SAS/STAT 15.3®
User's Guide

The FREQ Procedure

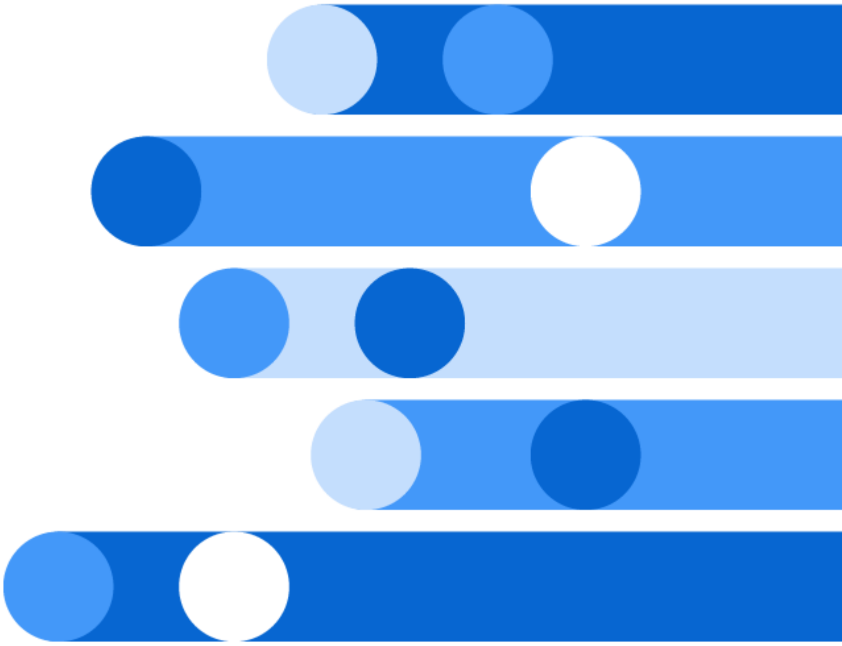

SAS® Documentation
February 26, 2025

This document is an individual chapter from *SAS/STAT® 15.3 User's Guide*.

The correct bibliographic citation for this manual is as follows: SAS Institute Inc. 2023. *SAS/STAT® 15.3 User's Guide*. Cary, NC: SAS Institute Inc.

SAS/STAT® 15.3 User's Guide

Copyright © 2023, SAS Institute Inc., Cary, NC, USA

All Rights Reserved. Produced in the United States of America.

For a hard-copy book: No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, or otherwise, without the prior written permission of the publisher, SAS Institute Inc.

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.

February 2025

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 trademarks of their respective companies.

SAS software may be provided with certain third-party software, including but not limited to open source software, which is licensed under its applicable third-party software license agreement. For license information about third-party software distributed with SAS software, refer to Third-Party Software Reference | SAS Support.

Chapter 47

The FREQ Procedure

Contents

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 within strata and across strata.

For one-way frequency tables, PROC FREQ provides 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, PROC FREQ computes chi-square tests. 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 stratification variables by computing statistics within and across 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, you can use PROC SURVEYFREQ. For more information, see Chapter 15, “Introduction to Survey Procedures.” When you want to fit models to categorical data, you can use a procedure such as CATMOD, GENMOD, GLIMMIX, LOGISTIC, PROBIT, or SURVEYLOGISTIC. For more information, see Chapter 9, “Introduction to Categorical Data Analysis Procedures.”

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 3240 for more information.

PROC FREQ uses ODS Graphics to create graphs as part of its output. For general information about ODS Graphics, see Chapter 24, “Statistical Graphics Using ODS.” For information about the statistical graphics that PROC FREQ produces, see the PLOTS= option in the TABLES statement and the section “ODS Graphics” on page 3244.

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 47.1 contains counts of the students who enrolled in the summer program by gender and whether they were assigned an internship slot.

Table 47.1 Summer Enrichment Data

		Enrollment		
Gender	**Internship**	**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=FREQ, 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 47.1 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 47.1 Crosstabulation Table

The FREQ Procedure

Frequency
Percent
Row Pct
Col Pct

Table of Internship by Enrollment

	Enrollment		
Internship	**yes**	**no**	**Total**
yes	67 30.04 63.21 50.00	39 17.49 36.79 43.82	106 47.53
no	67 30.04 57.26 50.00	50 22.42 42.74 56.18	117 52.47
Total	134 60.09	89 39.91	223 100.00

Figure 47.2 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 47.2 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 47.3, Figure 47.4, and Figure 47.5 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 47.4 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 47.6, Figure 47.7, and Figure 47.8 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 47.3 Crosstabulation Table for Boys

```
                     The FREQ Procedure

Frequency   Table 1 of Internship by Enrollment
Percent          Controlling for Gender=boys
Row Pct                     Enrollment
Col Pct     Internship        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 47.4 Frequency Plot for Boys

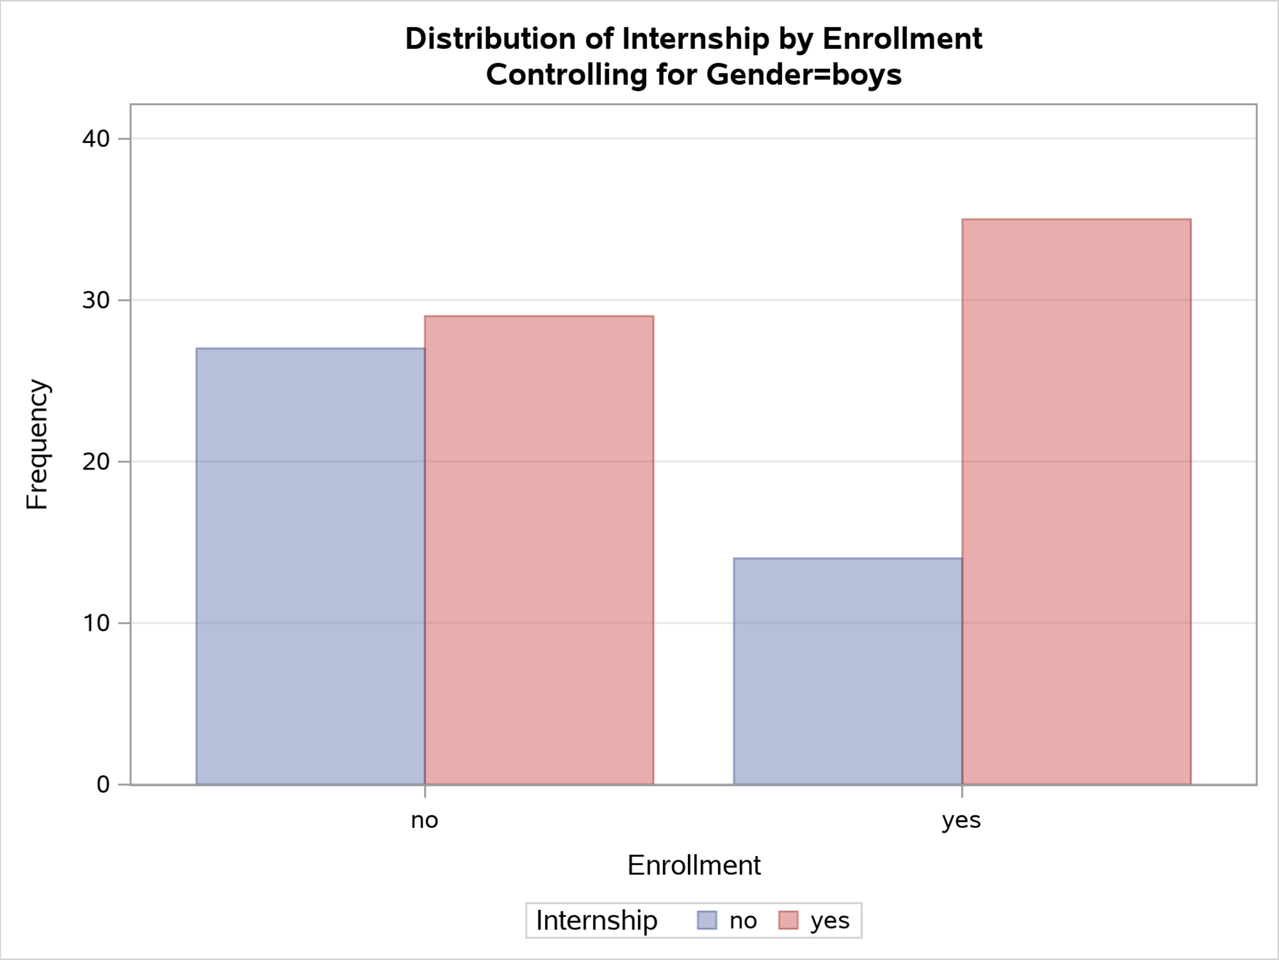

Figure 47.5 Chi-Square Statistics for Boys

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 47.6 Crosstabulation Table for Girls

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

Figure 47.7 Frequency Plot for Girls

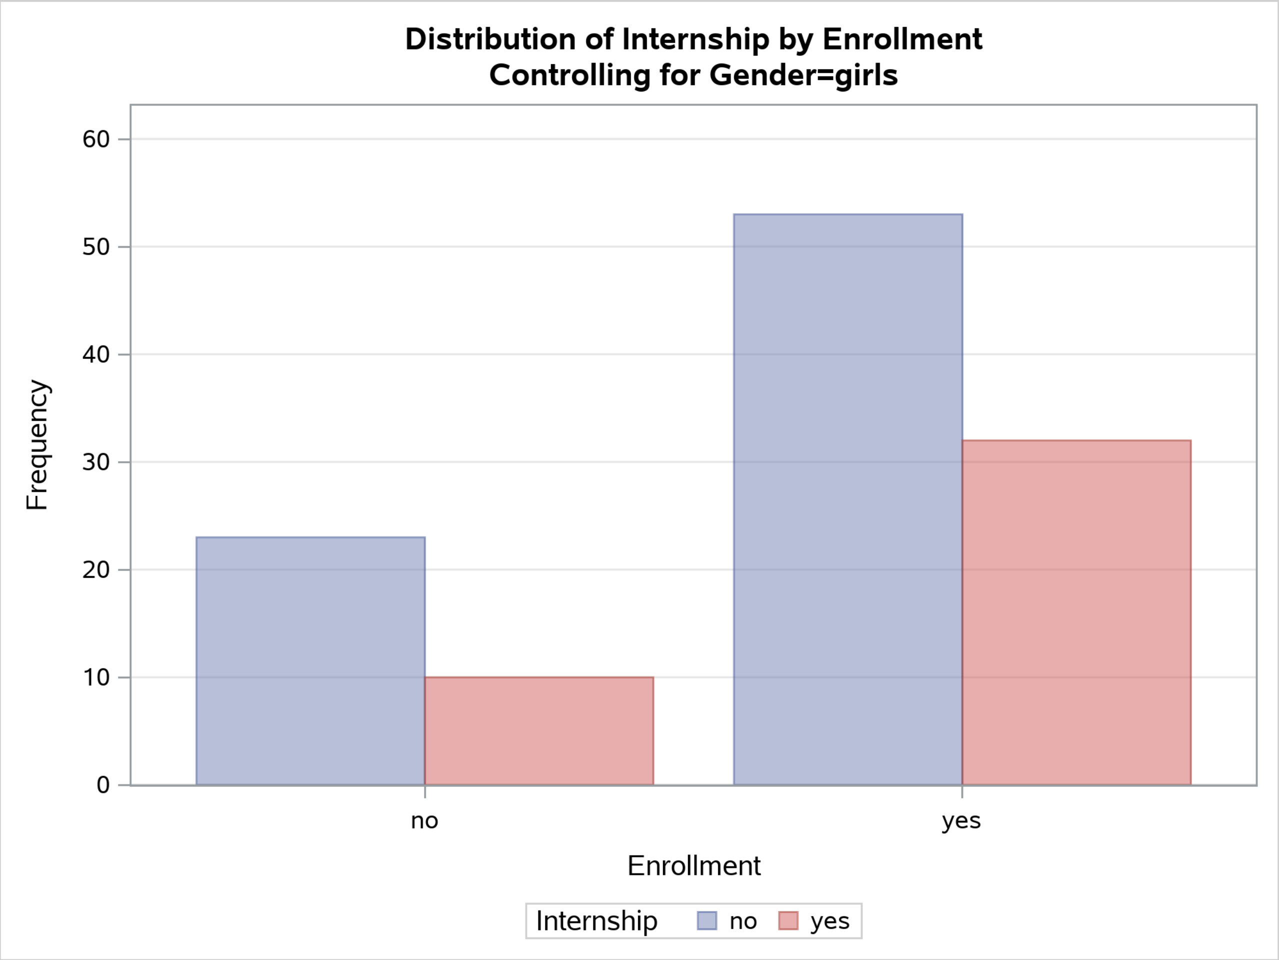

Figure 47.8 Chi-Square Statistics for Girls

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 47.9 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 47.9 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 47.2 contains the data.

Table 47.2 Skin Condition Data

	Dermatologist 2			
Dermatologist 1	**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 data set 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 is 0, which indicates 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 47.10 and Figure 47.11 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 0. The AGREE option also produces Bowker's symmetry test and the weighted kappa coefficient, but that output is not shown here. Figure 47.11 displays the agreement plot for the ratings of the two dermatologists.

Figure 47.10 Agreement Study

The FREQ Procedure

Statistics for Table of Derm1 by Derm2

Kappa Statistics

Statistic	Estimate	Standard Error	95% Confidence Limits	
Simple Kappa	0.3449	0.0724	0.2030	0.4868
Weighted Kappa	0.5082	0.0655	0.3798	0.6366

Test of H0: Kappa = 0

Estimate	H0 Std Err	Z	Pr > Z	Pr > |Z|
0.3449	0.0612	5.6366	<.0001	<.0001

Figure 47.11 Agreement Plot

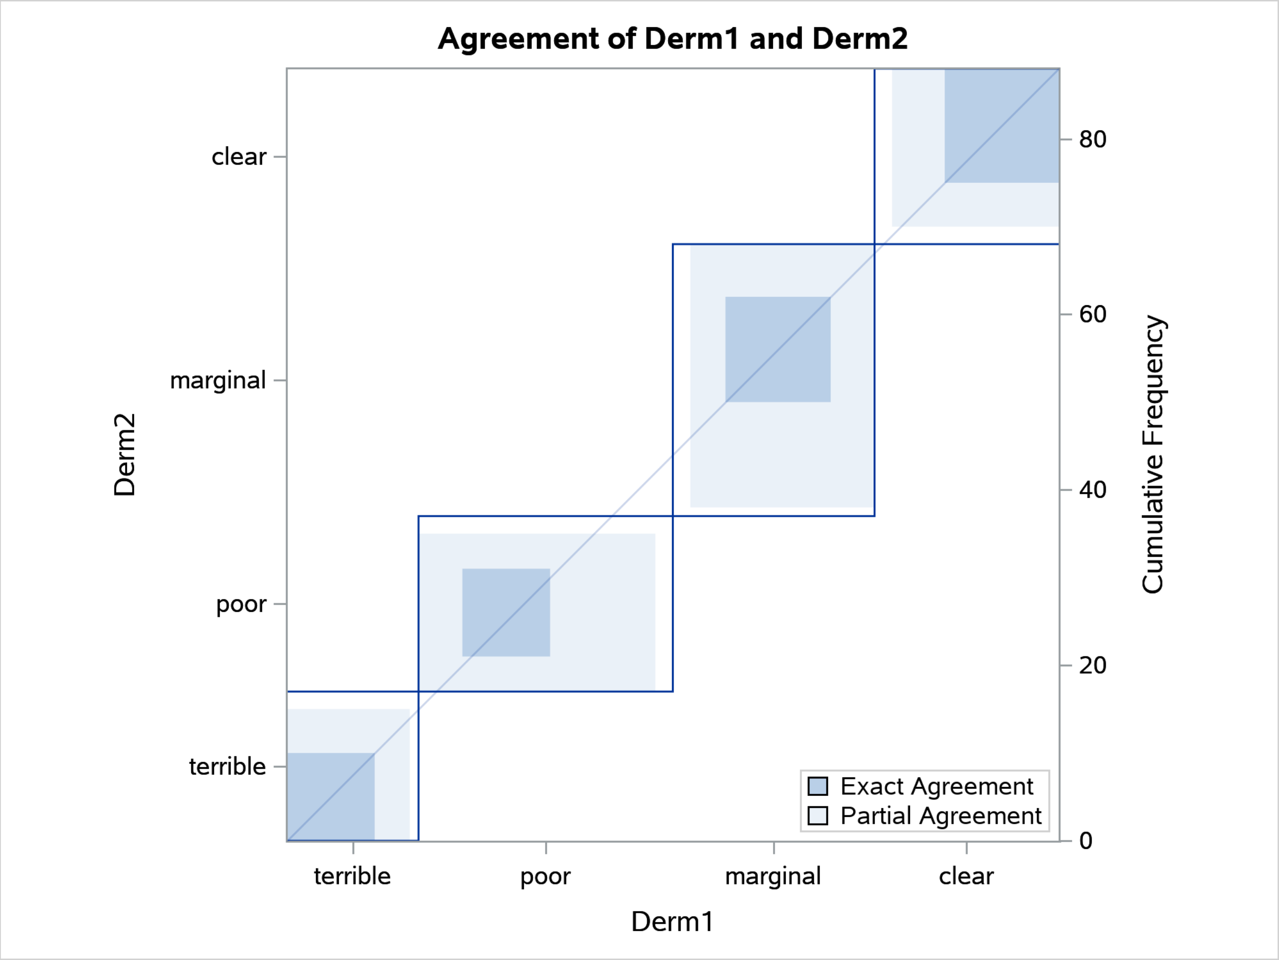

Syntax: FREQ Procedure

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;
```

Table 47.3 summarizes the basic functions of the procedure statements. The following sections provide 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 47.3 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 47.4 lists the *options* available in the PROC FREQ statement. Descriptions of the *options* follow in alphabetical order.

Table 47.4 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*:

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 use to construct cell outlines and dividers for crosstabulation tables. This option affects only the SAS monospace output destination.

PROC FREQ uses 3 of the 20 formatting characters that SAS provides. You can specify three characters in *formchar-string* to draw the vertical separators (1), the horizontal separators (2), and the vertical-horizontal intersections (7) in crosstabulation tables. By default, PROC FREQ uses FORMCHAR(1,2,7)='|-+'. Table 47.5 summarizes the formatting characters that PROC FREQ uses.

Table 47.5 Formatting Characters Used by PROC FREQ

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

To produce crosstabulation tables that have no cell outlines or dividers, you can specify all blanks in *formchar-string*—for example, FORMCHAR(1,2,7)='   ' (three blanks).

You can use any characters in *formchar-string*. If you use hexadecimal characters, you must put **x** after the closing quotation mark. For information about which hexadecimal codes to use for which characters, see the documentation for your hardware.

For more information about formatting characters, see the TABULATE procedure in the *Base SAS Procedures Guide*.

NLEVELS

displays the "Number of Variable Levels" table, which provides the number of levels for each variable named in the TABLES statements. For more information, see the section "Number of Variable Levels Table" on page 3230. PROC FREQ determines the variable levels from the formatted variable values, as described in the section "Grouping with Formats" on page 3147.

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 3227 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 23, "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 Programmers Guide: Essentials*.

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 in 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 in 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. For more information, see the section "Exact Statistics" on page 3222.

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 large 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. For more information, see the section "Computational Resources" on page 3224.

Statistic Options

The *statistic-options* specify which exact tests and confidence limits to compute. Table 47.6 lists the available *statistic-options* and the exact statistics that are computed. Descriptions of the *statistic-options* follow Table 47.6 in alphabetical order.

For one-way tables, PROC FREQ provides exact p-values for the binomial proportion test, the chi-square goodness-of-fit test, and the likelihood ratio chi-square test. PROC FREQ also provides exact (Clopper-Pearson) confidence limits 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, Cochran-Armitage test for trend, and the symmetry test. PROC FREQ also provides exact p-values 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 (proportion) 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 47.6 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 47.6.

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 47.6 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 | BIN	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 | ZELEN	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 | TAUB	Test for Kendall's tau-b
LRCHI	Likelihood ratio chi-square test (one-way and 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 | ODDSRATIO	Confidence limits for the odds ratio (for 2 × 2 tables)
PCHI	Pearson chi-square test (one-way and two-way tables)
PCORR	Test for the Pearson correlation coefficient
RELRISK	Confidence limits for the relative risk (for 2 × 2 tables)
RISKDIFF	Confidence limits for the risk difference (for 2 × 2 tables)
SCORR	Test for the Spearman correlation coefficient
SMDCR	Test for Somers' $D(C\|R)$
SMDRC	Test for Somers' $D(R\|C)$
STUTC | TAUC	Test for Stuart's tau-c
SYMMETRY | BOWKER	Symmetry test
TREND	Cochran-Armitage test for trend
WTKAPPA | WTKAP	Test for the weighted kappa coefficient

You can specify the following *statistic-options*:

AGREE

requests McNemar's exact test, an exact test for the simple kappa coefficient, and an exact test for the weighted kappa coefficient. For more information, see the sections "Tests and Measures of Agreement" on page 3204 and "Exact Statistics" on page 3222.

For McNemar's test, you can specify the null hypothesis ratio of discordant proportions by using the AGREE(MNULLRATIO=) option in the TABLES statement; by default, MNULLRATIO=1. For the weighted kappa coefficient, you can request Fleiss-Cohen weights by specifying the AGREE(WT=FC) option in the TABLES statement; by default, PROC FREQ computes the weighted kappa coefficient by using Cicchetti-Allison agreement weights.

McNemar's test is available for 2 × 2 tables. Kappa coefficients are defined only for square two-way tables, where the number of rows equals the number of columns. If your table is not square because some observations have weights of 0, you can specify the ZEROS option in the WEIGHT statement to include these observations in the analysis. For more information, see the section "Tables with Zero-Weight Rows or Columns" on page 3211.

For 2 × 2 tables, the weighted kappa coefficient is equivalent to the simple kappa coefficient, and PROC FREQ displays only analyses for the simple kappa coefficient.

BARNARD

requests Barnard's exact unconditional test for the risk (proportion) difference for 2 × 2 tables. For more information, see the section "Barnard's Unconditional Exact Test" on page 3186.

To request exact unconditional confidence limits for the risk difference. you can specify the RISKDIFF option in the EXACT statement. The RISKDIFF option in the TABLES statement provides asymptotic tests and several types of confidence limits for the risk difference. For more information, see the section "Risks and Risk Differences" on page 3176.

BINOMIAL

BIN

requests an exact test for the binomial proportion (for one-way tables). For more information, see the section "Binomial Tests" on page 3171. You can specify the null hypothesis proportion by using the BINOMIAL(P=) option in the TABLES statement; by default, P=0.5.

The BINOMIAL option in the TABLES statement provides exact (Clopper-Pearson) confidence limits for the binomial proportion by default. You can specify the BINOMIAL(CL=MIDP) option in the TABLES statement to request exact mid-p confidence limits for the binomial proportion. The BINOMIAL option in the TABLES statement also provides asymptotic (Wald) tests and several other confidence limit types for the binomial proportion. For more information, see the section "Binomial Proportion" on page 3168.

CHISQ

requests the following exact chi-square tests for two-way tables: Pearson chi-square, likelihood ratio chi-square, and Mantel-Haenszel chi-square. For more information, see the section "Chi-Square Tests and Statistics" on page 3153. The CHISQ option in the TABLES statement provides asymptotic tests for these statistics.

For one-way tables, the CHISQ option requests an exact chi-square goodness-of-fit test. You can specify null hypothesis proportions for this test by using the CHISQ(TESTP=) option in the TABLES statement. By default, the one-way chi-square test is based on the null hypothesis of equal proportions. For more information, see the section "Chi-Square Test for One-Way Tables" on page 3153.

COMOR

requests an exact test and exact confidence limits for the common odds ratio for multiway 2 × 2 tables. For more information, see the section "Exact Confidence Limits for the Common Odds Ratio" on page 3219. The CMH option in the TABLES statement provides Mantel-Haenszel and logit estimates of the common odds ratio along with their asymptotic confidence limits.

EQOR
ZELEN

requests Zelen's exact test for equal odds ratios for multiway 2 × 2 tables. For more information, see the section "Zelen's Exact Test for Equal Odds Ratios" on page 3218. The CMH option in the TABLES statement provides an (asymptotic) Breslow-Day test for homogeneity of odds ratios.

FISHER

requests Fisher's exact test. For more information, see the sections "Fisher's Exact Test" on page 3156 and "Exact Statistics" on page 3222. 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 an exact Jonckheere-Terpstra test. For more information, see the sections "Jonckheere-Terpstra Test" on page 3202 and "Exact Statistics" on page 3222. The JT option in the TABLES statement provides an asymptotic Jonckheere-Terpstra test.

KAPPA

requests an exact test for the simple kappa coefficient. For more information, see the sections "Simple Kappa Coefficient" on page 3205 and "Exact Statistics" on page 3222. 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 an asymptotic test for the simple kappa coefficient.

Kappa coefficients are defined only for square two-way tables, where the number of rows equals the number of columns. If your table is not square because some observations have weights of 0, you can specify the ZEROS option in the WEIGHT statement to include these observations in the analysis. For more information, see the section "Tables with Zero-Weight Rows or Columns" on page 3211.

KENTB
TAUB

requests an exact test for Kendall's tau-b. For more information, see the sections "Kendall's Tau-b" on page 3160 and "Exact Statistics" on page 3222. The MEASURES option in the TABLES statement provides an estimate and standard error of Kendall's tau-b. 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. For more information, see the sections "Likelihood Ratio Chi-Square Test" on page 3155 and "Exact Statistics" on page 3222. The CHISQ option in the TABLES statement provides an asymptotic likelihood ratio chi-square test for two-way tables.

For one-way tables, the LRCHI option requests an exact likelihood ratio goodness-of-fit test. You can specify null hypothesis proportions by using the CHISQ(TESTP=) option in the TABLES statement. By default, the one-way test is based on the null hypothesis of equal proportions. For more information, see the section "Likelihood Ratio Chi-Square Test for One-Way Tables" on page 3155.

MCNEM

requests an exact McNemar's test. For more information, see the sections "McNemar's Test" on page 3204 and "Exact Statistics" on page 3222. You can specify the null hypothesis ratio of discordant proportions by using the AGREE(MNULLRATIO=) option in the TABLES statement; by default, MNULLRATIO=1. The AGREE option in the TABLES statement provides an asymptotic McNemar's test.

MEASURES

requests exact tests for the Pearson and Spearman correlations. For more information, see the sections "Pearson Correlation Coefficient" on page 3162, "Spearman Rank Correlation Coefficient" on page 3163, and "Exact Statistics" on page 3222. 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. For more information, see the subsection Exact Confidence Limits in the section "Confidence Limits for the Odds Ratio" on page 3192. You can also request exact confidence limits for the odds ratio by specifying the OR option in the EXACT statement.

MHCHI

requests an exact test for the Mantel-Haenszel chi-square. For more information, see the sections "Mantel-Haenszel Chi-Square Test" on page 3156 and "Exact Statistics" on page 3222. The CHISQ option in the TABLES statement provides an asymptotic Mantel-Haenszel chi-square test.

OR

ODDSRATIO

requests exact confidence limits for the odds ratio for 2 × 2 tables. For more information, see the subsection "Exact Confidence Limits" in the section "Confidence Limits for the Odds Ratio" on page 3192.

You can request exact mid-*p* confidence limits for the odds ratio by specifying the OR(CL=MIDP) option in the TABLES statement. The OR(CL=) option in the TABLES statement also provides other types of confidence limits for the odds ratio. For more information, see the section "Confidence Limits for the Odds Ratio" on page 3192.

The ALPHA= option in the TABLES statement determines the confidence level of the exact confidence limits; by default, ALPHA=0.05, which produces 95% confidence limits for the odds ratio.

PCHI

requests an exact test for the Pearson chi-square for two-way tables. For more information, see the sections "Pearson Chi-Square Test for Two-Way Tables" on page 3154 and "Exact Statistics" on page 3222. The CHISQ option in the TABLES statement provides an asymptotic Pearson chi-square test.

For one-way tables, the PCHI option requests an exact chi-square goodness-of-fit test. You can specify null hypothesis proportions by using the CHISQ(TESTP=) option in the TABLES statement. By default, the goodness-of-fit test is based on the null hypothesis of equal proportions. For more information, see the section "Chi-Square Test for One-Way Tables" on page 3153.

PCORR

requests an exact test for the Pearson correlation coefficient. For more information, see the sections "Pearson Correlation Coefficient" on page 3162 and "Exact Statistics" on page 3222. The MEASURES option in the TABLES statement provides the estimate and standard error of the Pearson correlation. 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. By default (beginning in SAS/STAT 14.3), the exact confidence limits are computed by inverting two separate one-sided exact tests that are based on the score statistic (Chan and Zhang 1999). For more information, see the subsection "Exact Unconditional Confidence Limits" in the section "Confidence Limits for the Relative Risk" on page 3195.

The RELRISK(CL=) option in the TABLES statement provides additional types of confidence limits for the relative risk. For more information, see the section "Confidence Limits for the Risk Difference" on page 3178.

The ALPHA= option in the TABLES statement determines the confidence level; by default, ALPHA=0.05, which produces 95% confidence limits for the relative risk.

You can specify the following *options*:

COLUMN=1 | 2 | BOTH

specifies the table column of the relative risk. By default, COLUMN=1, which provides exact confidence limits for the column 1 relative risk. COLUMN=BOTH provides exact confidence limits for both column 1 and column 2 relative risks.

METHOD=NOSCORE | SCORE | SCORE2

specifies the computation method for the exact confidence limits. By default, METHOD=SCORE.

You can specify one of the following methods:

NOSCORE

computes the exact confidence limits by inverting two separate one-sided exact tests that are based on the unstandardized relative risk (Santner and Snell 1980). For more information, see the subsection "Exact Unconditional Confidence Limits" in the section "Confidence Limits for the Relative Risk" on page 3195. This method is the default in releases before SAS/STAT 14.3.

SCORE

computes the exact confidence limits by inverting two separate one-sided exact tests that are based on the score statistic (Chan and Zhang 1999). For more information, see the subsection "Exact Unconditional Confidence Limits" in the section "Confidence Limits for the Relative Risk" on page 3195. This method is the default beginning in SAS/STAT 14.3.

SCORE2

computes the exact confidence limits by inverting a single two-sided exact test that is based on the score statistic (Agresti and Min 2001). For more information, see the subsection "Exact Unconditional Confidence Limits" in the section "Confidence Limits for the Relative Risk" on page 3195.

RISKDIFF < (*options*) >

requests exact unconditional confidence limits for the risk difference for 2 × 2 tables. By default (beginning in SAS/STAT 14.3), the exact confidence limits are computed by inverting two separate one-sided exact tests that are based on the score statistic (Chan and Zhang 1999). For more information, see the subsection "Exact Unconditional Confidence Limits" in the section "Confidence Limits for the Risk Difference" on page 3178.

The RISKDIFF(CL=) option in the TABLES statement provides additional types of confidence limits for the risk difference. For more information, see the section "Confidence Limits for the Risk Difference" on page 3178.

The ALPHA= option in the TABLES statement determines the confidence level; by default, ALPHA=0.05, which produces 95% confidence limits for the risk difference.

You can specify the following *options*:

COLUMN=1 | 2 | BOTH

specifies the table column of the risk difference. By default, COLUMN=BOTH and the exact confidence limits are displayed in the 'Risk Estimates' tables. If you specify the RISKDIFF(NORISKS) option in the TABLES statement to suppress the 'Risk Estimates' tables, COLUMN=1 by default and the exact confidence limits are displayed in the 'Risk Difference Confidence Limits' table.

METHOD=NOSCORE | SCORE | SCORE2

specifies the computation method for the exact confidence limits. By default, METHOD=SCORE.

You can specify one of the following methods:

NOSCORE

computes the exact confidence limits by inverting two separate one-sided exact tests that are based on the unstandardized risk difference (Santner and Snell 1980). For more information, see the subsection "Exact Unconditional Confidence Limits" in the section "Confidence Limits for the Risk Difference" on page 3178. This method is the default in releases before SAS/STAT 14.3.

SCORE

computes the exact confidence limits by inverting two separate one-sided exact tests that are based on the score statistic (Chan and Zhang 1999). For more information, see the subsection "Exact Unconditional Confidence Limits" in the section "Confidence Limits for the Risk Difference" on page 3178. This method is the default beginning in SAS/STAT 14.3.

SCORE2

computes the exact confidence limits by inverting a single two-sided exact test that is based on the score statistic (Agresti and Min 2001). For more information, see the subsection "Exact Unconditional Confidence Limits" in the section "Confidence Limits for the Risk Difference" on page 3178.

SCORR

requests an exact test for the Spearman correlation coefficient. For more information, see the sections "Spearman Rank Correlation Coefficient" on page 3163 and "Exact Statistics" on page 3222. The MEASURES option in the TABLES statement provides the estimate and standard error of the Spearman correlation. The SCORR option in the TEST statement provides an asymptotic test for the Spearman correlation.

SMDCR

requests an exact test for Somers' $D(C|R)$. For more information, see the sections "Somers' *D*" on page 3161 and "Exact Statistics" on page 3222. The MEASURES option in the TABLES statement provides the estimate and standard error of Somers' $D(C|R)$. The SMDCR option in the TEST statement provides an asymptotic test for Somers' $D(C|R)$.

SMDRC

requests an exact test for Somers' $D(R|C)$. For more information, see the sections "Somers' *D*" on page 3161 and "Exact Statistics" on page 3222. The MEASURES option in the TABLES statement provides the estimate and standard error of Somers' $D(R|C)$. The SMDRC option in the TEST statement provides an asymptotic test for Somers' $D(C|R)$.

STUTC

TAUC

requests an exact test for Stuart's tau-*c*. For more information, see the sections "Stuart's Tau-c" on page 3161 and "Exact Statistics" on page 3222. The MEASURES option in the TABLES statement provides the estimate and standard error of Stuart's tau-*c*. The STUTC option in the TEST statement provides an asymptotic test for Stuart's tau-*c*.

SYMMETRY

BOWKER

requests an exact symmetry test. This test is available for square $R \times R$ two-way tables where the table dimension R is greater than 2. For more information, see the section "Exact Symmetry Test" on page 3205. The AGREE option in the TABLES statement provides an asymptotic symmetry test.

TREND

requests the exact Cochran-Armitage test for trend. For more information, see the sections "Cochran-Armitage Test for Trend" on page 3201 and "Exact Statistics" on page 3222. The TREND option in the TABLES statement provides an asymptotic Cochran-Armitage test for trend. This test is available for tables of dimensions $2 \times C$ or $R \times 2$.

WTKAPPA

WTKAP

requests an exact test for the weighted kappa coefficient. For more information, see the sections "Weighted Kappa Coefficient" on page 3207 and "Exact Statistics" on page 3222. By default, PROC FREQ computes the weighted kappa coefficient by using Cicchetti-Allison agreement weights. You can request Fleiss-Cohen agreement weights by specifying the AGREE(WT=FC) option in the TABLES statement.

Kappa coefficients are defined only for square two-way tables, where the number of rows equals the number of columns. If your table is not square because some observations have weights of 0, you can specify the ZEROS option in the WEIGHT statement to include these observations in the analysis. For more information, see the section "Tables with Zero-Weight Rows or Columns" on page 3211.

For 2 × 2 tables, the weighted kappa coefficient is equivalent to the simple kappa coefficient, and PROC FREQ displays only analyses for the simple kappa coefficient.

Computation Options

The *computation-options* specify options for computing exact statistics. You can specify the following *computation-options*:

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; a confidence level of α produces $100(1-\alpha)\%$ confidence limits. By default ALPHA=0.01, which produces 99% confidence limits for the Monte Carlo estimates.

This 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 maximum time *value* must be a positive number. This option is available for exact p-value computation and for Monte Carlo estimation of exact p-values. For more information, see the section "Computational Resources" on page 3224.

MC

requests Monte Carlo estimation of exact p-values instead of direct exact p-value computation. Monte Carlo estimation can be useful for large problems where exact computations require a substantial amount of time and memory but asymptotic approximations might not be sufficient. For more information, see the section "Monte Carlo Estimation" on page 3225.

This 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 always computes exact tests or confidence limits (not Monte Carlo estimates) for these statistics.

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

MIDP

requests exact mid p-values for the exact tests. The exact mid p-value is defined as the exact p-value minus half the exact point probability. For more information, see the section "Definition of p-Values" on page 3223.

The MIDP option is available for all EXACT statement *statistic-options* except the following: BARNARD, EQOR, OR, RELRISK, and RISKDIFF. You cannot specify both the MIDP option and the MC option.

N=*n*

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

This option invokes the MC option.

PFORMAT=*format-name* | EXACT

specifies the display format for exact *p*-values. PROC FREQ applies this format to one- and two-sided exact *p*-values, exact point probabilities, and exact mid *p*-values. By default, PROC FREQ displays exact *p*-values in the PVALUE6.4 format.

You can provide a *format-name* or you can specify PFORMAT=EXACT to control the format of exact *p*-values. The value of *format-name* can be any standard SAS numeric format or a user-defined format. The format length must not exceed 24. For information about formats, see the FORMAT procedure in the *Base SAS Procedures Guide* and the FORMAT statement and SAS format in *SAS Formats and Informats: Reference*.

If you specify PFORMAT=EXACT, PROC FREQ uses the 6.4 format to display exact *p*-values that are greater than or equal to 0.001; the procedure uses the E10.3 format to display values that are between 0.000 and 0.001.

POINT

requests exact point probabilities for the exact tests. The exact point probability is the exact probability that the test statistic equals the observed value. For more information, see the section "Definition of *p*-Values" on page 3223.

The POINT option is available for all EXACT statement *statistic-options* except the following: BARNARD, EQOR, OR, RELRISK, and RISKDIFF. You cannot specify both the POINT option and 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 0, PROC FREQ uses the time of day from the computer's clock to obtain the initial seed.

This 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 47.7 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 3228.

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 3227 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 3240 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 DATAn convention.

output-options

specify the statistics to include in the output data set. Table 47.7 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. For more information, see the section "Contents of the OUTPUT Statement Output Data Set" on page 3228.

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 47.7 lists the TABLES statement options that are required to produce the OUTPUT data set statistics.

Table 47.7 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 Q ($h \times 2 \times 2$ tables)	AGREE
AJCHI	Continuity-adjusted chi-square (2×2 tables)	CHISQ

Table 47.7 *continued*

Output Option	Output Data Set Statistics	Required TABLES Statement Option
ALL	CHISQ, MEASURES, and CMH statistics; N (number of nonmissing observations)	ALL
BDCHI	Breslow-Day test ($h \times 2 \times 2$ tables)	CMH, CMH1, or CMH2
BINOMIAL | BIN	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 V	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
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]
GAMMA	Gamma	MEASURES
GS | GAILSIMON	Gail-Simon test	CMH(GAILSIMON)
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

---

[1] CHISQ 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 47.7 *continued*

Output Option	Output Data Set Statistics	Required TABLES Statement Option
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 | ODDSRATIO	Odds ratio (2 × 2 tables)	MEASURES, OR, or RELRISK
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 symmetry test	AGREE
U	Symmetric uncertainty coefficient	MEASURES
UCR	Uncertainty coefficient $(C\|R)$	MEASURES

Table 47.7 *continued*

Output Option	Output Data Set Statistics	Required TABLES Statement Option
URC	Uncertainty coefficient $(R\|C)$	MEASURES
WTKAPPA | WTKAP	Weighted kappa coefficient	AGREE

You can specify the following *output-options*:

AGREE

includes the following tests and measures of agreement in the output data set: McNemar's test (for 2×2 tables), Bowker's symmetry test, 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. For more information, see the section "Tests and Measures of Agreement" on page 3204.

AGREE statistics are computed only for square tables, where the number of rows equals the number of columns. PROC FREQ provides Bowker's symmetry test 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. For more information, see the section "Continuity-Adjusted Chi-Square Test" on page 3156.

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

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. For more information, see the section "Breslow-Day Test for Homogeneity of the Odds Ratios" on page 3217.

BINOMIAL

BIN

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. For more information, see the section "Binomial Proportion" on page 3168.

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. For more information, see the section "Chi-Square Tests and Statistics" on page 3153. For one-way tables, CHISQ includes the chi-square goodness-of-fit test in the output data set. For more information, see the section "Chi-Square Test for One-Way Tables" on page 3153. 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. For more information, see the section "Cochran-Mantel-Haenszel Statistics" on page 3211.

If you specify the CMH(MANTELFLEISS) 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. For more information, see the section "Cochran-Mantel-Haenszel Statistics" on page 3211.

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. For more information, see the section "Cochran-Mantel-Haenszel Statistics" on page 3211.

