

7 Applying Statistical Techniques

Meta-Stat provides a wealth of statistical tools to help you examine your data.

Overview	137
Common Functions	141
Selecting Variables to be Analyzed	141
Deselecting a Variable	143
Weighting Your Data	145
Selecting Data with a SELECT IF Statement	146
Dealing with Missing Data	149
Using Dummy Variables	150
Displaying Descriptive Statistics	153
Results Screen	155
Second Results Screen for Unbiased Effect Size	157
Performing a Regression Analysis	159
First Results Screen: Summary Information	160
Second Results Screen: Coefficients and Standard Errors	162
Fourth Results Screen: Residuals	165
Recording Predicted Values	166
Using Dummy Variables to Explore Differences Between Groups	167
Testing Mean Differences Between Groups with Randomization	169
Showing Correlations Between Variables	173
Cross-Tabulating Groups	176
Results Screen	177
Cross-Tabulating for Other Groups	178

(continued)

Calculating Group Means	179
Results Screen	180
Results Screen for All Grouping Variables	182
Calculating Frequencies	183
Displaying Means for a Two-Way Breakdown of Groups	185
Results Screen	187
Showing Breakdowns for Other Groups	187

Overview

The appropriate role for inferential statistics in meta-analysis is not merely unclear, it is seen quite differently by different methodologists.

In 1981, in the first extended discussion of the topic, Glass, McGaw and Smith raised doubts about the applicability of inferential statistics for the meta-analysis problem. Inference at the level of persons within studies (of the type addressed by Rosenthal 1984) seemed quite unnecessary to them, since even a modest size synthesis will involve a few hundred persons (nested within studies) and lead to nearly automatic rejection of null hypotheses. Moreover the chances are remote that these persons or subjects within studies were drawn from defined populations with anything approaching probabilistic techniques; hence, probabilistic calculations advanced as if subjects had been randomly selected are dubious.

At the level of "studies," the question of the appropriateness of inferential statistics can be asked again, and Glass et al. seem to answer in the negative. They pointed out that there are two instances in which common inferential methods are clearly appropriate: when a defined population has been randomly sampled and when subjects have been randomly assigned to conditions in a controlled experiment. In the latter case, Fisher showed how the permutation test can be used to make inferences to the universe of all possible permutations. But this case is of little interest to meta-analysts who never assign units to treatments. Glass et al. claimed that the typical meta-analysis virtually never meets the condition of probabilistic sampling of a population¹. They took the position that inferential statistics has little role to play in meta-analysis: "The probability conclusions of inferential statistics depend on something like probabilistic sampling, or else they make no sense" (p. 199).

It is a common habit of thought to acknowledge that many data sets fail to meet probabilistic sampling conditions, but to argue that one might well treat the data in hand "as if" it were a random sample of some hypothetical population. Under this supposition, inferential techniques are applied and the results inspected. This circumlocution has neither logic nor common sense to support it. Indeed, it seems

¹ In one instance known to Smith, Glass & Miller (1980), the available population of drug treatment experiments was so large that it was in fact randomly sampled for the meta-analysis.

to be little short of a rationalization for performing statistics that one has gone to the trouble to learn whether they are appropriate or not. If the sample is fixed and the population is allowed to be hypothetical, then surely the data analyst will imagine a population that resembles the sample of data. Hence all of these "hypothetical populations" will be reflections of the samples and there will be no need for inferential statistics. The researcher runs the risk of generalizing to what well may be a fictitious, and hence irrelevant, population.

Hedges and Olkin (1985) developed inferential techniques that ignored the pro forma testing of null hypotheses and focused on the estimation of regression functions that estimate effects at different levels of study characteristics. They worried about both sources of statistical instability: that arising from persons within studies and that which arises from variation between studies. As they properly pointed out, the study based on 500 persons deserves greater weight than the study based on 5 persons in determining the response of the treatment condition to changes in study conditions. The techniques they present are based on traditional assumptions of random sampling and independence. It is, of course, unclear precisely how the validity of their methods are compromised by failure to achieve probabilistic sampling of persons and studies.

Rubin (1990) addressed most of these issues squarely and staked out a position that appeals to the authors of this manual: "...consider the idea that sampling and representativeness of the studies in a meta-analysis are important. I will claim that this is nonsense--we don't have to worry about representing a population but rather about other far more important things" (p. 155). These more important things to Rubin are the estimation of treatment effects under a set of standard or ideal study conditions. This process, as he outlined it, involves the fitting of response surfaces (a form of quantitative model building) between study effects (Y) and study conditions (X, W, Z etc.).

Where theorists disagree, technicians are well advised to leave open as many options as possible. Consequently, we have provided many methods of data analysis in *Meta-Stat* that address inferential issues, including Hedges's homogeneity tests, Rosenthal and Rubin's aggregate significance levels as well as conventional methods of analysis

like regression, and a newly rediscovered approximate randomization test².

Where conditions of random sampling are met, the inferential techniques in Meta-Stat will have clear and undisputed application. In all other instances, it will be up to the meta-analyst to decide whether to apply them and how to interpret them. We provide fair warning that the inferential statistical techniques in Meta-Stat are easy to misapply and misuse.

² We do not provide ANOVA as we could not identify any algorithms that are appropriate for use with effect-sizes. For Plain Statistics, you can use dummy coding with the REGRESSION routine to perform ANOVA, ANCOVA, and t-tests. If your criterion variable is UNBIASED, the group MEANS routine will provide Hedges's homogeneity tests -- tests which are directly analogous to ANOVA. If your criterion variable is UNBIASED and you select WEIGHT as a weighting variable, the REGRESSION routine will give you Hedges's Q_E and Q_B and the corrected standard errors for the regression weights.

Common Functions

Some functions are commonly available in the *Meta-Stat* statistical analyses. The following sections describe these functions.

You can also export data to analyzed with SPSS. See Chapter 4 for more information.

Selecting Variables to be Analyzed

After you select the statistical analysis you want to perform, *Meta-Stat* displays a selection window. This window contains fields where you identify the variables for the analysis.

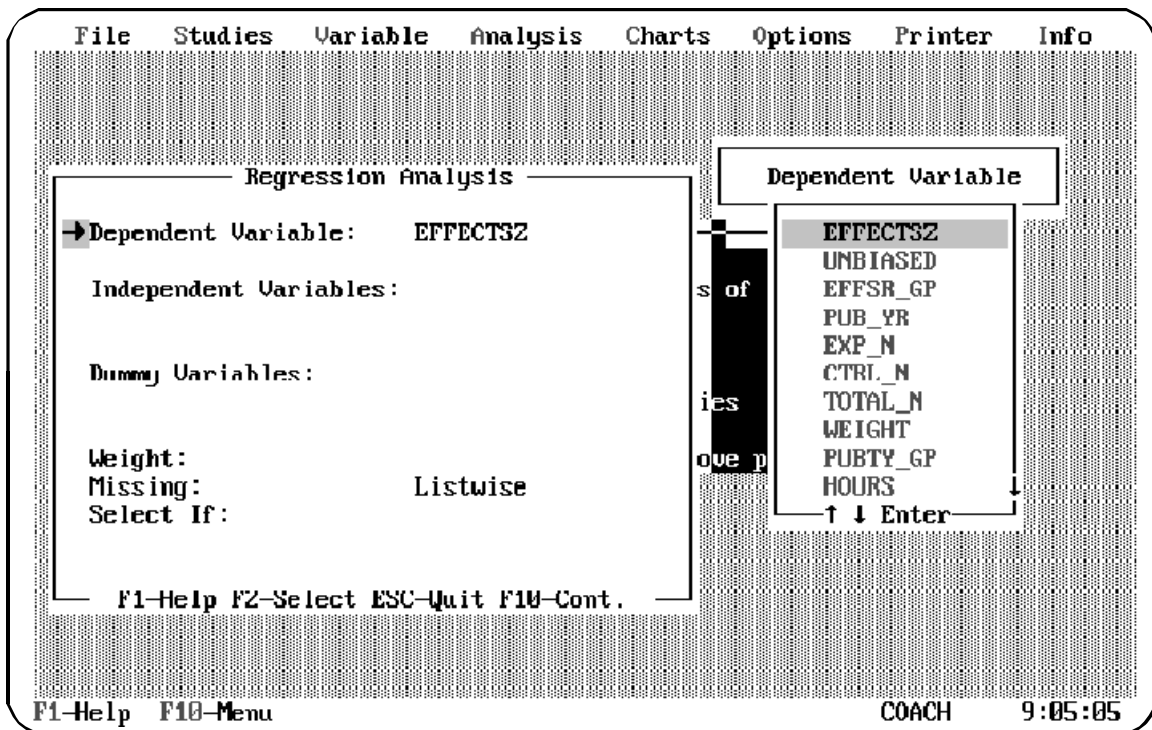
Selecting a Single Variable

For many analyses, you can select only a single variable in particular fields:

1. Highlight the desired field and press Enter.

Meta-Stat displays a window that lists the variables that are appropriate for that field.

In the following example from the Regression analysis, **Criterion variable** was highlighted when Enter was pressed. In the window, *Meta-Stat* displays a list of variables:



2. To select a variable, highlight the variable in the list by moving the mouse or using **9**. Then press Enter.

Meta-Stat returns you to the selection screen and fills in the field with the name of the variable you selected.

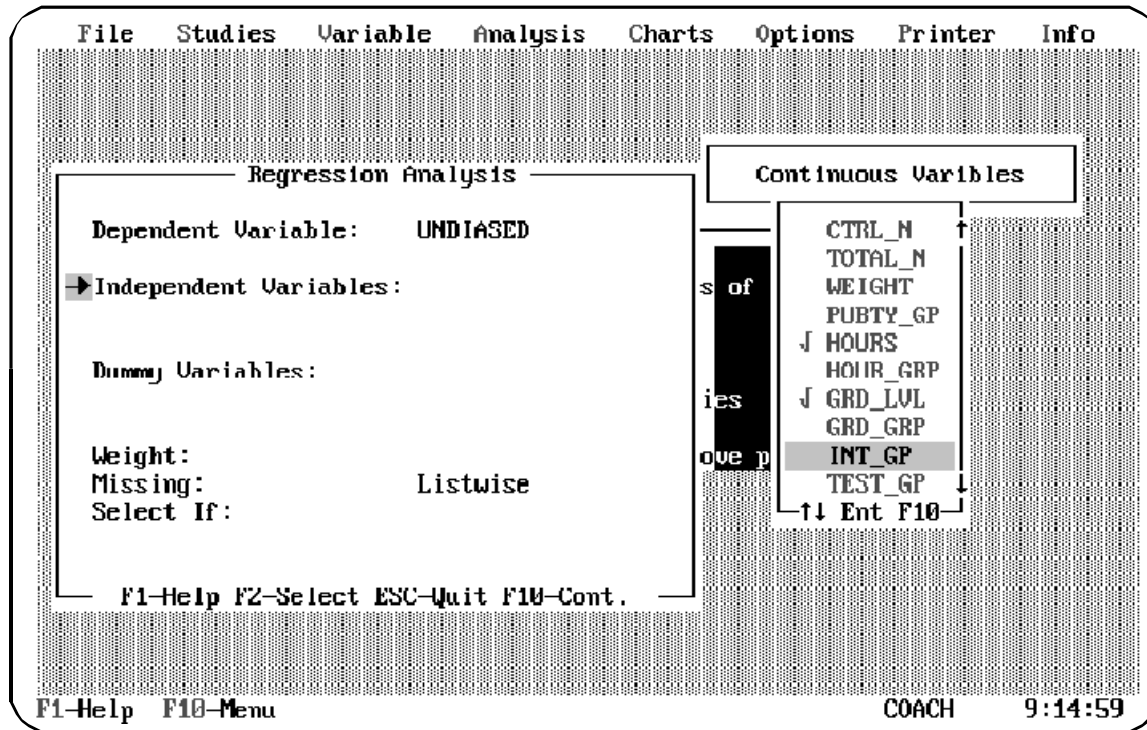
Selecting Multiple Variables

For some fields, you can select more than one variable. For example, you can identify more than one predictor variable for the Regression analysis. Follow these steps:

1. In the list of variables, move the pointer to the first variable you want to use and press Enter.

Meta-Stat marks the variable with a check-mark.

