

Using SPSS for One Way Analysis of Variance

This tutorial will show you how to use SPSS version 12 to perform a one-way, between- subjects analysis of variance and related post-hoc tests.

This tutorial assumes that you have:

- Downloaded the standard class [data set](#) (click on the link and save the data file)
- Started SPSS (click on Start | Programs | SPSS for Windows | SPSS 12.0 for Windows)
- Loaded the standard data set

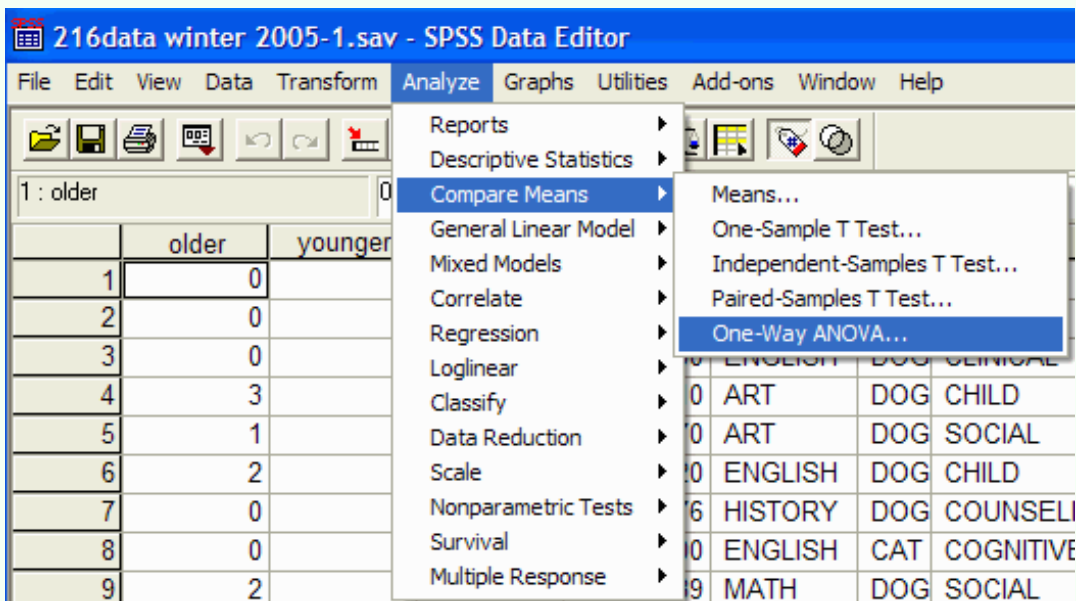
The one way analysis of variance (ANOVA) is an inferential statistical test that allows you to test if any of several means are different from each other. It assumes that the dependent variable has an interval or ratio scale, but it is often also used with ordinally scaled data.

In this example, we will test if the response to the question "If you could not be a psychology major, which of these majors would you choose? (Math, English, Visual Arts, or History)" influences the person's GPAs. We will follow the standard steps for performing hypothesis tests:

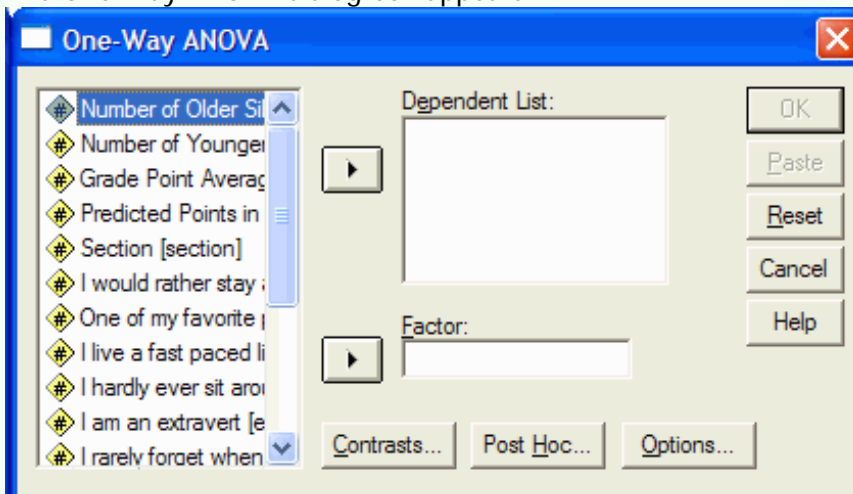
1. Write the null hypothesis:
 $H_0: \mu_{\text{Math}} = \mu_{\text{English}} = \mu_{\text{Visual Arts}} = \mu_{\text{History}}$
Where μ represents the mean GPA.
2. Write the alternative hypothesis:
 $H_1: \text{not } H_0$
(Remember that the alternative hypothesis must be mutually exclusive and *exhaustive* of the null hypothesis.)
3. Specify the α level: $\alpha = .05$
4. Determine the statistical test to perform: In this case, GPA is approximately ratio scaled, and we have multiple (4) groups, so the between-subjects ANOVA is appropriate.
5. Calculate the appropriate statistic:

SPSS assumes that the independent variable (technically a quasi-independent variable in this case) is represented numerically. In the sample data set, MAJOR is a string. So we must first convert MAJOR from a string variable to a numerical variable. See the tutorial on [transforming a variable](#) to learn how to do this. We need to automatically recode the MAJOR variable into a variable called MAJORNUM.