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. For more information, see the section "Correlation Statistic" on page 3213.

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. For more information, see the section "General Association Statistic" on page 3214.

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. For more information, see the section "ANOVA (Row Mean Scores) Statistic" on page 3213.

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. For more information, see the section "Cochran's Q Test" on page 3211.

CONTGY

includes the contingency coefficient in the output data set. The CHISQ option in the TABLES statement requests computation of the contingency coefficient. For more information, see the section "Contingency Coefficient" on page 3158.

CRAMV

includes Cramér's V in the output data set. The CHISQ option in the TABLES statement requests computation of Cramér's V. For more information, see the section "Cramér's V" on page 3158.

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. For more information, see the section "Tests for Equal Kappa Coefficients" on page 3210.

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. For more information, see the section "Zelen's Exact Test for Equal Odds Ratios" on page 3218.

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 $r > 2$. For more information, see the section "Tests for Equal Kappa Coefficients" on page 3210.

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. For more information, see the section "Fisher's Exact Test" on page 3156.

GAMMA

includes the gamma statistic in the output data set. The MEASURES option in the TABLES statement requests computation of the gamma statistic. For more information, see the section "Gamma" on page 3160.

GS

GAILSIMON

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. For more information, see the section "Gail-Simon Test for Qualitative Interactions" on page 3221.

JT

includes the Jonckheere-Terpstra test in the output data set. The JT option in the TABLES statement requests the Jonckheere-Terpstra test. For more information, see the section "Jonckheere-Terpstra Test" on page 3202.

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. For more information, see the sections "Simple Kappa Coefficient" on page 3205 and "Overall Kappa Coefficient" on page 3210.

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. For more information, see the section "Kendall's Tau-b" on page 3160.

LAMCR

includes the asymmetric lambda $\lambda(C|R)$ in the output data set. The MEASURES option in the TABLES statement requests computation of lambda. For more information, see the section "Lambda (Asymmetric)" on page 3165.

LAMDAS

includes the symmetric lambda in the output data set. The MEASURES option in the TABLES statement requests computation of lambda. For more information, see the section "Lambda (Symmetric)" on page 3166.

LAMRC

includes the asymmetric lambda $\lambda(R|C)$ in the output data set. The MEASURES option in the TABLES statement requests computation of lambda. For more information, see the section "Lambda (Asymmetric)" on page 3165.

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. For more information, see the section "Adjusted Odds Ratio and Relative Risk Estimates" on page 3215.

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. For more information, see the section "Adjusted Odds Ratio and Relative Risk Estimates" on page 3215.

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. For more information, see the section "Adjusted Odds Ratio and Relative Risk Estimates" on page 3215.

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. For more information, see the section "Likelihood Ratio Chi-Square Test" on page 3155.

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. For more information, see the section "McNemar's Test" on page 3204.

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. For more information, see the section "Measures of Association" on page 3158.

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. For more information, see the section "Mantel-Haenszel Chi-Square Test" on page 3156.

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. For more information, see the section "Adjusted Odds Ratio and Relative Risk Estimates" on page 3215.

MHRRC1

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. For more information, see the section "Adjusted Odds Ratio and Relative Risk Estimates" on page 3215.

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. For more information, see the section "Adjusted Odds Ratio and Relative Risk Estimates" on page 3215.

N

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

NMISS

includes the number of missing observations in the output data set. For more information, see the section “Missing Values” on page 3148.

OR

ODDSRATIO

RROR

includes the odds ratio (for 2 × 2 tables) in the output data set. The MEASURES, OR, and RELRISK options in the TABLES statement request this statistic. For more information, see the section “Odds Ratio” on page 3191.

PCHI

includes the Pearson chi-square in the output data set for two-way tables. For more information, see the section “Pearson Chi-Square Test for Two-Way Tables” on page 3154. For one-way tables, the PCHI option includes the chi-square goodness-of-fit test in the output data set. For more information, see the section “Chi-Square Test for One-Way Tables” on page 3153. 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. For more information, see the section “Pearson Correlation Coefficient” on page 3162.

PHI

includes the phi coefficient in the output data set. The CHISQ option in the TABLES statement requests computation of the phi coefficient. For more information, see the section “Phi Coefficient” on page 3157.

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. For more information, see the section “Polychoric Correlation” on page 3164.

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. For more information, see the section “Risks and Risk Differences” on page 3176.

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. For more information, see the section "Risks and Risk Differences" on page 3176.

RELRISK

includes the column 1 and column 2 relative risks (for 2 × 2 tables) in the output data set. The MEASURES and RELRISK options in the TABLES statement request these statistics. For more information, see the section "Relative Risks" on page 3194.

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. For more information, see the section "Risks and Risk Differences" on page 3176.

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. For more information, see the section "Risks and Risk Differences" on page 3176.

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. For more information, see the section "Risks and Risk Differences" on page 3176.

RRC1

RELRISK1

includes the column 1 relative risk in the output data set. The MEASURES and RELRISK options in the TABLES statement request relative risks, which are available for 2 × 2 tables. For more information, see the section "Odds Ratio and Relative Risks" on page 3191.

RRC2

RELRISK2

includes the column 2 relative risk in the output data set. The MEASURES and RELRISK options in the TABLES statement request relative risks, which are available for 2 × 2 tables. For more information, see the section "Odds Ratio and Relative Risks" on page 3191.

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. For more information, see the section "Risks and Risk Differences" on page 3176.

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. For more information, see the section "Risks and Risk Differences" on page 3176.

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. For more information, see the section "Risks and Risk Differences" on page 3176.

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. For more information, see the section "Risks and Risk Differences" on page 3176.

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. For more information, see the section "Risks and Risk Differences" on page 3176.

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. For more information, see the section "Risks and Risk Differences" on page 3176.

SCORR

includes the Spearman correlation coefficient in the output data set. The MEASURES option in the TABLES statement requests computation of the Spearman correlation. For more information, see the section "Spearman Rank Correlation Coefficient" on page 3163.

SMDCR

includes Somers' $D(C|R)$ in the output data set. The MEASURES option in the TABLES statement requests computation of Somers' D. For more information, see the section "Somers' D" on page 3161.

SMDRC

includes Somers' $D(R|C)$ in the output data set. The MEASURES option in the TABLES statement requests computation of Somers' D. For more information, see the section "Somers' D" on page 3161.

STUTC

TAUC

includes Stuart's tau-c in the output data set. The MEASURES option in the TABLES statement requests computation of tau-c. For more information, see the section "Stuart's Tau-c" on page 3161.

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 $R \times 2$. For more information, see the section "Cochran-Armitage Test for Trend" on page 3201.

TSYMM

BOWKER

includes Bowker's symmetry test in the output data set. The AGREE option in the TABLES statement requests computation of Bowker's test. For more information, see the section "Bowker's Symmetry Test" on page 3205.

U

includes the uncertainty coefficient (symmetric) in the output data set. The MEASURES option in the TABLES statement requests computation of the uncertainty coefficient. For more information, see the section "Uncertainty Coefficient (Symmetric)" on page 3167.

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. For more information, see the section "Uncertainty Coefficients (Asymmetric)" on page 3167.

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. For more information, see the section "Uncertainty Coefficients (Asymmetric)" on page 3167.

WTKAPPA

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 WTKAPPA option also includes the overall weighted kappa coefficient in the output data set. For more information, see the sections "Weighted Kappa Coefficient" on page 3207 and "Overall Kappa Coefficient" on page 3210.

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 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, and 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 47.8 illustrate grouping syntax.

Table 47.8 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 3147 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 any 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 following each table.

Options

Table 47.9 lists the *options* available in the TABLES statement. Descriptions of the *options* follow in alphabetical order.

Table 47.9 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 confidence level for confidence limits
BINOMIAL | BIN	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
COMMONRISKDIFF	Requests common risk difference for $h \times 2 \times 2$ tables
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
MEASURES	Requests measures of association
MISSING	Treats missing values as nonmissing
OR	Requests the odds ratio for 2×2 tables
PLCORR	Requests polychoric correlation
RELRISK	Requests relative risks for 2×2 tables
RISKDIFF	Requests risks and risk differences for 2×2 tables
SCORES=	Specifies type of row and column scores
SENSPEC	Requests sensitivity and specificity for 2×2 tables
TREND	Requests Cochran-Armitage test for trend
Control Additional Table Information	
CELLCHI2	Displays cell contributions to the Pearson chi-square statistic

Table 47.9 *continued*

Option	Description
CUMCOL	Displays cumulative column percentages
DEVIATION	Displays deviations of cell frequencies from expected values
EXPECTED	Displays expected cell frequencies
MISSPRINT	Displays missing value frequencies
PEARSONRES	Displays Pearson residuals in the CROSSLIST table
PRINTKWTS	Displays kappa coefficient weights
SCOROUT	Displays row and column scores
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 contents label for crosstabulation tables
CROSSLIST	Displays crosstabulation tables in ODS column format
FORMAT=	Formats 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
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
NOSPARSE	Suppresses zero-frequency levels in the CROSSLIST table, LIST table, and OUT= data set
NOWARN	Suppresses log warning message for the chi-square test
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*:

AGREE < (*agree-options*) >

requests tests and measures of classification agreement for square tables. This option provides the simple and weighted kappa coefficients along with their standard errors and confidence limits. For multiway tables, this option also produces the overall simple and weighted kappa coefficients (along with their standard errors and confidence limits) and tests for equal kappas among strata. For 2 × 2 tables, this option provides McNemar's test; for square tables that have more than two response categories (levels), this option provides Bowker's symmetry test. For multiway tables that have two response categories, this option also produces Cochran's Q test. For more information, see the section "Tests and Measures of Agreement" on page 3204.

Measures of agreement can be computed only for square tables, where the number of rows equals the number of columns. If your table is not square because some observations have weights of 0, you can specify the ZEROS option in the WEIGHT statement to include these observations in the analysis. For more information, see the section "Tables with Zero-Weight Rows or Columns" on page 3211.

For 2 × 2 tables, the weighted kappa coefficient is equivalent to the simple kappa coefficient, and PROC FREQ displays only analyses for the simple kappa coefficient.

You can specify the confidence level in the ALPHA= option. By default, ALPHA=0.05, which produces 95% confidence limits.

You can specify the EXACT statement to request McNemar's exact test (for 2 × 2 tables), an exact symmetry test, and exact tests for the simple and weighted kappa coefficients. For more information, see the section "Exact Statistics" on page 3222.

You can specify the following *agree-options*:

AC1

requests the AC1 agreement coefficient. For more information, see the section "AC1 Agreement Coefficient" on page 3210.

DFSYM=*df* | ADJUST

controls the degrees of freedom for Bowker's symmetry test. You can specify the value of the degrees of freedom (*df*), or you can specify DFSYM=ADJUST to adjust the degrees of freedom for empty table cells. The value of *df* must be a positive number. By default, *df* is $R(R-1)/2$, where R is the dimension of the two-way table.

When you specify DFSYM=ADJUST, the degrees of freedom are reduced by the number of off-diagonal table-cell pairs that have a total frequency of 0. By default, the degrees of freedom count all off-diagonal table-cell pairs. For more information, see the section "Bowker's Symmetry Test" on page 3205.

KAPPADETAILS

DETAILS

displays the "Kappa Details" table, which includes the following statistics for the simple kappa coefficient: observed agreement, chance-expected agreement, maximum kappa, and the B_n measure. If the two-way table is 2 × 2, the "Kappa Details" table also includes the prevalence index and the bias index. For more information, see the section "Simple Kappa Coefficient" on page 3205.

If the two-way table is larger than 2 × 2, this option also displays the "Weighted Kappa Details" table, which includes the observed agreement and chance-expected agreement components of the weighted kappa coefficient. For more information, see the section "Weighted Kappa Coefficient" on page 3207.

MNULLRATIO=*value*

specifies the null *value* of the ratio of discordant proportions for McNemar's test. By default, MNULLRATIO=1. For more information, see the section "McNemar's Test" on page 3204.

NULLKAPPA=*value*

requests the simple kappa coefficient test and specifies the null *value* for the test. The null value must be between –1 and 1. By default, NULLKAPPA=0. For more information, see the section "Simple Kappa Coefficient" on page 3205.

This option is not available when you specify the KAPPA option in the EXACT statement, which requests an exact test for the kappa coefficient.

NULLWTKAPPA=*value*

requests the weighted kappa coefficient test and specifies the null *value* for the test. The null value must be between –1 and 1. By default, NULLWTKAPPA=0. For more information, see the section "Weighted Kappa Coefficient" on page 3207.

This option is not available when you specify the WTKAPPA option in the EXACT statement, which requests an exact test for the weighted kappa coefficient.

PABAK

requests the prevalence-adjusted bias-adjusted kappa coefficient. For more information, see the section "Prevalence-Adjusted Bias-Adjusted Kappa" on page 3209.

PRINTKWTS

displays the agreement weights that PROC FREQ uses to compute the weighted kappa coefficient. Agreement weights reflect the relative agreement between pairs of variable levels. By default, PROC FREQ uses Cicchetti-Allison agreement weights. If you specify the WT=FC option, the procedure uses Fleiss-Cohen agreement weights. For more information, see the section "Weighted Kappa Coefficient" on page 3207.

TABLES=RESTORE

displays the following agreement tables (which are produced by the AGREE option) in factoid (label-value) format, which is the format of these tables in releases before SAS/STAT 14.2: "McNemar's Test," "Symmetry Test," "Monte Carlo Exact Symmetry Test," "Simple Kappa Coefficient," "Simple Kappa Test," "Weighted Kappa Coefficient," "Weighted Kappa Test," "Monte Carlo Estimates for the Exact Test," "Overall Kappa Coefficient," "Tests for Equal Kappa Coefficients," and "Cochran's Q."

By default (beginning in SAS/STAT 14.3), PROC FREQ displays all agreement tables in tabular form.

WT=FC

specifies Fleiss-Cohen agreement weights in the computation of the weighted kappa coefficient. Agreement weights reflect the relative agreement between pairs of variable levels. By default, PROC FREQ uses Cicchetti-Allison agreement weights to compute the weighted kappa coefficient. For more information, see the section "Weighted Kappa Coefficient" on page 3207.

WTKAPPADETAILS

displays the “Weighted Kappa Details” table, which includes the observed agreement and chance-expected agreement components of the weighted kappa coefficient. This information is available for two-way tables that are larger than 2×2. For more information, see the section “Weighted Kappa Coefficient” on page 3207.

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; a confidence level of α produces $100(1-\alpha)\%$ confidence limits. By default ALPHA=0.05, which produces 95% confidence limits.

This option applies to confidence limits that you request in the TABLES statement. The ALPHA= option in the EXACT statement applies to 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*) >

BIN < (*binomial-options*) >

requests the binomial proportion for one-way tables. When you specify this option, by default PROC FREQ 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 this option, PROC FREQ computes the proportion for the first level that appears in the one-way frequency table. The P= *binomial-option* specifies the null proportion for the binomial tests. If you do not specify this option, PROC FREQ uses 0.5 as the null proportion for the binomial tests.

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.

You can specify the level for the binomial confidence limits in the ALPHA= option. By default, ALPHA=0.05, which produces 95% confidence limits. As part of the noninferiority, superiority, and equivalence analyses, PROC FREQ provides null-based equivalence limits that have a confidence coefficient of $100(1 - 2\alpha)\%$ (Schuirmann 1999). In these analyses, the default of ALPHA=0.05 produces 90% equivalence limits. For more information, see the sections “Noninferiority Test” on page 3172 and “Equivalence Test” on page 3174.

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

For more information, see the section “Binomial Proportion” on page 3168.

Table 47.10 summarizes the *binomial-options*.

Table 47.10 BINOMIAL Options

Option	Description
CORRECT	Requests continuity correction
LEVEL=	Specifies the variable level
OUTLEVEL	Includes the level in the output data sets
P=	Specifies the null proportion
Request Confidence Limits	
CL=AGRESTICOULL | AC	Requests Agresti-Coull confidence limits
CL=BLAKER	Requests Blaker confidence limits
CL=EXACT | CLOPPERPEARSON	Requests exact (Clopper-Pearson) confidence limits
CL=JEFFREYS	Requests Jeffreys confidence limits
CL=LIKELIHOODRATIO | LR	Requests likelihood ratio confidence limits
CL=LOGIT	Requests logit confidence limits
CL=MIDP	Requests exact mid-p confidence limits
CL=WALD	Requests Wald confidence limits
CL=WILSON | SCORE	Requests Wilson (score) confidence limits
Request Tests	
EQUIV | EQUIVALENCE	Requests an equivalence test
MARGIN=	Specifies the test margin
NONINF | NONINFERIORITY	Requests a noninferiority test
SUP | SUPERIORITY	Requests a superiority test
VAR=NULL | SAMPLE	Specifies the test variance

You can specify the following *binomial-options*:

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. By default, ALPHA=0.05, which produces 95% confidence limits for the binomial proportion.

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 can specify in the MARGIN= *binomial-option*).

If you do not specify the CL= *binomial-option*, the BINOMIAL option displays Wald and exact (Clopper-Pearson) confidence limits in the “Binomial Proportion” table.

You can specify the following *types*:

AGRESTICOULL

AC

requests Agresti-Coull confidence limits for the binomial proportion. For more information, see the section "Agresti-Coull Confidence Limits" on page 3169.

BLAKER

requests Blaker confidence limits for the binomial proportion. For more information, see the section "Blaker Confidence Limits" on page 3169.

EXACT

CLOPPERPEARSON

requests exact (Clopper-Pearson) confidence limits for the binomial proportion. For more information, see the section "Exact (Clopper-Pearson) Confidence Limits" on page 3168.

If you do not specify the CL= *binomial-option*, PROC FREQ displays Wald and exact (Clopper-Pearson) confidence limits in the "Binomial Proportion" table. To request exact tests for the binomial proportion, you can specify the BINOMIAL option in the EXACT statement.

JEFFREYS

requests Jeffreys confidence limits for the binomial proportion. For more information, see the section "Jeffreys Confidence Limits" on page 3170.

LIKELIHOODRATIO

LR

requests likelihood ratio confidence limits for the binomial proportion. For more information, see the section "Likelihood Ratio Confidence Limits" on page 3170.

LOGIT

requests logit confidence limits for the binomial proportion. For more information, see the section "Logit Confidence Limits" on page 3170.

MIDP

requests exact mid-p confidence limits for the binomial proportion. For more information, see the section "Mid-p Confidence Limits" on page 3170.

WALD < (CORRECT) >

requests Wald confidence limits for the binomial proportion. For more information, see the section "Wald Confidence Limits" on page 3168.

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

If you do not specify the CL= *binomial-option*, PROC FREQ displays Wald and exact (Clopper-Pearson) confidence limits in the "Binomial Proportion" table.

WILSON < (CORRECT) >

SCORE < (CORRECT) >

requests Wilson confidence limits for the binomial proportion. These are also known as *score* confidence limits. For more information, see the section "Wilson (Score) Confidence Limits" on page 3171.

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. For more information, see the section "Equivalence Test" on page 3174. You can specify the equivalence test margins, the null proportion, and the variance type in the MARGIN=, P=, and VAR= *binomial-options*, respectively. To request an exact equivalence test, you can specify the BINOMIAL option in the EXACT statement.

LEVEL=*level-number* **|** **'***level-value***'**

specifies the variable level for the binomial proportion. You can specify the *level-number*, which is the order in which the level appears in the one-way frequency table. Or you can specify the *level-value*, which is the formatted value of the variable level. The *level-number* must be a positive integer. You must enclose the *level-value* in single quotes.

By default, PROC FREQ computes the binomial proportion for the first variable level that appears in the one-way frequency table.

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

specifies the margin for the noninferiority, superiority, and equivalence tests, which you can request by specifying the NONINF, SUP, and EQUIV *binomial-options*, respectively. By default, MARGIN=0.2.

For noninferiority and superiority tests, specify a single *value* in 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. PROC FREQ 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 in the P= *binomial-option*) and by the margin value. The noninferiority limit is the null proportion minus the margin. By default, the null proportion is 0.5 and the margin is 0.2, which produces a noninferiority limit of 0.3. The superiority limit is 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 and 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 in the P= *binomial-option*) and by the margin values. The lower equivalence limit is the null proportion plus the lower margin. By default, the null proportion is 0.5 and the lower margin is –0.2, which produces a lower equivalence limit of 0.3. The upper equivalence limit is the null proportion plus the upper margin, which is 0.7 by default.

For more information, see the sections "Noninferiority Test" on page 3172 and "Equivalence Test" on page 3174.

NONINF

NONINFERIORITY

requests a test of noninferiority for the binomial proportion. For more information, see the section "Noninferiority Test" on page 3172. You can specify the noninferiority test margin, the null proportion, and the variance type in the MARGIN=, P=, and VAR= *binomial-options*, respectively. To request an exact noninferiority test, you can specify the BINOMIAL option in the EXACT statement.

OUTLEVEL

includes the variables LevelNumber and LevelValue in all ODS output data sets that PROC FREQ produces when you specify the BINOMIAL option in the TABLES statement. The OUTLEVEL option also includes the variables LevelNumber and LevelValue in the statistics output data set that PROC FREQ produces when you specify the BINOMIAL option in the OUTPUT statement.

The LevelNumber and LevelValue variables identify the analysis variable level for which PROC FREQ computes the binomial proportion. The value of LevelNumber is the order of the level in the one-way frequency table. The value of LevelValue is the formatted value of the level. You can specify the OUTLEVEL *binomial-option* with or without the LEVEL= *binomial-option*.

P=*value*

specifies the null hypothesis proportion for the binomial tests. 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. PROC FREQ treats the value 1 as 1%. By default, P=0.5.

SUP

SUPERIORITY

requests a test of superiority for the binomial proportion. For more information, see the section "Superiority Test" on page 3174. You can specify the superiority test margin, the null proportion, and the variance type in the MARGIN=, P=, and VAR= *binomial-options*, respectively. To request an exact superiority test, you can specify the BINOMIAL option in the EXACT statement.

VAR=NULL | SAMPLE

specifies the type of variance to use in the Wald tests of noninferiority, superiority, and equivalence. If you specify VAR=SAMPLE, PROC FREQ computes the variance estimate by using the sample proportion. If you specify VAR=NULL, PROC FREQ computes a test-based variance by using the null hypothesis proportion (which you can specify in the P= *binomial-option*). For more information, see the sections "Noninferiority Test" on page 3172 and "Equivalence Test" on page 3174. By default, VAR=SAMPLE.

CELLCHI2

displays cell chi-squares in the crosstabulation table. A cell chi-square is the table cell's contribution to the Pearson chi-square statistic. The cell chi-square is computed as $(frequency - expected)^2 / expected$, where *frequency* is the table cell frequency (count) and *expected* is the expected cell frequency, which is computed under the null hypothesis that the row and column variables are independent. For more information, see the section "Pearson Chi-Square Test for Two-Way Tables" on page 3154.

This option applies to two-way and multiway tables that are displayed in the default crosstabulation cell format or in CROSSLIST format. This option has no effect for one-way frequency tables or for tables that are produced by 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. For more information, see the section "Chi-Square Tests and Statistics" on page 3153.

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 LRCHI *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 3153 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*:

DF=*df*

specifies the degrees of freedom for the chi-square tests. The value of *df* must not be 0. 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 3153 and "Chi-Square Tests and Statistics" on page 3153 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.

LRCHI

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 3155 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. For more information, see the section "Chi-Square Test for One-Way Tables" on page 3153. 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. For more information, see the section "Chi-Square Test for One-Way Tables" on page 3153. 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=*type* | (*types*)

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.

You can specify one or more of the following *types* in the WARN= option. If you specify more than one *type* 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 of association, which you can request by specifying the MEASURES option. For more information, see the sections "Measures of Association" on page 3158 and "Confidence Limits" on page 3158. You can set the level of the confidence limits by using the ALPHA= option; by default, ALPHA=0.05, which produces 95% confidence limits.

If you omit the MEASURES option, the CL option invokes MEASURES. The CL option is equivalent to the MEASURES(CL) option.

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*.) For more information, see the section "Cochran-Mantel-Haenszel Statistics" on page 3211.

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. For more information, see the section "Zelen's Exact Test for Equal Odds Ratios" on page 3218. 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. For more information, see the section "Exact Confidence Limits for the Common Odds Ratio" on page 3219.

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. For more information, see the section "Breslow-Day Test for Homogeneity of the Odds Ratios" on page 3217.

GAILSIMON < (COLUMN=1 | 2) >

GS < (COLUMN=1 | 2) >

requests the Gail-Simon test for qualitative interaction, which applies to stratified 2 × 2 tables. For more information, see the section "Gail-Simon Test for Qualitative Interactions" on page 3221.

The COLUMN= option specifies the column of the risk differences to use to compute 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.

I2

requests the I-square measure of heterogeneity for stratified 2 × 2 tables. I-square is computed from a Q test that is based on odds ratios. The I2 *cmh-option* invokes the QOR *cmh-option*. For more information, see the section "I-Square Measure of Heterogeneity" on page 3218.

MANTELFLEISS

MF

requests the Mantel-Fleiss criterion for the Mantel-Haenszel statistic for stratified 2 × 2 tables. For more information, see the section "Mantel-Fleiss Criterion" on page 3214.

QOR

requests a Q test for heterogeneity of odds ratios for stratified 2 × 2 tables. For more information, see the section "Q Test for Homogeneity of Odds Ratios" on page 3217.

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. For more information, see the description of the CMH option.

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. For more information, see the description of the CMH option.

COMMONRISKDIFF < *options* >

requests the common (stratified) risk difference for multiway 2 × 2 tables, where the risk difference is the difference between the row 1 proportion and the row 2 proportion in a 2 × 2 table. By default, this option provides Mantel-Haenszel and summary score estimates of the common risk difference, together with their confidence limits. For more information, see the section "Common Risk Difference" on page 3187.

You can specify the following *options* to request confidence limit types and tests for the common risk difference:

CL=*type* **|** **(***types***)**

requests confidence limits for the common risk difference. You can specify one or more *types* of confidence limits. When you specify only one *type*, you can omit the parentheses. You can specify CL=NONE to suppress the "Confidence Limits for the Common Risk Difference" table.

You can specify the confidence level in the ALPHA= option. By default, ALPHA=0.05, which produces 95% confidence limits for the common risk difference.

You can specify one or more of the following *types*:

K

KLINGENBERG

requests Klingenberg confidence limits for the Mantel-Haenszel common risk difference. For more information, see the section "Klingenberg Confidence Limits" on page 3188.

MH

requests Mantel-Haenszel confidence limits, which are computed by using Mantel-Haenszel stratum weights and the Sato variance estimator (Sato 1989). For more information, see the section "Mantel-Haenszel Confidence Limits and Test" on page 3187.

MR

MINRISK

requests minimum risk confidence limits, which are computed by using minimum risk weights. For more information, see the section "Minimum Risk Confidence Limits and Test" on page 3188.

NEWCOMBE

requests stratified Newcombe confidence limits that use Mantel-Haenszel weights to combine the stratum components. For more information, see the section "Stratified Newcombe Confidence Limits" on page 3190.

NEWCOMBEMR

requests stratified Newcombe confidence limits that use minimum risk weights to combine the stratum components. For more information, see the section "Stratified Newcombe Confidence Limits" on page 3190.

NONE

suppresses the "Confidence Limits for the Common Risk Difference" table.

SCORE

requests summary score confidence limits. For more information, see the section "Summary Score Confidence Limits" on page 3190.

COLUMN=1 | 2

specifies the table column for which to compute the common risk difference statistics. If you do not specify this option but you do specify the RISKDIFF(COLUMN=) option, PROC FREQ provides the common risk difference statistics for the column that you specify in the RISKDIFF(COLUMN=) option. If you do not specify either of these options, COLUMN=1 by default.

CORRECT=NO

removes the continuity correction in the minimum risk confidence limits and in the minimum risk test, which you can request by specifying the CL=MR and TEST=MR options, respectively. For more information, see the section "Minimum Risk Confidence Limits and Test" on page 3188.

PRINTWTS < =*type* | (*types*) >

displays the stratum weights together with the stratum risk differences and frequencies. By default, this option displays the weight type or types for the confidence limits and tests that you request. Optionally, you can specify the weight type to display.

You can specify one or more of the following *types*:

MH

displays Mantel-Haenszel stratum weights. For more information, see the section "Mantel-Haenszel Confidence Limits and Test" on page 3187.

MR

displays minimum risk stratum weights. For more information, see the section "Minimum Risk Confidence Limits and Test" on page 3188.

SCORE

displays summary score stratum weights. For more information, see the section "Summary Score Confidence Limits" on page 3190.

TEST < =*type* | (*types*) >

requests common risk difference tests. You can specify one or more *types*. When you specify only one *type*, you can omit the parentheses. If you do not specify *types*, this option provides tests that correspond to the confidence limit *types* that you specify in the CL= option.

You can specify one or more of the following *types*:

MH

requests a Mantel-Haenszel test, which is computed by using Mantel-Haenszel stratum weights and the Sato variance estimator (Sato 1989). For more information, see the section "Mantel-Haenszel Confidence Limits and Test" on page 3187.

MR < (VAR=SAMPLE) >

MINRISK < (VAR=SAMPLE) >

requests the minimum risk test, which is computed by using minimum risk weights. If you specify VAR=SAMPLE, PROC FREQ uses the sample (observed) variance estimate instead of a null variance estimate to compute the minimum risk test statistic. For more information, see the section "Minimum Risk Confidence Limits and Test" on page 3188.

SCORE

requests the summary score test. For more information, see the section "Summary Score Confidence Limits" on page 3190.

CONTENTS=*'string'*

specifies the label to use for crosstabulation tables in the contents file, the Results window, and the ODS trace record. For information about output presentation, see the *SAS Output Delivery System: User's Guide*. By default, the contents label for crosstabulation tables is "Cross-Tabular Freq Table".

PROC FREQ uses the same contents label for all crosstabulation tables that are produced by the same TABLES statement. To use different contents labels for different crosstabulation tables, you can request the tables in separate TABLES statements and specify the CONTENTS= option in each TABLES statement.

To remove the crosstabulation table entry from the contents file, you can specify CONTENTS=''.

The CONTENTS= option affects only contents labels for crosstabulation tables; this option does not affect contents labels for other tables that PROC FREQ produces. To specify the contents label for any table that PROC FREQ produces, you can use PROC TEMPLATE to create a customized table template. The CONTENTS_LABEL attribute in the DEFINE TABLE statement of PROC TEMPLATE specifies the contents label for the table. For more information, see the chapter "The TEMPLATE Procedure" in the *SAS Output Delivery System: User's Guide*.

CROSSLIST < (*options*) >

displays crosstabulation tables by using an ODS column format instead of the default crosstabulation table cell format. In CROSSLIST tables, the rows correspond to the crosstabulation table cells, and the columns correspond to descriptive statistics such as frequencies and percentages. For more information about the contents of CROSSLIST tables, see the section "Two-Way and Multiway Tables" on page 3232.

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. The FORMAT= and CUMCOL options have no effect on CROSSLIST tables.

By default, CROSSLIST tables display all levels of the column variable within each level of the row variable, including any levels that have frequencies of 0. By default, multiway CROSSLIST tables display all levels of the row variable within each stratum, including any row levels that have frequencies of 0 in the stratum. To suppress variable levels that have frequencies of 0, you can specify the NOSPARSE option.

You cannot specify both the LIST option and the CROSSLIST option in the same TABLES statement.

You can specify the following *options*:

PEARSONRES

displays Pearson residuals in the CROSSLIST table. The Pearson residual is the square root of the table cell's contribution to the Pearson chi-square statistic. The Pearson residual is computed as $(frequency - expected) / \sqrt{expected}$, where *frequency* is the table cell frequency (count) and *expected* is the expected table cell frequency, which is computed under the null hypothesis that the row and column variables are independent. For more information, see the section "Pearson Chi-Square Test for Two-Way Tables" on page 3154.

You can display the expected values, deviations, and cell chi-squares in the CROSSLIST table by specifying the EXPECTED, DEVIATION, and CELLCHI2 options, respectively.

STDRES

displays standardized residuals in the CROSSLIST table. The standardized residual is the ratio of (*frequency – expected*) to its standard error, where *frequency* is the table cell frequency (count) and *expected* is the expected table cell frequency, which is computed under the null hypothesis that the row and column variables are independent. For more information, see the section "Standardized Residuals" on page 3154. You can display the expected values and deviations in the CROSSLIST table by specifying the EXPECTED and DEVIATION options, respectively.

CUMCOL

displays the cumulative column percentages in the cells of the crosstabulation table. This option applies only to crosstabulation tables that are displayed in the default crosstabulation cell format.

DEVIATION

displays deviations in the crosstabulation table. A deviation is the difference between the observed table cell frequency and the expected frequency (*frequency* – *expected*). The expected frequencies are computed under the null hypothesis that the row and column variables are independent. For more information, see the section "Pearson Chi-Square Test for Two-Way Tables" on page 3154. You can display the expected values by specifying the EXPECTED option. This option applies to two-way and multiway tables that are displayed in the default crosstabulation cell format or in CROSSLIST format.

EXPECTED

displays expected frequencies in the crosstabulation table. Expected cell frequencies are computed as the product of the row marginals and the column marginals divided by the total frequency of the two-way table. These are the expected frequencies under the null hypothesis that the row and column variables are independent. For more information, see the section "Pearson Chi-Square Test for Two-Way Tables" on page 3154. This option applies to two-way and multiway tables that are displayed in the default crosstabulation cell format or in CROSSLIST format.

FISHER

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 3156 and "Exact Statistics" on page 3222 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 3224 for more information.

FORMAT=*format-name*

specifies the format for frequencies, expected frequencies, and deviations in crosstabulation tables.

You can specify any standard SAS numeric format, or you can specify a numeric format that is defined by the FORMAT procedure. The format length must not exceed 24.

This option applies to crosstabulation tables that are displayed in default crosstabulation cell form.

To change display formats in any table that PROC FREQ produces, you can use PROC TEMPLATE. For more information, see the chapter "The TEMPLATE Procedure" in the *SAS Output Delivery System: User's Guide*.

GAILSIMON < (COLUMN=1 | 2) >

GS < (COLUMN=1 | 2) >

requests the Gail-Simon test for qualitative interaction, which applies to stratified 2 × 2 tables. For more information, see the section "Gail-Simon Test for Qualitative Interactions" on page 3221.

The COLUMN= option specifies the column of the risk differences to use to compute 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. For more information, see the section "Jonckheere-Terpstra Test" on page 3202. 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 3222 for more information.

LIST

displays two-way and multiway tables by using a list format instead of the default crosstabulation cell format. This option displays an entire multiway table in a single table instead of separate two-way (stratum) tables. In list format, each row of the table corresponds to a single crosstabulation table cell. For more information, see the section "Two-Way and Multiway Tables" on page 3232.

By default, a table in list format does not display cells for which the frequency is 0 unless you specify the ZEROS option in the WEIGHT statement and do not specify the NOSPARSE option in the TABLES statement.

The LIST option is not available in a TABLES statement together with statistic options. When you specify statistic options in a TABLES statement, you must display the crosstabulation tables by using the default table cell format or the CROSSLIST format. You can specify the LIST option and statistic options in different TABLES statements in the same invocation of PROC FREQ.

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 < (CL) >

requests measures of association and their asymptotic standard errors. This option provides the following measures: 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). If you specify the CL option in parentheses after the

MEASURES option, PROC FREQ provides confidence limits for the measures of association. For more information, see the section “Measures of Association” on page 3158.

For 2 × 2 tables, the MEASURES option also provides the odds ratio, column 1 relative risk, column 2 relative risk, and their asymptotic Wald confidence limits. You can request the odds ratio and relative risks separately (without the other measures of association) by specifying the RELRISK option. You can request confidence limits for the odds ratio by specifying the OR(CL=) option.

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. For more information, see the descriptions of the TEST and EXACT statements and the section “Exact Statistics” on page 3222.

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 following the table. See the section “Missing Values” on page 3148 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 following the table. See the section “Missing Values” on page 3148 for more information.

NOCOL

suppresses the display of column percentages in crosstabulation tables. A table cell’s column percentage is computed as the cell frequency divided by the corresponding column frequency.

NOCUM

suppresses the display of cumulative frequencies and percentages in one-way frequency tables and in list-format crosstabulation tables (which you request by specifying the LIST option).

NOFREQ

suppresses the display of table cell frequencies in crosstabulation tables. This option also suppresses row total frequencies. This option applies only to crosstabulation tables that are displayed in the default crosstabulation cell format.

NOPERCENT

suppresses the display of percentages in crosstabulation tables. These percentages include table cell, row, and column percentages of the total two-way table frequency. To suppress the display of cell percentages of row or column totals, you can specify the NOROW or NOCOL option, respectively.

For one-way frequency tables and list-format crosstabulation tables, 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 tables. A table cell's row percentage is computed as the cell frequency divided by the corresponding row frequency.

NOSPARSE

suppresses the display of zero-frequency table cells in LIST and CROSSLIST tables. This option also excludes zero-frequency table cells from OUT= output data sets.

When you specify the ZEROS option in the WEIGHT statement, PROC FREQ includes observations that have weights of 0 in the analysis. By default, zero-frequency table cells are displayed in the LIST table and included in the OUT= data set. You can use the NOSPARSE option to exclude the zero-frequency table cells.

When you specify the CROSSLIST option, by default the CROSSLIST table displays all levels of the column variable within each level of the row variable (including any levels that have frequencies of 0). By default for multiway CROSSLIST tables, the CROSSLIST table displays all levels of the row variable within each stratum of the table (including any row levels that have frequencies of 0 in the stratum). You can use the NOSPARSE option to suppress the zero-frequency variable levels in the CROSSLIST table.

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% of the table 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.

OR < (CL=*type* | (*types*) >

ODDSRATIO < (CL=*type* | (*types*) >

requests the odds ratio and confidence limits for 2 × 2 tables. For more information, see the section "Odds Ratio" on page 3191.

You can specify one or more *types* of confidence limits. When you specify only one confidence limit *type*, you can omit the parentheses around the request. PROC FREQ displays the confidence limits in the "Confidence Limits for the Odds Ratio" table.

Specifying the OR option without the CL= option is equivalent to specifying the RELRISK option, which produces the “Odds Ratio and Relative Risks” table. For more information, see the description of the RELRISK option. When you specify the OR(CL=) option, PROC FREQ does not produce the “Odds Ratio and Relative Risks” table unless you also specify the RELRISK or MEASURES option.

The ALPHA= option determines the confidence level; by default, ALPHA=0.05, which produces 95% confidence limits for the odds ratio.

You can specify the following *types*:

EXACT

displays exact confidence limits for the odds ratio in the “Confidence Limits for the Odds Ratio” table. (By default, PROC FREQ displays the exact confidence limits in a separate table.) You must also request computation of the exact confidence limits by specifying the OR option in the EXACT statement. For more information, see the subsection “Exact Confidence Limits” in the section “Confidence Limits for the Odds Ratio” on page 3192.

LR

LIKELIHOODRATIO

requests likelihood ratio confidence limits for the odds ratio. For more information, see the subsection “Likelihood Ratio Confidence Limits” in the section “Confidence Limits for the Odds Ratio” on page 3192.

MIDP

requests exact mid-p confidence limits for the odds ratio. For more information, see the subsection “Exact Mid-p Confidence Limits” in the section “Confidence Limits for the Odds Ratio” on page 3192.

SCORE < (CORRECT=NO) >

requests score confidence limits for the odds ratio. For more information, see the subsection “Score Confidence Limits” in the section “Confidence Limits for the Odds Ratio” on page 3192. If you specify CORRECT=NO, PROC FREQ provides the uncorrected form of the score confidence limits.

WALD

requests asymptotic Wald confidence limits, which are based on a log transformation of the odds ratio. For more information, see the subsection “Wald Confidence Limits” in the section “Confidence Limits for the Odds Ratio” on page 3192.

WALDMODIFIED

requests Wald modified confidence limits for the odds ratio. For more information, see the subsection “Wald Modified Confidence Limits” in the section “Confidence Limits for the Odds Ratio” on page 3192.

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. For more information, see the section “Output Data Sets” on page 3227. 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. For more information, see the section "Output Data Sets" on page 3227. 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. For more information, see the section "Output Data Sets" on page 3227. 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 $n > 2$

For more information, see the section "Output Data Sets" on page 3227. The OUTPCT option has no effect for one-way tables.

PLCORR < (*options*) >

POLYCHORIC < (*options*) >

requests the polychoric correlation coefficient and its asymptotic standard error. For 2 × 2 tables, this statistic is more commonly known as the tetrachoric correlation coefficient and is labeled as such in the displayed output. For more information, see the section "Polychoric Correlation" on page 3164.