- Repeat step 1 until you have marked all the fields you want to use:

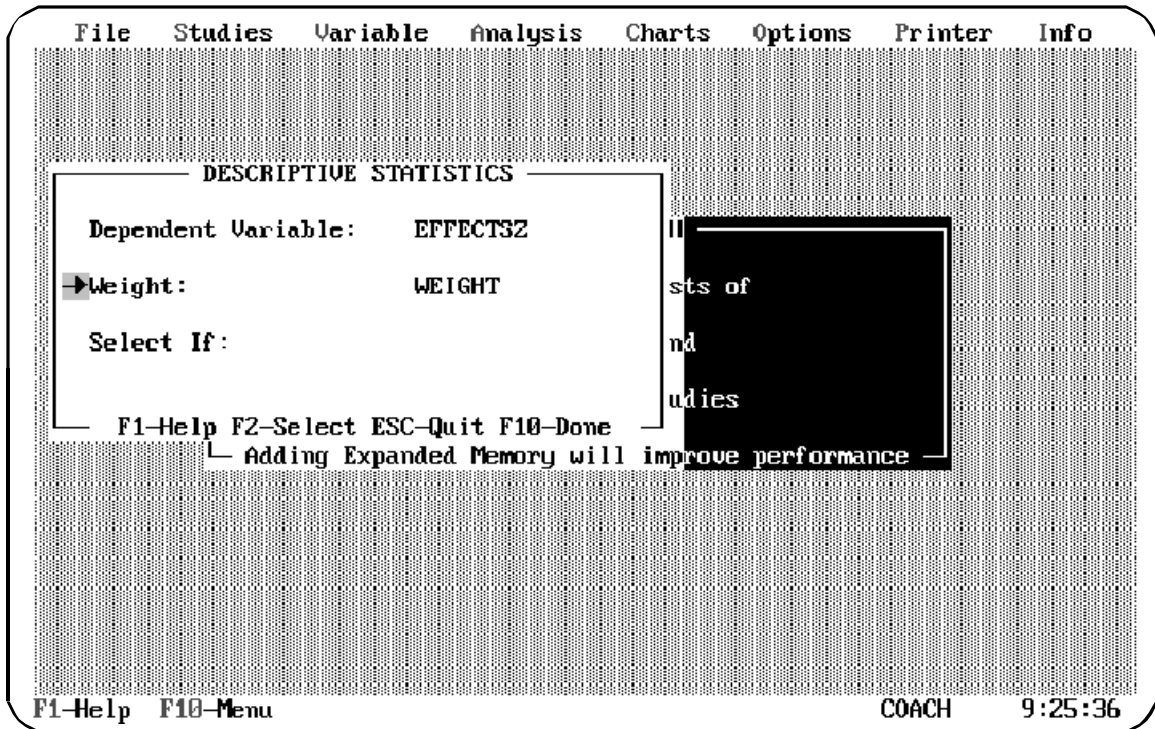


- Press **F10** when you have finished. *Meta-Stat* returns you to the selection screen and fills in the field with the variable names.

Deselecting a Variable

Meta-Stat remembers the last variable(s) you selected. To deselect one or more variables, call up the list of variables and press **Esc**.

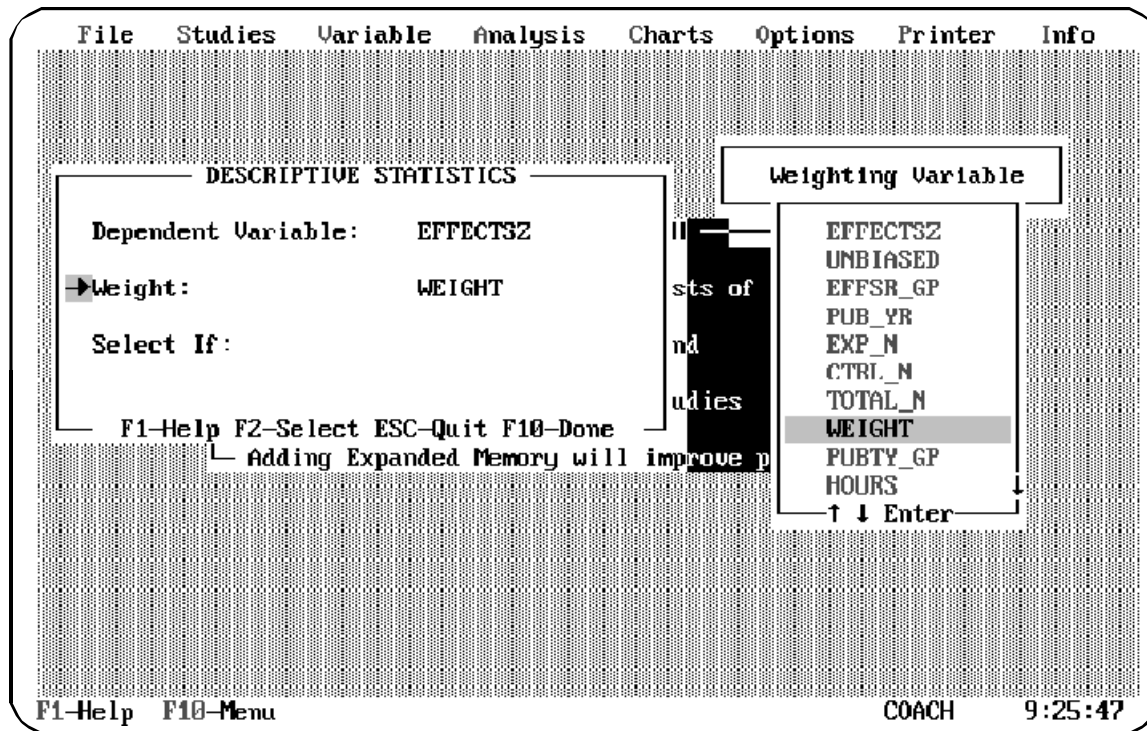
In the following example, the variable WEIGHT is currently selected as the weighting variable:



To deselect this variable:

1. Highlight the Weight field and press Enter

Meta-Stat displays a list of variables. The variable that is currently selected is highlighted.

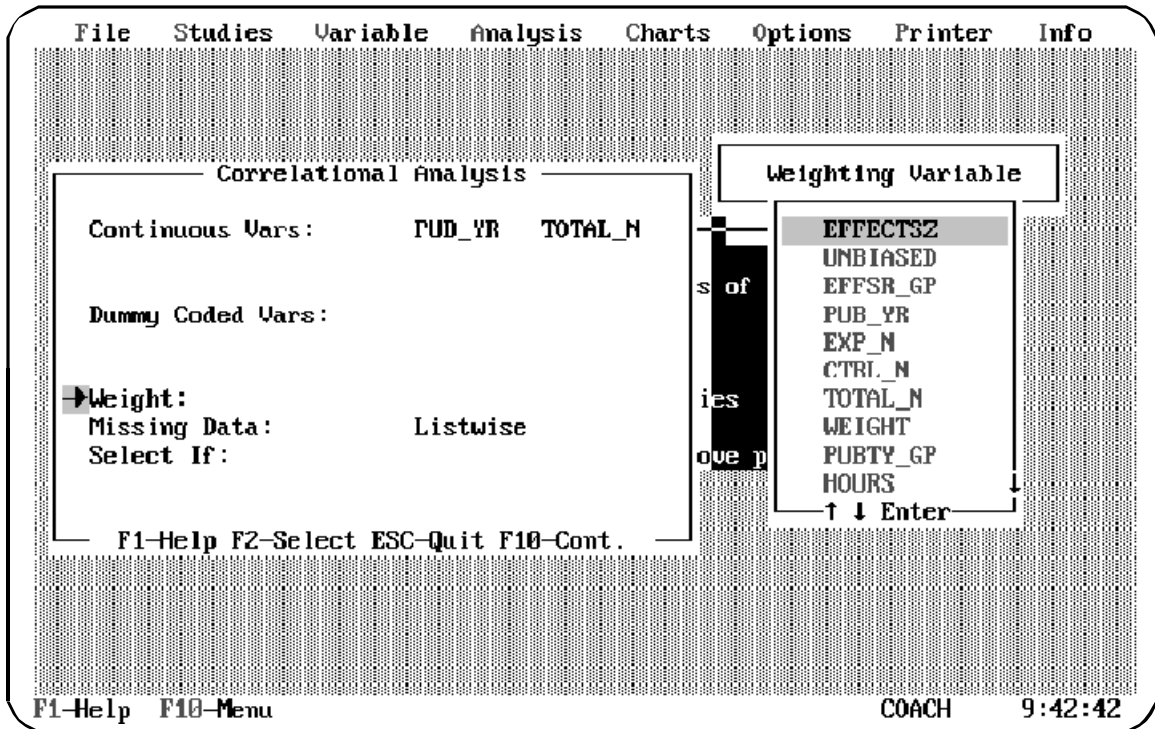


2. Press **Esc** to deselect the variable.

Weighting Your Data

You can weight your data using any variable. *Meta-Stat* includes a predefined variable called WEIGHT that you can use for this purpose. WEIGHT is an equation variable that equals the inverse of the variance. See Chapter 5 for more information about this variable.

To select the weighting variable, move the pointer to the **Weight** field and press **Enter**. Then select the variable you want to use:



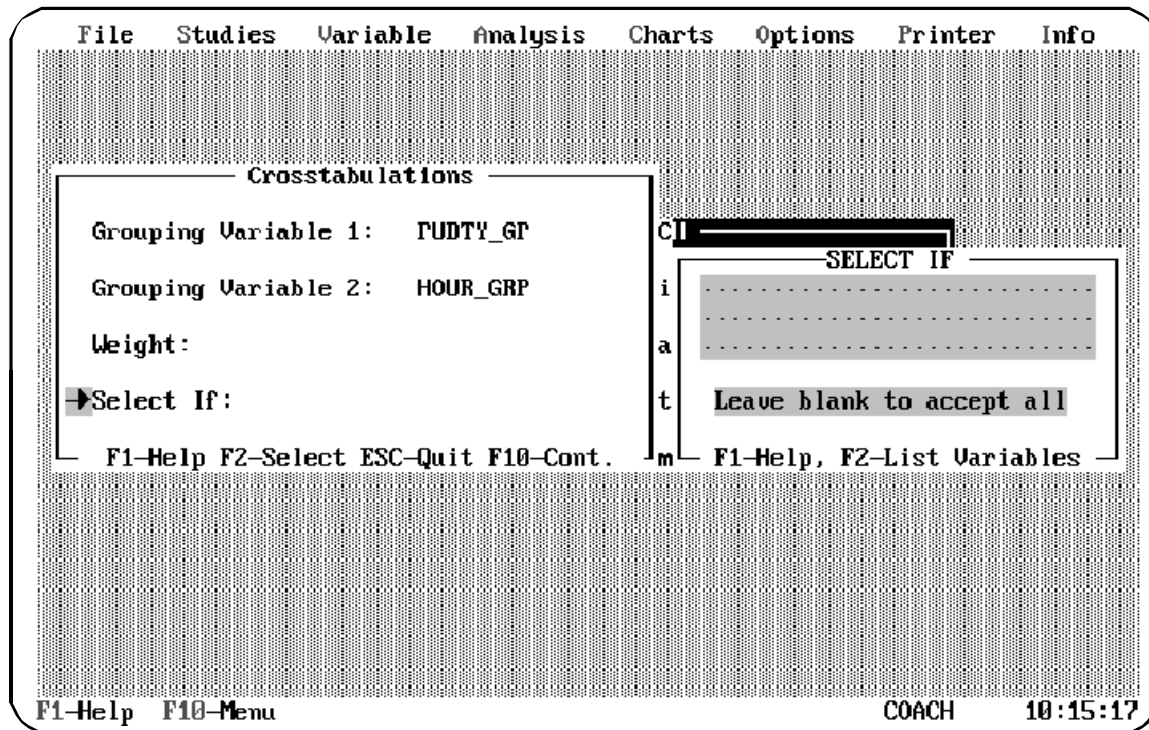
Selecting Data with a SELECT IF Statement

You can use a SELECT IF statement to limit the analysis to data meeting certain criteria. *Meta-Stat* will select for the analysis only the studies for which the IF statement is true.

For example, suppose you use a variable QUALITY to measure the quality of studies; a low-quality study is given a QUALITY value of 1. When you perform an analysis, you can use SELECT IF to trim the low-quality studies from the analysis by selecting only studies for which QUALITY is greater than 1.

