	DF	Value	Prob
1	Nonzero Correlation	1	0.0001	0.9933
2	Row Mean Scores Differ	3	0.5678	0.9038

Example 38.10: Cochran's Q Test

When a binary response is measured several times or under different conditions, Cochran's Q tests that the marginal probability of a positive response is unchanged across the times or conditions. When there are more than two response categories, you can use the CATMOD procedure to fit a repeated-measures model.

The data set `Drugs` contains data for a study of three drugs to treat a chronic disease (Agresti 2002). Forty-six subjects receive drugs A, B, and C. The response to each drug is either favorable ('F') or unfavorable ('U').

```
proc format;
  value $ResponseFmt 'F'='Favorable'
                    'U'='Unfavorable';
run;
```

```

data drugs;
  input Drug_A $ Drug_B $ Drug_C $ Count @@;
  datalines;
F F F 6   U F F 2
F F U 16  U F U 4
F U F 2   U U F 6
F U U 4   U U U 6
;

```

The following statements create one-way frequency tables of the responses to each drug. The AGREE option produces Cochran's Q and other measures of agreement for the three-way table. These statements produce [Output 38.10.1](#) through [Output 38.10.5](#).

```

proc freq data=Drugs;
  tables Drug_A Drug_B Drug_C / nocum;
  tables Drug_A*Drug_B*Drug_C / agree noprint;
  format Drug_A Drug_B Drug_C $ResponseFmt.;
  weight Count;
  title 'Study of Three Drug Treatments for a Chronic Disease';
run;

```

The one-way frequency tables in [Output 38.10.1](#) provide the marginal response for each drug. For drugs A and B, 61% of the subjects reported a favorable response while 35% of the subjects reported a favorable response to drug C. [Output 38.10.2](#) and [Output 38.10.3](#) display measures of agreement for the 'Favorable' and 'Unfavorable' levels of drug A, respectively. McNemar's test shows a strong discordance between drugs B and C when the response to drug A is favorable.

Output 38.10.1 One-Way Frequency Tables

Study of Three Drug Treatments for a Chronic Disease		
The FREQ Procedure		
Drug_A	Frequency	Percent

Favorable	28	60.87
Unfavorable	18	39.13
Drug_B	Frequency	Percent

Favorable	28	60.87
Unfavorable	18	39.13
Drug_C	Frequency	Percent

Favorable	16	34.78
Unfavorable	30	65.22

Output 38.10.2 Measures of Agreement for Drug A Favorable

McNemar's Test	

Statistic (S)	10.8889
DF	1
Pr > S	0.0010
Simple Kappa Coefficient	

Kappa	-0.0328
ASE	0.1167
95% Lower Conf Limit	-0.2615
95% Upper Conf Limit	0.1960

Output 38.10.3 Measures of Agreement for Drug A Unfavorable

McNemar's Test	

Statistic (S)	0.4000
DF	1
Pr > S	0.5271
Simple Kappa Coefficient	

Kappa	-0.1538
ASE	0.2230
95% Lower Conf Limit	-0.5909
95% Upper Conf Limit	0.2832

Output 38.10.4 displays the overall kappa coefficient. The small negative value of kappa indicates no agreement between drug B response and drug C response.

Output 38.10.4 Overall Measures of Agreement

Overall Kappa Coefficient	

Kappa	-0.0588
ASE	0.1034
95% Lower Conf Limit	-0.2615
95% Upper Conf Limit	0.1439
Test for Equal Kappa Coefficients	

Chi-Square	0.2314
DF	1
Pr > ChiSq	0.6305

Cochran's Q is statistically significant ($p=0.0145$ in [Output 38.10.5](#)), which leads to rejection of the hypothesis that the probability of favorable response is the same for the three drugs.

Output 38.10.5 Cochran's Q Test

Cochran's Q, for Drug_A by Drug_B by Drug_C	

Statistic (Q)	8.4706
DF	2
Pr > Q	0.0145

References

- Agresti, A. (1992), "A Survey of Exact Inference for Contingency Tables," *Statistical Science*, 7, 131–177.
- Agresti, A. (2002), *Categorical Data Analysis*, Second Edition, New York: John Wiley & Sons.
- Agresti, A. (2007), *An Introduction to Categorical Data Analysis*, Second Edition, New York: John Wiley & Sons.
- Agresti, A. and Caffo, B. (2000), "Simple and Effective Confidence Intervals for Proportions and Differences of Proportions Result from Adding Two Successes and Two Failures," *American Statistician*, 54, 280–288.
- Agresti, A. and Coull, B. A. (1998), "Approximate is Better Than 'Exact' for Interval Estimation of Binomial Proportions," *The American Statistician*, 52, 119–126.
- Agresti, A., Mehta, C. R., and Patel, N. R. (1990), "Exact Inference for Contingency Tables with Ordered Categories," *Journal of the American Statistical Association*, 85, 453–458.
- Agresti, A. and Min, Y. (2001), "On Small-Sample Confidence Intervals for Parameters in Discrete Distributions," *Biometrics*, 57, 963–971.
- Agresti, A., Wackerly, D., and Boyett, J. M. (1979), "Exact Conditional Tests for Cross-Classifications: Approximation of Attained Significance Levels," *Psychometrika*, 44, 75–83.
- Bangdiwala, S. I. (1988), *The Agreement Chart*, Technical report, University of North Carolina at Chapel Hill, Department of Biostatistics.
- Bangdiwala, S. I. and Bryan, H. E. (1987), "Using SAS Software Graphical Procedures for the Observer Agreement Chart," in *Proceedings of the Twelfth Annual SAS Users Group International Conference*, 1083–1088, Cary, NC: SAS Institute Inc.
- Bangdiwala, S. I., Haedo, A. S., Natal, M. L., and Villaveces, A. (2008), "The Agreement Chart as an Alternative to the Receiver-Operating Characteristic Curve for Diagnostic Tests," *Journal of Clinical Epidemiology*, 61, 866–874.

- Barker, L., Rolka, H., Rolka, D., and Brown, C. (2001), "Equivalence Testing for Binomial Random Variables: Which Test to Use?" *The American Statistician*, 55, 279–287.
- 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.
- Barnard, G. A. (1949), "Statistical Inference," *Journal of the Royal Statistical Society, Series B*, 11, 115–139.
- Berger, J. O. (1985), *Statistical Decision Theory and Bayesian Analysis*, Second Edition, New York: Springer-Verlag.
- Birch, M. W. (1965), "The Detection of Partial Association, II: The General Case," *Journal of the Royal Statistical Society, Series B*, 27, 111–124.
- Bishop, Y. M. M., Fienberg, S. E., and Holland, P. W. (1975), *Discrete Multivariate Analysis: Theory and Practice*, Cambridge, MA: MIT Press.
- Bowker, A. H. (1948), "Bowker's Test for Symmetry," *Journal of the American Statistical Association*, 43, 572–574.
- Breslow, N. E. (1996), "Statistics in Epidemiology: The Case-Control Study," *Journal of the American Statistical Association*, 91, 14–26.
- Breslow, N. E. and Day, N. E. (1980), *Statistical Methods in Cancer Research, Volume I: The Analysis of Case-Control Studies*, IARC Scientific Publications, No. 32, Lyon, France: International Agency for Research on Cancer.
- Breslow, N. E. and Day, N. E. (1987), *Statistical Methods in Cancer Research, Volume II: The Design and Analysis of Cohort Studies*, IARC Scientific Publications, No. 82, Lyon, France: International Agency for Research on Cancer.
- Bross, I. D. J. (1958), "How to Use Ridit Analysis," *Biometrics*, 14, 18–38.
- Brown, L. D., Cai, T. T., and DasGupta, A. (2001), "Interval Estimation for a Binomial Proportion," *Statistical Science*, 16, 101–133.
- Brown, M. B. and Benedetti, J. K. (1977), "Sampling Behavior of Tests for Correlation in Two-Way Contingency Tables," *Journal of the American Statistical Association*, 72, 309–315.
- Chan, I. S. F. (1998), "Exact Tests of Equivalence and Efficacy with a Non-Zero Lower Bound for Comparative Studies," *Statistics in Medicine*, 17, 1403–1413.
- Chan, I. S. F. (2003), "Proving Non-Inferiority or Equivalence of Two Treatments with Dichotomous Endpoints Using Exact Methods," *Statistical Methods in Medical Research*, 12, 37–58.
- Chan, I. S. F. and Zhang, Z. (1999), "Test-Based Exact Confidence Intervals for the Difference of Two Binomial Proportions," *Biometrics*, 55, 1202–1209.
- Chow, S.-C., Shao, J., and Wang, H. (2003), *Sample Size Calculations in Clinical Research*, Boca Raton, FL: CRC Press.