If you also specify the CL or MEASURES(CL) option, PROC FREQ provides confidence limits for the polychoric correlation. If you specify the PLCORR option in the TEST statement, the procedure provides Wald and likelihood ratio tests for the polychoric correlation. The PLCORR option invokes the MEASURES option.

You can specify the following *options*:

ADJUST

replaces a 2 × 2 table cell frequency of 0 by 0.5 before computing the tetrachoric correlation (Brown and Benedetti 1977a, p. 353). To maintain the row and column marginal frequencies, adjacent cell frequencies are decreased by 0.5 and the opposite cell frequency is increased by 0.5.

This option is available for 2 × 2 tables and is applied only when a single cell frequency is 0. It has no effect when both off-diagonal cell frequencies are 0 (and therefore the correlation is 1) or when both diagonal cell frequencies are 1 (and therefore the correlation is –1).

CONVERGE=*value*

specifies the convergence criterion. The *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= *number*, whichever happens first. For parameter values that are less than 0.01, PROC FREQ evaluates convergence by using the absolute difference instead of the relative difference. For more information, see the section "Polychoric Correlation" on page 3164.

MAXITER=*number*

specifies the maximum *number* of iterations. The value of *number* must be a positive integer. By default, MAXITER=50. Iterative computation of the polychoric correlation stops when the number of iterations exceeds the maximum *number* or when the convergence measure falls below the CONVERGE= *value*, whichever happens first. For more information, see the section "Polychoric Correlation" on page 3164.

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 oddsratioplot)
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 663 in Chapter 24, "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 47.11) 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 47.11 lists the available *plot-requests* together with their required TABLES statement options. Descriptions of the *plot-requests* follow the table in alphabetical order.

Table 47.11 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, or RELRISK ($h \times 2 \times 2$ table)
RELRISKPLOT	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, $r > 2$)

You can specify the following *plot-requests*:

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 47.12 lists the *plot-options* that are available for agreement plots. For descriptions of the *plot-options*, see the subsection "Plot Options" in this section.

Table 47.12 Plot Options for AGREEPLOT

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

*Default

If you specify the STATS *plot-option*, the agreement plot displays the values of the kappa coefficient, the weighted kappa coefficient, the B_n measure (Bangdiwala and Bryan 1987), and the sample size. PROC FREQ stores these statistics in an ODS table named BnMeasure, which is not displayed. For more information, see the section “ODS Table Names” on page 3240.

ALL

requests all plots that are associated with the specified analyses. Table 47.11 lists the available *plot-requests* and the corresponding analysis options. 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. (These plots are not produced 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 47.13 lists the *plot-options* that are available for cumulative frequency plots. For descriptions of the *plot-options*, see the subsection “Plot Options” in this section.

Table 47.13 Plot Options for CUMFREQPLOT

Plot Option	Description	Values
ORIENT=	Orientation	HORIZONTAL or VERTICAL*
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 47.14 lists the *plot-options* that are available for deviation plots. For descriptions of the *plot-options*, see the subsection “Plot Options” in this section.

Table 47.14 Plot Options for DEVIATIONPLOT

Plot Option	Description	Values
NOSTAT	No statistic	None
ORIENT=	Orientation	HORIZONTAL or VERTICAL*
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.

Table 47.15 lists the *plot-options* that are available for frequency plots. For descriptions of the *plot-options*, see the subsection "Plot Options" in this section.

Table 47.15 Plot Options for FREQPLOT

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

*Default

**For two-way tables

You can specify the following *plot-options* for all frequency plots: ORIENT=, SCALE=, and TYPE=. You can specify the following *plot-options* for frequency plots of two-way (and multiway) 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.

By default, PROC FREQ displays frequency plots as bar charts. To display frequency plots as dot plots, specify TYPE=DOTPLOT. To plot percentages instead of frequencies, specify SCALE=PERCENT. For two-way tables, there are four frequency plot layouts available, which you can request by specifying the TWOWAY= *plot-option*. For more information, see the subsection "Plot Options" in this section.

By default, graph cells in a two-way layout are first grouped by column variable levels; row variable levels are then displayed within the column variable levels. To group first by row variable levels, specify GROUPBY=ROW.

KAPPAPLOT < (*plot-options*) >

requests a plot of kappa statistics along 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 47.16 lists the *plot-options* that are available for kappa plots. For descriptions of the *plot-options*, see the subsection "Plot Options" in this section.

Table 47.16 Plot Options for KAPPAPLOT and WTKAPPAPLOT

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

*Default

MOSAICPLOT < (*plot-options*) >

requests a mosaic plot. Mosaic plots are available for two-way and multiway 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. For more information, see Friendly (2000).

By default, the colors of the tiles correspond to the row variable levels. If you specify the COLORSTAT= *plot-option*, the tiles are colored according to the values of the Pearson or standardized residuals.

You can specify the following *plot-options*:

COLORSTAT < =PEARSONRES | STDRES >

colors the mosaic plot tiles according to the values of residuals. If you specify COLORSTAT=PEARSONRES, the tiles are colored according to the Pearson residuals of the corresponding table cells. For more information, see the section "Pearson Chi-Square Test for Two-Way Tables" on page 3154. If you specify COLORSTAT=STDRES, the tiles are colored according to the standardized residuals of the corresponding table cells. For more information, see the section "Standardized Residuals" on page 3154. You can display the Pearson or standardized residuals in the CROSSLIST table by specifying the CROSSLIST(PEARSONRES) or CROSSLIST(STDRES) option, respectively.

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 along 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, or RELRISK option in the TABLES statement to compute the odds ratios.

Table 47.17 lists the *plot-options* that are available for odds ratio plots. For descriptions of the *plot-options*, see the subsection "Plot Options" in this section.

Table 47.17 Plot Options for ODDSRATIOPLOT, RELRISKPLOT, and RISKDIFFPLOT

Plot Option	Description	Values
CL=	Confidence limit type	Type
CLDISPLAY=	Error bar type	BAR, LINE, LINEARROW, SERIF*, or SERIFARROW
COLUMN=**	Risk column	1* or 2
COMMON=	Common value	NO or YES*
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	Statistic values	None

*Default

**Available for RELRISKPLOT and RISKDIFFPLOT

***Available for ODDSRATIOPLOT and RELRISKPLOT

You can specify one of the following confidence limit types for the odds ratio plot: exact (CL=EXACT), likelihood ratio (CL=LR), exact mid-p (CL=MIDP), score (CL=SCORE), Wald (CL=WALD), or Wald modified (CL=WALDMODIFIED). By default, the odds ratio plot displays Wald confidence limits. For more information, see the descriptions of the CL= *plot-option* and the OR(CL=) option.

To display exact confidence limits in the odds ratio plot, you must also request their computation by specifying the OR option in the EXACT statement.

When CL=WALD or CL=EXACT, the odds ratio plot displays the common odds ratio by default when it is available. To compute the common odds ratio along with Wald confidence limits, specify the CMH option in the TABLES statement. To compute the common odds ratio along with exact confidence limits, specify the COMOR option in the EXACT statement. To suppress display of the common odds ratio, specify COMMON=NO.

RELRISKPLOT < (*plot-options*) >

requests a plot of relative risks along 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.

Table 47.17 lists the *plot-options* that are available for relative risk plots. For descriptions of the *plot-options*, see the subsection "Plot Options" in this section.

You can specify one of the following confidence limit types for the relative risk plot: exact (CL=EXACT), likelihood ratio (CL=LR), score (CL=SCORE), Wald (CL=WALD), or Wald modified (CL=WALDMODIFIED). By default, the relative risk plot displays Wald confidence limits. For more information, see the descriptions of the CL= *plot-option* and the RELRISK(CL=) option.

To display exact confidence limits in the relative risk plot, you must also request their computation by specifying the RELRISK option in the EXACT statement. The risk column that you specify for the confidence limits must match the risk column that you specify for the plot.

The relative risk plot displays the common relative risk by default when you specify CL=WALD and the CMH option in the TABLES statement. To suppress display of the common relative risk, specify COMMON=NO.

In addition to the *plot-options* in Table 47.17, you can specify the following *plot-option*:

FOOTNOTE=NO

suppresses the footnote that identifies the column for which the relative risk is computed.

RISKDIFFPLOT < (*plot-options*) >

requests a plot of risk (proportion) differences along with confidence limits for multiway 2 × 2 tables. The risk difference plot displays the risk difference (with confidence limits) for each 2 × 2 table (stratum). Optionally, the plot also displays the common risk difference.

To produce a risk difference plot, you must also specify the RISKDIFF option in the TABLES statement to compute risk differences.

Table 47.17 lists the *plot-options* that are available for risk difference plots. For descriptions of the *plot-options*, see the subsection "Plot Options" in this section.

You can specify the confidence limit type for the stratum risk differences by using the CL= *plot-option*. You can specify one of the following confidence limit types: Agresti-Caffo (CL=AC), exact (CL=EXACT), Hauck-Anderson (CL=HA), Miettinen-Nurminen (score) (CL=MN), Newcombe (CL=NEWCOMBE), and Wald (CL=WALD). By default, the plot displays Wald confidence limits for the stratum risk differences. For more information, see the descriptions of the CL= *plot-option* and the RISKDIFF(CL=) option.

To display exact confidence limits in the risk difference plot, you must also request their computation by specifying the RISKDIFF option in the EXACT statement. The risk column that you specify for the confidence limits must match the risk column that you specify for the plot.

By default, the risk difference plot displays the common risk difference when you specify the RISKDIFF(COMMON) or COMMONRISKDIFF option unless you specify the CL=EXACT *plot-option*. To suppress display of the common risk difference, specify COMMON=NO.

In addition to the *plot-options* in Table 47.17, you can specify the following *plot-options*:

CLNOTE=NO

suppresses the note that identifies the confidence limit type.

COMMON=*type*

specifies the *type* of confidence limits to display for the common risk difference.

The default common confidence limit *type* depends on the stratum confidence limit *type*. By default, COMMON=NEWCOMBE if CL=NEWCOMBE and COMMON=SCORE if CL=SCORE. Otherwise, COMMON=MH by default.

You can specify one of the following *types*:

K

KLINGENBERG

displays Klingenberg confidence limits. For more information, see the COMMONRISKDIFF(CL=K) option and the section "Klingenberg Confidence Limits" on page 3188.

MH

displays Mantel-Haenszel confidence limits. For more information, see the COMMONRISKDIFF(CL=MH) option and the section "Mantel-Haenszel Confidence Limits and Test" on page 3187.

MR

MINRISK

displays minimum risk confidence limits. For more information, see the COMMONRISKDIFF(CL=MR) option and the section "Minimum Risk Confidence Limits and Test" on page 3188.

NEWCOMBE

displays stratified Newcombe confidence limits that use Mantel-Haenszel weights to combine the stratum components. For more information, see the COMMONRISKDIFF(CL=NEWCOMBE) option and the section "Stratified Newcombe Confidence Limits" on page 3190.

NEWCOMBEMR

displays stratified Newcombe confidence limits that use minimum risk weights to combine the stratum components. For more information, see the COMMONRISKDIFF(CL=NEWCOMBEMR) option and the section "Stratified Newcombe Confidence Limits" on page 3190.

NONE

suppresses the common risk difference in the risk difference plot.

SCORE

displays summary score confidence limits. For more information, see the COMMONRISKDIFF(CL=SCORE) option and the section "Summary Score Confidence Limits" on page 3190.

FOOTNOTE=NO

suppresses the footnote that identifies the column for which the risk difference is computed.

WTKAPPAPLOT < (*plot-options*) >

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

To produce a weighted kappa plot, you must specify the AGREE option in the TABLES statement to compute weighted kappa coefficients, and the table dimension must be greater than 1.

Table 47.16 lists the *plot-options* that are available for weighted kappa plots. For descriptions of the *plot-options*, see the subsection "Plot Options" in this section.

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=. For descriptions of these *plot-options*, see the subsection "Plot Options" in this section.

In addition to these *plot-options*, you can specify the following *global-plot-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. You can specify the CL= *plot-option* when you specify any of the following *plot-requests*: ODDSRATIOPLOT, RELRISKPLOT, and RISKDIFFPLOT.

For odds ratio plots (ODDSRATIOPLOT), the available confidence limit types include the following: exact (CL=EXACT), likelihood ratio (CL=LR), exact mid-p (CL=MIDP), score (CL=SCORE), Wald (CL=WALD), and Wald modified (CL=WALDMODIFIED). For more information, see the description of the OR(CL=) option and the section "Confidence Limits for the Odds Ratio" on page 3192. By default, CL=WALD. When you specify CL=EXACT to display exact confidence limits, you must also request computation of exact confidence limits by specifying the OR option in the EXACT statement.

For relative risk plots (RELRISKPLOT), the available confidence limit types include the following: exact (CL=EXACT), likelihood ratio (CL=LR), score (CL=SCORE), Wald (CL=WALD), and Wald modified (CL=WALDMODIFIED). For more information, see the description of the RELRISK(CL=) option and the section "Confidence Limits for the Relative Risk" on page 3195. By default, CL=WALD. When you specify CL=EXACT to display exact confidence limits, you must also request computation of exact confidence limits by specifying the RELRISK option in the EXACT statement.

For risk difference plots (RISKDIFFPLOT), the available confidence limit types include the following: Agresti-Caffo (CL=AC), exact (CL=EXACT), Hauck-Anderson (CL=HA), Miettinen-Nurminen (score) (CL=MN), Newcombe (CL=NEWCOMBE), and Wald (CL=WALD). For more information, see the description of the RISKDIFF(CL=) option and the section "Confidence Limits for the Risk Difference" on page 3178. By default, CL=WALD. When 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=BAR < ***width*** **> | LINE | LINEARROW | SERIF | SERIFARROW**

controls the appearance of the confidence limit error bars. You can specify the CLDISPLAY= *plot-option* when you specify the following *plot-requests*: KAPPAPLOT, ODDSRATIOPLOT, RELRISKPLOT, RISKDIFFPLOT, and WTKAPPAPLOT.

By default, CLDISPLAY=SERIF, which displays the confidence limits as lines with serifs. CLDISPLAY=LINE displays the confidence limits as plain lines without serifs. The CLDISPLAY=SERIFARROW and CLDISPLAY=LINEARROW *plot-options* display arrowheads on any error bars that are clipped by the RANGE= *plot-option*; if an entire error bar is cut from the plot, the plot displays an arrowhead that points toward the statistic.

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 table column to use to compute the risks (proportion) for the relative risk plot (RELRISKPLOT) and the risk difference plot (RISKDIFFPLOT). If you specify COLUMN=1, the plot displays the column 1 relative risks or the column 1 risk differences. Similarly, if you specify COLUMN=2, the plot displays the column 2 relative risks or risk differences.

By default for relative risk plots, COLUMN=1. By default for risk difference plots, COLUMN=1 if you request both column 1 and column 2 risk differences by specifying the RISKDIFF option. If you specify a column in the RISKDIFF(COLUMN=) option, the risk difference plot uses this column by default.

COMMON=NO | YES

controls the display of the common (overall) statistic in plots that display stratum (two-way table) statistics for multiway tables. You can specify the COMMON= *plot-option* when you specify the following *plot-requests*: KAPPAPLOT, ODDSRATIOPLOT, RELRISKPLOT, RISKDIFFPLOT, and WTKAPPAPLOT.

COMMON=NO suppresses display of the common statistic and its confidence limits. By default, COMMON=YES, which displays the common statistic and its confidence limits when these values are available. For more information, see the descriptions of the *plot-requests*.

EXACT

requests display of exact confidence limits instead of asymptotic confidence limits. You can specify the EXACT *plot-option* when you specify the following *plot-requests*: ODDSRATIOPLOT, RELRISKPLOT, and RISKDIFFPLOT. 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, which you can request by specifying the FREQPLOT *plot=request*. By default, 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 (RELRISKPLOT). This *plot-option* displays the odds ratio or relative risk axis on the log scale that you specify.

LEGEND=NO | YES

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

NOSTAT

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

NPANELPOS=*n*

divides the plot into multiple panels that display at most $|n|$ statistics or sections.

If n is positive, the number of statistics or sections 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 1 odds ratio in the second panel. This *plot-option* is available for all plots except mosaic plots and one-way weighted frequency plots.

For two-way frequency plots (FREQPLOT), NPANELPOS=n requests that panels display at most $|n|$ sections, where sections correspond to row or column variable levels, depending on the type of plot and the grouping. By default, n=4 and each panel includes at most four sections. This *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.

For plots that display statistics along with confidence limits, NPANELPOS=*n* requests that panels display at most $|n|$ statistics. By default, *n*=0 and all statistics are displayed in a single panel. This *plot-option* applies to the following plots: KAPPAPLOT, ODDSRATIOPLOT, RELRISKPLOT, RISKDIFFPLOT, and WTKAPPAPLOT.

ORDER=ASCENDING | DESCENDING

displays the two-way table (strata) statistics in order of the statistic value. You can specify the ORDER= *plot-option* when you specify the following *plot-requests*: KAPPAPLOT, ODDSRATIOPLOT, RELRISKPLOT, RISKDIFFPLOT, and WTKAPPAPLOT.

If you specify ORDER=ASCENDING or ORDER=DESCENDING, the plot displays the statistics in ascending or descending order, respectively. By default, the order of the statistics in the plot matches the order that the two-way table strata appear in the multiway table display.

ORIENT=HORIZONTAL | VERTICAL

controls the orientation of the plot. You can specify the ORIENT= *plot-option* when you specify the following *plot-requests*: 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=NO | YES

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. By default, PARTIAL=YES.

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

specifies the range of values to display. You can specify the RANGE= *plot-option* when you specify the following *plot-requests*: 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 | GROUPPERCENT | LOG | PERCENT | SQRT

specifies the scale of the frequencies to display. This *plot-option* is available for frequency plots (FREQPLOT) and cumulative frequency plots (CUMFREQPLOT).

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

SCALE=GROUPPERCENT is available for two-way frequency plots. This option displays the row or column percentages instead of the overall percentages (of the table frequency). By default (or when you specify the GROUPBY=COLUMN *plot-option*), SCALE=GROUPPERCENT displays the column percentages. If you specify the GROUPBY=ROW *plot-option*, the primary grouping of graph cells is by row variable level and the plot displays row percentages. For more information, see the description of the GROUPBY= *plot-option*.

SHOWSCALE=NO | YES

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. By default, SHOWSCALE=YES.

STATS

displays statistic values in the plot. For the following *plot-requests*, the STATS *plot-option* displays the statistics and their confidence limits on the right side of the plot: KAPPAPLOT, ODDSRATIOPLOT, RELRISKPLOT, RISKDIFFPLOT, and WTKAPPAPLOT.

For the agreement plot (AGREEPLOT), the STATS *plot-option* displays the values of the kappa statistic, the weighted kappa statistic, the B_n measure (Bangdiwala and Bryan 1987), and the sample size. PROC FREQ stores these statistics in an ODS table named BnMeasure, which is not displayed. For more information, see the section "ODS Table Names" on page 3240.

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

TWOWAY=CLUSTER | GROUPHORIZONTAL | GROUPVERTICAL | 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=GROUPVERTICAL, 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=GROUPHORIZONTAL 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*. By default, GROUPBY=COLUMN, which groups first by column variable.

TYPE=BARCHART | DOTPLOT

specifies the plot type (format) of the frequency (FREQPLOT), cumulative frequency (CUMFREQPLOT), and deviation plots (DEVIATIONPLOT). TYPE=BARCHART produces a bar chart and TYPE=DOTPLOT produces a dot plot. By default, TYPE=BARCHART.

PRINTKWTS

displays the agreement weights that PROC FREQ uses to compute the weighted kappa coefficient. Agreement weights reflect the relative agreement between pairs of variable levels. By default, PROC FREQ uses the Cicchetti-Allison form of agreement weights. If you specify the AGREE(WT=FC) option, the procedure uses the Fleiss-Cohen form of agreement weights. For more information, see the section "Weighted Kappa Coefficient" on page 3207.

This option has no effect unless you also specify the AGREE option to compute the weighted kappa coefficient. The PRINTKWTS option is equivalent to the AGREE(PRINTKWTS) option.

RELRISK < (*relrisk-options*) >

requests relative risk measures and their confidence limits for 2 × 2 tables. These measures include the odds ratio, the column 1 relative risk, and the column 2 relative risk. For more information, see the section "Odds Ratio and Relative Risks" on page 3191. By default, PROC FREQ displays the relative risk measures and their asymptotic Wald confidence limits in the "Odds Ratio and Relative Risks" table. You can also obtain this table by specifying the MEASURES option, which produces other measures of association in addition to the relative risks.

You can specify *relrisk-options* in parentheses after the RELRISK option to request tests and additional confidence limits for the column 1 or column 2 relative risk. Table 47.18 summarizes the *relrisk-options*.

When you request tests or additional confidence limit types for the relative risk, PROC FREQ does not display the "Odds Ratio and Relative Risks" table unless you also specify the PRINTALL *relrisk-option*.

Table 47.18 RELRISK (Relative Risk) Options

Option	Description
COLUMN=1 | 2	Specifies the risk column
PRINTALL	Displays "Odds Ratio and Relative Risks" table
Request Confidence Limits	
CL=EXACT	Displays exact confidence limits
CL=LR	Requests likelihood ratio confidence limits
CL=SCORE	Requests score confidence limits
CL=WALD	Requests Wald confidence limits
CL=WALDMODIFIED	Requests Wald modified confidence limits
Request Tests	
EQUAL(NULL=)	Requests an equality test
EQUIV | EQUIVALENCE	Requests an equivalence test
MARGIN=	Specifies the test margin

Table 47.18 *continued*

Option	Description
METHOD=	Specifies the test method
NONINF | NONINFERIORITY	Requests a noninferiority test
SUP | SUPERIORITY	Requests a superiority test

You can specify the following *relrisk-options*:

CL=*type* | (*types*)

specifies confidence limit types for the relative risk. You can specify one or more *types* of confidence limits. When you specify only one *type*, you can omit the parentheses around the request. When you specify the CL= *relrisk-option*, PROC FREQ displays the confidence limits in the "Confidence Limits for the Relative Risk" table.

The ALPHA= option determines the level of the confidence limits that the CL= *relrisk-option* provides. By default, ALPHA=0.05, which produces 95% confidence limits for the relative risk.

You can specify the following *types*:

EXACT

displays exact unconditional confidence limits for the relative risk in the "Confidence Limits for the Relative Risk" table. (By default, PROC FREQ displays the exact confidence limits in a separate table.) You must also request computation of the exact confidence limits by specifying the RELRISK option in the EXACT statement. For more information, see the subsection "Exact Unconditional Confidence Limits" in the section "Confidence Limits for the Relative Risk" on page 3195.

LR

LIKELIHOOD RATIO

requests likelihood ratio confidence limits for the relative risk. For more information, see the subsection "Likelihood Ratio Confidence Limits" in the section "Confidence Limits for the Relative Risk" on page 3195.

SCORE < (CORRECT=NO) >

requests score confidence limits for the relative risk. For more information, see the subsection "Score Confidence Limits" in the section "Confidence Limits for the Relative Risk" on page 3195. If you specify CORRECT=NO, PROC FREQ provides the uncorrected form of the confidence limits.

WALD

requests asymptotic Wald confidence limits, which are based on a log transformation of the relative risk. For more information, see the subsection "Wald Confidence Limits" in the section "Confidence Limits for the Relative Risk" on page 3195.

WALDMODIFIED

requests Wald modified confidence limits for the relative risk. For more information, see the subsection "Wald Modified Confidence Limits" in the section "Confidence Limits for the Relative Risk" on page 3195.

COLUMN=1 | 2

specifies the table column for which to compute the relative risk confidence limits (which you request by specifying the CL= *relrisk-option*) and the relative risk tests (EQUAL, EQUIV, NONINF, and SUP). By default, COLUMN=1.

This option has no effect on the "Odds Ratio and Relative Risks" table, which displays both column 1 and column 2 relative risks.

EQUAL < (NULL=*value*)>

requests an equality test for the relative risk. For more information, see the subsection "Equality Test" in the section "Relative Risk Tests" on page 3198.

You can specify the null hypothesis *value* of the relative risk in the NULL= option. The null *value* must be a positive number. By default, NULL=1. You can specify the equality test method in the METHOD= *relrisk-option*. By default, PROC FREQ produces a Wald test (METHOD=WALD).

EQUIV

EQUIVALENCE

requests an equivalence test for the relative risk. For more information, see the subsection "Equivalence Test" in the section "Relative Risk Tests" on page 3198. You can specify the test method in the METHOD= *relrisk-option*, and you can specify the test margins in the MARGIN= *relrisk-option*. By default, METHOD=WALD and MARGIN=(0.8,1.25).

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

specifies the margin for the noninferiority, superiority, and equivalence tests, which you request by specifying the NONINF, SUP, and EQUIV *relrisk-options*, respectively. By default, MARGIN=0.8 for noninferiority tests, MARGIN=1.25 for superiority tests, and MARGIN=(0.8,1.25) for equivalence tests.

For noninferiority and superiority tests, specify a single *value* in the MARGIN= option. The *value* must be a positive number. For a noninferiority test, the *value* should be less than 1; for a superiority test, the *value* should be greater than 1.

For an equivalence test, you can specify a single MARGIN= *value*, or you can specify both *lower* and *upper* values. All values must be positive numbers. If you specify a single *value*, PROC FREQ uses *value* as the lower margin and the inverse of *value* as the upper margin. If you specify both *lower* and *upper* values, the value of *lower* must be less than the value of *upper*.

METHOD=*method*

specifies the method to be used for the equality, equivalence, noninferiority, and superiority tests, which you request by specifying the EQUAL, EQUIV, NONINF, and SUP *relrisk-options*, respectively. By default, METHOD=WALD.

You can specify one of the following *methods*:

FM

SCORE

requests Farrington-Manning (score) tests for the equality, equivalence, noninferiority, and superiority analyses of the relative risk. For more information, see the subsection "Farrington-Manning (Score) Test" in the section "Relative Risk Tests" on page 3198.

LR

LIKELIHOODRATIO

requests likelihood ratio tests for the equality, equivalence, noninferiority, and superiority analyses of the relative risk. For more information, see the subsection "Likelihood Ratio Test" in the section "Relative Risk Tests" on page 3198.

WALD

requests Wald tests for the equality, equivalence, noninferiority, and superiority analyses of the relative risk. For more information, see the subsection "Wald Test" in the section "Relative Risk Tests" on page 3198.

WALDMODIFIED

requests Wald modified tests for the equality, equivalence, noninferiority, and superiority analyses of the relative risk. For more information, see the subsection "Wald Modified Test" in the section "Relative Risk Tests" on page 3198.

NONINF

NONINFERIORITY

requests a noninferiority test for the relative risk. For more information, see the subsection "Noninferiority Test" in the section "Relative Risk Tests" on page 3198. You can specify the test method in the METHOD= *relrisk-option*, and you can specify the margin in the MARGIN= *relrisk-option*. By default, METHOD=WALD and MARGIN=0.8.

PRINTALL

displays the "Odds Ratio and Relative Risks" table when you request tests or additional confidence limits by specifying *relrisk-options*. By default, PROC FREQ does not display this table when you request tests or additional confidence limits for the relative risk.

SUP

SUPERIORITY

requests a superiority test for the relative risk. For more information, see the subsection "Superiority Test" in the section "Relative Risk Tests" on page 3198. You can specify the test method in the METHOD= *relrisk-option*, and you can specify the margin in the MARGIN= *relrisk-option*. By default, METHOD=WALD and MARGIN=1.25.

RISKDIFF < (*riskdiff-options*) >

requests risks (binomial proportions) and risk differences for 2 × 2 tables. By default, this option 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; by default, this option also provides exact (Clopper-Pearson) confidence limits for the row 1, row 2, and total risks. You can request exact unconditional confidence limits for the risk difference by specifying the RISKDIFF option in the EXACT statement. PROC FREQ displays these results in the column 1 and column 2 "Risk Estimates" tables (which you can suppress by specifying the NORISKS *riskdiff-option*).

You can specify *riskdiff-options* in parentheses after the RISKDIFF option to request tests and additional confidence limits for the risk difference, in addition to estimates of the common risk difference for multiway 2 × 2 tables. Table 47.19 summarizes the *riskdiff-options*.

The CL= *riskdiff-option* requests confidence limits for the risk difference. Available confidence limit types include Agresti-Caffo, exact unconditional, Hauck-Anderson, Miettinen-Nurminen (score), 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 "Confidence Limits for the Risk Difference" table.

The CL=EXACT *riskdiff-option* displays exact unconditional confidence limits in the "Confidence Limits for the Risk Difference" table. When you specify CL=EXACT, you must also request computation of the exact confidence limits by specifying the RISKDIFF option in the EXACT statement.

The EQUAL, EQUIV, NONINF, and SUP *riskdiff-options* request tests of equality, equivalence, noninferiority, and superiority, respectively, for the risk difference. Available test methods include Farrington-Manning (score), Hauck-Anderson, and Wald. Newcombe (hybrid-score) confidence limits are available for the equivalence, noninferiority, and superiority analyses.

As part of the noninferiority, superiority, and equivalence analyses, PROC FREQ provides null-based equivalence limits that have a confidence coefficient of $100(1 - 2\alpha)\%$ (Schuirmann 1999). The ALPHA= option determines the confidence level; by default, ALPHA=0.05, which produces 90% equivalence limits for these analyses. For more information, see the sections "Noninferiority Tests" on page 3183 and "Equivalence Test" on page 3185.

Table 47.19 RISKDIFF (Proportion Difference) Options

Option	Description
COLUMN=1 | 2	Specifies the risk column
COMMON	Requests common risk difference
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 | SCORE	Requests Miettinen-Nurminen confidence limits
CL=NEWCOMBE	Requests Newcombe confidence limits
CL=WALD	Requests Wald confidence limits
Request Tests	
EQUAL(NULL=)	Requests an equality test
EQUIV | EQUIVALENCE	Requests an equivalence test
MARGIN=	Specifies the test margin
METHOD=	Specifies the test method
NONINF | NONINFERIORITY	Requests a noninferiority test
SUP | SUPERIORITY	Requests a superiority test
VAR=SAMPLE | NULL	Specifies the test variance

You can specify the following *riskdiff-options*:

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 "Confidence Limits for the Risk Difference" table.

The ALPHA= option determines the level of the confidence limits. By default, ALPHA=0.05, which produces 95% confidence limits for the risk difference.

You can specify the CL= *riskdiff-option* with or without requests for risk difference 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 can specify in the MARGIN= *riskdiff-option*).

You can specify the following *types*:

AC

AGRESTICAFFO

requests Agresti-Caffo confidence limits for the risk difference. For more information, see the subsection "Agresti-Caffo Confidence Limits" in the section "Confidence Limits for the Risk Difference" on page 3178.

EXACT

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

By default, PROC FREQ computes the exact confidence limits by inverting two separate one-sided exact tests that are based on the score statistic. For more information, see the RISKDIFF option in the EXACT statement and the subsection "Exact Unconditional Confidence Limits" in the section "Confidence Limits for the Risk Difference" on page 3178.

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. For more information, see the subsection "Hauck-Anderson Confidence Limits" in the section "Confidence Limits for the Risk Difference" on page 3178.

MN < (CORRECT=NO | MEE) >

SCORE < (CORRECT=NO | MEE) >

requests Miettinen-Nurminen (score) confidence limits for the risk difference. For more information, see the subsection "Miettinen-Nurminen (Score) Confidence Limits" in the section "Confidence Limits for the Risk Difference" on page 3178. 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 < (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. For more information, see the subsection "Newcombe Confidence Limits" in the section "Confidence Limits for the Risk Difference" on page 3178.

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. For more information, see the subsection "Wald Confidence Limits" in the section "Confidence Limits for the Risk Difference" on page 3178.

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 you request by specifying the CL= *riskdiff-option*). By default, COLUMN=1.

This 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*.

COMMON

requests estimates of the common (overall) risk difference for multiway 2 × 2 tables. This option provides Mantel-Haenszel and summary score estimates for the common risk difference, together with their confidence limits. If you specify the RISKDIFF(CL=NEWCOMBE) option, the RISKDIFF(COMMON) option also provides Newcombe confidence limits for the common risk difference. For more information, see the section "Common Risk Difference" on page 3187.

You can use the COMMONRISKDIFF option to request additional confidence limits and tests for the common risk difference.

If you do not specify the COLUMN= *riskdiff-option*, PROC FREQ provides the common risk difference for column 1 by default. If you specify COLUMN=2, PROC FREQ provides the common risk difference for column 2. COLUMN=BOTH does not apply to the common risk difference.

CORRECT

includes a continuity correction in the Wald confidence limits, Wald tests, and Newcombe confidence limits. For more information, see the section "Risks and Risk Differences" on page 3176.

EQUAL < (NULL=*value*)>

requests an equality test for the risk difference. For more information, see the section "Equality Tests" on page 3182.

You can specify the null hypothesis *value* of the risk difference in the NULL= option. By default, NULL=0. You can specify the null *value* in proportion form as a number between –1 and 1, or you can specify the null *value* in percentage form as a number between –100 and 100. When the *value* is between –100 and –1 or between 1 and 100, PROC FREQ converts the number to a proportion. PROC FREQ treats the values –1 and 1 as percentages.

You can specify the equality test method by using the METHOD= *riskdiff-option*. By default, PROC FREQ produces a Wald test (METHOD=WALD). By default, PROC FREQ uses the sample variance to compute the Wald test statistic; if you specify the VAR=NULL *riskdiff-option*, PROC FREQ uses the null (test-based) variance.

EQUIV

EQUIVALENCE

requests an equivalence test for the risk difference. For more information, see the section "Equivalence Test" on page 3185. You can specify the test method in the METHOD= *riskdiff-option*, and you can specify the margins in the MARGIN= *riskdiff-option*. By default, METHOD=WALD and MARGIN=0.2.

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

specifies the margin for the noninferiority, superiority, and equivalence tests, which you request by specifying the NONINF, SUP, and EQUIV *riskdiff-options*, respectively. By default, MARGIN=0.2.

For noninferiority and superiority tests, specify a single *value* in the MARGIN= option. The *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. PROC FREQ 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 *value*, it must be a positive number, as described previously. If you specify a single *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 and 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 to be used for the equality, equivalence, noninferiority, and superiority tests, which you request by specifying the EQUAL, EQUIV, NONINF, and SUP *riskdiff-options*, respectively. By default, METHOD=WALD.

You can specify the following *methods*:

FM

SCORE

requests Farrington-Manning (score) tests for the equality, equivalence, noninferiority, and superiority analyses. For more information, see the subsection "Farrington-Manning (Score) Test" in the section "Noninferiority Tests" on page 3183.

HA

requests Hauck-Anderson tests for the equality, equivalence, noninferiority, and superiority analyses. For more information, see the subsection "Hauck-Anderson Test" in the section "Noninferiority Tests" on page 3183.

NEWCOMBE

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. For more information, see the subsection "Newcombe Noninferiority Analysis" in the section "Noninferiority Tests" on page 3183.

WALD

requests Wald tests for the equality, equivalence, noninferiority, and superiority analyses. By default, PROC FREQ uses the sample variance to compute these Wald test statistics; if you specify the VAR=NULL *riskdiff-option*, PROC FREQ uses the null (test-based) variance. If you specify the CORRECT *riskdiff-option*, the Wald tests and confidence limits include a continuity correction. For more information, see the subsection "Wald Test" in the section "Noninferiority Tests" on page 3183.

NONINF

NONINFERIORITY

requests a noninferiority test for the risk difference. For more information, see the section "Noninferiority Tests" on page 3183. You can specify the test method in the METHOD= *riskdiff-option*, and you can specify the margin in the MARGIN= *riskdiff-option*. By default, METHOD=WALD and MARGIN=0.2.

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 superiority test for the risk difference. For more information, see the section "Superiority Test" on page 3185. You can specify the test method in the METHOD= *riskdiff-option*, and you can specify the margin in the MARGIN= *riskdiff-option*. By default, METHOD=WALD and MARGIN=0.2.

VAR=NULL | SAMPLE

specifies the type of variance to use in the Wald tests of equality, equivalence, noninferiority, and superiority. By default (or if you specify VAR=SAMPLE), PROC FREQ uses the sample variance to compute the Wald test statistics. If you specify VAR=NULL, PROC FREQ uses the null (test-based) variance. For more information, see the sections "Equality Tests" on page 3182 and "Noninferiority Tests" on page 3183.

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 3152 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. For more information, see the section “Scores” on page 3152. 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 by using the SCORES= option. For more information, see the section “Scores” on page 3152.

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 3240 for more information.

SENSPEC

requests estimates of sensitivity, specificity, positive predictive value, and negative predictive value for 2×2 tables. The “Sensitivity and Specificity” table provides these estimates together with their standard errors and confidence limits. For more information, see the section “Sensitivity and Specificity” on page 3176.

You can specify the confidence level in the ALPHA= option. By default, ALPHA=0.05, which produces 95% confidence limits.

SPARSE

reports all possible combinations of variable values in two-way and multiway tables, even if a combination does not occur in the data. This option applies only to crosstabulation tables that are displayed in LIST format and to OUT= output data sets.

When you specify the SPARSE option together with the LIST option, the LIST tables display all combinations of variable values, including levels that have a frequency of 0. By default, LIST tables do not display zero-frequency levels. When you specify the SPARSE option together with the OUT= option, the OUT= output data set includes empty (zero-frequency) crosstabulation table cells. By default, the OUT= output data set does not include zero-frequency table cells.

For more information, see the section “Missing Values” on page 3148.

TOTPCT

displays the percentage of the total multiway table frequency in multiway crosstabulation tables (n-way tables, where $n > 2$). By default, crosstabulation tables display percentages of the two-way table (stratum) frequency but do not display percentages of the total multiway table frequency. For more information, see the section “Two-Way and Multiway Tables” on page 3232.

By default, tables in list format (which you can request by specifying the LIST option) display the percentages of the total multiway table frequency. The variable `PERCENT` in the OUT= output data set also provides the percentages of the total multiway table frequency.

TREND

requests the Cochran-Armitage test for trend. The table must be $2 \times C$ or $R \times 2$ to compute the trend test. For more information, see the section "Cochran-Armitage Test for Trend" on page 3201. To request exact p-values for the trend test, specify the TREND option in the EXACT statement. See the section "Exact Statistics" on page 3222 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 47.20 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 is 0. The procedure displays the asymptotic standard error under the null hypothesis, the test statistic, and the one-sided and two-sided p-values. PROC FREQ also provides confidence limits for the measure. The ALPHA= option in the TABLES statement determines the confidence level; by default, ALPHA=0.05, which provides 95% confidence limits. For more information, see the sections "Asymptotic Tests" on page 3159 and "Confidence Limits" on page 3158. For information about the individual measures, see the sections "Measures of Association" on page 3158 and "Tests and Measures of Agreement" on page 3204.

You can also request exact tests for selected measures of association and agreement by using the EXACT statement. For more information, see the section "Exact Statistics" on page 3222.

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 specify in the TEST statement.

Table 47.20 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 | TAUB	Kendall's tau-b	ALL or MEASURES
MEASURES	Gamma, Kendall's tau-b, Stuart's tau-c, Somers' $D(C\|R)$, Somers' $D(R\|C)$, Pearson and Spearman correlations	ALL or MEASURES
PCORR	Pearson correlation coefficient	ALL or MEASURES

Table 47.20 *continued*

Test Option	Asymptotic Tests	Required TABLES Statement Option
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 | TAUC	Stuart's tau-c	ALL or MEASURES
WTKAPPA | WTKAP	Weighted kappa coefficient	AGREE

You can specify the following *test-options*:

AGREE

requests asymptotic tests for the simple kappa coefficient and the weighted kappa coefficient. For more information, see the sections "Simple Kappa Coefficient" on page 3205 and "Weighted Kappa Coefficient" on page 3207.

By default, these tests are based on null values of 0; you can specify nonzero null values for the simple kappa and weighted kappa tests by using the AGREE(NULLKAPPA=) and AGREE(NULLWTKAPPA=) options, respectively, in the TABLES statement.

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. For more information, see the section "Gamma" on page 3160. 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. For more information, see the section "Simple Kappa Coefficient" on page 3205.

By default, the null value of kappa for this test is 0; you can specify a nonzero null value by using the AGREE(NULLKAPPA=) option in the TABLES statement.

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. For more information, see the section "Kendall's Tau-b" on page 3160.

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. For more information, see the section "Measures of Association" on page 3158.

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. For more information, see the section "Pearson Correlation Coefficient" on page 3162.

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. For more information, see the section "Polychoric Correlation" on page 3164.

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. For more information, see the section "Spearman Rank Correlation Coefficient" on page 3163.

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' $D(C|R)$. For more information, see the section "Somers' D" on page 3161.

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' $D(R|C)$. For more information, see the section "Somers' D" on page 3161.

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*. For more information, see the section "Stuart's Tau-c" on page 3161.

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.