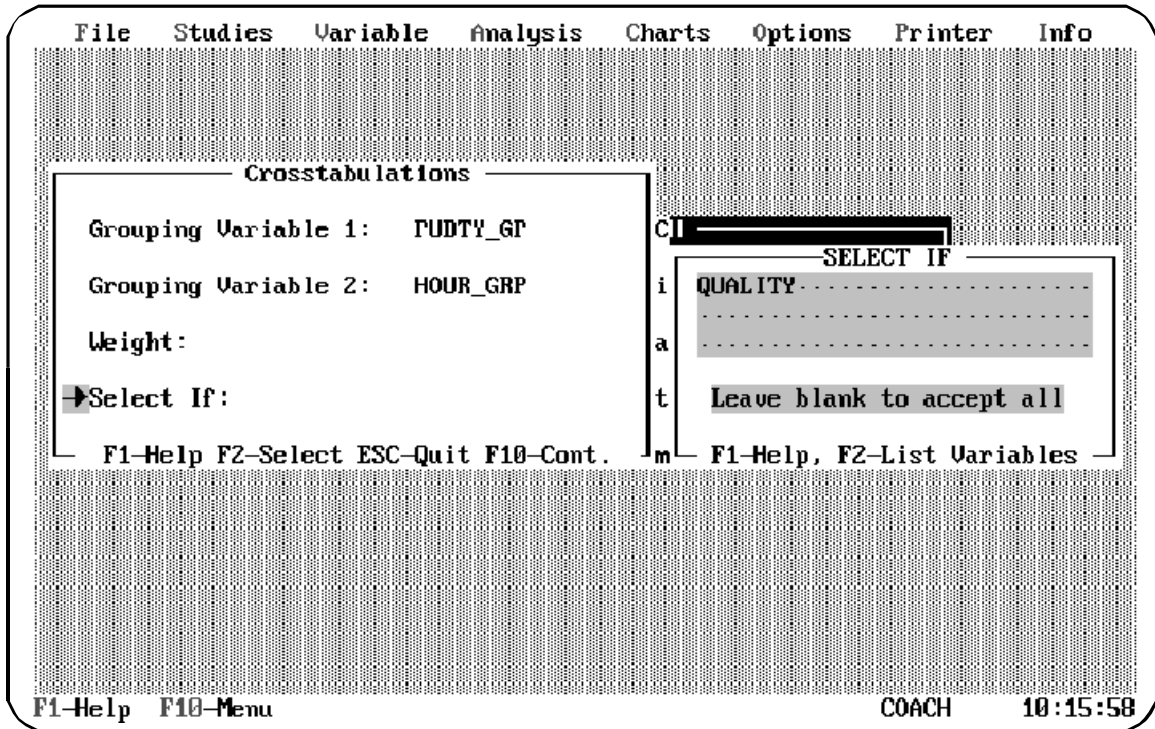
1. Move the pointer to Select If and press Enter.

Meta-Stat displays a window where you can enter selection criteria.



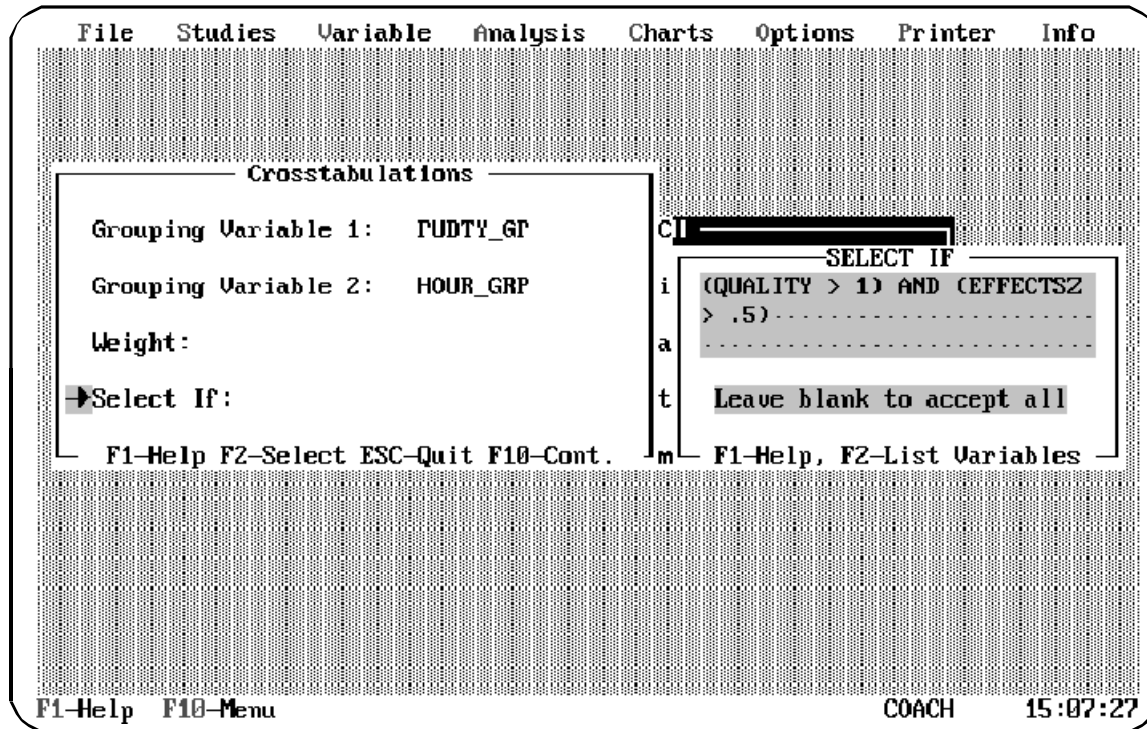
2. Press **F2** to display a list of variables.
3. In the list of variables, highlight the variable you want to use as part of the SELECT IF statement. Then press Enter.

Meta-Stat inserts the variable name in the selection criteria window:



4. Complete the IF statement by using parentheses, logical operators, and/or comparison operators. (Press **F1** to get more information about constructing an IF statement.)

In the following example, studies will be selected if the value of QUALITY is greater than 1 and if the value of EFFECTSZ is greater than 0.5:



Dealing with Missing Data

For the Regression and Correlation analyses, you can use the Missing Data field to tell *Meta-Stat* how to handle studies for which data are missing.

You can use two methods to handle missing data:

Method	Description
Listwise deletion	If a study does not have data for one of the variables, <i>Meta-Stat</i> ignores all other variables for that study. This is the default method.

Method	Description
Pairwise deletion	<p>If a study does not have data for one of the variables, <i>Meta-Stat</i> still uses data for other variables in the study.</p> <p>Caution: Be careful when performing regressions without complete data. It is generally valid to do this only when you can assume that the data are missing randomly throughout the studies.</p>

To identify the method you want to use, move the cursor to the Missing Data field and press Enter. Then select the method you want to use.

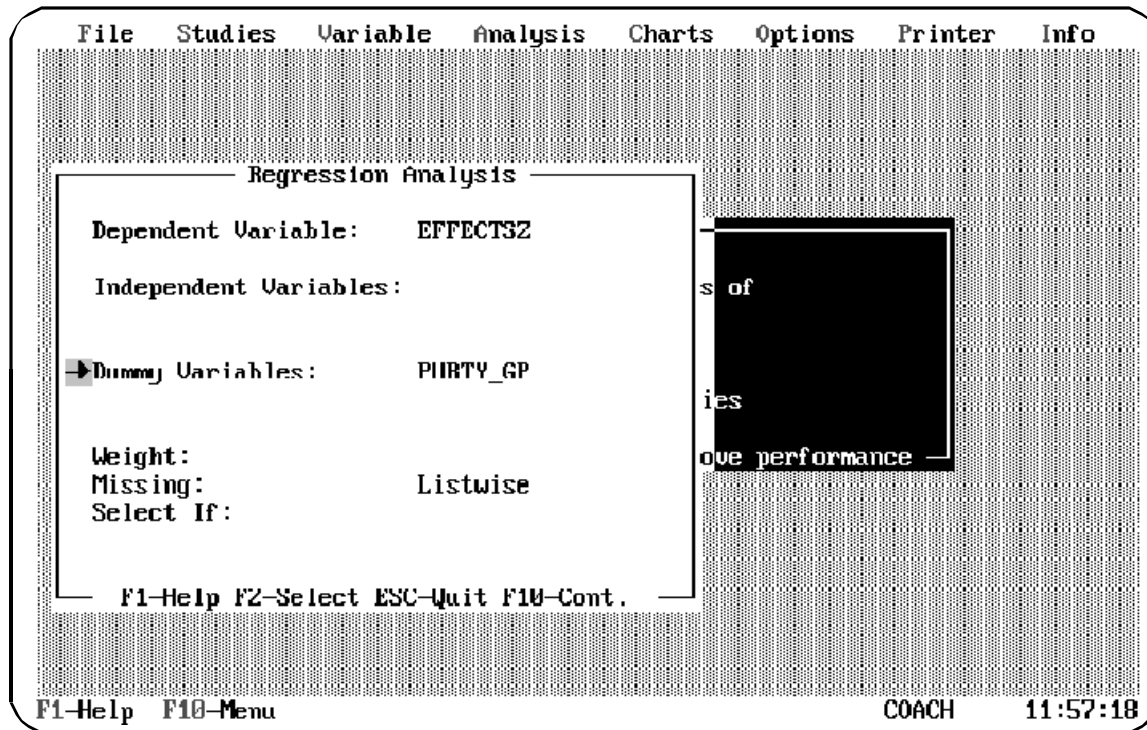
Using Dummy Variables

The Regression and Correlation analyses allow you to treat discrete and blocking data as dummy variables. For example, consider the discrete variable PUBTY_GP, which comes precoded with *Meta-Stat*. When you enter data for a study, you use PUBTY_GP to identify where the study was published—in a Journal, Dissertation, Paper, Book, or Other place.

During a Regression or Correlation analysis, *Meta-Stat* can treat these five categories as separate variables by assigning scores of 1 and 0 to them. For example, a study that was published in a Journal receives a score of 1 for the dummy variable Journal; it receives scores of 0 for the dummy variables Dissertation, Paper, Book, and Other.

Dummy variables take their name because they are created only as a convenience, so that they can be inserted into equations that are necessary for the analysis.

Use the field **Dummy Variables** to identify the discrete or blocking data from which you want to create dummy variables. The following example uses PUBTY_GP:



The following example shows a portion of Regression results for which dummy variables were created from PUBTY_GP:

META-STAT		REGRESSION					06/25/93
Variable		Beta	B	St Err	t	p	
PUBY_GP	journal	0.000	0.000	0.230	0.000	—	
PUBY_GP	dissert	0.204	0.346	0.217	1.597	0.120	
PUBY_GP	paper	0.020	0.047	0.217	0.216	0.025	
PUBY_GP	other	0.615	1.664	0.454	3.663	0.002	
Constant			0.536	0.124	4.334	0.000	

Ctrl-P Print F10 Menu PgUp, PgDn MoCOACH 11:57:53

Note: To avoid the multicollinearity problem in computing regression weight standard errors the first group is dropped. Usually, this is the group with missing data.

Displaying Descriptive Statistics

When examining your meta-analysis data, one of your first tasks will be to explore the distribution of many of your variables.

The Descriptive Statistics analysis provides a basic description of a variable. The mean shows you the central tendency for a variable, while the standard deviation and variance describe the dispersion of data around the mean. These same statistics are calculated for the Unbiased Effect Size variable.

Follow these steps to perform the analysis:

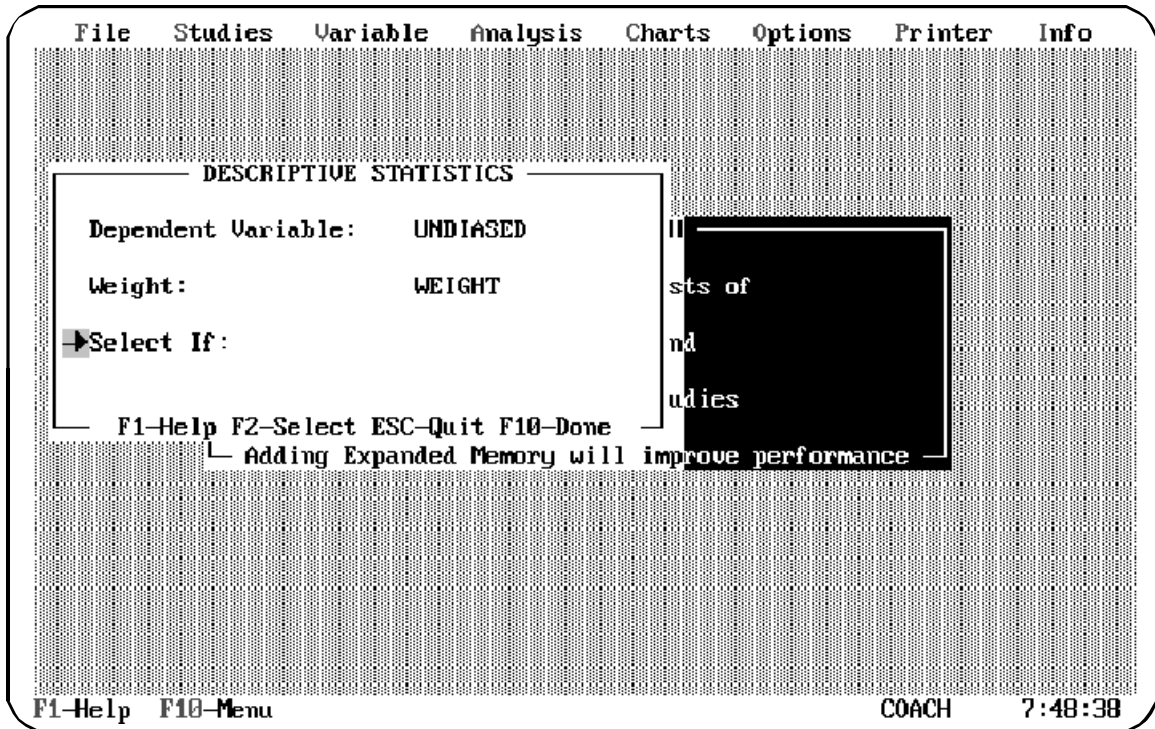
1. Select Analysis/Descriptive

Meta-Stat displays a screen where you can identify the variable you want to analyze.