- Cicchetti, D. V. and Allison, T. (1971), "A New Procedure for Assessing Reliability of Scoring EEG Sleep Recordings," *American Journal of EEG Technology*, 11, 101–109.
- Clopper, C. J. and Pearson, E. S. (1934), "The Use of Confidence or Fiducial Limits Illustrated in the Case of the Binomial," *Biometrika*, 26, 404–413.
- Cochran, W. G. (1950), "The Comparison of Percentages in Matched Samples," *Biometrika*, 37, 256–266.
- Cochran, W. G. (1954), "Some Methods for Strengthening the Common χ^2 Tests," *Biometrics*, 10, 417–451.
- Cohen, J. (1960), "A Coefficient of Agreement for Nominal Scales," *Educational and Psychological Measurement*, 20, 37–46.
- Collett, D. (1991), *Modelling Binary Data*, London: Chapman & Hall.
- Dann, R. S. and Koch, G. G. (2005), "Review and Evaluation of Methods for Computing Confidence Intervals for the Ratio of Two Proportions and Considerations for Non-inferiority Clinical Trials," *Journal of Biopharmaceutical Statistics*, 15, 85–107.
- Dmitrienko, A., Molenberghs, G., Chuang-Stein, C., and Offen, W. (2005), *Analysis of Clinical Trials Using SAS: A Practical Guide*, Cary, NC: SAS Institute Inc.
- Dragow, F. (1986), "Polychoric and Polyserial Correlations," in S. Kotz and N. L. Johnson, eds., *Encyclopedia of Statistical Sciences*, volume 7, 68–74, New York: John Wiley & Sons.
- Dunnnett, C. W. and Gent, M. (1977), "Significance Testing to Establish Equivalence between Treatments, with Special Reference to Data in the Form of 2×2 Tables," *Biometrics*, 33, 593–602.
- Farrington, C. P. and Manning, G. (1990), "Test Statistics and Sample Size Formulae for Comparative Binomial Trials with Null Hypothesis of Non-Zero Risk Difference or Non-Unity Relative Risk," *Statistics in Medicine*, 9, 1447–1454.
- Fienberg, S. E. (1980), *The Analysis of Cross-Classified Categorical Data*, Second Edition, Cambridge, MA: MIT Press.
- Fleiss, J. L. and Cohen, J. (1973), "The Equivalence of Weighted Kappa and the Intraclass Correlation Coefficient as Measures of Reliability," *Educational and Psychological Measurement*, 33, 613–619.
- Fleiss, J. L., Cohen, J., and Everitt, B. S. (1969), "Large-Sample Standard Errors of Kappa and Weighted Kappa," *Psychological Bulletin*, 72, 323–327.
- Fleiss, J. L., Levin, B., and Paik, M. C. (2003), *Statistical Methods for Rates and Proportions*, Third Edition, Hoboken, NJ: John Wiley & Sons.
- Freeman, G. H. and Halton, J. H. (1951), "Note on an Exact Treatment of Contingency, Goodness of Fit and Other Problems of Significance," *Biometrika*, 38, 141–149.
- Friendly, M. (2000), *Visualizing Categorical Data*, Cary, NC: SAS Institute Inc.
- Gail, M. and Mantel, N. (1977), "Counting the Number of $r \times c$ Contingency Tables with Fixed Margins," *Journal of the American Statistical Association*, 72, 859–862.
- Gail, M. and Simon, R. (1985), "Tests for Qualitative Interactions between Treatment Effects and Patient Subsets," *Biometrics*, 41, 361–372.

- Gart, J. J. (1971), "The Comparison of Proportions: A Review of Significance Tests, Confidence Intervals and Adjustments for Stratification," *Review of the International Statistical Institute*, 39, 148–169.
- Gart, J. J. and Nam, J. (1988), "Approximate Interval Estimation of the Ratio of Binomial Parameters: A Review and Corrections for Skewness," *Biometrics*, 44, 323–338.
- Goodman, L. A. and Kruskal, W. H. (1979), *Measures of Association for Cross Classification*, New York: Springer-Verlag.
- Greenland, S. and Robins, J. M. (1985), "Estimation of a Common Effect Parameter from Sparse Follow-up Data," *Biometrics*, 41, 55–68.
- Haldane, J. B. S. (1955), "The Estimation and Significance of the Logarithm of a Ratio of Frequencies," *Annals of Human Genetics*, 20, 309–314.
- Hauck, W. W. and Anderson, S. (1986), "A Comparison of Large-Sample Confidence Interval Methods for the Difference of Two Binomial Probabilities," *The American Statistician*, 40, 318–322.
- Hirji, K. F. (2006), *Exact Analysis of Discrete Data*, Boca Raton, FL: Chapman & Hall/CRC.
- Hirji, K. F., Vollset, S. E., Reis, I. M., and Afifi, A. A. (1996), "Exact Tests for Interaction in Several 2×2 Tables," *Journal of Computational and Graphical Statistics*, 5, 209–224.
- Hollander, M. and Wolfe, D. A. (1999), *Nonparametric Statistical Methods*, Second Edition, New York: John Wiley & Sons.
- Jones, M. P., O'Gorman, T. W., Lemka, J. H., and Woolson, R. F. (1989), "A Monte Carlo Investigation of Homogeneity Tests of the Odds Ratio under Various Sample Size Configurations," *Biometrics*, 45, 171–181.
- Kendall, M. (1955), *Rank Correlation Methods*, Second Edition, London: Charles Griffin and Co.
- Kendall, M. and Stuart, A. (1979), *The Advanced Theory of Statistics*, volume 2, New York: Macmillan.
- Kleinbaum, D. G., Kupper, L. L., and Morgenstern, H. (1982), *Epidemiologic Research: Principles and Quantitative Methods*, Research Methods Series, New York: Van Nostrand Reinhold.
- Landis, R. J., Heyman, E. R., and Koch, G. G. (1978), "Average Partial Association in Three-Way Contingency Tables: A Review and Discussion of Alternative Tests," *International Statistical Review*, 46, 237–254.
- Leemis, L. M. and Trivedi, K. S. (1996), "A Comparison of Approximate Interval Estimators for the Bernoulli Parameter," *The American Statistician*, 50, 63–68.
- Lehmann, E. L. (2006), *Nonparametrics: Statistical Methods Based on Ranks*, New York: Springer Science+Business Media.
- Liebetrau, A. M. (1983), *Measures of Association*, Quantitative Applications in the Social Sciences, vol. 32, Beverly Hills: Sage Publications.
- Mack, G. A. and Skillings, J. H. (1980), "A Friedman-Type Rank Test for Main Effects in a Two-Factor ANOVA," *Journal of the American Statistical Association*, 75, 947–951.

