

> SPSS Statistics Base 17.0 User's Guide



For more information about SPSS Inc. software products, please visit our Web site at <http://www.spss.com> or contact

SPSS Inc.
233 South Wacker Drive, 11th Floor
Chicago, IL 60606-6412
Tel: (312) 651-3000
Fax: (312) 651-3668

SPSS is a registered trademark and the other product names are the trademarks of SPSS Inc. for its proprietary computer software. No material describing such software may be produced or distributed without the written permission of the owners of the trademark and license rights in the software and the copyrights in the published materials.

The SOFTWARE and documentation are provided with RESTRICTED RIGHTS. Use, duplication, or disclosure by the Government is subject to restrictions as set forth in subdivision (c) (1) (ii) of The Rights in Technical Data and Computer Software clause at 52.227-7013. Contractor/manufacturer is SPSS Inc., 233 South Wacker Drive, 11th Floor, Chicago, IL 60606-6412. Patent No. 7,023,453

General notice: Other product names mentioned herein are used for identification purposes only and may be trademarks of their respective companies.

Windows is a registered trademark of Microsoft Corporation.

Apple, Mac, and the Mac logo are trademarks of Apple Computer, Inc., registered in the U.S. and other countries.

This product uses WinWrap Basic, Copyright 1993-2007, Polar Engineering and Consulting, <http://www.winwrap.com>.

Printed in the United States of America.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior written permission of the publisher.

ISBN-13: 978-1-56827-400-3

ISBN-10: 1-56827-400-9

1 2 3 4 5 6 7 8 9 0 11 10 09 08

Preface

SPSS Statistics 17.0

SPSS Statistics 17.0 is a comprehensive system for analyzing data. SPSS Statistics can take data from almost any type of file and use them to generate tabulated reports, charts and plots of distributions and trends, descriptive statistics, and complex statistical analyses.

This manual, the *SPSS Statistics Base 17.0 User's Guide*, documents the graphical user interface of SPSS Statistics. Examples using the statistical procedures found in SPSS Statistics Base 17.0 are provided in the Help system, installed with the software.

In addition, beneath the menus and dialog boxes, SPSS Statistics uses a command language. Some extended features of the system can be accessed only via command syntax. (Those features are not available in the Student Version.) Detailed command syntax reference information is available in two forms: integrated into the overall Help system and as a separate document in PDF form in the *Command Syntax Reference*, also available from the Help menu.

SPSS Statistics Options

The following options are available as add-on enhancements to the full (not Student Version) SPSS Statistics Base system:

Regression provides techniques for analyzing data that do not fit traditional linear statistical models. It includes procedures for probit analysis, logistic regression, weight estimation, two-stage least-squares regression, and general nonlinear regression.

Advanced Statistics focuses on techniques often used in sophisticated experimental and biomedical research. It includes procedures for general linear models (GLM), linear mixed models, variance components analysis, loglinear analysis, ordinal regression, actuarial life tables, Kaplan-Meier survival analysis, and basic and extended Cox regression.

Custom Tables creates a variety of presentation-quality tabular reports, including complex stub-and-banner tables and displays of multiple response data.

Forecasting performs comprehensive forecasting and time series analyses with multiple curve-fitting models, smoothing models, and methods for estimating autoregressive functions.

Categories performs optimal scaling procedures, including correspondence analysis.

Conjoint provides a realistic way to measure how individual product attributes affect consumer and citizen preferences. With Conjoint, you can easily measure the trade-off effect of each product attribute in the context of a set of product attributes—as consumers do when making purchasing decisions.

Exact Tests calculates exact p values for statistical tests when small or very unevenly distributed samples could make the usual tests inaccurate. This option is available only on Windows operating systems.

Missing Values describes patterns of missing data, estimates means and other statistics, and imputes values for missing observations.

Complex Samples allows survey, market, health, and public opinion researchers, as well as social scientists who use sample survey methodology, to incorporate their complex sample designs into data analysis.

Decision Trees creates a tree-based classification model. It classifies cases into groups or predicts values of a dependent (target) variable based on values of independent (predictor) variables. The procedure provides validation tools for exploratory and confirmatory classification analysis.

Data Preparation provides a quick visual snapshot of your data. It provides the ability to apply validation rules that identify invalid data values. You can create rules that flag out-of-range values, missing values, or blank values. You can also save variables that record individual rule violations and the total number of rule violations per case. A limited set of predefined rules that you can copy or modify is provided.

Neural Networks can be used to make business decisions by forecasting demand for a product as a function of price and other variables, or by categorizing customers based on buying habits and demographic characteristics. Neural networks are non-linear data modeling tools. They can be used to model complex relationships between inputs and outputs or to find patterns in data.

EZ RFM performs RFM (recency, frequency, monetary) analysis on transaction data files and customer data files.

Amos™ (analysis of moment structures) uses structural equation modeling to confirm and explain conceptual models that involve attitudes, perceptions, and other factors that drive behavior.

Installation

To install the Base system, run the License Authorization Wizard using the authorization code that you received from SPSS Inc. For more information, see the installation instructions supplied with the Base system.

Compatibility

SPSS Statistics is designed to run on many computer systems. See the installation instructions that came with your system for specific information on minimum and recommended requirements.

Serial Numbers

Your serial number is your identification number with SPSS Inc. You will need this serial number when you contact SPSS Inc. for information regarding support, payment, or an upgraded system. The serial number was provided with your Base system.

Customer Service

If you have any questions concerning your shipment or account, contact your local office, listed on the Web site at <http://www.spss.com/worldwide>. Please have your serial number ready for identification.

Training Seminars

SPSS Inc. provides both public and onsite training seminars. All seminars feature hands-on workshops. Seminars will be offered in major cities on a regular basis. For more information on these seminars, contact your local office, listed on the Web site at <http://www.spss.com/worldwide>.

Technical Support

Technical Support services are available to maintenance customers. Customers may contact Technical Support for assistance in using SPSS Statistics or for installation help for one of the supported hardware environments. To reach Technical Support, see the Web site at <http://www.spss.com>, or contact your local office, listed on the Web site at <http://www.spss.com/worldwide>. Be prepared to identify yourself, your organization, and the serial number of your system.

Additional Publications

The *SPSS Statistical Procedures Companion*, by Marija Norušis, has been published by Prentice Hall. A new version of this book, updated for SPSS Statistics 17.0, is planned. The *SPSS Advanced Statistical Procedures Companion*, also based on SPSS Statistics 17.0, is forthcoming. The *SPSS Guide to Data Analysis* for SPSS Statistics 17.0 is also in development. Announcements of publications available exclusively through Prentice Hall will be available on the Web site at <http://www.spss.com/estore> (select your home country, and then click Books).

Contents

1	Overview	1
	What's New in Version 17.0?	2
	Windows	3
	Designated Window versus Active Window.	4
	Status Bar	5
	Dialog Boxes	5
	Variable Names and Variable Labels in Dialog Box Lists	5
	Resizing Dialog Boxes.	6
	Dialog Box Controls	7
	Selecting Variables.	7
	Data Type, Measurement Level, and Variable List Icons	7
	Getting Information about Variables in Dialog Boxes.	8
	Basic Steps in Data Analysis	8
	Statistics Coach	8
	Finding Out More	9
2	Getting Help	10
	Getting Help on Output Terms	11
3	Data Files	12
	Opening Data Files	12
	To Open Data Files.	12
	Data File Types	13
	Opening File Options	13
	Reading Excel 95 or Later Files.	14
	Reading Older Excel Files and Other Spreadsheets	14
	Reading dBASE Files	14
	Reading Stata Files	15
	Reading Database Files	15
	Text Wizard	28
	Reading Dimensions Data	37
	File Information.	40

Saving Data Files	41
To Save Modified Data Files	41
Saving Data Files in External Formats	41
Saving Data Files in Excel Format	44
Saving Data Files in SAS Format	44
Saving Data Files in Stata Format	46
Saving Subsets of Variables	47
Exporting to a Database	47
Exporting to Dimensions	59
Protecting Original Data	60
Virtual Active File	60
Creating a Data Cache	62

4 Distributed Analysis Mode 64

Server Login	64
Adding and Editing Server Login Settings	65
To Select, Switch, or Add Servers	66
Searching for Available Servers	67
Opening Data Files from a Remote Server	67
File Access in Local and Distributed Analysis Mode	67
Availability of Procedures in Distributed Analysis Mode	68
Absolute versus Relative Path Specifications	69

5 Data Editor 70

Data View	70
Variable View	71
To Display or Define Variable Attributes	72
Variable Names	72
Variable Measurement Level	73
Variable Type	74
Variable Labels	76
Value Labels	76
Inserting Line Breaks in Labels	77
Missing Values	77
Column Width	78
Variable Alignment	78
Applying Variable Definition Attributes to Multiple Variables	78
Custom Variable Attributes	80

Customizing Variable View	82
Spell Checking	83
Entering Data	84
To Enter Numeric Data	84
To Enter Non-Numeric Data	85
To Use Value Labels for Data Entry	85
Data Value Restrictions in the Data Editor	85
Editing Data	85
Replacing or Modifying Data Values	86
Cutting, Copying, and Pasting Data Values	86
Inserting New Cases	87
Inserting New Variables	87
To Change Data Type	88
Finding Cases, Variables, or Imputations	88
Finding and Replacing Data and Attribute Values	90
Case Selection Status in the Data Editor	90
Data Editor Display Options	91
Data Editor Printing	92
To Print Data Editor Contents	92

6 Working with Multiple Data Sources 93

Basic Handling of Multiple Data Sources	93
Working with Multiple Datasets in Command Syntax	94
Copying and Pasting Information between Datasets	95
Renaming Datasets	95
Suppressing Multiple Datasets	96

7 Data Preparation 97

Variable Properties	97
Defining Variable Properties	97
To Define Variable Properties	98
Defining Value Labels and Other Variable Properties	99
Assigning the Measurement Level	101
Custom Variable Attributes	102
Copying Variable Properties	102
Multiple Response Sets	103
Defining Multiple Response Sets	104

Copying Data Properties	106
To Copy Data Properties	107
Selecting Source and Target Variables	107
Choosing Variable Properties to Copy	109
Copying Dataset (File) Properties	110
Results	113
Identifying Duplicate Cases	113
Visual Binning.	116
To Bin Variables.	117
Binning Variables.	117
Automatically Generating Binned Categories	119
Copying Binned Categories	121
User-Missing Values in Visual Binning	123

8 Data Transformations 124

Computing Variables.	124
Compute Variable: If Cases	126
Compute Variable: Type and Label	126
Functions	127
Missing Values in Functions	127
Random Number Generators	128
Count Occurrences of Values within Cases.	129
Count Values within Cases: Values to Count.	129
Count Occurrences: If Cases	130
Shift Values	131
Recoding Values.	132
Recode into Same Variables	132
Recode into Same Variables: Old and New Values	133
Recode into Different Variables	135
Recode into Different Variables: Old and New Values	135
Automatic Recode	137
Rank Cases.	140
Rank Cases: Types.	141
Rank Cases: Ties	142
Date and Time Wizard.	142
Dates and Times in SPSS Statistics	144
Create a Date/Time Variable from a String	145
Create a Date/Time Variable from a Set of Variables.	146
Add or Subtract Values from Date/Time Variables	148

Extract Part of a Date/Time Variable	155
Time Series Data Transformations	157
Define Dates	158
Create Time Series	159
Replace Missing Values	161
Scoring Data with Predictive Models	163
Loading a Saved Model	164
Displaying a List of Loaded Models	166
Additional Features Available with Command Syntax	166

9 File Handling and File Transformations 167

Sort Cases	167
Sort Variables	168
Transpose	169
Merging Data Files	170
Add Cases	170
Add Cases: Rename	173
Add Cases: Dictionary Information	173
Merging More Than Two Data Sources	173
Add Variables	173
Add Variables: Rename	175
Merging More Than Two Data Sources	175
Aggregate Data	175
Aggregate Data: Aggregate Function	178
Aggregate Data: Variable Name and Label	178
Split File	179
Select Cases	180
Select Cases: If	181
Select Cases: Random Sample	182
Select Cases: Range	183
Weight Cases	183
Restructuring Data	184
To Restructure Data	185
Restructure Data Wizard: Select Type	185
Restructure Data Wizard (Variables to Cases): Number of Variable Groups	189
Restructure Data Wizard (Variables to Cases): Select Variables	190
Restructure Data Wizard (Variables to Cases): Create Index Variables	192
Restructure Data Wizard (Variables to Cases): Create One Index Variable	194
Restructure Data Wizard (Variables to Cases): Create Multiple Index Variables	195
Restructure Data Wizard (Variables to Cases): Options	196

Restructure Data Wizard (Cases to Variables): Select Variables	197
Restructure Data Wizard (Cases to Variables): Sort Data	199
Restructure Data Wizard (Cases to Variables): Options	199
Restructure Data Wizard: Finish	201

10 Working with Output 203

Viewer	203
Showing and Hiding Results.	204
Moving, Deleting, and Copying Output	204
Changing Initial Alignment	205
Changing Alignment of Output Items	205
Viewer Outline	205
Adding Items to the Viewer	207
Finding and Replacing Information in the Viewer	208
Copying Output into Other Applications	209
To Copy and Paste Output Items into Another Application	209
Export Output	210
HTML Options	212
Word/RTF Options	213
Excel Options.	214
PowerPoint Options.	216
PDF Options.	217
Text Options.	219
Graphics Only Options	220
Graphics Format Options	221
Viewer Printing.	222
To Print Output and Charts	222
Print Preview.	222
Page Attributes: Headers and Footers	223
Page Attributes: Options	225
Saving Output.	226
To Save a Viewer Document	226

11 Pivot Tables 227

Manipulating a Pivot Table	227
Activating a Pivot Table	227
Pivoting a Table	227
Changing Display Order of Elements within a Dimension	228

Moving Rows and Columns within a Dimension Element	228
Transposing Rows and Columns.	229
Grouping Rows or Columns	229
Ungrouping Rows or Columns	229
Rotating Row or Column Labels	229
Working with Layers	230
Creating and Displaying Layers	230
Go to Layer Category	232
Showing and Hiding Items	232
Hiding Rows and Columns in a Table	233
Showing Hidden Rows and Columns in a Table.	233
Hiding and Showing Dimension Labels.	233
Hiding and Showing Table Titles.	233
TableLooks	234
To Apply or Save a TableLook.	234
To Edit or Create a TableLook	235
Table Properties	235
To Change Pivot Table Properties.	235
Table Properties: General	235
Table Properties: Footnotes	236
Table Properties: Cell Formats	237
Table Properties: Borders	239
Table Properties: Printing	240
Cell Properties	241
Font and Background.	242
Format Value	242
Alignment and Margins	243
Footnotes and Captions	244
Adding Footnotes and Captions	244
To Hide or Show a Caption	245
To Hide or Show a Footnote in a Table	245
Footnote Marker	245
Renumbering Footnotes.	245
Data Cell Widths.	246
Changing Column Width	246
Displaying Hidden Borders in a Pivot Table	246
Selecting Rows and Columns in a Pivot Table	247
Printing Pivot Tables	248
Controlling Table Breaks for Wide and Long Tables	248
Creating a Chart from a Pivot Table	248

12 Models **250**

Interacting with a Model.	250
Working with the Model Viewer.	250
Printing a Model.	252
Exporting a Model.	252

13 Working with Command Syntax **253**

Syntax Rules.	253
Pasting Syntax from Dialog Boxes	255
To Paste Syntax from Dialog Boxes	255
Copying Syntax from the Output Log	255
To Copy Syntax from the Output Log.	256
Using the Syntax Editor.	257
Syntax Editor Window	257
Terminology.	259
Auto-Completion	259
Color Coding	260
Breakpoints.	261
Bookmarks	262
Commenting Out Text.	263
Running Command Syntax	263
Unicode Syntax Files	264
Multiple Execute Commands.	265

14 Codebook **266**

Codebook Output Tab	268
Codebook Statistics Tab	270

15 Frequencies **273**

Frequencies Statistics	274
Frequencies Charts.	276
Frequencies Format	276

16 Descriptives 278

Descriptives Options 279
DESCRIPTIVES Command Additional Features 280

17 Explore 282

Explore Statistics 283
Explore Plots 284
 Explore Power Transformations 285
Explore Options 285
EXAMINE Command Additional Features 286

18 Crosstabs 287

Crosstabs Layers 288
Crosstabs Clustered Bar Charts 289
Crosstabs Statistics 289
Crosstabs Cell Display 291
Crosstabs Table Format 292

19 Summarize 294

Summarize Options 295
Summarize Statistics 296

20 Means 298

Means Options 300

21 OLAP Cubes 302

OLAP Cubes Statistics 303

OLAP Cubes Differences	305
OLAP Cubes Title	306
22 T Tests	307
Independent-Samples T Test	307
Independent-Samples T Test Define Groups	309
Independent-Samples T Test Options	309
Paired-Samples T Test	310
Paired-Samples T Test Options	311
One-Sample T Test	311
One-Sample T Test Options	313
T-TEST Command Additional Features	313
23 One-Way ANOVA	314
One-Way ANOVA Contrasts	315
One-Way ANOVA Post Hoc Tests	316
One-Way ANOVA Options	318
ONEWAY Command Additional Features	319
24 GLM Univariate Analysis	320
GLM Model	322
Build Terms	322
Sum of Squares	323
GLM Contrasts	324
Contrast Types	324
GLM Profile Plots	325
GLM Post Hoc Comparisons	326
GLM Save	328
GLM Options	329
UNIANOVA Command Additional Features	330

25	<i>Bivariate Correlations</i>	332
	Bivariate Correlations Options	334
	CORRELATIONS and NONPAR CORR Command Additional Features	334
26	<i>Partial Correlations</i>	335
	Partial Correlations Options	336
	PARTIAL CORR Command Additional Features	337
27	<i>Distances</i>	338
	Distances Dissimilarity Measures	340
	Distances Similarity Measures	341
	PROXIMITIES Command Additional Features	341
28	<i>Linear Regression</i>	343
	Linear Regression Variable Selection Methods	344
	Linear Regression Set Rule	345
	Linear Regression Plots	346
	Linear Regression: Saving New Variables	347
	Linear Regression Statistics	349
	Linear Regression Options	351
	REGRESSION Command Additional Features	352
29	<i>Ordinal Regression</i>	353
	Ordinal Regression Options	354
	Ordinal Regression Output	355
	Ordinal Regression Location Model	356
	Build Terms	358
	Ordinal Regression Scale Model	357
	Build Terms	358
	PLUM Command Additional Features	358

30 Curve Estimation 359

Curve Estimation Models 360
Curve Estimation Save 361

31 Partial Least Squares Regression 363

Model 365
Options 366

32 Nearest Neighbor Analysis 367

Neighbors 371
Features 372
Partitions 373
Save 375
Output 376
Options 377
Model View 378
 Feature Space 379
 Variable Importance 380
 Peers 381
 Nearest Neighbor Distances 382
 Quadrant Map 382
 Feature Selection Error Log 383
 k Selection Error Log 384
 k and Feature Selection Error Log 385
 Classification Table 385
 Error Summary 386

33 Discriminant Analysis 387

Discriminant Analysis Define Range 389
Discriminant Analysis Select Cases 389
Discriminant Analysis Statistics 390
Discriminant Analysis Stepwise Method 391
Discriminant Analysis Classification 392

Discriminant Analysis Save.	393
DISCRIMINANT Command Additional Features.	393
34 Factor Analysis	395
Factor Analysis Select Cases	396
Factor Analysis Descriptives.	397
Factor Analysis Extraction	398
Factor Analysis Rotation.	399
Factor Analysis Scores.	400
Factor Analysis Options	401
FACTOR Command Additional Features.	401
35 Choosing a Procedure for Clustering	403
36 TwoStep Cluster Analysis	404
TwoStep Cluster Analysis Options.	407
TwoStep Cluster Analysis Plots.	409
TwoStep Cluster Analysis Output	410
37 Hierarchical Cluster Analysis	412
Hierarchical Cluster Analysis Method.	413
Hierarchical Cluster Analysis Statistics.	414
Hierarchical Cluster Analysis Plots.	415
Hierarchical Cluster Analysis Save New Variables	416
CLUSTER Command Syntax Additional Features	416
38 K-Means Cluster Analysis	417
K-Means Cluster Analysis Efficiency.	418
K-Means Cluster Analysis Iterate	419
K-Means Cluster Analysis Save	419

K-Means Cluster Analysis Options	420
QUICK CLUSTER Command Additional Features	421

39 Nonparametric Tests 422

Chi-Square Test	422
Chi-Square Test Expected Range and Expected Values	424
Chi-Square Test Options	424
NPAR TESTS Command Additional Features (Chi-Square Test)	425
Binomial Test	425
Binomial Test Options	426
NPAR TESTS Command Additional Features (Binomial Test)	427
Runs Test	427
Runs Test Cut Point	428
Runs Test Options	428
NPAR TESTS Command Additional Features (Runs Test)	429
One-Sample Kolmogorov-Smirnov Test	429
One-Sample Kolmogorov-Smirnov Test Options	430
NPAR TESTS Command Additional Features (One-Sample Kolmogorov-Smirnov Test)	431
Two-Independent-Samples Tests	431
Two-Independent-Samples Test Types	432
Two-Independent-Samples Tests Define Groups	433
Two-Independent-Samples Tests Options	433
NPAR TESTS Command Additional Features (Two-Independent-Samples Tests)	434
Two-Related-Samples Tests	434
Two-Related-Samples Test Types	435
Two-Related-Samples Tests Options	435
NPAR TESTS Command Additional Features (Two Related Samples)	436
Tests for Several Independent Samples	436
Tests for Several Independent Samples Test Types	437
Tests for Several Independent Samples Define Range	437
Tests for Several Independent Samples Options	438
NPAR TESTS Command Additional Features (K Independent Samples)	438
Tests for Several Related Samples	438
Tests for Several Related Samples Test Types	439
Tests for Several Related Samples Statistics	440
NPAR TESTS Command Additional Features (K Related Samples)	440

40 Multiple Response Analysis **441**

Multiple Response Define Sets	441
Multiple Response Frequencies	443
Multiple Response Crosstabs	444
Multiple Response Crosstabs Define Ranges	446
Multiple Response Crosstabs Options.	446
MULT RESPONSE Command Additional Features	447

41 Reporting Results **448**

Report Summaries in Rows.	448
To Obtain a Summary Report: Summaries in Rows	449
Report Data Column/Break Format.	449
Report Summary Lines for/Final Summary Lines.	450
Report Break Options.	451
Report Options.	451
Report Layout	452
Report Titles	453
Report Summaries in Columns	454
To Obtain a Summary Report: Summaries in Columns.	454
Data Columns Summary Function.	455
Data Columns Summary for Total Column.	456
Report Column Format.	457
Report Summaries in Columns Break Options.	457
Report Summaries in Columns Options.	458
Report Layout for Summaries in Columns.	458
REPORT Command Additional Features.	458

42 Reliability Analysis **460**

Reliability Analysis Statistics	462
RELIABILITY Command Additional Features	463

43 Multidimensional Scaling **465**

Multidimensional Scaling Shape of Data.	467
---	-----

Multidimensional Scaling Create Measure	467
Multidimensional Scaling Model.	468
Multidimensional Scaling Options.	469
ALSCAL Command Additional Features	470
44 Ratio Statistics	471
Ratio Statistics	473
45 ROC Curves	475
ROC Curve Options	477
46 Overview of the Chart Facility	478
Building and Editing a Chart	478
Building Charts	478
Editing Charts	482
Chart Definition Options	485
Adding and Editing Titles and Footnotes.	485
Setting General Options.	485
47 Utilities	488
Variable Information	488
Data File Comments	489
Variable Sets	489
Define Variable Sets	489
Use Variable Sets	490
Reordering Target Variable Lists	492
48 Options	493
General Options	494
Viewer Options	496

Data Options	497
Changing the Default Variable View	499
Currency Options	499
To Create Custom Currency Formats	500
Output Label Options	500
Chart Options	502
Data Element Colors	503
Data Element Lines	503
Data Element Markers	504
Data Element Fills	504
Pivot Table Options	505
File Locations Options	506
Script Options	507
Syntax Editor Options	510
Multiple Imputations Options	512

49 Customizing Menus and Toolbars 514

Menu Editor	514
Customizing Toolbars	515
Show Toolbars	515
To Customize Toolbars	516
Toolbar Properties	517
Edit Toolbar	517
Create New Tool	518

50 Creating and Managing Custom Dialogs 520

Custom Dialog Builder Layout	521
Building a Custom Dialog	522
Dialog Properties	522
Specifying the Menu Location for a Custom Dialog	523
Laying Out Controls on the Canvas	524
Building the Syntax Template	525
Previewing a Custom Dialog	527
Managing Custom Dialogs	528
Control Types	529
Source List	530

Target List	531
Filtering Variable Lists	532
Check Box	532
Combo Box	533
Text Control	534
Number Control	534
Static Text Control	535
Item Group	535
Radio Group	536
Check Box Group	537
File Browser	538
Sub-dialog Button	539
Custom Dialogs for Extension Commands	540
Creating Localized Versions of Custom Dialogs	541

51 Production Jobs 544

HTML Options	547
PowerPoint Options	547
PDF Options	547
Text Options	547
Runtime Values	548
User Prompts	549
Running Production Jobs from a Command Line	549
Converting Production Facility Files	551

52 Output Management System 552

Output Object Types	554
Command Identifiers and Table Subtypes	556
Labels	557
OMS Options	558
Logging	563
Excluding Output Display from the Viewer	563
Routing Output to SPSS Statistics Data Files	564
Example: Single Two-Dimensional Table	564
Example: Tables with Layers	565
Data Files Created from Multiple Tables	566

Controlling Column Elements to Control Variables in the Data File	569
Variable Names in OMS-Generated Data Files	571
OXML Table Structure	571
OMS Identifiers	575
Copying OMS Identifiers from the Viewer Outline	576

53 Scripting Facility 578

Autoscripts	579
Creating Autoscripts	580
Associating Existing Scripts with Viewer Objects	581
Scripting with the Python Programming Language	582
Running Python Scripts and Python programs	583
Getting Started with Python Scripts	584
Getting Started with Autoscripts in Python	585
Running Python Scripts from Python Programs	586
Script Editor for the Python Programming Language	588
Scripting in Basic	588
Compatibility with Versions Prior to 16.0	589
The scriptContext Object	591
Startup Scripts	592

Appendix

A TABLES and IGRAPH Command Syntax Converter 593

Index 596

Overview

SPSS Statistics provides a powerful statistical-analysis and data-management system in a graphical environment, using descriptive menus and simple dialog boxes to do most of the work for you. Most tasks can be accomplished simply by pointing and clicking the mouse.

In addition to the simple point-and-click interface for statistical analysis, SPSS Statistics provides:

Data Editor. The Data Editor is a versatile spreadsheet-like system for defining, entering, editing, and displaying data.

Viewer. The Viewer makes it easy to browse your results, selectively show and hide output, change the display order results, and move presentation-quality tables and charts to and from other applications.

Multidimensional pivot tables. Your results come alive with multidimensional pivot tables. Explore your tables by rearranging rows, columns, and layers. Uncover important findings that can get lost in standard reports. Compare groups easily by splitting your table so that only one group is displayed at a time.

High-resolution graphics. High-resolution, full-color pie charts, bar charts, histograms, scatterplots, 3-D graphics, and more are included as standard features.

Database access. Retrieve information from databases by using the Database Wizard instead of complicated SQL queries.

Data transformations. Transformation features help get your data ready for analysis. You can easily subset data; combine categories; add, aggregate, merge, split, and transpose files; and more.

Online Help. Detailed tutorials provide a comprehensive overview; context-sensitive Help topics in dialog boxes guide you through specific tasks; pop-up definitions in pivot table results explain statistical terms; the Statistics Coach helps you find the procedures that you need; Case Studies provide hands-on examples of how to use statistical procedures and interpret the results.

Command language. Although most tasks can be accomplished with simple point-and-click gestures, SPSS Statistics also provides a powerful command language that allows you to save and automate many common tasks. The command language also provides some functionality that is not found in the menus and dialog boxes.

Complete command syntax documentation is integrated into the overall Help system and is available as a separate PDF document, *Command Syntax Reference*, which is also available from the Help menu.

What's New in Version 17.0?

New Syntax Editor. The syntax editor has been completely redesigned with features such as auto-completion, color coding, bookmarks, and breakpoints. Auto-completion provides you with a list of valid command names, subcommands, and keywords; so you'll spend less time referring to syntax charts. Color coding allows you to quickly spot unrecognized terms as well as some common syntactical errors. Bookmarks allow you to quickly navigate large command syntax files. Breakpoints allow you to stop execution at specified points so you can inspect data or output before proceeding. [For more information, see Using the Syntax Editor in Chapter 13 on p. 257.](#)

Custom Dialog Builder. The Custom Dialog Builder allows you to create and manage custom dialogs for generating command syntax. You can create custom dialogs to generate syntax from multiple commands, including custom extension commands implemented in Python or R. [For more information, see Creating and Managing Custom Dialogs in Chapter 50 on p. 520.](#)

Multiple language support. In addition to the ability to change the output language available in previous releases, you can now change the user interface language. [For more information, see General Options in Chapter 48 on p. 494.](#)

Codebook. The Codebook procedure reports the dictionary information — such as variable names, variable labels, value labels, missing values — and summary statistics for all or specified variables and multiple response sets in the active dataset. For nominal and ordinal variables and multiple response sets, summary statistics include counts and percents. For scale variables, summary statistics include mean, standard deviation, and quartiles. [For more information, see Codebook in Chapter 14 on p. 266.](#)

Nearest Neighbor analysis. Nearest Neighbor analysis is a method for classifying cases based on their similarity to other cases. In machine learning, it was developed as a way to recognize patterns of data without requiring an exact match to any stored patterns, or cases. Similar cases are near each other and dissimilar cases are distant from each other. Thus, the distance between two cases is a measure of their dissimilarity. [For more information, see Nearest Neighbor Analysis in Chapter 32 on p. 367.](#)

Multiple Imputation. The Multiple Imputation procedure performs multiple imputation of missing data values. Given a dataset containing missing values, it outputs one or more datasets in which missing values are replaced with plausible estimates. You can then obtain pooled results when running other procedures. The procedure also summarizes missing values in the working dataset. This feature is available in the Missing Values add-on option.

RFM analysis. RFM (recency, frequency, monetary) analysis is a technique used to identify existing customers who are most likely to respond to a new offer. This technique is commonly used in direct marketing. This feature is available in the EZ RFM add-on option.

Categorical Regression enhancements. Categorical Regression has been enhanced to include regularization and resampling methods to assess and improve prediction accuracy. Together, these new methods make it possible to create state-of-the-art models, even for high-volume data (where there are more variables than observations, such as in genomics). This feature is available in the Categories add-on option.

Graphboard. Graphboard visualizations are graphs, charts, and plots created from a visualization template. SPSS Statistics ships with built-in visualization templates. You can also use a separate product, SPSS Viz Designer, to create your own visualization templates. The new visualization templates are effectively custom visualization types.