2. Use Criterion variable to identify the variable for which you want to display the statistics.


3. If necessary, identify a weighting variable and construct a SELECT IF statement. For more information, see "Common Functions" earlier in this chapter.

In the following example, *Meta-Stat* will display descriptive statistics for the Unbiased Effect Size variable (UNBIASED), using the WEIGHT variable for weighting:



4. Press F10 to display the results.

Results Screen



This results screen:

- 1 Identifies the variables you selected for the analysis and the SELECT IF statement you used to select data. The criterion variable is the criterion variable you selected.
- 2 Lists the number of studies for which data were available for the criterion dependent) variable.
- 3 Provides weighted and unweighted statistics for the criterion variable.

Value	Description
Mean	Mean value, equal to the sum of the values divided by the number of cases (studies). The mean is the most common measure of central tendency.

Value	Description
Median	Median value, which is either the value of the case that lies exactly on the 50th percentile, once the cases have been ordered by rank from highest to lowest, or the value that is half way between the two inner most cases.
Variance	The dispersion of values around the mean, roughly equal to the average squared deviation from the mean. The smaller the variance, the more homogeneous the data. This is the variance for the presented dataset, i.e., N is used as the denominator .
Stand. Dev.	Standard deviation, equal to the positive square root of the variance ³
Min Value	Minimum value of the variable
Max Value	Maximum value of the variable
Range	The range between the minimum and maximum values

- 4 If you used the UNBIASED Effect Size as the criterion variable, *Meta-Stat* shows the results of Hedges's Homogeneity Test. This tests the hypothesis that the Effect Sizes for the studies are from different populations—that is, that they are not homogeneous. If the statistic is large (or equivalently if the value of p is low, say less than .05), Hedges's test indicates that there is a low probability that the Effect Sizes are from the same population. You can conclude that the effect sizes are homogeneous, i.e. the model of a single effect size is consistent with the data.

³ To convert from the variance of a population to the unbiased population estimate based on a sample, multiply by $n/(n-1)$.

Second Results Screen for Unbiased Effect Size

If you used the Unbiased Effect Size as the criterion variable, *Meta-Stat* shows a second screen of descriptive statistics. You can view this screen by pressing PgDn on the first screen.

META-STAT		DESCRIPTIVE STATISTICS		09/10/93
Criterion Variable:	UMBIASED	Unbiased Effect Size		1
Weighting Variable:	WEIGHT	Inverse of the variance		
Select If:				
Number of Valid Observations:	29			
Number of Missing Observations:	1			
		unweighted	weighted	
2	Mean ES	0.591	0.891	
	Fisher's Z	1.497 (p=0.13)	6.861 (p=0.00)	
	Equiv ES	0.579	0.912	
	Equiv R	0.278	0.415	
	Equiv R SQ	0.077	0.172	
	Failsafe N	0	326	
	Var Ratio	3.750	56.482	
Ctrl-P Print F10 Menu ↑↓ next variable PgUp basic				22:27:17

- 1 Some information is repeated from the first results screen.
- 2 The table presents the best estimate of the Effect Size for the population of studies.

Value	Description
Mean ES	Mean Effect Size

Value	Description
Fisher Z	A transformation that normalizes the Effect Size and provides a statistic that allows one to calculate the probability (p) of results under a null hypothesis of that the population Effect Size is zero.
Equiv ES	The equivalent Effect Size for the Fisher Z. This is the best estimate of the true Effect Size for this population of studies.
Equiv R	A correlation that is equivalent to Equiv ES.
Equiv R SQ	The amount of variance accounted for in the Effect size.
Failsafe N	The number of studies with contradictory results that would have to be collected in order to reduce the probability of the combined findings to statistical insignificance at $p=.05$ (Rosenthal, 1990).
Variance Ratio	The ratio of the error variance to the variance of the effect sizes. This provides a homogeneity measure. Schmidt and Hunter (1990, pages 281-338) state that you can conclude that the data are homogeneous, i.e., all the variance is due to forms of error, if the ratio is .75 or larger.

Performing a Regression Analysis

Use a regression analysis to investigate the relationship between predictor variables and a dependent or criterion variable, which is typically the Effect Size.

For example, in a meta-analysis of the effects of coaching on SAT scores, we might be interested in predicting the effects of coaching from predictor variables such as the number of hours of coaching and the grade level of the students. The regression analysis provides an equation that identifies how the values for hours and grade level must be weighted and summed to provide the best prediction of the coaching effects. The analysis also shows the accuracy of the equation and indicates how much of the variation in coaching effects is caused by hours and grade level.

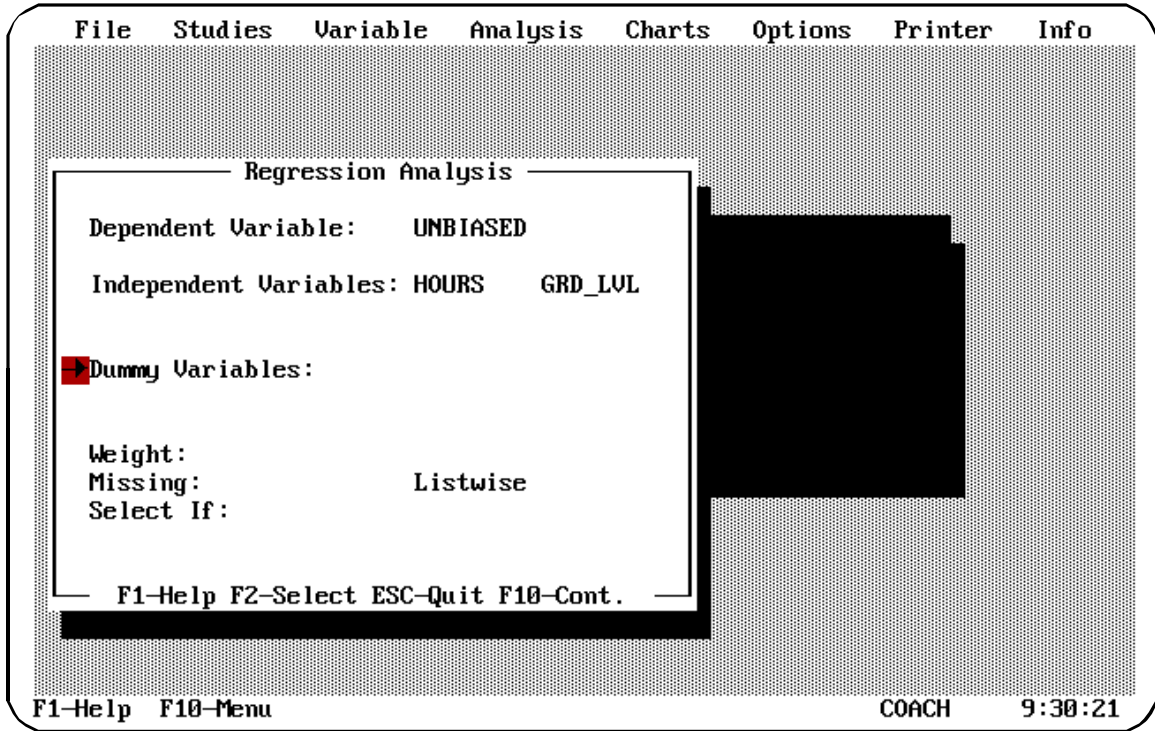
To perform a regression analysis, follow these steps:

1. Select Analysis/Regression

Meta-Stat displays a screen where you can identify the dependent and predictor variables.

2. Use Dummy Variables to treat each category within a blocking or discrete variable as a dummy variable. For more information, see the section "Using Dummy Variables" earlier in this chapter.
3. If necessary, identify a weighting variable, a method to handle missing data, and construct a SELECT IF statement. For more information, see "Common Functions" earlier in this chapter.

In the following example, a regression analysis will be performed to predict Unbiased Effect Size based on Hours and Grade Level:



4. Press F10 to display the first screen of the analysis.

First Results Screen: Summary Information

```

META-STAT                                09/10/93
----- REGRESSION -----
Criterion Variable:      UNBIASED Unbiased Effect Size  ①
Predictor Variables:    HOURS   GRD_LVL
Listwise deletion
Solution took 2 iterations

② Multiple R           0.8791      R Square           0.7727
Adjusted R-Square      0.7538      Standard Error      0.3176

          DF      Sum of squares      Mean square      F      p
③ Regression      2      8.2259      4.1150      40.803  0.0000
Residual          24      2.4204      0.1008
                                                    ④

Ctrl-P Print  F10 Menu                        PgDn More                        22:29:21
    
```

This screen:

- ① describes the variables you selected for the analysis and the method that *Meta-Stat* used to handle missing data.
- ② describes the strength of the relationship between the predictor variable(s) and the criterion variable.

Value	Description
Multiple R	Correlation between predictors and the criterion variable. The magnitude R is an index of the strength of the relationship.

R Square	The square of R, indicating the amount of variance accounted for by the predictor variables. In the example above, 77 percent of the variation in Unbiased Effect Size is explained by the predictor variables.
Adjusted R-Square	The adjusted R Square value.
Standard Error	The standard deviation of the differences between the actual and predicted values of the criterion variable (i.e., the standard deviation of the residuals).

- 3 The summary table describes actual variances (mean squares) of the predicted values and the residual values.
- 4 The p value is the probability that the observed degree of relationship, or one stronger, could have been obtained by sampling error from a population in which the multiple R is 0.

Second Results Screen: Coefficients and Standard Errors

To view the second screen of regression results, press PgDn on the first screen.

META-STAT 09/10/93

REGRESSION

Variable	Beta	B	St Err	t	p
① HOURS	0.168	0.191	0.170	1.125	0.271
GRD_LUL	0.745	0.095	0.019	4.985	0.000
Constant		-0.388	0.156	-2.483	0.019

②

Ctrl-P Print F10 Menu PgUp, PgDn More 22:29:42

Shows the following information for each prediction equation:

Value	Description
Constant	The predicted value of the criterion variable when the predictor variables are 0. In the example above, the predicted value for Unbiased Effect Size is -0.388 when the HOURS and GRD_LVL variables are 0.
Beta	<p>The beta weight, or the standardized regression coefficient:</p> <p>This value is calculated after the dependent and predictor variables have been standardized to have unit variance (that is, a standard deviation of 1 for all variables). The standardized coefficients are useful when the predictor variables are measured in different units, such as a duration in hours and a study date in years. Using these coefficients allows one to compare the relative influence of each predictor variable (see Darlington, 1968). In the example above, the predicted Unbiased Effect Size increases by 0.168 standard deviation units for every increase of one standard deviation unit in the HOURS variable. The predicted Unbiased Effect Size increases by 0.745 standard deviation units for every increase of one standard deviation unit in the GRD_LVL variable.</p>
B	<p>The nonstandardized regression coefficient.</p> <p>This value indicates the expected change in the criterion variable with a change of one unit in the associated predictor. In the example above, the predicted Unbiased Effect Size increases by 0.191 for every unit increase in the HOURS variable, while the predicted Unbiased Effect Size increases by 0.095 for every unit increase in the GRD_LVL variable.</p>
St Err	The standard error of each non-standardized regression coefficient. This value indicates the variability that would arise in B from repeated sampling and estimation of B with samples that are the same size as this one.

- 2 The p value shows the probability of obtaining a beta this size or more discrepant from zero in sampling from a population in which the relevant beta is zero.

Third Results Screen: Correlations

To view the third screen of regression results, press PgDn on the second screen.

