Chapter 12: Chi-Square Tests of Independence and Goodness-of-Fit

Chi-Square Test of Independence

Overview

Many variables analyzed in the social sciences are categorical, that is, measured on a nominal scale, and so take on values that differ only qualitatively from one another. For example, gender can be thought of as a categorical variable that can take on the values male and female. Nationality (e.g., Australian, Ethiopian, Canadian), personality type (type A, type B), and hair color are other examples.

When cases are categorized simultaneously into values of two nominally-scaled variables, the outcome can be represented in the form of a contingency table. For example, if one simultaneously classifies alcoholics according to their personality type (type A, type B) and gender (male, female) the results can be depicted in a contingency table such as that shown in Figure 12.1. Note that the numbers in the body of the table represent the number of cases that fall into each joint classification. Note also that although the contingency table shown in Figure 12.1 consists of two rows and two columns (i.e., it is a 2 x 2 table), the number of rows and columns in contingency tables may assume other values (e.g., 2 x 3, 9 x 4) depending on the number of categories of each variable.

Figure 12.1. Contingency table representing the number of alcoholic patients classified as having type A or type B personalities as a function of their gender (modeled after Bottlender, Preuss, & Soyka, 2006).

		Personality Type	
		Type A	Type B
Gender	Male	128	66
	Female	33	11

A question that is sometimes appropriate to investigate for data represented in contingency tables is whether, based on the sample of scores represented in the table, the variables are independent in the population. An analysis that addresses the question is the chi-square (χ^2) test of independence.

To test the independence of two categorical variables with the chi-square test of independence, one calculates the frequency of cases in each cell of a contingency table that is expected assuming the variables are independent, and then summarizes the degree to which the obtained frequency counts in cells of the table depart from the expected values. The summary measure, called the Pearson chi-square statistic, is compared with values of an appropriate theoretical chi-square distribution. If the value of the Pearson chi-square statistic falls above a cutoff value in the distribution beyond which only a small percent of theoretical values are found, the assumption that the variables are independent is rejected.

Underlying Assumptions

Two major assumptions underlie proper use of the chi-square test of independence. One is that all observations are independent. Thus, obtaining two measurements from the same person is precluded as are other sorts of dependencies that could arise when data come from spouses or siblings of participants. A second assumption, one that is difficult to verify, is that the sampling distribution of deviations of the actual and expected frequency counts is normal in form. Satisfying this assumption generally requires that the sample size be sufficiently large. Cochran's (1954) rule addresses the requirement. The rule is that there should be no expected frequency values under 1 and that no more than 20% of the expected frequency values should be under than 5. Some have suggested that that Cochran's rule may be too conservative. Camilli and Hopkins (1978) found that in a 2 x 2 table, as long as the total sample size exceeds 20, expected values as low as one or two in up to two cells produced acceptable results in terms of type I errors. For data represented in 2 x 2 and larger tables, Wickens (1989) suggests the general rule that total sample size be at least four or five times the number of cells.

Calculations Used in the Analysis

To obtain the value of sample chi-square, expected values of the frequency counts of each cell in the summary table must first be calculated that are consistent with the assumption the two variables are independent. Because two variables are independent if values of one variable do not influence values of the other variable, expected cell frequencies that reflect independence between variables will conform to the following rule: the ratio of marginal frequency counts for one variable will be repeated at each level of the other variable. Consider the table in Figure 12.2 that includes marginal sums obtained by adding frequency counts along rows and columns of the table. Notice that, as indicated in the marginal values circled to the right of the body of the table, there are more males in the data set than females. As shown in the row marginal values, of the 238 cases, 194 are males and 44 are females. According to our rule for establishing independence, the same ratio of males and females found in the marginal values should occur among participants with type A personalities and among participants with type B personalities. Because there are 161 cases with type A personalities, independence would require that 194/238 of the 161 cases be males, and 44/238 be females. Similarly, for the 77 type B cases, independence would require that 194/238 of the 77 be males and 44/238 of the 77 be females. The outcomes of these calculations are shown in cells of the table in the middle panel of Figure 12.2.

How discrepant are the data obtained in the experiment from the data expected if the variables are independent? As shown in the bottom panel of Figure 12.2, the difference between the two is determined by subtracting from each cell frequency count that was obtained in the study (referred to as the observed frequency count) the frequency expected if the variables are independent. The results are shown in the column titled d_i in

Figure 12.2. To change the d_i values to a form appropriate for comparison to a chi-square distribution, each is squared and divided by the value expected based on the assumption of independence between the variables. When summed, the result is the Pearson chi-square value for the sample. Symbolically,

$$\chi^2_{\text{sample}} = \sum \frac{\left(f_{\text{observed}} - f_{\text{expected}}\right)^2}{f_{\text{expected}}}. \qquad \text{Equation 12.1}$$

The value of the Pearson chi-square statistic for our example is 1.34.

Figure 12.2. Top panel shows a contingency table representing the number of alcoholic patients classified as being having type A or type B personalities as a function of their gender. Middle panel shows the calculations of expected frequencies if the variables are independent. Bottom panel shows the value of χ^2_{sample} is calculated.

		Personality Type		
		Type A	Type B	
Gender	Male	128	66	194
	Female	33	11	44
		161	77	Total = 238

If the variables are independent, the ratio of marginal frequencies shown here…,

will be repeated here and here.

		Personality Type		
		Type A	Type B	
Gender	Male	$\frac{194}{238} \times 161 = 131.24$	$\frac{194}{238} \times 77 = 62.76$	194
	Female	$\frac{44}{238} \times 161 = 29.76$	$\frac{44}{238} \times 77 = 14.24$	44

$f_{\text{observed}} - f_{\text{expected}}$	$= d_i$	d_i^2	$d_i^2 / f_{\text{expected}}$	
128 – 131.24	= -3.24	10.50	10.50/131.24	= 0.08
66 – 62.76	= 3.24	10.50	10.50/62.76	= 0.17
33 – 29.76	= 3.24	10.50	10.50/29.76	= 0.35
11 – 14.24	= -3.24	10.50	10.50/14.24	= 0.74
			χ^2_{sample}	= 1.34

The Chi-Square Distribution

The theoretical chi-square distributions that are used as a basis for making decisions in the chi-square test of independence change shape in accord with their degrees of freedom. The distributions were derived by F.R. Helmert in 1876 (David & Edwards, 2001) to describe the population of ratios formed by dividing the variance of each randomly drawn sample of N scores weighted by its degrees of freedom to the variance of the population being sampled:

$$\chi^2 = \frac{(N-1)s^2}{\sigma^2} \qquad \text{Equation 12.2}$$

Substituting the definition of sample variance into the formula, that is, the formula

$$s^2 = \frac{\sum(x_i - \bar{x})^2}{N-1},$$

gives

$$\chi^2 = \frac{\sum(x_i - \bar{x})^2}{\sigma^2}. \qquad \text{Equation 12.3}$$

Equation 12.3 makes it apparent that chi-square distributions can be expressed in terms of squared z-scores. If a single case is randomly sampled from a normal population and its value expressed in terms of z, the distribution of the squared z values is a chi-square distribution with degrees of freedom = 1. This and other chi-square distributions are shown in Figure 12.3. All chi-square distributions consist of positive values, are positively skewed, have a mean equal to the distribution's degrees of freedom and a variance that is twice the value of the degrees of freedom. As the chi-square distribution's degrees of freedom increases, its shape approaches that of a normal distribution.