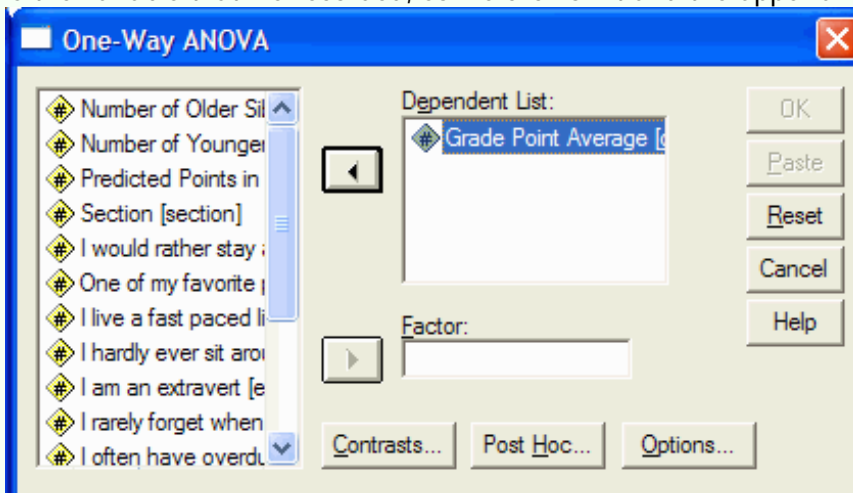
Once you have recoded the independent variable, you are ready to perform the ANOVA. Click on Analyze | Compare Means | One-Way ANOVA:



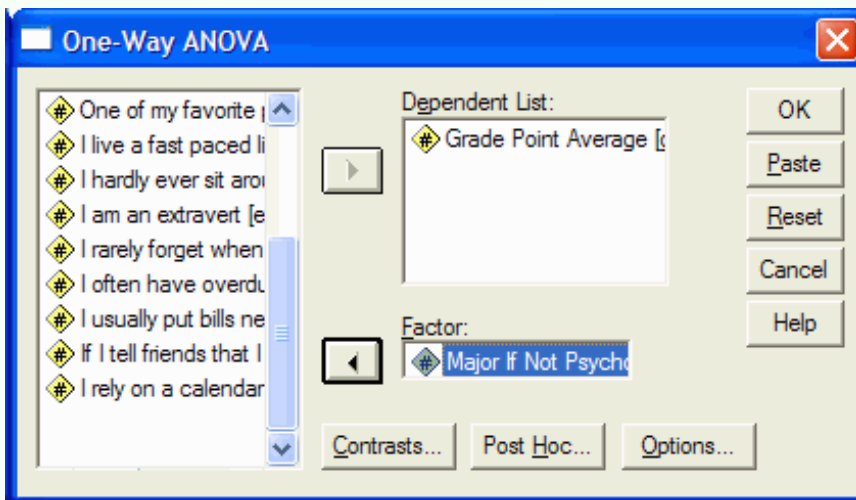
The One-Way ANOVA dialog box appears:



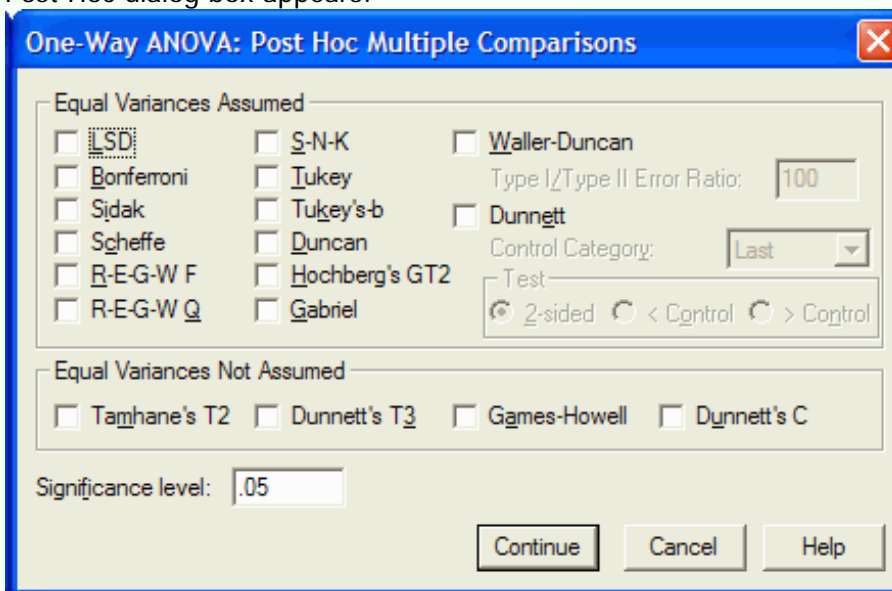
In the list at the left, click on the variable that corresponds to your dependent variable (the one that was measured.) Move it into the Dependent List by clicking on the upper arrow button. In this example, the GPA is the variable that we recorded, so we click on it and the upper arrow button:



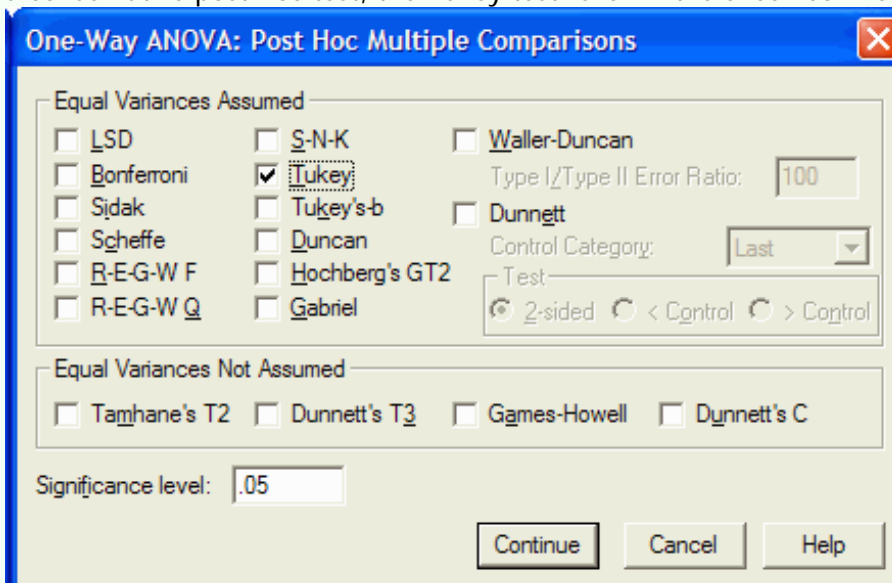
Now select the (quasi) independent variable from the list at the left and click on it. Move it into the Factor box by clicking on the lower arrow button. In this example, the quasi-independent variable is the recoded variable from above, MAJORNUM:



Click on the Post Hoc button to specify the type of multiple comparison that you would like to perform. The Post Hoc dialog box appears:

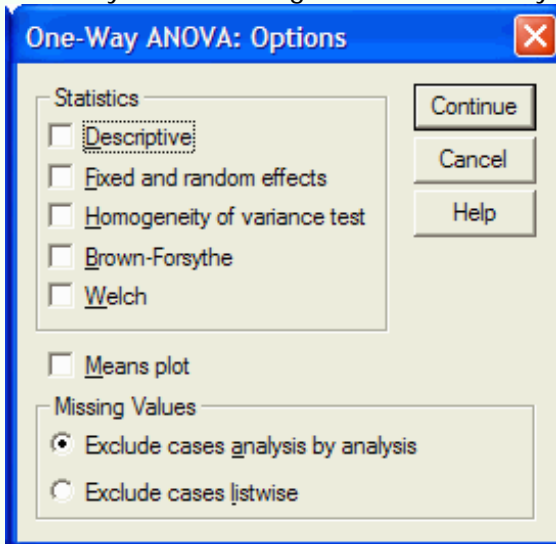


Consult your statistics text book to decide which post-hoc test is appropriate for you. In this example, I will use a conservative post-hoc test, the Tukey test. Click in the check box next to Tukey (not Tukey's-b):

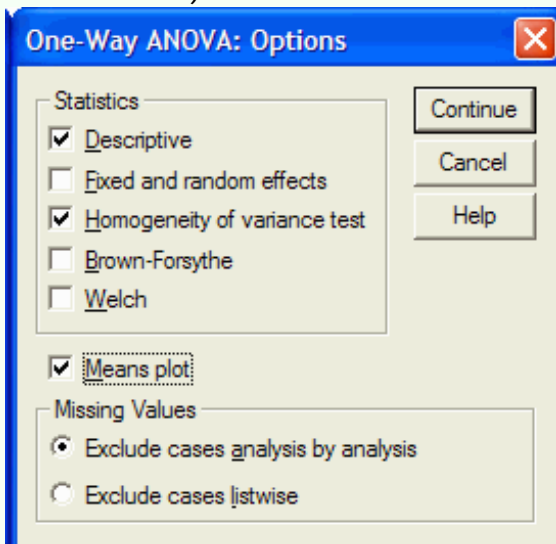


Click on the Continue Button to return to the One-Way ANOVA dialog box. Click on the Options button in the