Exporting Output. More output export format options and more control over exported content, including:

- Wrap or shrink wide table in Word documents. [For more information, see Word/RTF Options in Chapter 10 on p. 213.](#)
- Create new worksheets or append data to existing worksheets in an Excel workbook. [For more information, see Excel Options in Chapter 10 on p. 214.](#)
- Save output export specifications in the form of command syntax with the `OUTPUT EXPORT` command. All the features for exporting output in the Export Output dialog are now also available in command syntax; so you can save and re-run your export specifications and include them in automated production jobs.
- The Output Management System (OMS) now supports these additional output formats: Word, Excel, and PDF. [For more information, see Output Management System in Chapter 52 on p. 552.](#)

Shift Values. Shift Values creates new variables that contain the values of existing variables from preceding (lag) or subsequent (lead) cases. [For more information, see Shift Values in Chapter 8 on p. 131.](#)

Aggregate enhancements. You can now use the features of the Aggregate procedure without specifying a break variable. [For more information, see Aggregate Data in Chapter 9 on p. 175.](#)

Median function. A median function is now available for computing the median value across selected variables for each case.

Windows

There are a number of different types of windows in SPSS Statistics:

Data Editor. The Data Editor displays the contents of the data file. You can create new data files or modify existing data files with the Data Editor. If you have more than one data file open, there is a separate Data Editor window for each data file.

Viewer. All statistical results, tables, and charts are displayed in the Viewer. You can edit the output and save it for later use. A Viewer window opens automatically the first time you run a procedure that generates output.

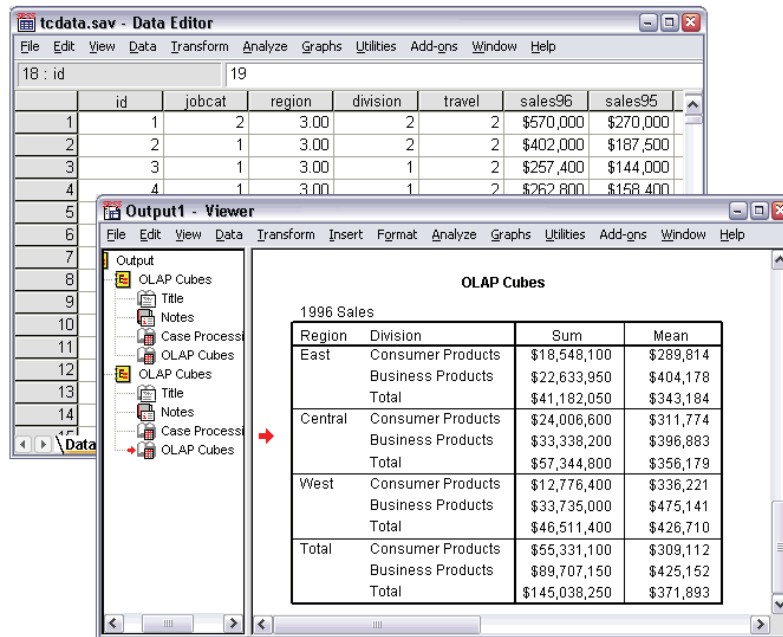
Pivot Table Editor. Output that is displayed in pivot tables can be modified in many ways with the Pivot Table Editor. You can edit text, swap data in rows and columns, add color, create multidimensional tables, and selectively hide and show results.

Chart Editor. You can modify high-resolution charts and plots in chart windows. You can change the colors, select different type fonts or sizes, switch the horizontal and vertical axes, rotate 3-D scatterplots, and even change the chart type.

Text Output Editor. Text output that is not displayed in pivot tables can be modified with the Text Output Editor. You can edit the output and change font characteristics (type, style, color, size).

Syntax Editor. You can paste your dialog box choices into a syntax window, where your selections appear in the form of command syntax. You can then edit the command syntax to use special features that are not available through dialog boxes. You can save these commands in a file for use in subsequent sessions.

Figure 1-1
Data Editor and Viewer



Designated Window versus Active Window

If you have more than one open Viewer window, output is routed to the **designated** Viewer window. If you have more than one open Syntax Editor window, command syntax is pasted into the designated Syntax Editor window. The designated windows are indicated by a plus sign in the icon in the title bar. You can change the designated windows at any time.

The designated window should not be confused with the **active** window, which is the currently selected window. If you have overlapping windows, the active window appears in the foreground. If you open a window, that window automatically becomes the active window and the designated window.

Changing the Designated Window

- Make the window that you want to designate the active window (click anywhere in the window).

- ▶ Click the Designate Window button on the toolbar (the plus sign icon).

or

- ▶ From the menus choose:
 - Utilities
 - Designate Window

Note: For Data Editor windows, the active Data Editor window determines the dataset that is used in subsequent calculations or analyses. There is no “designated” Data Editor window. [For more information, see Basic Handling of Multiple Data Sources in Chapter 6 on p. 93.](#)

Status Bar

The status bar at the bottom of each SPSS Statistics window provides the following information:

Command status. For each procedure or command that you run, a case counter indicates the number of cases processed so far. For statistical procedures that require iterative processing, the number of iterations is displayed.

Filter status. If you have selected a random sample or a subset of cases for analysis, the message Filter on indicates that some type of case filtering is currently in effect and not all cases in the data file are included in the analysis.

Weight status. The message Weight on indicates that a weight variable is being used to weight cases for analysis.

Split File status. The message Split File on indicates that the data file has been split into separate groups for analysis, based on the values of one or more grouping variables.

Dialog Boxes

Most menu selections open dialog boxes. You use dialog boxes to select variables and options for analysis.

Dialog boxes for statistical procedures and charts typically have two basic components:

Source variable list. A list of variables in the active dataset. Only variable types that are allowed by the selected procedure are displayed in the source list. Use of short string and long string variables is restricted in many procedures.

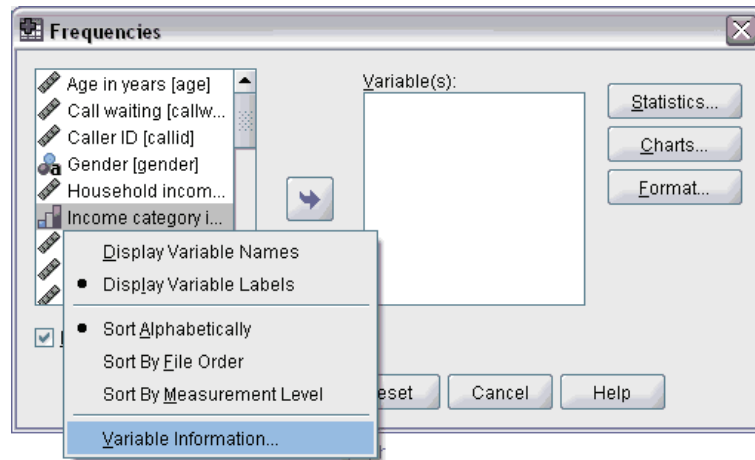
Target variable list(s). One or more lists indicating the variables that you have chosen for the analysis, such as dependent and independent variable lists.

Variable Names and Variable Labels in Dialog Box Lists

You can display either variable names or variable labels in dialog box lists, and you can control the sort order of variables in source variable lists.

- To control the default display attributes of variables in source lists, choose Options on the Edit menu. For more information, see [General Options in Chapter 48 on p. 494](#).
- To change the source variable list display attributes within a dialog box, right-click on any variable in the source list and select the display attributes from the context menu. You can display either variable names or variable labels (names are displayed for any variables without defined labels), and you can sort the source list by file order, alphabetical order, or measurement level. For more information on measurement level, see [Data Type, Measurement Level, and Variable List Icons on p. 7](#).

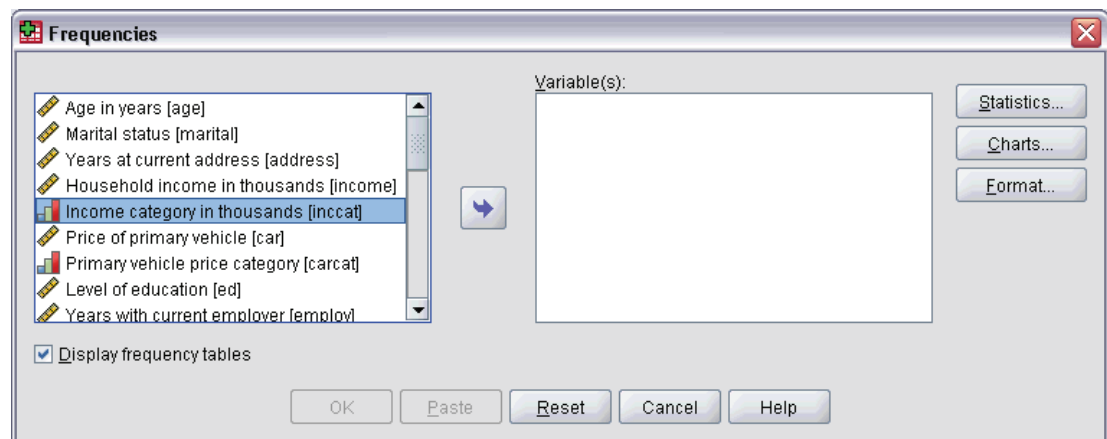
Figure 1-2
Variable labels displayed in a dialog box



Resizing Dialog Boxes

You can resize dialog boxes just like windows, by clicking and dragging the outside borders or corners. For example, if you make the dialog box wider, the variable lists will also be wider.

Figure 1-3
Resized dialog box



Dialog Box Controls

There are five standard controls in most dialog boxes:

OK. Runs the procedure. After you select your variables and choose any additional specifications, click OK to run the procedure and close the dialog box.

Paste. Generates command syntax from the dialog box selections and pastes the syntax into a syntax window. You can then customize the commands with additional features that are not available from dialog boxes.

Reset. Deselects any variables in the selected variable list(s) and resets all specifications in the dialog box and any subdialog boxes to the default state.

Cancel. Cancels any changes that were made in the dialog box settings since the last time it was opened and closes the dialog box. Within a session, dialog box settings are persistent. A dialog box retains your last set of specifications until you override them.

Help. Provides context-sensitive Help. This control takes you to a Help window that contains information about the current dialog box.

Selecting Variables












To select a single variable, simply select it in the source variable list and drag and drop it into the target variable list. You can also use arrow button to move variables from the source list to the target lists. If there is only one target variable list, you can double-click individual variables to move them from the source list to the target list.

You can also select multiple variables:

- To select multiple variables that are grouped together in the variable list, click the first variable and then Shift-click the last variable in the group.
- To select multiple variables that are not grouped together in the variable list, click the first variable, then Ctrl-click the next variable, and so on (Macintosh: Command-click).

Data Type, Measurement Level, and Variable List Icons

The icons that are displayed next to variables in dialog box lists provide information about the variable type and measurement level.

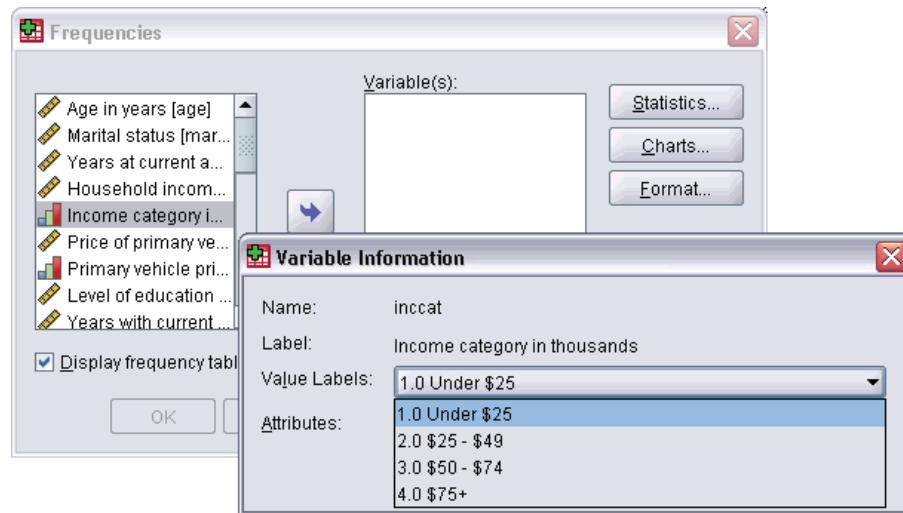
Measurement Level	Data Type			
	Numeric	String	Date	Time
Scale		n/a		
Ordinal				
Nominal				

- For more information on measurement level, see [Variable Measurement Level](#) on p. 73.
- For more information on numeric, string, date, and time data types, see [Variable Type](#) on p. 74.

Getting Information about Variables in Dialog Boxes

- ▶ Right-click a variable in the source or target variable list.
- ▶ Choose Variable Information.

Figure 1-4
Variable information



Basic Steps in Data Analysis

Analyzing data with SPSS Statistics is easy. All you have to do is:

Get your data into SPSS Statistics. You can open a previously saved SPSS Statistics data file, you can read a spreadsheet, database, or text data file, or you can enter your data directly in the Data Editor.

Select a procedure. Select a procedure from the menus to calculate statistics or to create a chart.

Select the variables for the analysis. The variables in the data file are displayed in a dialog box for the procedure.

Run the procedure and look at the results. Results are displayed in the Viewer.

Statistics Coach

If you are unfamiliar with SPSS Statistics or with the available statistical procedures, the Statistics Coach can help you get started by prompting you with simple questions, nontechnical language, and visual examples that help you select the basic statistical and charting features that are best suited for your data.

To use the Statistics Coach, from the menus in any SPSS Statistics window choose:

- Help
- Statistics Coach

The Statistics Coach covers only a selected subset of procedures in the Base system. It is designed to provide general assistance for many of the basic, commonly used statistical techniques.

Finding Out More

For a comprehensive overview of the basics, see the online tutorial. From any SPSS Statistics menu choose:

- Help
- Tutorial

Getting Help

Help is provided in many different forms:

Help menu. The Help menu in most windows provides access to the main Help system, plus tutorials and technical reference material.

- **Topics.** Provides access to the Contents, Index, and Search tabs, which you can use to find specific Help topics.
- **Tutorial.** Illustrated, step-by-step instructions on how to use many of the basic features. You don't have to view the whole tutorial from start to finish. You can choose the topics you want to view, skip around and view topics in any order, and use the index or table of contents to find specific topics.
- **Case Studies.** Hands-on examples of how to create various types of statistical analyses and how to interpret the results. The sample data files used in the examples are also provided so that you can work through the examples to see exactly how the results were produced. You can choose the specific procedure(s) that you want to learn about from the table of contents or search for relevant topics in the index.
- **Statistics Coach.** A wizard-like approach to guide you through the process of finding the procedure that you want to use. After you make a series of selections, the Statistics Coach opens the dialog box for the statistical, reporting, or charting procedure that meets your selected criteria. The Statistics Coach provides access to most statistical and reporting procedures in the Base system and many charting procedures.
- **Command Syntax Reference.** Detailed command syntax reference information is available in two forms: integrated into the overall Help system and as a separate document in PDF form in the *Command Syntax Reference*, available from the Help menu.
- **Statistical Algorithms.** The algorithms used for most statistical procedures are available in two forms: integrated into the overall Help system and as a separate document in PDF form available on the manuals CD. For links to specific algorithms in the Help system, choose Algorithms from the Help menu.

Context-sensitive Help. In many places in the user interface, you can get context-sensitive Help.

- **Dialog box Help buttons.** Most dialog boxes have a Help button that takes you directly to a Help topic for that dialog box. The Help topic provides general information and links to related topics.

- **Pivot table context menu Help.** Right-click on terms in an activated pivot table in the Viewer and choose What's This? from the context menu to display definitions of the terms.
- **Command syntax.** In a command syntax window, position the cursor anywhere within a syntax block for a command and press F1 on the keyboard. A complete command syntax chart for that command will be displayed. Complete command syntax documentation is available from the links in the list of related topics and from the Help Contents tab.

Other Resources

Technical Support Web site. Answers to many common problems can be found at <http://support.spss.com>. (The Technical Support Web site requires a login ID and password. Information on how to obtain an ID and password is provided at the URL listed above.)

Developer Central. Developer Central has resources for all levels of users and application developers. Download utilities, graphics examples, new statistical modules, and articles. Visit Developer Central at <http://www.spss.com/devcentral>.

Getting Help on Output Terms

To see a definition for a term in pivot table output in the Viewer:

- ▶ Double-click the pivot table to activate it.
- ▶ Right-click on the term that you want explained.
- ▶ Choose What's This? from the context menu.

A definition of the term is displayed in a pop-up window.

Figure 2-1

Activated pivot table glossary Help with right mouse button

The screenshot shows a pivot table with the following data:

				Gender		Total
				Female	Male	
Internet No	Owns PDA	No	Count	1897	1962	3859
			Expected Count	1903.4	1955.6	3859.0
	Total		% within Owns PDA	50.3%	49.7%	100.0%
			Count	2224	2285	4509
			Expected Count	2224.0	2285.0	4509.0
			% within Owns PDA	49.3%	50.7%	100.0%
Internet Yes	Owns PDA	No	Count	513	512	1025
			Expected Count	512.5	512.5	1025.0
	Total		% within Owns PDA	50.0%	50.0%	100.0%
			Count	305	306	611
			Expected Count	305.5	305.5	611.0
			% within Owns PDA	49.9%	50.1%	100.0%

Data Files

Data files come in a wide variety of formats, and this software is designed to handle many of them, including:

- Spreadsheets created with Excel and Lotus
- Database tables from many database sources, including Oracle, SQLServer, Access, dBASE, and others
- Tab-delimited and other types of simple text files
- Data files in SPSS Statistics format created on other operating systems
- SYSTAT data files
- SAS data files
- Stata data files

Opening Data Files

In addition to files saved in SPSS Statistics format, you can open Excel, SAS, Stata, tab-delimited, and other files without converting the files to an intermediate format or entering data definition information.

- Opening a data file makes it the active dataset. If you already have one or more open data files, they remain open and available for subsequent use in the session. Clicking anywhere in the Data Editor window for an open data file will make it the active dataset. [For more information, see Working with Multiple Data Sources in Chapter 6 on p. 93.](#)
- In distributed analysis mode using a remote server to process commands and run procedures, the available data files, folders, and drives are dependent on what is available on or from the remote server. The current server name is indicated at the top of the dialog box. You will not have access to data files on your local computer unless you specify the drive as a shared device and the folders containing your data files as shared folders. [For more information, see Distributed Analysis Mode in Chapter 4 on p. 64.](#)

To Open Data Files

- ▶ From the menus choose:
 - File
 - Open
 - Data...
- ▶ In the Open Data dialog box, select the file that you want to open.

- ▶ Click Open.

Optionally, you can:

- Automatically set the width of each string variable to the longest observed value for that variable using Minimize string widths based on observed values. This is particularly useful when reading code page data files in Unicode mode. [For more information, see General Options in Chapter 48 on p. 494.](#)
- Read variable names from the first row of spreadsheet files.
- Specify a range of cells to read from spreadsheet files.
- Specify a worksheet within an Excel file to read (Excel 95 or later).

For information on reading data from databases, see [Reading Database Files on p. 15](#). For information on reading data from text data files, see [Text Wizard on p. 28](#).

Data File Types

SPSS Statistics. Opens data files saved in SPSS Statistics format and also the DOS product SPSS/PC+.

SPSS/PC+. Opens SPSS/PC+ data files.

SYSTAT. Opens SYSTAT data files.

SPSS Statistics Portable. Opens data files saved in portable format. Saving a file in portable format takes considerably longer than saving the file in SPSS Statistics format.

Excel. Opens Excel files.

Lotus 1-2-3. Opens data files saved in 1-2-3 format for release 3.0, 2.0, or 1A of Lotus.

SYLK. Opens data files saved in SYLK (symbolic link) format, a format used by some spreadsheet applications.

dBASE. Opens dBASE-format files for either dBASE IV, dBASE III or III PLUS, or dBASE II. Each case is a record. Variable and value labels and missing-value specifications are lost when you save a file in this format.

SAS. SAS versions 6–9 and SAS transport files.

Stata. Stata versions 4–8.

Opening File Options

Read variable names. For spreadsheets, you can read variable names from the first row of the file or the first row of the defined range. The values are converted as necessary to create valid variable names, including converting spaces to underscores.

Worksheet. Excel 95 or later files can contain multiple worksheets. By default, the Data Editor reads the first worksheet. To read a different worksheet, select the worksheet from the drop-down list.

Range. For spreadsheet data files, you can also read a range of cells. Use the same method for specifying cell ranges as you would with the spreadsheet application.

Reading Excel 95 or Later Files

The following rules apply to reading Excel 95 or later files:

Data type and width. Each column is a variable. The data type and width for each variable are determined by the data type and width in the Excel file. If the column contains more than one data type (for example, date and numeric), the data type is set to string, and all values are read as valid string values.

Blank cells. For numeric variables, blank cells are converted to the system-missing value, indicated by a period. For string variables, a blank is a valid string value, and blank cells are treated as valid string values.

Variable names. If you read the first row of the Excel file (or the first row of the specified range) as variable names, values that don't conform to variable naming rules are converted to valid variable names, and the original names are used as variable labels. If you do not read variable names from the Excel file, default variable names are assigned.

Reading Older Excel Files and Other Spreadsheets

The following rules apply to reading Excel files prior to Excel 95 and other spreadsheet data:

Data type and width. The data type and width for each variable are determined by the column width and data type of the first data cell in the column. Values of other types are converted to the system-missing value. If the first data cell in the column is blank, the global default data type for the spreadsheet (usually numeric) is used.

Blank cells. For numeric variables, blank cells are converted to the system-missing value, indicated by a period. For string variables, a blank is a valid string value, and blank cells are treated as valid string values.

Variable names. If you do not read variable names from the spreadsheet, the column letters (*A*, *B*, *C*, ...) are used for variable names for Excel and Lotus files. For SYLK files and Excel files saved in R1C1 display format, the software uses the column number preceded by the letter *C* for variable names (*C1*, *C2*, *C3*, ...).

Reading dBASE Files

Database files are logically very similar to SPSS Statistics data files. The following general rules apply to dBASE files:

- Field names are converted to valid variable names.
- Colons used in dBASE field names are translated to underscores.
- Records marked for deletion but not actually purged are included. The software creates a new string variable, *D_R*, which contains an asterisk for cases marked for deletion.

Reading Stata Files

The following general rules apply to Stata data files:

- **Variable names.** Stata variable names are converted to SPSS Statistics variable names in case-sensitive form. Stata variable names that are identical except for case are converted to valid variable names by appending an underscore and a sequential letter (*_A*, *_B*, *_C*, ..., *_Z*, *_AA*, *_AB*, ..., etc.).
- **Variable labels.** Stata variable labels are converted to SPSS Statistics variable labels.
- **Value labels.** Stata value labels are converted to SPSS Statistics value labels, except for Stata value labels assigned to “extended” missing values.
- **Missing values.** Stata “extended” missing values are converted to system-missing values.
- **Date conversion.** Stata date format values are converted to SPSS Statistics `DATE` format (d-m-y) values. Stata “time-series” date format values (weeks, months, quarters, etc.) are converted to simple numeric (F) format, preserving the original, internal integer value, which is the number of weeks, months, quarters, etc., since the start of 1960.

Reading Database Files

You can read data from any database format for which you have a database driver. In local analysis mode, the necessary drivers must be installed on your local computer. In distributed analysis mode (available with SPSS Statistics Server), the drivers must be installed on the remote server. For more information, see [Distributed Analysis Mode in Chapter 4 on p. 64](#).

To Read Database Files

- ▶ From the menus choose:
 - File
 - Open Database
 - New Query...
- ▶ Select the data source.
- ▶ If necessary (depending on the data source), select the database file and/or enter a login name, password, and other information.
- ▶ Select the table(s) and fields. For OLE DB data sources (available only on Windows operating systems), you can only select one table.
- ▶ Specify any relationships between your tables.
- ▶ Optionally:
 - Specify any selection criteria for your data.
 - Add a prompt for user input to create a parameter query.
 - Save your constructed query before running it.

To Edit Saved Database Queries

- ▶ From the menus choose:
 - File
 - Open Database
 - Edit Query...
- ▶ Select the query file (*.spq) that you want to edit.
- ▶ Follow the instructions for creating a new query.

To Read Database Files with Saved Queries

- ▶ From the menus choose:
 - File
 - Open Database
 - Run Query...
- ▶ Select the query file (*.spq) that you want to run.
- ▶ If necessary (depending on the database file), enter a login name and password.
- ▶ If the query has an embedded prompt, enter other information if necessary (for example, the quarter for which you want to retrieve sales figures).

Selecting a Data Source

Use the first screen of the Database Wizard to select the type of data source to read.

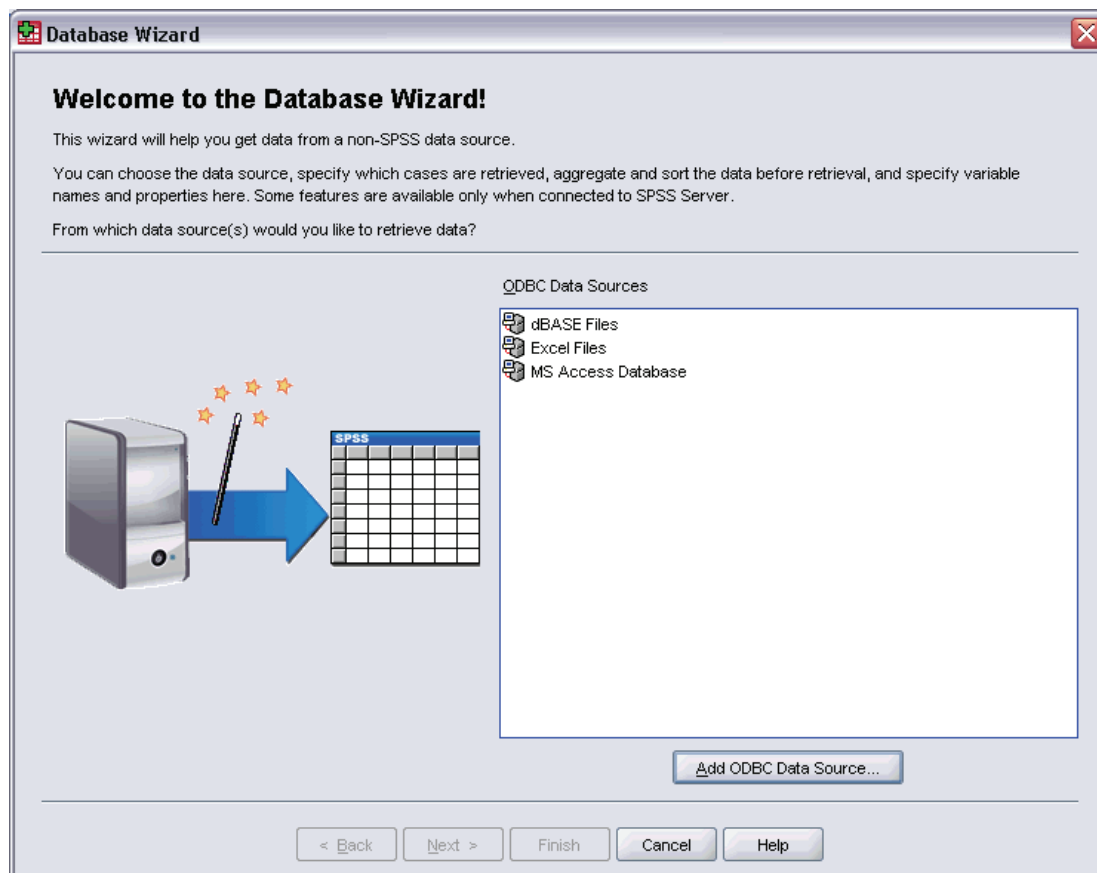
ODBC Data Sources

If you do not have any ODBC data sources configured, or if you want to add a new data source, click Add ODBC Data Source.

- On Linux operating systems, this button is not available. ODBC data sources are specified in *odbc.ini*, and the *ODBCINI* environment variables must be set to the location of that file. For more information, see the documentation for your database drivers.
- In distributed analysis mode (available with SPSS Statistics Server), this button is not available. To add data sources in distributed analysis mode, see your system administrator.

An ODBC data source consists of two essential pieces of information: the driver that will be used to access the data and the location of the database you want to access. To specify data sources, you must have the appropriate drivers installed. Drivers for a variety of database formats are available at <http://www.spss.com/drivers>.

Figure 3-1
Database Wizard



OLE DB Data Sources

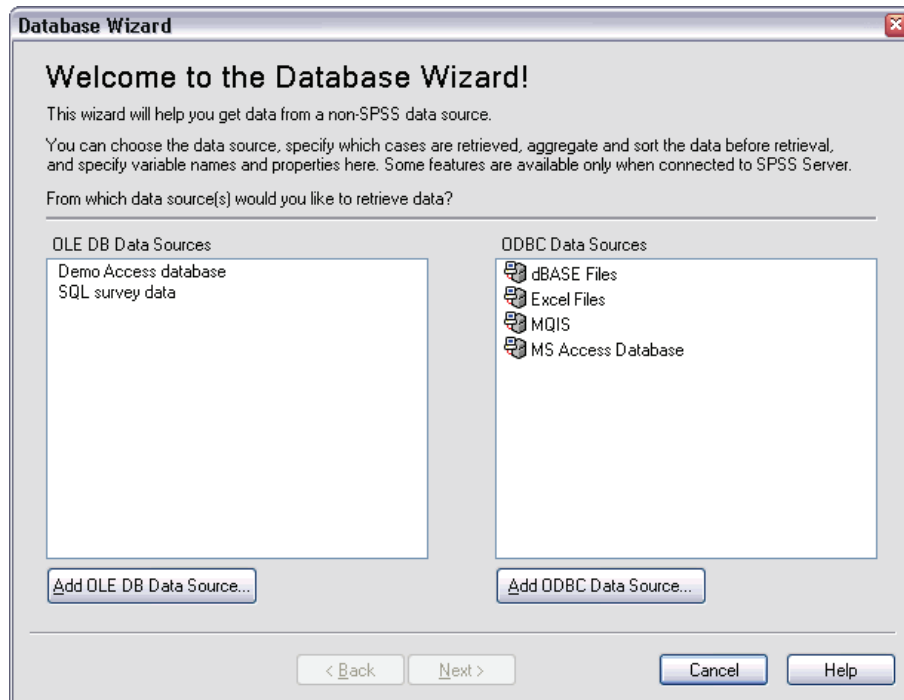
To access OLE DB data sources (available only on Microsoft Windows operating systems), you must have the following items installed:

- .NET framework. To obtain the most recent version of the .NET framework, go to <http://www.microsoft.com/net>.
- Dimensions Data Model and OLE DB Access. You can install this from the installation CD (disc 2).

The following limitations apply to OLE DB data sources:

- Table joins are not available for OLE DB data sources. You can read only one table at a time.
- You can add OLE DB data sources only in local analysis mode. To add OLE DB data sources in distributed analysis mode on a Windows server, consult your system administrator.
- In distributed analysis mode (available with SPSS Statistics Server), OLE DB data sources are available only on Windows servers, and both .NET and the Dimensions Data Model and OLE DB Access must be installed on the server.


Figure 3-2
Database Wizard with access to OLE DB data sources



To add an OLE DB data source:

- ▶ Click Add OLE DB Data Source.
- ▶ In Data Link Properties, click the Provider tab and select the OLE DB provider.
- ▶ Click Next or click the Connection tab.
- ▶ Select the database by entering the directory location and database name or by clicking the button to browse to a database. (A user name and password may also be required.)
- ▶ Click OK after entering all necessary information. (You can make sure the specified database is available by clicking the Test Connection button.)
- ▶ Enter a name for the database connection information. (This name will be displayed in the list of available OLE DB data sources.)