WTKAPPA

WTKAP

requests an asymptotic test for the weighted kappa coefficient. For more information, see the section "Weighted Kappa Coefficient" on page 3207.

By default, the null value of weighted kappa for this test is 0; you can specify a nonzero null value by using the AGREE(NULLWTKAPPA=) option in the TABLES statement.

The AGREE option in the TABLES statement provides the weighted kappa coefficient, its standard error, and confidence limits. You can request an exact test for the weighted kappa by specifying the WTKAPPA 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 3146 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 0, PROC FREQ ignores the observation unless you specify the ZEROS option, which includes observations that have weights of 0. If you do not specify a WEIGHT statement, PROC FREQ assigns a weight of 1 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*:

ZEROS

includes observations that have weights of 0. By default, PROC FREQ ignores observations that have weights of 0.

If you specify the ZEROS option, frequency and crosstabulation tables display levels that contain only zero-weight observations. If you do not specify the ZEROS option, PROC FREQ does not process observations that have weights of 0 and therefore does not display levels that contain only zero-weight observations.

When you specify the ZEROS option, PROC FREQ includes zero-weight levels in chi-square tests and binomial computations for one-way tables. This makes it possible to compute binomial tests and estimates for a reference level that 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 weights. For more information, see the section "Tables with Zero-Weight Rows or Columns" on page 3211.

Even when you specify the ZEROS option, PROC FREQ does not compute CHISQ or MEASURES statistics for two-way tables that contain a zero-weight row or column because most of these statistics are undefined in this case.

By default, the ZEROS option invokes the SPARSE option in the TABLES statement, which includes zero-weight table cells in the LIST table and OUT= data set. To suppress zero-weight cells, 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, 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 following 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 47.12 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 following 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 47.12 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				

Figure 47.12 *continued*

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 in a two-way table is missing, PROC FREQ assigns 0 to the frequency count of the corresponding table cell. By default, PROC FREQ does not include missing combinations in the LIST display or the OUT= output data set. To include missing combinations in the LIST display and the OUT= output data set, you can specify the SPARSE 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). PROC FREQ supports the following database management systems: Aster, DB2, Greenplum, Hadoop, HAWQ, Impala, Netazza, Oracle, SAP HANA, and Teradata. 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). Instead of transferring the entire data set over the network between the database and SAS software, in-database computation 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). In addition, 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. For information about the SQLGENERATION option and other options that affect in-database computation, see the section "In-Database Procedures" in *SAS/ACCESS for Relational Databases: Reference*. By default, PROC FREQ uses in-database computation when possible. PROC FREQ has no procedure options that control in-database computation.

PROC FREQ uses formatted values to group observations into the levels of frequency and crosstabulation tables. For more information, see the section "Grouping with Formats" on page 3147. If formats are available in the database, in-database summarization uses the formats. If formats are not available in the database, the in-database summarization uses 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 $X_i,\ i = 1, 2, \ldots, R$, and the column values by $Y_j,\ j = 1, 2, \ldots, C$. Let n_{ij} denote the frequency of the table cell in the *i*th row and *j*th column and define the following notation:

$$\begin{aligned}
n_{i\cdot} &= \sum_j n_{ij} && \text{(row totals)} \\
n_{\cdot j} &= \sum_i n_{ij} && \text{(column totals)} \\
n &= \sum_i \sum_j n_{ij} && \text{(overall total)} \\
p_{ij} &= n_{ij}/n && \text{(cell percentages)} \\
p_{i\cdot} &= n_{i\cdot}/n && \text{(row percentages of total)} \\
p_{\cdot j} &= n_{\cdot j}/n && \text{(column percentages of total)}
\end{aligned}$$

$$\begin{aligned}
R_i &= \text{score for row } i \\
C_j &= \text{score for column } j
\end{aligned}$$

$$\begin{aligned}
\bar{R} &= \sum_i n_{i\cdot} R_i / n && \text{(average row score)} \\
\bar{C} &= \sum_j n_{\cdot j} C_j / n && \text{(average column score)}
\end{aligned}$$

$$
\begin{aligned}
A_{ij} &= \sum_{k>i} \sum_{l>j} n_{kl} + \sum_{k<i} \sum_{l<j} n_{kl} \\
D_{ij} &= \sum_{k>i} \sum_{l<j} n_{kl} + \sum_{k<i} \sum_{l>j} n_{kl} \\
P &= \sum_{i} \sum_{j} n_{ij} A_{ij} \quad \text{(twice the number of concordances)} \\
Q &= \sum_{i} \sum_{j} n_{ij} D_{ij} \quad \text{(twice the number of discordances)}
\end{aligned}
$$

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

$$
\begin{aligned}
R^1_i &= \sum_{k<i} n_{k\cdot} + (n_{i\cdot} + 1)/2 \qquad i = 1, 2, \ldots, R \\
C^1_j &= \sum_{l<j} n_{\cdot l} + (n_{\cdot j} + 1)/2 \qquad j = 1, 2, \ldots, C
\end{aligned}
$$

where R^1_i is the rank score of row i, and C^1_j 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

$$
\begin{aligned}
R^2_i &= R^1_i / n \qquad i = 1, 2, \ldots, R \\
C^2_j &= C^1_j / n \qquad j = 1, 2, \ldots, C
\end{aligned}
$$

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

$$
\begin{aligned}
R^3_i &= R^1_i / (n + 1) \qquad i = 1, 2, \ldots, R \\
C^3_j &= C^1_j / (n + 1) \qquad j = 1, 2, \ldots, C
\end{aligned}
$$

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(LRCHI) 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 3222 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 $i = 1, 2, \ldots, C$. 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,

$$e_i = n \;/\; C \qquad \text{for} \;\; i = 1, 2, \ldots, C$$

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

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

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 3222 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,

$$e_{ij} = (n_{i\cdot} \times n_{\cdot j}) \; / \; n$$

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 3222 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. For more information, see Fienberg (1980).

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 3154 for more information.

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

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

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

$$e_{ij} = (n_{i\cdot} \times n_{\cdot j}) \; / \; n$$

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(LRCHI) 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

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

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 is the total sample size divided by the number of classes,

$$e_i = n \;/\; C \qquad \text{for} \;\; i = 1, 2, \ldots, C$$

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(TESTP=) option in the TABLES statement, PROC FREQ computes the expected frequencies as

$$e_i = p_i \times n \qquad \text{for} \;\; i = 1, 2, \ldots, C$$

where the proportions p_i are the TESTP= 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 LRCHI option in the EXACT statement. See the section "Exact Statistics" on page 3222 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

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

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 3222 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

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

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

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

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 3162. 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 1 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 3222 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 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 is 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,

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

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 is $(n_{11}n_{22}/n_{12}n_{21})$. The left-sided alternative is also 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, $p_{1|1} > p_{1|2}$. For more information, see Agresti (2007).

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 3222 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

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

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

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

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

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

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, is 0.05 and produces 95% confidence limits.

The confidence limits are computed as

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

where Est is the estimate of the measure, $z_{\alpha/2}$ is the $100(1-\alpha/2)$th 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 is 0. 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

$$z = \mathrm{Est} \ / \ \sqrt{\mathrm{Var_0(Est)}}$$

where Est is the estimate of the measure and $\mathrm{Var_0(Est)}$ is the variance of the estimate under the null hypothesis. Formulas for $\mathrm{Var_0(Est)}$ for the individual measures of association are given in the following sections.

Note that the ratio of Est to $\sqrt{\mathrm{Var_0(Est)}}$ 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 $H_0\colon \mathrm{gamma} = 0$ equal the *p*-values for the test of $H_0\colon \mathrm{tau} - b = 0$.

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 0, 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 0. When the test statistic is less than or equal to 0, 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 0. The one-sided *p*-value P_1 can be expressed as

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

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

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

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 is 0. For more information, see the section "Exact Statistics" on page 3222.

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, gamma tends to be close to 0. Gamma is computed as

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

and the asymptotic variance is

$$\mathrm{Var}(G) = \frac{16}{(P + Q)^4} \sum_i \sum_j n_{ij} (Q A_{ij} - P D_{ij})^2$$

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 0 is computed as

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

For more information, see Brown and Benedetti (1977b).

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

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

where

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

See Kendall (1955) for more information.

The variance under the null hypothesis that tau-*b* equals 0 is computed as

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

For more information, see Brown and Benedetti (1977b).

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 3222 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 $-1 \le \tau_c \le 1$. Stuart's tau-*c* is computed as

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

and the asymptotic variance is

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

where $m = \min(R, C)$ and $d_{ij} = A_{ij} - D_{ij}$. The variance under the null hypothesis that tau-*c* equals 0 is the same as the asymptotic variance

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

For more information, see Brown and Benedetti (1977b).

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 3222 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 $-1 \le D \le 1$. Somers' $D(C|R)$ is computed as

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

and its asymptotic variance is

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

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

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

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

The variance under the null hypothesis that $D(C|R)$ equals 0 is computed as

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

For more information, see Brown and Benedetti (1977b).

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 3222 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

$$r = v/w = s_{rc}/\sqrt{s_r s_c}$$

and its asymptotic variance is

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

where R_i and C_j are the row and column scores and

$$\begin{aligned}
s_r &= \sum_i\sum_j n_{ij}(R_i - \bar{R})^2 \\
s_c &= \sum_i\sum_j n_{ij}(C_j - \bar{C})^2 \\
s_{rc} &= \sum_i\sum_j n_{ij}(R_i - \bar{R})(C_j - \bar{C})
\end{aligned}$$

$$\begin{aligned}
b_{ij} &= (R_i - \bar{R})^2 s_c + (C_j - \bar{C})^2 s_r \\
v &= s_{rc} \\
w &= \sqrt{s_r s_c}
\end{aligned}$$

For more information, see Snedecor and Cochran (1989).

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 3152 for details about the available score types and how they are computed.

The variance under the null hypothesis that the correlation equals 0 is computed as

$$\mathrm{Var}_0(r) = \left(\sum_i \sum_j n_{ij} (R_i - \bar{R})^2 (C_j - \bar{C})^2 - s_{rc}^2 / n \right) \Big/ \, s_r s_c$$

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. For more information, see Brown and Benedetti (1977b).

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 3222 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 3152. This measure is appropriate only when both variables lie on an ordinal scale. The range of the Spearman correlation is $-1 \le \rho_s \le 1$. The Spearman correlation coefficient is computed as

$$r_s = v \, / \, w$$

and its asymptotic variance is

$$\mathrm{Var}(r_s) = \frac{1}{n^2 w^4} \sum_i \sum_j n_{ij} (z_{ij} - \bar{z})^2$$

where R_i and C_j are the row and column rank scores and

$$\begin{aligned}
v &= \sum_i \sum_j n_{ij} R(i) C(j) \\
w &= \frac{1}{12} \sqrt{FG} \\
F &= n^3 - \sum_i n_{i\cdot}^3 \\
G &= n^3 - \sum_j n_{\cdot j}^3 \\
R(i) &= R_i - n/2 \\
C(j) &= C_j - n/2 \\
\bar{z} &= \frac{1}{n} \sum_i \sum_j n_{ij} z_{ij} \\
z_{ij} &= w v_{ij} - v w_{ij}
\end{aligned}$$

$$
\begin{aligned}
v_{ij} &= n\Biggl(R(i)C(j) + \frac{1}{2}\sum_l n_{il}C(l) + \frac{1}{2}\sum_k n_{kj}R(k) + \\
&\qquad \sum_l \sum_{k>i} n_{kl}C(l) + \sum_k \sum_{l>j} n_{kl}R(k)\Biggr) \\
w_{ij} &= \frac{-n}{96w}\left(Fn_{\cdot j}^2 + Gn_{i\cdot}^2\right)
\end{aligned}
$$

For more information, see Snedecor and Cochran (1989).

The variance under the null hypothesis that the correlation equals 0 is computed as

$$
\mathrm{Var}_0(r_s) = \frac{1}{n^2w^2}\sum_i \sum_j n_{ij}(v_{ij} - \bar{v})^2
$$

where

$$
\bar{v} = \sum_i \sum_j n_{ij}v_{ij}/n
$$

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. For more information, see Brown and Benedetti (1977b).

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. For more information, see the section "Exact Statistics" on page 3222.

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

$$\hat{\rho} \;\pm\; (\, z_{\alpha/2} \times \mathrm{SE}(\hat{\rho})\,)$$

where $\hat{\rho}$ is the estimate of the polychoric correlation, $z_{\alpha/2}$ is the $100(1-\alpha/2)$th percentile of the standard normal distribution, and $\mathrm{SE}(\hat{\rho})$ 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 is 0. The Wald test statistic is computed as

$$z = \hat{\rho} \;/\; \mathrm{SE}(\hat{\rho})$$

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

$$G^2 = -2\ \log(L_0/L_1)$$

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 1 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 $0 \le \lambda(C|R) \le 1$. Asymmetric lambda $(C|R)$ is computed as

$$\lambda(C|R) = \frac{\sum_i r_i - r}{n - r}$$

and its asymptotic variance is

$$\mathrm{Var}(\lambda(C|R)) = \frac{n - \sum_i r_i}{(n-r)^3} \left(\sum_i r_i + r - 2 \sum_i (r_i \mid l_i = l) \right)$$

where

$$\begin{aligned} r_i &= \max_j(n_{ij}) \\ r &= \max_j(n_{\cdot j}) \\ c_j &= \max_i(n_{ij}) \\ c &= \max_i(n_{i\cdot}) \end{aligned}$$

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 $r = n_{\cdot j}$. 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 $r = n_{\cdot j}$.

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 $cs_j = -1$, l_i is defined to be the smallest such value of j, and cs_j is set equal to i. Otherwise, if $n_{il} = r_i$, l_i is defined to be equal to l. If neither condition is true, l_i is taken to be the smallest value of j such that $n_{ij} = r_i$.

The formulas for lambda asymmetric $(R|C)$ can be obtained by interchanging the indices.

For more information, see Goodman and Kruskal (1979).

Lambda (Symmetric)

The nondirectional lambda is the average of the two asymmetric lambdas, $\lambda(C|R)$ and $\lambda(R|C)$. Its range is $0 \le \lambda \le 1$. 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

$$\mathrm{Var}(\lambda) = \frac{1}{w^4}\Big(wvy - 2w^2\big(n - \sum_i \sum_j (n_{ij} \mid j = l_i, i = k_j)\big) - 2v^2(n - n_{kl})\Big)$$

where

$$\begin{aligned}
r_i &= \max_j(n_{ij}) \\
r &= \max_j(n_{\cdot j}) \\
c_j &= \max_i(n_{ij}) \\
c &= \max_i(n_{i\cdot}) \\
w &= 2n - r - c \\
v &= 2n - \sum_i r_i - \sum_j c_j \\
x &= \sum_i (r_i \mid l_i = l) \; + \; \sum_j (c_j \mid k_j = k) \; + \; r_k \; + \; c_l \\
y &= 8n - w - v - 2x
\end{aligned}$$

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 $(R|C)$.

For more information, see Goodman and Kruskal (1979).

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

$$\mathrm{Var}(U(C|R)) = \frac{1}{n^2 w^4} \sum_i \sum_j n_{ij} \left(H(Y) \log \left(\frac{n_{ij}}{n_{i\cdot}} \right) + (H(X) - H(XY)) \log \left(\frac{n_{\cdot j}}{n} \right) \right)^2$$

where

$$\begin{aligned} v &= H(X) + H(Y) - H(XY) \\ w &= H(Y) \\ H(X) &= -\sum_i \left(\frac{n_{i\cdot}}{n} \right) \log \left(\frac{n_{i\cdot}}{n} \right) \\ H(Y) &= -\sum_j \left(\frac{n_{\cdot j}}{n} \right) \log \left(\frac{n_{\cdot j}}{n} \right) \\ H(XY) &= -\sum_i \sum_j \left(\frac{n_{ij}}{n} \right) \log \left(\frac{n_{ij}}{n} \right) \end{aligned}$$

The formulas for the uncertainty coefficient $U(R|C)$ can be obtained by interchanging the indices.

For more information, see Theil (1972, pp. 115–120) and Goodman and Kruskal (1979).

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

$$\mathrm{Var}(U) = 4 \sum_i \sum_j \frac{n_{ij} \left(H(XY) \log \left(\frac{n_{i\cdot} n_{\cdot j}}{n^2} \right) - (H(X) + H(Y)) \log \left(\frac{n_{ij}}{n} \right) \right)^2}{n^2 \, (H(X) + H(Y))^4}$$

where $H(X)$, $H(Y)$, and $H(XY)$ are defined in the previous section. For more information, see Goodman and Kruskal (1979).

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

$$\mathrm{se}(\hat{p}) = \sqrt{\hat{p}\;(1-\hat{p})\;/\;n}$$

Binomial Confidence Limits

PROC FREQ provides Wald and exact (Clopper-Pearson) confidence limits for the binomial proportion. You can also request the following binomial confidence limit types by specifying the BINOMIAL(CL=) option: Agresti-Coull, Blaker, Jeffreys, exact mid-p, likelihood ratio, logit, and Wilson (score). For more information, see Brown, Cai, and DasGupta (2001), Agresti and Coull (1998), and Newcombe (1998b), in addition to the references cited for each confidence limit type.

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

$$\hat{p} \;\pm\; (\, z_{\alpha/2} \times \mathrm{se}(\hat{p}) \,)$$

where $z_{\alpha/2}$ is the $100(1-\alpha/2)$th percentile of the standard normal distribution. The confidence level α is determined by the ALPHA= option; by default, ALPHA=0.05, which 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

$$\hat{p} \;\pm\; (\, z_{\alpha/2} \times \mathrm{se}(\hat{p}) \;+\; (1/2n) \,)$$

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 $n_1 = 1, 2, \ldots n-1$:

$$\sum_{x=n_1}^{n} \binom{n}{x} P_L^{\,x} (1-P_L)^{\,n-x} \;=\; \alpha/2$$

$$\sum_{x=0}^{n_1} \binom{n}{x} P_U^{\,x} (1-P_U)^{\,n-x} \;=\; \alpha/2$$

The lower confidence limit is 0 when $n_1 = 0$, and the upper confidence limit is 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 $(\alpha/2)$th 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), and Leemis and Trivedi (1996).

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

$$\begin{aligned} \tilde{n}_1 &= n_1 + z^2_{\alpha/2}/2 \\ \tilde{n} &= n + z^2_{\alpha/2} \\ \tilde{p} &= \tilde{n}_1 \,/\, \tilde{n} \end{aligned}$$

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).

Blaker Confidence Limits If you specify the CL=BLAKER *binomial-option*, PROC FREQ computes Blaker confidence limits for the binomial proportion, which are constructed by inverting the two-sided exact Blaker test (Blaker 2000). The $100(1-\alpha)\%$ Blaker confidence interval consists of all values of the proportion p_0 for which the test statistic $B(p_0, n_1)$ falls in the acceptance region,

$$\{p_0 : B(p_0, n_1) > \alpha\}$$

where

$$B(p_0, n_1) \;=\; \text{Prob}(\, \gamma(p_0, X) \le \gamma(p_0, n_1) \mid p_0\,)$$

$$\gamma(p_0, n_1) \;=\; \min(\, \text{Prob}(\, X \ge n_1 \mid p_0\,), \; \text{Prob}(\, X \le n_1 \mid p_0\,)\,)$$

and X is a binomial random variable. For more information, see Blaker (2000).

Jeffreys Confidence Limits If you specify the CL=JEFFREYS *binomial-option*, PROC FREQ computes Jeffreys confidence limits for the binomial proportion as

$$\left(\beta(\alpha/2,\ n_1 + 1/2,\ n - n_1 + 1/2),\ \ \beta(1 - \alpha/2,\ n_1 + 1/2,\ n - n_1 + 1/2) \right)$$

where $\beta(\alpha, b, c)$ is the αth 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. For more information, see Brown, Cai, and DasGupta (2001). For information about using beta priors for inference on the binomial proportion, see Berger (1985).

Likelihood Ratio Confidence Limits If you specify the CL=LIKELIHOODRATIO *binomial-option*, PROC FREQ computes likelihood ratio confidence limits for the binomial proportion by inverting the likelihood ratio test. The likelihood ratio test statistic for the null hypothesis that the proportion equals p_0 can be expressed as

$$L(p_0) \ = \ -2\,(n_1 \log(\hat{p}/p_0) \ + \ (n - n_1) \log((1 - \hat{p})/(1 - p_0)))$$

The $100(1 - \alpha)\%$ likelihood ratio confidence interval consists of all values of p_0 for which the test statistic $L(p_0)$ falls in the acceptance region,

$$\{p_0 : L(p_0) < \chi^2_{1,\alpha}\}$$

where $\chi^2_{1,\alpha}$ is the $100(1 - \alpha)$th percentile of the chi-square distribution with 1 degree of freedom. PROC FREQ finds the confidence limits by iterative computation. For more information, see Fleiss, Levin, and Paik (2003), Brown, Cai, and DasGupta (2001), Agresti (2013), and Newcombe (1998b).

Logit Confidence Limits If you specify the CL=LOGIT *binomial-option*, PROC FREQ computes logit confidence limits for the binomial proportion, which are based on the logit transformation $Y = \log(\hat{p}/(1-\hat{p}))$. Approximate confidence limits for Y are computed as

$$Y_L \ = \ \log(\hat{p}/(1 - \hat{p})) \ - \ z_{\alpha/2} \sqrt{n/(n_1(n - n_1))}$$

$$Y_U \ = \ \log(\hat{p}/(1 - \hat{p})) \ + \ z_{\alpha/2} \sqrt{n/(n_1(n - n_1))}$$

The confidence limits for Y are inverted to produce $100(1 - \alpha)\%$ logit confidence limits P_L and P_U for the binomial proportion p as

$$P_L \ = \ \exp(Y_L/(1 + \exp(Y_L))$$

$$P_U \ = \ \exp(Y_U/(1 + \exp(Y_U))$$

For more information, see Brown, Cai, and DasGupta (2001) and Korn and Graubard (1998).

Mid-*p* Confidence Limits If you specify the CL=MIDP *binomial-option*, PROC FREQ computes exact mid-*p* confidence limits for the binomial proportion by inverting two one-sided binomial tests that include mid-*p* tail areas. The mid-*p* approach replaces the probability of the observed frequency by half of that probability in the Clopper-Pearson sum, which is described in the section "Exact (Clopper-Pearson) Confidence Limits" on page 3168. The exact mid-p confidence limits P_L and P_U are the solutions to the equations

$$\begin{aligned} \sum_{x=n_1+1}^{n} \binom{n}{x} P_L^{\,x} (1-P_L)^{\,n-x} \;+\; \frac{1}{2}\binom{n}{n_1} P_L^{\,n_1}(1-P_L)^{\,n-n_1} &= \alpha/2 \\ \sum_{x=0}^{n_1-1} \binom{n}{x} P_U^{\,x} (1-P_U)^{\,n-x} \;+\; \frac{1}{2}\binom{n}{n_1} P_U^{\,n_1}(1-P_U)^{\,n-n_1} &= \alpha/2 \end{aligned}$$

For more information, see Agresti and Gottard (2007), Agresti (2013), Newcombe (1998b), and Brown, Cai, and DasGupta (2001).

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

$$\left(\hat{p} \;+\; z^2_{\alpha/2}/2n \;\pm\; z_{\alpha/2} \sqrt{ \left(\hat{p}(1-\hat{p}) + z^2_{\alpha/2}/4n \right) / n } \;\right) \;/\; \left(1 + z^2_{\alpha/2}/n \right)$$

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), and Newcombe (1998b).

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

$$z = (\hat{p} - p_0)/\mathrm{se}$$

By default, the standard error is based on the null hypothesis proportion as

$$\mathrm{se} = \sqrt{p_0(1-p_0)/n}$$

If you specify the VAR=SAMPLE *binomial-option*, the standard error is computed from the sample proportion as

$$\mathrm{se} = \sqrt{\hat{p}(1-\hat{p})/n}$$

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. For more information, see Fleiss, Levin, and Paik (2003). 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 0 (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 0, 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

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

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 $H_0\colon p = p_0$. To compute the exact test, PROC FREQ uses the binomial probability function,

$$\mathrm{Prob}(X = x \mid p_0) = \binom{n}{x} {p_0}^{x} \ (1 - p_0)^{\ (n-x)} \qquad \text{for} \ \ x = 0, 1, 2, \ldots, n$$

where the variable X has a binomial distribution with parameters n and p_0. To compute the left-sided p-value, $\mathrm{Prob}(X \le n_1)$, PROC FREQ sums the binomial probabilities over x from 0 to n_1. To compute the right-sided p-value, $\mathrm{Prob}(X \ge n_1)$, 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,

$$P_1 = \min\left(\mathrm{Prob}(X \le n_1 \mid p_0), \ \mathrm{Prob}(X \ge n_1 \mid p_0) \right)$$

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

$$H_0\colon p - p_0 \le -\delta$$

versus the alternative

$$H_a\colon p - p_0 > -\delta$$

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

$$z = (\hat{p} - p_0^*) \ / \ \mathrm{se}$$

where p_0^* is the noninferiority limit,

$$p_0^* = p_0 - \delta$$

By default, the standard error is computed from the sample proportion as

$$\mathrm{se} = \sqrt{\hat{p}(1-\hat{p})/n}$$

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

$$\mathrm{se} = \sqrt{p_0^*(1-p_0^*)/n}$$

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

$$P_z = \mathrm{Prob}(Z > z)$$

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 3168 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,

$$P_x = \sum_{k=n_1}^{k=n} \binom{n}{k} (p_0^*)^{\ k} \ (1-p_0^*)^{\ (n-k)}$$

For more information, see Chow, Shao, and Wang (2003, p. 116). If you request exact binomial statistics, PROC FREQ also includes exact (Clopper-Pearson) confidence limits for the binomial proportion in the equivalence analysis display. For more information, see the section "Exact (Clopper-Pearson) Confidence Limits" on page 3168.

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

$$H_0\colon p - p_0 \le \delta$$

versus the alternative

$$H_a\colon p - p_0 > \delta$$

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 3172 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

$$H_0\colon p - p_0 \le \delta_L \quad \text{or} \quad p - p_0 \ge \delta_U$$

versus the alternative

$$H_a\colon \delta_L < p - p_0 < \delta_U$$

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

$$\mathrm{se} = \sqrt{p_L^*(1-p_L^*)/n}$$

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

$$P_{z,L} = \mathrm{Prob}(Z > z_L)$$

The asymptotic test for the upper margin is computed similarly. The Wald test statistic is

$$z_U = (\hat{p} - p_U^*) \;/\; \mathrm{se}$$

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

$$\mathrm{se} = \sqrt{p_U^*(1-p_U^*)/n}$$

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

$$P_{z,U} = \mathrm{Prob}(Z < z_U)$$

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 3168, 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 3172. 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). For more information, see the section "Exact (Clopper-Pearson) Confidence Limits" on page 3168.

Sensitivity and Specificity

The SENSPEC option in the TABLES statement provides estimates of sensitivity, specificity, positive predictive value, and negative predictive value for 2×2 tables. These measures are conditional (row and column) proportions in the 2×2 table crosstabulation. In sensitivity analysis, the row variable might represent a positive or negative diagnostic test result, and the column variable might represent the presence or absence of a condition. For more information, see Fleiss, Levin, and Paik (2003).

Sensitivity is defined as the column proportion for table cell (1,1), which is the ratio of the frequency in table cell (1,1) to the total frequency in column 1 of the 2×2 table. The estimate of sensitivity is computed as

$$\mathrm{SN} = n_{11} \,/\, n_{\cdot 1}$$

Specificity is defined as the column proportion for table cell (2,2), which is the ratio of the frequency in table cell (2,2) to the total frequency in column 2. The estimate of specificity is computed as

$$\mathrm{SP} = n_{22} \,/\, n_{\cdot 2}$$

The positive predictive value is the row proportion for table cell (1,1), which is computed as

$$\mathrm{PPV} = n_{11} \,/\, n_{1 \cdot}$$

The negative predictive value is the row proportion for table cell (2,2), which is computed as

$$\mathrm{NPV} = n_{22} \,/\, n_{2 \cdot}$$

The “Sensitivity and Specificity” table provides the estimates together with their standard errors and Wald confidence limits. PROC FREQ computes the standard errors and Wald confidence limits for these proportions as described in the section “Risks and Risk Differences” on page 3176. The value of the confidence coefficient α is determined by the ALPHA= option; by default, ALPHA=0.05, which produces 95% confidence limits.

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 3168.

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,

$$\hat{p} = n_{\cdot 1} \; / \; n$$

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

$$\mathrm{se}(\hat{p}_i) = \sqrt{\hat{p}_i \; (1 - \hat{p}_i) \; / \; n_{i\cdot}}$$

The standard error of the overall column 1 risk is computed as

$$\mathrm{se}(\hat{p}) = \sqrt{\hat{p} \; (1 - \hat{p}) \; / \; n}$$

Where the two rows represent independent binomial samples, the standard error of the column 1 risk difference is computed as

$$\mathrm{se}(\hat{d}) = \sqrt{\hat{p}_1(1 - \hat{p}_1)/n_{1\cdot} \; + \; \hat{p}_2(1 - \hat{p}_2)/n_{2\cdot}}$$

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. For more information, see the sections "Confidence Limits for the Risk Difference" on page 3178 and "Risk Difference Tests" on page 3182.

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). For more information, see the sections "Binomial Confidence Limits" on page 3168 and "Binomial Tests" on page 3171.

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

$$\mathrm{Est} \; \pm \; (\, z_{\alpha/2} \times \mathrm{se}(\mathrm{Est}) \,)$$

where Est is the estimate, $z_{\alpha/2}$ is the $100(1-\alpha/2)$th percentile of the standard normal distribution, and se(Est) is the standard error of the estimate. The confidence level α is determined by the value of the ALPHA= option; by default, ALPHA=0.05, which 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

$$\mathrm{Est} \ \pm \ (\, z_{\alpha/2} \times \mathrm{se}(\mathrm{Est}) + c \,)$$

where c is the continuity correction. For the row 1 risk, $c = (1/2n_{1\cdot})$; for the row 2 risk, $c = (1/2n_{2\cdot})$; for the overall risk, $c = (1/2n)$; and for the risk difference, $c = ((1/n_{1\cdot} + 1/n_{2\cdot})/2)$. The column 1 and column 2 risks use the same continuity correction.

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. For more information, see the section "Exact (Clopper-Pearson) Confidence Limits" on page 3168.

Confidence Limits for the Risk Difference PROC FREQ provides the following confidence limit types for the risk difference: Agresti-Caffo, exact unconditional, Hauck-Anderson, Miettinen-Nurminen (score), Newcombe (hybrid-score), and Wald confidence limits. Continuity-corrected forms of 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. By default, ALPHA=0.05, which 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). For more information, see the section "Risk Difference Tests" on page 3182.

Agresti-Caffo Confidence Limits

Agresti-Caffo confidence limits for the risk difference are computed as

$$\tilde{d} \ \pm \ (\, z_{\alpha/2} \times \mathrm{se}(\tilde{d}) \,)$$

where $\tilde{d} = \tilde{p}_1 - \tilde{p}_2$, $\tilde{p}_i = (n_{i1} + 1)/(n_{i\cdot} + 2)$,

$$\mathrm{se}(\tilde{d}) = \sqrt{\tilde{p}_1(1-\tilde{p}_2)/(n_{1\cdot}+2) \ + \ \tilde{p}_2(1-\tilde{p}_2)/(n_{2\cdot}+2)}$$

and $z_{\alpha/2}$ is the $100(1-\alpha/2)$th 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

Hauck-Anderson confidence limits for the risk difference are computed as

$$\hat{d} \ \pm \ (\, c \ + \ z_{\alpha/2} \times \mathrm{se}(\hat{d}) \,)$$

where $\hat{d} = \hat{p}_1 - \hat{p}_2$ and $z_{\alpha/2}$ is the $100(1-\alpha/2)$th percentile of the standard normal distribution. The standard error is computed from the sample proportions as

$$\mathrm{se}(\hat{d}) = \sqrt{\hat{p}_1(1-\hat{p}_1)/(n_{1\cdot}-1) \ + \ \hat{p}_2(1-\hat{p}_2)/(n_{2\cdot}-1)}$$

The Hauck-Anderson continuity correction c is computed as

$$c = 1 \ / \ \big(2 \ \min(n_{1\cdot}, \ n_{2\cdot})\big)$$

For more information, see Hauck and Anderson (1986). The subsection "Hauck-Anderson Test" in the section "Noninferiority Tests" on page 3183 describes the corresponding noninferiority test.

Miettinen-Nurminen (Score) Confidence Limits

Miettinen-Nurminen (score) confidence limits for the risk difference (Miettinen and Nurminen 1985) are computed by inverting score tests for the risk difference. A score-based test statistic for the null hypothesis that the risk difference equals δ can be expressed as

$$T(\delta) = (\hat{d} - \delta) / \sqrt{\widetilde{\mathrm{Var}}(\delta)}$$

where $\hat{d}$ is the observed value of the risk difference $(\hat{p}_1 - \hat{p}_2)$,

$$\widetilde{\mathrm{Var}}(\delta) = (n/(n-1)) \ \big(\ \tilde{p}_1(\delta)(1-\tilde{p}_1(\delta))/n_1 + \tilde{p}_2(\delta)(1-\tilde{p}_2(\delta))/n_2 \ \big)$$

and $\tilde{p}_1(\delta)$ and $\tilde{p}_2(\delta)$ are the maximum likelihood estimates of the row 1 and row 2 risks (proportions) under the restriction that the risk difference is δ. For more information, see Miettinen and Nurminen (1985, pp. 215–216) and Miettinen (1985, chapter 12).

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,

$$\{\delta : T(\delta) < z_{\alpha/2}\}$$

where $z_{\alpha/2}$ is the $100(1-\alpha/2)$th percentile of the standard normal distribution. 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 $\widetilde{\mathrm{Var}}(\delta)$ (Miettinen and Nurminen 1985, p. 216). For more information, see Newcombe and Nurminen (2011). If you specify the CL=MN(CORRECT=NO) *riskdiff-option*, PROC FREQ does not include the bias correction factor in this computation (Mee 1984). See also Agresti (2002, p. 77). The uncorrected confidence limits are labeled as "Miettinen-Nurminen-Mee" confidence limits in the displayed output.

The maximum likelihood estimates of p_1 and p_2, subject to the constraint that the risk difference is δ, are computed as

$$\tilde{p}_1 = 2u\cos(w) - b/3a \quad \text{and} \quad \tilde{p}_2 = \tilde{p}_1 - \delta$$

where

$$
\begin{aligned}
w &= (\pi + \cos^{-1}(v/u^3))/3 \\
v &= b^3/(3a)^3 - bc/6a^2 + d/2a \\
u &= \mathrm{sign}(v)\sqrt{b^2/(3a)^2 - c/3a} \\
a &= 1 + \theta \\
b &= -(1 + \theta + \hat{p}_1 + \theta\hat{p}_2 + \delta(\theta + 2)) \\
c &= \delta^2 + \delta(2\hat{p}_1 + \theta + 1) + \hat{p}_1 + \theta\hat{p}_2 \\
d &= -\hat{p}_1\delta(1 + \delta) \\
\theta &= n_{2\cdot}/n_{1\cdot}
\end{aligned}
$$

For more information, see Farrington and Manning (1990, p. 1453).

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

$$|p_i - \hat{p}_i| = z_{\alpha/2}\sqrt{p_i(1 - p_i)/n_{i\cdot}}$$

for $i = 1, 2$. The confidence limits are computed as

$$\left(\hat{p}_i \; + \; z^2_{\alpha/2}/2n_{i\cdot} \; \pm \; z_{\alpha/2}\sqrt{\left(\hat{p}_i(1 - \hat{p}_i) + z^2_{\alpha}/4n_{i\cdot}\right)/n_{i\cdot}} \right) \; / \; \left(1 + z^2_{\alpha/2}/n_{i\cdot}\right)$$

For more information, see the section “Wilson (Score) Confidence Limits” on page 3171.

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

Wald confidence limits for the risk difference are computed as

$$\hat{d} \ \pm \ (\ z_{\alpha/2} \times \mathrm{se}(\hat{d}) \ )$$

where $\hat{d} = \hat{p}_1 - \hat{p}_2$, $z_{\alpha/2}$ is the $100(1-\alpha/2)$th percentile of the standard normal distribution. and the standard error is computed from the sample proportions as

$$\mathrm{se}(\hat{d}) = \sqrt{\hat{p}_1(1-\hat{p}_1)/n_{1\cdot} \ + \ \hat{p}_2(1-\hat{p}_2)/n_{2\cdot}}$$

If you specify the CORRECT *riskdiff-option*, the Wald confidence limits include a continuity correction c,

$$\hat{d} \ \pm \ (\ c \ + \ z_{\alpha/2} \times \mathrm{se}(\hat{d}) \ )$$

where $c = (1/n_{1\cdot} + 1/n_{2\cdot})/2$.

The subsection "Wald Test" in the section "Noninferiority Tests" on page 3183 describes the corresponding noninferiority test.

Exact Unconditional Confidence Limits

If you specify the RISKDIFF option in the EXACT statement, PROC FREQ provides exact unconditional confidence limits for the risk difference ($d = p_1 - p_2$). The exact unconditional approach fixes the row margins of the 2×2 table and eliminates the nuisance parameter p_2 by using the maximum p-value (worst-case scenario) over all possible values of p_2 (Santner and Snell 1980). The conditional approach, which is described in the section "Exact Statistics" on page 3222, does not apply to the risk difference because of the nuisance parameter (Agresti 1992).

By default, PROC FREQ computes the confidence limits by the tail method, which inverts two separate one-sided exact tests of the risk difference, where the tests are based on the score statistic (Chan and Zhang 1999). The size of each one-sided exact test is at most $\alpha/2$, and the confidence coefficient is at least $(1-\alpha)$. If you specify the RISKDIFF(METHOD=NOSCORE) option in the EXACT statement, PROC FREQ computes the confidence limits by inverting two separate one-sided exact tests that are based on the unstandardized risk difference. If you specify the RISKDIFF(METHOD=SCORE2) option in the EXACT statement, PROC FREQ computes the confidence limits by inverting a single two-sided exact test that is based on the score statistic (Agresti and Min 2001).

The score statistic is a less discrete statistic than the unstandardized risk difference and produces less conservative confidence limits (Agresti and Min 2001). For more information, see Santner et al. (2007). The section "Miettinen-Nurminen (Score) Confidence Limits" describe computation of the risk difference score statistic. For more information, see Miettinen and Nurminen (1985) and Farrington and Manning (1990).

PROC FREQ computes the exact unconditional 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

$$f(n_{11}, n_{21}; n_1, n_2, d, p_2) = \binom{n_1}{n_{11}} (d + p_2)^{n_{11}} (1 - d - p_2)^{n_1 - n_{11}} \times \binom{n_2}{n_{21}} p_2^{n_{21}} (1 - p_2)^{n_2 - n_{21}}$$

For the tail method (which inverts two separate one-sided exact tests), 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

$$\begin{aligned} P_U(d_*) &= \sup_{p_2}\Big(\sum_{A,T(a)\geq t_0} f(n_{11},n_{21};n_1,n_2,d_*,p_2) \Big) \\ P_L(d_*) &= \sup_{p_2}\Big(\sum_{A,T(a)\leq t_0} f(n_{11},n_{21};n_1,n_2,d_*,p_2) \Big) \end{aligned}$$

The set A includes all 2×2 tables in which the row sums are n_1 and n_2, $T(a)$ denotes the value of the test statistic for table a in A, and t_0 is the value of the test statistic for the observed table. The test statistic is either the score statistic (by default) or the unstandardized risk difference. To compute $P_U(d_*)$, the sum includes probabilities of those tables for which $(T(a) \geq t_0)$. For a fixed value of d_*, $P_U(d_*)$ is defined as the maximum sum over all possible values of p_2.

The two-sided score method evaluates the p-values $P_U(d_*)$ and $P_L(d_*)$ by comparing $|T(a)|$ to $|t_0|$. To compute the confidence limits d_L and d_u, the two-sided method compares the p-values to α. For more information, see Agresti and Min (2001) and Santner et al. (2007).

Risk Difference Tests

PROC FREQ provides tests of equality, noninferiority, superiority, and equivalence for the risk (proportion) difference. The following analysis methods are available: Wald (with and without continuity correction), Hauck-Anderson, Farrington-Manning (score), and Newcombe (with and without continuity correction). You can specify the method by using the METHOD= *riskdiff-option*; by default, PROC FREQ provides Wald tests.

Equality Tests The equality test for the risk difference tests the null hypothesis that the risk difference equals the null value. You can specify a null value by using the EQUAL(NULL=) *riskdiff-option*; by default, the null value is 0. This test can be expressed as $H_0\colon d = d_0$ versus the alternative $H_a\colon d \neq d_0$, where $d = p_1 - p_2$ denotes the risk difference (for column 1 or column 2) and d_0 denotes the null value.