- Mantel, N. (1963), "Chi-Square Tests with One Degree of Freedom: Extensions of the Mantel-Haenszel Procedure," *Journal of the American Statistical Association*, 58, 690–700.
- Mantel, N. and Fleiss, J. L. (1980), "Minimum Expected Cell Size Requirements for the Mantel-Haenszel One-Degree-of-Freedom Chi-Square Test and a Related Rapid Procedure," *American Journal of Epidemiology*, 112, 129–134.
- Mantel, N. and Haenszel, W. (1959), "Statistical Aspects of Analysis of Data from Retrospective Studies of Disease," *Journal of the National Cancer Institute*, 22, 719–748.
- Margolin, B. H. (1988), "Test for Trend in Proportions," in S. Kotz and N. L. Johnson, eds., *Encyclopedia of Statistical Sciences*, volume 9, 334–336, New York: John Wiley & Sons.
- McNemar, Q. (1947), "Note on the Sampling Error of the Difference between Correlated Proportions or Percentages," *Psychometrika*, 12, 153–157.
- Mee, R. W. (1984), "Confidence Bounds for the Difference between Two Probabilities," *Biometrics*, 40, 1175–1176.
- Mehta, C. R. and Patel, N. R. (1983), "A Network Algorithm for Performing Fisher's Exact Test in $r \times c$ Contingency Tables," *Journal of the American Statistical Association*, 78, 427–434.
- Mehta, C. R., Patel, N. R., and Gray, R. (1985), "Computing an Exact Confidence Interval for the Common Odds Ratio in Several 2×2 Contingency Tables," *Journal of American Statistical Association*, 80, 969–973.
- Mehta, C. R., Patel, N. R., and Senchaudhuri, P. (1991), "Exact Stratified Linear Rank Tests for Binary Data," in E. M. Keramidas, ed., *Computing Science and Statistics: Proceedings of the 23rd Symposium on the Interface*, 200–207.
- Mehta, C. R., Patel, N. R., and Tsiatis, A. A. (1984), "Exact Significance Testing to Establish Treatment Equivalence with Ordered Categorical Data," *Biometrics*, 40, 819–825.
- Mehta, C. R. and Senchaudhuri, P. (2003), "Conditional versus Unconditional Exact Tests for Comparing Two Binomials," .
- Miettinen, O. and Nurminen, M. (1985), "Comparative Analysis of Two Rates," *Statistics in Medicine*, 4, 213–226.
- Newcombe, R. G. (1998a), "Interval Estimation for the Difference between Independent Proportions: Comparison of Eleven Methods," *Statistics in Medicine*, 17, 873–890.
- Newcombe, R. G. (1998b), "Two-Sided Confidence Intervals for the Single Proportion: Comparison of Seven Methods," *Statistics in Medicine*, 17, 857–872.
- Newcombe, R. G. and Nurminen, M. M. (2011), "In Defence of Score Intervals for Proportions and Their Differences," *Communications in Statistics—Theory and Methods*, 40, 1271–1282.
- Olsson, U. (1979), "Maximum Likelihood Estimation of the Polychoric Correlation Coefficient," *Psychometrika*, 12, 443–460.
- Pirie, W. (1983), "Jonckheere Tests for Ordered Alternatives," in S. Kotz and N. L. Johnson, eds., *Encyclopedia of Statistical Sciences*, volume 4, 315–318, New York: John Wiley & Sons.

- Radlow, R. and Alf, E. F. (1975), "An Alternate Multinomial Assessment of the Accuracy of the Chi-Square Test of Goodness of Fit," *Journal of the American Statistical Association*, 70, 811–813.
- Robins, J. M., Breslow, N., and Greenland, S. (1986), "Estimators of the Mantel-Haenszel Variance Consistent in Both Sparse Data and Large-Strata Limiting Models," *Biometrics*, 42, 311–323.
- Santner, T. J., Pradhan, V., Senchaudhuri, P., Mehta, C. R., and Tamhane, A. (2007), "Small-Sample Comparisons of Confidence Intervals for the Difference of Two Independent Binomial Proportions," *Computational Statistics and Data Analysis*, 51, 5791–5799.
- Santner, T. J. and Snell, M. K. (1980), "Small-Sample Confidence Intervals for $p_1 - p_2$ and p_1/p_2 in 2×2 Contingency Tables," *Journal of the American Statistical Association*, 75, 386–394.
- Schuurmann, D. J. (1987), "A Comparison of the Two One-Sided Tests Procedure and the Power Approach for Assessing the Equivalence of Average Bioavailability," *Journal of Pharmacokinetics and Biopharmaceutics*, 15, 657–680.
- Schuurmann, D. J. (1999), "Confidence Interval Methods for Bioequivalence Testing with Binomial Endpoints," in *Proceedings of the Biopharmaceutical Section*, 227–232, American Statistical Association.
- Silvapulle, M. J. (2001), "Tests Against Qualitative Interaction: Exact Critical Values and Robust Tests," *Biometrics*, 57, 1157–1165.
- Snedecor, G. W. and Cochran, W. G. (1989), *Statistical Methods*, Eighth Edition, Ames: Iowa State University Press.
- Somers, R. H. (1962), "A New Asymmetric Measure of Association for Ordinal Variables," *American Sociological Review*, 27, 799–811.
- Stokes, M. E., Davis, C. S., and Koch, G. G. (2012), *Categorical Data Analysis Using SAS*, Third edition Edition, Cary, NC: SAS Institute Inc.
- Suissa, S. and Shuster, J. J. (1985), "Exact Unconditional Sample Sizes for the 2x2 Binomial Trial," *Journal of the Royal Statistical Society, Series A*, 148, 317–327.
- Tarone, R. E. (1985), "On Heterogeneity Tests Based on Efficient Scores," *Biometrika*, 72, 1, 91–95.
- Theil, H. (1972), *Statistical Decomposition Analysis*, Amsterdam: North Holland Publishing.
- Thomas, D. G. (1971), "Algorithm AS-36. Exact Confidence Limits for the Odds Ratio in a 2×2 Table," *Applied Statistics*, 20, 105–110.
- Valz, P. D. and Thompson, M. E. (1994), "Exact Inference for Kendall's S and Spearman's ρ with Extensions to Fisher's Exact Test in $r \times c$ Contingency Tables," *Journal of Computational and Graphical Statistics*, 3, 459–472.
- van Elteren, P. H. (1960), "On the Combination of Independent Two-Sample Tests of Wilcoxon," *Bulletin of the International Statistical Institute*, 37, 351–361.
- Vollset, S. E., Hirji, K. F., and Elashoff, R. M. (1991), "Fast Computation of Exact Confidence Limits for the Common Odds Ratio in a Series of 2×2 Tables," *Journal of the American Statistical Association*, 86, 404–409.

Wilson, E. B. (1927), “Probable Inference, the Law of Succession, and Statistical Inference,” *Journal of the American Statistical Association*, 22, 209–212.

Woolf, B. (1955), “On Estimating the Relationship between Blood Group and Disease,” *Annals of Human Genetics*, 19, 251–253.

Zelen, M. (1971), “The Analysis of Several 2×2 Contingency Tables,” *Biometrika*, 58, 129–137.