Figure 12.3. Probability density functions of chi-square having 1, 2, 4, or 8 degrees of freedom.

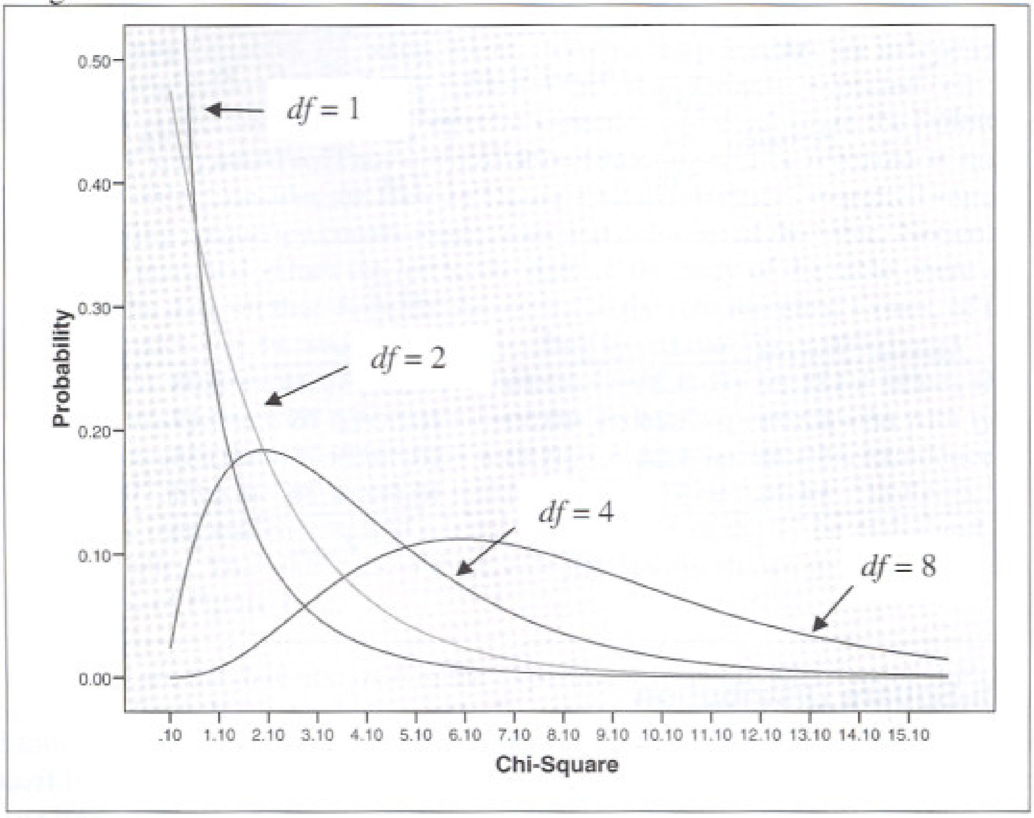

In the Pearson chi-square test of independence, degrees of freedom can be determined from the dimensions of the table used to represent the data. According to this method, degrees of freedom is given by the formula

$$df = (\#\text{ rows in the table} - 1) \times (\#\text{ columns in the table} - 1). \qquad \text{Equation 12.4}$$

Thus, for a 3 x 5 table, the degrees of freedom is 2 x 4 or 8. For the data shown in the 2 x 2 table in Figure 12.1, the degrees of freedom of the chi-square test of independence is 1. Because a table of chi-square values shows that, with one degree of freedom and an alpha of .05, the critical chi-square value is 3.84 our sample chi-square value of 1.34 is not sufficiently large to cause rejection of the assumption that the gender and personality type are unrelated.

Measuring Effect Size

Effect size for a chi-square test of independence may be summarized using the *phi* coefficient when the data conform to a 2 x 2 table and *Cramer's V* otherwise. The range of values of both coefficients is 0 to 1, values that correspond to complete independence and complete dependence of the variables, respectively. The formula for *phi* is

$$\phi = \sqrt{\frac{\chi^2}{N}}. \qquad \text{Equation 12.5}$$

In the formula for *Cramer's V* given below, *L* is the either the number of rows or the number of columns in the contingency table, whichever is smaller.

$$\mathit{Cramer's}\ V = \sqrt{\frac{\chi^2}{N(L-1)}} \qquad \text{Equation 12.6}$$

Cohen (1988) has suggested that *phi* values .10, .30, and .50 correspond to effects that could be described as small, medium, and large, respectively. Interpretation of the magnitude of *Cramer's V* is not as straightforward. For data represented in tables larger than 2 x 2 Cohen (1988) suggests that effect size be measured with the coefficient *w* (*w* is equivalent to *phi*). To convert *Cramer's V* to *w*, its *phi* equivalent, multiply its value by $\sqrt{(L-1)}$. Thus, a *Cramer's V* value of .14 obtained from a 3 x 4 table is equivalent to a *phi* or *w* value of .20 and so could be described as small to medium.

Preparing Data for the Analysis

A chi-square test of independence can be performed on data that represent values of two nominally-scaled variables for each case in a data file. For example, if the data of 238 participants were represented as individual cases in the **SPSS Data Editor**, and each case included a value of the variable **gender** (1 = male, 2 = female) and a value of the variable **perstype** (1 = type A, 2 = type B), a chi-square test of independence could be performed to determine if the variables are independent. A portion of such as file is shown below:

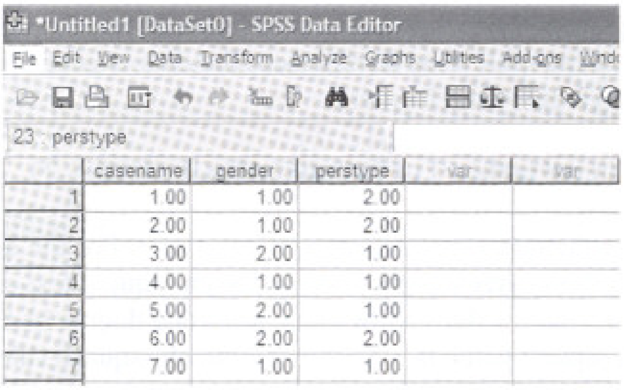

	casename	gender	perstype	var	var
1	1.00	1.00	2.00		
2	2.00	1.00	2.00		
3	3.00	2.00	1.00		
4	4.00	1.00	1.00		
5	5.00	2.00	1.00		
6	6.00	2.00	2.00		
7	7.00	1.00	1.00		

The data for each nominal scale can be represented using a variable declared as being either a string or a numeric type. Shown below are equivalent data represented with two string-type variables, **genderst** and **perstypest**:

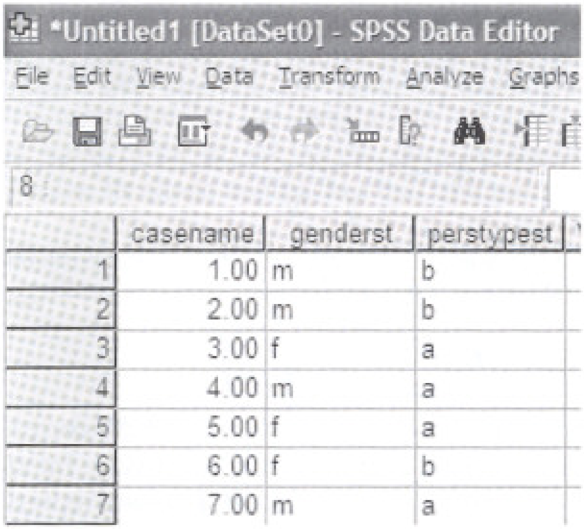

	casename	genderst	perstypest
1	1.00	m	b
2	2.00	m	b
3	3.00	f	a
4	4.00	m	a
5	5.00	f	a
6	6.00	f	b
7	7.00	m	a

An alternate representation of data suitable for analysis with a chi-square test of independence makes use of the **Weight Cases** function. The approach does not require coding each participant's data separately into a data file but, instead, weights the representation of each cell from a contingency table by its frequency count (i.e., the number of participants classified as falling into the joint classification represented by the cell). Three variables are required, two to represent each cell of the contingency table and one to represent the count in each cell. Data shown in the 2 x 2 table of Figure 12.1 is used to illustrate the coding scheme. Because the data are used to investigate the relation between gender and personality type, one variable in the **SPSS Data Editor** consists of the possible values of gender and a second consists of possible values of personality type. In effect, using this approach, each case is a cell in the table and there are as many cases as cells in the table. A third variable shows the frequency count of each cell identified by the other two variables.

The **SPSS Data Editor** spreadsheet shown below illustrates the coding system. The variable **gender**, using the values 1 (male) and 2 (female), and the variable **perstype,** using the values 1 (type A) and 2 (type B), identifies each cell in the contingency table. The variable **number** provides the frequency count of the cell.

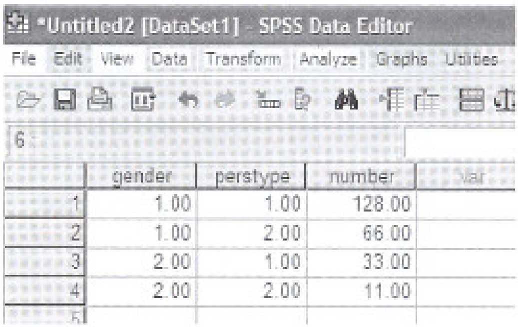

It is helpful to create variable and value labels for the variables used to identify cells. These are shown below for the data in our example:

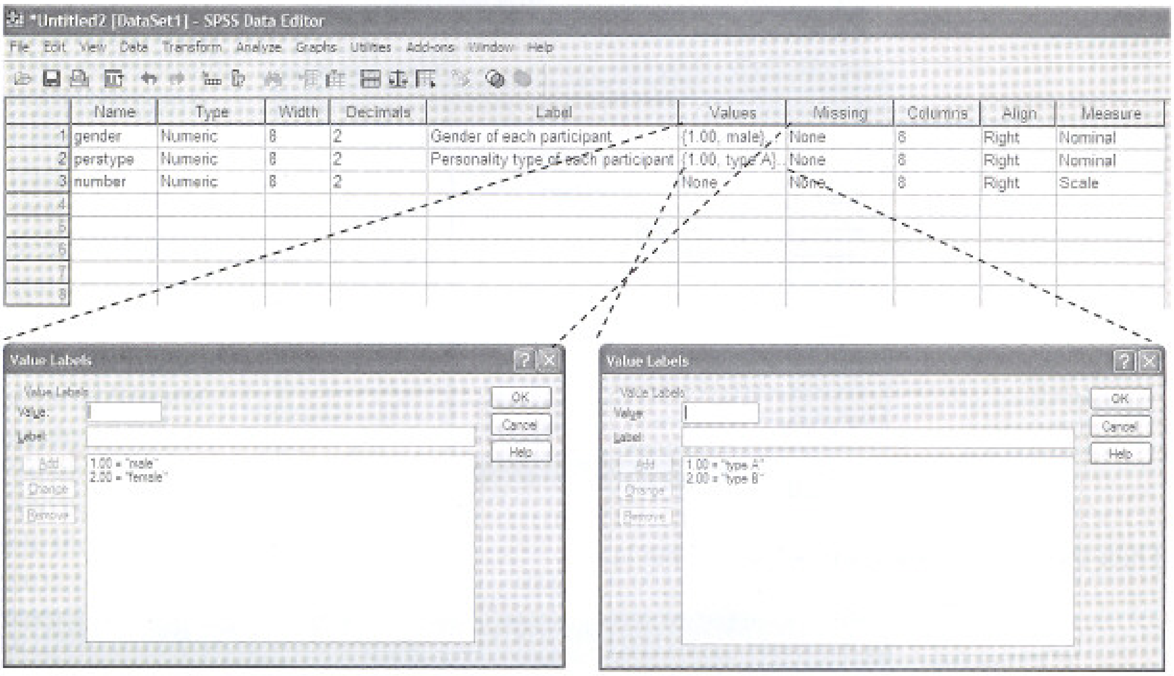

By selecting **Value Labels** in the **Data View**,

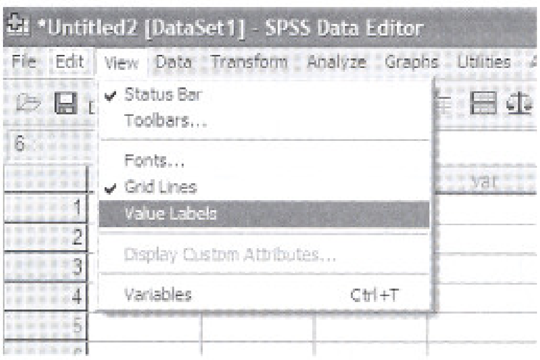

the cells become labeled and entry of the proper cell frequency count for the variable **number** is facilitated.

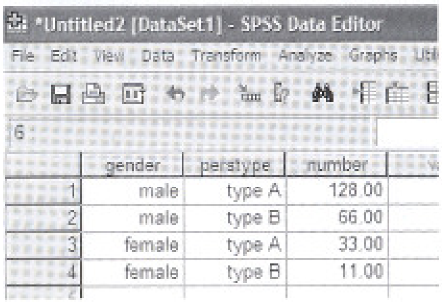

In the final step, the **Weight Cases** function is selected from the **Data** menu:

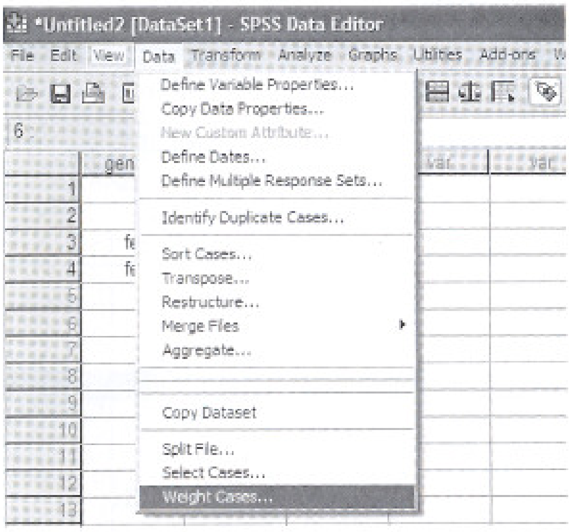

Then the selections shown below are made:

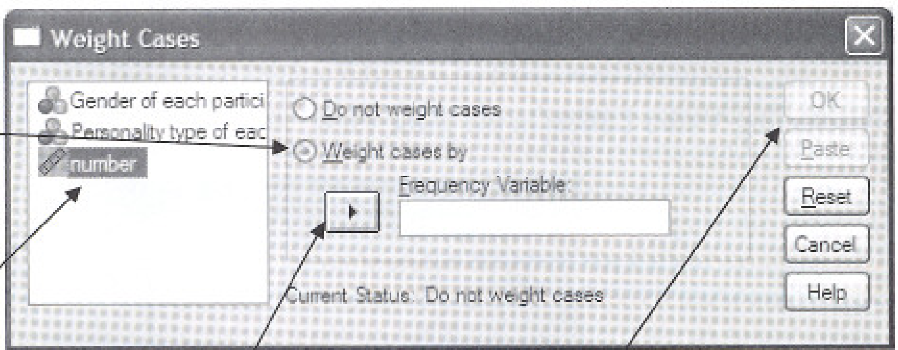

1. Click the **Weight cases by** button.

2. Click the variable name used to specify the frequency count of each cell.

3. Click to transfer the highlighted variable.

4. Click **OK**.

The data can now be analyzed using a chi-square test of independence.

Requesting a Chi-Square Test of Independence

The chi-square test of independence is accessed by selecting **Crosstabs** in the **Descriptive Statistics** menu as shown below.

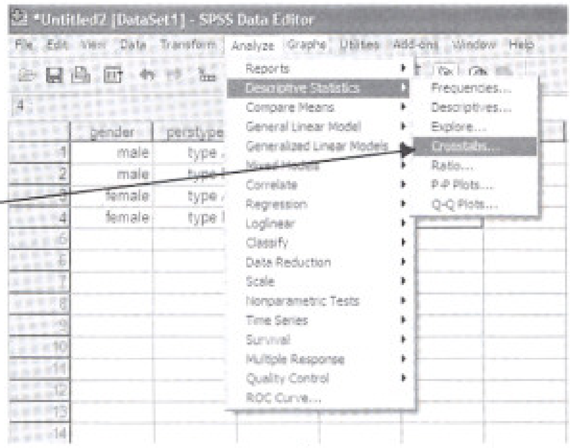

1. From the **Analyze** menu select **Descriptive Statistics** and then **Crosstabs**.

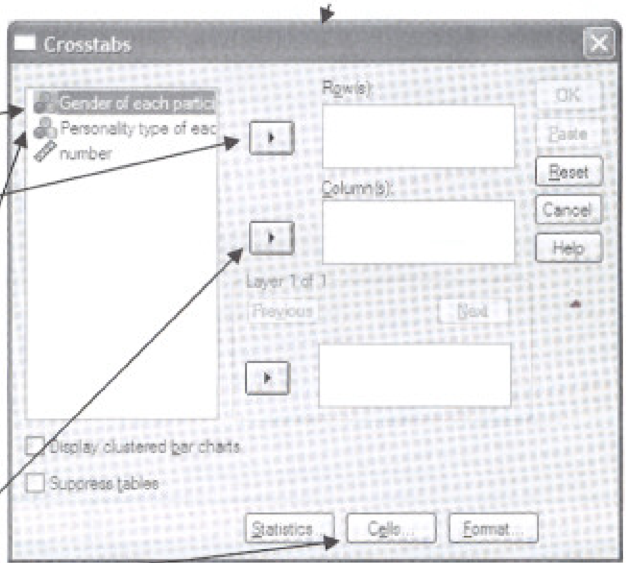

2. Click to highlight the variable name that identifies different rows in the contingency table then click this button to transfer it.

3. Click to highlight the variable name that identifies different columns in the contingency table then click this button to transfer it.