The test statistic is computed as

$$z = (\hat{d} - d_0)/\mathrm{se}(\hat{d})$$

where the standard error $\mathrm{se}(\hat{d})$ is computed by using the method that you specify. Available methods for the equality test include Wald (with and without continuity correction), Hauck-Anderson, and Farrington-Manning (score). For a description of the standard error computation, see the subsections "Wald Test," "Hauck-Anderson Test," and "Farrington-Manning (Score) Test," respectively, in the section "Noninferiority Tests" on page 3183.

PROC FREQ computes one-sided and two-sided p-values for equality tests. 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

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

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

$$H_0\colon p_1 - p_2 \le -\delta$$

versus the alternative

$$H_a\colon p_1 - p_2 > -\delta$$

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, –δ.

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

$$z = (\hat{d} + \delta) \;/\; \mathrm{se}(\hat{d})$$

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

$$\mathrm{se}(\hat{d}) = \sqrt{\hat{p}_1(1-\hat{p}_1)/n_{1\cdot} \;+\; \hat{p}_2(1-\hat{p}_2)/n_{2\cdot}}$$

If you specify the VAR=NULL *riskdiff-option*, the standard error is based on the null hypothesis that the risk difference equals –δ (Dunnett and Gent 1977). The standard error is computed as

$$\mathrm{se}(\hat{d}) = \sqrt{\tilde{p}(1-\tilde{p})/n_{2\cdot} \;+\; (\tilde{p}-\delta)(1-\tilde{p}+\delta)/n_{1\cdot}}$$

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 0; 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 $P_z = \mathrm{Prob}(Z > z)$, 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

$$z = (\hat{d} + \delta \ \pm \ c) \ / \ \mathrm{se}(\hat{d})$$

where $\hat{d} = \hat{p}_1 - \hat{p}_2$ and the standard error is computed from the sample proportions as

$$\mathrm{se}(\hat{d}) = \sqrt{\hat{p}_1(1-\hat{p}_1)/(n_{1\cdot}-1) \ + \ \hat{p}_2(1-\hat{p}_2)/(n_{2\cdot}-1)}$$

The Hauck-Anderson continuity correction c is computed as

$$c = 1 \ / \ \big(2 \ \min(n_{1\cdot}, \ n_{2\cdot})\big)$$

The p-value for the Hauck-Anderson noninferiority test is $P_z = \mathrm{Prob}(Z > z)$, where Z has a standard normal distribution. See Hauck and Anderson (1986) and Schuirmann (1999) for more information.

Farrington-Manning (Score) Test

If you specify the METHOD=FM *riskdiff-option*, PROC FREQ provides the Farrington-Manning (score) test of noninferiority for the risk difference. A score test statistic for the null hypothesis that the risk difference equals –δ can be expressed as

$$z = (\hat{d} + \delta) \ / \ \mathrm{se}(\hat{d})$$

where $\hat{d}$ is the observed value of the risk difference $(\hat{p}_1 - \hat{p}_2)$,

$$\mathrm{se}(\hat{d}) = \sqrt{\tilde{p}_1(1-\tilde{p}_1)/n_{1\cdot} \ + \ \tilde{p}_2(1-\tilde{p}_2)/n_{2\cdot}}$$

and $\tilde{p}_1$ and $\tilde{p}_2$ are the maximum likelihood estimates of the row 1 and row 2 risks (proportions) under the restriction that the risk difference is –δ. The p-value for the noninferiority test is $P_z = \mathrm{Prob}(Z > z)$, where Z has a standard normal distribution. For more information, see Miettinen and Nurminen (1985); Miettinen (1985); Farrington and Manning (1990); Dann and Koch (2005).

The maximum likelihood estimates of p_1 and p_1, subject to the constraint that the risk difference is –δ, are computed as

$$\tilde{p}_1 = 2u\cos(w) - b/3a \quad \text{and} \quad \tilde{p}_2 = \tilde{p}_1 + \delta$$

where

$$\begin{aligned}
w &= (\pi + \cos^{-1}(v/u^3))/3 \\
v &= b^3/(3a)^3 - bc/6a^2 + d/2a \\
u &= \mathrm{sign}(v)\sqrt{b^2/(3a)^2 - c/3a} \\
a &= 1 + \theta \\
b &= -(1 + \theta + \hat{p}_1 + \theta\hat{p}_2 - \delta(\theta + 2)) \\
c &= \delta^2 - \delta(2\hat{p}_1 + \theta + 1) + \hat{p}_1 + \theta\hat{p}_2 \\
d &= \hat{p}_1\delta(1 - \delta) \\
\theta &= n_{2\cdot}/n_{1\cdot}
\end{aligned}$$

For more information, see Farrington and Manning (1990, p. 1453).

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, –δ. If you specify the CORRECT *riskdiff-option*, the confidence limits includes a continuity correction. See the subsection "Newcombe Confidence Limits" in the section "Confidence Limits for the Risk Difference" on page 3178 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

$$H_0{:}: p_1 - p_2 \le \delta$$

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 3183 by replacing –δ by δ. See Chow, Shao, and Wang (2003) for more information.

Equivalence Test 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

$$H_0{:}\ p_1 - p_2 \le -\delta_L \quad \text{or} \quad p_1 - p_2 \ge \delta_U$$

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 –δ 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 (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.

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 3183 gives details about the Wald, Hauck-Anderson, Farrington-Manning (score), 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. These 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 test-based confidence limits to the equivalence limits, (δ_L, δ_U).

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 3156 and “Exact Statistics” on page 3222 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 $p_{\cdot 1}$ 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

$$\mathrm{Prob}(\pi) = \sum_{A, T(a) \ge t_0} f(n_{11}, n_{21}; n_1, n_2, \pi)$$

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 $(T(a) \ge t_0)$, where t_0 is the value of the test statistic for the observed table.

The sum $\mathrm{Prob}(\pi)$ 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 $\mathrm{Prob}(\pi)$ over all possible values of π,

$$\mathrm{Prob} = \sup_{(0 \le \pi \le 1)} (\mathrm{Prob}(\pi))$$

See Suissa and Shuster (1985) and Mehta and Senchaudhuri (2003).

Common Risk Difference

If you specify the COMMONRISKDIFF option in the TABLES statement, PROC FREQ provides estimates, confidence limits, and tests for the common (overall) risk difference for multiway 2 × 2 tables.

Mantel-Haenszel Confidence Limits and Test

PROC FREQ computes the Mantel-Haenszel estimate, confidence limits, and test for the common risk difference by using Mantel-Haenszel stratum weights (Mantel and Haenszel 1959) and the Sato variance estimator (Sato 1989). The Mantel-Haenszel estimate of the common risk difference is

$$\widehat{d}_{\mathrm{MH}} = \sum_h \widehat{d}_h \; w_h$$

where $\widehat{d}_h$ is the risk difference in stratum h and

$$w_h = \frac{n_{h1\cdot} n_{h2\cdot}}{n_h} \; / \; \sum_i \frac{n_{i1\cdot} n_{i2\cdot}}{n_i}$$

is the Mantel-Haenszel weight of stratum h. The column 1 risk difference in stratum (2 × 2 table) h is computed as

$$\widehat{d}_h = \hat{p}_{h1} - \hat{p}_{h2} = (n_{h11}/n_{h1\cdot}) - (n_{h21}/n_{h2\cdot})$$

where $\hat{p}_{h1}$ is the proportion of row 1 observations that are classified in column 1 and $\hat{p}_{h2}$ is the proportion or row 2 observations that are classified in column 1. The column 2 risk is computed in the same way. For more information, see Agresti (2013, p. 231).

PROC FREQ computes the variance of $\widehat{d}_{\mathrm{MH}}$ (Sato 1989) as

$$\widehat{\sigma}^2(\widehat{d}_{\mathrm{MH}}) = \left(\widehat{d}_{\mathrm{MH}} \sum_h P_h + \sum_h Q_h \right) / \left(\sum_h n_{h1\cdot} n_{h2\cdot} / n_h \right)^2$$

where

$$P_h = \left(n_{h1\cdot}^2 n_{h21} - n_{h2\cdot}^2 n_{h11} + n_{h1\cdot} n_{h2\cdot} (n_{h2\cdot} - n_{h1\cdot})/2 \right) / n_h^2$$

$$Q_h = \left(n_{h11} (n_{h2\cdot} - n_{h21}) + n_{h21} (n_{h1\cdot} - n_{h11}) \right) / 2n_h$$

The $100(1-\alpha)\%$ confidence limits for the common risk difference are

$$\widehat{d}_{\mathrm{MH}} \; \pm \; \left(z_{\alpha/2} \times \widehat{\sigma}(\widehat{d}_{\mathrm{MH}}) \right)$$

If you specify the COMMONRISKDIFF(TEST=MH) option, PROC FREQ provides a Mantel-Haenszel test of the null hypothesis that the common risk difference is 0, which is computed as $z_{\mathrm{MH}} = \widehat{d}_{\mathrm{MH}} \; / \; \widehat{\sigma}(\widehat{d}_{\mathrm{MH}})$. The two-sided p-value is $\mathrm{Prob}(|Z| > |z_{\mathrm{MH}}|)$, where Z has a standard normal distribution.

Klingenberg Confidence Limits

Klingenberg confidence limits (Klingenberg 2014) for the Mantel-Haenszel common risk difference are based on inverting a test of homogeneity that uses the null form of the Sato variance estimator (Sato 1989). For performance evaluation of Klingenberg confidence limits, see Fisher (2015) and Klingenberg (2014).

The $100(1-\alpha)\%$ Klingenberg confidence limits for the common risk difference are

$$\widehat{d}_{\mathrm{Mid}} \ \pm \ M_{\alpha/2}$$

where M (margin of error) is computed as

$$M_{\alpha/2} = \sqrt{\widehat{d}^{\,2}_{\mathrm{Mid}} - \widehat{d}^{\,2}_{\mathrm{MH}} + z^2_{\alpha/2} \ (Q/W^2)}$$

and the confidence interval midpoint is computed as

$$\widehat{d}_{\mathrm{Mid}} = \widehat{d}_{\mathrm{MH}} + 0.5 \ z^2_{\alpha/2} \ (P/W^2)$$

The values P, Q, and W are computed as

$$\begin{aligned} P &= \sum_h P_h \\ Q &= \sum_h Q_h \\ W &= \sum_h n_{h1\cdot} \, n_{h2\cdot} \ / \ n_h \end{aligned}$$

where h denotes the stratum, and P_h and Q_h are defined in the section "Mantel-Haenszel Confidence Limits and Test" on page 3187.

Minimum Risk Confidence Limits and Test

PROC FREQ computes the minimum risk estimate, confidence limits, and test for the common risk difference by using the method of Mehrotra and Railkar (2000). The stratum estimates are weighted by minimum risk weights, which minimize the mean square error of the estimate of the common risk difference. Minimum risk weights are designed to improve precision and reduce bias (compared to other weighting strategies) and can minimize the power loss that can occur when underlying assumptions are not met. For more information, see Mehrotra (2001) and Dmitrienko et al. (2005, section 1.3.3).

The minimum risk estimate of the common risk difference is

$$\widehat{d}_{\mathrm{MR}} = \sum_h \widehat{d}_h \ w^*_h$$

where $\widehat{d}_h$ is the risk difference in stratum h and w^*_h is the minimum risk weight of stratum h (which is described in the section "Minimum Risk Weights" on page 3189). The variance of $\widehat{d}_{\mathrm{MR}}$ is estimated by

$$\widehat{V}(\widehat{d}_{\mathrm{MR}}) = \sum_h w^{*2}_h \ \widehat{V}_h$$

where $\widehat{V}_h$ (the variance estimate of the stratum h risk difference) is computed as

$$\widehat{V}_h = \hat{p}_{h1}(1-\hat{p}_{h1})/n_{h1\cdot} \ + \ \hat{p}_{h2}(1-\hat{p}_{h2})/n_{h2\cdot}$$

The $100(1-\alpha)\%$ minimum risk confidence limits for the common risk difference are

$$\widehat{d}_{\mathrm{MR}} \ \pm \ \left(c \ + \ z_{\alpha/2} \ \sqrt{\widehat{V}(\widehat{d}_{\mathrm{MR}})} \right)$$

where the continuity correction is

$$c = 0.1875 \ / \ \sum_h (n_{h1\cdot} n_{h2\cdot} / n_h)$$

The continuity correction is applied only when $c < |\widehat{d}_{\mathrm{MR}}|$ (Fleiss, Levin, and Paik 2003). You can remove the continuity correction by specifying the COMMONRISKDIFF(CORRECT=NO) option.

By default, the minimum risk test is computed as

$$z_{\mathrm{MR}} = \left(\widehat{d}_{\mathrm{MR}} \pm c \right) \ / \ \sqrt{\widehat{V}_0(\widehat{d}_{\mathrm{MR}})}$$

The continuity correction c is subtracted from $\widehat{d}_{\mathrm{MR}}$ if $\widehat{d}_{\mathrm{MR}} > 0$ and added to $\widehat{d}_{\mathrm{MR}}$ if $\widehat{d}_{\mathrm{MR}} < 0$. The null variance of the common risk difference is estimated by

$$\widehat{V}_0(\widehat{d}_{\mathrm{MR}}) = \sum_h w_h^{*2} \ \widehat{V}_0(\widehat{d}_h)$$

where $\widehat{V}_0(\widehat{d}_h)$ (an estimate of the variance of the stratum h risk difference under the null hypothesis) is

$$\widehat{V}_0(\widehat{d}_h) = \bar{p}_h (1 - \bar{p}_h) \ (1/n_{h1\cdot} + 1/n_{h2\cdot})$$

and

$$\bar{p}_h = (n_{h1\cdot} \hat{p}_{h1\cdot} + n_{h2\cdot} \hat{p}_{h2\cdot}) \ / \ (n_{h1\cdot} + n_{h2\cdot})$$

The two-sided p-value is $\mathrm{Prob}(|Z| > |z_{\mathrm{MR}}|)$, where Z has a standard normal distribution.

If you specify the VAR=SAMPLE option for COMMONRISKDIFF(TEST=MR), PROC FREQ uses the sample variance estimate $\widehat{V}(\widehat{d}_{\mathrm{MR}})$ instead of the null variance estimate $\widehat{V}_0(\widehat{d}_{\mathrm{MR}})$ in the denominator of the test statistic z_{MR}. If you specify the COMMONRISKDIFF(CORRECT=NO) option, the continuity correction is not included in the test statistic.

Minimum Risk Weights The estimate of the minimum risk weight for stratum h is defined by Mehrotra and Railkar (2000) as

$$w_h^* = \frac{\beta_h}{\sum_i \widehat{V}_i^{-1}} - \left(\frac{\alpha_h \widehat{V}_h^{-1}}{\sum_i \widehat{V}_i^{-1} + \sum_i \alpha_i \widehat{d}_i \widehat{V}_i^{-1}} \right) \left(\frac{\sum_i \widehat{d}_i \beta_i}{\sum_i \widehat{V}_i^{-1}} \right)$$

where

$$\alpha_h = \widehat{d}_h \sum_i \widehat{V}_i^{-1} - \sum_i \widehat{d}_i \widehat{V}_i^{-1}$$

$$\beta_h = \widehat{V}_h^{-1} \left(1 + \alpha_h \sum_i f_i \widehat{d}_i \right)$$

and f_h is the fraction in stratum h

$$f_h = n_h \ / \ \sum_i n_h$$

All sums are over the s strata (2×2 tables) in the multiway table request, $\widehat{d}_i$ denotes the risk difference estimate in stratum i, and $\widehat{V}_i$ denotes the sample variance estimate of the risk difference in stratum i.

Summary Score Confidence Limits

PROC FREQ computes the summary score estimate of the common risk difference (Agresti 2013, p. 231) by using inverse-variance stratum weights and Miettinen-Nurminen (score) confidence limits for the stratum risk differences. For more information, see the section "Miettinen-Nurminen (Score) Confidence Limits."

The score confidence interval for the risk difference in stratum h can be expressed as $\widehat{d}'_h \pm z_{\alpha/2}\, s'_h$, where $\widehat{d}'_h$ is the midpoint of the score confidence interval and s'_h is the width of the confidence interval divided by $2z_{\alpha/2}$. The summary score estimate of the common risk difference is computed as

$$\widehat{d}_S = \sum_h \widehat{d}'_h\, w'_h$$

where

$$w'_h = (1/{s'_h}^2) \;/\; \sum_i (1/{s'_i}^2)$$

The variance of $\widehat{d}_S$ is computed as

$$\widehat{\sigma}^2(\widehat{d}_S) = 1 \;/\; \sum_h (1/{s'_h}^2)$$

The $100(1-\alpha)\%$ summary score confidence limits for the common risk difference are

$$\widehat{d}_S \;\pm\; \left(z_{\alpha/2} \times \widehat{\sigma}(\widehat{d}_S) \right)$$

If you specify the COMMONRISKDIFF(TEST=SCORE) option, PROC FREQ provides a summary score test of the null hypothesis that the common risk difference is 0. The test statistic is $z_S = \widehat{d}_S / \widehat{\sigma}(\widehat{d}_S)$ The two-sided p-value is $\mathrm{Prob}(|Z| > |z_S|)$ where Z has a standard normal distribution.

Stratified Newcombe Confidence Limits

PROC FREQ computes stratified Newcombe confidence limits for the common risk (proportion) difference by using the method of Yan and Su (2010). The stratified Newcombe confidence limits are constructed from stratified Wilson confidence limits for the common (overall) row proportions. By default, the strata are weighted by Mantel-Haenszel weights; if you specify the COMMONRISKDIFF(CL=NEWCOMBEMR) option, the strata are weighted by minimum risk weights.

PROC FREQ first computes individual Wilson confidence limits for the row proportions in each 2 × 2 table (stratum), as described in the section "Wilson (Score) Confidence Limits" on page 3171. These stratum Wilson confidence limits are then combined to form stratified Wilson confidence limits for the overall row proportions by using stratum weights (either Mantel-Haenszel or minimum risk). The confidence levels of the stratum Wilson confidence limits are chosen so that the overall confidence coefficient (for the stratified Wilson confidence limits) is $100(1-\alpha)\%$ (Yan and Su 2010).

Denote the lower and upper stratified Wilson score confidence limits for the common row 1 proportion as L_1 and U_1, respectively, and denote the lower and upper stratified Wilson confidence limits for the common row 2 proportion as L_2 and U_2, respectively. The $100(1-\alpha)\%$ stratified Newcombe confidence limits for the common risk (proportion) difference are

$$\begin{aligned} L &= \widehat{d} \;-\; z_{\alpha/2}\sqrt{\lambda_1 L_1(1-L_1) + \lambda_2 U_2(1-U_2)} \\ U &= \widehat{d} \;+\; z_{\alpha/2}\sqrt{\lambda_2 L_2(1-L_2) + \lambda_1 U_1(1-U_1)} \end{aligned}$$

where $\widehat{d}$ is the weighted estimate of the common risk difference and

$$\begin{aligned}\lambda_1 &= \sum_h w_h^2 \; / \; n_{h1\cdot} \\ \lambda_2 &= \sum_h w_h^2 \; / \; n_{h2\cdot}\end{aligned}$$

By default, the strata are weighted by Mantel-Haenszel weights, which are defined as

$$w_h = \frac{n_{h1\cdot} n_{h2\cdot}}{n_h} \; / \; \sum_i \frac{n_{i1\cdot} n_{i2\cdot}}{n_i}$$

and the weighted estimate of the common risk difference is $\widehat{d}_{\mathrm{MH}}$. For more information, see the section "Mantel-Haenszel Confidence Limits and Test" on page 3187. Optionally, the strata are weighted by minimum risk weights, and the weighted estimate of the common risk difference is $\widehat{d}_{\mathrm{MR}}$. For more information, see the section "Minimum Risk Confidence Limits and Test" on page 3188.

When there is a single stratum, the stratified Newcombe confidence interval is equivalent to the (unstratified) Newcombe confidence interval. For more information, see the subsection "Newcombe Confidence Limits" in the section "Confidence Limits for the Risk Difference" on page 3178. See also Kim and Won (2013).

Odds Ratio and Relative Risks

Odds Ratio

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. For more information, see Stokes, Davis, and Koch (2012), Agresti (2013), 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

$$\mathrm{OR} = \frac{n_{11}/n_{12}}{n_{21}/n_{22}} = \frac{n_{11}\, n_{22}}{n_{12}\, n_{21}}$$

The odds ratio can be any nonnegative number. When the row and column variables are independent, the true value of the odds ratio is 1. An odds ratio greater than 1 indicates that the odds of a positive response are higher in row 1 than in row 2. An odds ratio less than 1 indicates that the odds of a positive response are higher in row 2. The strength of association increases as the deviation from 1 increases.

The transformation $G = (\mathrm{OR} - 1)/(\mathrm{OR} + 1)$ transforms the odds ratio to the range (–1,1), where $G = 0$ 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.

Confidence Limits for the Odds Ratio The following types of confidence limits are available for the odds ratio: exact, exact mid-*p*, likelihood ratio, score, Wald, and Wald modified.

Wald Confidence Limits

The asymptotic Wald confidence limits are based on a log transformation of the odds ratio (Woolf 1955; Haldane 1956). PROC FREQ computes the Wald confidence limits as

$$\left(\mathrm{OR} \times \exp(-z\sqrt{v}), \;\; \mathrm{OR} \times \exp(z\sqrt{v}) \right)$$

where

$$v = \mathrm{Var}(\log(\mathrm{OR})) = 1/n_{11} + 1/n_{12} + 1/n_{21} + 1/n_{22}$$

and z is the $100(1-\alpha/2)$th percentile of the standard normal distribution. The confidence level α is determined by the ALPHA= option in the TABLES statement; by default, ALPHA=0.05, which produces 95% confidence limits for the odds ratio. If any of the four cell frequencies are 0, v is undefined and the Wald confidence limits cannot be computed. For more information, see Agresti (2013, p. 70).

Wald Modified Confidence Limits

PROC FREQ computes Wald modified confidence limits (Haldane 1956) for the odds ratio by replacing the n_{ij} by $(n_{ij} + 0.5)$ in the estimate and variance as follows:

$$\mathrm{OR}_m = \frac{(n_{11} + 0.5)\,(n_{22} + 0.5)}{(n_{12} + 0.5)\,(n_{21} + 0.5)}$$

$$v = \mathrm{Var}(\log(\mathrm{OR}_m)) = 1/(n_{11} + 0.5) + 1/(n_{12} + 0.5) + 1/(n_{21} + 0.5) + 1/(n_{22} + 0.5)$$

The modified confidence limits are computed as

$$\left(\mathrm{OR}_m \times \exp(-z\sqrt{v}), \;\; \mathrm{OR}_m \times \exp(z\sqrt{v}) \right)$$

where z is the $100(1-\alpha/2)$th percentile of the standard normal distribution. For more information, see Fleiss, Levin, and Paik (2003) and Agresti (2013).

Score Confidence Limits

Score confidence limits for the odds ratio (Miettinen and Nurminen 1985) are computed by inverting score tests for the odds ratio. A score-based chi-square test statistic for the null hypothesis that the odds ratio is θ can be expressed as

$$Q(\theta) = \{n_{1\cdot}\,(\hat{p}_1 - \tilde{p}_1)\}^2 \;/\; \{n/(n-1)\}\;\{1/\,(n_{1\cdot}\tilde{p}_1(1-\tilde{p}_1)) + 1/\,(n_{2\cdot}\tilde{p}_2(1-\tilde{p}_2))\}^{-1}$$

where $\hat{p}_1$ is the observed row 1 risk $(n_{11}/n_{1\cdot})$, and $\tilde{p}_1$ and $\tilde{p}_2$ are the maximum likelihood estimates of the row 1 and row 2 risks under the restriction that the odds ratio $(n_{11}n_{22}/n_{12}n_{21})$ is θ. For more information, see Miettinen and Nurminen (1985) and Miettinen (1985, chapter 14).

The $100(1-\alpha)$% score confidence interval for the odds ratio consists of all values of θ for which the test statistic $Q(\theta)$ falls in the acceptance region,

$$\{\theta : Q(\theta) < \chi^2_{1,\alpha}\}$$

where $\chi^2_{1,\alpha}$ is the $100(1-\alpha)$th percentile of the chi-square distribution with 1 degree of freedom. For more information about score confidence limits, see Agresti (2013).

By default, the score confidence limits include the bias correction factor $n/(n-1)$ in the denominator of $Q(\theta)$ (Miettinen and Nurminen 1985, p. 217). If you specify the CL=SCORE(CORRECT=NO) option, PROC FREQ does not include this factor in the computation.

The maximum likelihood estimates of p_1 and p_2, subject to the constraint that the odds ratio is θ, are computed as

$$\tilde{p}_2 = \left(-b + \sqrt{b^2 - 4ac}\right)/2a \quad \text{and} \quad \tilde{p}_1 = \tilde{p}_2\theta / \left(1 + \tilde{p}_2(\theta - 1)\right)$$

where

$$\begin{aligned} a &= n_{2\cdot}(\theta - 1) \\ b &= n_{1\cdot}\theta + n_{2\cdot} - \hat{p}_{\cdot 1}(\theta - 1) \\ c &= -\hat{p}_{\cdot 1} \end{aligned}$$

For more information, see Miettinen and Nurminen (1985, pp. 217–218) and Miettinen (1985, chapter 14).

Likelihood Ratio Confidence Limits

Likelihood ratio (profile likelihood) confidence limits for the odds ratio are computed by inverting likelihood ratio tests. The likelihood ratio test statistic for the null hypothesis that the odds ratio is θ can be expressed as

$$G^2(\theta) \;=\; 2\left(n_{11}\log(\hat{p}_1/\tilde{p}_1) + n_{12}\log((1-\hat{p}_1)/(1-\tilde{p}_1)) + n_{21}\log(\hat{p}_2/\tilde{p}_2) + n_{22}\log((1-\hat{p}_2)/(1-\tilde{p}_2))\right)$$

where $\hat{p}_i$ is the observed row *i* risk ($n_{11}/n_{1\cdot}$) and $\tilde{p}_i$ is the maximum likelihood estimate of the row *i* risk under the restriction that the odds ratio is θ. The computation of the maximum likelihood estimates is described in the subsection "Score Confidence Limits" in this section. For more information, see Agresti (2013), Miettinen and Nurminen (1985), and Miettinen (1985, chapter 14).

The $100(1-\alpha)\%$ likelihood ratio confidence interval for the odds ratio consists of all values of θ for which the test statistic $G^2(\theta)$ falls in the acceptance region,

$$\{\theta : G^2(\theta) < \chi^2_{1,\alpha}\}$$

where $\chi^2_{1,\alpha}$ is the $100(1-\alpha)$th percentile of the chi-square distribution with 1 degree of freedom.

Exact Confidence Limits

PROC FREQ computes exact confidence limits for the odds ratio by inverting two one-sided (equal-tail) exact tests that are based on the noncentral hypergeometric distribution, where the distribution is conditional on the observed marginal totals of the 2 × 2 table. The exact confidence limits ϕ_1 and ϕ_2 are the solutions to the equations

$$\begin{aligned} \sum_{i=n_{11}}^{n_{\cdot 1}} f(i : n_{\cdot 1},\, n_{1\cdot},\, n_{2\cdot},\, \phi_1) &= \alpha/2 \\ \sum_{i=0}^{n_{11}} f(i : n_{\cdot 1},\, n_{1\cdot},\, n_{2\cdot},\, \phi_2) &= \alpha/2 \end{aligned}$$

where

$$f(i : n_{\cdot 1},\ n_{1\cdot},\ n_{2\cdot},\ \phi) = \binom{n_{1\cdot}}{i}\binom{n_{2\cdot}}{n_{\cdot 1}-i}\phi^i \ \Big/ \ \sum_{i=0}^{n_{\cdot 1}} \binom{n_{1\cdot}}{i}\binom{n_{2\cdot}}{n_{\cdot 1}-i}\phi^i$$

For more information, see Fleiss, Levin, and Paik (2003), Thomas (1971), and Gart (1971).

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. For more information, see Agresti (1992).

When the odds ratio is 0, which occurs when either $n_{11} = 0$ or $n_{22} = 0$, PROC FREQ sets the lower exact confidence limit to 0 and determines the upper limit by using the level α (instead of $\alpha/2$). Similarly, when the odds ratio is 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 by using level α.

Exact Mid-p Confidence Limits

PROC FREQ computes exact mid-*p* confidence limits for the odds ratio by inverting two one-sided hypergeometric tests that include mid-*p* tail areas. The mid-*p* approach replaces the probability of the observed table by half of that probability in the hypergeometric probability sums, which are described in the subsection "Exact Confidence Limits" in this section. The exact mid-*p* confidence limits ϕ_1 and ϕ_2 are the solutions to the equations

$$\begin{aligned}
\sum_{i=n_{11}+1}^{n_{\cdot 1}} \big(f(i : n_{\cdot 1},\ n_{1\cdot},\ n_{2\cdot},\ \phi_1)\big) \ + \ (1/2) f(n_{11} : n_{\cdot 1},\ n_{1\cdot},\ n_{2\cdot},\ \phi_1) &= \alpha/2 \\
\sum_{i=0}^{n_{11}-1} \big(f(i : n_{\cdot 1},\ n_{1\cdot},\ n_{2\cdot},\ \phi_2)\big) \ + \ (1/2) f(n_{11} : n_{\cdot 1},\ n_{1\cdot},\ n_{2\cdot},\ \phi_2) &= \alpha/2
\end{aligned}$$

where

$$f(i : n_{\cdot 1}, n_{1\cdot}, n_{2\cdot}, \phi) = \binom{n_{1\cdot}}{i}\binom{n_{2\cdot}}{n_{\cdot 1}-i}\phi^i \ \Big/ \ \sum_{i=0}^{n_{\cdot 1}} \binom{n_{1\cdot}}{i}\binom{n_{2\cdot}}{n_{\cdot 1}-i}\phi^i$$

For more information, see Agresti (2013).

When the odds ratio is 0, which occurs when either $n_{11} = 0$ or $n_{22} = 0$, PROC FREQ sets the lower exact confidence limit to 0 and determines the upper limit by using the level α (instead of $\alpha/2$). Similarly, when the odds ratio is 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 by using level α.

Relative Risks

Relative risks are useful measures 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 risks are also useful in cross-sectional studies, where two variables are observed simultaneously. For more information, see Stokes, Davis, and Koch (2012) and Agresti (2007).

The relative risk is the ratio of the row 1 risk to the row 2 risk in a 2 × 2 table. The column 1 risk in row 1 is the proportion of row 1 observations that are classified in column 1, which can be expressed as

$$p_1 = n_{11} \;/\; n_{1\cdot}$$

Similarly, the column 1 risk in row 2 is

$$p_2 = n_{21} \;/\; n_{2\cdot}$$

The column 1 relative risk is defined as

$$R = p_1 \;/\; p_2$$

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 as the deviation from 1 increases.

Confidence Limits for the Relative Risk PROC FREQ provides the following types of confidence limits for the relative risk: exact unconditional, likelihood ratio, score, Wald, and Wald modified.

Wald Confidence Limits

The asymptotic Wald confidence limits are based on a log transformation of the relative risk. PROC FREQ computes the Wald confidence limits for the column 1 relative risk as

$$\left(\hat{r} \times \exp(-z\sqrt{v}), \;\; \hat{r} \times \exp(z\sqrt{v}) \right)$$

where $\hat{r}$ is the observed value of the relative risk, $\hat{p}_1 / \hat{p}_2$, and

$$v = \mathrm{Var}(\log(\hat{r})) = 1/n_{11} + 1/n_{21} - 1/n_{1\cdot} - 1/n_{2\cdot}$$

and z is the $100(1-\alpha/2)$th percentile of the standard normal distribution. The confidence level α is determined by the ALPHA= option in the TABLES statement; by default, ALPHA=0.05, which produces 95% confidence limits. If either cell frequency n_{11} or n_{21} is 0, v is undefined and the Wald confidence limits cannot be computed.

PROC FREQ computes the confidence limits for the column 2 relative risk in the same way.

Wald Modified Confidence Limits

PROC FREQ computes Wald modified confidence limits (Haldane 1956) for the relative risk by replacing the n_{ij} with $(n_{ij} + 0.5)$ and the $n_{i\cdot}$ with $(n_{i\cdot} + 0.5)$ in the estimate and variance as follows:

$$\hat{r}_m = \frac{(n_{11} + 0.5)/(n_{1\cdot} + 0.5)}{(n_{21} + 0.5)/(n_{2\cdot} + 0.5)}$$

$$v = \mathrm{Var}(\log(\hat{r}_m)) = 1/(n_{11} + 0.5) + 1/(n_{21} + 0.5) - 1/(n_{1\cdot} + 0.5) - 1/(n_{2\cdot} + 0.5)$$

The confidence limits are computed as

$$\left(\hat{r}_m \times \exp(-z\sqrt{v}), \;\; \hat{r}_m \times \exp(z\sqrt{v}) \right)$$

where z is the $100(1-\alpha/2)$th percentile of the standard normal distribution. For more information, see Fleiss, Levin, and Paik (2003) and Agresti (2013).

Score Confidence Limits

Score confidence limits (Miettinen and Nurminen 1985; Farrington and Manning 1990) are computed by inverting score tests for the relative risk. A score-based chi-square test statistic for the null hypothesis that the relative risk is r_0 can be expressed as

$$Q(r_0) = (\hat{p}_1 - r_0 \hat{p}_2)^2 \ / \ \widetilde{\mathrm{Var}}(r_0)$$

where $\hat{p}_1$ and $\hat{p}_2$ are the observed row 1 and row 2 risks (proportions), respectively,

$$\widetilde{\mathrm{Var}}(r_0) = (n/(n-1)) \ \left(\ \tilde{p}_1(1-\tilde{p}_1)/n_{1\cdot} \ + \ {r_0}^2 \ \tilde{p}_2(1-\tilde{p}_2)/n_{2\cdot} \right)$$

where $\tilde{p}_1$ and $\tilde{p}_2$ are the maximum likelihood estimates of p_1 and p_2, respectively, under the null hypothesis that the relative risk is r_0. For more information, see Miettinen and Nurminen (1985) and Miettinen (1985, chapter 13).

The $100(1-\alpha)\%$ score confidence interval for the relative risk consists of all values of r_0 for which the test statistic $Q(r_0)$ falls in the acceptance region,

$$\{ r_0 : Q(r_0) < \chi^2_{1,\alpha} \}$$

where $\chi^2_{1,\alpha}$ is the $100(1-\alpha)$th percentile of the chi-square distribution with 1 degree of freedom. For more information, see Agresti (2013).

By default, the score confidence limits include the bias correction factor $n/(n-1)$ in the denominator of $Q(r_0)$ (Miettinen and Nurminen 1985, p. 217). If you specify the CL=SCORE(CORRECT=NO) option, PROC FREQ does not include this factor in the computation.

The maximum likelihood estimates of p_1 and p_2, subject to the constraint that the relative risk is r_0, are computed as

$$\tilde{p}_1 = \left(-b - \sqrt{b^2 - 4ac} \right) / 2a \quad \text{and} \quad \tilde{p}_2 = \tilde{p}_1 / r_0$$

where

$$\begin{aligned} a &= 1 + \theta \\ b &= -\left(r_0(1 + \theta \hat{p}_2) + \theta + \hat{p}_1 \right) \\ c &= r_0(\hat{p}_1 + \theta \hat{p}_2) \\ \theta &= n_{2\cdot}/n_{1\cdot} \end{aligned}$$

For more information, see Farrington and Manning (1990, p. 1454) and Miettinen and Nurminen (1985, p. 217).

Likelihood Ratio Confidence Limits

Likelihood ratio (profile likelihood) confidence limits for the relative risk are computed by inverting likelihood ratio tests. The likelihood ratio test statistic for the null hypothesis that the relative risk ratio is r_0 can be expressed as

$$G^2(r_0) \ = \ 2 \left(n_{11} \log(\hat{p}_1/\tilde{p}_1) + n_{12} \log((1-\hat{p}_1)/(1-\tilde{p}_1)) + n_{21} \log(\hat{p}_2/\tilde{p}_2) + n_{22} \log((1-\hat{p}_2)/(1-\tilde{p}_2)) \right)$$

where $\hat{p}_i$ is the observed row *i* risk ($n_{i1}/n_{i\cdot}$) and $\tilde{p}_i$ is the maximum likelihood estimate of the row *i* risk under the restriction that the relative risk is r_0. Expressions for the maximum likelihood estimates $\tilde{p}_1$ and $\tilde{p}_2$ are given in the subsection "Score Confidence Limits" in this section. For more information, see Miettinen and Nurminen (1985) and Miettinen (1985, chapter 13).

The $100(1-\alpha)\%$ likelihood ratio confidence interval for the relative risk consists of all values of r_0 for which the test statistic $G^2(r_0)$ falls in the acceptance region,

$$\{\theta : G^2(r_0) < \chi^2_{1,\alpha}\}$$

where $\chi^2_{1,\alpha}$ is the $100(1-\alpha)$th percentile of the chi-square distribution with 1 degree of freedom.

Exact Unconditional Confidence Limits

If you specify the RELRISK option in the EXACT statement, PROC FREQ provides exact unconditional confidence limits for the relative risk. The exact unconditional approach fixes the row margins of the 2 × 2 table and eliminates the nuisance parameter p_2 by using the maximum *p*-value (worst-case scenario) over all possible values of p_2 (Santner and Snell 1980). The conditional approach, which is described in the section "Exact Statistics" on page 3222, does not apply to the relative risk because of the nuisance parameter (Agresti 1992).

By default, PROC FREQ computes the confidence limits by the tail method, which inverts two separate one-sided exact tests of the relative risk, where the tests are based on the score statistic (Chan and Zhang 1999). The size of each one-sided exact test is at most $\alpha/2$, and the confidence coefficient is at least $(1-\alpha)$. If you specify the RELRISK(METHOD=NOSCORE) option in the EXACT statement, PROC FREQ computes the confidence limits by inverting two separate one-sided exact tests that are based on the unstandardized relative risk. If you specify the RELRISK(METHOD=SCORE2) option in the EXACT statement, PROC FREQ computes the confidence limits by inverting a single two-sided exact test that is based on the score statistic (Agresti and Min 2001).

PROC FREQ uses the relative risk score statistic (or the modified form of the unstandardized relative risk) to compute the exact confidence limits as described in the subsection "Exact Unconditional Confidence Limits" in the section "Confidence Limits for the Risk Difference" on page 3178.

The score statistic is a less discrete statistic than the unstandardized risk difference and produces less conservative confidence limits (Agresti and Min 2001). For more information, see Santner et al. (2007). The relative risk score statistic (Miettinen and Nurminen 1985; Farrington and Manning 1990) is computed as

$$z(r_0) = (\hat{p}_1 - r_0\hat{p}_2) \ / \ \mathrm{se}(r_0)$$

where

$$\mathrm{se}(r_0) = \sqrt{\tilde{p}_1(1-\tilde{p}_1)/n_{1\cdot} \ + \ r_0{}^2 \ \tilde{p}_2(1-\tilde{p}_2)/n_{2\cdot}}$$

where $\tilde{p}_1$ and $\tilde{p}_2$ are the maximum likelihood estimates of p_1 and p_2 under the restriction that the relative risk is r_0. Expressions for the maximum likelihood estimates $\tilde{p}_1$ and $\tilde{p}_2$ are given in the subsection "Score Confidence Limits" in this section. For more information, see Farrington and Manning (1990, p. 1454) and Miettinen and Nurminen (1985, p. 217).

When the confidence limits are computed by using the unstandardized relative risk as the test statistic (METHOD=NOSCORE), PROC FREQ uses a modified form of the relative risk to ensure that the statistic is defined when there are zero-frequency table cells. The modified form adds 0.5 to the table cell and row frequencies (Gart and Nam 1988) and is computed as

$$\hat{r}_m = \frac{(n_{11} + 0.5)/(n_{1\cdot} + 0.5)}{(n_{21} + 0.5)/(n_{2\cdot} + 0.5)}$$

For more information, see the subsection "Wald Modified Confidence Limits" in this section.

Relative Risk Tests PROC FREQ provides tests of equality, noninferiority, superiority, and equivalence for the relative risk. The following analysis methods are available: Wald (which is based on a log transformation), Wald modified, score, and likelihood ratio. You can specify the method by using the METHOD= *relrisk-option*; by default, PROC FREQ provides Wald tests.

Equality Test

An equality test for the relative risk can be expressed as

$$H_0\colon R = r_0$$

versus the alternative

$$H_a\colon R \neq r_0$$

where $R = p_1 / p_2$ denotes the relative risk (for column 1 or column 2) and r_0 denotes the null value. You can specify a null value by using the EQUAL(NULL=) *relrisk-option*; by default, the null value is 1.