Subject Index

- adjusted odds ratio
 - FREQ procedure, 2596
- adjusted relative risks
 - FREQ procedure, 2597
- agreement plots
 - FREQ procedure, 2484
- agreement, measures of
 - FREQ procedure, 2588
- Agresti-Caffo confidence limits
 - risk difference (FREQ), 2574
- Agresti-Coull confidence limits
 - proportions (FREQ), 2566
- ANOVA (row mean scores) statistic
 - Mantel-Haenszel (FREQ), 2594
- association, measures of
 - FREQ procedure, 2556
- bar charts
 - FREQ procedure, 2533
- Barnard's test
 - FREQ procedure, 2581
- binomial proportions
 - Clopper-Pearson test (FREQ), 2568
 - confidence limits (FREQ), 2566
 - equivalence tests (FREQ), 2570
 - exact test (FREQ), 2568
 - FREQ procedure, 2565
 - noninferiority tests (FREQ), 2569
 - superiority tests (FREQ), 2570
 - tests (FREQ), 2568
 - TOST (FREQ), 2571
- Bowker's test of symmetry
 - FREQ procedure, 2588
- Breslow-Day test
 - FREQ procedure, 2598
- Tarone's adjustment (FREQ), 2598
- case-control studies
 - odds ratio (FREQ), 2582
- categorical data analysis
 - FREQ procedure, 2472
- cell count data
 - example (FREQ), 2622
 - FREQ procedure, 2544
- chi-square goodness-of-fit test
 - FREQ procedure, 2551
- chi-square tests
 - FREQ procedure, 2550
- Cicchetti-Allison weights
 - kappa coefficient (FREQ), 2591
- Clopper-Pearson confidence limits
 - proportions (FREQ), 2567
- Cochran's Q test
 - FREQ procedure, 2588, 2592
- Cochran-Armitage test for trend
 - FREQ procedure, 2585
- Cochran-Mantel-Haenszel statistics
 - FREQ procedure, 2592
- cohort studies
 - relative risks (FREQ), 2583
- common odds ratio
 - exact confidence limits (FREQ), 2599
 - exact test (FREQ), 2599
 - logit (FREQ), 2596
 - Mantel-Haenszel (FREQ), 2596
- common relative risks
 - logit (FREQ), 2597
 - Mantel-Haenszel (FREQ), 2597
- concordant observations
 - FREQ procedure, 2556
- confidence limits
 - exact (FREQ), 2487
 - measures of association (FREQ), 2556
 - proportions (FREQ), 2566
- contingency coefficient
 - FREQ procedure, 2555
- contingency tables
 - FREQ procedure, 2472, 2505
- continuity-adjusted chi-square test
 - FREQ procedure, 2553
- correlation statistic
 - Mantel-Haenszel (FREQ), 2594
- Cramér's V statistic
 - FREQ procedure, 2555
- crosstabulation tables
 - FREQ procedure, 2472, 2505, 2611
- discordant observations
 - FREQ procedure, 2556
- dot plots
 - FREQ procedure, 2533, 2625
- equivalence tests
 - binomial proportions (FREQ), 2570
 - risk difference (FREQ), 2580
- exact confidence limits
 - odds ratio (FREQ), 2583
 - proportion difference (FREQ), 2580

- proportions (FREQ), 2567
- ratio of proportions (FREQ), 2584
- relative risks (FREQ), 2584
- risk difference (FREQ), 2580
- exact p -values
 - FREQ procedure, 2603
- exact tests
 - computational algorithms (FREQ), 2602
 - computational resources (FREQ), 2604
 - FREQ procedure, 2487, 2601, 2643
 - Monte Carlo estimation (FREQ), 2493
 - network algorithm (FREQ), 2602
- Farrington-Manning test
 - risk difference (FREQ), 2578
- Fisher's exact test
 - FREQ procedure, 2554
- Fleiss-Cohen weights
 - kappa coefficient (FREQ), 2591
- Freeman-Halton test
 - FREQ procedure, 2555
- FREQ procedure
 - adjusted odds ratio (Mantel-Haenszel), 2596
 - adjusted relative risks (Mantel-Haenszel), 2597
 - Agresti-Caffo confidence limits, 2574
 - Agresti-Coull confidence limits, 2566
 - ANOVA (row mean scores) statistic, 2594
 - bar charts, 2533
 - Barnard's test, 2581
 - binomial proportions, 2565
 - Bowker's test of symmetry, 2588
 - Breslow-Day test, 2598
 - cell count data, 2544
 - chi-square goodness-of-fit test, 2551
 - chi-square tests, 2550
 - Clopper-Pearson confidence limits, 2567
 - Cochran's Q test, 2588, 2592
 - Cochran-Armitage test for trend, 2585
 - common odds ratio, 2599
 - computational resources, 2606
 - computational resources (exact tests), 2604
 - contingency coefficient, 2555
 - continuity-adjusted chi-square test, 2553
 - correlation statistic, 2594
 - Cramér's V statistic, 2555
 - crosstabulation tables, 2611
 - default tables, 2505
 - displayed output, 2609
 - dot plots, 2533, 2625
 - equivalence tests, 2570
 - equivalence tests (risk difference), 2580
 - exact confidence limits, 2487
 - exact p -values, 2603
 - exact tests, 2487, 2601, 2643
 - exact unconditional confidence limits, 2580
 - Farrington-Manning test, 2578
 - Fisher's exact test, 2554
 - Freeman-Halton test, 2555
 - Friedman's chi-square test, 2646
 - Gail-Simon test, 2601
 - gamma statistic, 2556, 2557
 - general association statistic, 2595
 - grouping with formats, 2544
 - Hauck-Anderson confidence limits, 2574
 - in-database computation, 2548
 - input data sets, 2485, 2544
 - introductory examples, 2474
 - Jeffreys confidence limits, 2566
 - Jonckheere-Terpstra test, 2586
 - kappa coefficient, 2588, 2589
 - Kendall's tau- b statistic, 2556, 2558
 - lambda asymmetric, 2556, 2563
 - lambda symmetric, 2556, 2563
 - likelihood ratio chi-square test, 2553
 - Mantel-Fleiss criterion, 2595
 - Mantel-Haenszel chi-square test, 2553
 - Mantel-Haenszel statistics, 2592
 - maximum time (exact tests), 2493
 - McNemar's test, 2588
 - measures of agreement, 2588
 - measures of association, 2556
 - Miettinen-Nurminen confidence limits, 2575
 - missing values, 2545
 - Monte Carlo estimation (exact tests), 2487, 2493, 2604
 - mosaic plots, 2527
 - multiway tables, 2611
 - network algorithm, 2602
 - Newcombe confidence limits, 2575, 2579
 - noninferiority tests, 2569
 - noninferiority tests (risk difference), 2577
 - odds ratio, 2582
 - ODS graph names, 2621
 - ODS table names, 2617
 - one-way frequency tables, 2610
 - ordering of levels, 2486
 - output data sets, 2494, 2606
 - overall kappa coefficient, 2591
 - Pearson chi-square test, 2551
 - Pearson correlation coefficient, 2556, 2559
 - phi coefficient, 2555
 - polychoric correlation coefficient, 2556, 2562
 - relative risks, 2583
 - risk difference, 2572
 - scores, 2549
 - simple kappa coefficient, 2589
 - Somers' D statistics, 2556, 2559