Variable	R	Mean	St Dev	N
UNBIASED	1.0000	0.5940	0.6400	27
HOURS	0.7331	1.2926	0.5629	27
GRD_LUL	0.8722	7.7407	5.0199	27

1

META-STAT REGRESSION 09/10/93

Ctrl-P Print F10 Menu PgUp, PgDn More 22:30:17

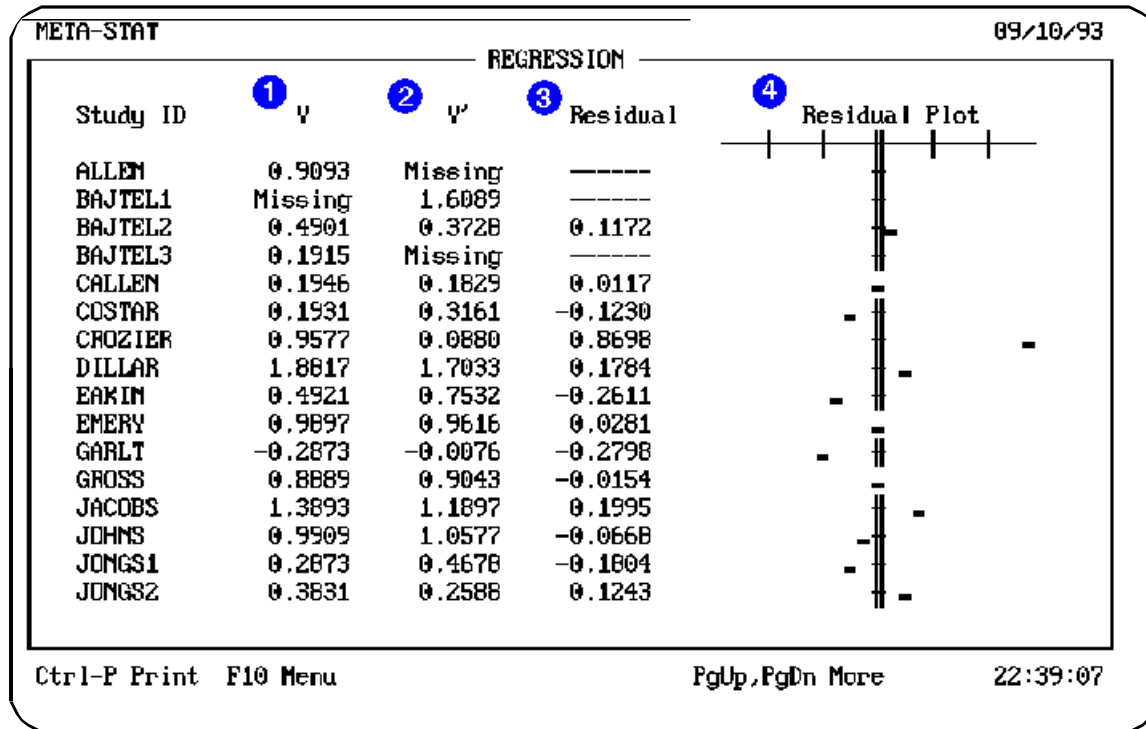
- 1 Provides a way to evaluate the contributions of each of multiple predictor variable to the *variation* in the criterion variable.

Value	Description
R	The correlation between the criterion and predictor variables. Larger numbers, approaching +1.0 or -1.0 indicate a stronger correlation. The sign of R indicates a positive or negative relationship.
Mean	The average value of the variable.
St Dev	The standard deviation for the variable. This is the square root of the unbiased population variance estimator (N-1 is the denominator).
N	The number of studies for which data were available.

Fourth Results Screen: Residuals

To display the fourth results screen, press PgDn. *Meta-Stat* asks if you want to record predicted values. If you answer Y, for "Yes", *Meta-Stat* will save the

predicted values for the criterion variable in a new variable called Y_HAT. For more information, see the section below called "Recording Predicted Values."



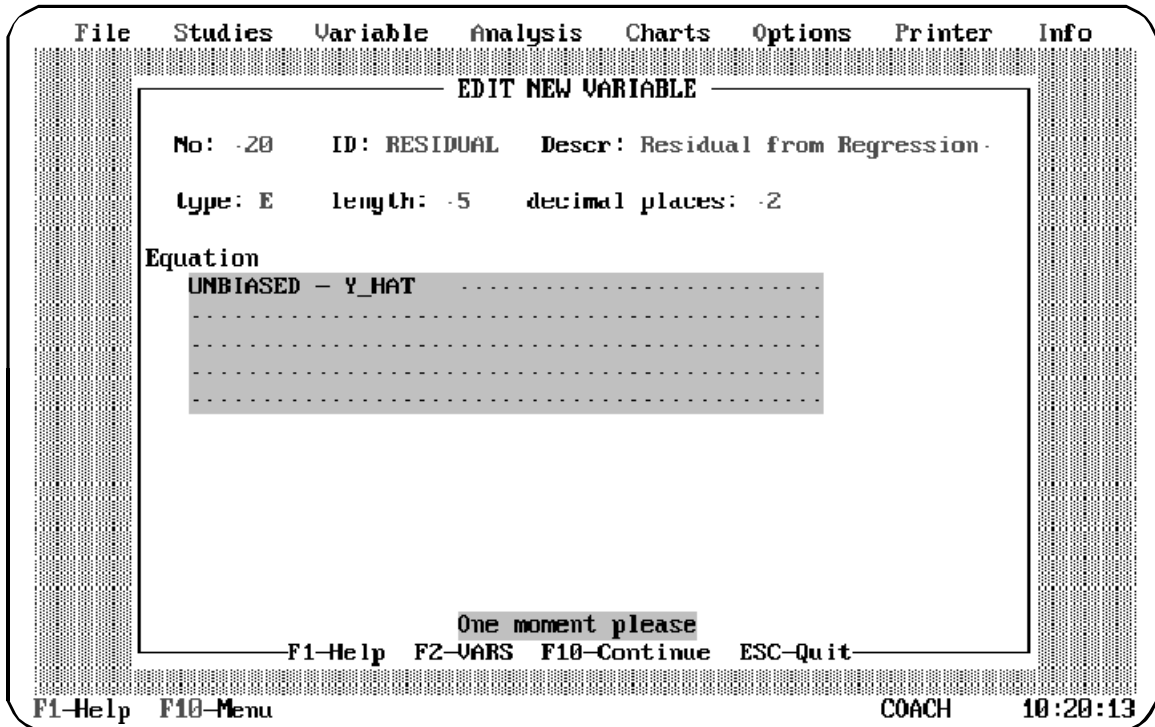
- 1 Y is the actual value of the criterion variable.
- 2 Y' is the estimated or predicted value for the criterion variable, based on the predictors. This predicted value is saved in the new variable Y_HAT, if you told *Meta-Stat* to save the values.
- 3 Residual is the difference between Y and Y'. This is the error in prediction. The closer the values are to the vertical axis, the smaller the error.
- 4 *Meta-Stat* shows a plot of the residual values.

Recording Predicted Values

On the fourth results screen for the regression analysis, *Meta-Stat* asks:

Do you want to record predicted values?

If you answer Y, for "Yes", *Meta-Stat* will save the predicted values for the criterion variable—typically the Effect Size—in a new variable called Y_HAT.



You may then want to create a second new variable to hold the residual amount—that is, the difference between the predicted value for Effect Size (Y_HAT) and the actual value.

Plots of the residuals can help you assess whether your data meets the assumptions required of least squares regression analysis. We recommend the following plots (after Draper and Smith, 1966):

1. Bar Graph of the frequencies of the residuals values. If the model is correct, the plot should resemble observations drawn from a normal distribution with a zero mean. (Since an intercept is fit, the mean must be zero; but are the residuals normally distributed?)

2. Scatterplot of the residuals against Y_HAT.
3. Scatterplots of the residuals against each of the predictor variables.

If your plots show outliers or indicate possible curvilinear relationships, you may want to try one or more of the following:

- a) Checking your data entry and evaluating the veracity of some data points.
- b) Constructing a SELECT IF to trim suspicious datapoints.
- c) Adding or removing predictor variables.
- d) Transforming the criterion variable.
- e) Transforming one or more predictor variables.
- f) Adding a weighting variable.
- g) Using Non-linear regression.

Using Dummy Variables to Explore Differences Between Groups

If you tell *Meta-Stat* to create dummy variables from a discrete or blocking variable, you can explore the differences between the groups identified by that variable.

The following example shows regression results for which dummy variables were created from the variable PUBTY_GP. The results permit the comparison of studies based on the place where the studies were published:

META-STAT		REGRESSION					06/25/93
Variable		Beta	B	St Err	t	p	
PUBTY_GP	journal	0.000	0.000	0.230	0.000	—	
PUBTY_GP	dissert	0.204	0.346	0.217	1.597	0.120	
PUBTY_GP	paper	0.020	0.047	0.217	0.216	0.025	
PUBTY_GP	other	0.615	1.664	0.454	3.663	0.002	
Constant			0.536	0.124	4.334	0.000	

Ctrl-P Print F10 Menu PgUp, PgDn MoCOACH 11:57:53

The summary table obtained when using Dummy coding is the same as the summary table that would be obtained from an analysis of variance or covariance. For more information about dummy variables, see the section "Using Dummy Variables" earlier in this chapter.

NOTE: If the criterion variable is UNBIASED and the weighting variable is WEIGHT, *Meta-Stat* follows the regression procedures for effect-sizes outlined in Hedges and Olkin (1985, 167-188). The sums of squares are reported as homogeneity statistics and are individually tested. The standard errors for the regression weights are corrected by dividing by the positive square root of the mean square error.

Testing Mean Differences Between Groups with Randomization

Using the Randomization analysis, you can test the mean differences between groups that are represented by discrete or blocking variables. Randomization is a useful tool because it makes no assumptions about the distribution of data, such as the normality assumption. Instead, the actual group of studies in hand is used for the analysis.

During the analysis, *Meta-Stat* calculates the mean value of the criterion variable for each group, then calculates the difference in means between the groups. Then *Meta-Stat* conducts up to 5,000 randomizing iterations to determine whether this actual difference is in fact unusual. During each iteration, *Meta-Stat* randomly draws groups from your group from the aggregate group of studies and calculates the differences in mean values for the groups.

After completing all iterations, *Meta-Stat* displays the probability of getting a difference in means that is equal to or larger than the actual difference. If the probability is high, you can assume that the actual difference in means is not unusual, because there is a high probability of finding this difference or one larger using randomly partitioned data. But if the probability is low, you can assume that the single obtained difference is unusual.

To use the Randomization analysis, follow these steps:

1. Select Analysis/Randomization

Meta-Stat displays a screen where you can identify the dependent and predictor variables.

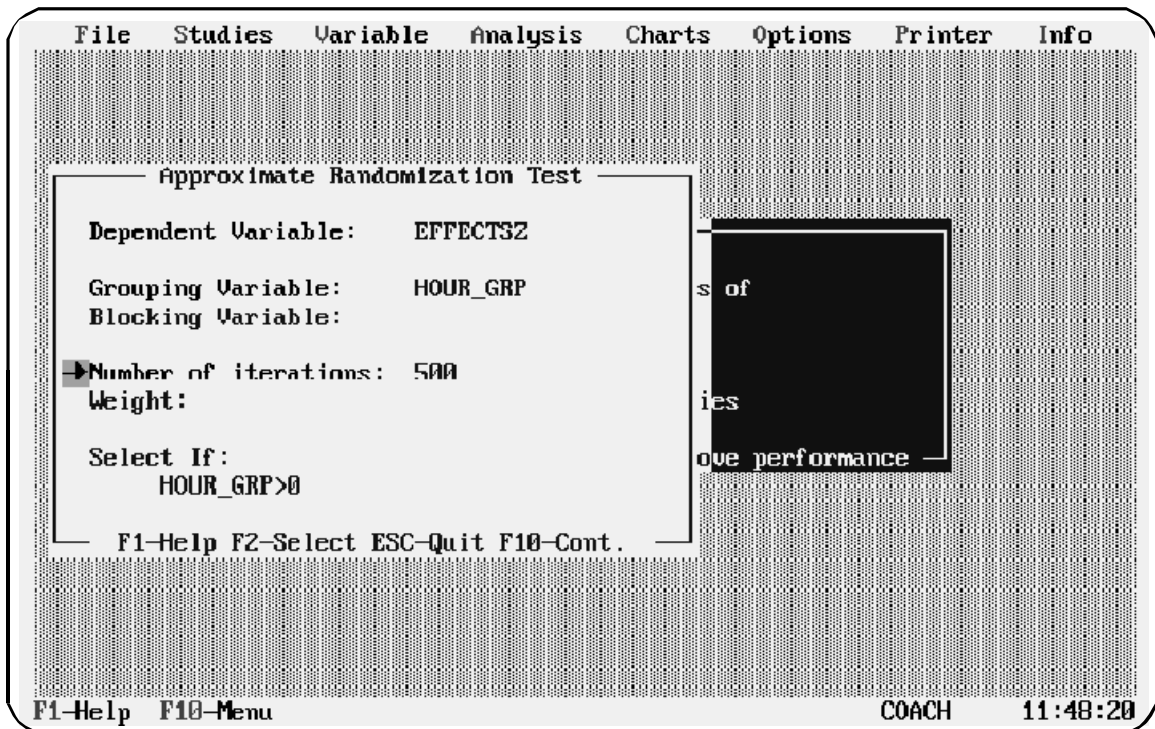
2. Use **Criterion variable** to identify the variable that you want to measure for the groups.
3. Use **Grouping Variable** to identify the discrete or blocking variable for which you want to test differences between groups.
4. Optionally, use **Blocking Variable** to control for the effects of another variable without showing the variable on the results screen.