The test statistic is computed by the method that you specify; by default, PROC FREQ uses the Wald test. For information about test statistic computation, see the subsections "Wald Test," "Wald Modified Test," "Farrington-Manning (Score) Test," and "Likelihood Ratio Test" in this section.

For the Wald and score methods, the test statistics z have standard normal distributions under the null hypothesis. For the likelihood ratio test, the test statistic G^2 has a chi-square distribution with 1 degree of freedom under the null hypothesis.

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

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

where Z has a standard normal distribution. The two-sided p-value is computed as $P_2 = 2 \times P_1$.

Noninferiority Test

A noninferiority test for the relative risk can be expressed as

$$H_0\colon R \leq \delta$$

versus the alternative

$$H_a\colon R > \delta$$

where $R = p_1 / p_2$ denotes the relative risk (for column 1 or column 2) and δ denotes the noninferiority margin (limit). You can specify the margin by using the MARGIN= *relrisk-option*; by default, the noninferiority margin is 0.8. The noninferiority margin for a relative risk test should be less than 1. Rejection of the null hypothesis indicates that the row 1 risk is not inferior to the row 2 risk. For more information, see Chow, Shao, and Wang (2008).

The test statistic z is computed by the method that you specify. For information about test statistic computation, see the subsections "Wald Test," "Wald Modified Test," "Farrington-Manning (Score) Test," and "Likelihood Ratio Test" in this section. The test statistic z is computed by using the noninferiority margin (limit) as the null value of the relative risk. Under the null hypothesis, the test statistic has a standard normal distribution. The p-value for the noninferiority test is the right-sided p-value (the probability that $Z > z$).

As part of the noninferiority analysis, PROC FREQ also provides confidence limits for the relative risk. The confidence coefficient is $100(1 - 2\alpha)\%$ (Schuirmann 1999). The confidence level α is determined by the ALPHA= option in the TABLES statement; by default, ALPHA=0.05, which produces 90% confidence limits for the noninferiority analysis. You can compare the confidence limits to the value of the noninferiority limit δ.

Superiority Test

A superiority test for the relative risk can be expressed as

$$H_0\colon R \le \delta$$

versus the alternative

$$H_a\colon R > \delta$$

where $R = p_1 / p_2$ denotes the relative risk (for column 1 or column 2) and δ denotes the superiority margin (limit). You can specify the margin by using the MARGIN= *relrisk-option*; by default, the superiority margin is 1.25. The superiority margin for a relative risk test should be greater than 1. Rejection of the null hypothesis indicates that the row 1 risk is superior to the row 2 risk. For more information, see Chow, Shao, and Wang (2008).

The test statistic z is computed by using the superiority margin (limit) as the null value of the relative risk. Under the null hypothesis, the test statistic has a standard normal distribution. The p-value for the superiority test is the right-sided p-value (the probability that $Z > z$).

The computations for the superiority analysis are the same as the computations for the noninferiority analysis, which are described in the subsection "Noninferiority Test" in this section.

Equivalence Test

An equivalence test for the relative risk can be expressed as

$$H_0\colon R \le \delta_L \quad \text{or} \quad R \ge \delta_U$$

versus the alternative

$$H_a\colon \delta_L < R < \delta_U$$

where δ_L is the lower margin and δ_U is the upper margin. Rejection of the null hypothesis indicates that the two risks are equivalent. For more information, see Chow, Shao, and Wang (2008).

You can specify the margins by using the MARGIN= *relrisk-option*; by default, the lower margin is 0.8 and the upper margin is 1.25. If you specify a single margin value, PROC FREQ uses this value as the lower margin for the equivalence test and computes the upper margin as the inverse of the lower margin.

PROC FREQ computes two one-sided tests (TOST) for equivalence analysis (Schuirmann 1987), which include a right-sided test for the lower margin δ_L and a left-sided test for the upper margin δ_U. The lower test statistic uses the lower margin as the null relative risk value, and the p-value is the right-sided probability $(Z > z_L)$. The upper test statistic uses the upper margin as the null value, and the p-value is the left-sided probability $(Z < z_U)$. The overall p-value is taken to be the larger of the two p-values for the lower and upper tests.

The test statistics are computed by the method that you specify. For more information about the test statistic computation, see the subsections "Wald Test," "Wald Modified Test," "Farrington-Manning (Score) Test," and "Likelihood Ratio Test" in this section.

As part of the equivalence analysis, PROC FREQ also provides confidence limits for the relative risk. The confidence coefficient is $100(1-2\alpha)\%$ (Schuirmann 1999). The confidence level α is determined by the ALPHA= option in the TABLES statement; by default, ALPHA=0.05, which produces 90% confidence limits for the equivalence analysis. You can compare the confidence limits to the equivalence limits, which are δ_L and δ_U.

Wald Test

The Wald test statistic (which is based on a log transformation of the relative risk) is computed as $z(r_0) = (\log(\hat{r}) - \log(r_0))/\sqrt{v}$, where $\hat{r}$ is the relative risk estimate $(\hat{p}_1/\hat{p}_2)$, r_0 is the null value of the relative risk, and

$$v = \mathrm{Var}(\log(\hat{r})) = 1/n_{11} + 1/n_{21} - 1/n_{1\cdot} - 1/n_{2\cdot}$$

The null value is determined by the type of test (equality, noninferiority, superiority, or equivalence) and the null or margin values that you specify. The side of the p-value and the interpretation of the test are also determined by the type of test; for more information, see the subsections "Equality Test," "Noninferiority Test," "Superiority Test," and "Equivalence Test" in this section.

Wald Modified Test

The Wald modified test statistic is computed by replacing the n_{ij} with $(n_{ij}+0.5)$ and the $n_{i\cdot}$ with $(n_{i\cdot}+0.5)$ in the relative risk estimate and variance. The test statistic is computed as $z(r_0) = (\log(\hat{r}_m) - \log(r_0))/\sqrt{v}$, where r_0 is the null value of the relative risk,

$$\hat{r}_m = \frac{(n_{11}+0.5)/(n_{1\cdot}+0.5)}{(n_{21}+0.5)/(n_{2\cdot}+0.5)}$$

$$v = \mathrm{Var}(\log(\hat{r}_m)) = 1/(n_{11}+0.5) + 1/(n_{21}+0.5) - 1/(n_{1\cdot}+0.5) - 1/(n_{2\cdot}+0.5)$$

The null value is determined by the type of test (equality, noninferiority, superiority, or equivalence) and the null or margin values that you specify. The side of the p-value and the interpretation of the test are also determined by the type of test; for more information, see the subsections "Equality Test," "Noninferiority Test," "Superiority Test," and "Equivalence Test" in this section.

Farrington-Manning (Score) Test

The relative risk score test statistic (Miettinen and Nurminen 1985; Farrington and Manning 1990) for the null value r_0 is computed as

$$z(r_0) = (\hat{p}_1 - r_0\hat{p}_2) \;/\; \mathrm{se}(r_0)$$

where

$$\mathrm{se}(r_0) = \sqrt{\tilde{p}_1(1-\tilde{p}_1)/n_{1\cdot} \ + \ {r_0}^2 \ \tilde{p}_2(1-\tilde{p}_2)/n_{2\cdot}}$$

where $\tilde{p}_1$ and $\tilde{p}_2$ are the maximum likelihood estimates of p_1 and p_2 under the null value r_0. Expressions for the maximum likelihood estimates $\tilde{p}_1$ and $\tilde{p}_2$ are given in the subsection "Score Confidence Limits" in this section.

The null value is determined by the type of test (equality, noninferiority, superiority, or equivalence) and the null or margin values that you specify. The side of the *p*-value and the interpretation of the test are also determined by the type of test; for more information, see the subsections "Equality Test," "Noninferiority Test," "Superiority Test," and "Equivalence Test" in this section.

Likelihood Ratio Test

The likelihood ratio statistic for the null relative risk value r_0 is computed as

$$G^2(r_0) \ = \ 2\left(n_{11}\log(\hat{p}_1/\tilde{p}_1) + n_{12}\log((1-\hat{p}_1)/(1-\tilde{p}_1)) + n_{21}\log(\hat{p}_2/\tilde{p}_2) + n_{22}\log((1-\hat{p}_2)/(1-\tilde{p}_2))\right)$$

where $\tilde{p}_1$ and $\tilde{p}_2$ are the maximum likelihood estimates of p_1 and p_2 under the null value r_0. Expressions for the maximum likelihood estimates $\tilde{p}_1$ and $\tilde{p}_2$ are given in the subsection "Score Confidence Limits" in this section. For more information, see Miettinen and Nurminen (1985) and Miettinen (1985, chapter 13).

PROC FREQ computes the likelihood ratio test statistic $z(r_0)$ for the noninferiority, superiority, and equivalence tests as $\sqrt{G^2(r_0)}$, where the sign is positive if the estimate is greater than the null value ($\hat{r} \geq r_0$) and negative otherwise ($\hat{r} < r_0$).

The null value is determined by the type of test (equality, noninferiority, superiority, or equivalence) and the null or margin values that you specify. The side of the *p*-value and the interpretation of the test are also determined by the type of test; for more information, see the subsections "Equality Test," "Noninferiority Test," "Superiority Test," and "Equivalence Test" in this section.

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. For more information, see Margolin (1988) and Agresti (2002). 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). By default, SCORES=TABLE. For more information, see the section “Scores” on page 3152. 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 $p_{i1} = n_{i1}/n_{i\cdot}$ 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 0, 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 0, 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

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

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 3222 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 $\tau_1 \le \tau_2 \le \cdots \le \tau_R$ (or $\tau_1 \ge \tau_2 \ge \cdots \ge \tau_R$), 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. By default, 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,

$$\begin{aligned} M_{i,i'} \quad = & \quad \{ \quad \text{number of times } X_{i,j} < X_{i',j'}, \quad j = 1,\ldots,n_{i\cdot};\;\; j' = 1,\ldots,n_{i'\cdot} \} \\ + \tfrac{1}{2} & \quad \{ \quad \text{number of times } X_{i,j} = X_{i',j'}, \quad j = 1,\ldots,n_{i\cdot};\;\; j' = 1,\ldots,n_{i'\cdot} \} \end{aligned}$$

where $X_{i,j}$ is response *j* in row *i*. The Jonckheere-Terpstra test statistic is computed as

$$J = \sum_{1 \le i <} \; \sum_{i' \le R} M_{i,i'}$$

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

$$J^* = (J - \mathrm{E}_0(J)) \; / \; \sqrt{\mathrm{Var}_0(J)}$$

where $\mathrm{E}_0(J)$ and $\mathrm{Var}_0(J)$ are the expected value and variance of the test statistic under the null hypothesis,

$$\mathrm{E}_0(J) = \left(n^2 - \sum_i n_{i\cdot}^2 \right) / 4$$

$$\mathrm{Var}_0(J) = A/72 + B/\left(36n(n-1)(n-2)\right) + C/\left(8n(n-1)\right)$$

where

$$A = n(n-1)(2n+5) - \sum_i n_{i\cdot}(n_{i\cdot}-1)(2n_{i\cdot}+5) - \sum_j n_{\cdot j}(n_{\cdot j}-1)(2n_{\cdot j}+5)$$

$$B = \left(\sum_i n_{i\cdot}(n_{i\cdot}-1)(n_{i\cdot}-2) \right) \left(\sum_j n_{\cdot j}(n_{\cdot j}-1)(n_{\cdot j}-2) \right)$$

$$C = \left(\sum_i n_{i\cdot}(n_{i\cdot}-1) \right) \left(\sum_j n_{\cdot j}(n_{\cdot j}-1) \right)$$

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 0, 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 0, 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

$$P_1 = \begin{cases} \mathrm{Prob}(Z > J^*) & \text{if} \quad J^* > 0 \\ \mathrm{Prob}(Z < J^*) & \text{if} \quad J^* \le 0 \end{cases}$$

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

$$P_2 = \mathrm{Prob}(|Z| > |J^*|)$$

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 3222 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 (for which the number of rows equals the number of columns). By default, these statistics include McNemar's test for 2×2 tables, Bowker's symmetry test, the simple kappa coefficient, and the weighted kappa coefficient. For multiple strata (n-way tables, where $n > 2$), the AGREE option provides the overall simple and weighted kappa coefficients, in addition to tests for equal kappas (simple and weighted) among strata. For multiple strata of 2×2 tables, the AGREE option provides Cochran's Q test.

Optionally, PROC FREQ provides kappa tests and other agreement statistics. 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. For more information, see the section "Exact Statistics" on page 3222.

The following sections provide the formulas that PROC FREQ uses to compute agreement 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 (McNemar 1947) for 2×2 tables when you specify the AGREE option. This test is appropriate when you are analyzing data from matched pairs of subjects with a dichotomous (yes-no) response. By default, the null hypothesis for McNemar's test is marginal homogeneity, which can be expressed as $p_{1\cdot} = p_{\cdot 1}$; this is equivalent to a discordant proportion ratio (p_{12}/p_{21}) of 1. The corresponding test statistic 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 1 degree of freedom.

Optionally, you can specify the null ratio of discordant proportions (p_{12}/p_{21}) by using the AGREE(MNULLRATIO=) option. When the null ratio is r, McNemar's test is computed as

$$Q_M(r) = (n_{12} - e_{12})^2)/e_{12} \;+\; (n_{21} - e_{21})^2/e_{21}$$

where $e_{12} = D/(1 + 1/r)$, $e_{21} = D/(1 + r)$, and D is the number of discordant pairs, $(n_{12} + n_{21})$. Under the null hypothesis, $Q_M(r)$ has an asymptotic chi-square distribution with 1 degree of freedom.

PROC FREQ also computes an exact p-value for McNemar's test when you specify the MCNEM option in the EXACT statement.

Bowker's Symmetry Test

The null hypothesis for Bowker's symmetry test (Bowker 1948) is symmetric table-cell proportions, which can be expressed as $p_{ij} = p_{ji}$ for all off-diagonal pairs of table cells. For 2 × 2 tables, Bowker's test is identical to McNemar's test; therefore, PROC FREQ provides Bowker's test only for square tables that are larger than 2 × 2.

Bowker's symmetry test 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, where R is the dimension of the square, two-way table.

By default, the number of degrees of freedom for this test ($R(R-1)/2$) is the number of off-diagonal table-cell comparisons. You can specify the number of degrees of freedom in the AGREE(DFSYM=) option. Alternatively, you can specify the AGREE(DFSYM=ADJUST) option, which reduces the degrees of freedom by the number of off-diagonal table-cell pairs that have a total frequency of 0. For more information, see Hoenig, Morgan, and Brown (1995).

Exact Symmetry Test When you specify the SYMMETRY option in the EXACT statement, PROC FREQ provides an exact symmetry test by using the method of Krauth (1973). This exact test is computed by conditioning on the observed frequency sums of the complementary off-diagonal table-cell pairs $(n_{ij} + n_{ji})$. PROC FREQ evaluates the symmetry test statistic for all tables in the reference set, which includes all possible tables in which the frequency sums of the off-diagonal table-cell pairs match the corresponding frequency sums in the observed table. The exact p-value is then computed as the sum of the table probabilities for those tables for which the symmetry test statistic is greater than or equal to the observed test statistic. The table probabilities are computed as products of $R(R-1)/2$ binomial probabilities (which correspond to the off-diagonal table-cell pairs in tables of dimension R) by using the binomial proportion 0.5 under the null hypothesis of symmetry. For more information, see the section "Exact Statistics" on page 3222.

Alternatively, you can request a Monte Carlo estimate of the exact p-value by specifying the SYMMETRY option together with the MC *computation-option* in the EXACT statement. The Monte Carlo computation for the exact symmetry test is conditional on the same reference set that the exact test uses (tables in which the frequency sums of the off-diagonal table-cell pairs match the corresponding sums in the observed table). For more information, see the section "Monte Carlo Estimation" on page 3225.

Simple Kappa Coefficient

The simple kappa coefficient (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 $P_e = \sum_i p_{i\cdot}p_{\cdot i}$. The component P_o is the proportion of observed agreement, and the component P_e represents the proportion of chance-expected agreement.

If the two response variables are viewed as two independent ratings of the n subjects, the kappa coefficient is +1 when there is complete agreement of the raters. When the observed agreement exceeds the chance-expected agreement, the kappa coefficient is positive, and its magnitude reflects the strength of agreement. When the observed agreement is less than the chance-expected agreement, the kappa coefficient is negative. The minimum value of kappa is between –1 and 0, depending on the marginal proportions of the table.

PROC FREQ computes the asymptotic variance of the simple kappa coefficient as

$$\mathrm{Var}(\hat{\kappa}) = (A + B - C) \ / \ (1 - P_e)^2 \ n$$

where

$$\begin{aligned} A &= \sum_i p_{ii} \, (1 - (p_{i\cdot} + p_{\cdot i})(1 - \hat{\kappa}))^2 \\ B &= (1 - \hat{\kappa})^2 \sum\sum_{i \neq j} \ p_{ij} (p_{\cdot i} + p_{j\cdot})^2 \\ C &= (\ \hat{\kappa} - P_e (1 - \hat{\kappa})\ )^2 \end{aligned}$$

For more information, see Fleiss, Cohen, and Everitt (1969).

Confidence limits for the simple kappa coefficient are computed as

$$\hat{\kappa} \ \pm \ \left(z_{\alpha/2} \times \sqrt{\mathrm{Var}(\hat{\kappa})} \ \right)$$

where $z_{\alpha/2}$ is the $100(1-\alpha/2)$th percentile of the standard normal distribution. The value of α is determined by the ALPHA= option; by default ALPHA=0.05, which produces 95% confidence limits.

PROC FREQ provides an asymptotic test for the simple kappa coefficient. By default, the null hypothesis value of kappa is 0; alternatively, you can specify a nonzero null value of kappa (by using the AGREE(NULLKAPPA=) option in the TABLES statement). When the null value of kappa is nonzero, PROC FREQ computes the test statistic as

$$z = (\hat{\kappa} - \kappa_0) \ / \ \sqrt{\mathrm{Var}(\hat{\kappa})}$$

where κ_0 is the null value that you specify and $\mathrm{Var}(\hat{\kappa})$ is the variance of the kappa coefficient.

When the null value of kappa is 0, PROC FREQ computes the test statistic as

$$z = \hat{\kappa} \ / \ \sqrt{\mathrm{Var}_0(\hat{\kappa})}$$

where $\mathrm{Var}_0(\hat{\kappa})$ is the variance of the kappa coefficient under the null hypothesis (that kappa is 0) and is computed as

$$\mathrm{Var}_0(\hat{\kappa}) = \left(P_e + P_e^2 - \sum_i p_{i\cdot} p_{\cdot i} (p_{i\cdot} + p_{\cdot i}) \right) \ / \ (1 - P_e)^2 \ n$$

This test statistic has an asymptotic standard normal distribution under the null hypothesis. For more information, see Fleiss, Levin, and Paik (2003).

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. For more information, see the section "Exact Statistics" on page 3222.

Kappa Details When you specify the AGREE(KAPPADETAILS) option, PROC FREQ displays the "Kappa Details" table, which includes the observed agreement P_o, chance-expected agreement P_e, maximum kappa, and B_n measure.

The maximum kappa, which is the maximum possible value of the kappa coefficient given the marginal proportions of the two-way table, is computed as

$$\max(\kappa) = (\max(P_o) - P_e) \ / \ (1 - P_e)$$

where

$$\max(P_o) = \left(\sum_i \min(n_{i\cdot}, n_{\cdot i}) \right) \ / \ n$$

The B_n measure (Bangdiwala 1988; Bangdiwala et al. 2008) is computed as

$$B_n = \left(\sum_i n_{ii}^2 \right) / \left(\sum_i \sum_j n_{i\cdot} n_{\cdot i} \right)$$

For 2 × 2 tables, the "Kappa Details" table also includes the prevalence index and the bias index. The prevalence index is the absolute difference between the agreement proportions, $|p_{11} - p_{22}|$. The bias index is the absolute difference between the disagreement proportions, $|p_{12} - p_{21}|$. For more information, see Sim and Wright (2005) and Byrt, Bishop, and Carlin (1993).

Weighted Kappa Coefficient

The weighted kappa coefficient is a generalization of the simple kappa coefficient that uses weights to quantify the relative differences between categories. For 2 × 2 tables, the weighted kappa coefficient is equivalent to the simple kappa coefficient; therefore, 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 section "Kappa Weights" on page 3209. The kappa 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

$$\hat{\kappa}_w = \left(P_{o(w)} - P_{e(w)} \right) \ / \ \left(1 - P_{e(w)} \right)$$

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 component $P_{o(w)}$ is the proportion of observed (weighted) agreement, and the component $P_{e(w)}$ represents the proportion of chance-expected (weighted) agreement. When you specify the AGREE(WTKAPDETAILS) option, PROC FREQ displays these components in the "Weighted Kappa Details" table.

PROC FREQ computes the asymptotic variance of the weighted kappa coefficient as

$$\mathrm{Var}(\hat{\kappa}_w) = \left(\sum_i \sum_j p_{ij} \left(w_{ij} - (\overline{w}_{i\cdot} + \overline{w}_{\cdot j})(1 - \hat{\kappa}_w) \right)^2 - \left(\hat{\kappa}_w - P_{e(w)}(1 - \hat{\kappa}_w) \right)^2 \right) \; / \; (1 - P_{e(w)})^2 \, n$$

where

$$\overline{w}_{i\cdot} = \sum_j p_{\cdot j} w_{ij}$$

$$\overline{w}_{\cdot j} = \sum_i p_{i\cdot} w_{ij}$$

For more information, see Fleiss, Cohen, and Everitt (1969).

Confidence limits for the weighted kappa coefficient are computed as

$$\hat{\kappa}_w \; \pm \; \left(z_{\alpha/2} \times \sqrt{\mathrm{Var}(\hat{\kappa}_w)} \right)$$

where $z_{\alpha/2}$ is the $100(1 - \alpha/2)$th percentile of the standard normal distribution. The value of α is determined by the ALPHA= option; by default ALPHA=0.05, which produces 95% confidence limits.

PROC FREQ provides an asymptotic test for the weighted kappa coefficient. By default, the null hypothesis value of weighted kappa is 0; alternatively, you can specify a nonzero null value of weighted kappa (by using the AGREE(NULLWTKAPPA=) option in the TABLES statement). When the null value of weighted kappa is nonzero, PROC FREQ computes the test statistic as

$$z = (\hat{\kappa}_w - \kappa_{w(0)}) \; / \; \sqrt{\mathrm{Var}(\hat{\kappa}_w)}$$

where $\kappa_{w(0)}$ is the null value that you specify and $\mathrm{Var}(\hat{\kappa}_w)$ is the variance of the weighted kappa coefficient.

When the null value of weighted kappa is 0, PROC FREQ computes the test statistic as

$$z = \hat{\kappa}_w \; / \; \sqrt{\mathrm{Var}_0(\hat{\kappa}_w)}$$

where $\mathrm{Var}_0(\hat{\kappa}_w)$ is the variance of the weighted kappa coefficient under the null hypothesis (that weighted kappa is 0) and is computed as

$$\mathrm{Var}_0(\hat{\kappa}_w) = \left(\sum_i \sum_j p_{i\cdot} p_{\cdot j} \left(w_{ij} - (\overline{w}_{i\cdot} + \overline{w}_{\cdot j}) \right)^2 - P_{e(w)}^2 \right) \; / \; (1 - P_{e(w)})^2 \, n$$

This test statistic has an asymptotic standard normal distribution under the null hypothesis. For more information, see Fleiss, Levin, and Paik (2003).

PROC FREQ also provides an exact test for the weighted kappa coefficient. You can request the exact test by specifying the KAPPA or AGREE option in the EXACT statement. For more information, see the section “Exact Statistics” on page 3222.

Kappa 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 the AGREE(WT=FC) option, 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. For more information, see Cicchetti and Allison (1971).

The SCORES= option in the TABLES statement determines the type of column scores used to compute the kappa weights (and other score-based statistics). By default, SCORES=TABLE. For more information, see the section "Scores" on page 3152. 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 the AGREE(WT=FC) option in the TABLES statement, PROC FREQ computes Fleiss-Cohen kappa coefficient weights as

$$w_{ij} = 1 - \frac{(C_i - C_j)^2}{(C_C - C_1)^2}$$

For more information, see Fleiss and Cohen (1973).

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.

Prevalence-Adjusted Bias-Adjusted Kappa

When you specify the AGREE(PABAK) option, PROC FREQ provides the prevalence-adjusted bias-adjusted kappa coefficient (PABAK) (Byrt, Bishop, and Carlin 1993). This coefficient is computed as

$$\hat{\kappa}_a = (P_o - 1/R) \ / \ (1 - 1/R)$$

where $P_o = \sum_i p_{ii}$ and R is the dimension of the square, two-way table. The component P_o is the proportion of observed agreement, and the component $1/R$ represents the chance-expected agreement. When the table is 2×2, $\hat{\kappa}_a = 2P_o - 1$. For more information, see Sim and Wright (2005), Xie (2013), and Holley and Guilford (1964).

PROC FREQ computes the variance of the prevalence-adjusted bias-adjusted kappa as

$$\text{Var}(\hat{\kappa}_a) = (R/(R-1))^2 \ (P_o(1 - P_o)/n)$$

Confidence limits are computed as

$$\hat{\kappa}_a \ \pm \ \left(z_{\alpha/2} \times \sqrt{\text{Var}(\hat{\kappa}_a)} \right)$$

where $z_{\alpha/2}$ is the $100(1-\alpha/2)$th percentile of the standard normal distribution. The value of α is determined by the ALPHA= option; by default ALPHA=0.05, which produces 95% confidence limits.

AC1 Agreement Coefficient

When you specify the AGREE(AC1) option, PROC FREQ provides Gwet's first-order agreement coefficient, AC1 (Gwet 2008). This coefficient is computed as

$$\hat{\gamma} = \left(P_o - P_{e(\gamma)}\right) \ / \ \left(1 - P_{e(\gamma)}\right)$$

where $P_o = \sum_i p_{ii}$, $P_e = \sum_i e_i(1-e_i)/(R-1)$, and $e_i = (p_{i\cdot} + p_{\cdot i})/2$ The component P_o is the proportion of observed agreement, and the component $P_{e(\gamma)}$ represents the proportion of chance-expected agreement. For more information, see Xie (2013) and Blood and Spratt (2007).

PROC FREQ computes the variance of AC1 as

$$\mathrm{Var}(\hat{\gamma}) = \left(P_o(1-P_o) - 4(1-\hat{\gamma})A + 4(1-\hat{\gamma}^2)B\right) \ / \ n(1-P_{e(\gamma)})^2$$

where

$$A = \sum_i p_{ii}(1-e_i)/(R-1) - P_o P_{e(\gamma)}$$

$$B = \sum_i \sum_j p_{ij} \left(1 - (e_i + e_j)/2\right)^2 / (R-1)^2 - P_{e(\gamma)}^2$$

Confidence limits for AC1 are computed as

$$\hat{\gamma} \ \pm \ \left(z_{\alpha/2} \times \sqrt{\mathrm{Var}(\hat{\gamma})} \right)$$

where $z_{\alpha/2}$ is the $100(1-\alpha/2)$th percentile of the standard normal distribution. The value of α is determined by the ALPHA= option; by default ALPHA=0.05, which produces 95% confidence limits.

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 $h = 1, 2, \ldots, q$, and let $\mathrm{Var}(\hat{\kappa}_h)$ denote the variance of $\hat{\kappa}_h$. The estimate of the overall kappa coefficient is computed as

$$\hat{\kappa}_T = \sum_{h=1}^{q} \frac{\hat{\kappa}_h}{\mathrm{Var}(\hat{\kappa}_h)} \ / \ \sum_{h=1}^{q} \frac{1}{\mathrm{Var}(\hat{\kappa}_h)}$$

For more information, see Fleiss, Levin, and Paik (2003).

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:

$$Q_K = \sum_{h=1}^{q} (\hat{\kappa}_h - \hat{\kappa}_T)^2 \ / \ \mathrm{Var}(\hat{\kappa}_h)$$

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

$$Q_C = m(m-1) \left(\sum_{j=1}^{m} T_j^2 - T^2 \right) \Bigg/ \left(mT - \sum_{k=1}^{N} S_k^2 \right)$$

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. For more information, see Cochran (1950). 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-Weight Rows or Columns

The AGREE statistics are defined only for square tables, where the number of rows equals the number of columns; if a table is not square, PROC FREQ does not compute AGREE statistics for the table. 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 contains r rows but only r–1 columns, the table is not square and PROC FREQ does not compute AGREE statistics. To create a square table in this situation, you can use the ZEROS option in the WEIGHT statement, which includes zero-weight observations in the analysis. You can 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. When you use this input data set and specify the ZEROS option, the analysis 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 in the WEIGHT statement. By default, PROC FREQ does not process observations that have weights of 0 because these observations do not contribute to the total frequency count, and because many of the tests and measures of association are undefined for tables that contain zero-weight rows or columns. However, kappa statistics are defined for tables that contain zero-weight rows or columns, and the ZEROS option enables you to input zero-weight observations and construct the tables needed to compute kappa statistics.

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, in addition to 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 $h = 1, 2, \ldots, q$. 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 $n_{hi.}$ and $n_{h.j}$, 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 ($'$).

$$
\begin{array}{lcll}
\mathbf{n}'_{hi} & = & (n_{hi1}, n_{hi2}, \ldots, n_{hiC}) & (1 \times C) \\
\mathbf{n}'_{h} & = & (\mathbf{n}'_{h1}, \mathbf{n}'_{h2}, \ldots, \mathbf{n}'_{hR}) & (1 \times RC) \\
p_{hi\cdot} & = & n_{hi\cdot} \;/\; n_h & (1 \times 1) \\
p_{h\cdot j} & = & n_{h\cdot j} \;/\; n_h & (1 \times 1) \\
\mathbf{P}'_{h*\cdot} & = & (p_{h1\cdot}, p_{h2\cdot}, \ldots, p_{hR\cdot}) & (1 \times R) \\
\mathbf{P}'_{h\cdot *} & = & (p_{h\cdot 1}, p_{h\cdot 2}, \ldots, p_{h\cdot C}) & (1 \times C)
\end{array}
$$

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,

$$\mathbf{m}_h = \mathbf{E}[\mathbf{n}_h \mid H_0] = n_h(\mathbf{P}_{h\cdot *} \otimes \mathbf{P}_{h*\cdot})$$

$$\mathbf{Var}[\mathbf{n}_h \mid H_0] = c\,\big(\,(\mathbf{D}_{\mathbf{P}h\cdot *} - \mathbf{P}_{h\cdot *}\mathbf{P}'_{h\cdot *}) \otimes (\mathbf{D}_{\mathbf{P}h*\cdot} - \mathbf{P}_{h*\cdot}\mathbf{P}'_{h*\cdot})\,\big)$$

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

$$Q_{\mathrm{CMH}} = \mathbf{G}'\mathbf{V_G}^{-1}\mathbf{G}$$

where

$$
\begin{array}{lcl}
\mathbf{G} & = & \sum_h \mathbf{B}_h(\mathbf{n}_h - \mathbf{m}_h) \\
\mathbf{V_G} & = & \sum_h \mathbf{B}_h\,(\mathbf{Var}[\mathbf{n}_h \mid H_0])\,\mathbf{B}'_h
\end{array}
$$

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 1 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, 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 3152 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, the ANOVA statistic is a Kruskal-Wallis test.

When there is more than one stratum, 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 47.9 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, 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, 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 $q(R\text{–}1)(C\text{–}1)$ degrees of freedom. The CMH statistic requires only a large overall sample size because it has only $(R\text{–}1)(C\text{–}1)$ 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(MANTELFLEISS) 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 more information, 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

$$\text{MF} = \min\left(\left[\sum_h m_{h11} - \sum_h (n_{h11})_L \right], \ \left[\sum_h (n_{h11})_U - \sum_h m_{h11} \right] \right)$$

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,

$$\begin{aligned} m_{h11} &= n_{h1\cdot}\, n_{h\cdot 1} \,/\, n_h \\ (n_{h11})_L &= \max(\, 0,\ \ n_{h1\cdot} - n_{h\cdot 2}\,) \\ (n_{h11})_U &= \min(\, n_{h\cdot 1},\ \ n_{h1\cdot}\,) \end{aligned}$$

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" on page 3191 for more information on these measures.

Throughout this section, z denotes the $100(1-\alpha/2)$th 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

$$\mathrm{OR}_{\mathrm{MH}} = \left(\sum_h n_{h11}\, n_{h22} / n_h \right) \; / \; \left(\sum_h n_{h12}\, n_{h21} / n_h \right)$$

It is always computed unless the denominator is 0. For more information, see Mantel and Haenszel (1959) and Agresti (2002).

To compute confidence limits for the common odds ratio, PROC FREQ uses the Robins, Breslow, and Greenland (1986) variance estimate for $\log(\mathrm{OR}_{\mathrm{MH}})$. The $100(1-\alpha/2)\%$ confidence limits for the common odds ratio are

$$\left(\, \mathrm{OR}_{\mathrm{MH}} \times \exp(-z\hat{\sigma}), \;\; \mathrm{OR}_{\mathrm{MH}} \times \exp(z\hat{\sigma}) \, \right)$$

where

$$\begin{aligned}
\hat{\sigma}^2 &= \widehat{\mathrm{Var}}(\, \log(\mathrm{OR}_{\mathrm{MH}})\,) \\
&= \frac{\sum_h (n_{h11} + n_{h22})(n_{h11}\, n_{h22}) / n_h^2}{2 \left(\sum_h n_{h11}\, n_{h22} / n_h \right)^2} \\
&\quad + \frac{\sum_h [(n_{h11} + n_{h22})(n_{h12}\, n_{h21}) + (n_{h12} + n_{h21})(n_{h11}\, n_{h22})] / n_h^2}{2 \left(\sum_h n_{h11}\, n_{h22} / n_h \right) \left(\sum_h n_{h12}\, n_{h21} / n_h \right)} \\
&\quad + \frac{\sum_h (n_{h12} + n_{h21})(n_{h12}\, n_{h21}) / n_h^2}{2 \left(\sum_h n_{h12}\, n_{h21} / n_h \right)^2}
\end{aligned}$$

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

$$\mathrm{OR_L} = \exp\left(\sum_h w_h \log(\mathrm{OR}_h) \;/\; \sum_h w_h \right)$$

and the corresponding $100(1-\alpha)\%$ confidence limits are

$$\left(\mathrm{OR_L} \times \exp\left(-z/\sqrt{\sum_h w_h} \right), \;\; \mathrm{OR_L} \times \exp\left(z/\sqrt{\sum_h w_h} \right) \right)$$

where OR_h is the odds ratio for stratum h, and

$$w_h = 1/\mathrm{Var}(\log(\mathrm{OR}_h))$$

If any table cell frequency in a stratum h is 0, PROC FREQ adds 0.5 to each cell frequency in the stratum before computing OR_h and w_h (Haldane 1956) for the logit estimate. The procedure provides 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

$$\mathrm{RR_{MH}} = \left(\sum_h n_{h11}\, n_{h2\cdot} \;/\; n_h \right) \;/\; \left(\sum_h n_{h21}\, n_{h1\cdot} \;/\; n_h \right)$$

It is always computed unless the denominator is 0. 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 $\log(\mathrm{RR_{MH}})$. The $100(1-\alpha/2)\%$ confidence limits for the common relative risk are

$$\left(\mathrm{RR_{MH}} \times \exp(-z\hat{\sigma}), \;\; \mathrm{RR_{MH}} \times \exp(z\hat{\sigma}) \right)$$

where

$$\hat{\sigma}^2 = \widehat{\mathrm{Var}}(\,\log(\mathrm{RR_{MH}})\,) = \frac{\sum_h (n_{h1\cdot}\, n_{h2\cdot}\, n_{h\cdot1} - n_{h11}\, n_{h21}\, n_h)/n_h^2}{\left(\sum_h n_{h11}\, n_{h2\cdot}/n_h\right)\left(\sum_h n_{h21}\, n_{h1\cdot}/n_h\right)}$$

Logit Estimator

The adjusted logit estimate of the common relative risk for column 1 is computed as

$$\mathrm{RR_L} = \exp\left(\sum_h w_h \log(\mathrm{RR}_h) \;/\; \sum w_h \right)$$

and the corresponding $100(1-\alpha)\%$ confidence limits are

$$\left(\mathrm{RR_L} \times \exp\left(-z \;/\; \sqrt{\sum_h w_h} \right),\;\; \mathrm{RR_L} \times \exp\left(z \;/\; \sqrt{\sum_h w_h} \right) \right)$$

where RR_h is the column 1 relative risk estimate for stratum h and

$$w_h = 1 \;/\; \mathrm{Var}(\log(\mathrm{RR}_h))$$

If n_{h11} or n_{h21} is 0, PROC FREQ adds 0.5 to each cell frequency in the stratum before computing RR_h and w_h for the logit estimate. The procedure prints a warning when this occurs. For more information, see Kleinbaum, Kupper, and Morgenstern (1982, Sections 17.4 and 17.5).

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

$$Q_{\mathrm{BD}} = \sum_h (n_{h11} - \mathrm{E}(n_{h11} \mid \mathrm{OR_{MH}}))^2 \;/\; \mathrm{Var}(n_{h11} \mid \mathrm{OR_{MH}})$$

where E and Var denote expected value and variance, respectively. The summation does not include any table that contains a row or column that has a total frequency of 0. If $\mathrm{OR_{MH}}$ is 0 or undefined, 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

$$Q_{\mathrm{BDT}} = Q_{\mathrm{BD}} - \left(\sum_h (n_{h11} - \mathrm{E}(n_{h11} \mid \mathrm{OR_{MH}})) \right)^2 \;/\; \sum_h \mathrm{Var}(n_{h11} \mid \mathrm{OR_{MH}})$$

See Tarone (1985); Jones et al. (1989); Breslow (1996) for more information.

Q Test for Homogeneity of Odds Ratios

PROC FREQ computes a Q test for homogeneity of odds ratios as

$$Q = \sum_h w_h (\theta_h - \bar{\theta})^2$$

where θ_h is the log odds ratio in stratum h and $\bar{\theta}$ is the logit estimate of the common log odds ratio. The stratum weights w_h are

$$w_h = 1/\mathrm{Var}(\theta_h)$$

where

$$\mathrm{Var}(\theta_h) = 1/n_{h11} + 1/n_{h12} + 1/n_{h21} + 1/n_{h22}$$

If any table cell frequency in a stratum is 0, PROC FREQ adds 0.5 to each cell frequency in the stratum before computing θ_h and w_h. For more information, see the sections "Odds Ratio" on page 3191 and "Adjusted Odds Ratio and Relative Risk Estimates" on page 3215.

Under the null hypothesis of homogeneity, the Q statistic has approximately a chi-square distribution with k–1 degrees of freedom, where k is the number of strata.

I-Square Measure of Heterogeneity

The I-square statistic (Higgins and Thompson 2002) is a measure of heterogeneity among strata for stratified 2×2 tables. I-square is expressed in percentage form and can be interpreted as the proportion of total variability that is due to between-strata variability. For more information, see Higgins et al. (2003) and Thorlund et al. (2012).

PROC FREQ computes I-square for the Q test for odds ratios as

$$I^2 = \max\left(\, 100\% \times (Q - (k-1))/Q,\ 0\,\right)$$

where k is the number of strata and Q is described in the section "Q Test for Homogeneity of Odds Ratios" on page 3217.

PROC FREQ computes uncertainty limits for I-square by using the test-based method of Higgins and Thompson (2002). This method constructs confidence limits for H, where $H^2 = Q/(k-1)$. When $Q > k$ or $k = 2$, the standard error of $\log(H)$ is computed as

$$\mathrm{SE}_1(\log(H)) = (\,\log(Q) - \log(k-1)\,)\ /\ 2\left(\sqrt{2Q} - \sqrt{2k-3}\,\right)$$

When $Q \le k$ and $k > 2$, the standard error of $\log(H)$ is computed as

$$\mathrm{SE}_0(\log(H)) = \sqrt{\left(\,1 - (1/3(k-2)^2)\,\right)\ /\ 2(k-2)}$$

The $100(1-\alpha)\%$ confidence limits for H are

$$\left(\,H \times \exp(-z_{\alpha/2} \times \mathrm{SE}(\log(H)),\ \ H \times \exp(z_{\alpha/2} \times \mathrm{SE}(\log(H))\,\right)$$

The uncertainty limits for I^2 are computed by transforming the confidence limits for H, where $I^2 = 1 - (1/H^2)$.

When I^2 is 0, PROC FREQ sets the lower confidence limit to 0 and determines the upper limit by using the level α (instead of $\alpha/2$).