Figure 3-3
Save OLE DB Connection Information As dialog box



- ▶ Click OK.

This takes you back to the first screen of the Database Wizard, where you can select the saved name from the list of OLE DB data sources and continue to the next step of the wizard.

Deleting OLE DB Data Sources

To delete data source names from the list of OLE DB data sources, delete the UDL file with the name of the data source in:

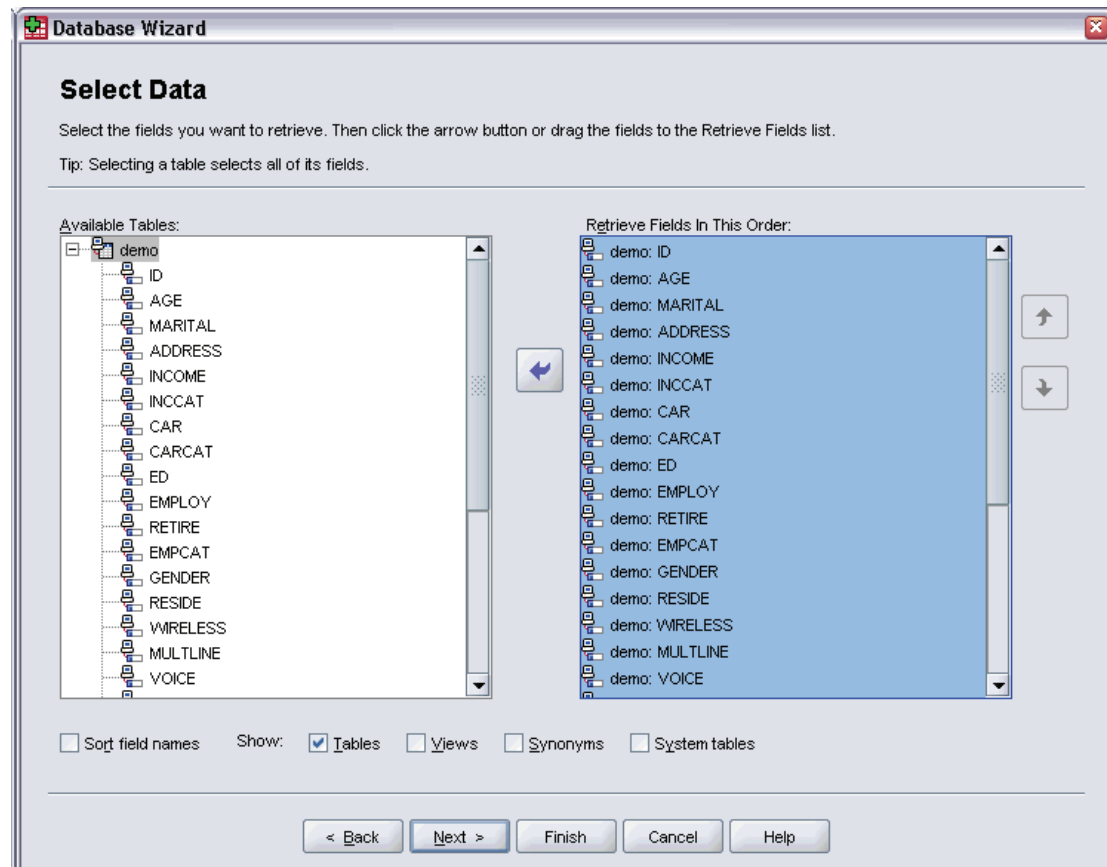
[drive]:\Documents and Settings\[user login]\Local Settings\Application Data\SPSS\UDL

Selecting Data Fields

The Select Data step controls which tables and fields are read. Database fields (columns) are read as variables.

If a table has any field(s) selected, all of its fields will be visible in the following Database Wizard windows, but only fields that are selected in this step will be imported as variables. This enables you to create table joins and to specify criteria by using fields that you are not importing.

Figure 3-4
Database Wizard, selecting data



Displaying field names. To list the fields in a table, click the plus sign (+) to the left of a table name. To hide the fields, click the minus sign (–) to the left of a table name.

To add a field. Double-click any field in the Available Tables list, or drag it to the Retrieve Fields In This Order list. Fields can be reordered by dragging and dropping them within the fields list.

To remove a field. Double-click any field in the Retrieve Fields In This Order list, or drag it to the Available Tables list.

Sort field names. If this check box is selected, the Database Wizard will display your available fields in alphabetical order.

By default, the list of available tables displays only standard database tables. You can control the type of items that are displayed in the list:

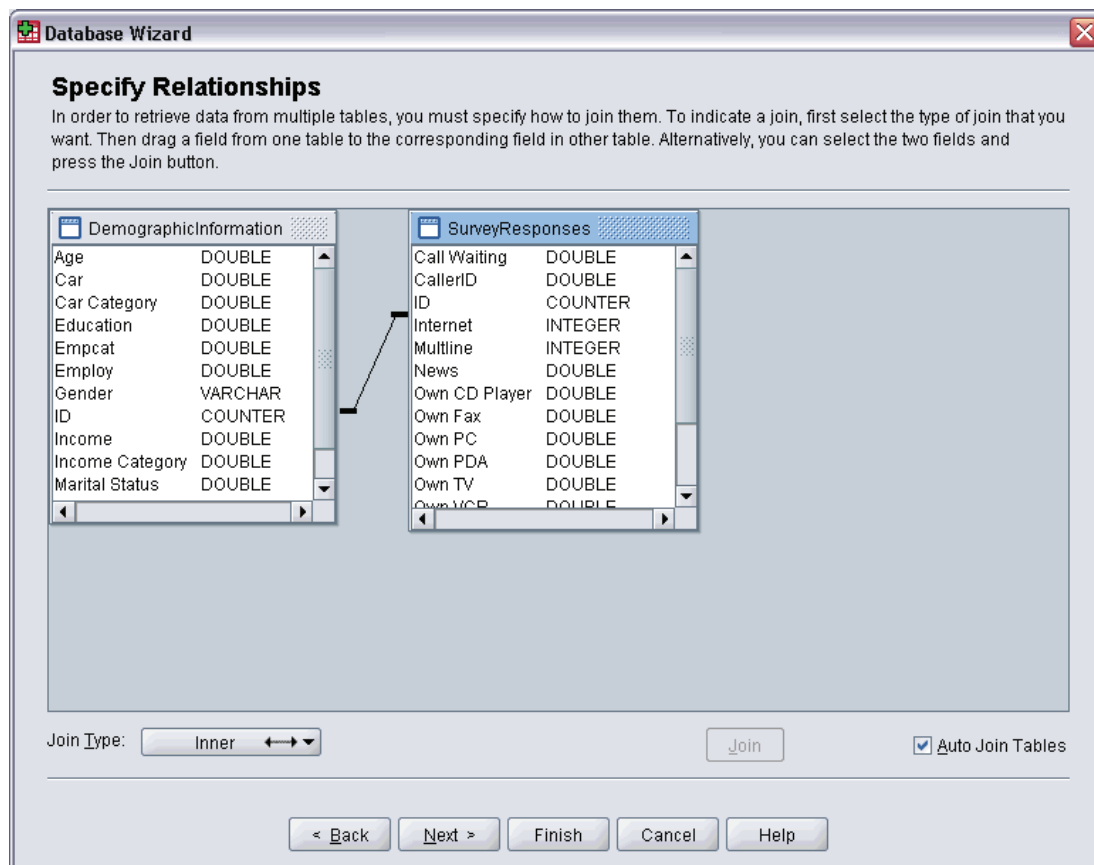
- **Tables.** Standard database tables.
- **Views.** Views are virtual or dynamic “tables” defined by queries. These can include joins of multiple tables and/or fields derived from calculations based on the values of other fields.
- **Synonyms.** A synonym is an alias for a table or view, typically defined in a query.
- **System tables.** System tables define database properties. In some cases, standard database tables may be classified as system tables and will only be displayed if you select this option. Access to real system tables is often restricted to database administrators.

Note: For OLE DB data sources (available only on Windows operating systems), you can select fields only from a single table. Multiple table joins are not supported for OLE DB data sources.

Creating a Relationship between Tables

The Specify Relationships step allows you to define the relationships between the tables for ODBC data sources. If fields from more than one table are selected, you must define at least one join.

Figure 3-5
Database Wizard, specifying relationships



Establishing relationships. To create a relationship, drag a field from any table onto the field to which you want to join it. The Database Wizard will draw a **join line** between the two fields, indicating their relationship. These fields must be of the same data type.

Auto Join Tables. Attempts to automatically join tables based on primary/foreign keys or matching field names and data type.

Join Type. If outer joins are supported by your driver, you can specify inner joins, left outer joins, or right outer joins.

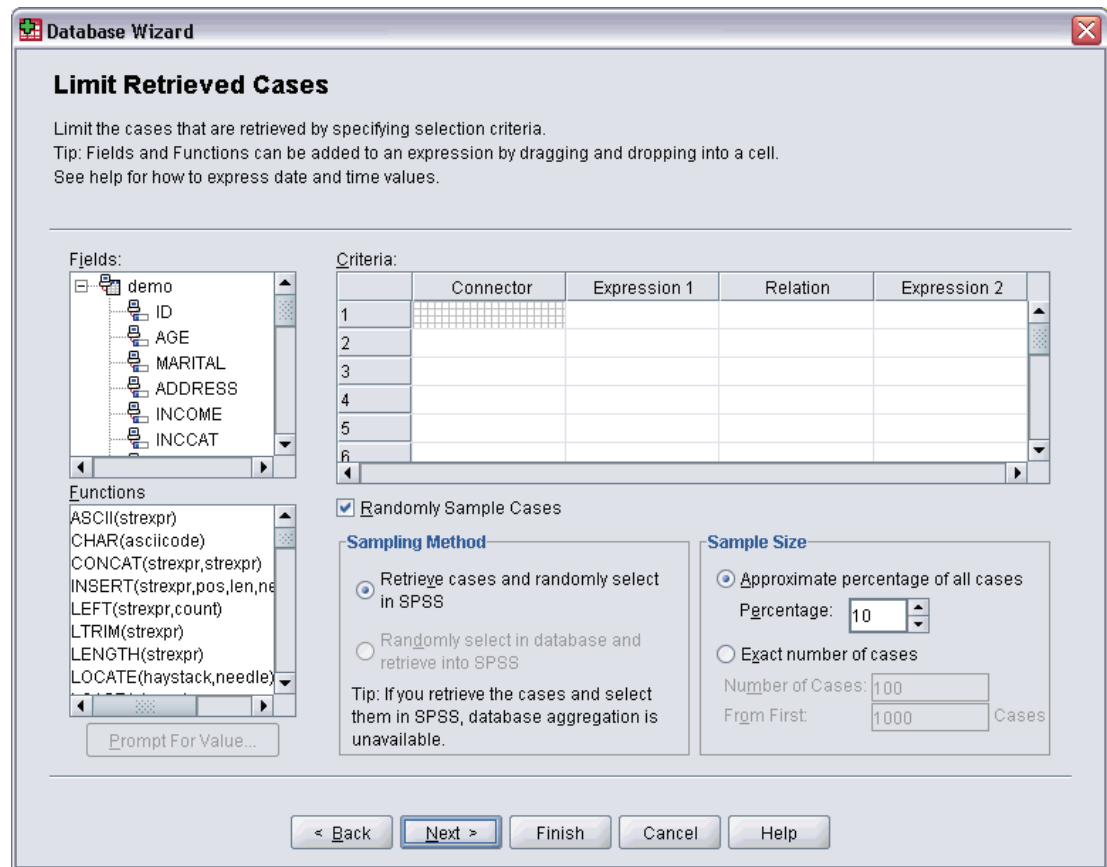
- **Inner joins.** An inner join includes only rows where the related fields are equal. In this example, all rows with matching ID values in the two tables will be included.
- **Outer joins.** In addition to one-to-one matching with inner joins, you can also use outer joins to merge tables with a one-to-many matching scheme. For example, you could match a table in which there are only a few records representing data values and associated descriptive labels with values in a table containing hundreds or thousands of records representing survey respondents. A left outer join includes all records from the table on the left and, from the table on the right, includes only those records in which the related fields are equal. In a right outer join, the join imports all records from the table on the right and, from the table on the left, imports only those records in which the related fields are equal.

Limiting Retrieved Cases

The Limit Retrieved Cases step allows you to specify the criteria to select subsets of cases (rows). Limiting cases generally consists of filling the criteria grid with criteria. Criteria consist of two expressions and some relation between them. The expressions return a value of *true*, *false*, or *missing* for each case.

- If the result is *true*, the case is selected.
- If the result is *false* or *missing*, the case is not selected.
- Most criteria use one or more of the six relational operators (<, >, <=, >=, =, and <>).
- Expressions can include field names, constants, arithmetic operators, numeric and other functions, and logical variables. You can use fields that you do not plan to import as variables.

Figure 3-6
Database Wizard, limiting retrieved cases



To build your criteria, you need at least two expressions and a relation to connect the expressions.

- ▶ To build an expression, choose one of the following methods:
 - In an Expression cell, type field names, constants, arithmetic operators, numeric and other functions, or logical variables.
 - Double-click the field in the Fields list.

- Drag the field from the Fields list to an Expression cell.
 - Choose a field from the drop-down menu in any active Expression cell.
- To choose the relational operator (such as = or >), put your cursor in the Relation cell and either type the operator or choose it from the drop-down menu.

If the SQL contains `WHERE` clauses with expressions for case selection, dates and times in expressions need to be specified in a special manner (including the curly braces shown in the examples):

- Date literals should be specified using the general form `{d 'yyyy-mm-dd'}`.
- Time literals should be specified using the general form `{t 'hh:mm:ss'}`.
- Date/time literals (timestamps) should be specified using the general form `{ts 'yyyy-mm-dd hh:mm:ss'}`.
- The entire date and/or time value must be enclosed in single quotes. Years must be expressed in four-digit form, and dates and times must contain two digits for each portion of the value. For example January 1, 2005, 1:05 AM would be expressed as:

```
{ts '2005-01-01 01:05:00'}
```

Functions. A selection of built-in arithmetic, logical, string, date, and time SQL functions is provided. You can drag a function from the list into the expression, or you can enter any valid SQL function. See your database documentation for valid SQL functions. A list of standard functions is available at:

<http://msdn2.microsoft.com/en-us/library/ms711813.aspx>

Use Random Sampling. This option selects a random sample of cases from the data source. For large data sources, you may want to limit the number of cases to a small, representative sample, which can significantly reduce the time that it takes to run procedures. Native random sampling, if available for the data source, is faster than SPSS Statistics random sampling, because SPSS Statistics random sampling must still read the entire data source to extract a random sample.

- **Approximately.** Generates a random sample of approximately the specified percentage of cases. Since this routine makes an independent pseudorandom decision for each case, the percentage of cases selected can only approximate the specified percentage. The more cases there are in the data file, the closer the percentage of cases selected is to the specified percentage.
- **Exactly.** Selects a random sample of the specified number of cases from the specified total number of cases. If the total number of cases specified exceeds the total number of cases in the data file, the sample will contain proportionally fewer cases than the requested number.

Note: If you use random sampling, aggregation (available in distributed mode with SPSS Statistics Server) is not available.

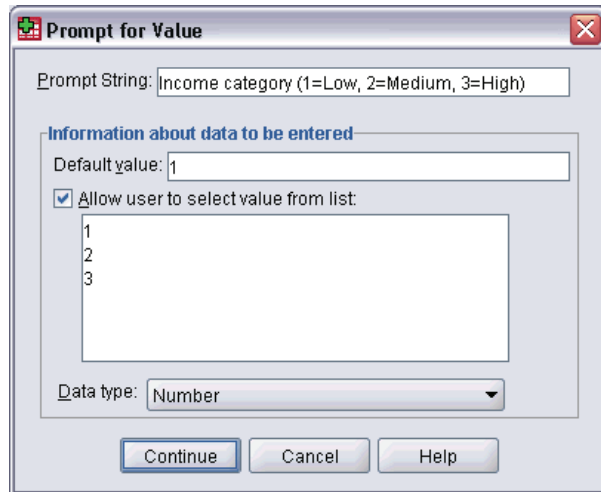
Prompt For Value. You can embed a prompt in your query to create a **parameter query**. When users run the query, they will be asked to enter information (based on what is specified here). You might want to do this if you need to see different views of the same data. For example, you may want to run the same query to see sales figures for different fiscal quarters.

- Place your cursor in any Expression cell, and click Prompt For Value to create a prompt.

Creating a Parameter Query

Use the Prompt for Value step to create a dialog box that solicits information from users each time someone runs your query. This feature is useful if you want to query the same data source by using different criteria.

Figure 3-7
Prompt for Value




To build a prompt, enter a prompt string and a default value. The prompt string is displayed each time a user runs your query. The string should specify the kind of information to enter. If the user is not selecting from a list, the string should give hints about how the input should be formatted. An example is as follows: Enter a Quarter (Q1, Q2, Q3, ...).

Allow user to select value from list. If this check box is selected, you can limit the user to the values that you place here. Ensure that your values are separated by returns.

Data type. Choose the data type here (Number, String, or Date).

The final result looks like this:

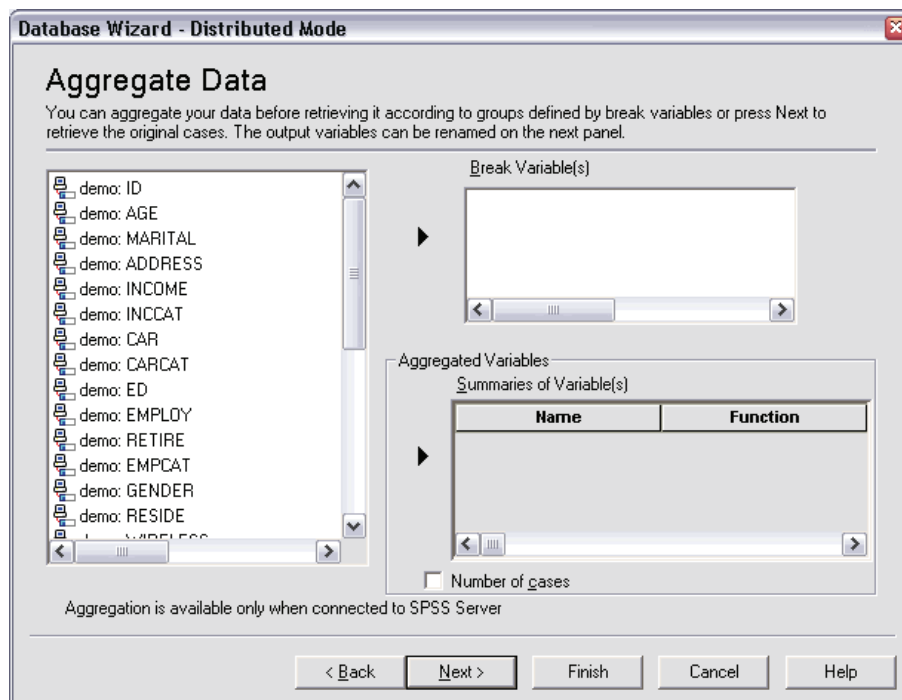
Figure 3-8
User-defined prompt



Aggregating Data

If you are in distributed mode, connected to a remote server (available with SPSS Statistics Server), you can aggregate the data before reading it into SPSS Statistics.

Figure 3-9
Database Wizard, aggregating data



You can also aggregate data after reading it into SPSS Statistics, but preaggregating may save time for large data sources.

- ▶ To create aggregated data, select one or more break variables that define how cases are grouped.
- ▶ Select one or more aggregated variables.
- ▶ Select an aggregate function for each aggregate variable.
- ▶ Optionally, create a variable that contains the number of cases in each break group.

Note: If you use SPSS Statistics random sampling, aggregation is not available.

Defining Variables

Variable names and labels. The complete database field (column) name is used as the variable label. Unless you modify the variable name, the Database Wizard assigns variable names to each column from the database in one of two ways:

- If the name of the database field forms a valid, unique variable name, the name is used as the variable name.
- If the name of the database field does not form a valid, unique variable name, a new, unique name is automatically generated.

Click any cell to edit the variable name.

Converting strings to numeric values. Select the Recode to Numeric box for a string variable if you want to automatically convert it to a numeric variable. String values are converted to consecutive integer values based on alphabetical order of the original values. The original values are retained as value labels for the new variables.

Width for variable-width string fields. This option controls the width of variable-width string values. By default, the width is 255 bytes, and only the first 255 bytes (typically 255 characters in single-byte languages) will be read. The width can be up to 32,767 bytes. Although you probably don't want to truncate string values, you also don't want to specify an unnecessarily large value, which will cause processing to be inefficient.

Minimize string widths based on observed values. Automatically set the width of each string variable to the longest observed value.

Figure 3-10
Database Wizard, defining variables

Define Variables
Edit SPSS variable names and properties.
Names will be automatically generated if left blank.
To convert a string variable to numeric using the original values as value labels, check Recode to Numeric.

	Result Variable Name	Data Type	Recode to Numeric
demo: ID	ID	Numeric	
demo: AGE	AGE	Numeric	
demo: MARITAL	MARITAL	Numeric	
demo: ADDRESS	ADDRESS	Numeric	
demo: INCOME	INCOME	Numeric	
demo: INCCAT	INCCAT	Numeric	
demo: CAR	CAR	Numeric	
demo: CARCAT	CARCAT	Numeric	
demo: ED	ED	Numeric	
demo: EMPLOY	EMPLOY	Numeric	
demo: RETIRE	RETIRE	Numeric	
demo: EMPCAT	EMPCAT	Numeric	
demo: GENDER	GENDER	String	<input checked="" type="checkbox"/>
demo: RESIDE	RESIDE	Numeric	
demo: WIRELESS	WIRELESS	Numeric	
demo: MULTLINE	MULTLINE	Numeric	
demo: VOICE	VOICE	Numeric	
demo: PAGER	PAGER	Numeric	

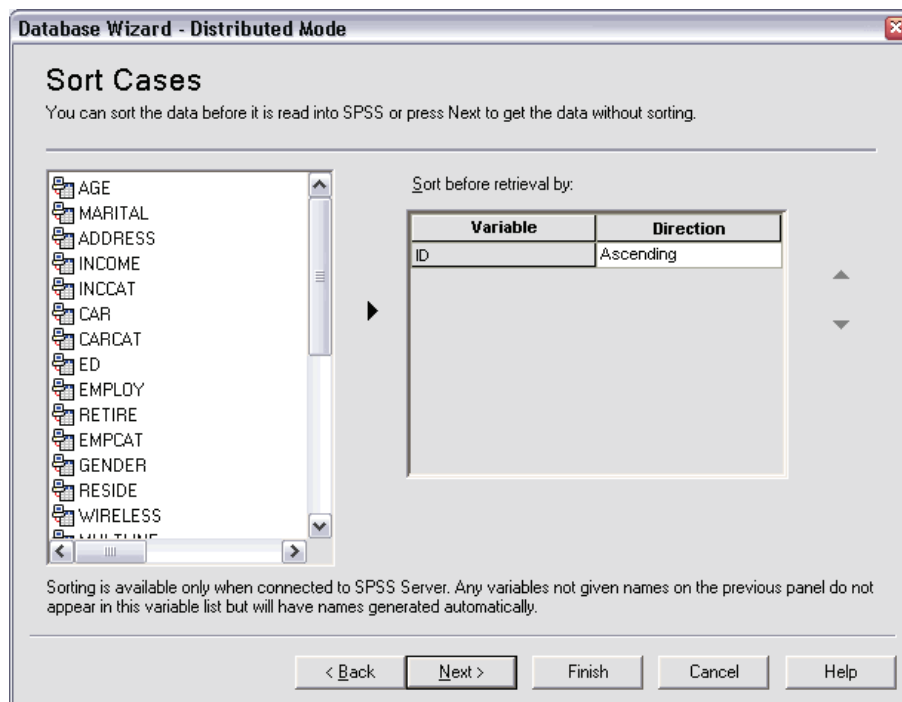
Width for variable-width string fields: Minimize string widths based on observed values

< Back Next > Finish Cancel Help

Sorting Cases

If you are in distributed mode, connected to a remote server (available with SPSS Statistics Server), you can sort the data before reading it into SPSS Statistics.

Figure 3-11
Database Wizard, sorting cases



You can also sort data after reading it into SPSS Statistics, but presorting may save time for large data sources.

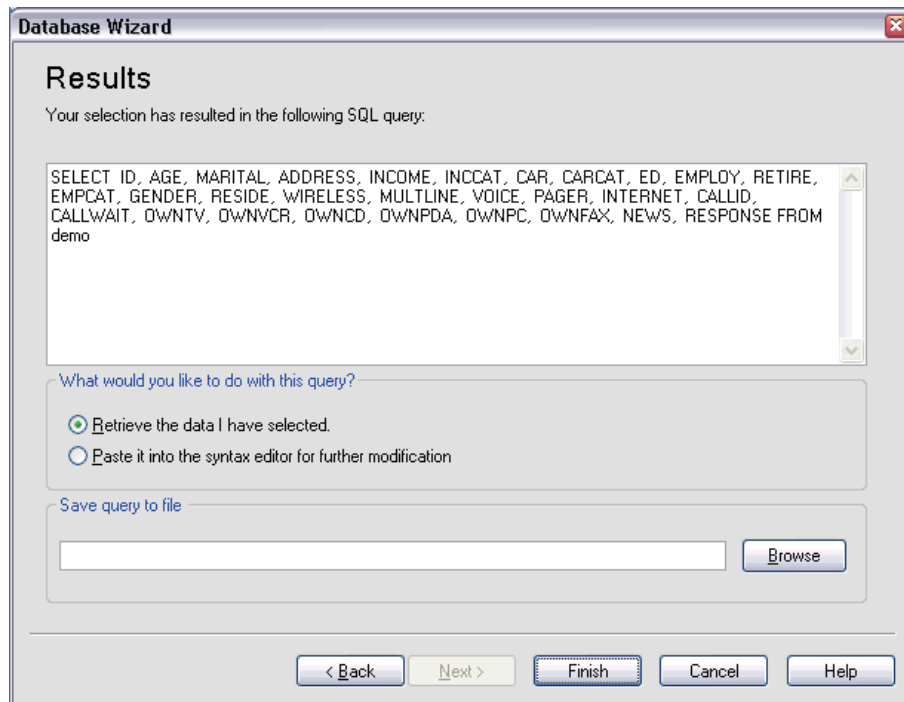
Results

The Results step displays the SQL Select statement for your query.

- You can edit the SQL Select statement before you run the query, but if you click the Back button to make changes in previous steps, the changes to the Select statement will be lost.
- To save the query for future use, use the Save query to file section.
- To paste complete GET DATA syntax into a syntax window, select Paste it into the syntax editor for further modification. Copying and pasting the Select statement from the Results window will not paste the necessary command syntax.

Note: The pasted syntax contains a blank space before the closing quote on each line of SQL that is generated by the wizard. These blanks are not superfluous. When the command is processed, all lines of the SQL statement are merged together in a very literal fashion. Without the space, there would be no space between the last character on one line and first character on the next line.

Figure 3-12
Database Wizard, results panel



Text Wizard

The Text Wizard can read text data files formatted in a variety of ways:

- Tab-delimited files
- Space-delimited files
- Comma-delimited files
- Fixed-field format files

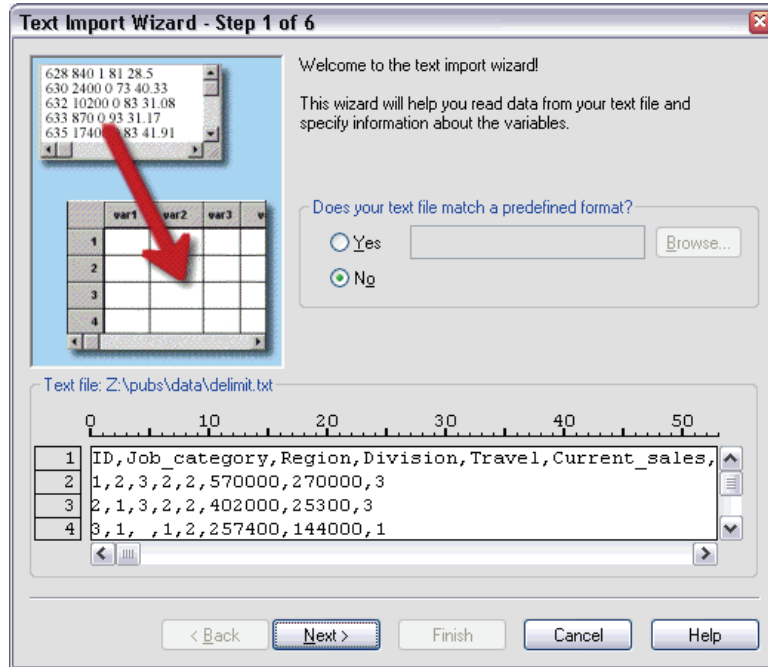
For delimited files, you can also specify other characters as delimiters between values, and you can specify multiple delimiters.

To Read Text Data Files

- ▶ From the menus choose:
 - File
 - Read Text Data...
- ▶ Select the text file in the Open Data dialog box.
- ▶ Follow the steps in the Text Wizard to define how to read the data file.

Text Wizard: Step 1

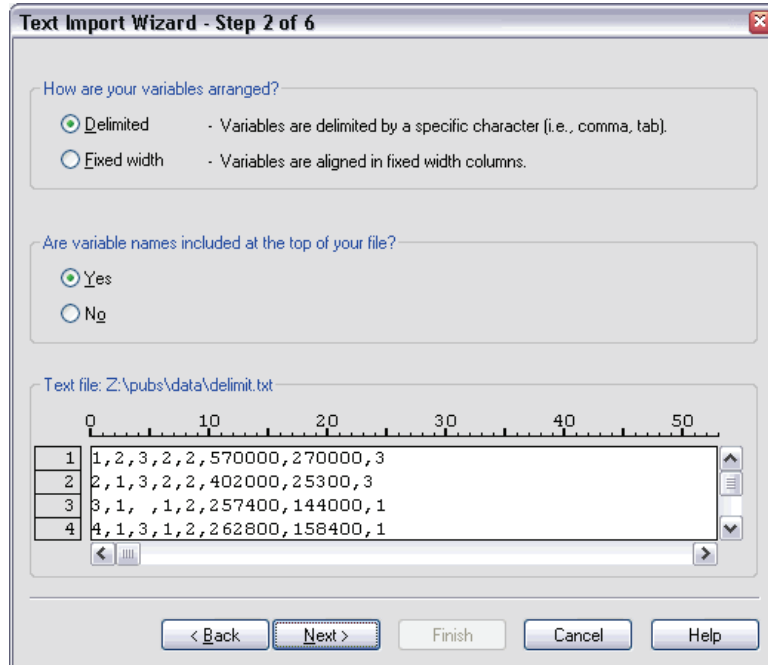
Figure 3-13
Text Wizard: Step 1



The text file is displayed in a preview window. You can apply a predefined format (previously saved from the Text Wizard) or follow the steps in the Text Wizard to specify how the data should be read.

Text Wizard: Step 2

Figure 3-14
Text Wizard: Step 2



This step provides information about variables. A variable is similar to a field in a database. For example, each item in a questionnaire is a variable.

How are your variables arranged? To read your data properly, the Text Wizard needs to know how to determine where the data value for one variable ends and the data value for the next variable begins. The arrangement of variables defines the method used to differentiate one variable from the next.