5. Select the **Number of Iterations** you want to use. Use fewer iterations when you are exploring. Use more iterations to get more precise results.

Note that your results will be slightly different each time you perform the analysis, due to differences caused by randomization.

6. If necessary, identify a weighting variable, a method to handle missing data, and a SELECT IF statement. For more information, see "Common Functions" earlier in this chapter.

In the following example, *Meta-Stat* uses 500 iterations to measure differences in EFFECTSZ between groups identified by the HOUR_GRP variable. In other words, the analysis will test whether the variable HOUR_GRP has some relationship to the value of EFFECTSZ.



7. Press F10 to display the results of the analysis.

```

META-STAT                                09/10/93
----- APPROXIMATE RANDOMIZATION TEST OF MEAN DIFFERENCES -----
Criterion Variable  UNBIASED Unbiased Effect Size  ①
Grouping Variable  PUBYRGRP Year of publication
30 Iterations

                                Mean Differences and (probabilities)
                                Means and sample size on diagonal
                                ②
      60-70      60-70      76-80      71-75      81+
      1.1459      1.1459      0.5885      0.6136      1.0489
                   (5)      (0.129)      (0.129)      (0.032)
      76-80      0.5574      0.0251      0.4605
                   (15)      (0.968)      (0.226)
      71-75      0.5322      0.4353
                   (5)      (0.258)
      81+      0.0969
                   (4)

                                Done .

Ctrl-P Print  F10 Menu                                22:39:59
    
```

During the analysis, *Meta-Stat* first calculated the mean Effect Size that was obtained for each group (as represented by the HOUR_GRP variable). *Meta-Stat* then calculated the difference in obtained means between each group.

Next, for each iteration of the analysis, *Meta-Stat* randomly selected studies from the entire collect of studies (28 in this example) to create three groups: one with 7 studies, a second with 16 studies, and a third with 5 studies. (Using this random partition ensures that, for the subsequent calculations, there is no difference between the groups.) For each group, *Meta-Stat* calculated the average Effect Size; it then calculated the differences between these averages. *Meta-Stat* then noted whether the difference was larger than the obtained difference between the actual groups. After the last iteration, *Meta-Stat* calculated the percentage of the iterations where the calculated difference was greater than the obtained difference. This is the probability that the obtained difference is a random occurrence.

- 1 Identifies the variables and iterations you selected.

- 2 Shows the obtained means, sample sizes, differences between means, and probabilities for the groups identified by the grouping variable. Referring to the example above:
- Obtained means and sample sizes for each group are shown along the diagonal. The actual mean Effect Sizes for the Long, Normal, and Short groups are 1.4250, 0.4953, and 0.3540, respectively.

The sample sizes for the Long, Normal, and Short groups are 7, 16, and 5.
 - The difference between the obtained mean Effect Size for the Normal group (0.4953) and the obtained mean Effect Size for the Short group (0.3540) is (0.1413). In the Randomization analysis, *Meta-Stat* calculated that, over 500 iterations, there is a probability of approximately 65 percent that you would obtain this large a difference in Effect Size with randomly drawn data. In other words, the observed difference in Effect Size between these two groups is not unusual; there is a high probability of seeing this difference using randomly partitioned data.
 - By contrast, the difference between the obtained mean Effect Size for the Normal group (0.4953) and the obtained mean Effect Size for the Long group (1.4250) is (0.9297). In the Randomization analysis, *Meta-Stat* calculated that, over 500 iterations, there is a probability of well under 1 percent (0.002) that one would obtain a difference this large with randomly partitioned data. In other words, the actual observed difference between these two groups is unusual, and one can assume that there is in fact some difference between these groups.

Note: The approximate randomization test will tell whether the variables are independent. It is only appropriate for making inferences about a population when the data have been randomly selected.

Showing Correlations Between Variables

Use a Correlation analysis to investigate the relationship between variables. The analysis shows the degree to which variation in one variable is related to variation in another.

Correlations can be positive or negative. A positive correlation means that a pair of variables tends to increase and decrease together. A negative correlation indicates an inverse relationship, in which one variable increases as another decreases. The size of the correlation measures the strength of the relationship. A correlation of 1 indicates a perfect linear relationship between the variables. Correlations that are close to 0 indicate little or no linear relationship.

Correlation analysis provides a good starting point for other analyses, including Regression. After using Correlation to discover strong relationships between variables, you can use Regression to calculate prediction equations for the variables.

To perform a correlation analysis, follow these steps:

1. Select Analysis/Correlation

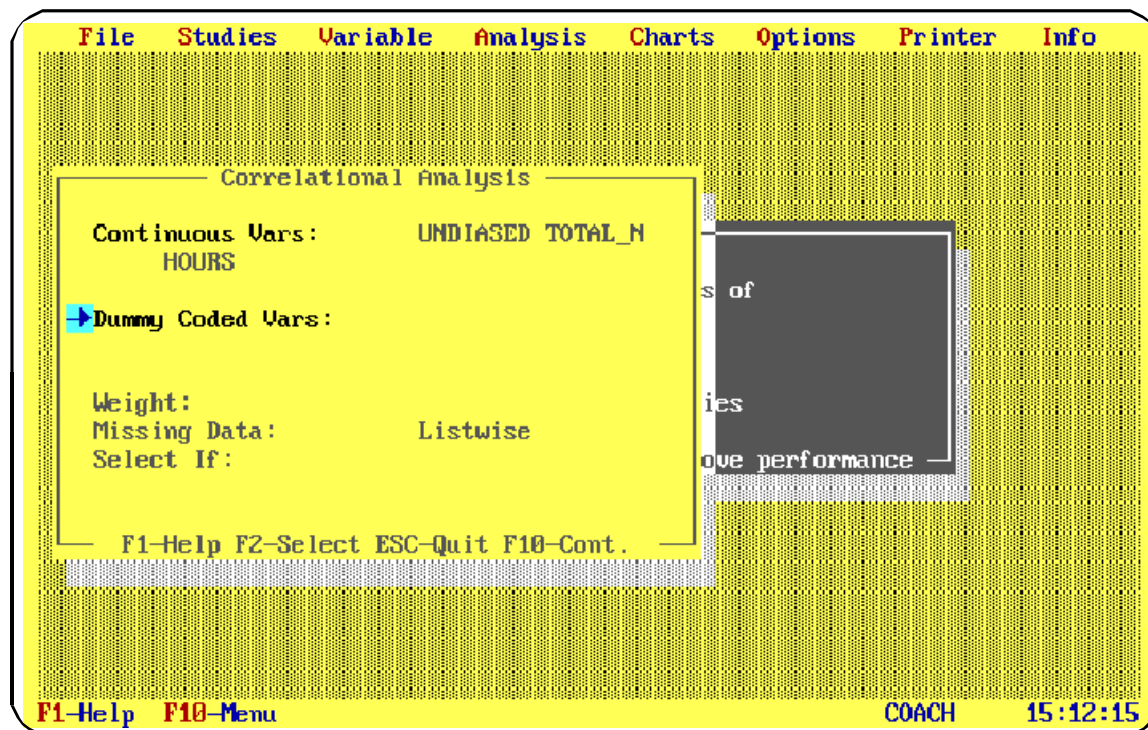
Meta-Stat displays a screen where you can identify the variables you want to use.

2. Use Continuous Vars to identify continuous or equation variables for the analysis.

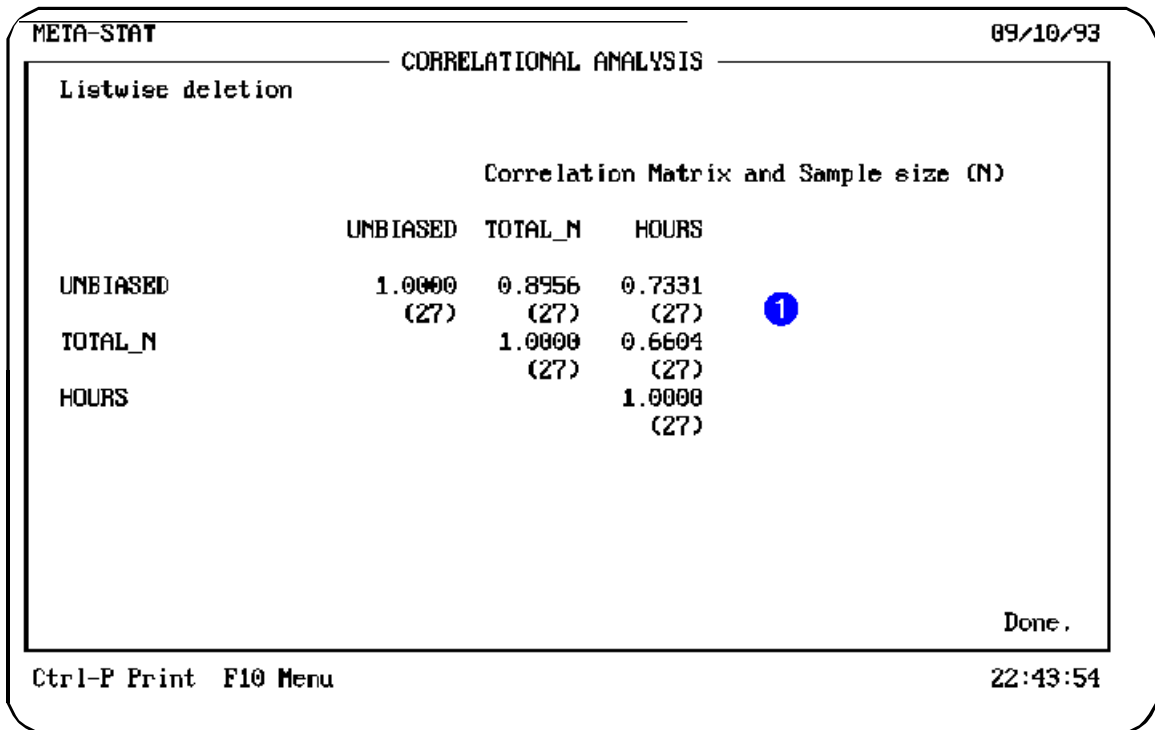
Use **Dummy Coded Vars** to identify blocking or discrete variables for the analysis. *Meta-Stat* will treat each category within the variable as a separate dummy variable for the analysis. For more information, see the section "Using Dummy Variables" earlier in this chapter.

3. If necessary, identify a weighting variable, a method to handle missing data, and a SELECT IF statement. For more information, see "Common Functions" earlier in this chapter.

In the following example, a Correlation analysis will be performed on Unbiased Effect Size, Hours, and Total Sample Size (TOTAL_N):



4. Press **F10** to display the results.



1 The screen shows the correlation coefficient between each pair of variables. In the example above, the correlation between Hours and Unbiased Effect Size is relatively strong (0.5406), while the correlation between TOTAL_N (Total Sample Size) and Unbiased Effect Size is weak (0.2423). Larger numbers indicate a stronger correlation, with 1 meaning a perfect correlation and 0 meaning no linear correlation.

The number in parentheses identifies how many studies were used for the sample. This number is determined by the number of studies in your meta-analysis and the method you used to handle missing data. For more information, see "Common Functions" earlier in this chapter.

HINT: The correlation coefficient measures only linear (i.e., straight line) relationship between two variables. It is possible, though unlikely, that two variables can be perfectly curvilinearly related and yet their correlation equals 0.

To detect curvilinear relationships, you should inspect the Scatterplot option (see Chapter 8). When variables are curvilinearly related, you

may wish to consider a transformation to "linearize" them before performing multiple regression analyses.

Cross-Tabulating Groups

Use a Cross-Tabulation analysis to investigate the relationship between two discrete or blocking variables.

The analysis displays a two-way table that shows the joint frequency distribution. You can use the table to review the distribution of values for the two groups. The analysis also provides a measure of statistical significance, to help you determine if the variables are statistically independent.