4. Click **Cells**.

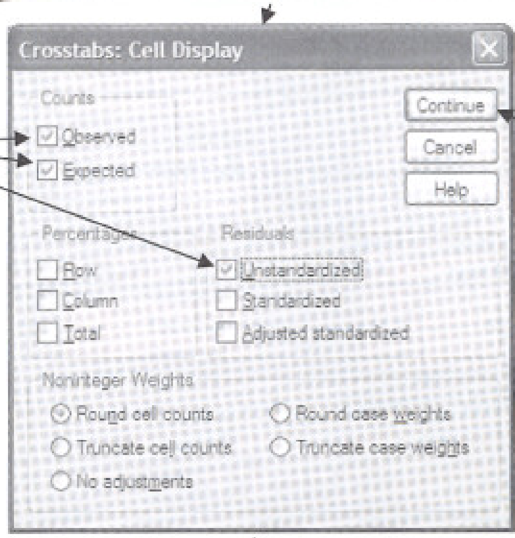

5. Click to select these options.

6. Click **Continue.**

7. Click **Statistics**.

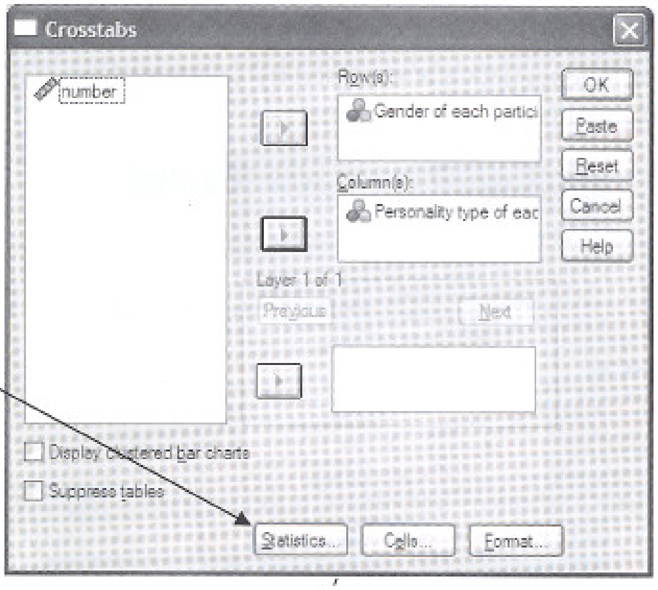

8. Click to select these options.

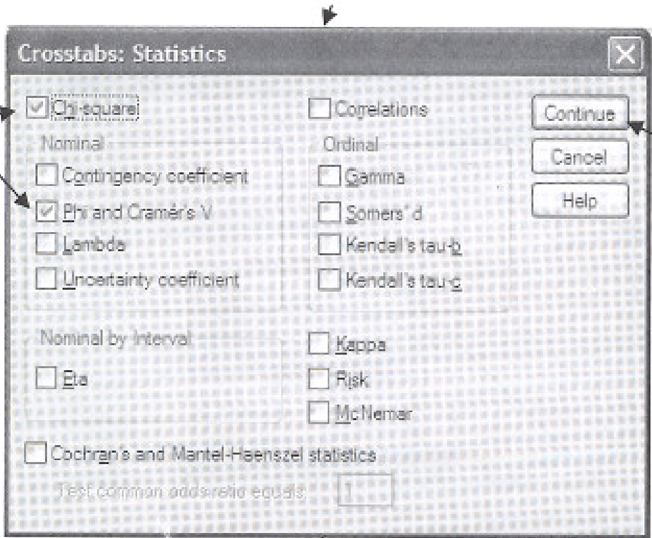

9. Click **Continue.**

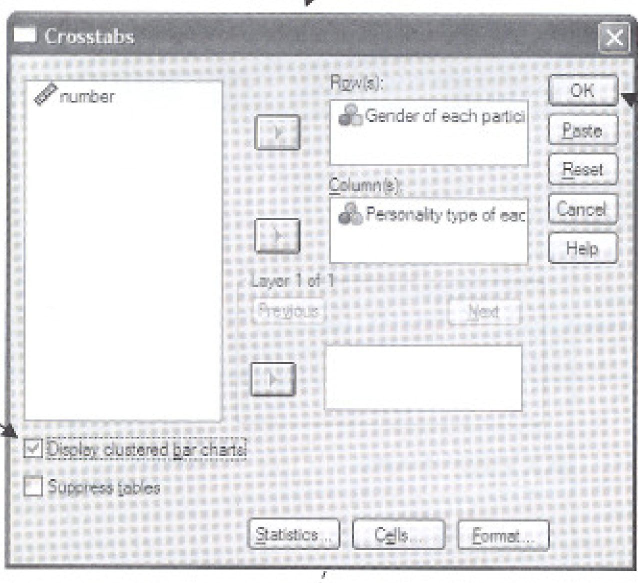

11. Click **OK**.

10. Click here to select this.

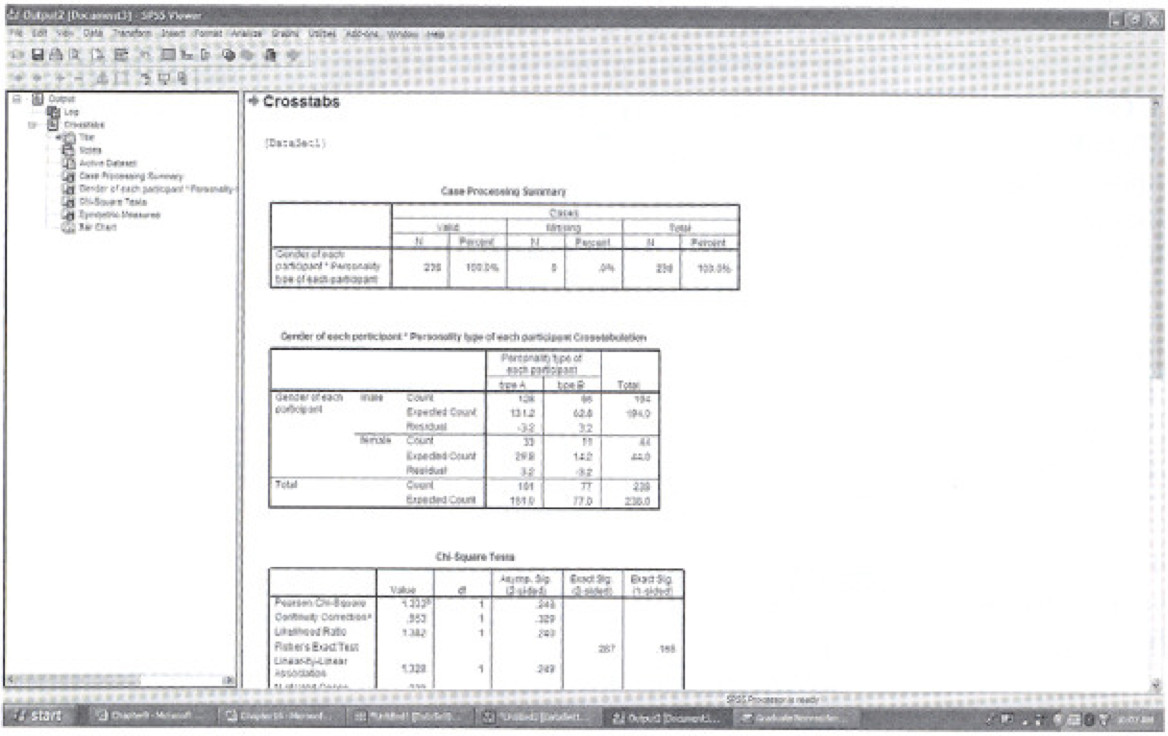

Interpreting the Output

The output includes a contingency table with the obtained and expected number of cases in each cell, a table that includes the value and significance level of the sample Pearson chi-square, a table that provides the value of *phi*, and a bar chart that shows the number of cases as a function of the category values used to identify each cell in the contingency table.

Case Processing Summary

	Cases					
	Valid		Missing		Total	
	N	Percent	N	Percent	N	Percent
Gender of each participant * Personality type of each participant	238	100.0%	0	.0%	238	100.0%

1. Inspect the contingency table to verify the data are entered correctly.

Gender of each participant * Personality type of each participant Crosstabulation

			Personality type of each participant		
			type A	type B	Total
Gender of each participant	male	Count			194
		Expected Count	131.2	62.8	194.0
		Residual	-3.2	3.2	
	female	Count			44
		Expected Count	29.8	14.2	44.0
		Residual	3.2	-3.2	
Total		Count	161	77	238
		Expected Count	161.0	77.0	238.0

These are the observed frequency counts.

3. Obtain the value of the Pearson chi-square statistic and its significance from this table.

Chi-Square Tests

	Value	df	Asymp. Sig. (2-sided)	Exact Sig. (2-sided)	Exact Sig. (1-sided)
Pearson Chi-Square		1			
Continuity Correction[a]	.953	1	.329		
Likelihood Ratio	1.382	1	.240		
Fisher's Exact Test				.287	.165
Linear-by-Linear Association	1.328	1	.249		
N of Valid Cases	238				

a. Computed only for a 2x2 table

b. 0 cells (.0%) have expected count less than 5. The minimum expected count is 14.24.

Because the significance is > .05, the two variables are not significantly related.

4. Obtain the value of the *phi* from this table.

Symmetric Measures

		Value	Approx. Sig.
Nominal by Nominal	Phi		.248
	Cramer's V	.075	.248
N of Valid Cases		238	

a. Not assuming the null hypothesis.

b. Using the asymptotic standard error assuming the null hypothesis.

Although *phi* can assume a negative value, it is generally reported as the positive value.

The significance of *phi* will be identical to that of the Pearson chi-square.