- **Delimited.** Spaces, commas, tabs, or other characters are used to separate variables. The variables are recorded in the same order for each case but not necessarily in the same column locations.
- **Fixed width.** Each variable is recorded in the same column location on the same record (line) for each case in the data file. No delimiter is required between variables. In fact, in many text data files generated by computer programs, data values may appear to run together without even spaces separating them. The column location determines which variable is being read.

Are variable names included at the top of your file? If the first row of the data file contains descriptive labels for each variable, you can use these labels as variable names. Values that don't conform to variable naming rules are converted to valid variable names.

Text Wizard: Step 3 (Delimited Files)

Figure 3-15
Text Wizard: Step 3 (for delimited files)

This step provides information about cases. A case is similar to a record in a database. For example, each respondent to a questionnaire is a case.

The first case of data begins on which line number? Indicates the first line of the data file that contains data values. If the top line(s) of the data file contain descriptive labels or other text that does not represent data values, this will *not* be line 1.

How are your cases represented? Controls how the Text Wizard determines where each case ends and the next one begins.

- **Each line represents a case.** Each line contains only one case. It is fairly common for each case to be contained on a single line (row), even though this can be a very long line for data files with a large number of variables. If not all lines contain the same number of data values, the number of variables for each case is determined by the line with the greatest number of data values. Cases with fewer data values are assigned missing values for the additional variables.
- **A specific number of variables represents a case.** The specified number of variables for each case tells the Text Wizard where to stop reading one case and start reading the next. Multiple cases can be contained on the same line, and cases can start in the middle of one line and be continued on the next line. The Text Wizard determines the end of each case based on the number of values read, regardless of the number of lines. Each case must contain data values (or missing values indicated by delimiters) for all variables, or the data file will be read incorrectly.

How many cases do you want to import? You can import all cases in the data file, the first n cases (n is a number you specify), or a random sample of a specified percentage. Since the random sampling routine makes an independent pseudo-random decision for each case, the percentage of cases selected can only approximate the specified percentage. The more cases there are in the data file, the closer the percentage of cases selected is to the specified percentage.

Text Wizard: Step 3 (Fixed-Width Files)

Figure 3-16
Text Wizard: Step 3 (for fixed-width files)

The screenshot shows the 'Text Import Wizard - Fixed Width Step 3 of 6' dialog box. It contains the following fields and options:

- 'The first case of data begins on which line number?' with a spinner box set to 1.
- 'How many lines represent a case?' with a spinner box set to 1.
- 'How many cases do you want to import?' with three radio button options:
 - All of the cases
 - The first 1000 cases.
 - A percentage of the cases: 10 %
- 'Data preview' section with a horizontal axis from 0 to 50 and a table showing the first four lines of data.

1	12	3.0022	\$570,000	\$270,000	3
2	21	3.0022	\$402,000	\$25,300	3
3	31	12	\$257,400	\$144,000	1
4	41	3.0012	\$262,800	\$158,400	1

At the bottom of the dialog are buttons for '< Back', 'Next >', 'Finish', 'Cancel', and 'Help'.

This step provides information about cases. A case is similar to a record in a database. For example, each respondent to questionnaire is a case.

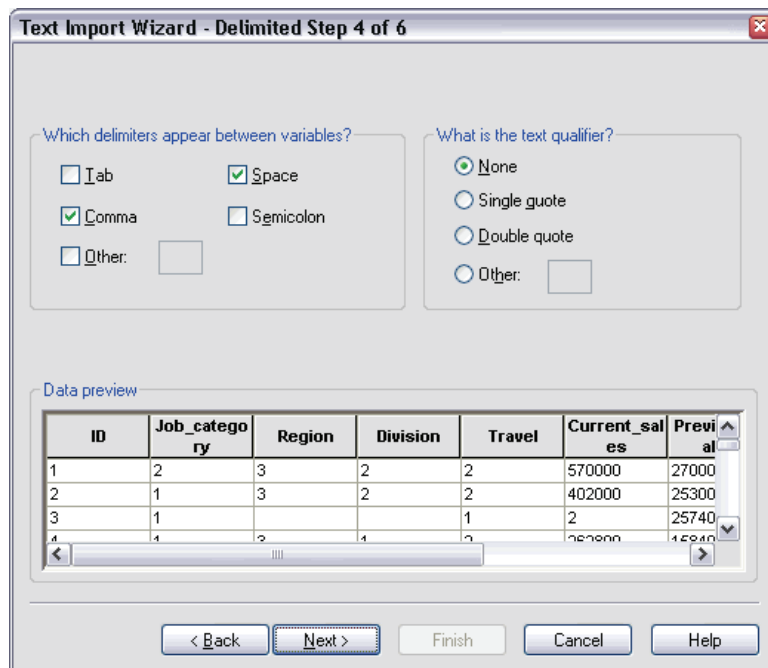
The first case of data begins on which line number? Indicates the first line of the data file that contains data values. If the top line(s) of the data file contain descriptive labels or other text that does not represent data values, this will *not* be line 1.

How many lines represent a case? Controls how the Text Wizard determines where each case ends and the next one begins. Each variable is defined by its line number within the case and its column location. You need to specify the number of lines for each case to read the data correctly.

How many cases do you want to import? You can import all cases in the data file, the first n cases (n is a number you specify), or a random sample of a specified percentage. Since the random sampling routine makes an independent pseudo-random decision for each case, the percentage of cases selected can only approximate the specified percentage. The more cases there are in the data file, the closer the percentage of cases selected is to the specified percentage.

Text Wizard: Step 4 (Delimited Files)

Figure 3-17
Text Wizard: Step 4 (for delimited files)



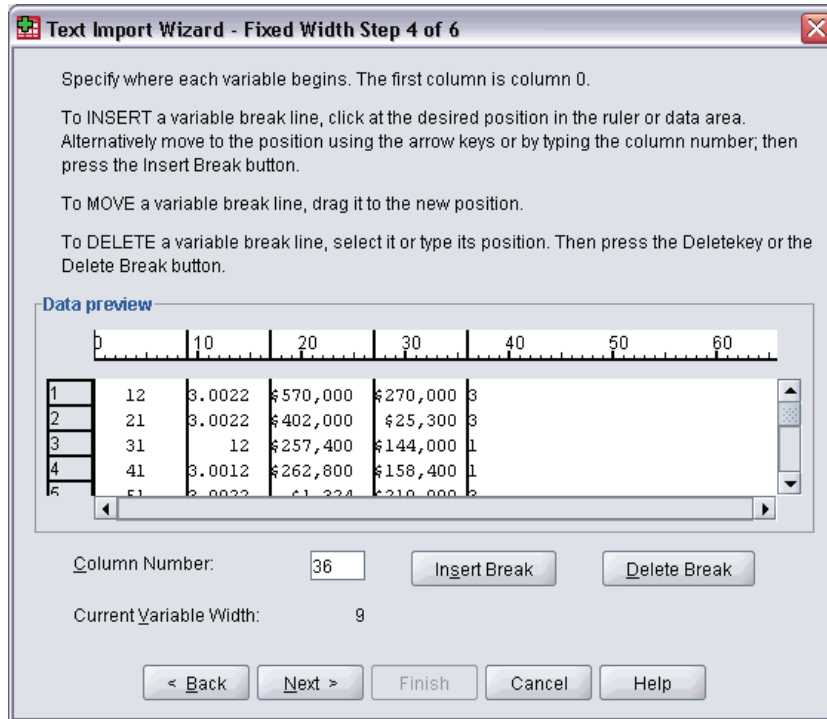
This step displays the Text Wizard's best guess on how to read the data file and allows you to modify how the Text Wizard will read variables from the data file.

Which delimiters appear between variables? Indicates the characters or symbols that separate data values. You can select any combination of spaces, commas, semicolons, tabs, or other characters. Multiple, consecutive delimiters without intervening data values are treated as missing values.

What is the text qualifier? Characters used to enclose values that contain delimiter characters. For example, if a comma is the delimiter, values that contain commas will be read incorrectly unless there is a text qualifier enclosing the value, preventing the commas in the value from being interpreted as delimiters between values. CSV-format data files exported from Excel use a double quotation mark (") as a text qualifier. The text qualifier appears at both the beginning and the end of the value, enclosing the entire value.

Text Wizard: Step 4 (Fixed-Width Files)

Figure 3-18
Text Wizard: Step 4 (for fixed-width files)



This step displays the Text Wizard's best guess on how to read the data file and allows you to modify how the Text Wizard will read variables from the data file. Vertical lines in the preview window indicate where the Text Wizard currently thinks each variable begins in the file.

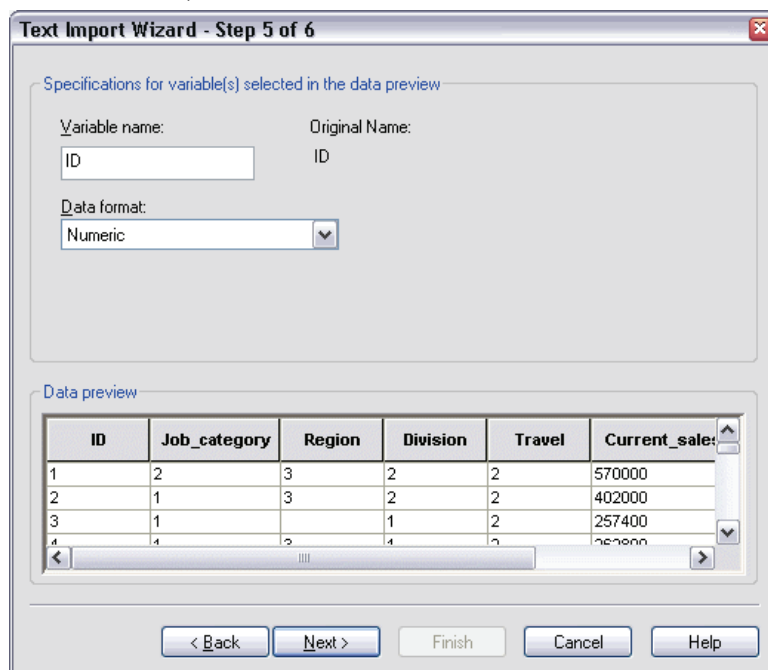
Insert, move, and delete variable break lines as necessary to separate variables. If multiple lines are used for each case, the data will be displayed as one line for each case, with subsequent lines appended to the end of the line.

Notes:

For computer-generated data files that produce a continuous stream of data values with no intervening spaces or other distinguishing characteristics, it may be difficult to determine where each variable begins. Such data files usually rely on a data definition file or some other written description that specifies the line and column location for each variable.

Text Wizard: Step 5

Figure 3-19
Text Wizard: Step 5



This step controls the variable name and the data format that the Text Wizard will use to read each variable and which variables will be included in the final data file.

Variable name. You can overwrite the default variable names with your own variable names. If you read variable names from the data file, the Text Wizard will automatically modify variable names that don't conform to variable naming rules. Select a variable in the preview window and then enter a variable name.

Data format. Select a variable in the preview window and then select a format from the drop-down list. Shift-click to select multiple contiguous variables or Ctrl-click to select multiple noncontiguous variables.

Text Wizard Formatting Options

Formatting options for reading variables with the Text Wizard include:

Do not import. Omit the selected variable(s) from the imported data file.

Numeric. Valid values include numbers, a leading plus or minus sign, and a decimal indicator.

String. Valid values include virtually any keyboard characters and embedded blanks. For delimited files, you can specify the number of characters in the value, up to a maximum of 32,767. By default, the Text Wizard sets the number of characters to the longest string value encountered for the selected variable(s). For fixed-width files, the number of characters in string values is defined by the placement of variable break lines in step 4.

Date/Time. Valid values include dates of the general format *dd-mm-yyyy*, *mm/dd/yyyy*, *dd.mm.yyyy*, *yyyy/mm/dd*, *hh:mm:ss*, and a variety of other date and time formats. Months can be represented in digits, Roman numerals, or three-letter abbreviations, or they can be fully spelled out. Select a date format from the list.

Dollar. Valid values are numbers with an optional leading dollar sign and optional commas as thousands separators.

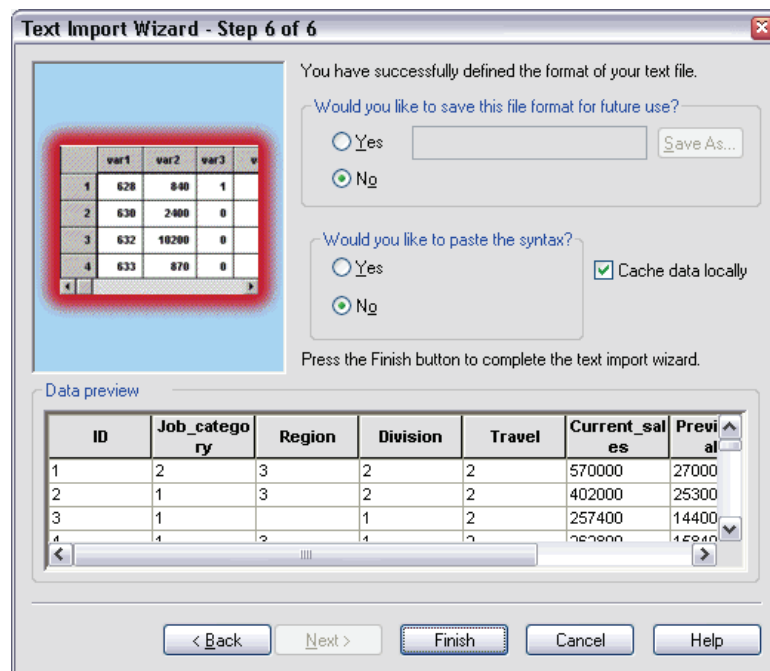
Comma. Valid values include numbers that use a period as a decimal indicator and commas as thousands separators.

Dot. Valid values include numbers that use a comma as a decimal indicator and periods as thousands separators.

Note: Values that contain invalid characters for the selected format will be treated as missing. Values that contain any of the specified delimiters will be treated as multiple values.

Text Wizard: Step 6

Figure 3-20
Text Wizard: Step 6



This is the final step of the Text Wizard. You can save your specifications in a file for use when importing similar text data files. You can also paste the syntax generated by the Text Wizard into a syntax window. You can then customize and/or save the syntax for use in other sessions or in production jobs.

Cache data locally. A data cache is a complete copy of the data file, stored in temporary disk space. Caching the data file can improve performance.

Reading Dimensions Data

On Microsoft Windows operating systems, you can read data from Dimensions products, including Quanvert, Quancept, and mrInterview. (*Note:* This feature is only available with SPSS Statistics installed on Microsoft Windows operating systems.)

To read Dimensions data sources, you must have the following items installed:

- .NET framework
- Dimensions Data Model and OLE DB Access

Versions of these components that are compatible with this release can be installed from the installation CD and are available on the AutoPlay menu. You can read Dimensions data sources only in local analysis mode. This feature is not available in distributed analysis mode using SPSS Statistics Server.

To read data from a Dimensions data source:

- ▶ In any open SPSS Statistics window, from the menus choose:
 - File
 - Open Dimensions Data
- ▶ On the Connection tab of Data Link Properties, specify the metadata file, the case data type, and the case data file.
- ▶ Click OK.
- ▶ In the Dimensions Data Import dialog box, select the variables that you want to include and select any case selection criteria.
- ▶ Click OK to read the data.

Data Link Properties Connection Tab

To read a Dimensions data source, you need to specify:

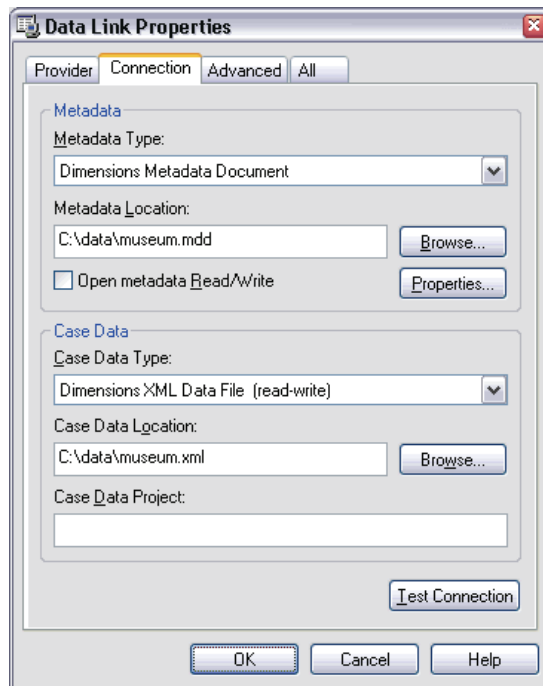
Metadata Location. The metadata document file (*.mdd*) that contains questionnaire definition information.

Case Data Type. The format of the case data file. Available formats include:

- Quancept Data File (DRS). Case data in a Quancept *.drs*, *.drz*, or *.dru* file.
- Quanvert Database. Case data in a Quanvert database.
- Dimensions Database (MS SQL Server). Case data in a relational database in SQL Server. This option can be used to read data that are collected by using mrInterview.
- Dimensions XML Data File. Case data in an XML file.

Case Data Location. The file that contains the case data. The format of this file must be consistent with the selected case data type.

Figure 3-21
Data Link Properties: Connection tab



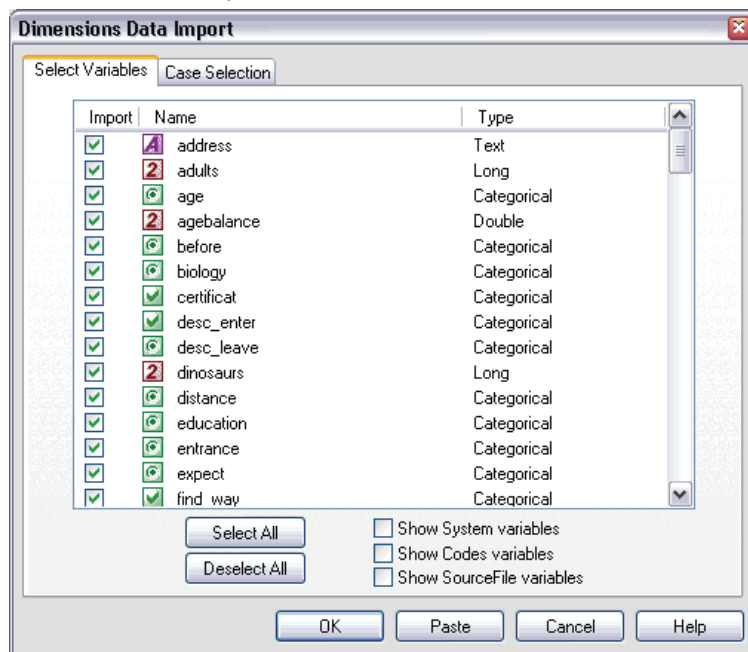
Note: The extent to which other settings on the Connection tab or any settings on the other Data Link Properties tabs may or may not affect reading Dimensions data into SPSS Statistics is not known, so we recommend that you do not change any of them.

Select Variables Tab

You can select a subset of variables to read. By default, all standard variables in the data source are displayed and selected.

- **Show System variables.** Displays any “system” variables, including variables that indicate interview status (*in progress*, *completed*, *finish date*, and so on). You can then select any system variables that you want to include. By default, all system variables are excluded.
- **Show Codes variables.** Displays any variables that represent codes that are used for open-ended “Other” responses for categorical variables. You can then select any Codes variables that you want to include. By default, all Codes variables are excluded.
- **Show SourceFile variables.** Displays any variables that contain filenames of images of scanned responses. You can then select any SourceFile variables that you want to include. By default, all SourceFile variables are excluded.

Figure 3-22
Dimensions Data Import: Select Variables tab



Case Selection Tab

For Dimensions data sources that contain system variables, you can select cases based on a number of system variable criteria. You do not need to include the corresponding system variables in the list of variables to read, but the necessary system variables must exist in the source data to apply the selection criteria. If the necessary system variables do not exist in the source data, the corresponding selection criteria are ignored.

Data collection status. You can select respondent data, test data, or both. You can also select cases based on any combination of the following interview status parameters:

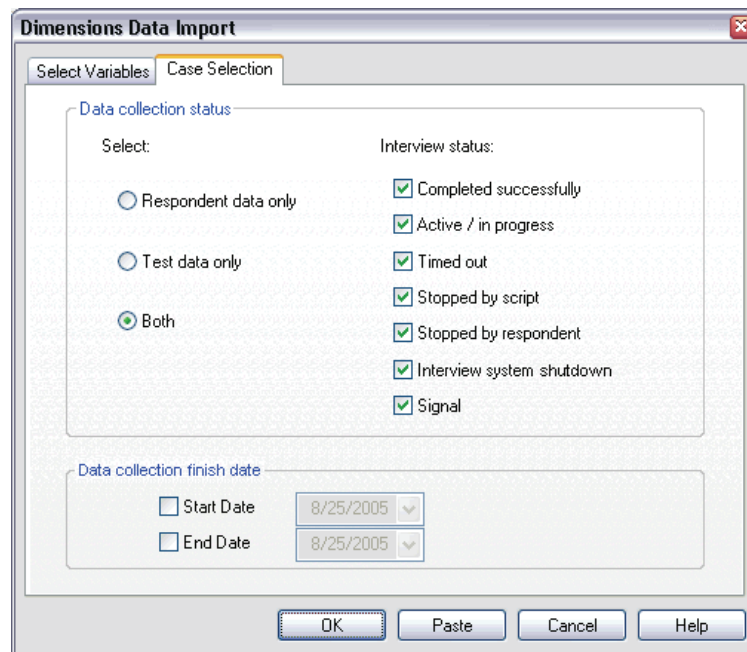
- Completed successfully
- Active/in progress
- Timed out
- Stopped by script
- Stopped by respondent
- Interview system shutdown
- Signal (terminated by a signal statement in the script)

Data collection finish date. You can select cases based on the data collection finish date.

- **Start Date.** Cases for which data collection finished on or after the specified date are included.

- **End Date.** Cases for which data collection finished before the specified date are included. This does *not* include cases for which data collection finished on the end date.
- If you specify both a start date and end date, this defines a range of finish dates from the start date to (but not including) the end date.

Figure 3-23
Dimensions Data Import: Case Selection tab



File Information

A data file contains much more than raw data. It also contains any variable definition information, including:

- Variable names
- Variable formats
- Descriptive variable and value labels

This information is stored in the dictionary portion of the data file. The Data Editor provides one way to view the variable definition information. You can also display complete dictionary information for the active dataset or any other data file.

To Display Data File Information

- ▶ From the menus in the Data Editor window choose:
 - File
 - Display Data File Information
- ▶ For the currently open data file, choose Working File.
- ▶ For other data files, choose External File, and then select the data file.

The data file information is displayed in the Viewer.

Saving Data Files

In addition to saving data files in SPSS Statistics format, you can save data in a wide variety of external formats, including:

- Excel and other spreadsheet formats
- Tab-delimited and CSV text files
- SAS
- Stata
- Database tables

To Save Modified Data Files

- ▶ Make the Data Editor the active window (click anywhere in the window to make it active).
- ▶ From the menus choose:
 - File
 - Save

The modified data file is saved, overwriting the previous version of the file.

Note: A data file saved in Unicode mode cannot be read by versions of SPSS Statistics prior to 16.0. To save a Unicode data file in a format that can be read by earlier releases, open the file in code page mode and re-save it. The file will be saved in the encoding based on the current locale. Some data loss may occur if the file contains characters not recognized by the current locale. For information on switching between Unicode mode and code page mode, see [General Options](#) on p. 494.

Saving Data Files in External Formats

- ▶ Make the Data Editor the active window (click anywhere in the window to make it active).
- ▶ From the menus choose:
 - File
 - Save As...
- ▶ Select a file type from the drop-down list.
- ▶ Enter a filename for the new data file.

To write variable names to the first row of a spreadsheet or tab-delimited data file:

- ▶ Click Write variable names to spreadsheet in the Save Data As dialog box.

To save value labels instead of data values in Excel files:

- ▶ Click Save value labels where defined instead of data values in the Save Data As dialog box.

To save value labels to a SAS syntax file (active only when a SAS file type is selected):

- ▶ Click Save value labels into a .sas file in the Save Data As dialog box.

For information on exporting data to database tables, see [Exporting to a Database on p. 47](#).

For information on exporting data for use in Dimensions applications, see [Exporting to Dimensions on p. 59](#).

Saving Data: Data File Types

You can save data in the following formats:

SPSS Statistics (*.sav). SPSS Statistics format.

- Data files saved in SPSS Statistics format cannot be read by versions of the software prior to version 7.5. Data files saved in Unicode mode cannot be read by releases of SPSS Statistics prior to version 16.0 [For more information, see General Options in Chapter 48 on p. 494](#).
- When using data files with variable names longer than eight bytes in version 10.x or 11.x, unique, eight-byte versions of variable names are used—but the original variable names are preserved for use in release 12.0 or later. In releases prior to 10.0, the original long variable names are lost if you save the data file.
- When using data files with string variables longer than 255 bytes in versions prior to release 13.0, those string variables are broken up into multiple 255-byte string variables.

Version 7.0 (*.sav). Version 7.0 format. Data files saved in version 7.0 format can be read by version 7.0 and earlier versions but do not include defined multiple response sets or Data Entry for Windows information.

SPSS/PC+ (*.sys). SPSS/PC+ format. If the data file contains more than 500 variables, only the first 500 will be saved. For variables with more than one defined user-missing value, additional user-missing values will be recoded into the first defined user-missing value.

SPSS Statistics Portable (*.por). Portable format that can be read by other versions of SPSS Statistics and versions on other operating systems. Variable names are limited to eight bytes and are automatically converted to unique eight-byte names if necessary. In most cases, saving data in portable format is no longer necessary, since SPSS Statistics data files should be platform/operating system independent. You cannot save data files in portable file in Unicode mode. [For more information, see General Options in Chapter 48 on p. 494](#).

Tab-delimited (*.dat). Text files with values separated by tabs. (*Note:* Tab characters embedded in string values are preserved as tab characters in the tab-delimited file. No distinction is made between tab characters embedded in values and tab characters that separate values.)

Comma-delimited (*.csv). Text files with values separated by commas or semicolons. If the current SPSS Statistics decimal indicator is a period, values are separated by commas. If the current decimal indicator is a comma, values are separated by semicolons.

Fixed ASCII (*.dat). Text file in fixed format, using the default write formats for all variables. There are no tabs or spaces between variable fields.

Excel 2007 (*.xlsx). Microsoft Excel 2007 XLSX-format workbook. The maximum number of variables is 16,000; any additional variables beyond the first 16,000 are dropped. If the dataset contains more than one million cases, multiple sheets are created in the workbook.

Excel 97 through 2003 (*.xls). Microsoft Excel 97 workbook. The maximum number of variables is 256; any additional variables beyond the first 256 are dropped. If the dataset contains more than 65,356 cases, multiple sheets are created in the workbook.

Excel 2.1 (*.xls). Microsoft Excel 2.1 spreadsheet file. The maximum number of variables is 256, and the maximum number of rows is 16,384.

1-2-3 Release 3.0 (*.wk3). Lotus 1-2-3 spreadsheet file, release 3.0. The maximum number of variables that you can save is 256.

1-2-3 Release 2.0 (*.wk1). Lotus 1-2-3 spreadsheet file, release 2.0. The maximum number of variables that you can save is 256.

1-2-3 Release 1.0 (*.wks). Lotus 1-2-3 spreadsheet file, release 1A. The maximum number of variables that you can save is 256.

SYLK (*.slk). Symbolic link format for Microsoft Excel and Multiplan spreadsheet files. The maximum number of variables that you can save is 256.

dBASE IV (*.dbf). dBASE IV format.

dBASE III (*.dbf). dBASE III format.

dBASE II (*.dbf). dBASE II format.

SAS v7+ Windows short extension (*.sd7). SAS versions 7–8 for Windows short filename format.

SAS v7+ Windows long extension (*.sas7bdat). SAS versions 7–8 for Windows long filename format.

SAS v7+ for UNIX (*.ssd01). SAS v8 for UNIX.

SAS v6 for Windows (*.sd2). SAS v6 file format for Windows/OS2.

SAS v6 for UNIX (*.ssd01). SAS v6 file format for UNIX (Sun, HP, IBM).

SAS v6 for Alpha/OSF (*.ssd04). SAS v6 file format for Alpha/OSF (DEC UNIX).

SAS Transport (*.xpt). SAS transport file.

Stata Version 8 Intercooled (*.dta).

Stata Version 8 SE (*.dta).

Stata Version 7 Intercooled (*.dta).

Stata Version 7 SE (*.dta).

Stata Version 6 (*.dta).

Stata Versions 4–5 (*.dta).

Saving File Options

For spreadsheet, tab-delimited files, and comma-delimited files, you can write variable names to the first row of the file.

Saving Data Files in Excel Format

You can save your data in one of three Microsoft Excel file formats. Excel 2.1, Excel 97, and Excel 2007.

- Excel 2.1 and Excel 97 are limited to 256 columns; so only the first 256 variables are included.
- Excel 2007 is limited to 16,000 columns; so only the first 16,000 variables are included.
- Excel 2.1 is limited to 16,384 rows; so only the first 16,384 cases are included.
- Excel 97 and Excel 2007 also have limits on the number of rows per sheet, but workbooks can have multiple sheets, and multiple sheets are created if the single-sheet maximum is exceeded.

Variable Types

The following table shows the variable type matching between the original data in SPSS Statistics and the exported data in Excel.

SPSS Statistics Variable Type	Excel Data Format
Numeric	0.00; #,##0.00; ...
Comma	0.00; #,##0.00; ...
Dollar	\$#,##0_); ...
Date	d-mmm-yyyy
Time	hh:mm:ss
String	General

Saving Data Files in SAS Format

Special handling is given to various aspects of your data when saved as a SAS file. These cases include:

- Certain characters that are allowed in SPSS Statistics variable names are not valid in SAS, such as @, #, and \$. These illegal characters are replaced with an underscore when the data are exported.
- SPSS Statistics variable names that contain multibyte characters (for example, Japanese or Chinese characters) are converted to variables names of the general form *Vnnn*, where *nnn* is an integer value.
- SPSS Statistics variable labels containing more than 40 characters are truncated when exported to a SAS v6 file.
- Where they exist, SPSS Statistics variable labels are mapped to the SAS variable labels. If no variable label exists in the SPSS Statistics data, the variable name is mapped to the SAS variable label.
- SAS allows only one value for system-missing, whereas SPSS Statistics allows numerous system-missing values. As a result, all system-missing values in SPSS Statistics are mapped to a single system-missing value in the SAS file.

Save Value Labels

You have the option of saving the values and value labels associated with your data file to a SAS syntax file. For example, when the value labels for the *cars.sav* data file are exported, the generated syntax file contains:

```
libname library '\name\' ;

proc format library = library ;
  value ORIGIN /* Country of Origin */
    1 = 'American'
    2 = 'European'
    3 = 'Japanese' ;
  value CYLINDER /* Number of Cylinders */
    3 = '3 Cylinders'
    4 = '4 Cylinders'
    5 = '5 Cylinders'
    6 = '6 Cylinders'
    8 = '8 Cylinders' ;
  value FILTER__ /* cylrec = 1 | cylrec = 2 (FILTER) */
    0 = 'Not Selected'
    1 = 'Selected' ;

proc datasets library = library ;
modify cars;
  format  ORIGIN ORIGIN.;
  format  CYLINDER CYLINDER.;
  format  FILTER__ FILTER__.;
quit;
```

This feature is not supported for the SAS transport file.