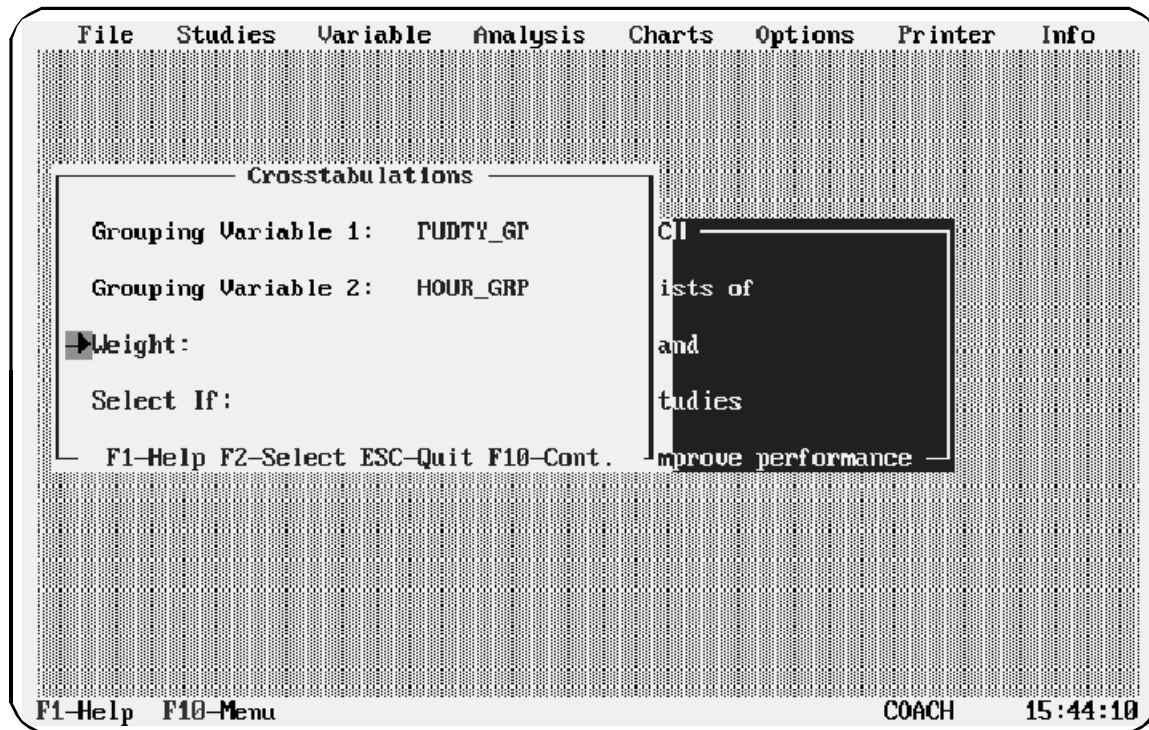
To perform a cross-tabulation, follow these steps:

1. Select Analysis/crossTabs

Meta-Stat displays a screen where you can identify the variables you want to use.

2. Use Grouping Variable 1 to identify the discrete or blocking variable that you want to use as the column variable in the two-way table.
3. Use Grouping Variable 2 to identify the discrete or blocking variable that you want to use as the row variable in the table.
4. Optionally, identify a weighting variable and construct a SELECT IF statement. For more information, see "Common Functions" earlier in this chapter.

In the following example, a cross-tabulation will be performed using grouping variables for Publication Type and Hours:



4. Press **F10** to display the results.

Results Screen

META-STAT 09/10/93

CROSSTABULATIONS

PUBTY_GRP Publication Type Valid Observations: 30 1
 by HOUR_GRP Duration of program

Observed (and Expected) Frequencies
 Publication Type

Duration of	0 Missing	1 journal	2 dissert	3 paper	5 other	Row Total
0 Missing	0 (0.07)	0 (1.00)	1 (0.40)	1 (0.40)	0 (0.13)	2.00
1 short	2 0 (0.17)	2 (2.50)	1 (1.00)	2 (1.00)	0 (0.33)	5.00
2 average	1 (0.50)	11 (7.50)	2 (3.00)	1 (3.00)	0 (1.00)	15.00
3 long	0 (0.27)	2 (4.00)	2 (1.60)	2 (1.60)	2 (0.53)	8.00
Column Total	1.00	15.00	6.00	6.00	2.00	30.00

Ctrl-P Print F10 Menu ↑↓ new grp 1 ↔ new grp 2 22:44:28

- 1 Identifies the variables you selected for the analysis, along with the number of studies for which data were available.
- 2 Shows the tabulated results, with row and column totals. The numbers in parentheses show the values that would be expected, given the row and column totals, if there were no relationship between the variables. Large differences between the observed and expected values may indicate a systematic relationship between the variables.

Cross-Tabulating for Other Groups

From the results screen, you can quickly flip through cross-tabulations for other groups in your meta-analysis. Press **8** or **9** to select a new column variable. Press the **7** or **6** cursor key to select a new row variable.

Calculating Group Means

Use a Group Means analysis to calculate means for members of categories, or groups, that are identified by a discrete or blocking variable. The analysis provides you with the mean, standard deviation, and 95 percent confidence intervals. If you use Unbiased Effect Size as the criterion variable, the Group Means analysis also provides homogeneity measures.

To use the Group Means analysis, follow these steps:

1. Select Analysis/group Means

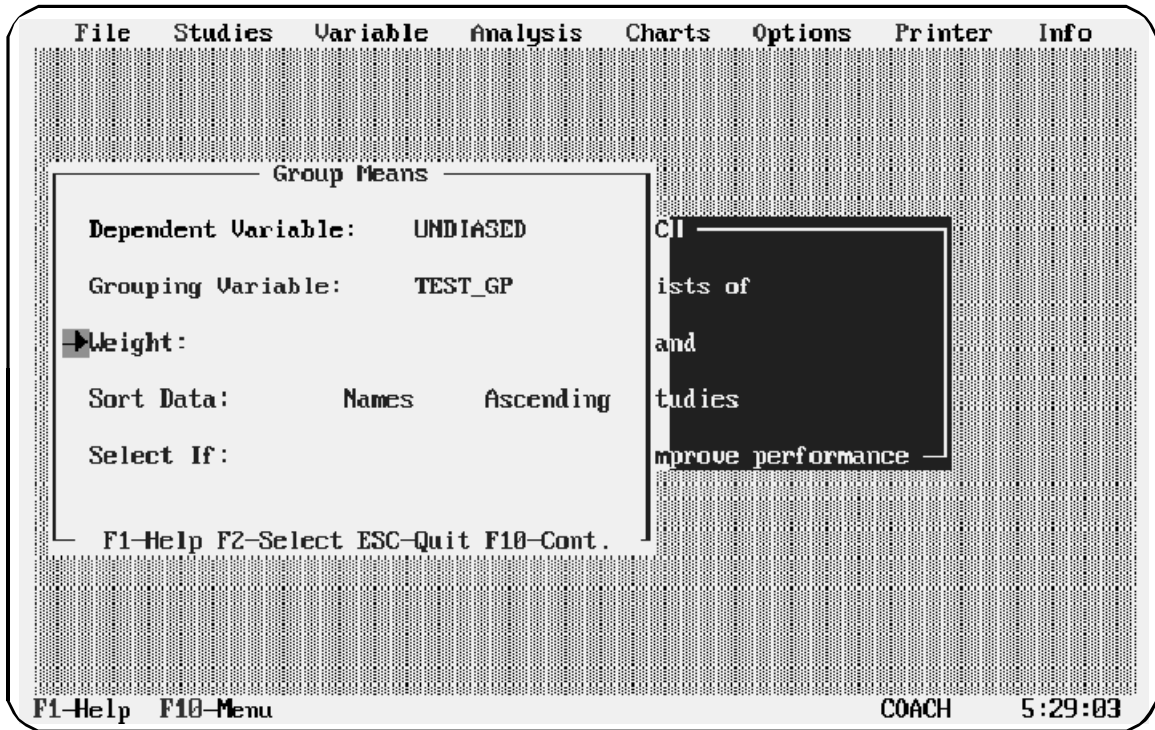
Meta-Stat displays a screen where you can identify the variables you want to use.

2. Use Criterion variable to identify the variable that you want to measure for each group. Typically, this is the effect size.
3. Use Grouping Variable to identify the discrete or blocking variable for which you want to calculate the means.

Leave this field blank to calculate means for all discrete and blocking variables.

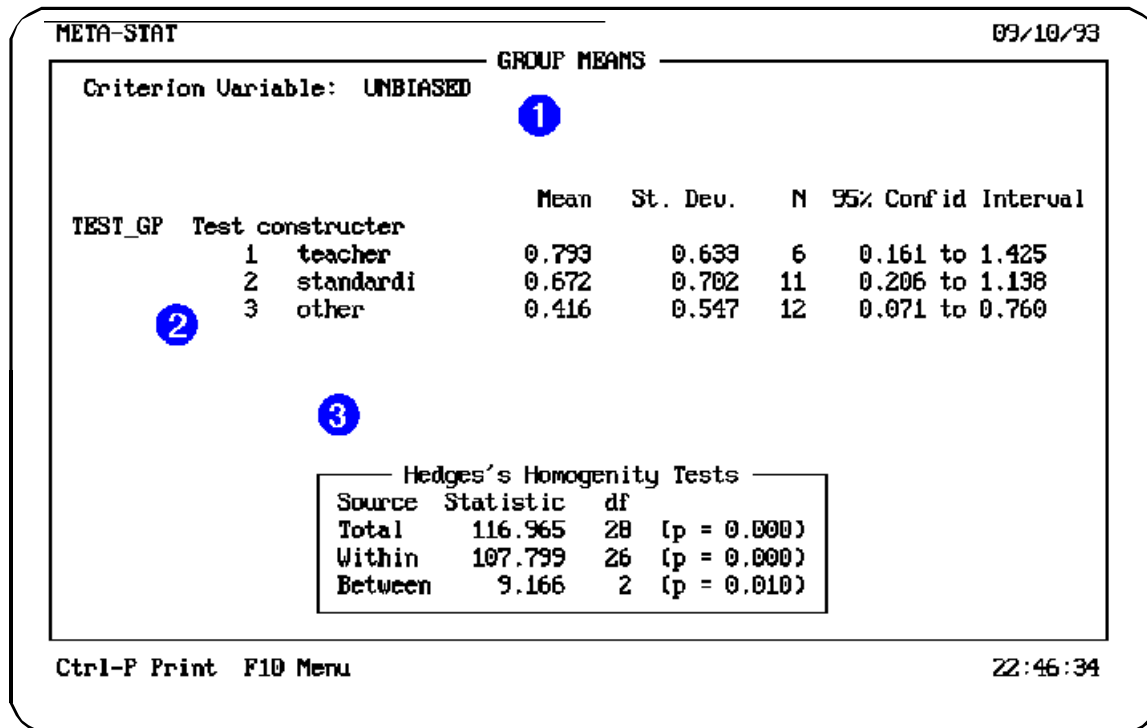
4. If you left Grouping Variable blank, use the Sort Data field to identify a sorting sequence for the group members that *Meta-Stat* will display.
5. Optionally, identify a weighting variable and construct a SELECT IF statement. For more information, see "Common Functions" earlier in this chapter.

In the following example, *Meta-Stat* will calculate the mean Effect Size for the TEST_GP variable:



6. Press **F10** to display the results.

Results Screen



- ① This screen identifies the criterion variable you are measuring for each group member.

- 2 For each group member, *Meta-Stat* displays the following information:

Value	Description
Mean	Mean value, equal to the sum of the values divided by the number of cases (studies). The mean is the most common measure of central tendency.

Stand. Dev.	Standard deviation, equal to the positive square root of the variance. (The variance is a measure of the dispersion of values around the mean, roughly equal to the average squared deviation from the mean.) The smaller the standard deviation, the more homogeneous the data.
N	Number of group members
95% Confid Int.	The 95 percent confidence intervals for the group means.

- 3 Hedges's Homogeneity Test measures the total variation among the groups, as well as the variation within and between the groups, to test the hypothesis that the Effect Sizes come from the same population—that is, that they are homogeneous. If the statistic is large (or equivalently if the value of p is low, say less than .05), Hedges's test indicates that there is a low probability that the Effect Sizes are from the same population.

Hedges and Olkin note (1985, pages 157-159) that the homogeneity statistics can be used for fitting models to effect size for groups in a procedure that is analogous to the procedure used to fit hierarchical log-linear models to contingency tables. They outline a strategy similar to the following:

Step 1: Look at the total homogeneity statistic (Q_T). If the value is not large, i.e. is not statistically significant, then you are done. You can conclude that the model of a single effect size fits the data.