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

$$P(\,S = s \mid n_{h1\cdot},\, n_{h\cdot 1},\, n_{h\cdot 2};\;\; h = 1, \ldots, q\,) = C_s\,\phi^{\,s} \;/\; \sum_{x\,=\,l}^{x\,=\,u} C_x\,\phi^{\,x}$$

where

$$C_s = \sum_{s_1 + \cdots + s_q \,=\, s} \left(\prod_h C_{s_h} \right)$$

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:

$$\sum_{x = s_0}^{x = u} C_x \, {\phi_1}^x \; / \; \sum_{x = l}^{x = u} C_x \, {\phi_1}^x \;=\; \alpha/2$$

$$\sum_{x = l}^{x = s_0} C_x \, {\phi_2}^x \; / \; \sum_{x = l}^{x = u} C_x \, {\phi_2}^x \;=\; \alpha/2$$

When the observed sum s_0 equals the lower bound l, PROC FREQ sets the lower confidence limit to 0 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

$$P_0(\, S = s \mid n_{h1\cdot}, \, n_{h\cdot1}, \, n_{h\cdot2}; \;\; h = 1, \ldots, q \,) = C_s \; / \; \sum_{x = l}^{x = u} \; C_x$$

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

$$\mathrm{E}_0(S) = \sum_{x = l}^{x = u} x \, C_x \; / \; \sum_{x = l}^{x = u} C_x$$

The one-sided exact p-value is computed from the conditional distribution as $P_0(S >= s_0)$ or $P_0(S \le s_0)$, depending on whether the observed sum s_0 is greater or less than $\mathrm{E}_0(S)$,

$$P_1 \;=\; P_0(\, S >= s_0 \,) = \sum_{x = s_0}^{x = u} C_x \; / \; \sum_{x = l}^{x = u} C_x \quad \text{if} \;\; s_0 > \mathrm{E}_0(S)$$

$$P_1 \;=\; P_0(\, S <= s_0 \,) = \sum_{x = l}^{x = s_0} C_x \; / \; \sum_{x = l}^{x = u} C_x \quad \text{if} \;\; s_0 \le \mathrm{E}_0(S)$$

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,

$${P_2}^a = 2 \times P_1$$

In addition, 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, $\;l \le s \le u$,

$${P_2}^b = \sum_{l \le s \le u : P_0(s) \le P_0(s_0)} P_0(s)$$

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,

$${P_2}^c = P_0 \; (\, |S - \mathrm{E}_0(S)| \;\ge\; |s_0 - \mathrm{E}_0(S)| \,)$$

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 more information, 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. For more information, see the section "Risks and Risk Differences" on page 3176. By default, PROC FREQ uses column 1 risks to compute the Gail-Simon test. If you specify the GAILSIMON(COLUMN=2) option, PROC FREQ uses column 2 risks.

PROC FREQ computes the Gail-Simon test statistics as described in Gail and Simon (1985),

$$\begin{aligned} Q- &= \sum_h (d_h/s_h)^2 \; I(d_h > 0) \\ Q+ &= \sum_h (d_h/s_h)^2 \; I(d_h < 0) \\ Q &= \min\,(Q-,\; Q+) \end{aligned}$$

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 $h = 1, 2, \ldots, q$.

The p-values for the Gail-Simon statistics are computed as

$$\begin{aligned} P(Q-) &= \sum_h (1 - F_h(Q-)) \; B(h; n = q, p = 0.5) \\ P(Q+) &= \sum_h (1 - F_h(Q+)) \; B(h; n = q, p = 0.5) \\ P(Q) &= \sum_{h=1}^{q-1} (1 - F_h(Q)) \; B(h; n = (q-1), p = 0.5) \end{aligned}$$

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 therefore the asymptotic p-values might not be 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. Asymptotic results might also be unreliable when the distribution of the data is sparse, skewed, or heavily tied. For more information, see Agresti (2007) and Bishop, Fienberg, and Holland (1975). Exact computations are based on the statistical theory of exact conditional inference for contingency tables, which is reviewed by Agresti (1992).

In addition to the computation of exact p-values, PROC FREQ provides the option to estimate exact p-values by Monte Carlo simulation. This can be useful for large problems where exact computations require a substantial amount of time and memory but asymptotic approximations might not be sufficient.

Exact p-values are available for many tests that PROC FREQ performs. For one-way tables, PROC FREQ provides exact p-values for the binomial proportion test, the chi-square goodness-of-fit test, and the likelihood ratio chi-square test. PROC FREQ also provides exact (Clopper-Pearson) confidence limits 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, Cochran-Armitage test for trend, and the symmetry test. PROC FREQ also provides exact p-values 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 (proportion) 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.

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, which was 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.

To implement the network algorithm, PROC FREQ defines a reference set from the input data. For most exact tests that PROC FREQ provides, the reference set includes all tables that have the same marginal row and column sums as the observed table. Corresponding to the 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 (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, no 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 further analysis. Otherwise, the algorithm continues to process this node and the associated paths. The algorithm finishes when all nodes have been accounted for.

PROC FREQ performs the network algorithm by using full numerical precision to represent all statistics, row and column scores, and other quantities in the computations. Although it is possible to use rounding to improve the speed and memory requirements of the algorithm, PROC FREQ does not use rounding because it might reduce the accuracy of the results.

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, PROC FREQ calculates the exact chi-square test based on that null hypothesis.

Other exact computations are described in sections about the individual statistics. For information about the computation of exact confidence limits and tests for the binomial proportion, see the section "Binomial Proportion" on page 3168. For information about computation of exact confidence limits for the odds ratio, see the subsection "Exact Confidence Limits" in the section "Confidence Limits for the Odds Ratio" on page 3192. For information about other exact computations, see the subsection "Exact Unconditional Confidence Limits" in the section "Confidence Limits for the Risk Difference" on page 3178, the subsection "Exact Unconditional Confidence Limits" in the section "Confidence Limits for the Relative Risk" on page 3195, and the sections "Exact Symmetry Test" on page 3205, "Exact Confidence Limits for the Common Odds Ratio" on page 3219 and "Zelen's Exact Test for Equal Odds Ratios" on page 3218.

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, the symmetry 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

$$P_1 = \begin{cases} \text{Prob(Test Statistic} \geq t \text{)} & \text{if} \quad t > \mathrm{E}_0(T) \\ \text{Prob(Test Statistic} \leq t \text{)} & \text{if} \quad t \leq \mathrm{E}_0(T) \end{cases}$$

where t is the observed value of the test statistic and $\mathrm{E}_0(T)$ 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

$$P_2 = \text{Prob (}|\text{Test Statistic} - \mathrm{E}_0(T)| \geq |t - \mathrm{E}_0(T)|)$$

If you specify the POINT option in the EXACT statement, PROC FREQ provides exact point probabilities for the exact tests. The exact point probability is the exact probability that the test statistic equals the observed value.

If you specify the MIDP option in the EXACT statement, PROC FREQ provides exact mid-p-values. The exact mid p-value is defined as the exact p-value minus half the exact point probability, which equals the average of $\text{Prob(Test Statistic} \geq t)$ and $\text{Prob(Test Statistic} > t)$ for a right-sided test. The exact mid p-value is smaller and less conservative than the non-adjusted exact p-value. For more information, see Agresti (2013, section 1.1.4) and Hirji (2006, sections 2.5 and 2.11.1).

Computational Resources

PROC FREQ uses relatively fast and efficient algorithms for exact computations. These algorithms, together with improvements in computing power, make it feasible to perform exact computations for data where previously only asymptotic methods could be applied. Nevertheless, depending on your available computing resources, exact computations for some very large problems might require a prohibitive amount of time and memory. For such large problems, consider whether exact methods are really needed or whether asymptotic methods might give results that are very close to the exact results while requiring much less computing time and memory. When asymptotic methods might not be sufficient for such large problems, consider using Monte Carlo estimation of exact p-values, which is described in the section “Monte Carlo Estimation” on page 3225.

There is no formula that can predict in advance how much time and memory are needed to compute an exact p-value for a specific data set and test. The time and memory requirements depend on several factors, which include the following: the total number of observations, the number of rows and columns in the table, the particular arrangement of the observations into table cells, and the test to be performed. Generally, larger problems (in terms of total sample size, number of rows, and number of columns) tend to require more time and memory. For a fixed total sample size, time and memory requirements tend to increase as the number of rows and number of columns increase because of the corresponding increase in the number of reference set tables. For a fixed sample size, time and memory requirements also tend to increase as the marginal row and column totals become more homogeneous. For more information, see Agresti, Mehta, and Patel (1990) and Gail and Mantel (1977).

While PROC FREQ is computing an exact p-value, you can terminate the computation by pressing the system interrupt key sequence and choosing to stop computations. For more information, see the *SAS Companion* for your system. After you terminate an exact computation, PROC FREQ completes all other remaining tasks. The procedure reports missing values for any exact p-values that were not computed before termination.

To limit the amount of time that PROC FREQ uses for exact computations, you can specify the MAXTIME= option in the EXACT statement. This option sets 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 an exact computation in the time that you specify, the procedure terminates the computation and completes the remaining tasks.

Monte Carlo Estimation

When you specify the MC option in the EXACT statement, PROC FREQ computes Monte Carlo estimates of exact p-values. Monte Carlo estimation can be useful for large problems where exact computations require a substantial amount of time and memory but asymptotic approximations might not be sufficient. Monte Carlo estimates are available for all exact tests that PROC FREQ provides except the binomial proportion test and those tests that apply only to 2×2 or $h \times 2 \times 2$ tables.

To describe the precision of a Monte Carlo estimate, PROC FREQ provides the asymptotic standard error and $100(1-\alpha)\%$ confidence limits. You can specify the confidence level α in the ALPHA= option in the EXACT statement; by default, ALPHA=0.01, which produces 99% confidence limits.

You can specify the number of Monte Carlo samples by using the N=n option in the EXACT statement. By default, PROC FREQ uses 10,000 samples to compute a Monte Carlo estimate. To improve the precision of the Monte Carlo estimates, you can specify a larger value of n; this increases the computation time because more samples are generated. To reduce the computation time, you can specify a smaller value of n.

PROC FREQ computes a Monte Carlo estimate of an exact p-value by generating a random sample of tables from the reference set for the exact test. For most exact tests that PROC FREQ provides, the reference set includes tables that have the same total sample size, row sums, and column sums as the observed table. (For the exact symmetry test, the reference set includes tables that have the same total sample size as the observed table and the same frequency sums of the off-diagonal table cell pairs.)

PROC FREQ generates a random sample of tables from the reference set by using 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 of the test statistic for the observed table. To estimate 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. The estimate of the p-value is the number of these tables divided by the total number of sample tables, which can be expressed as

$$\begin{aligned}
\hat{p}_{\mathrm{mc}} &= m \;/\; n \\
m &= \text{number of samples for which (test statistic} \geq t_o) \\
n &= \text{total number of samples} \\
t_o &= \text{observed test statistic}
\end{aligned}$$

PROC FREQ computes estimates of left-sided and two-sided exact p-values similarly. For left-sided exact p-values, PROC FREQ evaluates whether the sample test statistics are less than or equal to the observed test statistic. For two-sided exact p-values, PROC FREQ compares sample test statistics to the observed test statistic by using the definition of the two-sided p-value (P_2) for the test. For more information, see the section “Definition of p-Values” on page 3223 and descriptions of the individual tests.

The variable m has a binomial distribution with n trials and success probability p. The asymptotic standard error of the Monte Carlo estimate is

$$\mathrm{se}(\hat{p}_{\mathrm{mc}}) = \sqrt{\hat{p}_{\mathrm{mc}} \; (1 - \hat{p}_{\mathrm{mc}}) \; / \; (n-1)}$$

PROC FREQ constructs asymptotic confidence limits for the exact p-value as

$$\hat{p}_{\mathrm{mc}} \; \pm \; \left(z_{\alpha/2} \; \times \; \mathrm{se}(\hat{p}_{\mathrm{mc}}) \right)$$

where $z_{\alpha/2}$ is the $100(1-\alpha/2)$th 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 $\hat{p}_{\mathrm{mc}}$ is 0, PROC FREQ computes confidence limits for the p-value as

$$\left(\, 0, \;\; 1 - \alpha^{(1/n)} \, \right)$$

When the Monte Carlo estimate $\hat{p}_{\mathrm{mc}}$ is 1, PROC FREQ computes confidence limits for the p-value as

$$\left(\, \alpha^{(1/n)}, \;\; 1 \, \right)$$

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 cell frequencies of 0. 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. In addition, 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 3147.

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 3240 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 0 or missing, 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 47.7 lists the statistics that can be stored in the output data set. You identify which statistics to include by specifying *output-options*. For more information, see the description of the OUTPUT statement.

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 47.21 lists the prefixes and their descriptions.

Table 47.21 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_	*p*-value
P2_	Two-sided *p*-value
PL_	Left-sided *p*-value
PR_	Right-sided *p*-value
XP_	Exact *p*-value
XP2_	Exact two-sided *p*-value
XPL_	Exact left-sided *p*-value
XPR_	Exact right-sided *p*-value
XPT_	Exact point probability
XMP_	Exact mid *p*-value
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. For variables that were available in the output data set in releases before SAS/STAT 8.2, PROC FREQ truncates the variable name to eight characters when the length of the prefix plus the *output-option* name exceeds 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. For more information, see the section "Grouping with Formats" on page 3147. 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 MIDP option in the EXACT statement, PROC FREQ displays the exact mid *p*-value for the chi-square test.

- 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 specify the MIDP option in the EXACT statement, PROC FREQ displays the exact mid *p*-value for the binomial proportion test.

- If you request binomial confidence limits by specifying the BINOMIAL(CL=) option, PROC FREQ displays the “Binomial Confidence Limits” table, which includes the Lower and Upper Confidence Limits for each confidence limit Type that you request. In addition to Wald and Clopper-Pearson (Exact) confidence limits, you can request the following confidence limit types for the binomial proportion: Agresti-Coull, Blaker, Jeffreys, Likelihood Ratio, Logit, Mid-p, and Wilson (score).

- If you request a binomial noninferiority or superiority test by specifying the NONINF or SUP *binomial-option*, PROC FREQ displays a Noninferiority Analysis or Superiority Analysis table that contains 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 an Equivalence Analysis table that contains the following information: 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.

Two-Way and Multiway Tables

PROC FREQ displays all multiway 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 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
- Standardized Residual, if you specify the CROSSLIST(STDRES) option
- Pearson Residual, if you specify the CROSSLIST(PEARSONRES) 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 Two-Way and 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 POINT option in the EXACT statement, PROC FREQ displays the exact point probability for Fisher's exact test. If you specify the MIDP option in the EXACT statement, PROC FREQ displays the Mid p-Value for the test.
- 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. If you specify the POINT option in the EXACT statement, PROC FREQ displays the exact point probability for Fisher's exact test. If you specify the MIDP option in the EXACT statement, PROC FREQ displays the Mid p-Value for the test.

- 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 MIDP option in the EXACT statement, PROC FREQ displays the exact mid p-value for each test. 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 MC option in the EXACT statement, PROC FREQ displays Monte Carlo estimates for all exact p-values that you request 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 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 polychoric correlation and its asymptotic standard error (ASE). For 2×2 tables, this statistic is known as the tetrachoric correlation (and is labeled as such in the displayed output). If you specify the CL option, PROC FREQ also displays confidence limits for the polychoric correlation. If you specify the PLCORR option in the TEST statement, PROC FREQ displays the polychoric correlation, asymptotic standard error (ASE), confidence limits, and the following: the standardized test statistic (Z), the corresponding one-sided and two-sided probability values, the likelihood ratio (LR) chi-square, and the probability value (Pr > ChiSq).

- 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 H0, 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 H0, 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 MIDP option in the EXACT statement, PROC FREQ displays the exact Mid p-Value for each test.

- If you specify the SENSPEC option for 2 × 2 tables, PROC FREQ displays the "Sensitivity and Specificity" table. This table displays the Estimate, Standard Error, and Confidence Limits for the following statistics: Sensitivity, Specificity, Positive Predictive Value, and Negative Predictive Value.

- 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. You can suppress this table by specifying the RISKDIFF(NORISKS) option.

- If you specify the RISKDIFF(CL=) option for 2 × 2 tables, PROC FREQ displays the "Confidence Limits for the Proportion (Risk) Difference" table, which includes the Lower and Upper Confidence Limits for each confidence limit Type that you request (Agresti-Caffo, Exact, Hauck-Anderson, Miettinen-Nurminen, Newcombe, or Wald).

- If you specify the RISKDIFF(NONINF) option for 2×2 tables, PROC FREQ displays the "Noninferiority Analysis for the Risk Difference" table, which includes the Risk Difference, test ASE, standardized test statistic Z, probability value (Pr > Z), Noninferiority Limit, and (test-based) Confidence Limits.

- If you specify the RISKDIFF(SUP) option for 2 × 2 tables, PROC FREQ displays the "Superiority Analysis for the Risk Difference" table, which includes the Risk Difference, test ASE, standardized test statistic Z, probability value (Pr > Z), Superiority Limit, and (test-based) Confidence Limits.

- If you specify the RISKDIFF(EQUIV) option for 2 × 2 tables, PROC FREQ displays the "Equivalence Analysis for the Risk Difference" table, which includes the Risk Difference, test ASE, Equivalence Limits, and (test-based) Confidence Limits. PROC FREQ also displays the "Two One-Sided Tests (TOST)" table, which includes test statistics (Z) and P-Values for the Lower Margin and Upper Margin tests, along with the Overall P-Value.

- If you specify the RISKDIFF(EQUAL) option for 2 × 2 tables, PROC FREQ displays the "Risk Difference Test" table, which includes the Risk Difference, test ASE, standardized test statistic Z, One-sided probability value (Pr > Z or Pr < Z), and Two-sided probability value (Pr > |Z|).

- If you specify the MEASURES option or the RELRISK option for 2 × 2 tables, PROC FREQ displays the "Odds Ratio and Relative Risks" table, which includes the following statistics with their confidence limits: Odds Ratio, Relative Risk (Column 1), and Relative Risk (Column 2). If you specify the OR option in the EXACT statement, PROC FREQ also displays the "Exact Confidence Limits for the Odds Ratio" table. If you specify the RELRISK option in the EXACT statement, PROC FREQ displays the "Exact Confidence Limits for the Relative Risk" table.

- If you specify the OR(CL=) option for 2 × 2 tables, PROC FREQ displays the "Confidence Limits for the Odds Ratio" table, which includes the Lower and Upper Confidence Limits for each confidence limit Type that you request (Exact, Mid-p, Likelihood Ratio, Score, Wald, or Wald Modified).

- If you specify the RELRISK(CL=) option for 2 × 2 tables, PROC FREQ displays the "Confidence Limits for the Relative Risk" table, which includes the Lower and Upper Confidence Limits for each confidence limit Type that you request (Exact, Likelihood Ratio, Score, Wald, or Wald Modified).

- If you specify the RELRISK(NONINF) option, PROC FREQ displays the "Noninferiority Analysis for the Relative Risk" table, which includes the Relative Risk, standardized test statistic Z, probability value (Pr > Z), Noninferiority Limit, and Confidence Limits.

- If you specify the RELRISK(SUP) option, PROC FREQ displays the “Superiority Analysis for the Relative Risk” table, which includes the Relative Risk, standardized test statistic Z, probability value (Pr > Z), Superiority Limit, and Confidence Limits.

- If you specify the RELRISK(EQUIV) option, PROC FREQ displays the “Equivalence Analysis for the Relative Risk” table, which includes the Relative Risk, Equivalence Limits, and Confidence Limits. PROC FREQ also displays the “Two One-Sided Tests(TOST)” table, which includes test statistics (Z) and P-Values for the Lower Margin and Upper Margin tests, along with the Overall P-Value.

- If you specify the RELRISK(EQUAL) option, PROC FREQ displays the “Relative Risk Test” table, which includes the Relative Risk, standardized test statistic Z, One-sided probability value (Pr > Z or Pr < Z), and Two-sided probability value (Pr > |Z|).

- 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 MIDP option in the EXACT statement, PROC FREQ displays the exact Mid p-Value for the trend test.

- 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 MIDP option in the EXACT statement, PROC FREQ displays the exact Mid *p*-Value for the Jonckheere-Terpstra test.

- If you specify the AGREE option for a 2×2 table, PROC FREQ displays the “McNemar’s Test” table. This table includes the McNemar test statistic (chi-square), the degrees of freedom, and the *p*-value. If you specify the MCNEM option in the EXACT statement, this table also includes the exact *p*-value. If you specify the POINT option or the MIDP option in the EXACT statement, the “McNemar’s Test” table includes the exact point probability or the exact mid *p*-value, respectively.

- If you specify the AGREE option for a square table of dimension greater than 2, PROC FREQ displays the “Symmetry Test” table. This table displays Bowker’s symmetry test statistic (chi-square), the degrees of freedom, and the *p*-value. If you specify the SYMMETRY option in the EXACT statement, this table also includes the exact *p*-value. If you specify the POINT option or the MIDP option in the EXACT statement, the “Symmetry Test” table includes the exact point probability or the exact mid *p*-value, respectively.

- The AGREE option also produces the “Kappa Statistics” table, which displays the simple kappa coefficient. If the dimension of the two-way table is greater than 2, the “Kappa Statistics” table includes the weighted kappa coefficient. If you specify the AGREE(AC1) option or the AGREE(PABAK) option, this table includes the AC1 agreement coefficient or the prevalence-adjusted bias-adjusted kappa (PABAK), respectively. The “Kappa Statistics” table displays the standard error and confidence limits for each agreement statistic.

- If you specify the AGREE(KAPPADETAILS) option, PROC FREQ displays the “Kappa Details” table, which includes the observed agreement, the chance-expected agreement, the maximum kappa, and the

B_N measure. For 2 × 2 tables, the "Kappa Details" table also includes the prevalence index and the bias index.

- If you specify the AGREE(WTKAPPADETAILS) or AGREE(KAPPADETAILS) option for a square table of dimension greater than 2, PROC FREQ produces the "Weighted Kappa Details" table, which displays the observed agreement and the chance-expected agreement components of the weighted kappa coefficient.

- If you specify the AGREE(PRINTKWTS) option for a square table of dimension greater than 2, PROC FREQ displays the matrix of agreement weights in the "Kappa Coefficient Weights" table.

- If you request a simple kappa coefficient test, PROC FREQ produces the "Kappa Test" table. You can request this test by specifying the KAPPA option in the TEST statement, the KAPPA option in the EXACT statement, or the AGREE(NULLKAPPA=) option in the TABLES statement. The "Kappa Test" table displays the kappa coefficient, null test value, standard error (when the null value is 0), standardized test statistic (Z), and one-sided and two-sided *p*-values.

  If you request an exact test (by specifying the KAPPA option in the EXACT statement), the "Kappa Test" table also includes the exact one-sided and two-sided *p*-values. If you specify the POINT option or the MIDP option in the EXACT statement, the "Kappa Test" table includes the point probability or the exact mid *p*-value, respectively.

- If you request a weighted kappa coefficient test for a square table of dimension greater than 2, PROC FREQ produces the "Weighted Kappa Test" table. You can request this test by specifying the WTKAPPA option in the TEST statement, the WTKAPPA option in the EXACT statement, or the AGREE(NULLWTKAPPA=) option in the TABLES statement. The "Weighted Kappa Test" table displays the weighted kappa coefficient, null test value, standard error (when the null value is 0), standardized test statistic (Z), and one-sided and two-sided *p*-values.

  If you request an exact test (by specifying the WTKAPPA option in the EXACT statement), the "Weighted Kappa Test" table also includes the exact one-sided and two-sided *p*-values. If you specify the POINT option or the MIDP option in the EXACT statement, the "Weighted Kappa Test" table includes the point probability or the exact mid *p*-value, respectively.

- If you specify the AGREE option for a multiway square table, PROC FREQ displays the "Overall Kappa Coefficients" table, which includes the overall simple kappa coefficient together with its standard error and confidence limits. This table also includes the overall weighted kappa coefficient if the two-way table dimension is greater than 2.

- For multiway square tables, the AGREE option also produces the "Tests for Equal Kappa Coefficients" table. This table includes the chi-square statistic, degrees of freedom, and *p*-value for the test of equal simple kappa coefficients (over all strata). If the two-way table dimension is greater than 2, this table also includes the test for equal weighted kappa coefficients.

- For multiway 2 × 2 tables, the AGREE option displays the "Cochran's Q" table, which includes Cochran's Q statistic (to test for marginal homogeneity), the degrees of freedom, and the *p*-value.

- If you specify the COMMONRISKDIFF option for a multiway 2 × 2 table, PROC FREQ displays the "Confidence Limits for the Common Risk Difference" table, which includes the Method, Value of the common risk difference, Standard Error, and Confidence Limits for each confidence limit type that you request (Klingenberg, Mantel-Haenszel, Minimum Risk, Newcombe, Newcombe MR, or Summary Score).

- If you specify the COMMONRISKDIFF(TEST) option for a multiway 2 × 2 table, PROC FREQ displays the “Common Risk Difference Tests” table, which includes Method, Risk Difference, Z, and Pr > |Z| for each test that you request (Mantel-Haenszel, Minimum Risk, or Summary Score).

- If you specify the COMMONRISKDIFF(PRINTWTS) option for a multiway 2 × 2 table, PROC FREQ displays the “Stratum Weights” table, which includes the following information for each stratum (2 × 2 table): Stratum index, variable levels, Risk Difference, Frequency, Fraction, and the stratum weights that you request (Mantel-Haenszel, Minimum Risk, or Summary Score Weights).

- 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 (Pr > ChiSq).

- If you specify the CMH(QOR) option for a stratified 2 × 2 table, PROC FREQ displays the “Q Test for Homogeneity of Odds Ratios” table, which includes the Chi-Square, the degrees of freedom (DF), and the probability value (Pr > ChiSq).

- If you specify the CMH(I2) option for a stratified 2 × 2 table, PROC FREQ displays the “I-Square Measure of Heterogeneity” table, which includes the I-Square, the degrees of freedom (DF), and the Confidence Limits.

- If you specify the CMH option in the TABLES statement and also specify the COMOR option in the EXACT statement for a multiway 2×2 table, PROC FREQ displays exact confidence limits for the Common Odds Ratio. PROC FREQ also displays the Exact Test of H0: Common Odds Ratio = 1. The test output includes the Cell (1,1) Sum (S), Mean of S Under H0, One-sided Pr <= S, and Point Pr = S. PROC FREQ also provides exact two-sided probability values for the test, computed according to the following three methods: 2 * One-sided, Sum of probabilities <= Point probability, and Pr >= |S - Mean|. If you specify the MIDP option in the EXACT statement, PROC FREQ provides the exact Mid *p*-Value for the common odds ratio test.

- If you specify the CMH option in the TABLES statement and also specify the EQOR option in the EXACT statement for a multiway 2 × 2 table, PROC FREQ computes Zelen’s exact test for equal odds ratios. 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 Pr <= 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 23, "Using the Output Delivery System."

Table 47.22 lists the ODS table names together with their descriptions and the options that are required to produce the tables.

Table 47.22 ODS Tables Produced by PROC FREQ

ODS Table Name	Description	Statement	Option
BarnardsTest	Barnard's exact test	EXACT	BARNARD
Binomial	Binomial proportion	TABLES	BINOMIAL
BinomialCLs	Binomial confidence limits	TABLES	BINOMIAL(CL=)
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)
BinomialSup	Binomial superiority test	TABLES	BINOMIAL(SUP)
BinomialTest	Binomial proportion test	TABLES	BINOMIAL
BnMeasure	Agreement measures	TABLES	PLOTS=AGREEPLOT(STATS)
BreslowDayTest	Breslow-Day test	TABLES	CMH ($h \times 2 \times 2$ table)
ChiSq	Chi-square tests	TABLES	CHISQ
CMH	Cochran-Mantel-Haenszel statistics	TABLES	CMH
CochransQ	Cochran's Q	TABLES	AGREE ($h \times 2 \times 2$ table)
ColScores	Column scores	TABLES	SCOROUT
CommonOdds-RatioCl	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)
CommonPdiff	Common risk difference confidence limits	TABLES	COMMONRISKDIFF ($h \times 2 \times 2$ table)
CommonPdiffTests	Common risk difference tests	TABLES	COMMONRISKDIFF(TESTS) ($h \times 2 \times 2$ 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 ($h \times r \times r$ table, $r > 2$)
EqualOddsRatios	Tests for equal odds ratios	EXACT	EQOR ($h \times 2 \times 2$ table)

Table 47.22 *continued*

ODS Table Name	**Description**	**Statement**	**Option**
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
I2	I-square measure	TABLES	CMH(I2) ($h \times 2 \times 2$ table)
JTTest	Jonckheere-Terpstra test	TABLES	JT
JTTestMC	Monte Carlo estimates for Jonckheere-Terpstra exact test	EXACT	JT / MC
KappaDetails	Kappa details	TABLES	AGREE(KAPPADETAILS)
KappaMC	Monte Carlo exact test for simple kappa coefficient	EXACT	KAPPA / MC
KappaStatistics	Kappa statistics	TABLES	AGREE
KappaTest	Simple kappa test	TEST or EXACT or TABLES	KAPPA KAPPA AGREE(NULLKAPPA=)
KappaWeights	Kappa weights	TABLES	AGREE(PRINTKWTS)
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
MantelFleiss	Mantel-Fleiss criterion	TABLES	CMH(MANTELFLEISS) ($h \times 2 \times 2$ 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
OddsRatioCLs	Odds ratio confidence limits	TABLES	OR(CL=) (2×2 table)
OddsRatioExactCL	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)

Table 47.22 *continued*

ODS Table Name	Description	Statement	Option
OneWayFreqs	One-way frequencies	PROC or TABLES	(no TABLES stmt) (one-way table)
OneWayLRChiSq	One-way likelihood ratio chi-square test	TABLES	CHISQ(LRCHI) (one-way table)
OverallKappa	Overall simple kappa	TABLES	AGREE ($h \times 2 \times 2$ table)
OverallKappas	Overall kappa coefficients	TABLES	AGREE ($h \times r \times r$ table, $r > 2$)
PdiffCLs	Risk difference confidence limits	TABLES	RISKDIFF(CL=) (2×2 table)
PdiffEquiv	Equivalence analysis for the risk difference	TABLES	RISKDIFF(EQUIV) (2×2 table)
PdiffEquivTest	Equivalence test for the risk difference	TABLES	RISKDIFF(EQUIV) (2×2 table)
PdiffNoninf	Noninferiority test for the risk difference	TABLES	RISKDIFF(NONINF) (2×2 table)
PdiffSup	Superiority test for the risk difference	TABLES	RISKDIFF(SUP) (2×2 table)
PdiffTest	Risk 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
PearsonCorrTest	Pearson correlation test	TEST or EXACT	PCORR PCORR
PlCorr	Polychoric correlation	TEST	PLCORR
PlCorrTest	Polychoric correlation test	TEST	PLCORR
QOR	Q test for odds ratios	TABLES	CMH(QOR) ($h \times 2 \times 2$ table)
RelativeRiskCLs	Relative risk confidence limits	TABLES	RELRISK(CL=) (2×2 table)
RelativeRisks	Relative risk estimates	TABLES	RELRISK or MEASURES (2×2 table)
RelRisk1ExactCL	Exact confidence limits for column 1 relative risk	EXACT	RELRISK (2×2 table)
RelRisk2ExactCL	Exact confidence limits for column 2 relative risk	EXACT	RELRISK (2×2 table)
RelriskEquiv	Equivalence analysis for the relative risk	TABLES	RELRISK(EQUIV) (2×2 table)

Table 47.22 *continued*

ODS Table Name	Description	Statement	Option
RelriskEquivTest	Equivalence test for the relative risk	TABLES	RELRISK(EQUIV) (2×2 table)
RelriskNoninf	Noninferiority test for the relative risk	TABLES	RELRISK(NONINF) (2×2 table)
RelriskSup	Superiority test for the relative risk	TABLES	RELRISK(SUP) (2×2 table)
RelriskTest	Relative risk test	TABLES	RELRISK(EQUAL) (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
SenSpec	Sensitivity and specificity	TABLES	SENSPEC (2×2 table)
SomersDCR	Somers' $D(C\|R)$	TEST or EXACT	SMDCR SMDCR
SomersDCRMC	Monte Carlo exact test for Somers' $D(C\|R)$	EXACT	SMDCR / MC
SomersDCRTest	Somers' $D(C\|R)$ test	TEST or EXACT	SMDCR SMDCR
SomersDRC	Somers' $D(R\|C)$	TEST or EXACT	SMDRC SMDRC
SomersDRCMC	Monte Carlo exact test for Somers' $D(R\|C)$	EXACT	SMDRC / MC
SomersDRCTest	Somers' $D(R\|C)$ test	TEST or EXACT	SMDRC SMDRC
SpearmanCorr	Spearman correlation	TEST or EXACT	SCORR SCORR
SpearmanCorrMC	Monte Carlo exact test for Spearman correlation	EXACT	SCORR / MC
SpearmanCorrTest	Spearman correlation test	TEST or EXACT	SCORR SCORR
StratumWeights	Stratum weights and risk differences	TABLES	COMMONRISKDIFF ($h \times 2 \times 2$ table)
SymmetryMC	Monte Carlo exact symmetry test	EXACT	SYMMETRY / MC
SymmetryTest	Symmetry test	TABLES	AGREE
TauB	Kendall's tau-*b*	TEST or EXACT	KENTB KENTB
TauBMC	Monte Carlo exact test for Kendall's tau-*b*	EXACT	KENTB / MC
TauBTest	Kendall's tau-*b* test	TEST or EXACT	KENTB KENTB
TauC	Stuart's tau-*c*	TEST or EXACT	STUTC STUTC

Table 47.22 *continued*

ODS Table Name	Description	Statement	Option
TauCMC	Monte Carlo exact test for Stuart's tau-c	EXACT	STUTC / MC
TauCTest	Stuart's tau-c test	TEST or EXACT	STUTC STUTC
TrendTest	Cochran-Armitage trend test	TABLES	TREND
TrendTestMC	Monte Carlo exact test for trend	EXACT	TREND / MC
WtKappaDetails	Weighted kappa details	TABLES	AGREE(WTKAPPADETAILS)
WtKappaMC	Monte Carlo exact test for weighted kappa coefficient	EXACT	WTKAPPA / MC
WtKappaTest	Weighted kappa test	TEST or EXACT or TABLES	WTKAPPA WTKAPPA AGREE(NULLWTKAPPA=)

ODS Graphics

Statistical procedures use ODS Graphics to create graphs as part of their output. ODS Graphics is described in detail in Chapter 24, "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 663 in Chapter 24, "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 662 in Chapter 24, "Statistical Graphics Using ODS."

When ODS Graphics is enabled, you can request specific plots by specifying 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. For more information, see the description of the PLOTS= option.