One-Way ANOVA dialog box. The One-Way ANOVA Options dialog box appears:



Click in the check box to the left of Descriptives (to get descriptive statistics), Homogeneity of Variance (to get a test of the assumption of homogeneity of variance) and Means plot (to get a graph of the means of the conditions.):



Click on the Continue button to return to the One-Way ANOVA dialog box. In the One Way ANOVA dialog box, click on the OK button to perform the analysis of variance. The SPSS output window will appear. The output consists of six major sections. First, the descriptive section appears:

Descriptives

Grade Point Average

	N	Mean	Std. Deviation	Std. Error	95% Confidence Interval for Mean		Minimum	Maximum
					Lower Bound	Upper Bound		
ART	15	3.190	.4474	.1155	2.943	3.438	2.4	4.0
ENGLISH	16	2.937	.5781	.1445	2.629	3.245	2.0	4.0
HISTORY	7	2.918	.6669	.2521	2.301	3.534	2.1	3.9
MATH	7	3.144	.4962	.1876	2.685	3.602	2.7	3.9
Total	45	3.051	.5361	.0799	2.890	3.212	2.0	4.0

For each dependent variable (e.g. GPA), the descriptives output gives the sample size, mean, standard deviation, minimum, maximum, standard error, and confidence interval for each level of the (quasi) independent variable. In this example, there were 7 people who responded that they would be a math major if they could not be a psychology major, and their mean GPA was 3.144, with a standard deviation of 0.496.

There were 16 people who would be an English major if they could not be a psychology major, and their mean GPA was 2.937 with a standard deviation of 0.5788.

Test of Homogeneity of Variances

Grade Point Average

Levene Statistic	df1	df2	Sig.
1.001	3	41	.402

The Test of Homogeneity of Variances output tests $H_0: \sigma^2_{\text{Math}} = \sigma^2_{\text{English}} = \sigma^2_{\text{Art}} = \sigma^2_{\text{History}}$. This is an important assumption made by the analysis of variance. To interpret this output, look at the column labeled Sig. This is the p value. If the p value is less than or equal to your α level for this test, then you can reject the H_0 that the variances are equal. If the p value is greater than α level for this test, then we fail to reject H_0 which increases our confidence that the variances are equal and the homogeneity of variance assumption has been met. The p value is .402. Because the p value is greater than the α level, we fail to reject H_0 implying that there is little evidence that the variances are not equal and the homogeneity of variance assumption may be reasonably satisfied.

ANOVA

Grade Point Average

	Sum of Squares	df	Mean Square	F	Sig.
Between Groups	.683	3	.228	.781	.511
Within Groups	11.961	41	.292		
Total	12.645	44			

The ANOVA output gives us the analysis of variance summary table. There are six columns in the output:

Column	Description
Unlabeled (Source of variance)	The first column describes each row of the ANOVA summary table. It tells us that the first row corresponds to the between-groups estimate of variance (the estimate that measures the effect and error). The between-groups estimate of variance forms the numerator of the F ratio. The second row corresponds to the within-groups estimate of variance (the estimate of error). The within-groups estimate of variance forms the denominator of the F ratio. The final row describes the total variability in the data.
Sum of Squares	The Sum of squares column gives the sum of squares for each of the estimates of variance. The sum of squares corresponds to the numerator of the variance ratio.
df	The third column gives the degrees of freedom for each estimate of variance. The degrees of freedom for the between-groups estimate of variance is given by the number of levels of the IV - 1. In this example there are four levels of the quasi-IV, so there are $4 - 1 = 3$ degrees of freedom for the between-groups estimate of variance. The degrees of freedom for the within-groups estimate of variance is calculated by subtracting one from the number of people in each condition / category and summing across the conditions / categories. In this example, there are 2 people in the Math category, so that category has $7 - 1 = 6$ degrees of freedom. There are 16 people in the English category, so that category has $16 - 1 = 15$ degrees of freedom. For art, there are $15 - 1 = 14$ degrees of freedom. For history there are $7 - 1 = 6$ degrees of freedom. Summing the dfs together, we find there are $6 + 15 + 14 + 6 = 41$ degrees of freedom for the within-groups estimate of variance. The final row gives the total degrees of freedom which is given by the total number of scores - 1.

	There are 45 scores, so there are 44 total degrees of freedom.
Mean Square	The fourth column gives the estimates of variance (the mean squares.) Each mean square is calculated by dividing the sum of square by its degrees of freedom. $MS_{\text{Between-groups}} = SS_{\text{Between-groups}} / df_{\text{Between-groups}}$ $MS_{\text{Within-groups}} = SS_{\text{Within-groups}} / df_{\text{Within-groups}}$
F	The fifth column gives the F ratio. It is calculated by dividing mean square between-groups by mean square within-groups. $F = MS_{\text{Between-groups}} / MS_{\text{Within-groups}}$
Sig.	The final column gives the significance of the F ratio. This is the p value. If the p value is less than or equal your α level, then you can reject H_0 that all the means are equal. In this example, the p value is .511 which is greater than the α level, so we fail to reject H_0 . That is, there is insufficient evidence to claim that some of the means may be different from each other.