Variable Types

The following table shows the variable type matching between the original data in SPSS Statistics and the exported data in SAS.

SPSS Statistics Variable Type	SAS Variable Type	SAS Data Format
Numeric	Numeric	12
Comma	Numeric	12
Dot	Numeric	12
Scientific Notation	Numeric	12

SPSS Statistics Variable Type	SAS Variable Type	SAS Data Format
Date	Numeric	(Date) for example, MMDDYY10, ...
Date (Time)	Numeric	Time18
Dollar	Numeric	12
Custom Currency	Numeric	12
String	Character	\$8

Saving Data Files in Stata Format

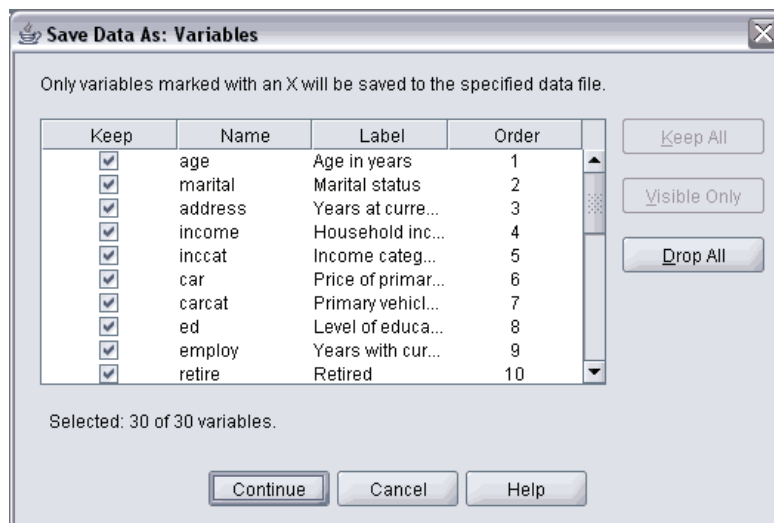
- Data can be written in Stata 5–8 format and in both Intercooled and SE format (versions 7 and 8 only).
- Data files that are saved in Stata 5 format can be read by Stata 4.
- The first 80 bytes of variable labels are saved as Stata variable labels.
- For numeric variables, the first 80 bytes of value labels are saved as Stata value labels. For string variables, value labels are dropped.
- For versions 7 and 8, the first 32 bytes of variable names in case-sensitive form are saved as Stata variable names. For earlier versions, the first eight bytes of variable names are saved as Stata variable names. Any characters other than letters, numbers, and underscores are converted to underscores.
- SPSS Statistics variable names that contain multibyte characters (for example, Japanese or Chinese characters) are converted to variable names of the general form *Vnnn*, where *nnn* is an integer value.
- For versions 5–6 and Intercooled versions 7–8, the first 80 bytes of string values are saved. For Stata SE 7–8, the first 244 bytes of string values are saved.
- For versions 5–6 and Intercooled versions 7–8, only the first 2,047 variables are saved. For Stata SE 7–8, only the first 32,767 variables are saved.

SPSS Statistics Variable Type	Stata Variable Type	Stata Data Format
Numeric	Numeric	g
Comma	Numeric	g
Dot	Numeric	g
Scientific Notation	Numeric	g
Date*, Datetime	Numeric	D_m_Y
Time, DTime	Numeric	g (number of seconds)
Wkday	Numeric	g (1–7)
Month	Numeric	g (1–12)
Dollar	Numeric	g
Custom Currency	Numeric	g
String	String	s

*Date, Adate, Edate, SDate, Jdate, Qyr, Moyr, Wkyr

Saving Subsets of Variables

Figure 3-24
Save Data As Variables dialog box



The Save Data As Variables dialog box allows you to select the variables that you want saved in the new data file. By default, all variables will be saved. Deselect the variables that you don't want to save, or click Drop All and then select the variables that you want to save.

Visible Only. Selects only variables in variable sets currently in use. [For more information, see Use Variable Sets in Chapter 47 on p. 490.](#)

To Save a Subset of Variables

- ▶ Make the Data Editor the active window (click anywhere in the window to make it active).
- ▶ From the menus choose:
 - File
 - Save As...
- ▶ Click Variables.
- ▶ Select the variables that you want to save.

Exporting to a Database

You can use the Export to Database Wizard to:

- Replace values in existing database table fields (columns) or add new fields to a table.
- Append new records (rows) to a database table.
- Completely replace a database table or create a new table.

To export data to a database:

- ▶ From the menus in the Data Editor window for the dataset that contains the data you want to export, choose:
 - File
 - Export to Database
- ▶ Select the database source.
- ▶ Follow the instructions in the export wizard to export the data.

Creating Database Fields from SPSS Statistics Variables

When creating new fields (adding fields to an existing database table, creating a new table, replacing a table), you can specify field names, data type, and width (where applicable).

Field name. The default field names are the same as the SPSS Statistics variable names. You can change the field names to any names allowed by the database format. For example, many databases allow characters in field names that aren't allowed in variable names, including spaces. Therefore, a variable name like *CallWaiting* could be changed to the field name *Call Waiting*.

Type. The export wizard makes initial data type assignments based on the standard ODBC data types or data types allowed by the selected database format that most closely matches the defined SPSS Statistics data format—but databases can make type distinctions that have no direct equivalent in SPSS Statistics, and vice versa. For example, most numeric values in SPSS Statistics are stored as double-precision floating-point values, whereas database numeric data types include float (double), integer, real, and so on. In addition, many databases don't have equivalents to SPSS Statistics time formats. You can change the data type to any type available in the drop-down list.

As a general rule, the basic data type (string or numeric) for the variable should match the basic data type of the database field. If there is a data type mismatch that cannot be resolved by the database, an error results and no data are exported to the database. For example, if you export a string variable to a database field with a numeric data type, an error will result if any values of the string variable contain non-numeric characters.

Width. You can change the defined width for string (char, varchar) field types. Numeric field widths are defined by the data type.

By default, SPSS Statistics variable formats are mapped to database field types based on the following general scheme. Actual database field types may vary, depending on the database.

SPSS Statistics Variable Format	Database Field Type
Numeric	Float or Double
Comma	Float or Double
Dot	Float or Double
Scientific Notation	Float or Double
Date	Date or Datetime or Timestamp
Datetime	Datetime or Timestamp
Time, DTime	Float or Double (number of seconds)
Wkday	Integer (1–7)
Month	Integer (1–12)

SPSS Statistics Variable Format

Dollar
 Custom Currency
 String

Database Field Type

Float or Double
 Float or Double
 Char or Varchar

User-Missing Values

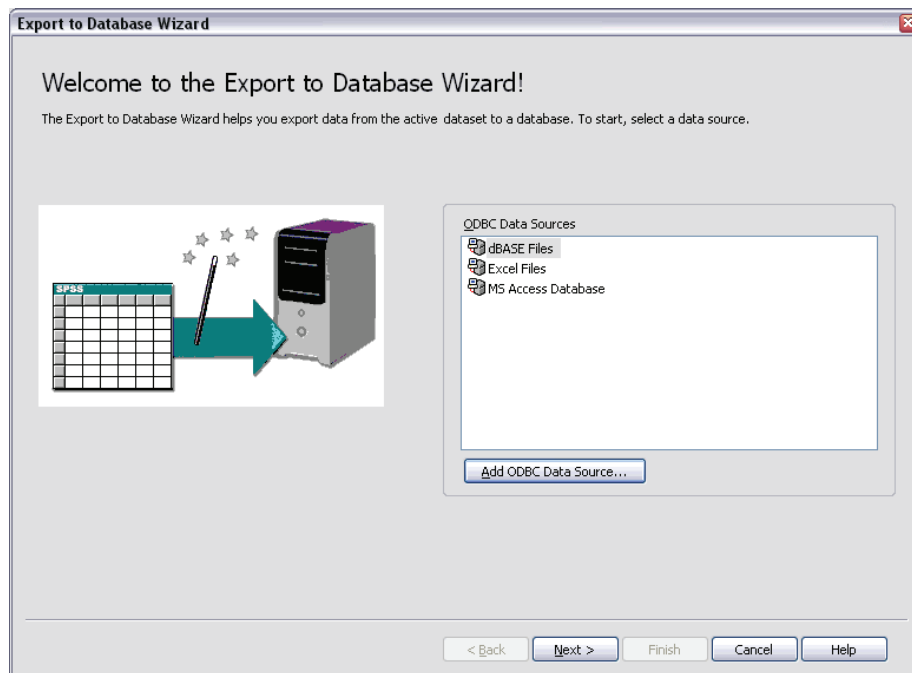
There are two options for the treatment of user-missing values when data from variables are exported to database fields:

- **Export as valid values.** User-missing values are treated as regular, valid, nonmissing values.
- **Export numeric user-missing as nulls and export string user-missing values as blank spaces.** Numeric user-missing values are treated the same as system-missing values. String user-missing values are converted to blank spaces (strings cannot be system-missing).

Selecting a Data Source

In the first panel of the Export to Database Wizard, you select the data source to which you want to export data.

Figure 3-25
 Export to Database Wizard, selecting a data source



You can export data to any database source for which you have the appropriate ODBC driver. (Note: Exporting data to OLE DB data sources is not supported.)

If you do not have any ODBC data sources configured, or if you want to add a new data source, click Add ODBC Data Source.

- On Linux operating systems, this button is not available. ODBC data sources are specified in *odbc.ini*, and the *ODBCINI* environment variables must be set to the location of that file. For more information, see the documentation for your database drivers.
- In distributed analysis mode (available with SPSS Statistics Server), this button is not available. To add data sources in distributed analysis mode, see your system administrator.

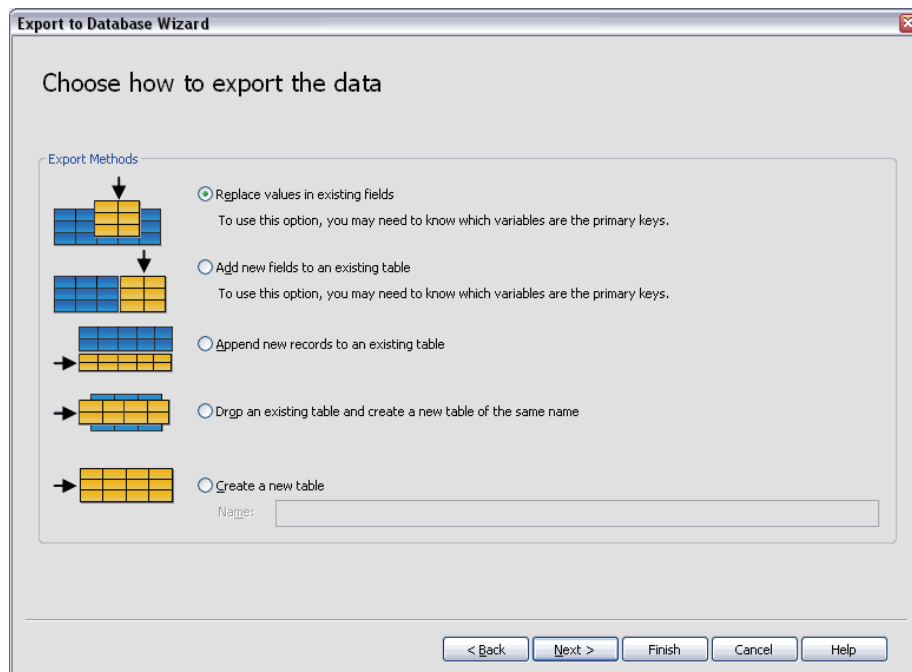
An ODBC data source consists of two essential pieces of information: the driver that will be used to access the data and the location of the database you want to access. To specify data sources, you must have the appropriate drivers installed. Drivers for a variety of database formats are available at <http://www.spss.com/drivers>.

Some data sources may require a login ID and password before you can proceed to the next step.

Choosing How to Export the Data

After you select the data source, you indicate the manner in which you want to export the data.

Figure 3-26
Export to Database Wizard, choosing how to export



The following choices are available for exporting data to a database:

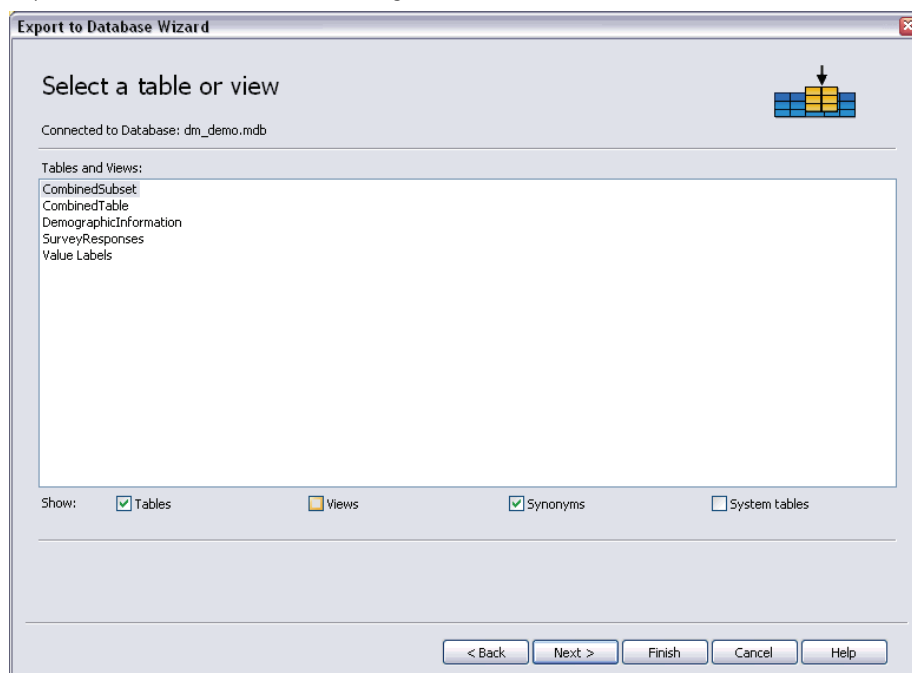
- **Replace values in existing fields.** Replaces values of selected fields in an existing table with values from the selected variables in the active dataset. [For more information, see Replacing Values in Existing Fields on p. 54.](#)
- **Add new fields to an existing table.** Creates new fields in an existing table that contain the values of selected variables in the active dataset. [For more information, see Adding New Fields on p. 55.](#) This option is not available for Excel files.

- **Append new records to an existing table.** Adds new records (rows) to an existing table containing the values from cases in the active dataset. [For more information, see *Appending New Records \(Cases\)* on p. 56.](#)
- **Drop an existing table and create a new table of the same name.** Deletes the specified table and creates a new table of the same name that contains selected variables from the active dataset. All information from the original table, including definitions of field properties (for example, primary keys, data types) is lost. [For more information, see *Creating a New Table or Replacing a Table* on p. 57.](#)
- **Create a new table.** Creates a new table in the database containing data from selected variables in the active dataset. The name can be any value that is allowed as a table name by the data source. The name cannot duplicate the name of an existing table or view in the database. [For more information, see *Creating a New Table or Replacing a Table* on p. 57.](#)

Selecting a Table

When modifying or replacing a table in the database, you need to select the table to modify or replace. This panel in the Export to Database Wizard displays a list of tables and views in the selected database.

Figure 3-27
Export to Database Wizard, selecting a table or view



By default, the list displays only standard database tables. You can control the type of items that are displayed in the list:

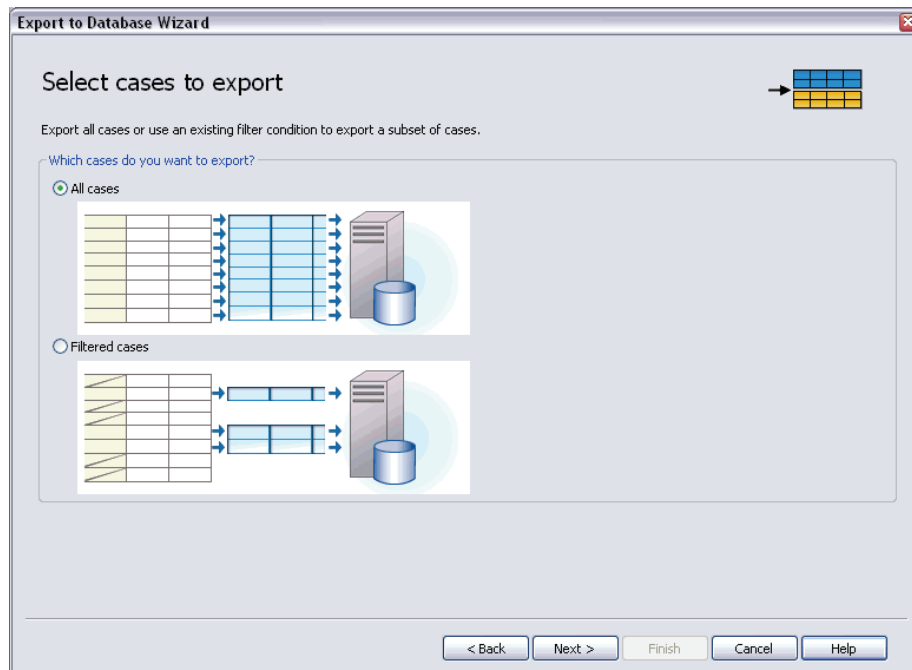
- **Tables.** Standard database tables.

- **Views.** Views are virtual or dynamic “tables” defined by queries. These can include joins of multiple tables and/or fields derived from calculations based on the values of other fields. You can append records or replace values of existing fields in views, but the fields that you can modify may be restricted, depending on how the view is structured. For example, you cannot modify a derived field, add fields to a view, or replace a view.
- **Synonyms.** A synonym is an alias for a table or view, typically defined in a query.
- **System tables.** System tables define database properties. In some cases, standard database tables may be classified as system tables and will be displayed only if you select this option. Access to real system tables is often restricted to database administrators.

Selecting Cases to Export

Case selection in the Export to Database Wizard is limited either to all cases or to cases selected using a previously defined filter condition. If no case filtering is in effect, this panel will not appear, and all cases in the active dataset will be exported.

Figure 3-28
Export to Database Wizard, selecting cases to export



For information on defining a filter condition for case selection, see [Select Cases on p. 180](#).

Matching Cases to Records

When adding fields (columns) to an existing table or replacing the values of existing fields, you need to make sure that each case (row) in the active dataset is correctly matched to the corresponding record in the database.

- In the database, the field or set of fields that uniquely identifies each record is often designated as the **primary key**.
- You need to identify which variable(s) correspond to the primary key field(s) or other fields that uniquely identify each record.
- The fields don't have to be the primary key in the database, but the field value or combination of field values must be unique for each case.

To match variables with fields in the database that uniquely identify each record:

- ▶ Drag and drop the variable(s) onto the corresponding database fields.

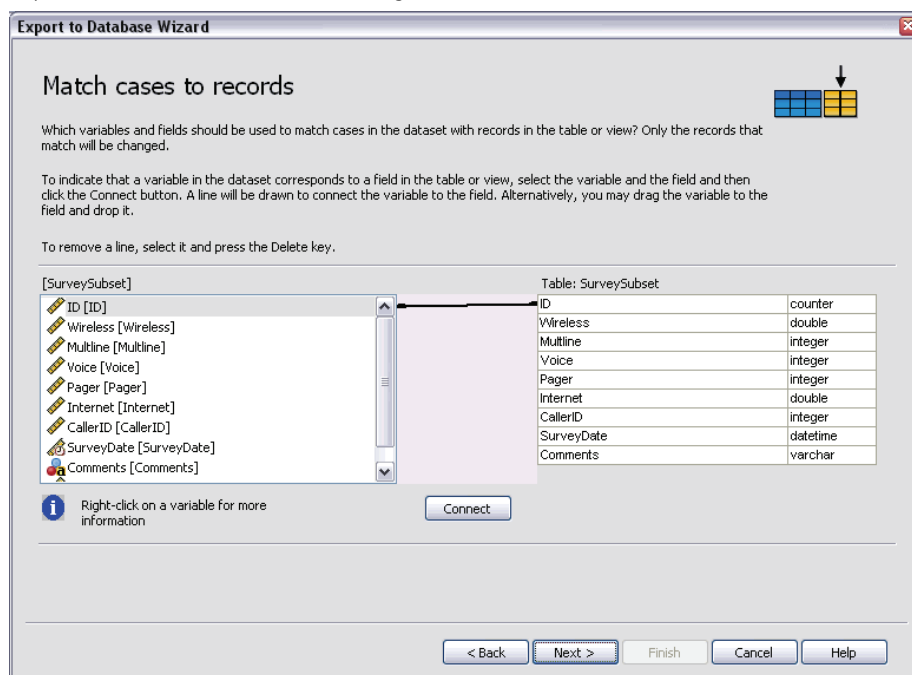
or

- ▶ Select a variable from the list of variables, select the corresponding field in the database table, and click Connect.

To delete a connection line:

- ▶ Select the connection line and press the Delete key.

Figure 3-29
Export to Database Wizard, matching cases to records



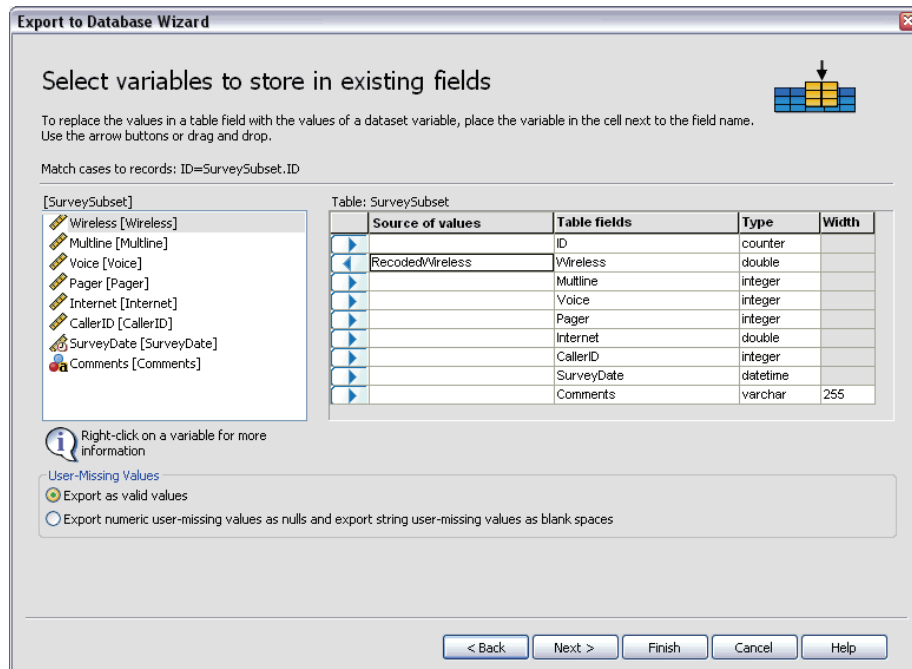
Note: The variable names and database field names may not be identical (since database field names may contain characters not allowed in SPSS Statistics variable names), but if the active dataset was created from the database table you are modifying, either the variable names or the variable labels will usually be at least similar to the database field names.

Replacing Values in Existing Fields

To replace values of existing fields in a database:

- ▶ In the Choose how to export the data panel of the Export to Database Wizard, select Replace values in existing fields.
- ▶ In the Select a table or view panel, select the database table.
- ▶ In the Match cases to records panel, match the variables that uniquely identify each case to the corresponding database field names.
- ▶ For each field for which you want to replace values, drag and drop the variable that contains the new values into the Source of values column, next to the corresponding database field name.

Figure 3-30
Export to Database Wizard, replacing values of existing fields



- As a general rule, the basic data type (string or numeric) for the variable should match the basic data type of the database field. If there is a data type mismatch that cannot be resolved by the database, an error results and no data is exported to the database. For example, if you export a string variable to a database field with a numeric data type (for example, double, real, integer), an error will result if any values of the string variable contain non-numeric characters. The letter *a* in the icon next to a variable denotes a string variable.
- You cannot modify the field name, type, or width. The original database field attributes are preserved; only the values are replaced.

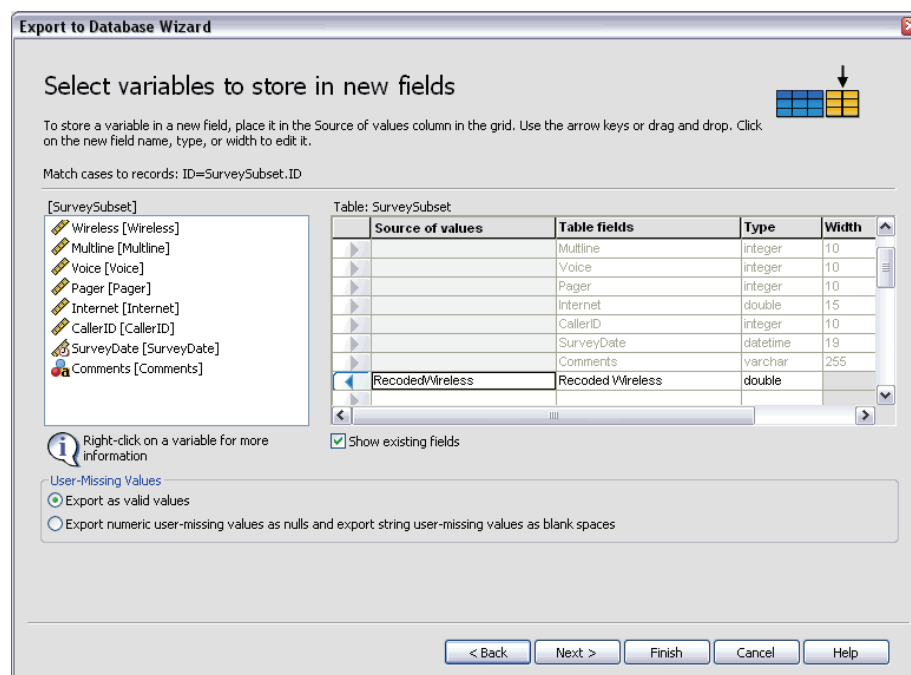
Adding New Fields

To add new fields to an existing database table:

- ▶ In the Choose how to export the data panel of the Export to Database Wizard, select Add new fields to an existing table.
- ▶ In the Select a table or view panel, select the database table.
- ▶ In the Match cases to records panel, match the variables that uniquely identify each case to the corresponding database field names.
- ▶ Drag and drop the variables that you want to add as new fields to the Source of values column.

Figure 3-31

Export to Database Wizard, adding new fields to an existing table



For information on field names and data types, see the section on creating database fields from SPSS Statistics variables in [Exporting to a Database on p. 47](#).

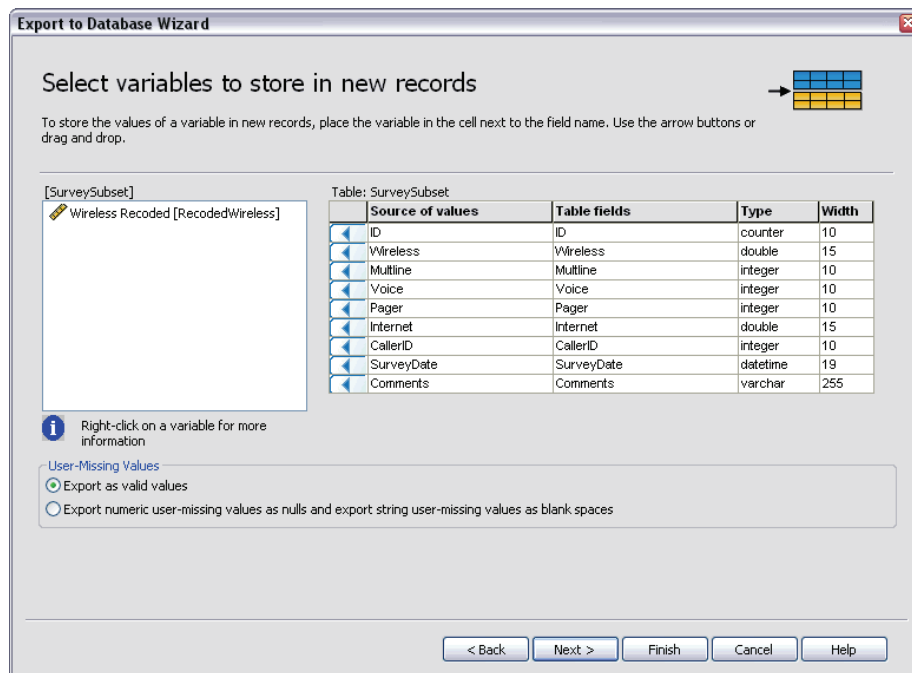
Show existing fields. Select this option to display a list of existing fields. You cannot use this panel in the Export to Database Wizard to replace existing fields, but it may be helpful to know what fields are already present in the table. If you want to replace the values of existing fields, see [Replacing Values in Existing Fields on p. 54](#).

Appending New Records (Cases)

To append new records (cases) to a database table:

- ▶ In the Choose how to export the data panel of the Export to Database Wizard, select Append new records to an existing table.
- ▶ In the Select a table or view panel, select the database table.
- ▶ Match variables in the active dataset to table fields by dragging and dropping variables to the Source of values column.

Figure 3-32
Export to Database Wizard, adding records (cases) to a table



The Export to Database Wizard will automatically select all variables that match existing fields, based on information about the original database table stored in the active dataset (if available) and/or variable names that are the same as field names. This initial automatic matching is intended only as a guide and does not prevent you from changing the way in which variables are matched with database fields.

When adding new records to an existing table, the following basic rules/limitations apply:

- All cases (or all selected cases) in the active dataset are added to the table. If any of these cases duplicate existing records in the database, an error may result if a duplicate key value is encountered. For information on exporting only selected cases, see [Selecting Cases to Export](#) on p. 52.

- You can use the values of new variables created in the session as the values for existing fields, but you cannot add new fields or change the names of existing fields. To add new fields to a database table, see [Adding New Fields on p. 55](#).
- Any excluded database fields or fields not matched to a variable will have no values for the added records in the database table. (If the Source of values cell is empty, there is no variable matched to the field.)

