

Chapter 12 – Comparing Counts

Count data are analyzed primarily for three different purposes: whether or not data agree with a specified distribution (a goodness-of-fit test), whether or not observed distributions collected at different times or places are consistent with one another (a test of homogeneity), and whether or not data classified according to two categorical variables indicate the categorical variables are related or not (a test of independence). All of these tests use a probability distribution called the χ^2 (chi-squared) distribution. The first test, goodness-of-fit, is not a built-in function on the TI-83, but can be done easily enough. This is a built-in function with the latest operating system version on the TI-84 and on the 89 as well. The other two tests are built into all calculators and the mechanics are exactly the same whether testing homogeneity or independence; what is different is the setting and conclusions which can be made.

The χ^2 statistic is defined to be $\sum \frac{(Obs - Exp)^2}{Exp}$ where the sum is taken over all the cells in a table. The quantities

$\frac{(Obs - Exp)}{\sqrt{Exp}}$ are the standardized residuals which are examined in the event the null hypothesis is rejected to determine which cells deviated most from what is expected.

TESTING GOODNESS OF FIT

Does your zodiac sign determine how successful you will be in later life? Fortune magazine collected the zodiac signs of 256 heads of the largest 400 companies. Here are the number of births for each sign.

Births	Sign
23	Aries
20	Taurus
18	Gemini
23	Cancer
20	Leo
19	Virgo
18	Libra
21	Scorpio
19	Sagittarius
22	Capricorn
24	Aquarius
29	Pisces

We can see some variation in the number of births per sign- Pisces had the most and Gemini the least, but is it enough to claim that successful people are more likely to be born under certain star signs than others? If there is no difference between the signs, each should have (roughly) 1/12 of the births. That's the null hypothesis in this situation: births are evenly distributed across the year. The alternate is that the null is wrong: births are not evenly distributed across the year. These executives were a convenience sample, but we have no reason to suspect there was bias in their selection. Further, executives should be independent of one another and each category has expected frequency $(1/12)*256 = 21.333 > 5$. We may proceed with our inference.

TI-83 Calculations

Here, the observed counts have been entered into L1 and the (hypothesized) probability for each birth sign (1/12) has been entered into L2. The calculator displays the decimal equivalent for the entered fraction.

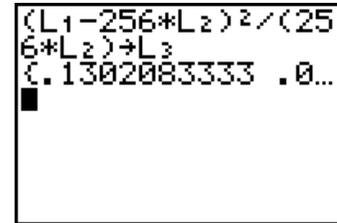
Note: Make sure both lists are the same length!

L1	L2	L3	Z
18	.08333		
21	.08333		
19	.08333		
22	.08333		
24	.08333		
29	.08333		
-----	*****		
L2(13) =			

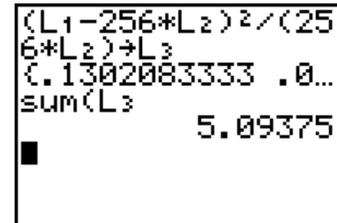
This next step calculates the quantities to be summed into the χ^2 statistic. The expected values for each cell are the proportions times the total number of individuals in the study (or 256 in this example). One needs to be careful in entering the command – the parentheses are necessary! The general form of the command is $(\text{Obslist} - n * \text{Problist})^2 / (n * \text{Problist})$ **STO** Newlist. Here the command was entered by pressing the following:

$((2 \text{nd} 1 - 2 5 6 \times 2 \text{nd} 2) x^2 \div ((2 5 6 \times 2 \text{nd} 2)) \text{STO} \blacktriangleright 2 \text{nd} 3$.

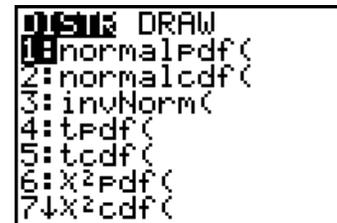
After pressing **ENTER** the first few entries are displayed.



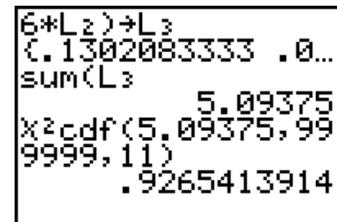
We need to add the entries in L3 to find the χ^2 value. Press **2nd****STAT** (List), arrow to **MATH** then press **5** to select option 5: sum(. The command shell is transferred to the home screen. Press **2nd****3** for L3 followed by **ENTER**. The χ^2 statistic is 5.09375.