We would write the F ratio as: The one-way, between-subjects analysis of variance failed to reveal a reliable effect of other major on GPA, $F(3, 41) = 0.781$, $p = .511$, $MS_{\text{error}} = 0.292$, $\alpha = .05$.

The 3 is the between-groups degrees of freedom, 41 is the within-groups degrees of freedom, 0.781 is the F ratio from the F column, .511 is the value in the Sig. column (the p value), and 0.292 is the within-groups mean square estimate of variance.

6. Decide whether to reject H_0 : If the p value associated with the F ratio is less than or equal to the α level, then you can reject the null hypothesis that all the means are equal. In this case, the p value equals .511, which is greater than the α level (.05), so we fail to reject H_0 .

When the F ratio is statistically significant, we need to look at the multiple comparisons output. Even though our F ratio is not statistically significant, we will look at the multiple comparisons to see how they are interpreted.

Multiple Comparisons

Dependent Variable: Grade Point Average

Tukey HSD

(I) Major If Not Psychology	(J) Major If Not Psychology	Mean Difference (I-J)	Std. Error	Sig.	95% Confidence Interval	
					Lower Bound	Upper Bound
ART	ENGLISH	.2532	.1941	.565	-.267	.773
	HISTORY	.2729	.2472	.689	-.389	.935
	MATH	.0469	.2472	.998	-.615	.709
ENGLISH	ART	-.2532	.1941	.565	-.773	.267
	HISTORY	.0197	.2448	1.000	-.636	.675
	MATH	-.2063	.2448	.834	-.862	.449
HISTORY	ART	-.2729	.2472	.689	-.935	.389
	ENGLISH	-.0197	.2448	1.000	-.675	.636
	MATH	-.2260	.2887	.862	-.999	.547
MATH	ART	-.0469	.2472	.998	-.709	.615
	ENGLISH	.2063	.2448	.834	-.449	.862
	HISTORY	.2260	.2887	.862	-.547	.999

The Multiple Comparisons output gives the results of the Post-Hoc tests that you requested. In this example, I requested Tukey multiple comparisons, so the output reflects that choice. Different people have different

opinions about when to look at the multiple comparisons output. One of the leading opinions is that the multiple comparison output is only meaningful if the overall F ratio is statistically significant. In this example, it is not statistically significant, so technically I should not check the multiple comparisons output.

The output includes a separate row for each level of the independent variable. In this example, there are four rows corresponding to the four levels of the quasi-IV. Lets consider the first row, the one with major equal to art. There are three sub-rows within in this row. Each sub-row corresponds to one of the other levels of the quasi-IV. Thus, there are three comparisons described in this row:

Comparison	H ₀	H ₁
Art vs English	H ₀ : $\mu_{\text{Art}} = \mu_{\text{English}}$	H ₁ : $\mu_{\text{Art}} \neq \mu_{\text{English}}$
Art vs History	H ₀ : $\mu_{\text{Art}} = \mu_{\text{History}}$	H ₁ : $\mu_{\text{Art}} \neq \mu_{\text{History}}$
Art vs Math	H ₀ : $\mu_{\text{Art}} = \mu_{\text{Math}}$	H ₁ : $\mu_{\text{Art}} \neq \mu_{\text{Math}}$

The second column in the output gives the difference between the means. In this example, the difference between the GPA of the people who would be art majors and those who would be English majors is 0.2532. The third column gives the standard error of the mean. The fourth column is the p value for the multiple comparison. In this example, the p value for comparing the GPAs of people who would be art majors with those those who would be English majors is 0.565, meaning that it is unlikely that these means are different (as you would expect given that the difference (0.2532) is small.) If the p values is less than or equal to the α level, then you can reject the corresponding H₀. In this example, the p value is .565 which is larger than the α level of .05, so we fail to reject H₀ that the mean GPA of the people who would be art majors is different from the mean GPA of the people who would be English majors. The final two columns give you the 95% confidence interval.