Step 2: If Q_T is large, then look at the within homogeneity statistic, Q_W . If Q_W is not statistically significant, then you can conclude that the model of a different effect size for each of your groups adequately fits the data. The between fit statistic, Q_B , indicates the extent to which the effect sizes differ among groups.

A large value of Q_W indicates that effect sizes are not homogeneous within the groups. You should redefine your groups, either through secondary classifications, using a different grouping variable, creating new groups, or using the SELECT IF feature. Then again, compute Q_W .

Results Screen for All Grouping Variables

If you did not specify a grouping variable, *Meta-Stat* shows results for all discrete and blocking variables. Use PgDn to move from one screen to the next:

NOTE: If you are interested in analyzing the UNBIASED effect size as the criterion variable, *Meta-Stat* will compute Hedges's homogeneity statistics. You do not need to specify WEIGHT as a weighting variable as this is done internally to perform the necessary calculations.

Calculating Frequencies

Use Frequency analysis to count the number of different values for a variable and display the frequency at which each value occurs. This analysis gives you a clear picture of the breakdown of data for a variable.

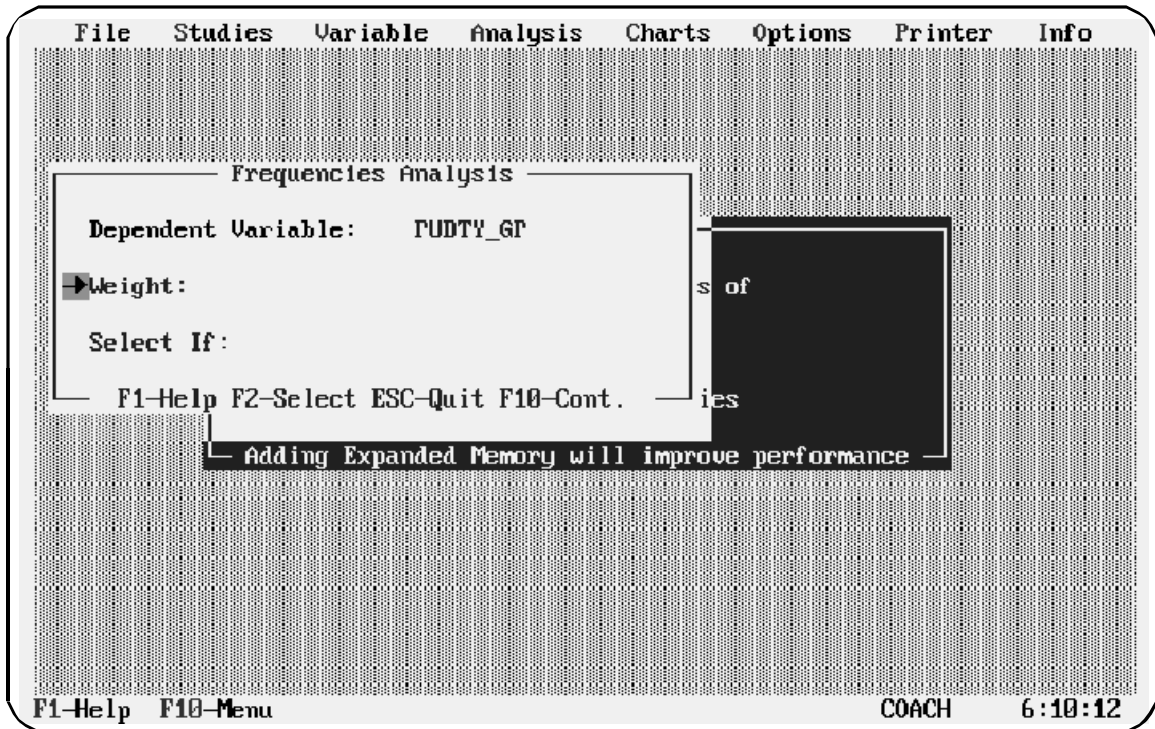
Follow these steps:

1. Select Analysis/Frequencies

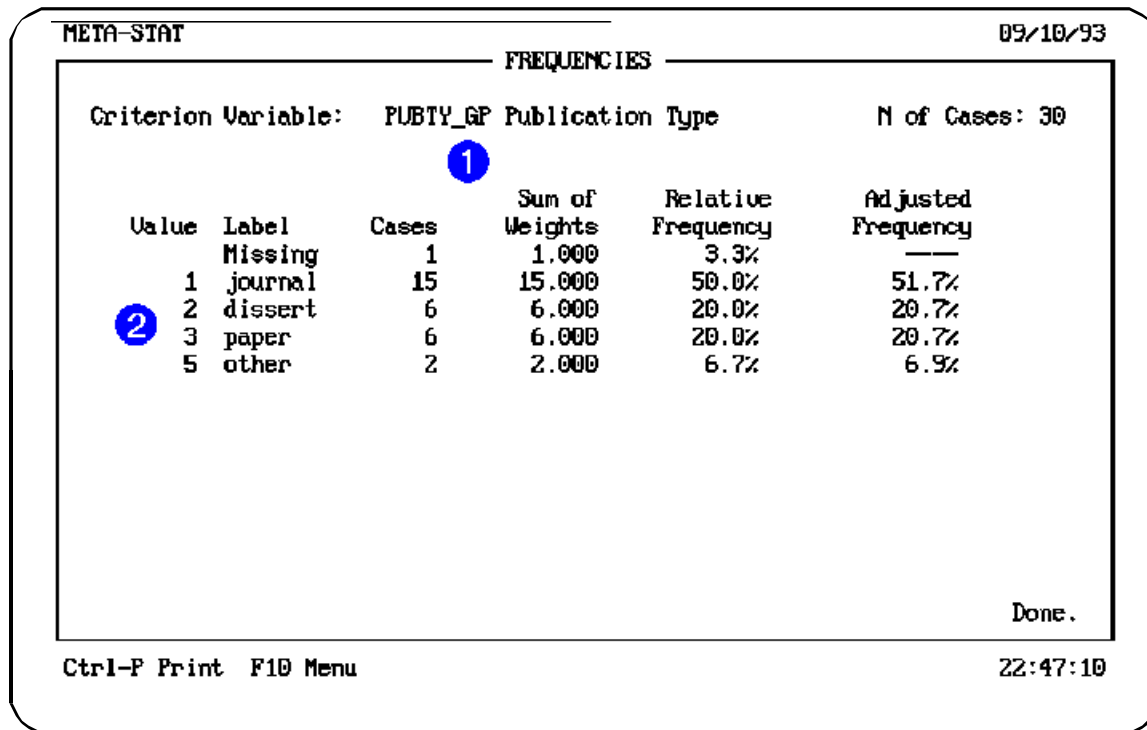
Meta-Stat displays a screen where you can identify the variable you want to count.

2. Use Criterion variable to identify the variable that you want to count. This can be any type of variable.
3. If necessary, identify a weighting variable and construct a SELECT IF statement. For more information, see "Common Functions" earlier in this chapter.

In the following example, *Meta-Stat* will count PUBTY_GP:



- Press **F10** to display the results.



The results screen:

- Identifies the variable you are measuring.
- Displays the following information for each value of the variable:

Value	Description
Value	A value for the variable.
Label	For discrete and blocking variables only, shows the label attached to the value.
Sum of Weights	The sum of the weighting data that you used, if any.
Relative Frequency	The relative frequency with which the value occurs.
Adjusted Frequency	The relative frequency with which the value occurs, adjusted after removing missing values.

Displaying Means for a Two-Way Breakdown of Groups

The Breakdown analysis is similar to the Cross-Tabulation analysis in that it provides a two-way breakdown of grouping variables. However, the Breakdown analysis provides more detail on the data by calculating, for each two-way breakdown, the mean for the criterion variable. For example, where Cross-Tabulation provides you with the number of cases for each two-way breakdown, the Breakdown analysis calculates means for the criterion variable for each two-way breakdown.

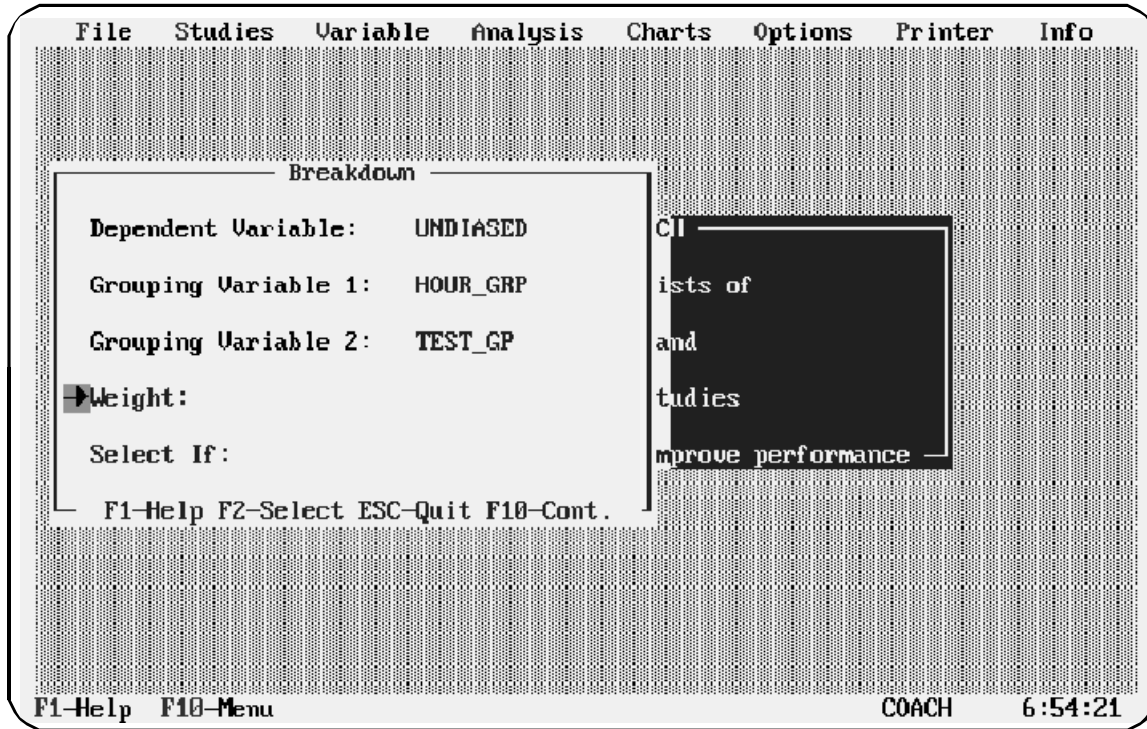
To perform the analysis, follow these steps:

1. Select Analysis/Breakdown

Meta-Stat displays a screen where you can identify the variables you want to use.

2. Use **Criterion Variable** to identify the variable you want to count.
3. Use Grouping Variable 1 to identify the discrete or blocking variable that you want to use as the column variable in the two-way table.
4. Use Grouping Variable 2 to identify the discrete or blocking variable that you want to use as the row variable.
5. Optionally, identify a weighting variable and constructing a SELECT IF statement. For more information, see "Common Functions" earlier in this chapter.

In the following example, *Meta-Stat* will calculate the mean Effect Size for each cell in the two-way breakdown of HOUR_GRP and TEST_GP:



5. Press **F10** to display the results.

Results Screen

META-STAT		BREAKDOWN				09/10/93
UMBIASED Unbiased Effect Size by HOUR_GRP Duration of program TEST_GRP Test constructor		Valid Observations: 29				
		Means (and Cell Sizes)				
		Duration of program				
Test constr	0 Missing	1 short	2 average	3 long	Row Total	
1 teacher	0.909 (1)		0.325 (3)	1.436 (2)	0.793 (6)	
2 standard	0.192 (1)		0.462 (7)	1.322 (3)	0.672 (11)	
3 other		0.178 (5)	0.428 (5)	0.980 (2)	0.416 (12)	
Column Total	0.550 (2)	0.178 (5)	0.423 (15)	1.257 (7)	0.591 (29)	

Ctrl-P Print F10 Menu ↑ new grp 1 ↔ new grp 2 22:47:44

The results screen:

- 1 Identifies the variables you selected for the analysis, along with the number of studies for which data were available.
- 2 Shows the tabulated results for each breakdown, along with means for each row and column. The numbers in parentheses show the cell sizes. For example, in studies where teachers devised tests, the mean Effect Size was 1.108. However, where the duration of the program was long, the mean Effect Size was 1.958. Long durations, in general, were associated with higher Effect Sizes (a mean of 1.475).

Showing Breakdowns for Other Groups

From the results screen, you can quickly flip through breakdowns for other groups in your meta-analysis. Press **8** or **9** to select a new column variable. Press **6** or **7** to select a new row variable.