- Spearman rank correlation coefficient, 2556, 2560
- standardized residuals, 2552
- Stuart's tau-*c* statistic, 2556, 2558
- superiority tests, 2570
- superiority tests (risk difference), 2579
- tetrachoric correlation coefficient, 2562
- uncertainty coefficients, 2556, 2564, 2565
- Wald confidence limits (risk difference), 2576
- weighted kappa coefficient, 2588, 2590
- Wilson confidence limits, 2567
- Yule's *Q* statistic, 2557
- Zelen's exact test, 2598
- frequency plots
 - FREQ procedure, 2476
- frequency tables
 - FREQ procedure, 2472, 2505
 - one-way (FREQ), 2610
- Friedman's chi-square test
 - FREQ procedure, 2646
- Gail-Simon test
 - FREQ procedure, 2601
- gamma statistic
 - FREQ procedure, 2556, 2557
- general association statistic
 - Mantel-Haenszel (FREQ), 2595
- Hauck-Anderson confidence limits
 - risk difference (FREQ), 2574
- hypothesis tests
 - exact (FREQ), 2487
- in-database computation
 - FREQ procedure, 2548
- Jeffreys confidence limits
 - proportions (FREQ), 2566
- Jonckheere-Terpstra test
 - FREQ procedure, 2586
- kappa coefficient
 - FREQ procedure, 2588, 2589
 - weights (FREQ), 2591
- Kendall's tau-*b* statistic
 - FREQ procedure, 2556, 2558
- lambda asymmetric
 - FREQ procedure, 2556, 2563
- lambda symmetric
 - FREQ procedure, 2556, 2563
- likelihood ratio chi-square test
 - FREQ procedure, 2553
- Mantel-Fleiss criterion
 - FREQ procedure, 2595
- Mantel-Haenszel chi-square test
 - FREQ procedure, 2553
- Mantel-Haenszel statistics
 - ANOVA (row mean scores) statistic (FREQ), 2594
 - correlation statistic (FREQ), 2594
 - FREQ procedure, 2592
 - general association statistic (FREQ), 2595
 - Mantel-Fleiss criterion (FREQ), 2595
- McNemar's test
 - FREQ procedure, 2588
- measures of association
 - exact tests (FREQ), 2557
 - tests (FREQ), 2556
- Mehta-Patel network algorithm
 - exact tests (FREQ), 2602
- Miettinen-Nurminen confidence limits
 - risk difference (FREQ), 2575
- missing values
 - FREQ procedure, 2545
- modified ridit scores
 - FREQ procedure, 2550
- Monte Carlo estimation
 - exact tests (FREQ), 2487, 2493, 2604
- mosaic plots
 - FREQ procedure, 2527
- multiway tables
 - FREQ procedure, 2472, 2505, 2611
- network algorithm
 - exact tests (FREQ), 2602
- Newcombe confidence limits
 - risk difference (FREQ), 2575, 2579
- noninferiority tests
 - binomial proportions (FREQ), 2569
 - risk difference (FREQ), 2577
- odds ratio
 - Breslow-Day test (FREQ), 2598
 - case-control studies (FREQ), 2582
 - exact confidence limits (FREQ), 2583
 - FREQ procedure, 2582
 - logit adjusted (FREQ), 2596
 - Mantel-Haenszel adjusted (FREQ), 2596
 - Zelen's exact test (FREQ), 2598
- ODS graph names
 - FREQ procedure, 2621
- Pearson chi-square test
 - FREQ procedure, 2551
- Pearson correlation coefficient
 - FREQ procedure, 2556, 2559
- phi coefficient
 - FREQ procedure, 2555

- polychoric correlation coefficient
 - FREQ procedure, 2556, 2562
- proportion difference
 - FREQ procedure, 2572
- proportions, *see* binomial proportions (FREQ)
- rank scores
 - FREQ procedure, 2550
- relative risks
 - cohort studies (FREQ), 2583
 - exact confidence limits (FREQ), 2584
 - FREQ procedure, 2583
 - logit adjusted (FREQ), 2597
 - Mantel-Haenszel adjusted (FREQ), 2597
- ridit scores
 - FREQ procedure, 2550
- risk difference
 - confidence limits (FREQ), 2574
 - equivalence tests (FREQ), 2580
 - exact confidence limits (FREQ), 2580
 - FREQ procedure, 2572
 - noninferiority tests (FREQ), 2577
 - superiority tests (FREQ), 2579
 - tests (FREQ), 2576
 - TOST (FREQ), 2580
- risks, *see also* binomial proportions (FREQ)
 - FREQ procedure, 2572
- row mean scores statistic
 - Mantel-Haenszel (FREQ), 2594
- Somers' D statistics
 - FREQ procedure, 2556, 2559
- Spearman rank correlation coefficient
 - FREQ procedure, 2556, 2560
- standardized residuals
 - FREQ procedure, 2552
- stratified analysis
 - FREQ procedure, 2472, 2505
- Stuart's tau- c statistic
 - FREQ procedure, 2556, 2558
- superiority tests
 - binomial proportions (FREQ), 2570
 - risk difference (FREQ), 2579
- table scores
 - FREQ procedure, 2549
- tables
 - contingency (FREQ), 2472
 - crosstabulation (FREQ), 2472, 2611
 - multiway (FREQ), 2472, 2611
 - one-way frequency (FREQ), 2472, 2610
- Tarone's adjustment
 - Breslow-Day test (FREQ), 2598
- tetrachoric correlation coefficient
 - FREQ procedure, 2562
- TOST
 - equivalence tests (FREQ), 2571, 2580
- trend test
 - FREQ procedure, 2585
- uncertainty coefficients
 - FREQ procedure, 2556, 2564, 2565
- Wald confidence limits
 - risk difference (FREQ), 2576
- weighted kappa coefficient
 - FREQ procedure, 2588, 2590
- Wilson confidence limits
 - proportions (FREQ), 2567
- Yule's Q statistic
 - FREQ procedure, 2557
- Zelen's test
 - equal odds ratios (FREQ), 2598
- zeros, structural and random
 - agreement statistics (FREQ), 2592

Syntax Index

- AGREE option
 - EXACT statement (FREQ), 2489
 - OUTPUT statement (FREQ), 2497
 - TABLES statement (FREQ), 2508
 - TEST statement (FREQ), 2541
- AJCHI option
 - OUTPUT statement (FREQ), 2498
- ALL option
 - OUTPUT statement (FREQ), 2498
 - TABLES statement (FREQ), 2509
- ALPHA= option
 - EXACT statement (FREQ), 2493
 - TABLES statement (FREQ), 2509
- BARNARD option
 - EXACT statement (FREQ), 2489
- BDCHI option
 - OUTPUT statement (FREQ), 2498
- BDT option (CMH)
 - TABLES statement (FREQ), 2516
- BINOMIAL option
 - EXACT statement (FREQ), 2489
 - OUTPUT statement (FREQ), 2498
 - TABLES statement (FREQ), 2509
- BOWKER option
 - OUTPUT statement (FREQ), 2504
- BY statement
 - FREQ procedure, 2486
- CELLCHI2 option
 - TABLES statement (FREQ), 2513
- CHISQ option
 - EXACT statement (FREQ), 2489, 2635
 - OUTPUT statement (FREQ), 2498
 - TABLES statement (FREQ), 2513, 2635
- CL option
 - TABLES statement (FREQ), 2515
- CL= option (BINOMIAL)
 - TABLES statement (FREQ), 2510
- CL= option (RISKDIFF)
 - TABLES statement (FREQ), 2535
- CL=AGRESTICAFFO option (RISKDIFF)
 - TABLES statement (FREQ), 2536
- CL=AGRESTICOULL option (BINOMIAL)
 - TABLES statement (FREQ), 2510
- CL=ALL option (BINOMIAL)
 - TABLES statement (FREQ), 2510
- CL=CLOPPERPEARSON option (BINOMIAL)
 - TABLES statement (FREQ), 2511
- CL=EXACT option (BINOMIAL)
 - TABLES statement (FREQ), 2511
- CL=EXACT option (RISKDIFF)
 - TABLES statement (FREQ), 2536
- CL=HA option (RISKDIFF)
 - TABLES statement (FREQ), 2536
- CL=JEFFREYS option (BINOMIAL)
 - TABLES statement (FREQ), 2511
- CL=NEWCOMBE option (RISKDIFF)
 - TABLES statement (FREQ), 2536
- CL=WALD option (BINOMIAL)
 - TABLES statement (FREQ), 2511
- CL=WALD option (RISKDIFF)
 - TABLES statement (FREQ), 2536
- CL=WILSON option (BINOMIAL)
 - TABLES statement (FREQ), 2511
- CMH option
 - OUTPUT statement (FREQ), 2498
 - TABLES statement (FREQ), 2516
- CMH1 option
 - OUTPUT statement (FREQ), 2499
 - TABLES statement (FREQ), 2516
- CMH2 option
 - OUTPUT statement (FREQ), 2499
 - TABLES statement (FREQ), 2517
- CMHCOR option
 - OUTPUT statement (FREQ), 2499
- CMHGA option
 - OUTPUT statement (FREQ), 2499
- CMHRMS option
 - OUTPUT statement (FREQ), 2499
- COCHQ option
 - OUTPUT statement (FREQ), 2499
- COLUMN= option (REL RISK)
 - EXACT statement (FREQ), 2491
- COLUMN= option (RISKDIFF)
 - EXACT statement (FREQ), 2492
 - TABLES statement (FREQ), 2536
- COMOR option
 - EXACT statement (FREQ), 2489
 - OUTPUT statement (FREQ), 2501
- COMPRESS option
 - PROC FREQ statement, 2485
- CONTENTS= option
 - TABLES statement (FREQ), 2517
- CONTGY option
 - OUTPUT statement (FREQ), 2499
- CONVERGE= option (PLCORR)