Grade Point Average

Tukey HSD ^{a,b}

		Subset for alpha = .05
Major If Not Psychology	N	1
HISTORY	7	2.918
ENGLISH	16	2.937
MATH	7	3.144
ART	15	3.190
Sig.		.686

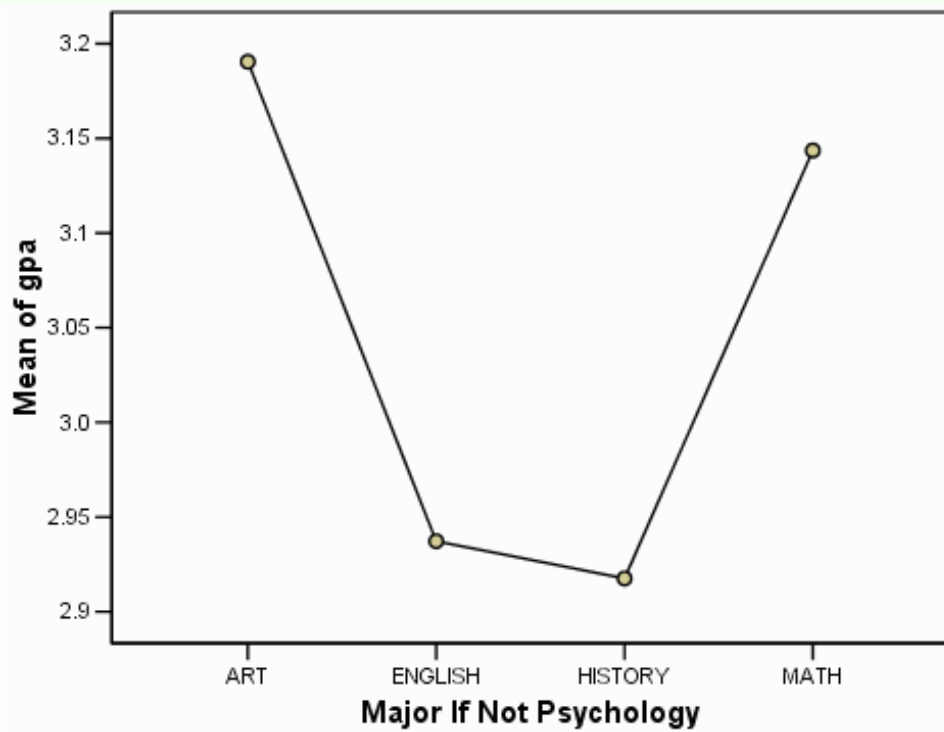
Means for groups in homogeneous subsets are displayed.

a. Uses Harmonic Mean Sample Size = 9.641.

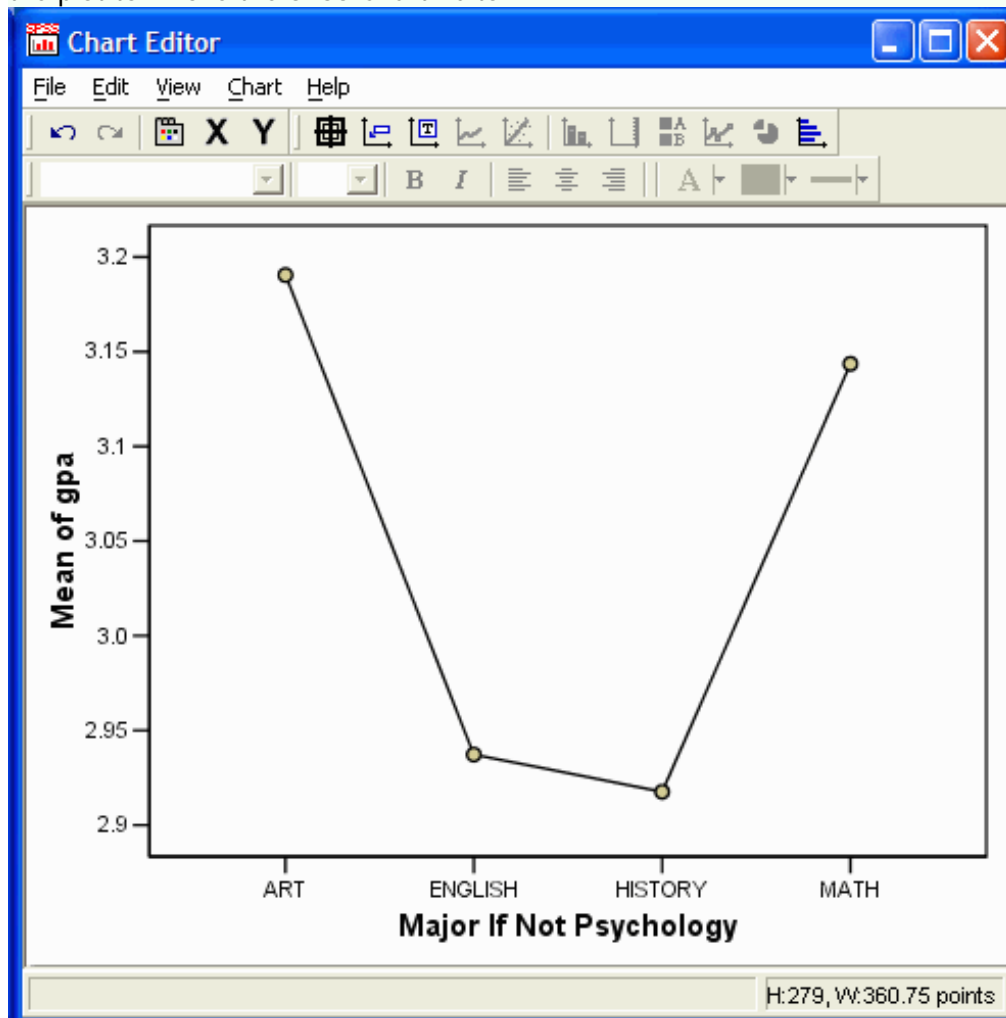
b. The group sizes are unequal. The harmonic mean of the group sizes is used. Type I error levels are not guaranteed.