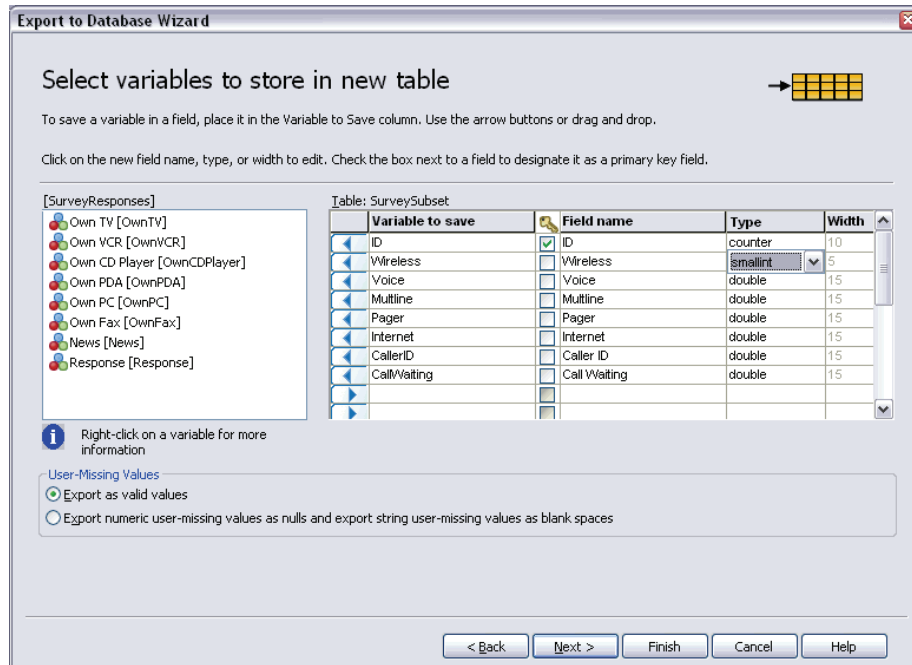
Creating a New Table or Replacing a Table

To create a new database table or replace an existing database table:

- ▶ In the Choose how to export the data panel of the export wizard, select Drop an existing table and create a new table of the same name or select Create a new table and enter a name for the new table.
- ▶ If you are replacing an existing table, in the Select a table or view panel, select the database table.
- ▶ Drag and drop variables into the Variables to save column.
- ▶ Optionally, you can designate variables/fields that define the primary key, change field names, and change the data type.

Figure 3-33

Export to Database Wizard, selecting variables for a new table



Primary key. To designate variables as the primary key in the database table, select the box in the column identified with the key icon.

- All values of the primary key must be unique or an error will result.

- If you select a single variable as the primary key, every record (case) must have a unique value for that variable.
- If you select multiple variables as the primary key, this defines a composite primary key, and the *combination* of values for the selected variables must be unique for each case.

For information on field names and data types, see the section on creating database fields from SPSS Statistics variables in [Exporting to a Database](#) on p. 47.

Completing the Database Export Wizard

The last panel of the Export to Database Wizard provides a summary that indicates what data will be exported and how it will be exported. It also gives you the option of either exporting the data or pasting the underlying command syntax to a syntax window.

Figure 3-34
Export to Database Wizard, finish panel



Summary Information

- **Dataset.** The SPSS Statistics session name for the dataset that will be used to export data. This information is primarily useful if you have multiple open data sources. Data sources opened using the graphical user interface (for example, the Database Wizard) are automatically assigned names such as *DataSet1*, *DataSet2*, etc. A data source opened using command syntax will have a dataset name only if one is explicitly assigned.
- **Table.** The name of the table to be modified or created.
- **Cases to Export.** Either all cases are exported or cases selected by a previously defined filter condition are exported. [For more information, see Selecting Cases to Export](#) on p. 52.

- **Action.** Indicates how the database will be modified (for example, create a new table, add fields or records to an existing table).
- **User-Missing Values.** User-missing values can be exported as valid values or treated the same as system-missing for numeric variables and converted to blank spaces for string variables. This setting is controlled in the panel in which you select the variables to export.

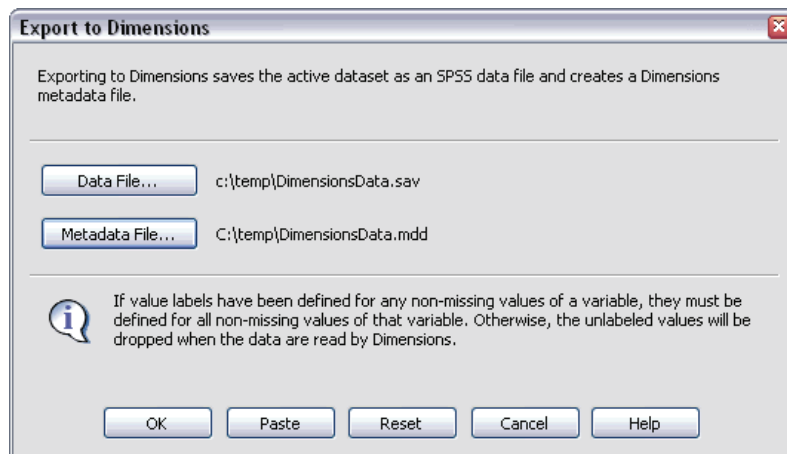
Exporting to Dimensions

The Export to Dimensions dialog box creates SPSS Statistics data files and Dimensions metadata files that you can use to read the data into Dimensions applications such as mrInterview and mrTables. This is particularly useful when “roundtripping” data between SPSS Statistics and Dimensions applications. For example, you can read an mrInterview data source into SPSS Statistics, calculate some new variables, and then save the data in a form that can be read by mrTables, without loss of any of the original metadata attributes.

To export data for use in Dimensions applications:

- ▶ From the menus in the Data Editor window that contains the data you want to export, choose:
File
Export to Dimensions
- ▶ Click Data File to specify the name and location of the SPSS Statistics data file.
- ▶ Click Metadata File to specify the name and location of the Dimensions metadata file.

Figure 3-35
Export to Dimensions dialog box



For new variables and datasets not created from Dimensions data sources, SPSS Statistics variable attributes are mapped to Dimensions metadata attributes in the metadata file according to the methods described in the SAV DSC documentation in the Dimensions Development Library.

If the active dataset was created from a Dimensions data source:

- The new metadata file is created by merging the original metadata attributes with metadata attributes for any new variables, plus any changes to original variables that might affect their metadata attributes (for example, addition of, or changes to, value labels).
- For original variables read from the Dimensions data source, any metadata attributes not recognized by SPSS Statistics are preserved in their original state. For example, SPSS Statistics converts grid variables to regular SPSS Statistics variables, but the metadata that defines these grid variables is preserved when you save the new metadata file.
- If any Dimensions variables were automatically renamed to conform to SPSS Statistics variable naming rules, the metadata file maps the converted names back to the original Dimensions variable names.

The presence or absence of value labels can affect the metadata attributes of variables and consequently the way those variables are read by Dimensions applications. If value labels have been defined for any nonmissing values of a variable, they should be defined for all nonmissing values of that variable; otherwise, the unlabeled values will be dropped when the data file is read by Dimensions.

Protecting Original Data

To prevent the accidental modification or deletion of your original data, you can mark the file as read-only.

- ▶ From the Data Editor menus choose:

File
Mark File Read Only

If you make subsequent modifications to the data and then try to save the data file, you can save the data only with a different filename, so the original data are not affected.

You can change the file permissions back to read-write by choosing Mark File Read Write from the File menu.

Virtual Active File

The virtual active file enables you to work with large data files without requiring equally large (or larger) amounts of temporary disk space. For most analysis and charting procedures, the original data source is reread each time you run a different procedure. Procedures that modify the data require a certain amount of temporary disk space to keep track of the changes, and some actions always require enough disk space for at least one entire copy of the data file.

Figure 3-36
Temporary disk space requirements

Action	GET FILE = 'v1-5.sav'. FREQUENCIES...	COMPUTE v6 = ... RECODE v4... REGRESSION... /SAVE ZPRED.	SORT CASES BY... or CACHE																																																																																																																																					
Virtual Active File	<table border="1"> <thead> <tr><th>v1</th><th>v2</th><th>v3</th><th>v4</th><th>v5</th></tr> </thead> <tbody> <tr><td>11</td><td>12</td><td>13</td><td>14</td><td>15</td></tr> <tr><td>21</td><td>22</td><td>23</td><td>24</td><td>25</td></tr> <tr><td>31</td><td>32</td><td>33</td><td>34</td><td>35</td></tr> <tr><td>41</td><td>42</td><td>43</td><td>44</td><td>45</td></tr> <tr><td>51</td><td>52</td><td>53</td><td>54</td><td>55</td></tr> <tr><td>61</td><td>62</td><td>63</td><td>64</td><td>65</td></tr> </tbody> </table>	v1	v2	v3	v4	v5	11	12	13	14	15	21	22	23	24	25	31	32	33	34	35	41	42	43	44	45	51	52	53	54	55	61	62	63	64	65	<table border="1"> <thead> <tr><th>v1</th><th>v2</th><th>v3</th><th>v4</th><th>v5</th><th>v6</th><th>zpre</th></tr> </thead> <tbody> <tr><td>11</td><td>12</td><td>13</td><td>14</td><td>15</td><td>16</td><td>1</td></tr> <tr><td>21</td><td>22</td><td>23</td><td>24</td><td>25</td><td>26</td><td>2</td></tr> <tr><td>31</td><td>32</td><td>33</td><td>34</td><td>35</td><td>36</td><td>3</td></tr> <tr><td>41</td><td>42</td><td>43</td><td>44</td><td>45</td><td>46</td><td>4</td></tr> <tr><td>51</td><td>52</td><td>53</td><td>54</td><td>55</td><td>56</td><td>5</td></tr> <tr><td>61</td><td>62</td><td>63</td><td>64</td><td>65</td><td>66</td><td>6</td></tr> </tbody> </table>	v1	v2	v3	v4	v5	v6	zpre	11	12	13	14	15	16	1	21	22	23	24	25	26	2	31	32	33	34	35	36	3	41	42	43	44	45	46	4	51	52	53	54	55	56	5	61	62	63	64	65	66	6	<table border="1"> <thead> <tr><th>v1</th><th>v2</th><th>v3</th><th>v4</th><th>v5</th><th>v6</th><th>zpre</th></tr> </thead> <tbody> <tr><td>11</td><td>12</td><td>13</td><td>14</td><td>15</td><td>16</td><td>1</td></tr> <tr><td>21</td><td>22</td><td>23</td><td>24</td><td>25</td><td>26</td><td>2</td></tr> <tr><td>31</td><td>32</td><td>33</td><td>34</td><td>35</td><td>36</td><td>3</td></tr> <tr><td>41</td><td>42</td><td>43</td><td>44</td><td>45</td><td>46</td><td>4</td></tr> <tr><td>51</td><td>52</td><td>53</td><td>54</td><td>55</td><td>56</td><td>5</td></tr> <tr><td>61</td><td>62</td><td>63</td><td>64</td><td>65</td><td>66</td><td>6</td></tr> </tbody> </table>	v1	v2	v3	v4	v5	v6	zpre	11	12	13	14	15	16	1	21	22	23	24	25	26	2	31	32	33	34	35	36	3	41	42	43	44	45	46	4	51	52	53	54	55	56	5	61	62	63	64	65	66	6
v1	v2	v3	v4	v5																																																																																																																																				
11	12	13	14	15																																																																																																																																				
21	22	23	24	25																																																																																																																																				
31	32	33	34	35																																																																																																																																				
41	42	43	44	45																																																																																																																																				
51	52	53	54	55																																																																																																																																				
61	62	63	64	65																																																																																																																																				
v1	v2	v3	v4	v5	v6	zpre																																																																																																																																		
11	12	13	14	15	16	1																																																																																																																																		
21	22	23	24	25	26	2																																																																																																																																		
31	32	33	34	35	36	3																																																																																																																																		
41	42	43	44	45	46	4																																																																																																																																		
51	52	53	54	55	56	5																																																																																																																																		
61	62	63	64	65	66	6																																																																																																																																		
v1	v2	v3	v4	v5	v6	zpre																																																																																																																																		
11	12	13	14	15	16	1																																																																																																																																		
21	22	23	24	25	26	2																																																																																																																																		
31	32	33	34	35	36	3																																																																																																																																		
41	42	43	44	45	46	4																																																																																																																																		
51	52	53	54	55	56	5																																																																																																																																		
61	62	63	64	65	66	6																																																																																																																																		
Data Stored in Temporary Disk Space	None	<table border="1"> <thead> <tr><th>v4</th><th>v6</th><th>zpre</th></tr> </thead> <tbody> <tr><td>14</td><td>16</td><td>1</td></tr> <tr><td>24</td><td>26</td><td>2</td></tr> <tr><td>34</td><td>36</td><td>3</td></tr> <tr><td>44</td><td>46</td><td>4</td></tr> <tr><td>54</td><td>56</td><td>5</td></tr> <tr><td>64</td><td>66</td><td>6</td></tr> </tbody> </table>	v4	v6	zpre	14	16	1	24	26	2	34	36	3	44	46	4	54	56	5	64	66	6	<table border="1"> <thead> <tr><th>v1</th><th>v2</th><th>v3</th><th>v4</th><th>v5</th><th>v6</th><th>zpre</th></tr> </thead> <tbody> <tr><td>11</td><td>12</td><td>13</td><td>14</td><td>15</td><td>16</td><td>1</td></tr> <tr><td>21</td><td>22</td><td>23</td><td>24</td><td>25</td><td>26</td><td>2</td></tr> <tr><td>31</td><td>32</td><td>33</td><td>34</td><td>35</td><td>36</td><td>3</td></tr> <tr><td>41</td><td>42</td><td>43</td><td>44</td><td>45</td><td>46</td><td>4</td></tr> <tr><td>51</td><td>52</td><td>53</td><td>54</td><td>55</td><td>56</td><td>5</td></tr> <tr><td>61</td><td>62</td><td>63</td><td>64</td><td>65</td><td>66</td><td>6</td></tr> </tbody> </table>	v1	v2	v3	v4	v5	v6	zpre	11	12	13	14	15	16	1	21	22	23	24	25	26	2	31	32	33	34	35	36	3	41	42	43	44	45	46	4	51	52	53	54	55	56	5	61	62	63	64	65	66	6																																																															
v4	v6	zpre																																																																																																																																						
14	16	1																																																																																																																																						
24	26	2																																																																																																																																						
34	36	3																																																																																																																																						
44	46	4																																																																																																																																						
54	56	5																																																																																																																																						
64	66	6																																																																																																																																						
v1	v2	v3	v4	v5	v6	zpre																																																																																																																																		
11	12	13	14	15	16	1																																																																																																																																		
21	22	23	24	25	26	2																																																																																																																																		
31	32	33	34	35	36	3																																																																																																																																		
41	42	43	44	45	46	4																																																																																																																																		
51	52	53	54	55	56	5																																																																																																																																		
61	62	63	64	65	66	6																																																																																																																																		

Actions that don't require any temporary disk space include:

- Reading SPSS Statistics data files
- Merging two or more SPSS Statistics data files
- Reading database tables with the Database Wizard
- Merging SPSS Statistics data files with database tables
- Running procedures that read data (for example, Frequencies, Crosstabs, Explore)

Actions that create one or more columns of data in temporary disk space include:

- Computing new variables
- Recoding existing variables
- Running procedures that create or modify variables (for example, saving predicted values in Linear Regression)

Actions that create an entire copy of the data file in temporary disk space include:

- Reading Excel files
- Running procedures that sort data (for example, Sort Cases, Split File)
- Reading data with GET TRANSLATE or DATA LIST commands
- Using the Cache Data facility or the CACHE command
- Launching other applications from SPSS Statistics that read the data file (for example, AnswerTree, DecisionTime)

Note: The GET DATA command provides functionality comparable to DATA LIST without creating an entire copy of the data file in temporary disk space. The SPLIT FILE command in command syntax does not sort the data file and therefore does not create a copy of the data file. This command, however, requires sorted data for proper operation, and the dialog box interface

for this procedure will automatically sort the data file, resulting in a complete copy of the data file. (Command syntax is not available with the Student Version.)

Actions that create an entire copy of the data file by default:

- Reading databases with the Database Wizard
- Reading text files with the Text Wizard

The Text Wizard provides an optional setting to automatically cache the data. By default, this option is selected. You can turn it off by deselecting `Cache data locally`. For the Database Wizard, you can paste the generated command syntax and delete the `CACHE` command.

Creating a Data Cache

Although the virtual active file can vastly reduce the amount of temporary disk space required, the absence of a temporary copy of the “active” file means that the original data source has to be reread for each procedure. For large data files read from an external source, creating a temporary copy of the data may improve performance. For example, for data tables read from a database source, the SQL query that reads the information from the database must be reexecuted for any command or procedure that needs to read the data. Since virtually all statistical analysis procedures and charting procedures need to read the data, the SQL query is reexecuted for each procedure you run, which can result in a significant increase in processing time if you run a large number of procedures.

If you have sufficient disk space on the computer performing the analysis (either your local computer or a remote server), you can eliminate multiple SQL queries and improve processing time by creating a data cache of the active file. The data cache is a temporary copy of the complete data.

Note: By default, the Database Wizard automatically creates a data cache, but if you use the `GET DATA` command in command syntax to read a database, a data cache is not automatically created. (Command syntax is not available with the Student Version.)

To Create a Data Cache

- ▶ From the menus choose:
 - File
 - Cache Data...
- ▶ Click OK or Cache Now.

OK creates a data cache the next time the program reads the data (for example, the next time you run a statistical procedure), which is usually what you want because it doesn’t require an extra data pass. Cache Now creates a data cache immediately, which shouldn’t be necessary under most circumstances. Cache Now is useful primarily for two reasons:

- A data source is “locked” and can’t be updated by anyone until you end your session, open a different data source, or cache the data.
- For large data sources, scrolling through the contents of the Data View tab in the Data Editor will be much faster if you cache the data.

To Cache Data Automatically

You can use the `SET` command to automatically create a data cache after a specified number of changes in the active data file. By default, the active data file is automatically cached after 20 changes in the active data file.

- ▶ From the menus choose:

- File
 - New
 - Syntax

- ▶ In the syntax window, type `SET CACHE n` (where n represents the number of changes in the active data file before the data file is cached).

- ▶ From the menus in the syntax window choose:

- Run
- All

Note: The cache setting is not persistent across sessions. Each time you start a new session, the value is reset to the default of 20.

Distributed Analysis Mode

Distributed analysis mode allows you to use a computer other than your local (or desktop) computer for memory-intensive work. Because remote servers that are used for distributed analysis are typically more powerful and faster than your local computer, distributed analysis mode can significantly reduce computer processing time. Distributed analysis with a remote server can be useful if your work involves:

- Large data files, particularly data read from database sources.
- Memory-intensive tasks. Any task that takes a long time in local analysis mode may be a good candidate for distributed analysis.

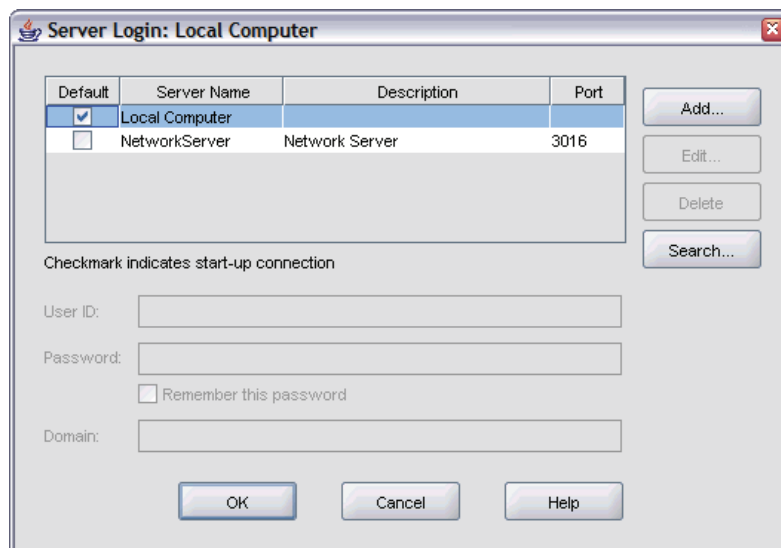
Distributed analysis affects only data-related tasks, such as reading data, transforming data, computing new variables, and calculating statistics. Distributed analysis has no effect on tasks related to editing output, such as manipulating pivot tables or modifying charts.

Note: Distributed analysis is available only if you have both a local version and access to a licensed server version of the software that is installed on a remote server.

Server Login

The Server Login dialog box allows you to select the computer that processes commands and runs procedures. You can select your local computer or a remote server.

Figure 4-1
Server Login dialog box



You can add, modify, or delete remote servers in the list. Remote servers usually require a user ID and password, and a domain name may also be necessary. Contact your system administrator for information about available servers, a user ID and password, domain names, and other connection information.

You can select a default server and save the user ID, domain name, and password that are associated with any server. You are automatically connected to the default server when you start a new session.

If you are licensed to use the Adaptor for Predictive Enterprise Services and your site is running Predictive Enterprise Services 3.5 or later, you can click Search... to view a list of servers that are available on your network. If you are not logged on to a Predictive Enterprise Repository, you will be prompted to enter connection information before you can view the list of servers.

Adding and Editing Server Login Settings

Use the Server Login Settings dialog box to add or edit connection information for remote servers for use in distributed analysis mode.

Figure 4-2
Server Login Settings dialog box



Contact your system administrator for a list of available servers, port numbers for the servers, and additional connection information. Do not use the Secure Socket Layer unless instructed to do so by your administrator.

Server Name. A server “name” can be an alphanumeric name that is assigned to a computer (for example, NetworkServer) or a unique IP address that is assigned to a computer (for example, 202.123.456.78).

Port Number. The port number is the port that the server software uses for communications.

Description. You can enter an optional description to display in the servers list.

Connect with Secure Socket Layer. Secure Socket Layer (SSL) encrypts requests for distributed analysis when they are sent to the remote server. Before you use SSL, check with your administrator. For this option to be enabled, SSL must be configured on your desktop computer and the server.

To Select, Switch, or Add Servers

- ▶ From the menus choose:

- File
 - Switch Server...

To select a default server:

- ▶ In the server list, select the box next to the server that you want to use.
- ▶ Enter the user ID, domain name, and password that were provided by your administrator.
Note: You are automatically connected to the default server when you start a new session.

To switch to another server:

- ▶ Select the server from the list.
- ▶ Enter your user ID, domain name, and password (if necessary).
Note: When you switch servers during a session, all open windows are closed. You will be prompted to save changes before the windows are closed.

To add a server:

- ▶ Get the server connection information from your administrator.
- ▶ Click Add to open the Server Login Settings dialog box.
- ▶ Enter the connection information and optional settings, and then click OK.

To edit a server:

- ▶ Get the revised connection information from your administrator.
- ▶ Click Edit to open the Server Login Settings dialog box.
- ▶ Enter the changes and click OK.

To search for available servers:

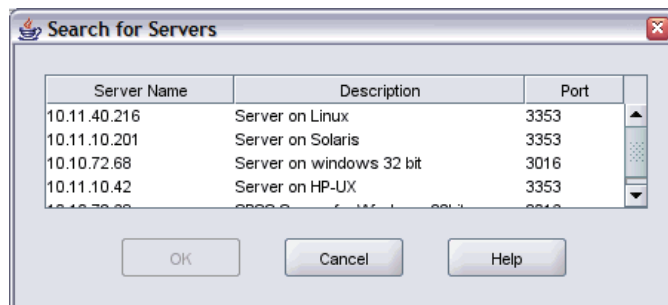
Note: The ability to search for available servers is available only if you are licensed to use the Adaptor for Predictive Enterprise Services and your site is running Predictive Enterprise Services 3.5 or later.

- ▶ Click Search... to open the Search for Servers dialog box. If you are not logged on to a Predictive Enterprise Repository, you will be prompted for connection information.
- ▶ Select one or more available servers and click OK. The servers will now appear in the Server Login dialog box.
- ▶ To connect to one of the servers, follow the instructions “To switch to another server.”

Searching for Available Servers

Use the Search for Servers dialog box to select one or more servers that are available on your network. This dialog box appears when you click Search... on the Server Login dialog box.

Figure 4-3
Search for Servers dialog box



Select one or more servers and click OK to add them to the Server Login dialog box. Although you can manually add servers in the Server Login dialog box, searching for available servers lets you connect to servers without requiring that you know the correct server name and port number. This information is automatically provided. However, you still need the correct logon information, such as user name, domain, and password.

Opening Data Files from a Remote Server

In distributed analysis mode, the Open Remote File dialog box replaces the standard Open File dialog box.

- The contents of the list of available files, folders, and drives depends on what is available on or from the remote server. The current server name is indicated at the top of the dialog box.
- In distributed analysis mode, you will not have access to files on your local computer unless you specify the drive as a shared device or specify the folders containing your data files as shared folders. Consult the documentation for your operating system for information on how to “share” folders on your local computer with the server network.
- If the server is running a different operating system (for example, you are running Windows and the server is running UNIX), you probably won’t have access to local data files in distributed analysis mode even if they are in shared folders.

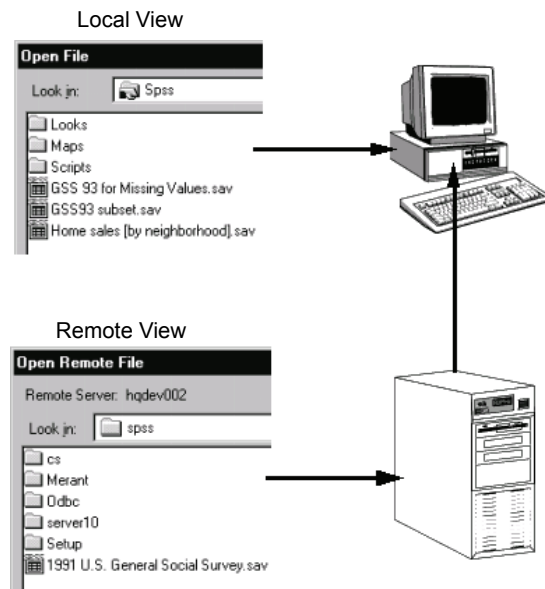
File Access in Local and Distributed Analysis Mode

The view of data folders (directories) and drives for both your local computer and the network is based on the computer that you are currently using to process commands and run procedures—which is not necessarily the computer in front of you.

Local analysis mode. When you use your local computer as your “server,” the view of data files, folders, and drives in the file access dialog box (for opening data files) is similar to what you see in other applications or in Windows Explorer. You can see all of the data files and folders on your computer and any files and folders on mounted network drives.

Distributed analysis mode. When you use another computer as a “remote server” to run commands and procedures, the view of data files, folders, and drives represents the view from the remote server computer. Although you may see familiar folder names (such as *Program Files*) and drives (such as *C*), these items are *not* the folders and drives on your computer; they are the folders and drives on the remote server.

Figure 4-4
Local and remote views



In distributed analysis mode, you will not have access to data files on your local computer unless you specify the drive as a shared device or specify the folders containing your data files as shared folders. If the server is running a different operating system (for example, you are running Windows and the server is running UNIX), you probably won't have access to local data files in distributed analysis mode even if they are in shared folders.

Distributed analysis mode is not the same as accessing data files that reside on another computer on your network. You can access data files on other network devices in local analysis mode or in distributed analysis mode. In local mode, you access other devices from your local computer. In distributed mode, you access other network devices from the remote server.

If you're not sure if you're using local analysis mode or distributed analysis mode, look at the title bar in the dialog box for accessing data files. If the title of the dialog box contains the word *Remote* (as in *Open Remote File*), or if the text *Remote Server: [server name]* appears at the top of the dialog box, you're using distributed analysis mode.

Note: This situation affects only dialog boxes for accessing data files (for example, *Open Data*, *Save Data*, *Open Database*, and *Apply Data Dictionary*). For all other file types (for example, *Viewer files*, *syntax files*, and *script files*), the local view is always used.

Availability of Procedures in Distributed Analysis Mode

In distributed analysis mode, procedures are available for use only if they are installed on both your local version and the version on the remote server.

If you have optional components installed locally that are not available on the remote server and you switch from your local computer to a remote server, the affected procedures will be removed from the menus and the corresponding command syntax will result in errors. Switching back to local mode will restore all affected procedures.

Absolute versus Relative Path Specifications

In distributed analysis mode, relative path specifications for data files and command syntax files are relative to the current server, not relative to your local computer. A relative path specification such as `/mydocs/mydata.sav` does not point to a directory and file on your local drive; it points to a directory and file on the remote server's hard drive.

Windows UNC Path Specifications

If you are using a Windows server version, you can use universal naming convention (UNC) specifications when accessing data and syntax files with command syntax. The general form of a UNC specification is:

```
\\servername\sharename\path\filename
```

- *Servername* is the name of the computer that contains the data file.
- *Sharename* is the folder (directory) on that computer that is designated as a shared folder.
- *Path* is any additional folder (subdirectory) path below the shared folder.
- *Filename* is the name of the data file.

An example is as follows:

```
GET FILE='\\hqdev001\public\july\sales.sav'.
```

If the computer does not have a name assigned to it, you can use its IP address, as in:

```
GET FILE='\\204.125.125.53\public\july\sales.sav'.
```

Even with UNC path specifications, you can access data and syntax files only from devices and folders that are designated as shared. When you use distributed analysis mode, this situation includes data and syntax files on your local computer.

UNIX Absolute Path Specifications

For UNIX server versions, there is no equivalent to the UNC path, and all directory paths must be absolute paths that start at the root of the server; relative paths are not allowed. For example, if the data file is located in `/bin/data` and the current directory is also `/bin/data`, `GET FILE='sales.sav'` is not valid; you must specify the entire path, as in:

```
GET FILE='/bin/sales.sav'.  
INSERT FILE='/bin/salesjob.sps'.
```

Data Editor

The Data Editor provides a convenient, spreadsheet-like method for creating and editing data files. The Data Editor window opens automatically when you start a session.

The Data Editor provides two views of your data:

- **Data View.** This view displays the actual data values or defined value labels.
- **Variable View.** This view displays variable definition information, including defined variable and value labels, data type (for example, string, date, or numeric), measurement level (nominal, ordinal, or scale), and user-defined missing values.

In both views, you can add, change, and delete information that is contained in the data file.

Data View

Figure 5-1
Data View

	age	marital	address	income	inccat	car
1	55	Married	12	72.00	\$50 - \$74	36.
2	56	Unmarried	29	153.00	\$75+	76.
3	28	Married	9	28.00	\$25 - \$49	13.
4	24	Married	4	26.00	\$25 - \$49	12.
5	25	Unmarried	2	23.00	Under \$25	11.
6	45	Married	9	76.00	\$75+	37.
7	42	Unmarried	19	40.00	\$25 - \$49	19.
8	35	Unmarried	15	57.00	\$50 - \$74	28.
9	46	Unmarried	26	24.00	Under \$25	12.
10	34	Married	0	89.00	\$75+	46.
11	55	Married	17	72.00	\$50 - \$74	35.

Many of the features of Data View are similar to the features that are found in spreadsheet applications. There are, however, several important distinctions:

- Rows are cases. Each row represents a case or an observation. For example, each individual respondent to a questionnaire is a case.
- Columns are variables. Each column represents a variable or characteristic that is being measured. For example, each item on a questionnaire is a variable.

- Cells contain values. Each cell contains a single value of a variable for a case. The cell is where the case and the variable intersect. Cells contain only data values. Unlike spreadsheet programs, cells in the Data Editor cannot contain formulas.
- The data file is rectangular. The dimensions of the data file are determined by the number of cases and variables. You can enter data in any cell. If you enter data in a cell outside the boundaries of the defined data file, the data rectangle is extended to include any rows and/or columns between that cell and the file boundaries. There are no “empty” cells within the boundaries of the data file. For numeric variables, blank cells are converted to the system-missing value. For string variables, a blank is considered a valid value.