TABLES statement (FREQ), 2517
 CORRECT option (BINOMIAL)
 TABLES statement (FREQ), 2511
 CORRECT option (RISKDIFF)
 TABLES statement (FREQ), 2537
 CRAMV option
 OUTPUT statement (FREQ), 2499
 CROSSLIST option
 TABLES statement (FREQ), 2517
 CUMCOL option
 TABLES statement (FREQ), 2518

 DATA= option
 PROC FREQ statement, 2485
 DEVIATION option
 TABLES statement (FREQ), 2518
 DF= option (CHISQ)
 TABLES statement (FREQ), 2513

 EQKAP option
 OUTPUT statement (FREQ), 2499
 EQOR option
 EXACT statement (FREQ), 2490
 OUTPUT statement (FREQ), 2499
 EQUAL option (RISKDIFF)
 TABLES statement (FREQ), 2537
 EQUIVALENCE option (BINOMIAL)
 TABLES statement (FREQ), 2511
 EQUIVALENCE option (RISKDIFF)
 TABLES statement (FREQ), 2537
 EQWKP option
 OUTPUT statement (FREQ), 2500
 EXACT option
 OUTPUT statement (FREQ), 2500
 TABLES statement (FREQ), 2518
 EXACT statement
 FREQ procedure, 2487
 EXPECTED option
 TABLES statement (FREQ), 2518

 FISHER option
 EXACT statement (FREQ), 2490
 OUTPUT statement (FREQ), 2500
 TABLES statement (FREQ), 2518
 FORMAT= option
 TABLES statement (FREQ), 2519
 FORMCHAR= option
 PROC FREQ statement, 2485
 FREQ procedure
 syntax, 2484
 FREQ procedure, BY statement, 2486
 FREQ procedure, EXACT statement, 2487
 AGREE option, 2489
 ALPHA= option, 2493
 BARNARD option, 2489
 BINOMIAL option, 2489
 CHISQ option, 2489, 2635
 COLUMN= option (REL RISK), 2491
 COLUMN= option (RISKDIFF), 2492
 COMOR option, 2489
 EQOR option, 2490
 FISHER option, 2490
 JT option, 2490
 KAPPA option, 2490
 KENTB option, 2490
 LRCHI option, 2490
 MAXTIME= option, 2493
 MC option, 2493
 MCNEM option, 2490
 MEASURES option, 2490
 METHOD= option (REL RISK), 2491
 METHOD= option (RISKDIFF), 2492
 MHCHI option, 2491
 N= option, 2494
 OR option, 2491, 2635
 PCHI option, 2491
 PCORR option, 2491
 POINT option, 2494
 REL RISK option, 2491
 RISKDIFF option, 2492
 SEED= option, 2494
 SMDCR option, 2492
 SMDRC option, 2492
 STUTC option, 2492
 TREND option, 2493, 2643
 WTKAP option, 2493
 ZELEN option, 2490
 FREQ procedure, OUTPUT statement, 2494
 AGREE option, 2497
 AJCHI option, 2498
 ALL option, 2498
 BDCHI option, 2498
 BINOMIAL option, 2498
 BOWKER option, 2504
 CHISQ option, 2498
 CMH option, 2498
 CMH1 option, 2499
 CMH2 option, 2499
 CMHCOR option, 2499
 CMHGA option, 2499
 CMHRMS option, 2499
 COCHQ option, 2499
 COMOR option, 2501
 CONTGY option, 2499
 CRAMV option, 2499
 EQKAP option, 2499
 EQOR option, 2499
 EQWKP option, 2500
 EXACT option, 2500

- FISHER option, 2500
- GAILSIMON option, 2500
- GAMMA option, 2500
- JT option, 2500
- KAPPA option, 2500
- KENTB option, 2500
- LAMCR option, 2500
- LAMDAS option, 2500
- LAMRC option, 2500
- LGOR option, 2501
- LGRRC1 option, 2501
- LGRRC2 option, 2501
- LRCHI option, 2501
- MCNEM option, 2501
- MEASURES option, 2501
- MHCHI option, 2501
- MHOR option, 2501
- MHRRC1 option, 2501
- MHRRC2 option, 2502
- N option, 2502
- NMISS option, 2502
- OR option, 2502
- OUT= option, 2495
- output-options, 2495
- PCHI option, 2502
- PCORR option, 2502
- PLCORR option, 2502
- RDIF1 option, 2502
- RDIF2 option, 2503
- RELRISK option, 2503
- RISK1 option, 2503
- RISK11 option, 2503
- RISK12 option, 2504
- RISK2 option, 2504
- RISK21 option, 2504
- RISK22 option, 2504
- RISKDIFF option, 2503
- RISKDIFF1 option, 2503
- RISKDIFF2 option, 2503
- RRC1 option, 2503
- RRC2 option, 2503
- SCORR option, 2504
- SMDCR option, 2504
- SMDRC option, 2504
- STUTC option, 2504
- TAUB option, 2500
- TAUC option, 2504
- TREND option, 2504
- TSYMM option, 2504
- U option, 2505
- UCR option, 2505
- URC option, 2505
- WTKAP option, 2505
- ZELEN option, 2499

- FREQ procedure, PROC FREQ statement, 2484
 - COMPRESS option, 2485
 - DATA= option, 2485
 - FORMCHAR= option, 2485
 - NLEVELS option, 2486
 - NOPRINT option, 2486
 - ORDER= option, 2486
 - PAGE option, 2486
- FREQ procedure, TABLES statement, 2505
 - AGREE option, 2508
 - ALL option, 2509
 - ALPHA= option, 2509
 - BDT option (CMH), 2516
 - BINOMIAL option, 2509
 - CELLCHI2 option, 2513
 - CHISQ option, 2513, 2635
 - CL option, 2515
 - CL= option (BINOMIAL), 2510
 - CL= option (RISKDIFF), 2535
 - CL=AGRESTICAFFO option (RISKDIFF), 2536
 - CL=AGRESTICOULL option (BINOMIAL), 2510
 - CL=ALL option (BINOMIAL), 2510
 - CL=CLOPPERPEARSON option (BINOMIAL), 2511
 - CL=EXACT option (BINOMIAL), 2511
 - CL=EXACT option (RISKDIFF), 2536
 - CL=HA option (RISKDIFF), 2536
 - CL=JEFFREYS option (BINOMIAL), 2511
 - CL=NEWCOMBE option (RISKDIFF), 2536
 - CL=WALD option (BINOMIAL), 2511
 - CL=WALD option (RISKDIFF), 2536
 - CL=WILSON option (BINOMIAL), 2511
 - CMH option, 2516
 - CMH1 option, 2516
 - CMH2 option, 2517
 - COLUMN= option (RISKDIFF), 2536
 - CONTENTS= option, 2517
 - CONVERGE= option (PLCORR), 2517
 - CORRECT option (BINOMIAL), 2511
 - CORRECT option (RISKDIFF), 2537
 - CROSSLIST option, 2517
 - CUMCOL option, 2518
 - DEVIATION option, 2518
 - DF= option (CHISQ), 2513
 - EQUAL option (RISKDIFF), 2537
 - EQUIVALENCE option (BINOMIAL), 2511
 - EQUIVALENCE option (RISKDIFF), 2537
 - EXACT option, 2518
 - EXPECTED option, 2518
 - FISHER option, 2518
 - FORMAT= option, 2519
 - GAILSIMON option, 2519