Bar Chart

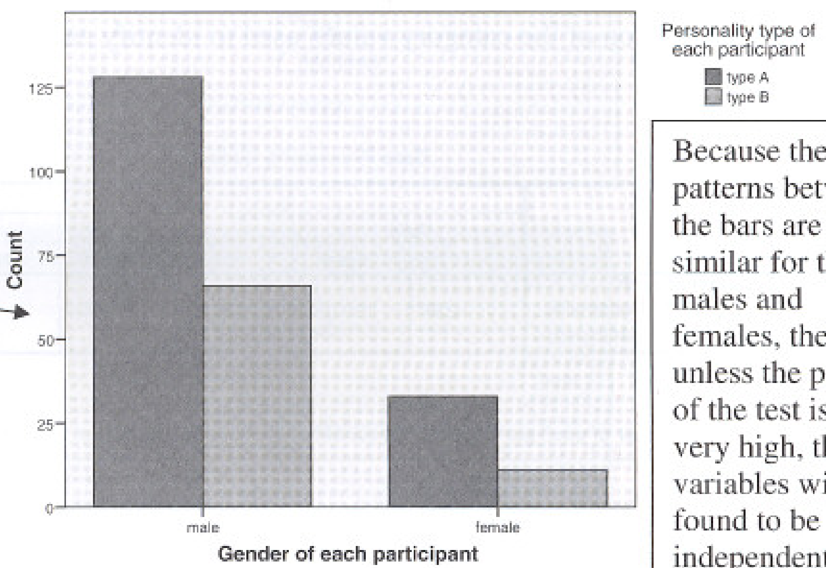

2. Inspect the bar chart to determine if the pattern of frequency counts among the bars within each cluster is similar over all the clusters. If it is, the variables are probably independent.

Because the patterns between the bars are so similar for the males and females, then, unless the power of the test is very high, the variables will be found to be independent.

Reporting the Results

The outcome of the test could be described as follows:

> To determine whether there was a relation between personality type (type A, type B) and gender (male, female) among detoxified inpatients at an alcohol dependence treatment ward, 238 patients (194 males, 44 females) were categorized as having personality types A or B. Of the 194 males the personalities of 128 were categorized as type A and 66 as type B; of the 44 females, the number of patients classified as type A and type B were 33 and 11, respectively. A chi-square test of independence indicated the relation between gender and personality type was not significant, χ^2 (1, N = 238) = 1.33, p > .20.

Notice that when the Pearson chi-square is not significant, there is no need to mention the value of *phi* because any difference from zero is attributed to chance. Analysis of data shown in Figure 12.4, however, produces a Pearson chi-square value sufficiently high to reject the assumption the variables are unrelated. The description of the outcome includes a report of the effect size.

Figure 12.4. Contingency table representing the number of undergraduate students reporting whether or not they ever cheated in a class as a function of their attitude towards the instructor.

		Attitude Towards Instructor	
		Liked	Disliked
Cheated?	Yes	240	41
	No	425	36

Tables that contain output that should be included in a description of the outcome are shown below:

Haveyou cheated? * Did you like the instructor? Crosstabulation

			Did you like the instructor?		Total
			like	no	
Haveyou cheated?	yes	Count	240	41	281
		Expected Count	251.8	29.2	281.0
		Residual	-11.8	11.8	
	no	Count	425	36	461
		Expected Count	413.2	47.8	461.0
		Residual	11.8	-11.8	
Total		Count	665	77	742
		Expected Count	665.0	77.0	742.0

This is a version of Pearson chi-square for 2 x 2 tables that makes the test more conservative (by subtracting .5 from each $f_{\text{observed}} - f_{\text{expected}}$ before squaring) and is rarely used.

This is an alternate statistic to the Pearson chi-square that also has a chi-square distribution. The values of the two statistics become more similar as the sample size increases.

This statistic is appropriate when each variable is ordinal rather than nominal.

Chi-Square Tests

	Value	df	Asymp. Sig. (2-sided)	Exact Sig. (2-sided)	Exact Sig. (1-sided)
Pearson Chi-Square	8.633[b]	1	.003		
	7.919	1	.005		
	8.372	1	.004		
Fisher's Exact Test				.004	.003
Association	8.621	1	.003		
N of Valid Cases	742				

a. Computed only for a 2x2 table

b. 0 cells (.0%) have expected count less than 5. The minimum expected count is 29.16.

Symmetric Measures

		Value	Approx. Sig.
Nominal by Nominal	Phi	-.108	.003
	Cramer's V	.108	.003
N of Valid Cases		742	

a. Not assuming the null hypothesis.

b. Using the asymptotic standard error assuming the null hypothesis.

A summary of the outcome is shown next.

> To determine if student cheating in a class is related to student attitude towards an instructor, students in a large introductory undergraduate class were asked to think of the class they took the previous year at the same or a similar time of day and to indicate both whether or not they cheated in any test or assignment for the class and whether they liked or disliked the instructor. Of the 665 students who indicated they liked the instructor, 240 (36%) reported cheating; of the 77 students who indicated they disliked the instructor 41 (53%) reported cheating. The outcome of a chi-square test of independence indicated that cheating was significantly related to attitude towards an instructor, χ^2 (1, N = 742) = 8.63, p < .01, ϕ = .11.

Chi-Square Goodness-of-Fit Test

Overview

A chi-square goodness-of-fit test is used to determine if the distribution of frequency counts among categories of a nominally-scaled variable (or a variable treated as though it were nominally-scaled) matches a theoretical or reference distribution. Although the test can be applied to nominally-scaled variables that have just two possible values, the preferred test for such dichotomous variables is the binomial test because it can provide exact probabilities for small sample sizes. However, for variables with more than two categories, the chi-square goodness-of-fit test may be appropriate. Consider, for example, the outcome of a survey intended to determine the attitude of a random sample of college students towards inviting a controversial speaker to campus. Of 78 students polled, 53 indicate they favor making the invitation, 20 are against making the invitation, and 5 are undecided. Consider, further, that among the entire faculty of the university, 80% had previously indicated they favored making the invitation, 10% were against making the invitation, and 10% were undecided. A chi-square goodness-of-fit test can be used to determine if student attitudes are compatible with those of the faculty.

The procedure is similar to that for the chi-square test of independence. In this case, however, the expected value of each cell is determined by the percent of cases in the corresponding cell of the reference distribution. For example, because 80% of the faculty members favored inviting the controversial speaker to campus, the expected frequency among the 78 students is .80 x 78 = 62.40. Calculation of expected frequencies for the two other attitude categories and calculations needed to obtain the value of chi-square are shown in Figure 12.5. The value of chi-square obtained for the sample is compared to a critical value from a distribution with degrees of freedom that equals the number of categories of the nominally-scaled variable – 1. Thus, in our example, because there are three possible attitudes, degrees of freedom is 2.

Figure 12.5. Top panel shows the number of students who favor, oppose, or are undecided about offering an invitation to visit the campus to a controversial speaker and the percent of faculty who hold each of these opinions. Bottom panel shows calculations for a chi-square goodness-of-fit test to assess whether student and faculty attitudes are identical.