PROC FREQ assigns a name to each graph that it creates by using ODS Graphics. You can use these names to refer to the graphs. Table 47.23 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 47.23 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, or RELRISK ($h \times 2 \times 2$ table)
RelRiskPlot	Relative risk plot	RELRISKPLOT	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 47.1: 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 cell frequencies of 0 in the output data set. The WEIGHT statement specifies that the variable Count contains the observation weights. These statements create Output 47.1.1 through Output 47.1.3.

```
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 47.1.1 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 47.1.2, with eye color defining the table rows and hair color defining the table columns. A cell frequency of 0 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 47.1.3) 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 that has a frequency of 0. The variable Expected contains the expected frequencies, as requested by the OUTEXPECT option.

Output 47.1.1 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 47.1.2 Crosstabulation Table

Frequency
Percent
Row Pct
Col Pct

Table of Eyes by Hair

Eyes(Eye Color)	**Hair(Hair Color)** 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 47.1.3 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 47.2: Frequency Dot Plots

This example produces frequency dot plots for the children's eye and hair color data from Example 47.1.

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 47.2.1, Output 47.2.2, and Output 47.2.3 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 47.2.2 and Output 47.2.3) is GROUPVERTICAL. Two-way layouts STACKED and GROUPHORIZONTAL are also available.

Output 47.2.1 One-Way Frequency Dot Plot

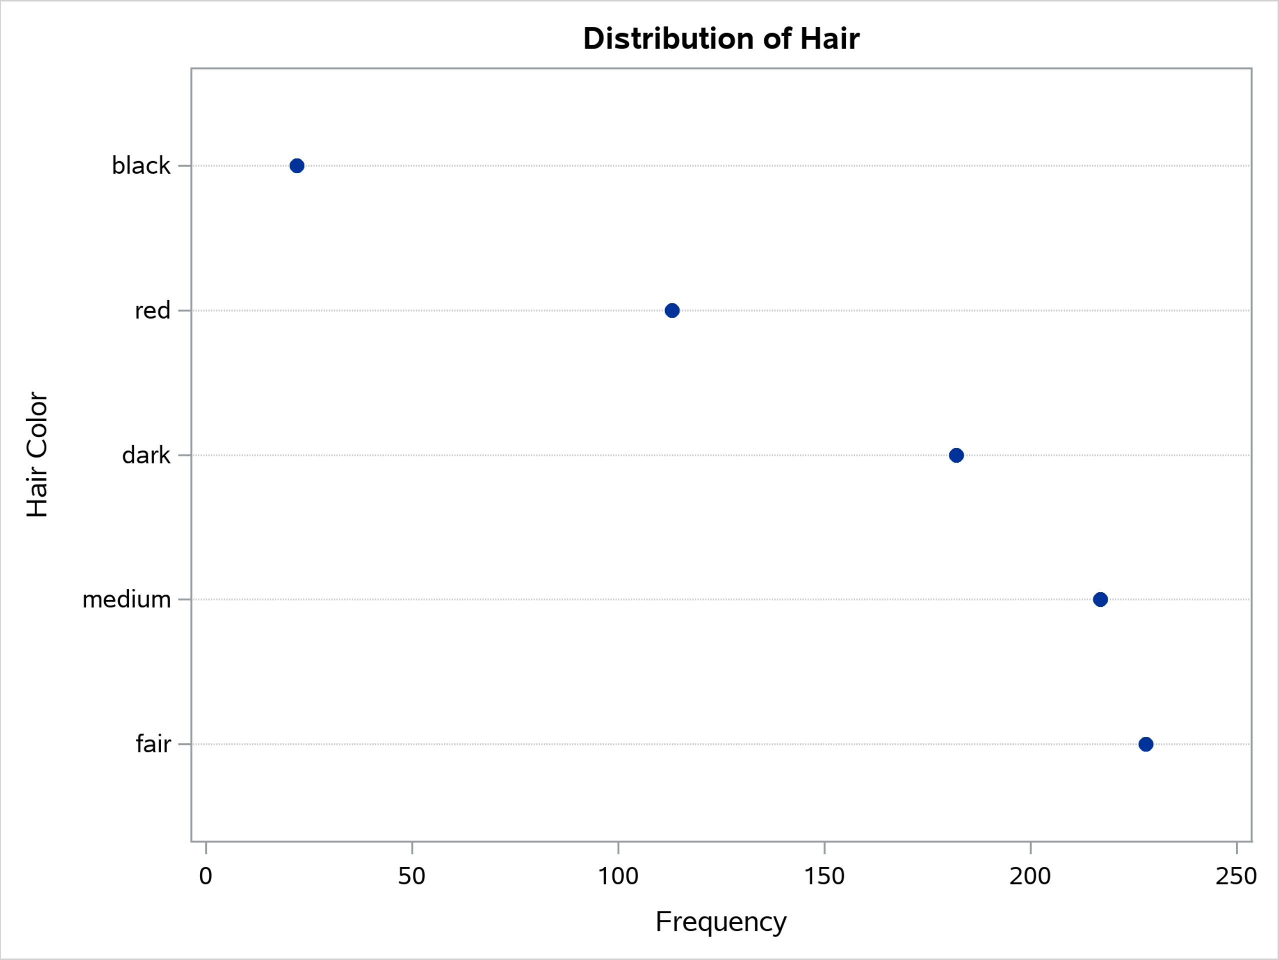

Output 47.2.2 Two-Way Frequency Dot Plot

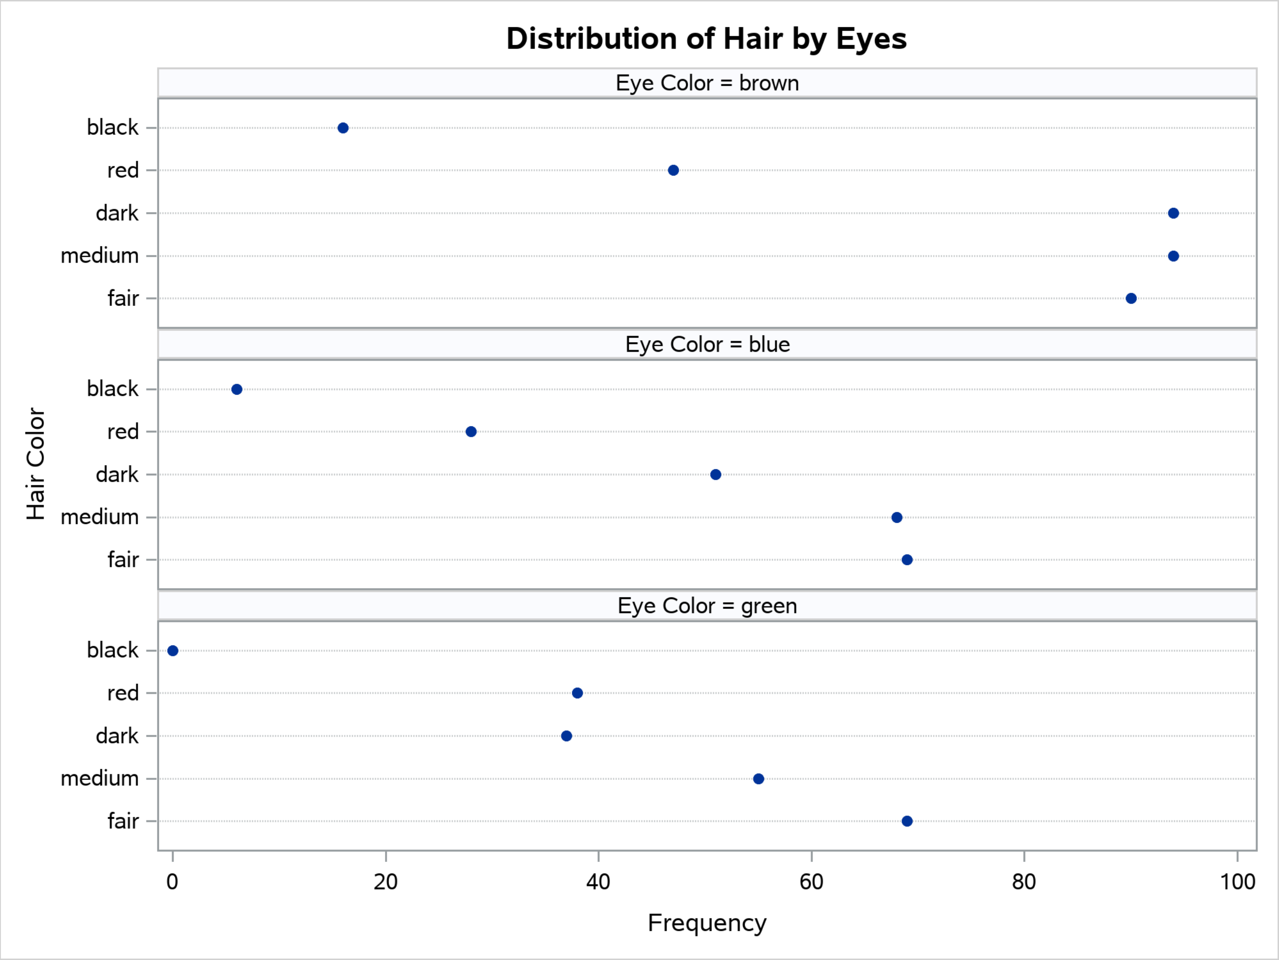

Output 47.2.3 Two-Way Percent Dot Plot

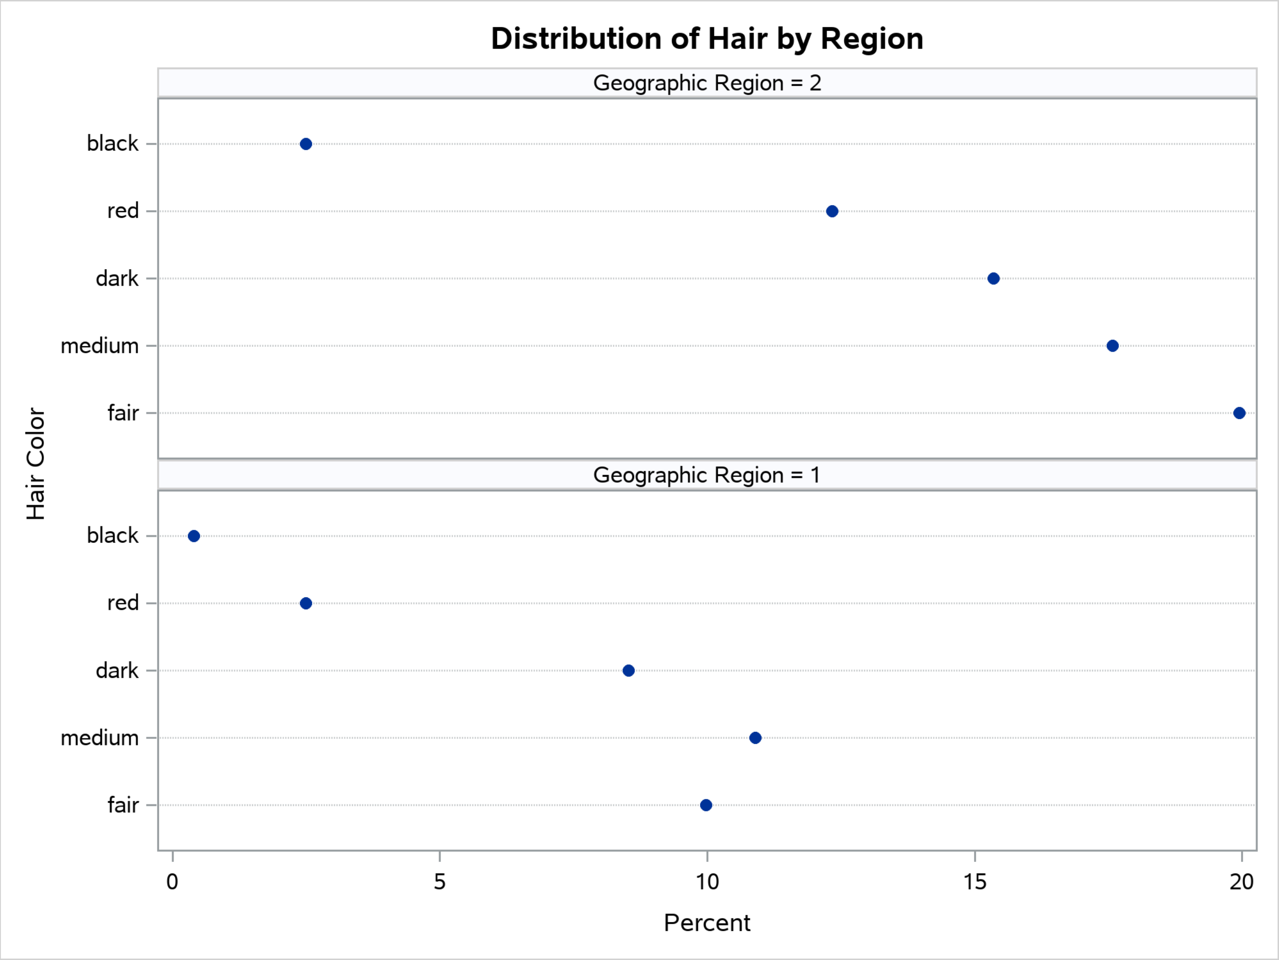

Example 47.3: Chi-Square Goodness-of-Fit Tests

This example examines whether the children's hair color (from Example 47.1) 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 47.3.1 through Output 47.3.4.

```
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 47.3.1 Frequency Table and Chi-Square Test for Region 1

Hair Color of European Children

The FREQ Procedure

Geographic Region=1

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 47.3.1 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 47.3.2 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 47.3.2 Deviation Plot for Region 1

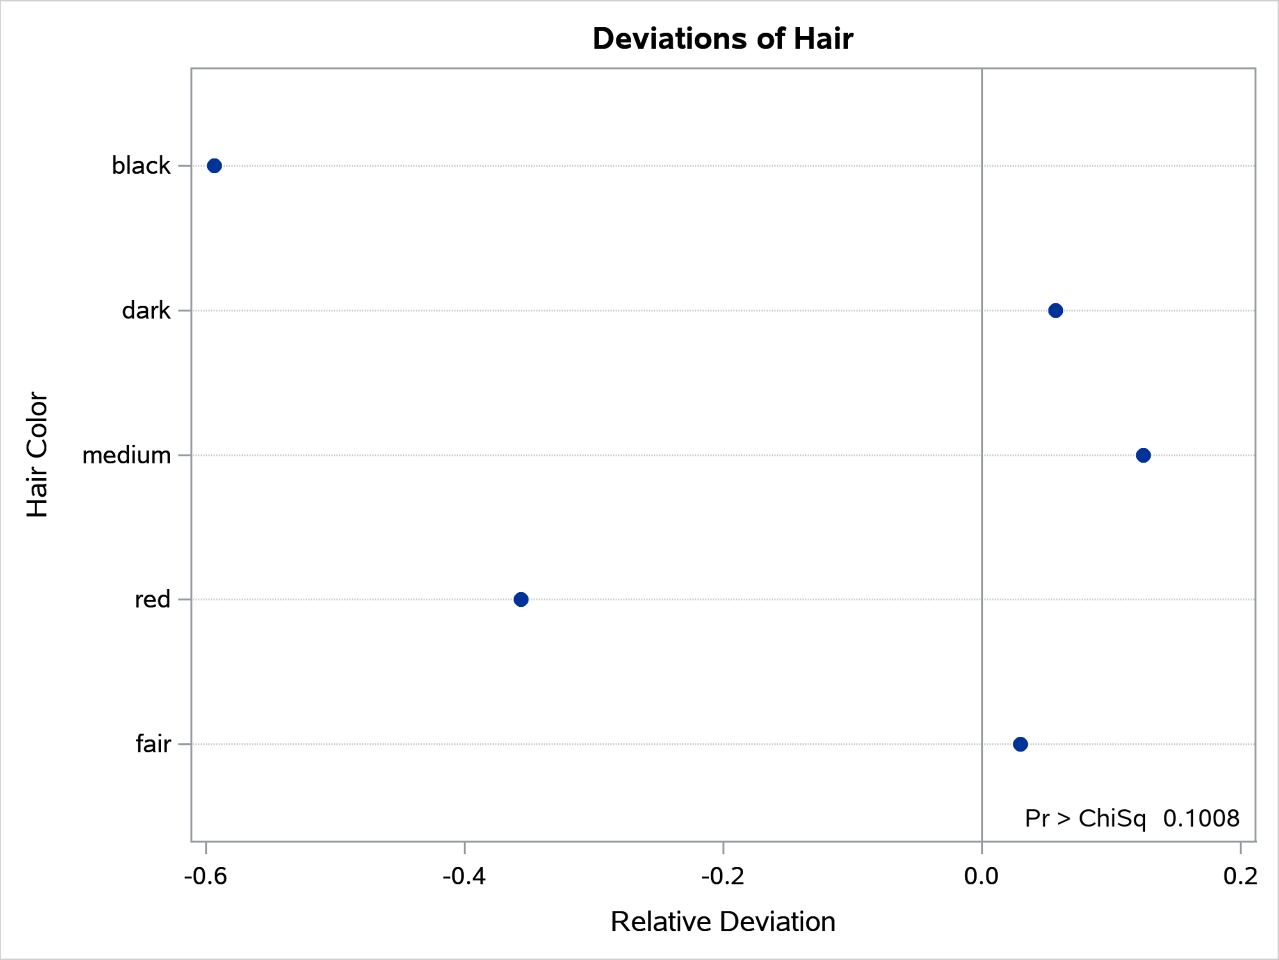

Output 47.3.3 and Output 47.3.4 show the results for Region 2. PROC FREQ computes a chi-square statistic for each region. The chi-square statistic is significant at the 0.05 level for Region 2 (p=0.0003) but not for Region 1. This indicates a significant departure from the hypothesized percentages in Region 2.

Output 47.3.3 Frequency Table and Chi-Square Test for Region 2

Hair Color of European Children

The FREQ Procedure

Geographic Region=2

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

Output 47.3.3 *continued*

Geographic Region=2

Chi-Square Test for Specified Proportions	
Chi-Square	21.3824
DF	4
Pr > ChiSq	0.0003

Output 47.3.4 Deviation Plot for Region 2

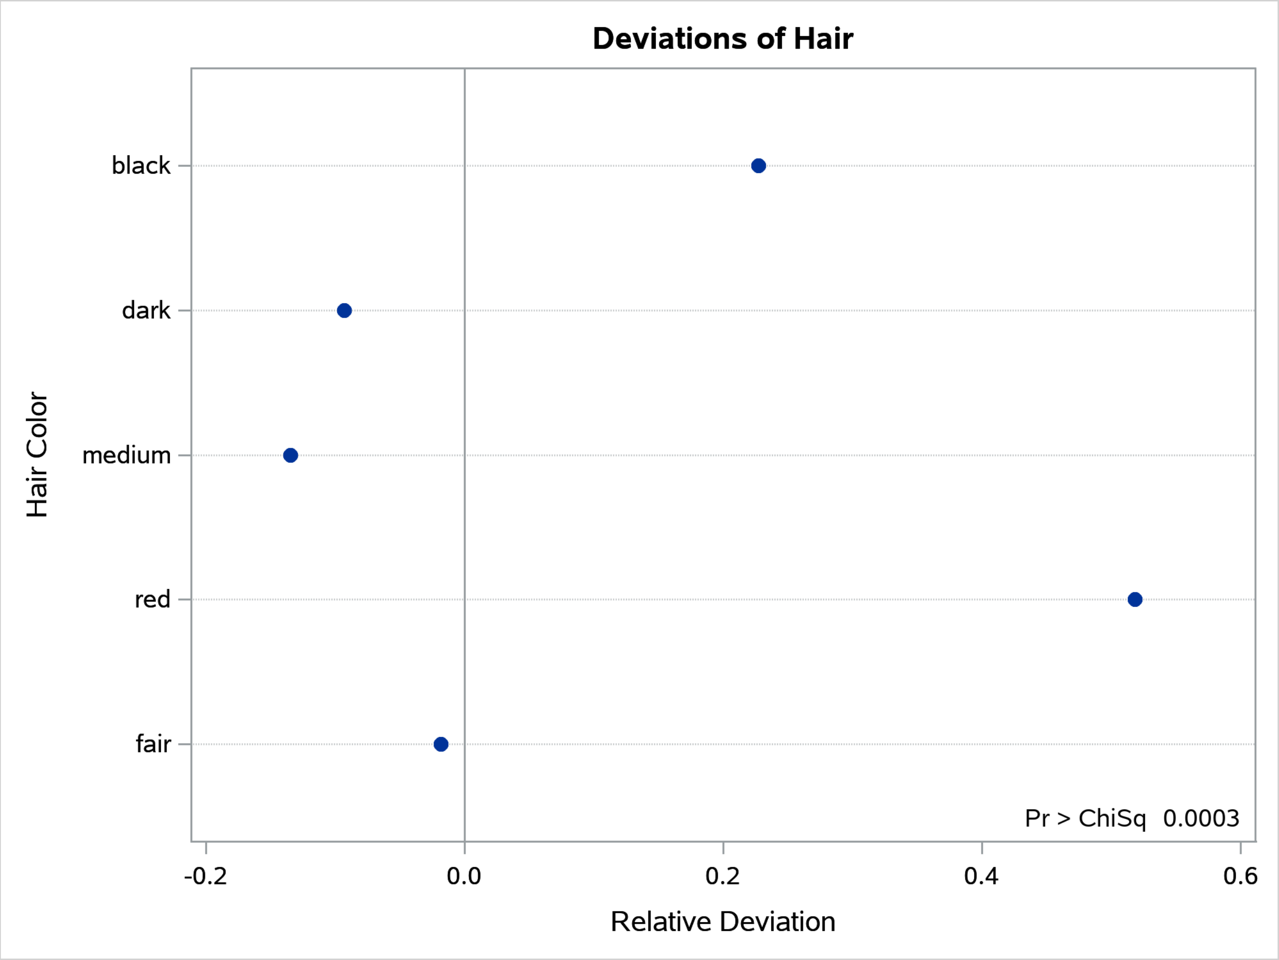

Example 47.4: Binomial Proportions

In this example, PROC FREQ computes binomial proportions, confidence limits, and tests. The example uses the eye and hair color data from Example 47.1. 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 47.1) and test the null hypothesis that the population proportion equals 50%. These statements also compute an equivalence test 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 is 0.5. You can specify a different test proportion in 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 47.4.1 displays the results for eye color, and Output 47.4.2 displays the results for hair color.

Output 47.4.1 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

Output 47.4.1 *continued*

Confidence Limits for the Binomial Proportion

Proportion = 0.4475

Type	90% Confidence Limits	
Agresti-Coull	0.4181	0.4773
Clopper-Pearson (Exact)	0.4174	0.4779
Wilson	0.4181	0.4773

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 47.4.1 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 47.4.1 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 0, 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 47.4.2 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 47.4.2 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 - p0 <= Lower Margin or >= Upper Margin

Ha: Lower Margin < P - p0 < Upper Margin

p0 = 0.28 Lower Margin = -0.1 Upper Margin = 0.1

Proportion	ASE (Sample)
0.2992	0.0166

Output 47.4.2 *continued*

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 47.5: 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, and 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 47.5.1 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 47.5.1 Contingency Table

```
             Case-Control Study of High Fat/Cholesterol Diet

                          The FREQ Procedure

Frequency   |          Table of Exposure by Response
Percent     |                         Response(Heart Disease)
Row Pct     |
Col Pct     |     Exposure               Yes        No     Total
------------+---------------------------------------------------
            | High Cholesterol Diet       11         4        15
            |                          47.83     17.39     65.22
            |                          73.33     26.67
            |                          84.62     40.00
            |
            | Low Cholesterol Diet         2         6         8
            |                           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 47.5.2 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, and therefore a one-sided test is appropriate. 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 47.5.3, 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 47.5.2 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 47.5.3 Relative Risk

Odds Ratio and Relative Risks

Statistic	Value	95% Confidence Limits	
Odds Ratio	8.2500	1.1535	59.0029
Relative Risk (Column 1)	2.9333	0.8502	10.1204
Relative Risk (Column 2)	0.3556	0.1403	0.9009

Odds Ratio

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 47.6: Output Data Set of Chi-Square Statistics

This example uses the Color data from Example 47.1 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 47.6.1 and Output 47.6.2. The contingency table in Output 47.6.1 displays eye and hair color in the order in which they appear in the Color data set. The Pearson chi-square statistic in Output 47.6.2 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 47.6.3 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 47.6.2.

Output 47.6.1 Contingency Table

Chi-Square Tests for 3 by 5 Table of Eye and Hair Color

The FREQ Procedure

Frequency Expected Cell Chi-Square Percent

Table of Eyes by Hair

	Hair(Hair Color)					
Eyes(Eye Color)	**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 47.6.2 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 47.6.3 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 47.7: 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 47.7.1 through Output 47.7.4 show the results of the analysis. The relative risk plot (Output 47.7.1) displays the relative risks and confidence limits for the two levels of Gender and for the overall (common) relative risk. Output 47.7.2 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 47.7.3. Because this is a prospective study, the relative risk estimate assesses the effectiveness of the new drug; the “Cohort (Col1 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 47.7.4 indicates no significant gender difference in the odds ratios.

Output 47.7.1 Relative Risk Plot

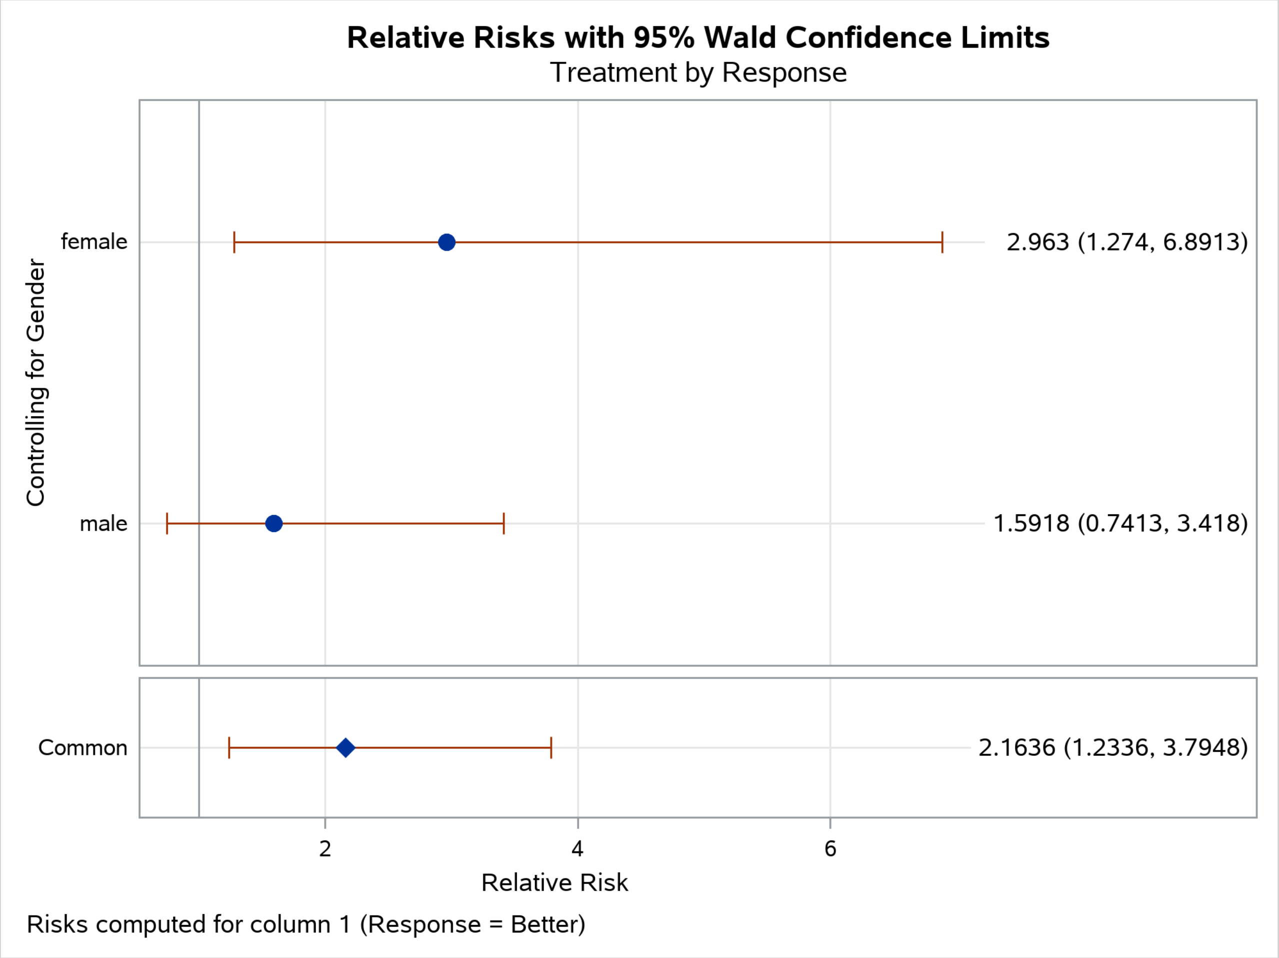

Output 47.7.2 Cochran-Mantel-Haenszel Statistics

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 47.7.3 CMH Option: Common Relative Risks

Common Odds Ratio and Relative Risks

Statistic	Method	Value	95% Confidence Limits	
Odds Ratio	Mantel-Haenszel	3.3132	1.4456	7.5934
	Logit	3.2941	1.4182	7.6515
Relative Risk (Column 1)	Mantel-Haenszel	2.1636	1.2336	3.7948
	Logit	2.1059	1.1951	3.7108
Relative Risk (Column 2)	Mantel-Haenszel	0.6420	0.4705	0.8761
	Logit	0.6613	0.4852	0.9013

Output 47.7.4 CMH Option: Breslow-Day Test

Breslow-Day Test for Homogeneity of Odds Ratios	
Chi-Square	1.4929
DF	1
Pr > ChiSq	0.2218

Example 47.8: 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 47.8.1 through Output 47.8.4 display the results of the analysis. The "Col Pct" values in Output 47.8.1 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 47.8.2) also shows this increasing trend.

Output 47.8.1 Contingency Table

Clinical Trial for Treatment of Pain

The FREQ Procedure

Frequency
Percent
Row Pct
Col Pct

Table of Adverse by Dose

	Dose					
Adverse	**0**	**1**	**2**	**3**	**4**	**Total**
No	26 16.15 25.49 81.25	26 16.15 25.49 78.79	23 14.29 22.55 71.88	18 11.18 17.65 56.25	9 5.59 8.82 28.13	102 63.35
Yes	6 3.73 10.17 18.75	7 4.35 11.86 21.21	9 5.59 15.25 28.13	14 8.70 23.73 43.75	23 14.29 38.98 71.88	59 36.65
Total	32 19.88	33 20.50	32 19.88	32 19.88	32 19.88	161 100.00

Output 47.8.2 Mosaic Plot

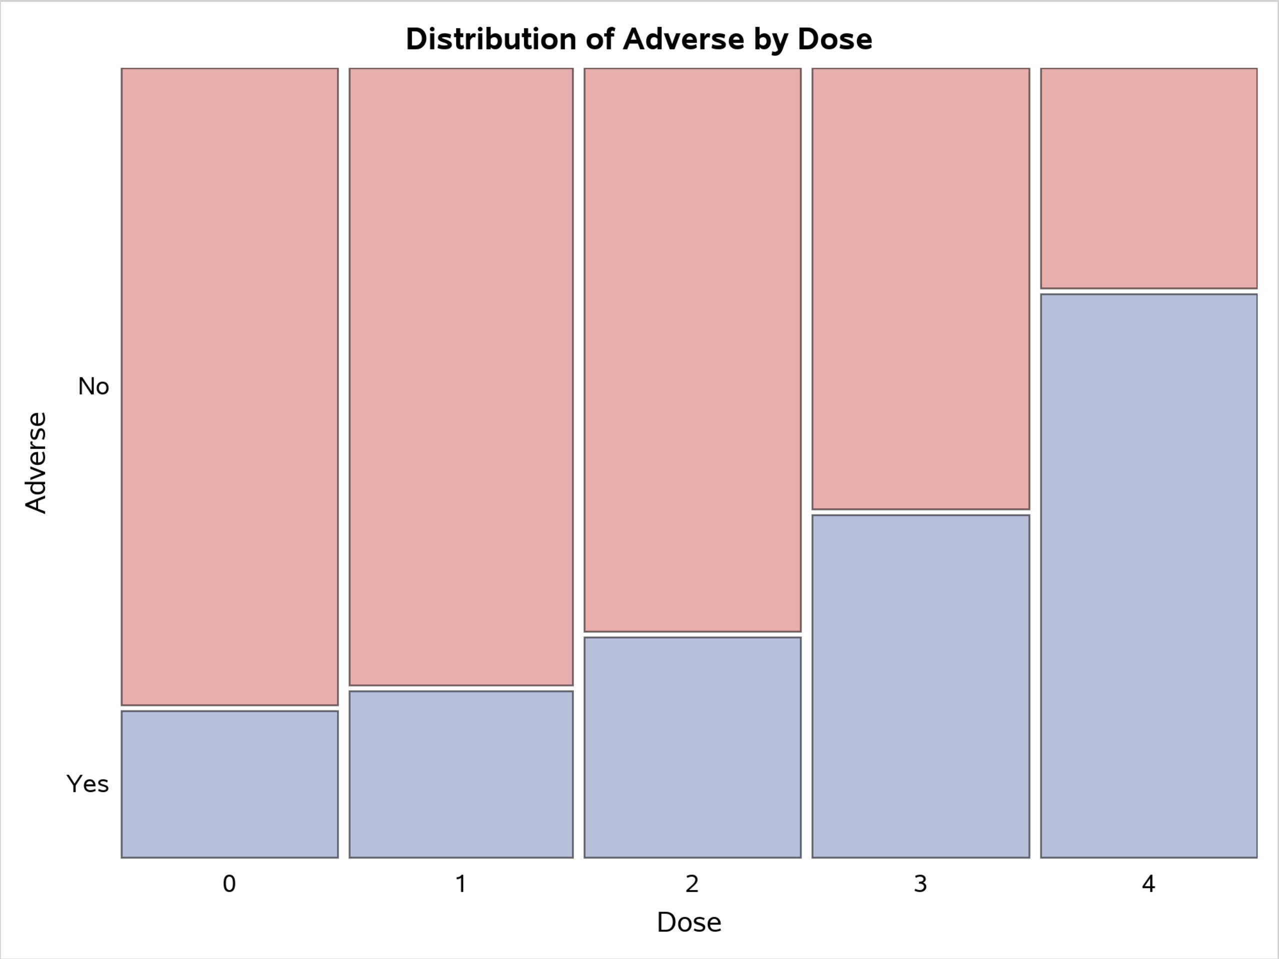

Output 47.8.3 displays the measures of association produced by the MEASURES option. Somers' $D(R|C)$

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 0, 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 47.8.4) 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 47.8.3 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 47.8.4 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	<.0001
Two-sided Pr >= |Z|	<.0001

Example 47.9: 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, 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 and D'Abrera 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 47.9.1 and Output 47.9.2.

```
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 47.9.1 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 47.9.2). Thus, adjusting for subject is critical to reducing the background variation due to subject differences.

Output 47.9.1 CMH Statistics: Stratifying by Subject

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 47.9.2 CMH Statistics: No Stratification

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 47.10: 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 47.10.1 through Output 47.10.5.

```
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 47.10.1 provide the marginal response for each drug. For drugs A and B, 61% of the subjects reported a favorable response; for drug C, 35% of the subjects reported a favorable response. Output 47.10.2 and Output 47.10.3 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 47.10.1 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

Output 47.10.1 *continued*

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 47.10.2 Measures of Agreement for Drug A Favorable

McNemar's Test

Chi-Square	DF	Pr > ChiSq
10.8889	1	0.0010

Simple Kappa Coefficient

Estimate	Standard Error	95% Confidence Limits	
-0.0328	0.1167	-0.2615	0.1960

Output 47.10.3 Measures of Agreement for Drug A Unfavorable

McNemar's Test

Chi-Square	DF	Pr > ChiSq
0.4000	1	0.5271

Simple Kappa Coefficient

Estimate	Standard Error	95% Confidence Limits	
-0.1538	0.2230	-0.5909	0.2832

Output 47.10.4 displays the overall kappa coefficient. The small negative value of kappa indicates no agreement between drug B response and drug C response.

Output 47.10.4 Overall Measures of Agreement

Overall Kappa Coefficient

Estimate	Standard Error	95% Confidence Limits	
-0.0588	0.1034	-0.2615	0.1439

Test for Equal Kappas

Chi-Square	DF	Pr > ChiSq
0.2314	1	0.6305

Cochran's Q is statistically significant (p=0.0145 in Output 47.10.5), which leads to rejection of the hypothesis that the probability of favorable response is the same for the three drugs.

Output 47.10.5 Cochran's Q Test

Cochran's Q, for Drug_A by Drug_B by Drug_C

Chi-Square	DF	Pr > ChiSq
8.4706	2	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*. 2nd ed. New York: John Wiley & Sons.

Agresti, A. (2007). *An Introduction to Categorical Data Analysis*. 2nd ed. New York: John Wiley & Sons.

Agresti, A. (2013). *Categorical Data Analysis*. 3rd ed. Hoboken, NJ: 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.” *American Statistician* 52:119–126.

Agresti, A., and Gottard, A. (2007). “Nonconservative Exact Small-Sample Inference for Discrete Data.” *Computational Statistics and Data Analysis* 51:6447–6458.

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, Department of Biostatistics, University of North Carolina at Chapel Hill.

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?” *American Statistician* 55:279–287.

Barnard, G. A. (1945). “A New Test for 2 × 2 Tables.” *Nature* 156:177.

Barnard, G. A. (1947). “Significance Tests for 2 × 2 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*. 2nd ed. New York: Springer-Verlag.

Birch, M. W. (1965). “The Detection of Partial Association, Part 2: 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.

Blaker, H. (2000). “Confidence Curves and Improved Exact Confidence Intervals for Discrete Distributions.” *Canadian Journal of Statistics* 28:783–798.

Blood, E., and Spratt, K. F. (2007). “Disagreement on Agreement: Two Alternative Agreement Coefficients.” In *Proceedings of the SAS Global Forum 2007 Conference*. Cary, NC: SAS Institute Inc. `http://www2.sas.com/proceedings/forum2007/186-2007.pdf`.

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). *The Analysis of Case-Control Studies*. Statistical Methods in Cancer Research, IARC Scientific Publications, vol. 1, no. 32. Lyon: International Agency for Research on Cancer.

Breslow, N. E., and Day, N. E. (1987). *The Design and Analysis of Cohort Studies.* Statistical Methods in Cancer Research, IARC Scientific Publications, vol. 2, no. 82. Lyon: 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. (1977a). “On the Mean and Variance of the Tetrachoric Correlation Coefficient.” *Psychometrika* 42:347–355.

Brown, M. B., and Benedetti, J. K. (1977b). “Sampling Behavior of Tests for Correlation in Two-Way Contingency Tables.” *Journal of the American Statistical Association* 72:309–315.

Byrt, T., Bishop, J., and Carlin, J. B. (1993). “Bias, Prevalence, and Kappa.” *Journal of Clinical Epidemiology* 46:423–429.

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.

Chow, S.-C., Shao, J., and Wang, H. (2008). *Sample Size Calculations in Clinical Research.* 2nd ed. Boca Raton, FL: Chapman & Hall/CRC.

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.

Drasgow, F. (1986). "Polychoric and Polyserial Correlations." In *Encyclopedia of Statistical Sciences*, vol. 7, edited by S. Kotz, N. L. Johnson, and C. B. Read. New York: John Wiley & Sons.

Dunnett, 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.* 2nd ed. Cambridge, MA: MIT Press.

Fisher, K. (2015). "Evaluating Methods of Estimating Common Risk Difference for Stratified Binomial Trials for Less Common Events." Poster presented at the Duke-Industry Statistics Symposium, Durham, NC. `https://sites.duke.edu/biostatsworkshop2015/files/2015/10/P8_Poster_for_Duke_Symposium_Kate_Fisher_101915_Final.pdf`.

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*. 3rd ed. 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. H., 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. H., 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.

Gwet, K. L. (2008). "Computing Inter-rater Reliability and Its Variance in the Presence of High Agreement." *British Journal of Mathematical and Statistical Psychology* 61:29–48.

Haldane, J. B. S. (1956). "The Estimation and Significance of the Logarithm of a Ratio of Frequencies." *Annals of Human Genetics* 20:309–311.

Hauck, W. W., and Anderson, S. (1986). "A Comparison of Large-Sample Confidence Interval Methods for the Difference of Two Binomial Probabilities." *American Statistician* 40:318–322.

Higgins, J. P. T., and Thompson, S. G. (2002). "Quantifying Heterogeneity in a Meta-analysis." *Statistics in Medicine* 21:1539–1558.

Higgins, J. P. T., Thompson, S. G., Deeks, J. J., and Altman, D. G. (2003). "Measuring Inconsistency in Meta-analyses." *BMJ* 327:557–560.

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.

Hoenig, J. M., Morgan, M. J., and Brown, C. A. (1995). "Analysing Differences between Two Age Determination Methods by Tests of Symmetry." *Canadian Journal of Fisheries and Aquatic Sciences* 52:364–368.

Hollander, M., and Wolfe, D. A. (1999). *Nonparametric Statistical Methods*. 2nd ed. New York: John Wiley & Sons.

Holley, J. W., and Guilford, J. P. (1964). "A Note on the G Index of Agreement." *Educational and Psychological Measurement* 24:749–753.

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. G. (1955). *Rank Correlation Methods*. 2nd ed. London: Charles Griffin.

Kendall, M. G., and Stuart, A. (1979). *The Advanced Theory of Statistics.* 4th ed. Vol. 2. New York: Macmillan.

Kim, Y., and Won, S. (2013). "Adjusted Proportion Difference and Confidence Interval in Stratified Randomized Trials." In *Proceedings of PharmaSUG 2013 (Pharmaceutical Industry SAS Users Group).* Paper SP04. Cary, NC: SAS Institute Inc.

Kleinbaum, D. G., Kupper, L. L., and Morgenstern, H. (1982). *Epidemiologic Research: Principles and Quantitative Methods*. New York: Van Nostrand Reinhold.

Klingenberg, B. (2014). "A New and Improved Confidence Interval for the Mantel-Haenszel Risk Difference." *Statistics in Medicine* 33:2968–2983.

Korn, E. L., and Graubard, B. I. (1998). "Confidence Intervals for Proportions with Small Expected Number of Positive Counts Estimated from Survey Data." *Survey Methodology* 24:193–201.

Krauth, J. (1973). "Nichtparametrische Ansätze zur Auswertung von Verlaufskurven [Nonparametric approaches to analyzing time effect curves]." *Biometrische Zeitschrift* 15:557–566.

Landis, J. R., 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." *American Statistician* 50:63–68.

Lehmann, E. L., and D'Abrera, H. J. M. (2006). *Nonparametrics: Statistical Methods Based on Ranks*. Rev. ed. New York: Springer.

Liebetrau, A. M. (1983). *Measures of Association.* Vol. 32 of Quantitative Applications in the Social Sciences. Beverly Hills, CA: 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 *Encyclopedia of Statistical Sciences*, vol. 9, edited by S. Kotz, N. L. Johnson, and C. B. Read. 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.

Mehrotra, D. V. (2001). “Stratification Issues with Binary Endpoints.” *Drug Information Journal* 35:1343–1350.

Mehrotra, D. V., and Railkar, R. (2000). “Minimum Risk Weights for Comparing Treatments in Stratified Binomial Trials.” *Statistics in Medicine* 19:811–825.

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. J. (1985). “Computing an Exact Confidence Interval for the Common Odds Ratio in Several 2×2 Contingency Tables.” *Journal of the 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 *Computing Science and Statistics: Proceedings of the Twenty-Third Symposium on the Interface*, edited by E. M. Keramidas, 200–207. Fairfax Station, VA: Interface Foundation.

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.” Cambridge, MA: Cytel Software Corporation.

Miettinen, O. S. (1985). *Theoretical Epidemiology: Principles of Occurrence in Research Medicine*. New York: John Wiley & Sons.

Miettinen, O. S., and Nurminen, M. 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 *Encyclopedia of Statistical Sciences*, vol. 4, edited by S. Kotz, N. L. Johnson, and C. B. Read. 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.

Sato, T. (1989). "On the Variance Estimator of the Mantel-Haenszel Risk Difference." *Biometrics* 45:1323–1324. Letter to the editor.

Schuirmann, 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.

Schuirmann, D. J. (1999). "Confidence Interval Methods for Bioequivalence Testing with Binomial Endpoints." In *Proceedings of the Biopharmaceutical Section*, 227–232. Alexandria, VA: American Statistical Association.

Silvapulle, M. J. (2001). "Tests against Qualitative Interaction: Exact Critical Values and Robust Tests." *Biometrics* 57:1157–1165.

Sim, J., and Wright, C. C. (2005). "The Kappa Statistic in Reliability Studies: Use, Interpretation, and Sample Size Requirements." *Physical Therapy* 85:257–268.

Snedecor, G. W., and Cochran, W. G. (1989). *Statistical Methods*. 8th ed. 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*. 3rd ed. Cary, NC: SAS Institute Inc.

Suissa, S., and Shuster, J. J. (1985). "Exact Unconditional Sample Sizes for the 2×2 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* 72:91–95.

Theil, H. (1972). *Statistical Decomposition Analysis*. Amsterdam: North-Holland.

Thomas, D. G. (1971). "Algorithm AS-36: Exact Confidence Limits for the Odds Ratio in a 2×2 Table." *Journal of the Royal Statistical Society, Series C* 20:105–110.

Thorlund, K., Imberger, G., Johnston, B. C., Walsh, M., Awad, T., Thabane, L., Gluud, C., et al. (2012). "Evolution of Heterogeneity (I^2) Estimates and Their 95% Confidence Intervals in Large Meta-analyses." *PLoS ONE* 7:1–8. `https://doi.org/10.1371/journal.pone.0039471`.

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.

Xie, Q. (2013). "Agree or Disagree? A Demonstration of an Alternative Statistic to Cohen's Kappa for Measuring the Extent and Reliability of Agreement between Observers." Paper presented at FCSM Research Conference, Nov. 7–9, Washington, DC.

Yan, X., and Su, X. G. (2010). "Stratified Wilson and Newcombe Confidence Intervals for Multiple Binomial Proportions." *Statistics in Biopharmaceutical Research* 2:329–335.

Zelen, M. (1971). "The Analysis of Several 2 × 2 Contingency Tables." *Biometrika* 58:129–137.