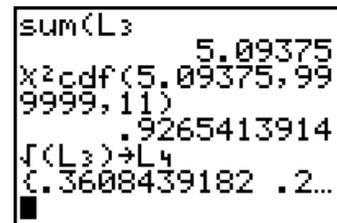
Now we need to get a p-value for the test. These tests are always one-tailed, so the p-value corresponds to area between 5.09375 and ∞ under a curve with $k - 1$ degrees of freedom, where k is the number of categories (here there are 12 categories, so 11 degrees of freedom). Press **2nd****VAR**(DISTR). We want choice 7: χ^2 cdf(. Either arrow to the selection and press **ENTER** or press **7**. The command shell is transferred to the home screen.



The parameters are low end (5.09375), high end (properly **1****2nd****.****9****9** for infinity, but practically, several 9's will work), then the degrees of freedom. Be sure to separate the parameters with a comma. The p-value for the test is 0.927 which is very large. We will not reject the null hypothesis and conclude that these data indicate no significant birth sign differences among the executives.



If the null were rejected, an examination of standardized residuals would show which cells were most different. The standardized residuals are just the square root of the entries in our list which had the components of χ^2 . To find them, from the home screen press **2nd****x²****2nd****3** **STO** **2nd****4** followed by **ENTER**. To further examine these, use the Statistics list editor. We see (not surprisingly) the birth sign most different from its expected value is Pisces, since it was the largest entry in the list.



TI-84/89 calculations

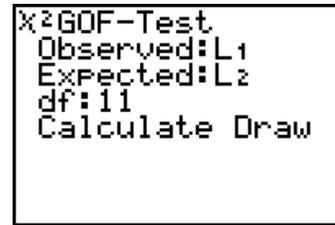
Using a TI-84 or -89 calculator, one can carry out the computations explicitly as indicated above for an 83, or use the built-in function. If you want to use the built-in function, the following section details its use. It is similar on both of these calculators.

The observed counts have been entered into L1 and the expected counts for each birth sign (1/12 of 256, or 256/12) has been entered into L2. The calculator displays the decimal equivalent for the entered fraction. If you are using a TI-89, enter the expected counts the same way.

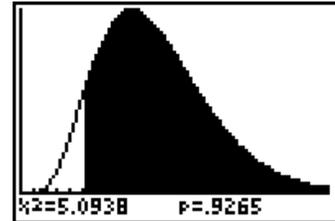
L1	L2	L3	Z
23	21.333	-----	
20	21.333		
18	21.333		
23	21.333		
20	21.333		
19	21.333		
18	21.333		

L2(1)=21.33333333...

To perform the test, select χ^2 GOF-Test from the STAT, TESTS menu. Specify the two lists, and the degrees of freedom for the test, which are $k - 1$ where k is the number of categories. Since we have 12 birth signs, the degrees of freedom are 11. Notice we have a choice of Calculate or Draw, just as with other hypothesis tests. Calculate gives more decimal places of output, and Draw highlights the area under the distribution curve which corresponds to the p-value of the test.



These are the results when Draw is selected. The calculated test statistic is $\chi^2 = 5.09$ and the p-value of the test is 0.9265. With these results, we fail to reject the null hypothesis and conclude, based on this data, the birth signs of executives are distributed evenly across the year – the observed differences were due to randomness.



If you are using a TI-89, the calculator automatically stores a new list of the components of the χ^2 statistic. The square roots of these values are the standardized residuals (after appropriate positive and negative signs are attached). Scanning down this list we see (yet again) the birth sign most different from its expected value is Pisces, the last entry in the list.

F1- Tools	F2- Plots	F3- List	F4- Calc	F5- Distr	F6- Tests	F7- Ints
list5	list6	zscor...	compl...			
						.13021
						.08333
						.52083
						.13021
						.08333
						.25521
complist[1]=.130208333333...						
MAIN RAD AUTO FUNC B/B						

TESTS OF HOMOGENEITY

Many colleges and universities survey graduating classes to determine their plans for the future. We might wonder whether plans depend on the student’s major through their undergraduate college in a university. Here is a summary table from several universities for 2006 graduates. Each cell of the table shows how many students from each college (the columns) had that particular type of plan (the rows).