JT option, 2519
 LEVEL= option (BINOMIAL), 2511
 LIST option, 2519
 LRCHISQ option (CHISQ), 2514
 MANTELLEIS option (CMH), 2516
 MARGIN= option (BINOMIAL), 2512
 MARGIN= option (RISKDIFF), 2537
 MAXITER= option (PLCORR), 2519
 MAXLEVELS= option, 2519
 MEASURES option, 2520
 METHOD= option (RISKDIFF), 2537
 METHOD=FM option (RISKDIFF), 2538
 METHOD=HA option (RISKDIFF), 2538
 METHOD=NEWCOMBE option (RISKDIFF), 2538
 METHOD=WALD option (RISKDIFF), 2538
 MISSING option, 2520
 MISSPRINT option, 2520
 NOCOL option, 2521
 NOCUM option, 2521
 NOFREQ option, 2521
 NONINFERIORITY option (BINOMIAL), 2512
 NONINFERIORITY option (RISKDIFF), 2538
 NOPERCENT option, 2521
 NOPRINT option, 2521
 NORISKS option (RISKDIFF), 2538
 NOROW option, 2521
 NOSPARSE option, 2521
 NOWARN option, 2521
 OR option, 2534
 OUT= option, 2522
 OUTCUM option, 2522
 OUTEXPECT option, 2522, 2623
 OUTPCT option, 2522
 P= option (BINOMIAL), 2512
 PLCORR option, 2522
 PLOTS= option, 2523
 PLOTS=AGREEPLOT option, 2524
 PLOTS=CUMFREQPLOT option, 2525
 PLOTS=DEVIATIONPLOT option, 2525
 PLOTS=FREQPLOT option, 2525
 PLOTS=KAPPAPLOT option, 2526
 PLOTS=MOSAICPLOT option, 2527
 PLOTS=NONE option, 2527
 PLOTS=ODDSRATIOPLOT option, 2527
 PLOTS=RELRIKSPLOT option, 2528
 PLOTS=RISKDIFFPLOT option, 2528
 PLOTS=WTKAPPAPLOT, 2529
 PRINTKWT option, 2534
 RELRISK option, 2534, 2635
 RISKDIFF option, 2534
 SCORES= option, 2538, 2646
 SCOROUT option, 2539
 SPARSE option, 2539, 2623

STDRES option (CROSSLIST), 2518
 SUPERIORITY option (BINOMIAL), 2513
 SUPERIORITY option (RISKDIFF), 2538
 TESTF= option, 2551
 TESTF= option (CHISQ), 2514
 TESTP= option, 2551, 2629
 TESTP= option (CHISQ), 2514
 TOTPCT option, 2539
 TREND option, 2539, 2643
 VAR= option (BINOMIAL), 2513
 VAR= option (RISKDIFF), 2538
 WARN= option (CHISQ), 2515
 FREQ procedure, TEST statement, 2540
 AGREE option, 2541
 GAMMA option, 2541
 KAPPA option, 2541
 KENTB option, 2541
 MEASURES option, 2541
 PCORR option, 2541
 PLCORR option, 2542
 SCORR option, 2542
 SMDCR option, 2542, 2643
 SMDRC option, 2542
 STUTC option, 2542
 TAUB option, 2541
 TAUC option, 2542
 WTKAP option, 2542
 FREQ procedure, WEIGHT statement, 2543
 ZEROS option, 2543
 GAILSIMON option
 OUTPUT statement (FREQ), 2500
 TABLES statement (FREQ), 2519
 GAMMA option
 OUTPUT statement (FREQ), 2500
 TEST statement (FREQ), 2541
 JT option
 EXACT statement (FREQ), 2490
 OUTPUT statement (FREQ), 2500
 TABLES statement (FREQ), 2519
 KAPPA option
 EXACT statement (FREQ), 2490
 OUTPUT statement (FREQ), 2500
 TEST statement (FREQ), 2541
 KENTB option
 EXACT statement (FREQ), 2490
 OUTPUT statement (FREQ), 2500
 TEST statement (FREQ), 2541
 LAMCR option
 OUTPUT statement (FREQ), 2500
 LAMDAS option
 OUTPUT statement (FREQ), 2500

- LAMRC option
 - OUTPUT statement (FREQ), 2500
- LEVEL= option (BINOMIAL)
 - TABLES statement (FREQ), 2511
- LGOR option
 - OUTPUT statement (FREQ), 2501
- LGRRC1 option
 - OUTPUT statement (FREQ), 2501
- LGRRC2 option
 - OUTPUT statement (FREQ), 2501
- LIST option
 - TABLES statement (FREQ), 2519
- LRCHI option
 - EXACT statement (FREQ), 2490
 - OUTPUT statement (FREQ), 2501
- LRCHISQ option (CHISQ)
 - TABLES statement (FREQ), 2514
- MANTELFLEISS option (CMH)
 - TABLES statement (FREQ), 2516
- MARGIN= option (BINOMIAL)
 - TABLES statement (FREQ), 2512
- MARGIN= option (RISKDIFF)
 - TABLES statement (FREQ), 2537
- MAXITER= option (PLCORR)
 - TABLES statement (FREQ), 2519
- MAXLEVELS= option
 - TABLES statement (FREQ), 2519
- MAXTIME= option
 - EXACT statement (FREQ), 2493
- MC option
 - EXACT statement (FREQ), 2493
- MCNEM option
 - EXACT statement (FREQ), 2490
 - OUTPUT statement (FREQ), 2501
- MEASURES option
 - EXACT statement (FREQ), 2490
 - OUTPUT statement (FREQ), 2501
 - TABLES statement (FREQ), 2520
 - TEST statement (FREQ), 2541
- METHOD= option (RELRISK)
 - EXACT statement (FREQ), 2491
- METHOD= option (RISKDIFF)
 - EXACT statement (FREQ), 2492
 - TABLES statement (FREQ), 2537
- METHOD=FM option (RISKDIFF)
 - TABLES statement (FREQ), 2538
- METHOD=HA option (RISKDIFF)
 - TABLES statement (FREQ), 2538
- METHOD=NEWCOMBE option (RISKDIFF)
 - TABLES statement (FREQ), 2538
- METHOD=WALD option (RISKDIFF)
 - TABLES statement (FREQ), 2538
- MHCHI option
 - EXACT statement (FREQ), 2491
 - OUTPUT statement (FREQ), 2501
- MHOR option
 - OUTPUT statement (FREQ), 2501
- MHRRC1 option
 - OUTPUT statement (FREQ), 2501
- MHRRC2 option
 - OUTPUT statement (FREQ), 2502
- MISSING option
 - TABLES statement (FREQ), 2520
- MISSPRINT option
 - TABLES statement (FREQ), 2520
- N option
 - OUTPUT statement (FREQ), 2502
- N= option
 - EXACT statement (FREQ), 2494
- NLEVELS option
 - PROC FREQ statement, 2486
- NMISS option
 - OUTPUT statement (FREQ), 2502
- NOCOL option
 - TABLES statement (FREQ), 2521
- NOCUM option
 - TABLES statement (FREQ), 2521
- NOFREQ option
 - TABLES statement (FREQ), 2521
- NONINFERIORITY option (BINOMIAL)
 - TABLES statement (FREQ), 2512
- NONINFERIORITY option (RISKDIFF)
 - TABLES statement (FREQ), 2538
- NOPERCENT option
 - TABLES statement (FREQ), 2521
- NOPRINT option
 - PROC FREQ statement, 2486
 - TABLES statement (FREQ), 2521
- NORISKS option (RISKDIFF)
 - TABLES statement (FREQ), 2538
- NOROW option
 - TABLES statement (FREQ), 2521
- NOSPARSE option
 - TABLES statement (FREQ), 2521
- NOWARN option
 - TABLES statement (FREQ), 2521
- OR option
 - EXACT statement (FREQ), 2491, 2635
 - OUTPUT statement (FREQ), 2502
 - TABLES statement (FREQ), 2534
- ORDER= option
 - PROC FREQ statement, 2486
- OUT= option
 - OUTPUT statement (FREQ), 2495
 - TABLES statement (FREQ), 2522