	Attitude about Invitation			
	Favor	Oppose	Undecided	
Students	53	20	5	N = 78
Faculty	80%	10%	10%	
f_{expected}	.80 x 78 = 62.40	.10 x 78 = 7.80	.10 x 78 = 7.80	

$f_{\text{observed}} - f_{\text{expected}}$	= d_i	d_i^2	$d_i^2 / f_{\text{expected}}$	
53 – 62.40	= -9.40	88.36	88.36/62.40	= 1.42
20 – 7.80	= 12.20	148.84	148.84/7.80	= 19.08
5 – 7.80	= -2.80	7.84	7.84/7.80	= 1.01
			χ^2_{sample}	= 21.51

Measuring Effect Size

Effect size for a chi-square goodness-of-fit test can be measured using the statistic *w* suggested by Cohen (1988) and given by the formula

$$w = \sqrt{\sum_{i=1}^{m} \frac{\left(p_{observed_i} - p_{expected_i}\right)}{p_{expected_i}}} \qquad \text{Equation 12.7}$$

where m = the number of cells
$p_{observed\ i}$ = the proportion of cases found for cell i
$p_{expected\ i}$ = the proportion of cases expected for cell i.

The value of w can be most easily obtained from the chi-square statistic:

$$w = \sqrt{\frac{\chi^2}{N}} \qquad \text{Equation 12.8}$$

where N = the total sample size.

Interpretation of the magnitude of w is identical to that of *phi*: the values .10, .30, and .50 correspond to small, medium, and large effects, respectively.

Preparing Data for the Analysis

The most efficient way of coding data for our example into the **SPSS Data Editor** spreadsheet utilizes a table format as shown below:

Attitude is a numeric variable whose value labels are displayed using the **Value Labels** option in the **View** menu.

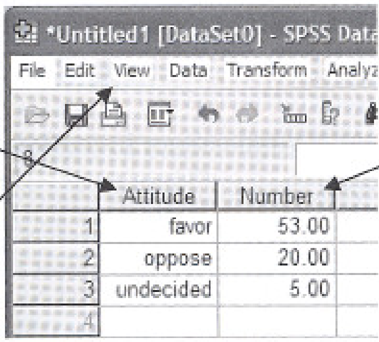

Number is a numeric variable that gives the frequency of each category of **Attitude**.

The cells must then be weighted using the value specified in the variable **Number**:

1. From the **Data** menu select **Weight Cases**

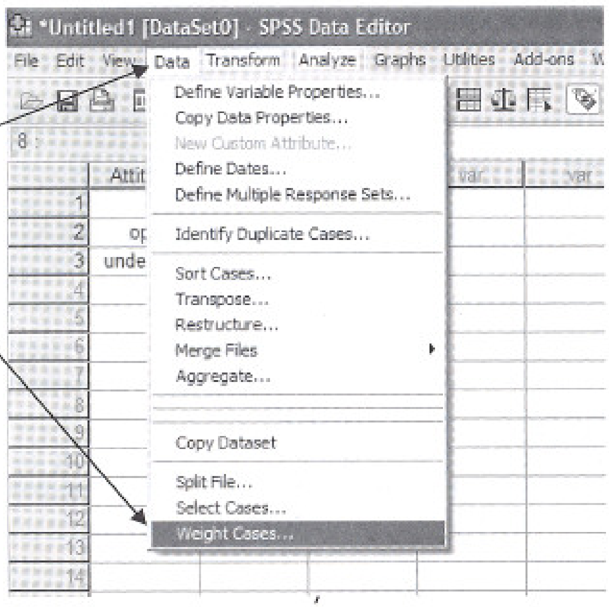

2. Click the **Weight cases by** button.

3. Click the variable name used to specify the frequency count of each cell.

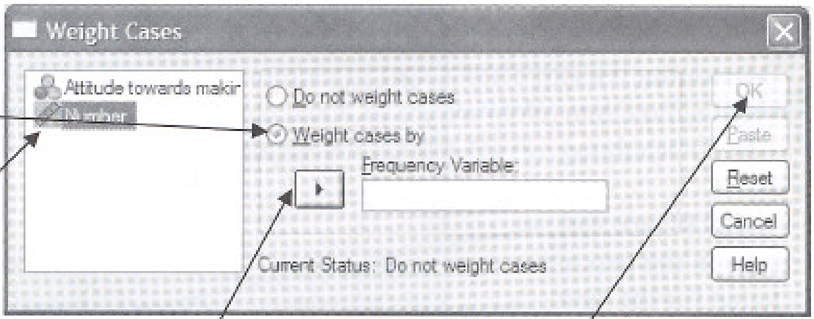

4. Click to transfer the highlighted variable.

5. Click **OK**.

Requesting a Chi-Square Goodness-of-Fit Test

The procedure for requesting a chi-square goodness-of-fit test is shown below:

1. From the **Analyze** menu select **Nonparametric Tests** and then **Chi-Square**.

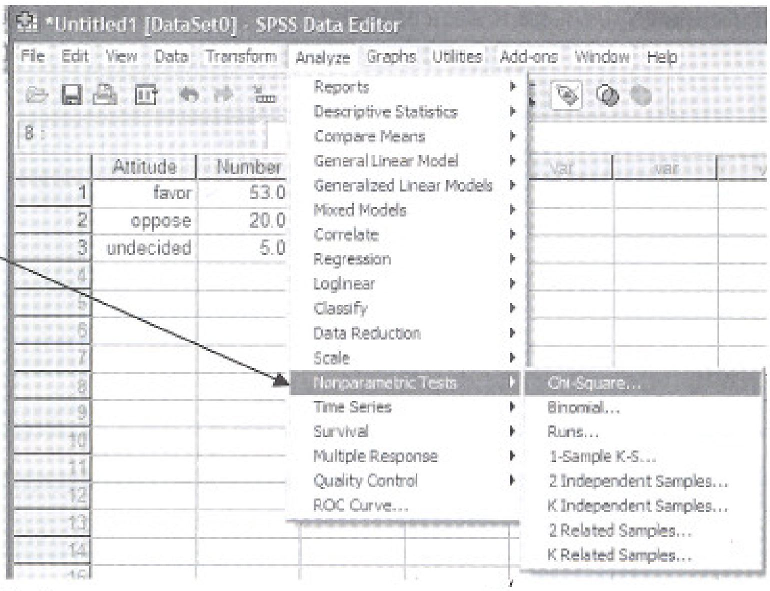

2. Click to highlight the variable name that identifies different rows in the contingency table then click this button to transfer it.