	Agriculture	Arts & Sciences	Engineering	Social Sciences	Total
Employed	379	305	243	125	1052
Grad School	186	238	202	96	722
Other	104	123	37	58	322
Total	669	666	482	279	2096

Visually, choices do not appear to be the same (look at the row for employment) but is the difference real or is it due perhaps to different size classes between the colleges? Since we have really the same distribution (future plans) for different colleges, this is a test of homogeneity: the null hypothesis is that the distribution of students’ plans is the same across colleges, the alternate hypothesis is that the distributions are not the same.

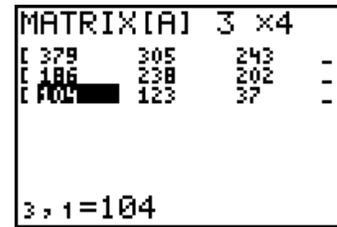
Entering Data into a matrix: TI-83/84

We will first enter the numbers in the body of the table into a matrix. If you have a TI-83, press **[MATRIX]**. If you have a TI-83+ or TI-84, press **[2nd][x⁻¹]**. Arrow to EDIT. Press **[ENTER]** to select matrix A.

First we need to give the size of the matrix. The body of our table had three rows and four columns, so the matrix is 3 x 4. (We don’t consider the column or row of totals as part of the data.) Press **[4][ENTER][3][ENTER]** to change the size of the matrix. Now type in the entries in the body of the table following each by **[ENTER]**. The process goes left to right, top to bottom.



Here is part of the filled-in matrix. Press $\text{2nd}[\text{MODE}](\text{QUIT})$ to leave the Matrix Editor. Notice the small dashes at the right of the screen. These indicate there is more data “hidden” off to the right. To see it, press D .



Entering Data into a matrix: TI-89

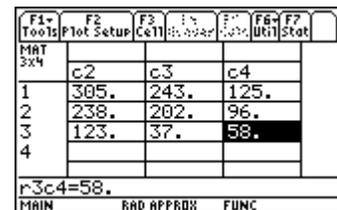
We will first enter the numbers in the body of the table into a matrix. From the home screen, press APPS Arrow to choice 6:Data/Matrix Editor and press the right arrow. Select choice 3:New. On the TI-89 Titanium, select the Data/Matrix Editor from the basic Applications screen.



We first define the type. Press the right arrow and select Matrix. Leave the folder set as main unless you really want to change it. Press the down arrow and give the matrix a name. Here, to keep myself straight, I’ve named the matrix mat1. The cursor in this field is set to alpha by default, so pressing any number keys will result in their letter equivalent. If a number is desired, press alpha to change out of alpha mode. The body of our table had three rows and four columns, so the matrix is 3 x 4. Press ENTER .

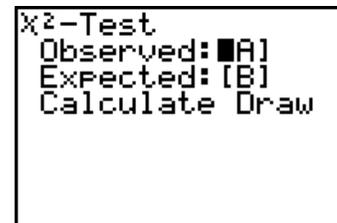


Cells are entered across rows. Press ENTER after each cell entry to proceed to the next. Here is part of the filled-in matrix. Press $\text{2nd}[\text{ESC}](\text{QUIT})$ to leave the Matrix Editor.

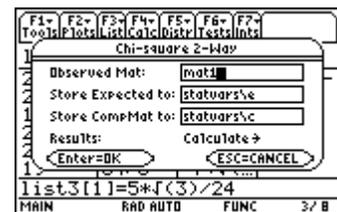


Performing the test

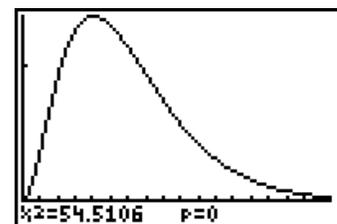
Now we’re ready to perform the test. On a TI-83/84, press STAT then arrow to TESTS. We want choice C: χ^2 -Test. Either press $\text{ALPHA}[\text{PRGM}](\text{C})$ or press the up arrow to find the option followed by ENTER . All we need to tell the calculator is where the observed counts were entered and where to store expected counts. These default to matrix A for Observed and matrix B for Expected. If you need to change them, press MATRX or $\text{2nd}[\text{x}^{-1}]$ then select the desired matrix. Notice we again have choices Calculate and Draw.