Variable View

Figure 5-2
Variable View

	Name	Type	Width	Decimals	Label	Value
1	age	Numeric	4	0	Age in years	None
2	marital	Numeric	4	0	Marital status	{0, Unmar
3	address	Numeric	4	0	Years at curre	None
4	income	Numeric	8	2	Household inc	None
5	inccat	Numeric	8	2	Income catego	{1.00, Unc
6	car	Numeric	8	2	Price of primar	None
7	carcat	Numeric	8	2	Primary vehicl	{1.00, Ecc
8	ed	Numeric	4	0	Level of educat	{1, Did no
9	employ	Numeric	4	0	Years with curr	None
10	retire	Numeric	4	0	Retired	{0, No}...
11	empcat	Numeric	4	0	Years with curr	{1, Less t
12	jobsat	Numeric	4	0	Job satisfactio	{1, Highly
13	gender	String	2	0	Gender	{f, Female
14	reside	Numeric	4	0	Number of peo	None
15	wireless	Numeric	4	0	Wireless servi	{0, No}...
16	multiline	Numeric	4	0	Multiple lines	{0, No}...

Variable View contains descriptions of the attributes of each variable in the data file. In Variable View:

- Rows are variables.
- Columns are variable attributes.

You can add or delete variables and modify attributes of variables, including the following attributes:

- Variable name
- Data type
- Number of digits or characters
- Number of decimal places
- Descriptive variable and value labels

- User-defined missing values
- Column width
- Measurement level

All of these attributes are saved when you save the data file.

In addition to defining variable properties in Variable View, there are two other methods for defining variable properties:

- The Copy Data Properties Wizard provides the ability to use an external SPSS Statistics data file or another dataset that is available in the current session as a template for defining file and variable properties in the active dataset. You can also use variables in the active dataset as templates for other variables in the active dataset. Copy Data Properties is available on the Data menu in the Data Editor window.
- Define Variable Properties (also available on the Data menu in the Data Editor window) scans your data and lists all unique data values for any selected variables, identifies unlabeled values, and provides an auto-label feature. This method is particularly useful for categorical variables that use numeric codes to represent categories—for example, 0 = *Male*, 1 = *Female*.

To Display or Define Variable Attributes

- ▶ Make the Data Editor the active window.
- ▶ Double-click a variable name at the top of the column in Data View, or click the Variable View tab.
- ▶ To define new variables, enter a variable name in any blank row.
- ▶ Select the attribute(s) that you want to define or modify.

Variable Names

The following rules apply to variable names:

- Each variable name must be unique; duplication is not allowed.
- Variable names can be up to 64 bytes long, and the first character must be a letter or one of the characters @, #, or \$. Subsequent characters can be any combination of letters, numbers, nonpunctuation characters, and a period (.). In code page mode, sixty-four bytes typically means 64 characters in single-byte languages (for example, English, French, German, Spanish, Italian, Hebrew, Russian, Greek, Arabic, and Thai) and 32 characters in double-byte languages (for example, Japanese, Chinese, and Korean). Many string characters that only take one byte in code page mode take two or more bytes in Unicode mode. For example, é is one byte in code page format but is two bytes in Unicode format; so *résumé* is six bytes in a code page file and eight bytes in Unicode mode.

Note: Letters include any nonpunctuation characters used in writing ordinary words in the languages supported in the platform's character set.

- Variable names cannot contain spaces.

- A # character in the first position of a variable name defines a scratch variable. You can only create scratch variables with command syntax. You cannot specify a # as the first character of a variable in dialog boxes that create new variables.
- A \$ sign in the first position indicates that the variable is a system variable. The \$ sign is not allowed as the initial character of a user-defined variable.
- The period, the underscore, and the characters \$, #, and @ can be used within variable names. For example, *A_.\$@#1* is a valid variable name.
- Variable names ending with a period should be avoided, since the period may be interpreted as a command terminator. You can only create variables that end with a period in command syntax. You cannot create variables that end with a period in dialog boxes that create new variables.
- Variable names ending in underscores should be avoided, since such names may conflict with names of variables automatically created by commands and procedures.
- Reserved keywords cannot be used as variable names. Reserved keywords are ALL, AND, BY, EQ, GE, GT, LE, LT, NE, NOT, OR, TO, and WITH.
- Variable names can be defined with any mixture of uppercase and lowercase characters, and case is preserved for display purposes.
- When long variable names need to wrap onto multiple lines in output, lines are broken at underscores, periods, and points where content changes from lower case to upper case.

Variable Measurement Level

You can specify the level of measurement as scale (numeric data on an interval or ratio scale), ordinal, or nominal. Nominal and ordinal data can be either string (alphanumeric) or numeric.

- **Nominal.** A variable can be treated as nominal when its values represent categories with no intrinsic ranking (for example, the department of the company in which an employee works). Examples of nominal variables include region, zip code, and religious affiliation.
- **Ordinal.** A variable can be treated as ordinal when its values represent categories with some intrinsic ranking (for example, levels of service satisfaction from highly dissatisfied to highly satisfied). Examples of ordinal variables include attitude scores representing degree of satisfaction or confidence and preference rating scores.
- **Scale.** A variable can be treated as scale when its values represent ordered categories with a meaningful metric, so that distance comparisons between values are appropriate. Examples of scale variables include age in years and income in thousands of dollars.

Note: For ordinal string variables, the alphabetic order of string values is assumed to reflect the true order of the categories. For example, for a string variable with the values of *low*, *medium*, *high*, the order of the categories is interpreted as *high*, *low*, *medium*, which is not the correct order. In general, it is more reliable to use numeric codes to represent ordinal data.

New numeric variables created during a session are assigned the scale measurement level. For data read from external file formats and SPSS Statistics data files that were created prior to version 8.0, default assignment of measurement level is based on the following rules:

- Numeric variables with fewer than 24 unique values and string variables are set to nominal.
- Numeric variables with 24 or more unique values are set to scale.

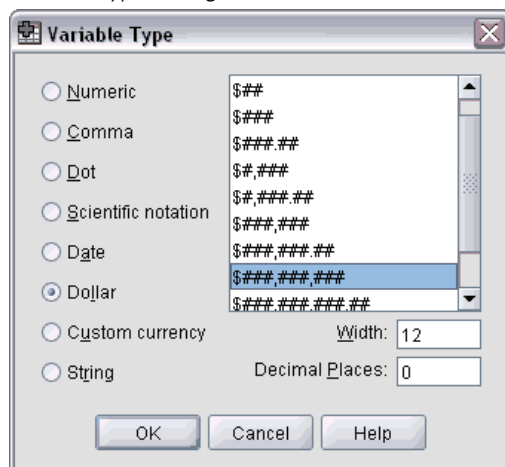
You can change the scale/nominal cutoff value for numeric variables in the Options dialog box. For more information, see [Data Options in Chapter 48 on p. 497](#).

The Define Variable Properties dialog box, available from the Data menu, can help you assign the correct measurement level. For more information, see [Assigning the Measurement Level in Chapter 7 on p. 101](#).

Variable Type

Variable Type specifies the data type for each variable. By default, all new variables are assumed to be numeric. You can use Variable Type to change the data type. The contents of the Variable Type dialog box depend on the selected data type. For some data types, there are text boxes for width and number of decimals; for other data types, you can simply select a format from a scrollable list of examples.

Figure 5-3
Variable Type dialog box



The available data types are as follows:

Numeric. A variable whose values are numbers. Values are displayed in standard numeric format. The Data Editor accepts numeric values in standard format or in scientific notation.

Comma. A numeric variable whose values are displayed with commas delimiting every three places and displayed with the period as a decimal delimiter. The Data Editor accepts numeric values for comma variables with or without commas or in scientific notation. Values cannot contain commas to the right of the decimal indicator.

Dot. A numeric variable whose values are displayed with periods delimiting every three places and with the comma as a decimal delimiter. The Data Editor accepts numeric values for dot variables with or without periods or in scientific notation. Values cannot contain periods to the right of the decimal indicator.

Scientific notation. A numeric variable whose values are displayed with an embedded E and a signed power-of-10 exponent. The Data Editor accepts numeric values for such variables with or without an exponent. The exponent can be preceded by E or D with an optional sign or by the sign alone—for example, 123, 1.23E2, 1.23D2, 1.23E+2, and 1.23+2.

Date. A numeric variable whose values are displayed in one of several calendar-date or clock-time formats. Select a format from the list. You can enter dates with slashes, hyphens, periods, commas, or blank spaces as delimiters. The century range for two-digit year values is determined by your Options settings (from the Edit menu, choose Options, and then click the Data tab).

Dollar. A numeric variable displayed with a leading dollar sign (\$), commas delimiting every three places, and a period as the decimal delimiter. You can enter data values with or without the leading dollar sign.

Custom currency. A numeric variable whose values are displayed in one of the custom currency formats that you have defined on the Currency tab of the Options dialog box. Defined custom currency characters cannot be used in data entry but are displayed in the Data Editor.

String. A variable whose values are not numeric and therefore are not used in calculations. The values can contain any characters up to the defined length. Uppercase and lowercase letters are considered distinct. This type is also known as an alphanumeric variable.

To Define Variable Type

- ▶ Click the button in the *Type* cell for the variable that you want to define.
- ▶ Select the data type in the Variable Type dialog box.
- ▶ Click OK.

Input versus Display Formats

Depending on the format, the display of values in Data View may differ from the actual value as entered and stored internally. Following are some general guidelines:

- For numeric, comma, and dot formats, you can enter values with any number of decimal positions (up to 16), and the entire value is stored internally. The Data View displays only the defined number of decimal places and rounds values with more decimals. However, the complete value is used in all computations.
- For string variables, all values are right-padded to the maximum width. For a string variable with a maximum width of three, a value of *No* is stored internally as 'N O ' and is not equivalent to ' N O '.
- For date formats, you can use slashes, dashes, spaces, commas, or periods as delimiters between day, month, and year values, and you can enter numbers, three-letter abbreviations, or complete names for month values. Dates of the general format *dd-mmm-yy* are displayed with dashes as delimiters and three-letter abbreviations for the month. Dates of the general

format *dd/mm/yy* and *mm/dd/yy* are displayed with slashes for delimiters and numbers for the month. Internally, dates are stored as the number of seconds from October 14, 1582. The century range for dates with two-digit years is determined by your Options settings (from the Edit menu, choose Options, and then click the Data tab).

- For time formats, you can use colons, periods, or spaces as delimiters between hours, minutes, and seconds. Times are displayed with colons as delimiters. Internally, times are stored as a number of seconds that represents a time interval. For example, 10:00:00 is stored internally as 36000, which is 60 (seconds per minute) x 60 (minutes per hour) x 10 (hours).

Variable Labels

You can assign descriptive variable labels up to 256 characters (128 characters in double-byte languages). Variable labels can contain spaces and reserved characters that are not allowed in variable names.

To Specify Variable Labels

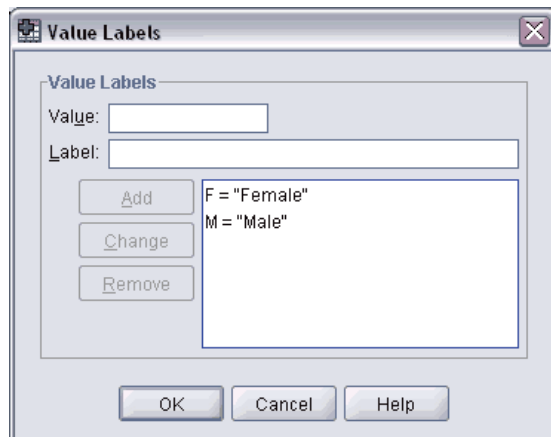
- ▶ Make the Data Editor the active window.
- ▶ Double-click a variable name at the top of the column in Data View, or click the Variable View tab.
- ▶ In the *Label* cell for the variable, enter the descriptive variable label.

Value Labels

You can assign descriptive value labels for each value of a variable. This process is particularly useful if your data file uses numeric codes to represent non-numeric categories (for example, codes of 1 and 2 for *male* and *female*).

- Value labels can be up to 120 bytes.

Figure 5-4
Value Labels dialog box



To Specify Value Labels

- ▶ Click the button in the *Values* cell for the variable that you want to define.
- ▶ For each value, enter the value and a label.
- ▶ Click Add to enter the value label.
- ▶ Click OK.

Inserting Line Breaks in Labels

Variable labels and value labels automatically wrap to multiple lines in pivot tables and charts if the cell or area isn't wide enough to display the entire label on one line, and you can edit results to insert manual line breaks if you want the label to wrap at a different point. You can also create variable labels and value labels that will *always* wrap at specified points and be displayed on multiple lines.

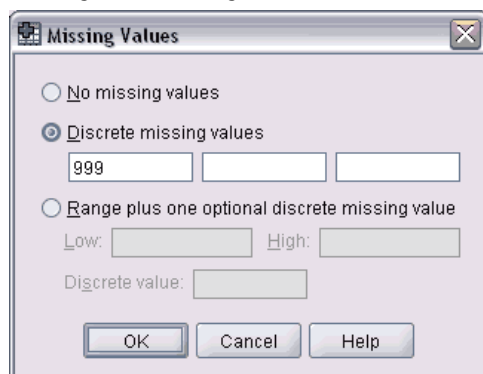
- ▶ For variable labels, select the *Label* cell for the variable in Variable View in the Data Editor.
- ▶ For value labels, select the *Values* cell for the variable in Variable View in the Data Editor, click the button in the cell, and select the label that you want to modify in the Value Labels dialog box.
- ▶ At the place in the label where you want the label to wrap, type `\n`.

The `\n` is not displayed in pivot tables or charts; it is interpreted as a line break character.

Missing Values

Missing Values defines specified data values as **user-missing**. For example, you might want to distinguish between data that are missing because a respondent refused to answer and data that are missing because the question didn't apply to that respondent. Data values that are specified as user-missing are flagged for special treatment and are excluded from most calculations.

Figure 5-5
Missing Values dialog box



- You can enter up to three discrete (individual) missing values, a range of missing values, or a range plus one discrete value.
- Ranges can be specified only for numeric variables.

- All string values, including null or blank values, are considered to be valid unless you explicitly define them as missing.
- Missing values for string variables cannot exceed eight bytes. (There is no limit on the defined width of the string variable, but defined missing values cannot exceed eight bytes.)
- To define null or blank values as missing for a string variable, enter a single space in one of the fields under the Discrete missing values selection.

To Define Missing Values

- ▶ Click the button in the *Missing* cell for the variable that you want to define.
- ▶ Enter the values or range of values that represent missing data.

Column Width

You can specify a number of characters for the column width. Column widths can also be changed in Data View by clicking and dragging the column borders.

- Column width for proportional fonts is based on average character width. Depending on the characters used in the value, more or fewer characters may be displayed in the specified width.
- Column width affect only the display of values in the Data Editor. Changing the column width does not change the defined width of a variable.

Variable Alignment

Alignment controls the display of data values and/or value labels in Data View. The default alignment is right for numeric variables and left for string variables. This setting affects only the display in Data View.

Applying Variable Definition Attributes to Multiple Variables

After you have defined variable definition attributes for a variable, you can copy one or more attributes and apply them to one or more variables.

Basic copy and paste operations are used to apply variable definition attributes. You can:

- Copy a single attribute (for example, value labels) and paste it to the same attribute cell(s) for one or more variables.
- Copy all attributes from one variable and paste them to one or more other variables.
- Create multiple new variables with all the attributes of a copied variable.

Applying Variable Definition Attributes to Other Variables

To Apply Individual Attributes from a Defined Variable

- ▶ In Variable View, select the attribute cell that you want to apply to other variables.

- ▶ From the menus choose:
 - Edit
 - Copy
- ▶ Select the attribute cell(s) to which you want to apply the attribute. (You can select multiple target variables.)
- ▶ From the menus choose:
 - Edit
 - Paste

If you paste the attribute to blank rows, new variables are created with default attributes for all attributes except the selected attribute.

To Apply All Attributes from a Defined Variable

- ▶ In Variable View, select the row number for the variable with the attributes that you want to use. (The entire row is highlighted.)
- ▶ From the menus choose:
 - Edit
 - Copy
- ▶ Select the row number(s) for the variable(s) to which you want to apply the attributes. (You can select multiple target variables.)
- ▶ From the menus choose:
 - Edit
 - Paste

Generating Multiple New Variables with the Same Attributes

- ▶ In Variable View, click the row number for the variable that has the attributes that you want to use for the new variable. (The entire row is highlighted.)
- ▶ From the menus choose:
 - Edit
 - Copy
- ▶ Click the empty row number beneath the last defined variable in the data file.
- ▶ From the menus choose:
 - Edit
 - Paste Variables...
- ▶ In the Paste Variables dialog box, enter the number of variables that you want to create.
- ▶ Enter a prefix and starting number for the new variables.
- ▶ Click OK.

The new variable names will consist of the specified prefix plus a sequential number starting with the specified number.

Custom Variable Attributes

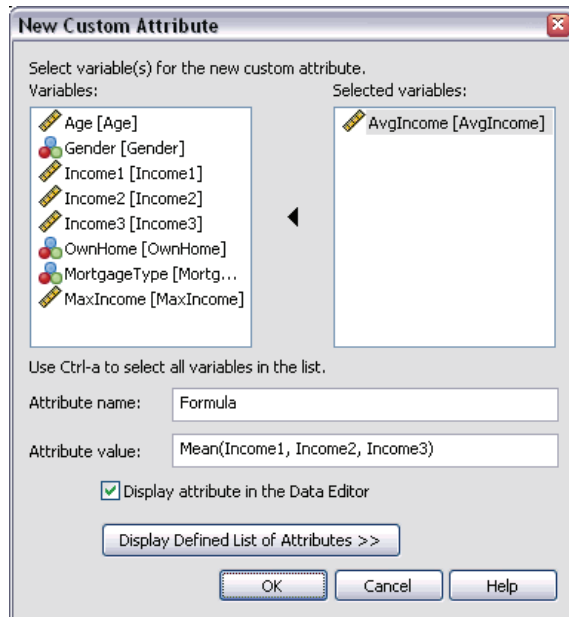
In addition to the standard variable attributes (for example, value labels, missing values, measurement level), you can create your own custom variable attributes. Like standard variable attributes, these custom attributes are saved with SPSS Statistics data files. Therefore, you could create a variable attribute that identifies the type of response for survey questions (for example, single selection, multiple selection, fill-in-the-blank) or the formulas used for computed variables.

Creating Custom Variable Attributes

To create new custom attributes:

- ▶ In Variable View, from the menus choose:
Data
New Custom Attribute...
- ▶ Drag and drop the variables to which you want to assign the new attribute to the Selected Variables list.
- ▶ Enter a name for the attribute. Attribute names must follow the same rules as variable names. [For more information, see Variable Names on p. 72.](#)
- ▶ Enter an optional value for the attribute. If you select multiple variables, the value is assigned to all selected variables. You can leave this blank and then enter values for each variable in Variable View.

Figure 5-6
New Custom Attribute dialog box



Display attribute in the Data Editor. Displays the attribute in Variable View of the Data Editor. For information on controlling the display of custom attributes, see [Displaying and Editing Custom Variable Attributes](#) below.

Display Defined List of Attributes. Displays a list of custom attributes already defined for the dataset. Attribute names that begin with a dollar sign (\$) are reserved attributes that cannot be modified.

Displaying and Editing Custom Variable Attributes

Custom variable attributes can be displayed and edited in the Data Editor in Variable View.

Figure 5-7
Custom variable attributes displayed in Variable View

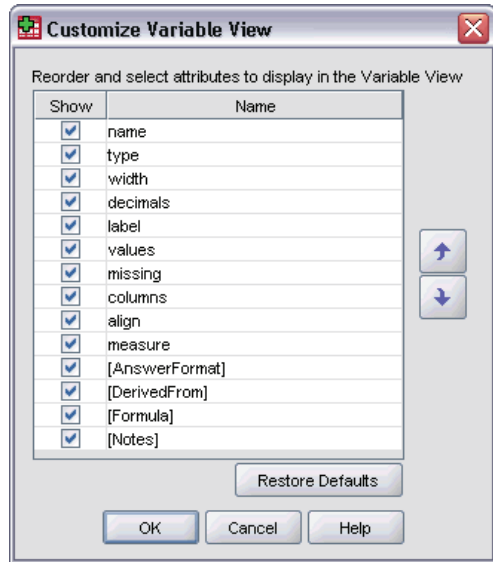
	[AnswerFormat]	[DerivedFrom]	[Formula]	[Notes]
1	Fill-in			Empty
2	Single select			Empty
3	Fill-in			Empty
4	Fill-in			Empty
5	Fill-in			Empty
6	Single select			Empty
7	Single select			Empty
8	Computed value	Array...	mean(Income1, Income2, Income3)	Empty
9	Computed value	Array...	max(Income1, Income2, Income3)	Empty
10				
11				
12				

- Custom variable attribute names are enclosed in square brackets.
- Attribute names that begin with a dollar sign are reserved and cannot be modified.
- A blank cell indicates that the attribute does not exist for that variable; the text *Empty* displayed in a cell indicates that the attribute exists for that variable but no value has been assigned to the attribute for that variable. Once you enter text in the cell, the attribute exists for that variable with the value you enter.
- The text *Array...* displayed in a cell indicates that this is an attribute array—an attribute that contains multiple values. Click the button in the cell to display the list of values.

To Display and Edit Custom Variable Attributes

- ▶ In Variable View, from the menus choose:
View
Customize Variable View...
- ▶ Select (check) the custom variable attributes you want to display. (The custom variable attributes are the ones enclosed in square brackets.)

Figure 5-8
Customize Variable View

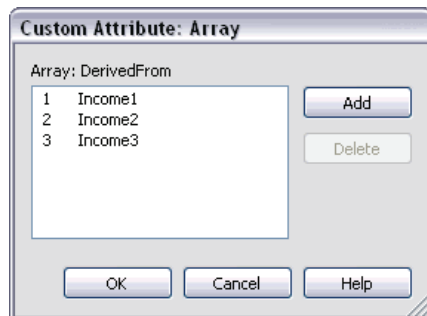


Once the attributes are displayed in Variable View, you can edit them directly in the Data Editor.

Variable Attribute Arrays

The text *Array...* (displayed in a cell for a custom variable attribute in Variable View or in the Custom Variable Properties dialog box in Define Variable Properties) indicates that this is an **attribute array**, an attribute that contains multiple values. For example, you could have an attribute array that identifies all of the source variables used to compute a derived variable. Click the button in the cell to display and edit the list of values.

Figure 5-9
Custom Attribute Array dialog box



Customizing Variable View

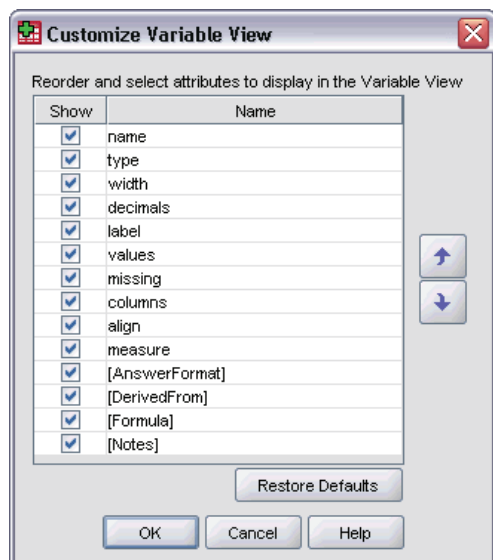
You can use Customize Variable View to control which attributes are displayed in Variable View (for example, name, type, label) and the order in which they are displayed.

- Any custom variable attributes associated with the dataset are enclosed in square brackets. [For more information, see Creating Custom Variable Attributes on p. 80.](#)
- Customized display settings are saved with SPSS Statistics data files.
- You can also control the default display and order of attributes in Variable View. [For more information, see Changing the Default Variable View in Chapter 48 on p. 499.](#)

To Customize Variable View

- ▶ In Variable View, from the menus choose:
View
Customize Variable View...
- ▶ Select (check) the variable attributes you want to display.
- ▶ Use the up and down arrow buttons to change the display order of the attributes.

Figure 5-10
Customize Variable View dialog box



Restore Defaults. Apply the default display and order settings.

Spell Checking

Variable and Value Labels

To check the spelling of variable labels and value labels:

- ▶ Select the Variable View tab in the Data Editor window.
- ▶ Right-click on the *Labels* or *Values* column and from the context menu choose:
Spelling

or

- ▶ In Variable View, from the menus choose:
 - Utilities
 - Spelling

or

 - ▶ In the Value Labels dialog box, click Spelling. (This limits the spell checking to the value labels for a particular variable.)
- Spell checking is limited to variable labels and value labels in Variable View of the Data Editor.

String Data Values

To check the spelling of string data values:

- ▶ Select the Data View tab of the Data Editor.
- ▶ Optionally, select one or more variables (columns) to check. To select a variable, click the variable name at the top of the column.
- ▶ From the menus choose:
 - Utilities
 - Spelling
 - If there are no selected variables in Data View, all string variables will be checked.
 - If there are no string variables in the dataset or the none of the selected variables is a string variable, the Spelling option on the Utilities menu is disabled.

Entering Data

In Data View, you can enter data directly in the Data Editor. You can enter data in any order. You can enter data by case or by variable, for selected areas or for individual cells.

- The active cell is highlighted.
- The variable name and row number of the active cell are displayed in the top left corner of the Data Editor.
- When you select a cell and enter a data value, the value is displayed in the cell editor at the top of the Data Editor.
- Data values are not recorded until you press Enter or select another cell.
- To enter anything other than simple numeric data, you must define the variable type first.

If you enter a value in an empty column, the Data Editor automatically creates a new variable and assigns a variable name.

To Enter Numeric Data

- ▶ Select a cell in Data View.
- ▶ Enter the data value. (The value is displayed in the cell editor at the top of the Data Editor.)
- ▶ To record the value, press Enter or select another cell.

To Enter Non-Numeric Data

- ▶ Double-click a variable name at the top of the column in Data View or click the Variable View tab.
- ▶ Click the button in the *Type* cell for the variable.
- ▶ Select the data type in the Variable Type dialog box.
- ▶ Click OK.
- ▶ Double-click the row number or click the Data View tab.
- ▶ Enter the data in the column for the newly defined variable.

To Use Value Labels for Data Entry

- ▶ If value labels aren't currently displayed in Data View, from the menus choose:
View
Value Labels
- ▶ Click the cell in which you want to enter the value.
- ▶ Choose a value label from the drop-down list.

The value is entered, and the value label is displayed in the cell.

Note: This process works only if you have defined value labels for the variable.

Data Value Restrictions in the Data Editor

The defined variable type and width determine the type of value that can be entered in the cell in Data View.

- If you type a character that is not allowed by the defined variable type, the character is not entered.
- For string variables, characters beyond the defined width are not allowed.
- For numeric variables, integer values that exceed the defined width can be entered, but the Data Editor displays either scientific notation or a portion of the value followed by an ellipsis (...) to indicate that the value is wider than the defined width. To display the value in the cell, change the defined width of the variable.

Note: Changing the column width does not affect the variable width.

Editing Data

With the Data Editor, you can modify data values in Data View in many ways. You can:

- Change data values
- Cut, copy, and paste data values
- Add and delete cases

- Add and delete variables
- Change the order of variables

Replacing or Modifying Data Values

To Delete the Old Value and Enter a New Value

- ▶ In Data View, double-click the cell. (The cell value is displayed in the cell editor.)
- ▶ Edit the value directly in the cell or in the cell editor.
- ▶ Press Enter or select another cell to record the new value.

Cutting, Copying, and Pasting Data Values

You can cut, copy, and paste individual cell values or groups of values in the Data Editor. You can:

- Move or copy a single cell value to another cell
- Move or copy a single cell value to a group of cells
- Move or copy the values for a single case (row) to multiple cases
- Move or copy the values for a single variable (column) to multiple variables
- Move or copy a group of cell values to another group of cells

Data Conversion for Pasted Values in the Data Editor

If the defined variable types of the source and target cells are not the same, the Data Editor attempts to convert the value. If no conversion is possible, the system-missing value is inserted in the target cell.

Converting numeric or date into string. Numeric (for example, numeric, dollar, dot, or comma) and date formats are converted to strings if they are pasted into a string variable cell. The string value is the numeric value as displayed in the cell. For example, for a dollar format variable, the displayed dollar sign becomes part of the string value. Values that exceed the defined string variable width are truncated.

Converting string into numeric or date. String values that contain acceptable characters for the numeric or date format of the target cell are converted to the equivalent numeric or date value. For example, a string value of 25/12/91 is converted to a valid date if the format type of the target cell is one of the day-month-year formats, but the value is converted to system-missing if the format type of the target cell is one of the month-day-year formats.

Converting date into numeric. Date and time values are converted to a number of seconds if the target cell is one of the numeric formats (for example, numeric, dollar, dot, or comma). Because dates are stored internally as the number of seconds since October 14, 1582, converting dates to numeric values can yield some extremely large numbers. For example, the date 10/29/91 is converted to a numeric value of 12,908,073,600.

Converting numeric into date or time. Numeric values are converted to dates or times if the value represents a number of seconds that can produce a valid date or time. For dates, numeric values that are less than 86,400 are converted to the system-missing value.

Inserting New Cases

Entering data in a cell in a blank row automatically creates a new case. The Data Editor inserts the system-missing value for all other variables for that case. If there are any blank rows between the new case and the existing cases, the blank rows become new cases with the system-missing value for all variables. You can also insert new cases between existing cases.

To Insert New Cases between Existing Cases

- ▶ In Data View, select any cell in the case (row) below the position where you want to insert the new case.
- ▶ From the menus choose:
 - Edit
 - Insert Cases

A new row is inserted for the case, and all variables receive the system-missing value.

Inserting New Variables

Entering data in an empty column in Data View or in an empty row in Variable View automatically creates a new variable with a default variable name (the prefix *var* and a sequential number) and a default data format type (numeric). The Data Editor inserts the system-missing value for all cases for the new variable. If there are any empty columns in Data View or empty rows in Variable View between the new variable and the existing variables, these rows or columns also become new variables with the system-missing value for all cases. You can also insert new variables between existing variables.

To Insert New Variables between Existing Variables

- ▶ Select any cell in the variable to the right of (Data View) or below (Variable View) the position where you want to insert the new variable.
- ▶ From the menus choose:
 - Edit
 - Insert Variable

A new variable is inserted with the system-missing value for all cases.

To Move Variables

- ▶ To select the variable, click the variable name in Data View or the row number for the variable in Variable View.
- ▶ Drag and drop the variable to the new location.

- ▶ If you want to place the variable between two existing variables: In Data View, drop the variable on the variable column to the right of where you want to place the variable, or in Variable View, drop the variable on the variable row below where you want to place the variable.

To Change Data Type

You can change the data type for a variable at any time by using the Variable Type dialog box in Variable View. The Data Editor will attempt to convert existing values to the new type. If no conversion is possible, the system-missing value is assigned. The conversion rules are the same as the rules for pasting data values to a variable with a different format type. If the change in data format may result in the loss of missing-value specifications or value labels, the Data Editor displays an alert box and asks whether you want to proceed with the change or cancel it.

Finding Cases, Variables, or Imputations

The Go To dialog box finds the specified case (row) number or variable name in the Data Editor.

Cases

- ▶ For cases, from the menus choose:
Edit
Go to Case...

- ▶ Enter an integer value that represents the current row number in Data View.

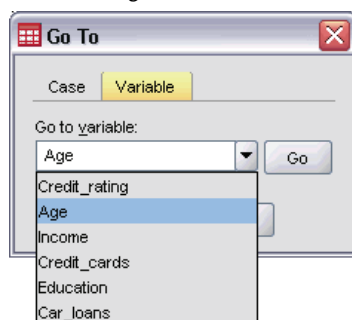
Note: The current row number for a particular case can change due to sorting and other actions.