The next part of the SPSS output (shown above) summarizes the results of the multiple comparisons procedure. Often there are several subset columns in this section of the output. The means listed in each subset column are not statistically reliably different from each other. In this example, all four means are listed in a single subset column, so none of the means are reliably different from any of the other means. That is not to say that the means are not different from each other, but only that we failed to observe a difference between any of the means. This is consistent with the fact that we failed to reject the null hypothesis of the ANOVA.

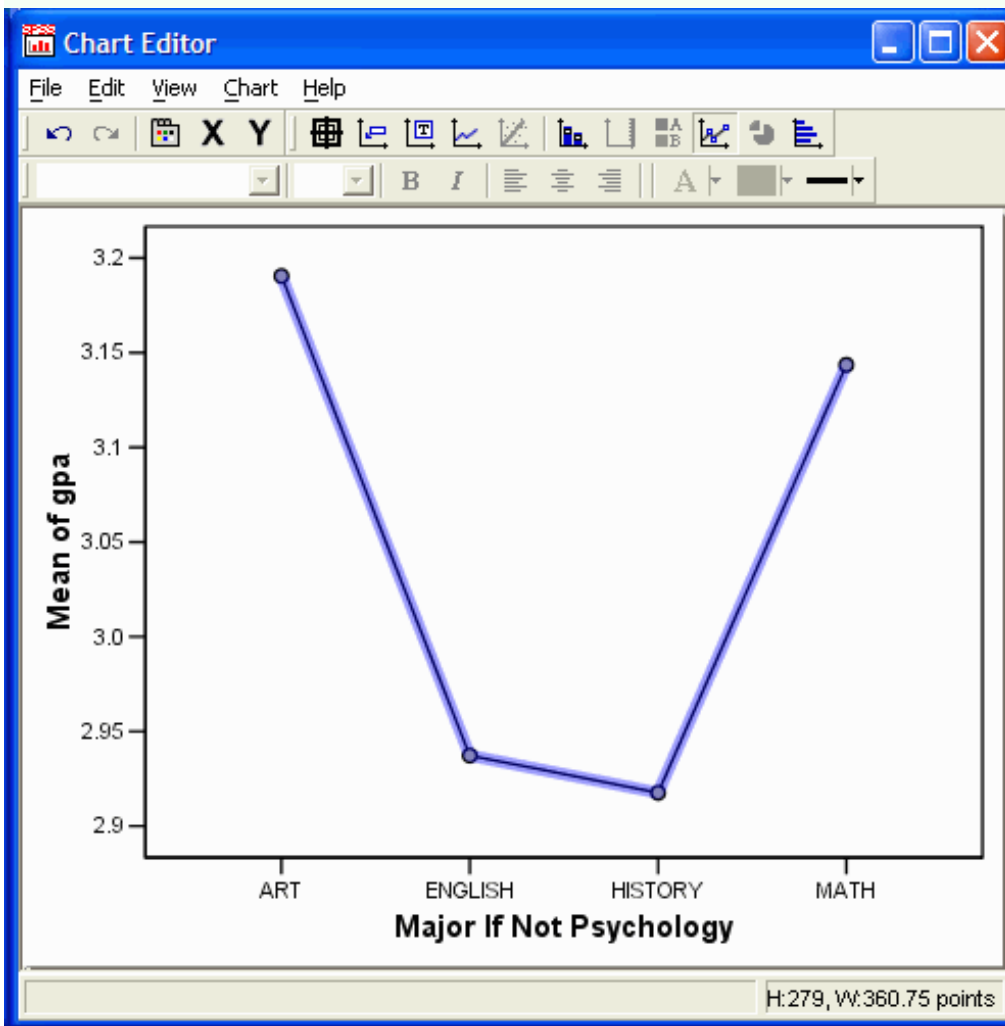
The final part of the SPSS output is a graph showing the dependent variable (GPA) on the Y axis and the (quasi) independent variable (other major) on the X axis:



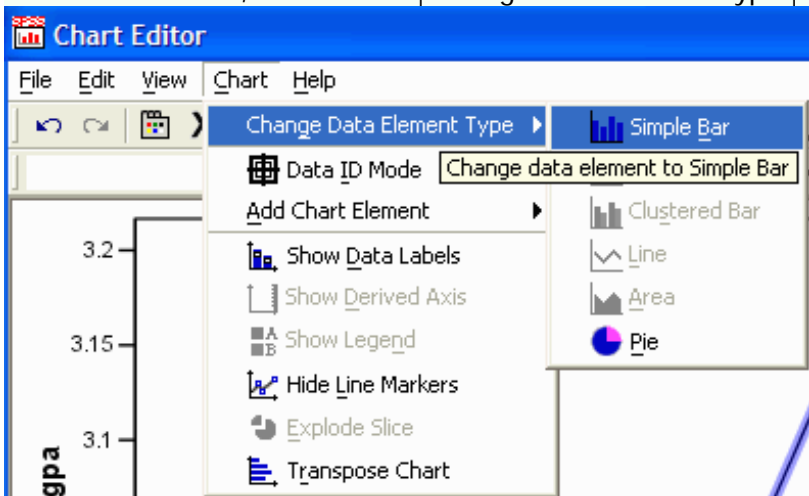
Because the quasi-independent variable is nominally scaled, the plot really should be a bar plot. Double click on the plot to invoke the SPSS Chart Editor:



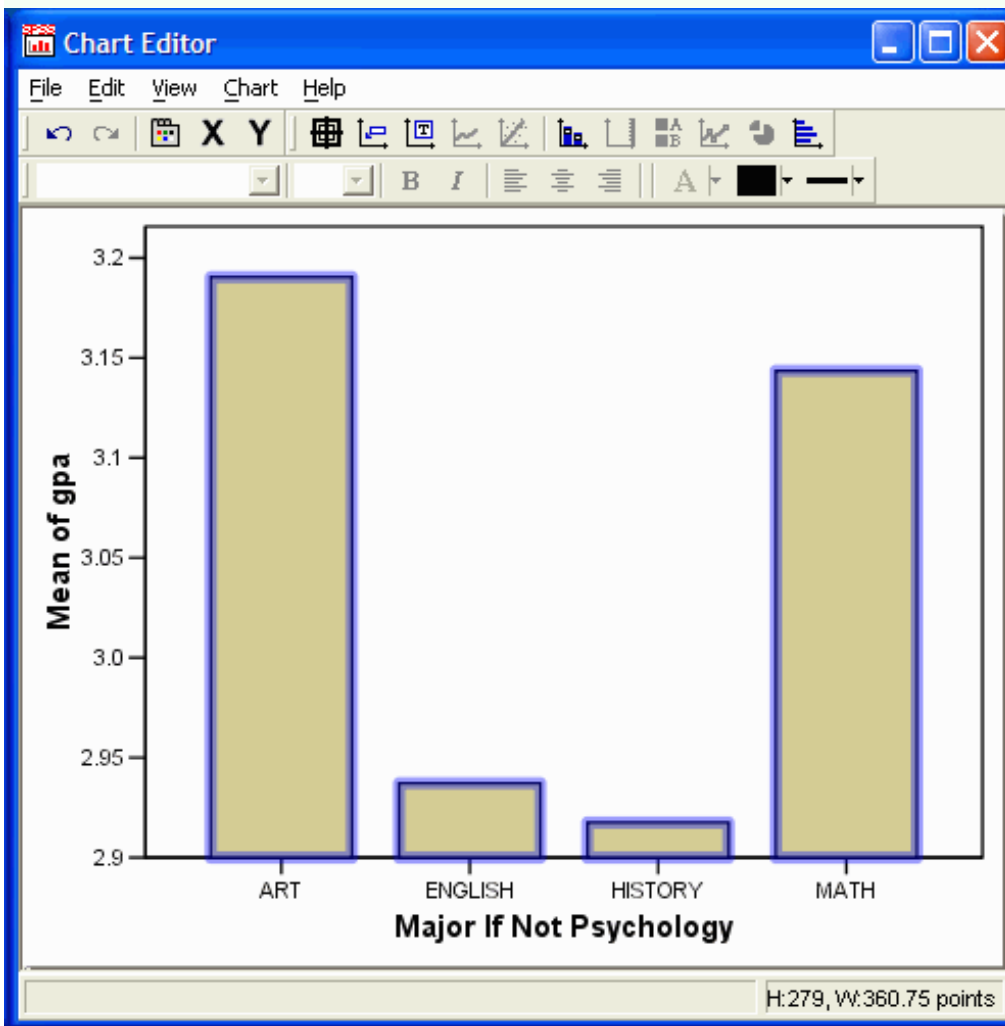
In the Chart Editor, click on one of the data points:



In the Chart Editor, select Chart | Change Data Element Type | Simple Bar:



The new bar graph appears in the editor:



Make any other changes to the bar graph that you want. (See the [tutorial on editing graphs](#) if you don't remember how to make changes.)

Close the Chart Editor by selecting File | Close in the chart editor.