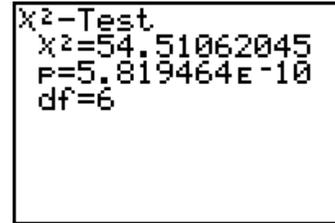
On an 89, return to the Statistics app and press $\text{2nd}[\text{F1}]$ for the Tests menu. Select choice 8:Chi2 2-way. To enter the name of the observed matrix, Go to the VAR-LINK screen ($\text{2nd}[\text{=}]$) and find the name you gave it. The calculator will automatically store expected counts and components of the χ^2 statistic in the named statistics matrices. Just as with the goodness-of-fit test, the (appropriately signed) square root of each component will be the standardized residual.



Selecting Draw gives the screen at right. We see three things here: the χ^2 statistic (54.5106) and the p-value to four decimal places (0). We also see a different sort of distribution curve. χ^2 curves are not symmetric; they are right skewed, and the shape varies with degrees of freedom (compare this curve to the one on page 71.)

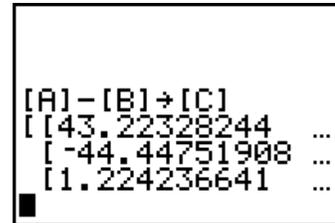


If we had chosen **Calculate** instead of **Draw** we would have this screen. It shows the same statistic value, but a little more exact p-value (but this is still essentially zero). We are also given the degrees of freedom for the test which are $(rows - 1)(cols - 1)$, so $(3-1)(4-1) = 6$.

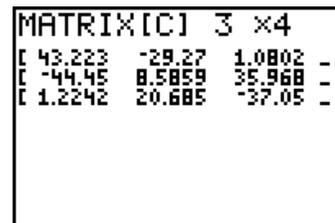


Since the p-value is so small, we reject the null hypothesis and conclude the distributions of post-graduation plans are not the same for all colleges.

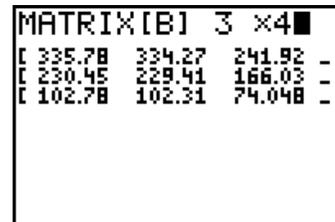
Where are the differences? We'd like to get the standardized residuals. Unfortunately, the TI-83/84 won't give these easily. We can, however, compute the matrix Obs-Exp. The Observed counts are in matrix A and the Expected counts are in matrix B. From the home screen, press **[MATRX]**, press **[ENTER]** to select matrix A, then **[]** press **[MATRX]** press the down arrow and select matrix B by pressing **[ENTER]** then press **[STO] [MATRX]** and select a new matrix (probably C). Lastly, another **[ENTER]** will perform the calculation and display some of the results.



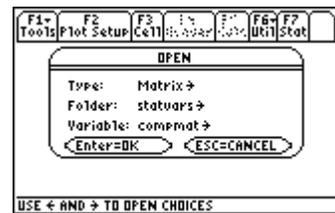
Return to the matrix editor to see the entire matrix. The signs here indicate whether or not the observed count was more than the expected (positive) or less than expected (negative). Clearly the largest entries (in absolute value) are those for agriculture graduates. Many more of these than expected plan to be employed, while far fewer than expected have plans for graduate school.



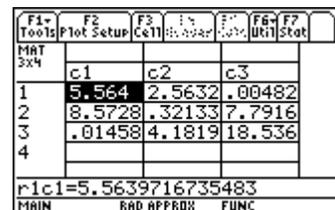
Here is matrix B which contains the expected counts. If we divide the entries in matrix C by the square root of the corresponding entry in matrix B we have the standardized residuals. The standardized residual agricultural graduates with employment plans is $43.223/\sqrt{335.78} = 2.360$. Similarly we find the standardized residual for agriculture graduates with graduate school plans is -2.928 .



If you are using a TI-89, to see the components matrix, return to the Data/Matrix editor application, and select **Open**. On the next screen, you will be asked what type of data you want to open, the folder it is in, and the variable name. Here, I have indicated that I want to see a **matrix** in the **statvars** folder, namely, **compmat**, the components matrix. Use the right arrow to expand each selection, and the down arrow to change a value. Press **[ENTER]** to actually make the selection and proceed to the next entry. Pressing **[ENTER]** when done will display the matrix.