OUTCUM option
 TABLES statement (FREQ), 2522
 OUTEXPECT option
 TABLES statement (FREQ), 2522, 2623
 OUTPCT option
 TABLES statement (FREQ), 2522
 OUTPUT statement
 FREQ procedure, 2494

 P= option (BINOMIAL)
 TABLES statement (FREQ), 2512
 PAGE option
 PROC FREQ statement, 2486
 PCHI option
 EXACT statement (FREQ), 2491
 OUTPUT statement (FREQ), 2502
 PCORR option
 EXACT statement (FREQ), 2491
 OUTPUT statement (FREQ), 2502
 TEST statement (FREQ), 2541
 PLCORR option
 OUTPUT statement (FREQ), 2502
 TABLES statement (FREQ), 2522
 TEST statement (FREQ), 2542
 PLOTS= option
 TABLES statement (FREQ), 2523
 PLOTS=AGREEPLOT option
 TABLES statement (FREQ), 2524
 PLOTS=CUMFREQPLOT option
 TABLES statement (FREQ), 2525
 PLOTS=DEVIATIONPLOT option
 TABLES statement (FREQ), 2525
 PLOTS=FREQPLOT option
 TABLES statement (FREQ), 2525
 PLOTS=KAPPAPLOT option
 TABLES statement (FREQ), 2526
 PLOTS=MOSAICPLOT option
 TABLES statement (FREQ), 2527
 PLOTS=NONE option
 TABLES statement (FREQ), 2527
 PLOTS=ODDSRATIOPLOT option
 TABLES statement (FREQ), 2527
 PLOTS=RELRISKPLOT option
 TABLES statement (FREQ), 2528
 PLOTS=RISKDIFFPLOT option
 TABLES statement (FREQ), 2528
 PLOTS=WTKAPPAPLOT option
 TABLES statement (FREQ), 2529
 POINT option
 EXACT statement (FREQ), 2494
 PRINTKWT option
 TABLES statement (FREQ), 2534
 PROC FREQ statement, *see* FREQ procedure

 RDIF1 option
 OUTPUT statement (FREQ), 2502
 RDIF2 option
 OUTPUT statement (FREQ), 2503
 RELRISK option
 EXACT statement (FREQ), 2491
 OUTPUT statement (FREQ), 2503
 TABLES statement (FREQ), 2534, 2635
 RISK1 option
 OUTPUT statement (FREQ), 2503
 RISK11 option
 OUTPUT statement (FREQ), 2503
 RISK12 option
 OUTPUT statement (FREQ), 2504
 RISK2 option
 OUTPUT statement (FREQ), 2504
 RISK21 option
 OUTPUT statement (FREQ), 2504
 RISK22 option
 OUTPUT statement (FREQ), 2504
 RISKDIFF option
 EXACT statement (FREQ), 2492
 OUTPUT statement (FREQ), 2503
 TABLES statement (FREQ), 2534
 RISKDIFF1 option
 OUTPUT statement (FREQ), 2503
 RISKDIFF2 option
 OUTPUT statement (FREQ), 2503
 RRC1 option
 OUTPUT statement (FREQ), 2503
 RRC2 option
 OUTPUT statement (FREQ), 2503

 SCORES= option
 TABLES statement (FREQ), 2538, 2646
 SCOROUT option
 TABLES statement (FREQ), 2539
 SCORR option
 OUTPUT statement (FREQ), 2504
 TEST statement (FREQ), 2542
 SEED= option
 EXACT statement (FREQ), 2494
 SMDCR option
 EXACT statement (FREQ), 2492
 OUTPUT statement (FREQ), 2504
 TEST statement (FREQ), 2542, 2643
 SMDRC option
 EXACT statement (FREQ), 2492
 OUTPUT statement (FREQ), 2504
 TEST statement (FREQ), 2542
 SPARSE option
 TABLES statement (FREQ), 2539, 2623
 STDRES option (CROSSLIST)
 TABLES statement (FREQ), 2518
 STUTC option

- EXACT statement (FREQ), 2492
- OUTPUT statement (FREQ), 2504
- TEST statement (FREQ), 2542
- SUPERIORITY option (BINOMIAL)
 - TABLES statement (FREQ), 2513
- SUPERIORITY option (RISKDIFF)
 - TABLES statement (FREQ), 2538
- TABLES statement
 - FREQ procedure, 2505
- TAUB option
 - OUTPUT statement (FREQ), 2500
 - TEST statement (FREQ), 2541
- TAUC option
 - OUTPUT statement (FREQ), 2504
 - TEST statement (FREQ), 2542
- TEST statement
 - FREQ procedure, 2540
- TESTF= option
 - TABLES statement (FREQ), 2551
- TESTF= option (CHISQ)
 - TABLES statement (FREQ), 2514
- TESTP= option
 - TABLES statement (FREQ), 2551, 2629
- TESTP= option (CHISQ)
 - TABLES statement (FREQ), 2514
- TOTPCT option
 - TABLES statement (FREQ), 2539
- TREND option
 - EXACT statement (FREQ), 2493, 2643
 - OUTPUT statement (FREQ), 2504
 - TABLES statement (FREQ), 2539, 2643
- TSYMM option
 - OUTPUT statement (FREQ), 2504
- U option
 - OUTPUT statement (FREQ), 2505
- UCR option
 - OUTPUT statement (FREQ), 2505
- URC option
 - OUTPUT statement (FREQ), 2505
- VAR= option (BINOMIAL)
 - TABLES statement (FREQ), 2513
- VAR= option (RISKDIFF)
 - TABLES statement (FREQ), 2538
- WARN= option (CHISQ)
 - TABLES statement (FREQ), 2515
- WEIGHT statement
 - FREQ procedure, 2543
- WTKAP option
 - EXACT statement (FREQ), 2493
 - OUTPUT statement (FREQ), 2505
 - TEST statement (FREQ), 2542

- ZELLEN option
 - EXACT statement (FREQ), 2490
 - OUTPUT statement (FREQ), 2499
- ZEROS option
 - WEIGHT statement (FREQ), 2543

Your Turn

We welcome your feedback.

- If you have comments about this book, please send them to **`yourturn@sas.com`**. Include the full title and page numbers (if applicable).
- If you have comments about the software, please send them to **`suggest@sas.com`**.

SAS® Publishing Delivers!

Whether you are new to the work force or an experienced professional, you need to distinguish yourself in this rapidly changing and competitive job market. SAS® Publishing provides you with a wide range of resources to help you set yourself apart. Visit us online at support.sas.com/bookstore.

SAS® Press

Need to learn the basics? Struggling with a programming problem? You'll find the expert answers that you need in example-rich books from SAS Press. Written by experienced SAS professionals from around the world, SAS Press books deliver real-world insights on a broad range of topics for all skill levels.

support.sas.com/saspress

SAS® Documentation

To successfully implement applications using SAS software, companies in every industry and on every continent all turn to the one source for accurate, timely, and reliable information: SAS documentation. We currently produce the following types of reference documentation to improve your work experience:

- Online help that is built into the software.
- Tutorials that are integrated into the product.
- Reference documentation delivered in HTML and PDF – **free** on the Web.
- Hard-copy books.

support.sas.com/publishing

SAS® Publishing News

Subscribe to SAS Publishing News to receive up-to-date information about all new SAS titles, author podcasts, and new Web site features via e-mail. Complete instructions on how to subscribe, as well as access to past issues, are available at our Web site.

support.sas.com/spn



**THE
POWER
TO KNOW®**