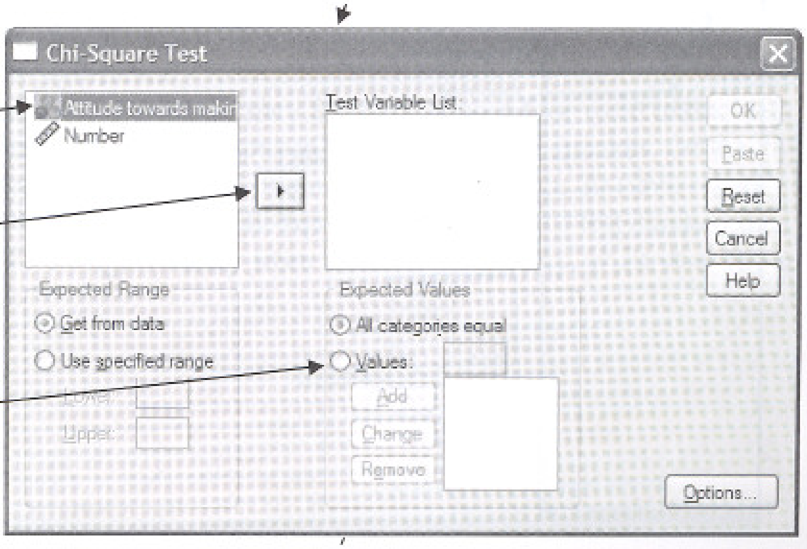

3. Click here to select **Values**.

4. Click here and type in the percent value for the first cell in the reference distribution, then click **Add**.

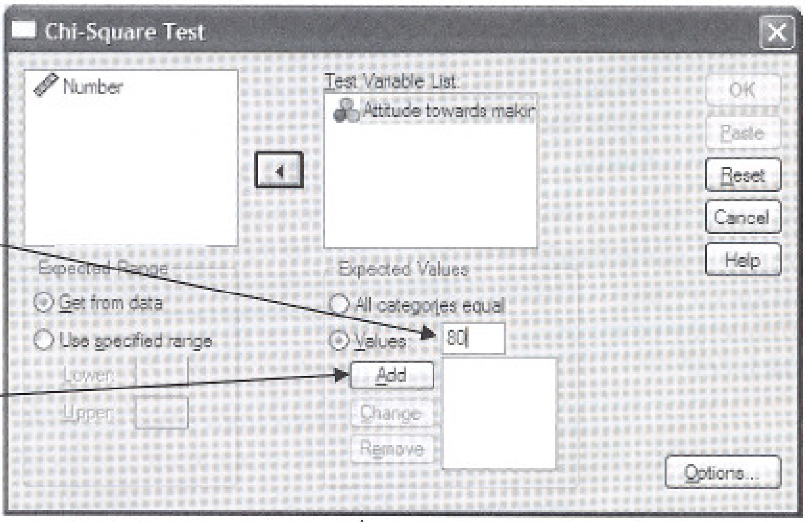

5. Repeat step 4 to enter the % for each cell of the reference distribution.

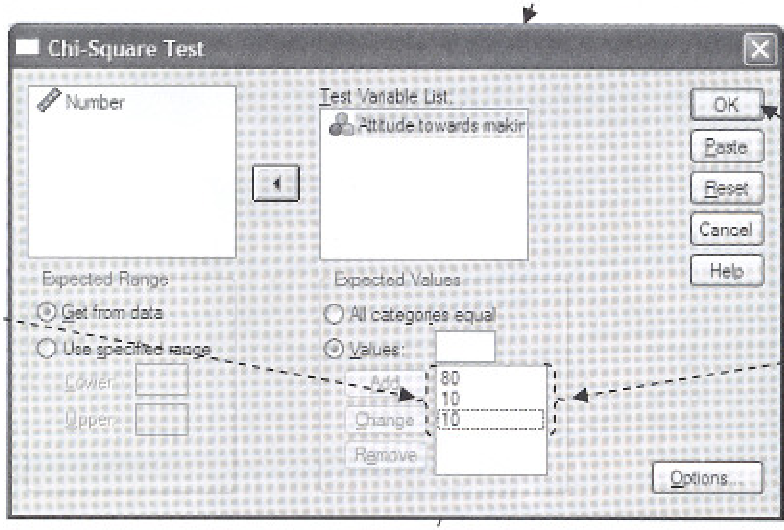

6. Click **OK**.

Note: Other values proportional to these could be used (e.g., 8, 1, 1).

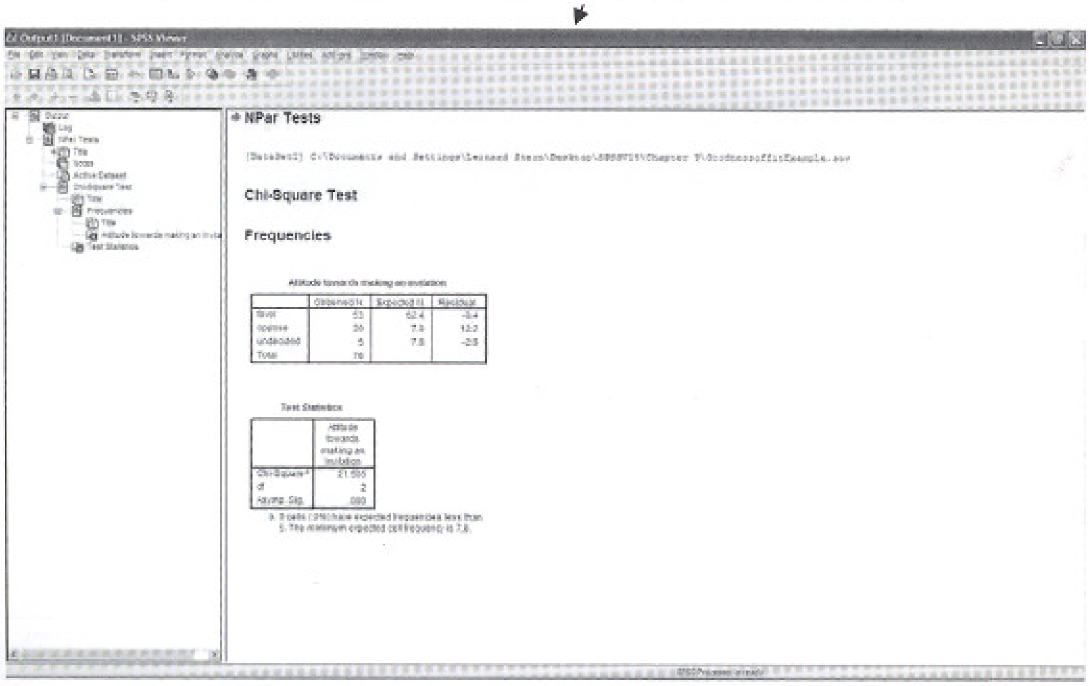

Interpreting the Output

The output includes a table with the obtained number of cases, the expected number of cases, and the difference between the two for each cell of the variable, and a table that gives the value and significance level of the sample chi-square for the goodness-of-fit test.

1. Inspect the observed N values in this table to verify the data are entered correctly.

Attitude towards making an invitation

	Observed N	Expected N	Residual
favor	53	62.4	-9.4
oppose	20	7.8	12.2
undecided	5	7.8	-2.8
Total	78		

2. Find the value of chi-square and its level of significance from this table.

Test Statistics

	Attitude towards making an invitation
Chi-Square[a]	21.503
df	2
Asymp. Sig.	

a. 0 cells (.0%) have expected frequencies less than 5. The minimum expected cell frequency is 7.8.

Because the significance level is < .05, the distribution in the population tested differs significantly from that in the reference group.

Reporting the Results

The outcome of the goodness-of-fit test can be described as follows:

> Data reported by the faculty senate indicated that among the faculty in the university, 80% favored extending an invitation to a controversial speaker to visit campus, 10% opposed extending an invitation, and 10% were undecided. To determine whether student attitudes differed significantly from those of the faculty, a randomly selected sample of 78 students was asked whether they favored, opposed, or were undecided about extending an invitation to visit campus to the controversial speaker. Of the 78 students, 53 (68%) indicated they favored extending the invitation, 20 (26%) were opposed, and 5 (6%) were undecided. A chi-square goodness-of-fit test showed the attitudes of students differed significantly from those of the faculty, χ^2 (2, N = 78) = 21.51, $p > .001$, $w = .53$.