The components matrix shows the largest contribution to the χ^2 statistic is in the third column, third row. This was the entry corresponding to engineering graduates with "other" plans (many fewer than expected). The next largest value was for agriculture graduates intending to pursue graduate school (again many fewer than expected). It seems that plans of university graduates are highly dependent on their field of study.



TESTING INDEPENDENCE

Tests of independence are used when the same individuals are classified according to two categorical variables. A study from the University of Texas Southwestern Medical Center examined whether the risk of Hepatitis C was affected by whether people had tattoos and by where they got their tattoos. The data from this study can be summarized in a two-way table as follows.

	Hepatitis C	No Hepatitis C
Tattoo, Parlor	17	35
Tattoo, Elsewhere	8	53
No Tattoo	22	491

Is the chance of having hepatitis C independent of (not related to) tattoo status? Our null hypothesis is that the two are not related. The alternate hypothesis is that there is a relationship.

Enter the numbers from the body of the table into a 3 x 2 matrix as discussed above. Press $\text{2nd}[\text{MODE}]$ (Quit) to exit the matrix editor.

```
MATRIX[A] 3 x2
[ 17   35 ]
[  8   53 ]
[ 22  491 ]

3, 2=491
```

The χ^2 test is performed just as it is for the test of homogeneity. Here are the results. The p-value is extremely small. We will reject the null hypothesis and conclude there is a relationship between hepatitis C status and tattoos.

```
 $\chi^2$ -Test
 $\chi^2=57.91217384$ 
 $P=2.657855E-13$ 
 $df=2$ 
```

Here are the results of computing Obs-Exp and storing the result into a new matrix as above. The largest standardized residual is $13.096/\sqrt{3.9042} = 6.628$. People with tattoos from tattoo parlors are more likely than normal to have hepatitis C. Could it be that tattoo parlors are a source of hepatitis C?

```
MATRIX[C] 3 x2
[ 13.096  -13.1 ]
[  3.4201  -3.42 ]
[ -16.52  16.516 ]
```

Here is the matrix of expected cell counts. Not all of them are more than 5. This means the conclusions based on the χ^2 test we just performed may not be valid. Since the largest standardized residual is for one of these cells, this is a problem. A common solution is to combine cells in some manner to overcome the problem. In this case, both rows for tattooed people could be combined into one.

```
MATRIX[B] 3 x2
[ 3.9042  48.096 ]
[ 4.5799  56.42 ]
[ 38.516  474.48 ]
```

Race and Traffic Stops

Are Black drivers stopped more than others? Data from traffic stops in Cincinnati reported the race of the driver and whether the stop resulted in a search of the vehicle.¹ Is race a factor in vehicle searches? Race was recorded as Black, White, or Other. Since people who were stopped were classified both by race and whether or not they were searched, this is a test of independence.

¹ Police Vehicle Stops in Cincinnati, John E. Eck, Lin Liu, Growette Bostaph, Oct. 1, 2003, available at www.cincinnati-oh.gov.

	Black	White	Other	Total
Not searched	787	594	27	1408
Searched	813	293	19	1125
Total	1600	887	46	2533

This matrix is 2 x 3 because we have two rows (searched or not) and three columns for race.

```
MATRIX[A] 2 x3
[ 787  594  27  ]
[ 813  293  19  ]

z, z=19
```

The χ^2 statistic is 73.25 with a p-value 1×10^{-16} (essentially 0.) We have extremely strong evidence that whether or not a car was searched depends on the race of the driver. How are the variables related?

```
 $\chi^2$ -Test
 $\chi^2=73.25312847$ 
P=1.239611E-16
df=2
```

It is clear from looking at the differences Obs – Exp that Black drivers are more likely to be searched than white drivers. These are not the residuals (divide the entries in the matrix by the square root of the expected counts) but the direction of the relationship is clear.

```
MATRIX[C] 2 x3
[ -102.4  100.95  1.4303 ]
[ 102.38  -100.9  -1.43  ]
```

WHAT CAN GO WRONG?

Expected cell counts less than 5.

Check the computed matrix of expected cell counts. If they are not all greater than 5 the analysis may be invalid.

Missing or misplaced parentheses.

When computing elements for the goodness-of-fit test the parentheses are crucial.

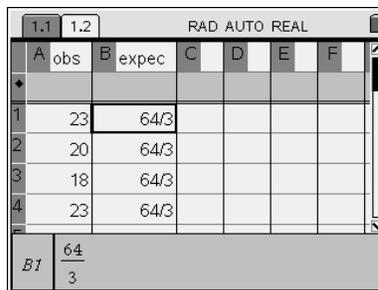
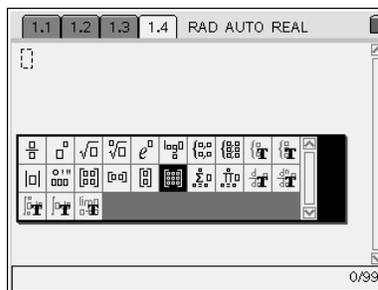
Overusing the test.

These tests are so easy to do and data from surveys and such are commonly analyzed this way. The problem that arises here is that in this situation the temptation is to check many questions to see if relationships exist; but performing many tests on *dependent* data (the answers came from the same individuals) such as this is dangerous. In addition, remember that, just by random sampling, when dealing at $\alpha = 5\%$ we'll expect to see something "significant" 5% of the time when it really isn't. This danger is magnified when using repeated tests — it's called the problem of multiple comparisons.

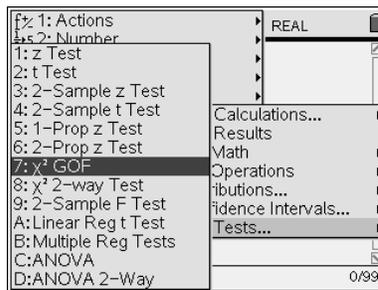
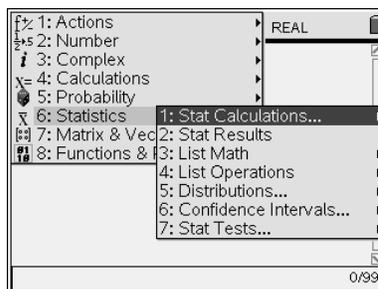
Commands for the TI-Nspire™ Handheld Calculator

For a chi-square goodness of fit test, first create lists of observed and expected values. The sign of zodiac example has expected counts of 256/12 or 64/3.

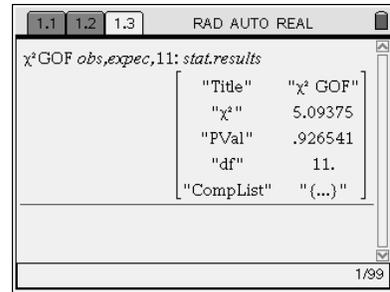
Open a Calculator page, press , and then select Statistics, Stat Tests, and χ^2 GOF.



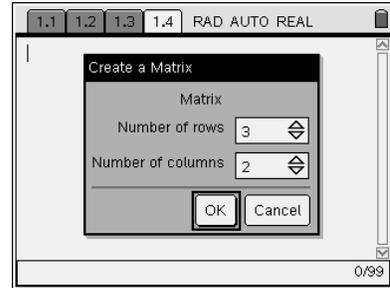
Select the list names and enter the degrees of freedom in the input box.



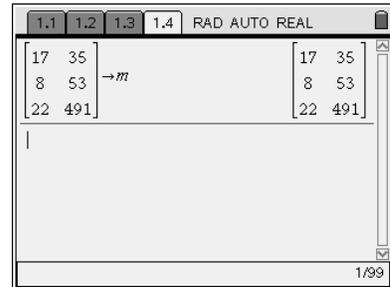
For chi-square tests with two-way tables, start on a calculator page. Press ctrl [x] to open templates. Highlight the matrix template as shown and [enter] .



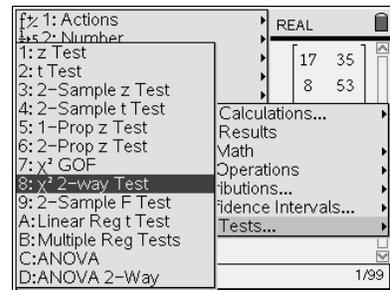
Type the dimensions of the matrix. For the tattoo example use 3 by 2.



Place the values in the matrix. Use [tab] to move to the next element. After all values are entered press \blacktriangleright to move the cursor outside the brackets. Press ctrl [var] for the store arrow, type a name for the matrix and [enter] .



Now press [menu] and select Statistics, Stat Tests, and χ^2 2-way Test.



In the input box select the name of the matrix.