Variables

- ▶ For variables, from the menus choose:
Edit
Go to Variable...

- ▶ Enter the variable name or select the variable from the drop-down list.

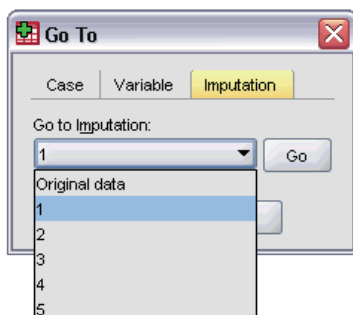
Figure 5-11
Go To dialog box



Imputations

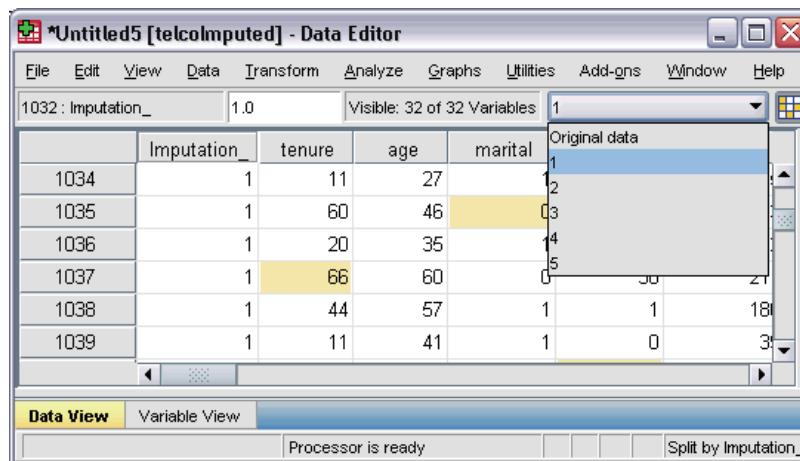
- ▶ From the menus choose:
Edit
Go to Imputation...
- ▶ Select the imputation (or Original data) from the drop-down list.

Figure 5-12
Go To dialog box



Alternatively, you can select the imputation from the drop-down list in the edit bar in Data View of the Data Editor.

Figure 5-13
Data Editor with imputation markings ON



Relative case position is preserved when selecting imputations. For example, if there are 1000 cases in the original dataset, case 1034, the 34th case in the first imputation, displays at the top of the grid. If you select imputation 2 in the dropdown, case 2034, the 34th case in imputation 2, would display at the top of the grid. If you select Original data in the dropdown, case 34 would display at the top of the grid. Column position is also preserved when navigating between imputations, so that it is easy to compare values between imputations.

Finding and Replacing Data and Attribute Values

To find and/or replace data values in Data View or attribute values in Variable View:

- ▶ Click a cell in the column you want to search. (Finding and replacing values is restricted to a single column.)

- ▶ From the menus choose:

Edit
Find

or

Edit
Replace

Data View

- You cannot search up in Data View. The search direction is always down.
- For dates and times, the formatted values as displayed in Data View are searched. For example, a date displayed as 10/28/2007 will not be found by a search for a date of 10-28-2007.
- For other numeric variables, Contains, Begins with, and Ends with search formatted values. For example, with the Begins with option, a search value of \$123 for a Dollar format variable will find both \$123.00 and \$123.40 but not \$1,234. With the Entire cell option, the search value can be formatted or unformatted (simple F numeric format), but only exact numeric values (to the precision displayed in the Data Editor) are matched.
- If value labels are displayed for the selected variable column, the label text is searched, not the underlying data value, and you cannot replace the label text.

Variable View

- Find is only available for the *Name*, *Label*, *Values*, *Missing*, and custom variable attribute columns.
- Replace is only available for the *Label*, *Values*, and custom attribute columns.
- In the *Values* (value labels) column, the search string can match either the data value or a value label.

Note: Replacing the data value will delete any previous value label associated with that value.

Case Selection Status in the Data Editor

If you have selected a subset of cases but have not discarded unselected cases, unselected cases are marked in the Data Editor with a diagonal line (slash) through the row number.

Figure 5-14
Filtered cases in the Data Editor

	age	marital	address	income	inccat	car
1	18	Unmarried	0	13.00	Under \$25	6.30
2	18	Unmarried	0	17.00	Under \$25	8.60
3	19	Married	0	14.00	Under \$25	7.10
4	19	Married	0	15.00	Under \$25	7.20
5	19	Married	0	17.00	Under \$25	8.30
6	19	Married	0	14.00	Under \$25	7.20
7	19	Unmarried	0	19.00	Under \$25	9.40
8	20	Unmarried	1	14.00	Under \$25	7.00
9	20	Married	1	17.00	Under \$25	8.60
10	20	Married	0	19.00	Under \$25	9.40
11	20	Married	0	24.00	Under \$25	11.40
12	20	Married	1	14.00	Under \$25	6.90

Data Editor Display Options

The View menu provides several display options for the Data Editor:

Fonts. This option controls the font characteristics of the data display.

Grid Lines. This option toggles the display of grid lines.

Value Labels. This option toggles between the display of actual data values and user-defined descriptive value labels. This option is available only in Data View.

Using Multiple Views

In Data View, you can create multiple views (panes) by using the splitters that are located below the horizontal scroll bar and to the right of the vertical scroll bar.

You can also use the Window menu to insert and remove pane splitters. To insert splitters:

- ▶ In Data View, from the menus choose:
Window
Split

Splitters are inserted above and to the left of the selected cell.

- If the top left cell is selected, splitters are inserted to divide the current view approximately in half, both horizontally and vertically.

- If any cell other than the top cell in the first column is selected, a horizontal pane splitter is inserted above the selected cell.
- If any cell other than the first cell in the top row is selected, a vertical pane splitter is inserted to the left of the selected cell.

Data Editor Printing

A data file is printed as it appears on the screen.

- The information in the currently displayed view is printed. In Data View, the data are printed. In Variable View, data definition information is printed.
- Grid lines are printed if they are currently displayed in the selected view.
- Value labels are printed in Data View if they are currently displayed. Otherwise, the actual data values are printed.

Use the View menu in the Data Editor window to display or hide grid lines and toggle between the display of data values and value labels.

To Print Data Editor Contents

- ▶ Make the Data Editor the active window.
- ▶ Click the tab for the view that you want to print.
- ▶ From the menus choose:
 - File
 - Print...

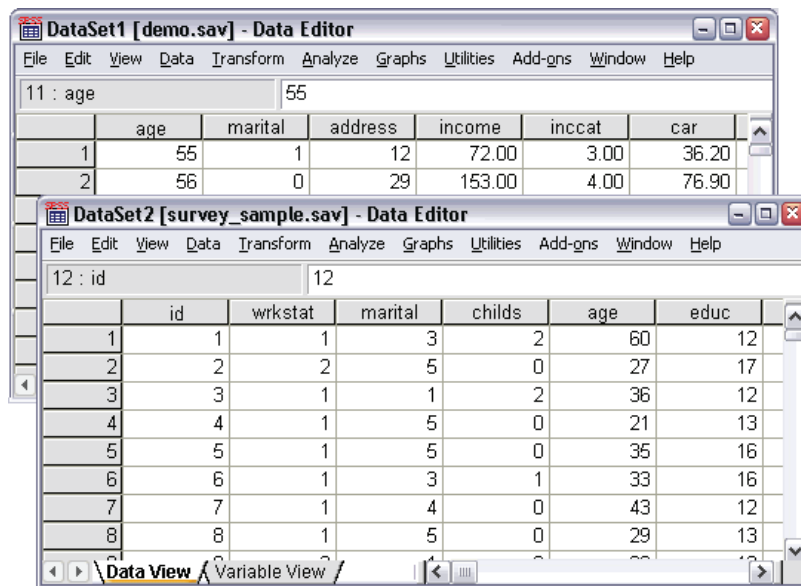
Working with Multiple Data Sources

Starting with version 14.0, multiple data sources can be open at the same time, making it easier to:

- Switch back and forth between data sources.
- Compare the contents of different data sources.
- Copy and paste data between data sources.
- Create multiple subsets of cases and/or variables for analysis.
- Merge multiple data sources from various data formats (for example, spreadsheet, database, text data) without saving each data source first.

Basic Handling of Multiple Data Sources

Figure 6-1
Two data sources open at same time

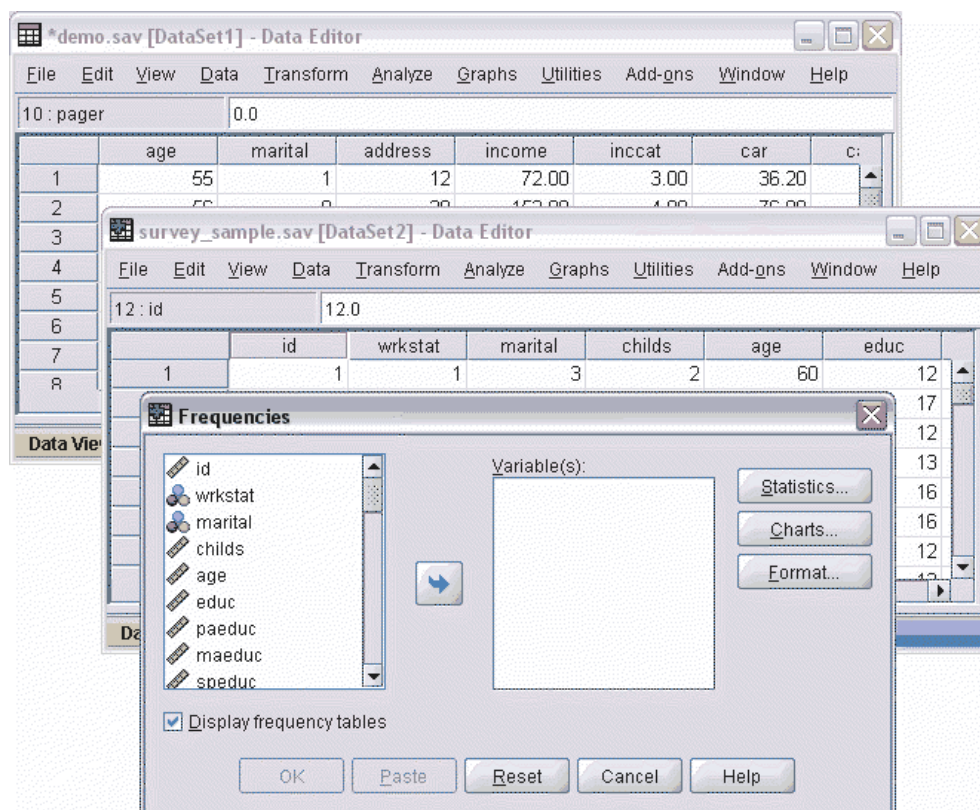


By default, each data source that you open is displayed in a new Data Editor window. (See [General Options](#) for information on changing the default behavior to only display one dataset at a time, in a single Data Editor window.)

- Any previously open data sources remain open and available for further use.
- When you first open a data source, it automatically becomes the **active dataset**.

- You can change the active dataset simply by clicking anywhere in the Data Editor window of the data source that you want to use or by selecting the Data Editor window for that data source from the Window menu.
- Only the variables in the active dataset are available for analysis.

Figure 6-2
Variable list containing variables in the active dataset



- You cannot change the active dataset when any dialog box that accesses the data is open (including all dialog boxes that display variable lists).
- At least one Data Editor window must be open during a session. When you close the last open Data Editor window, SPSS Statistics automatically shuts down, prompting you to save changes first.

Working with Multiple Datasets in Command Syntax

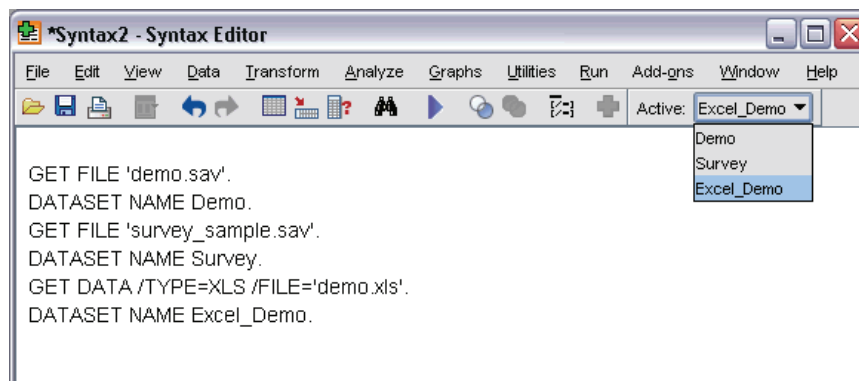
If you use command syntax to open data sources (for example, `GET FILE`, `GET DATA`), you need to use the `DATASET NAME` command to name each dataset explicitly in order to have more than one data source open at the same time.

When working with command syntax, the active dataset name is displayed on the toolbar of the syntax window. All of the following actions can change the active dataset:

- Use the `DATASET ACTIVATE` command.

- Click anywhere in the Data Editor window of a dataset.
- Select a dataset name from the toolbar in the syntax window.

Figure 6-3
Open datasets displayed on syntax window toolbar



Copying and Pasting Information between Datasets

You can copy both data and variable definition attributes from one dataset to another dataset in basically the same way that you copy and paste information within a single data file.

- Copying and pasting selected data cells in Data View pastes only the data values, with no variable definition attributes.
- Copying and pasting an entire variable in Data View by selecting the variable name at the top of the column pastes all of the data and all of the variable definition attributes for that variable.
- Copying and pasting variable definition attributes or entire variables in Variable View pastes the selected attributes (or the entire variable definition) but does not paste any data values.

Renaming Datasets

When you open a data source through the menus and dialog boxes, each data source is automatically assigned a dataset name of *DataSetn*, where *n* is a sequential integer value, and when you open a data source using command syntax, no dataset name is assigned unless you explicitly specify one with `DATASET NAME`. To provide more descriptive dataset names:

- ▶ From the menus in the Data Editor window for the dataset whose name you want to change choose:
 - File
 - Rename Dataset...
- ▶ Enter a new dataset name that conforms to variable naming rules. [For more information, see Variable Names in Chapter 5 on p. 72.](#)

Suppressing Multiple Datasets

If you prefer to have only one dataset available at a time and want to suppress the multiple dataset feature:

- ▶ From the menus choose:
 - Edit
 - Options...
- ▶ Click the General tab.

Select (check) Open only one dataset at a time.

For more information, see [General Options in Chapter 48 on p. 494](#).

Data Preparation

Once you've opened a data file or entered data in the Data Editor, you can start creating reports, charts, and analyses without any additional preliminary work. However, there are some additional data preparation features that you may find useful, including the ability to:

- Assign variable properties that describe the data and determine how certain values should be treated.
- Identify cases that may contain duplicate information and exclude those cases from analyses or delete them from the data file.
- Create new variables with a few distinct categories that represent ranges of values from variables with a large number of possible values.

Variable Properties

Data entered in the Data Editor in Data View or read from an external file format (such as an Excel spreadsheet or a text data file) lack certain variable properties that you may find very useful, including:

- Definition of descriptive value labels for numeric codes (for example, 0 = *Male* and 1 = *Female*).
- Identification of missing values codes (for example, 99 = *Not applicable*).
- Assignment of measurement level (nominal, ordinal, or scale).

All of these variable properties (and others) can be assigned in Variable View in the Data Editor. There are also several utilities that can assist you in this process:

- **Define Variable Properties** can help you define descriptive value labels and missing values. This is particularly useful for categorical data with numeric codes used for category values.
- **Copy Data Properties** provides the ability to use an existing SPSS Statistics data file as a template for file and variable properties in the current data file. This is particularly useful if you frequently use external-format data files that contain similar content (such as monthly reports in Excel format).

Defining Variable Properties

Define Variable Properties is designed to assist you in the process of creating descriptive value labels for categorical (nominal, ordinal) variables. Define Variable Properties:

- Scans the actual data values and lists all unique data values for each selected variable.

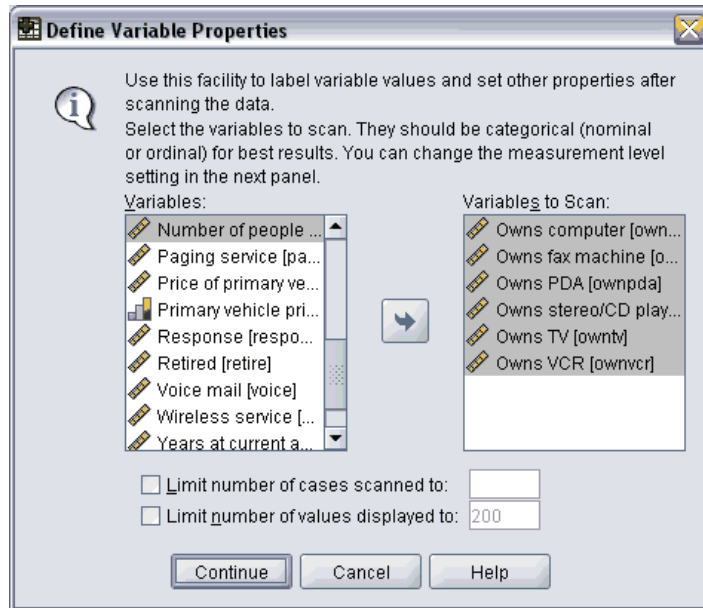
- Identifies unlabeled values and provides an “auto-label” feature.
- Provides the ability to copy defined value labels from another variable to the selected variable or from the selected variable to multiple additional variables.

Note: To use Define Variable Properties without first scanning cases, enter 0 for the number of cases to scan.

To Define Variable Properties

- ▶ From the menus choose:
 - Data
 - Define Variable Properties...

Figure 7-1
Initial dialog box for selecting variables to define

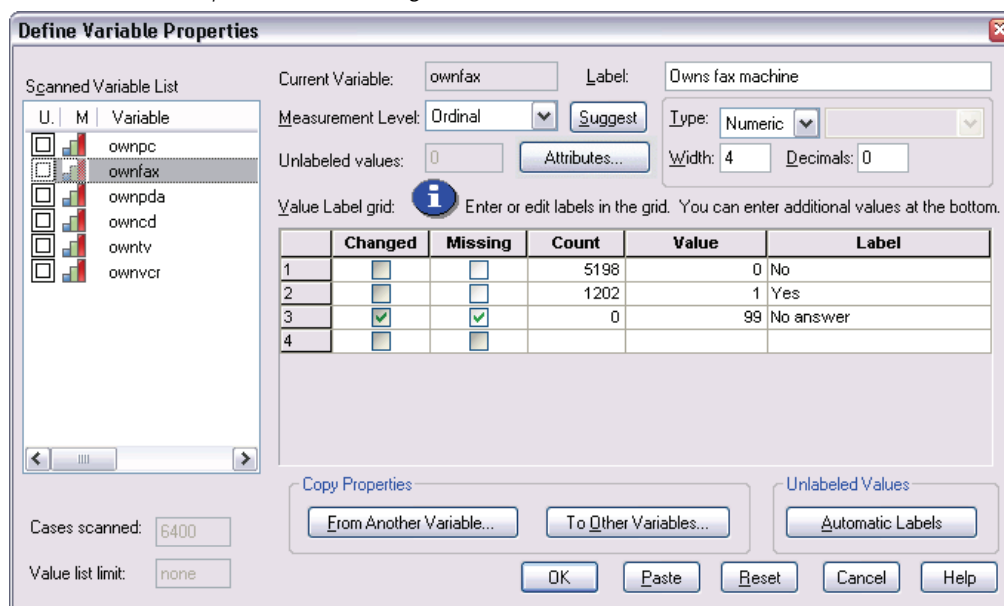


- ▶ Select the numeric or string variables for which you want to create value labels or define or change other variable properties, such as missing values or descriptive variable labels.
- ▶ Specify the number of cases to scan to generate the list of unique values. This is particularly useful for data files with a large number of cases for which a scan of the complete data file might take a significant amount of time.
- ▶ Specify an upper limit for the number of unique values to display. This is primarily useful to prevent listing hundreds, thousands, or even millions of values for scale (continuous interval, ratio) variables.
- ▶ Click Continue to open the main Define Variable Properties dialog box.
- ▶ Select a variable for which you want to create value labels or define or change other variable properties.
- ▶ Enter the label text for any unlabeled values that are displayed in the Value Label grid.

- ▶ If there are values for which you want to create value labels but those values are not displayed, you can enter values in the *Value* column below the last scanned value.
- ▶ Repeat this process for each listed variable for which you want to create value labels.
- ▶ Click OK to apply the value labels and other variable properties.

Defining Value Labels and Other Variable Properties

Figure 7-2
Define Variable Properties, main dialog box



The Define Variable Properties main dialog box provides the following information for the scanned variables:

Scanned Variable List. For each scanned variable, a check mark in the *Unlabeled (U.)* column indicates that the variable contains values without assigned value labels.

To sort the variable list to display all variables with unlabeled values at the top of the list:

- ▶ Click the *Unlabeled* column heading under Scanned Variable List.

You can also sort by variable name or measurement level by clicking the corresponding column heading under Scanned Variable List.

Value Label Grid

- **Label.** Displays any value labels that have already been defined. You can add or change labels in this column.
- **Value.** Unique values for each selected variable. This list of unique values is based on the number of scanned cases. For example, if you scanned only the first 100 cases in the data file, then the list reflects only the unique values present in those cases. If the data file has already

been sorted by the variable for which you want to assign value labels, the list may display far fewer unique values than are actually present in the data.

- **Count.** The number of times each value occurs in the scanned cases.
- **Missing.** Values defined as representing missing data. You can change the missing values designation of the category by clicking the check box. A check indicates that the category is defined as a user-missing category. If a variable already has a range of values defined as user-missing (for example, 90-99), you cannot add or delete missing values categories for that variable with Define Variable Properties. You can use Variable View in the Data Editor to modify the missing values categories for variables with missing values ranges. [For more information, see Missing Values in Chapter 5 on p. 77.](#)
- **Changed.** Indicates that you have added or changed a value label.

Note: If you specified 0 for the number of cases to scan in the initial dialog box, the Value Label grid will initially be blank, except for any preexisting value labels and/or defined missing values categories for the selected variable. In addition, the Suggest button for the measurement level will be disabled.

Measurement Level. Value labels are primarily useful for categorical (nominal and ordinal) variables, and some procedures treat categorical and scale variables differently; so it is sometimes important to assign the correct measurement level. However, by default, all new numeric variables are assigned the scale measurement level. Thus, many variables that are in fact categorical may initially be displayed as scale.

If you are unsure of what measurement level to assign to a variable, click Suggest.

Copy Properties. You can copy value labels and other variable properties from another variable to the currently selected variable or from the currently selected variable to one or more other variables.

Unlabeled Values. To create labels for unlabeled values automatically, click Automatic Labels.

Variable Label and Display Format

You can change the descriptive variable label and the display format.

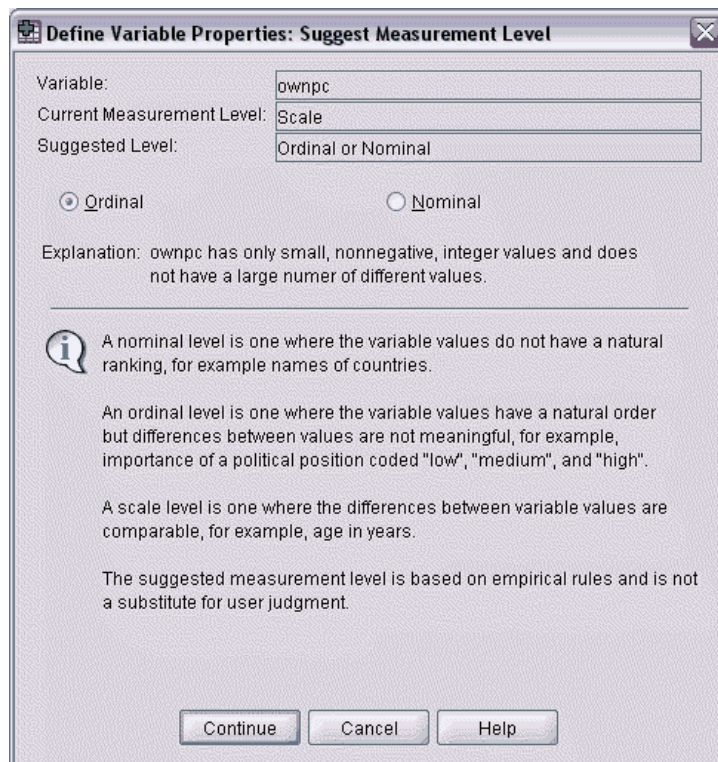
- You cannot change the variable's fundamental type (string or numeric).
- For string variables, you can change only the variable label, not the display format.
- For numeric variables, you can change the numeric type (such as numeric, date, dollar, or custom currency), width (maximum number of digits, including any decimal and/or grouping indicators), and number of decimal positions.
- For numeric date format, you can select a specific date format (such as dd-mm-yyyy, mm/dd/yy, and yyyyddd)
- For numeric custom format, you can select one of five custom currency formats (CCA through CCE). [For more information, see Currency Options in Chapter 48 on p. 499.](#)

- An asterisk is displayed in the *Value* column if the specified width is less than the width of the scanned values or the displayed values for preexisting defined value labels or missing values categories.
- A period (.) is displayed if the scanned values or the displayed values for preexisting defined value labels or missing values categories are invalid for the selected display format type. For example, an internal numeric value of less than 86,400 is invalid for a date format variable.

Assigning the Measurement Level

When you click **Suggest** for the measurement level in the Define Variable Properties main dialog box, the current variable is evaluated based on the scanned cases and defined value labels, and a measurement level is suggested in the Suggest Measurement Level dialog box that opens. The Explanation area provides a brief description of the criteria used to provide the suggested measurement level.

Figure 7-3
Suggest Measurement Level dialog box



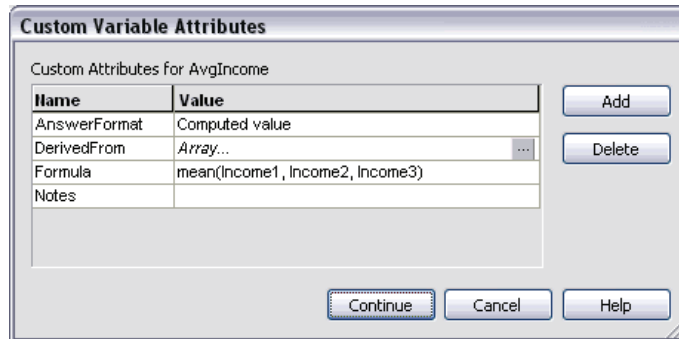
Note: Values defined as representing missing values are not included in the evaluation for measurement level. For example, the explanation for the suggested measurement level may indicate that the suggestion is in part based on the fact that the variable contains no negative values, whereas it may in fact contain negative values—but those values are already defined as missing values.

- Click **Continue** to accept the suggested level of measurement or **Cancel** to leave the measurement level unchanged.

Custom Variable Attributes

The Attributes button in Define Variable Properties opens the Custom Variable Attributes dialog box. In addition to the standard variable attributes, such as value labels, missing values, and measurement level, you can create your own custom variable attributes. Like standard variable attributes, these custom attributes are saved with SPSS Statistics data files.

Figure 7-4
Custom Variable Attributes



Name. Attribute names must follow the same rules as variable names. [For more information, see Variable Names in Chapter 5 on p. 72.](#)

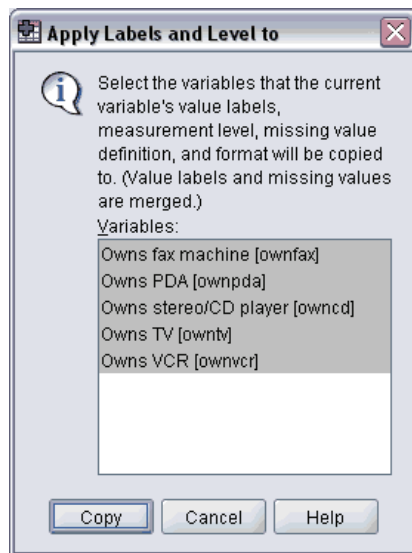
Value. The value assigned to the attribute for the selected variable.

- Attribute names that begin with a dollar sign are reserved and cannot be modified. You can view the contents of a reserved attribute by clicking on the button in the desired cell.
- The text *Array...*, displayed in a Value cell, indicates that this is an **attribute array**, an attribute that contains multiple values. Click the button in the cell to display the list of values.

Copying Variable Properties

The Apply Labels and Level dialog box is displayed when you click From Another Variable or To Other Variables in the Define Variable Properties main dialog box. It displays all of the scanned variables that match the current variable's type (numeric or string). For string variables, the defined width must also match.

Figure 7-5
Apply Labels and Level dialog box



- ▶ Select a single variable from which to copy value labels and other variable properties (except variable label).

or

- ▶ Select one or more variables to which to copy value labels and other variable properties.
- ▶ Click Copy to copy the value labels and the measurement level.
 - Existing value labels and missing value categories for target variable(s) are not replaced.
 - Value labels and missing value categories for values not already defined for the target variable(s) are added to the set of value labels and missing value categories for the target variable(s).
 - The measurement level for the target variable(s) is always replaced.
 - If either the source or target variable has a defined range of missing values, missing values definitions are not copied.

Multiple Response Sets

Custom Tables and the Chart Builder support a special kind of “variable” called a **multiple response set**. Multiple response sets aren’t really “variables” in the normal sense. You can’t see them in the Data Editor, and other procedures don’t recognize them. Multiple response sets use multiple variables to record responses to questions where the respondent can give more than one answer. Multiple response sets are treated like categorical variables, and most of the things you can do with categorical variables, you can also do with multiple response sets.

Multiple response sets are constructed from multiple variables in the data file. A multiple response set is a special construct within a data file. You can define and save multiple response sets in SPSS Statistics data files, but you cannot import or export multiple response sets from/to other file formats. You can copy multiple response sets from other SPSS Statistics data files using

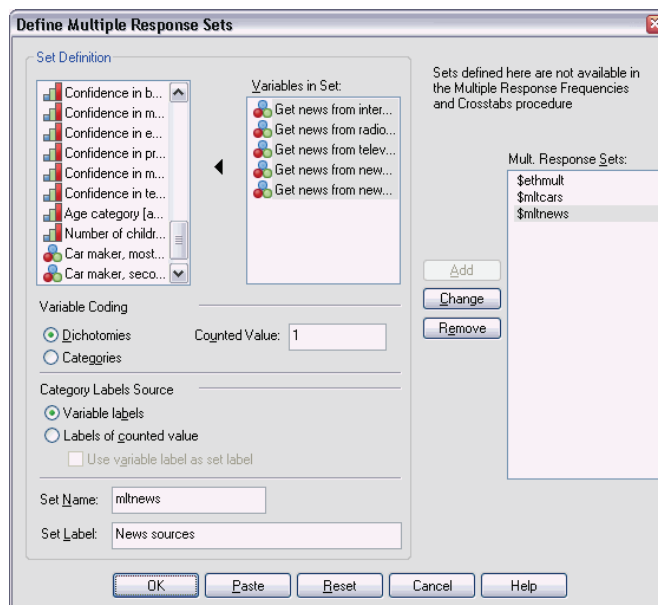
Copy Data Properties, which is accessed from the Data menu in the Data Editor window. For more information, see [Copying Data Properties on p. 106](#).

Defining Multiple Response Sets

To define multiple response sets:

- ▶ From the menus, choose:
 - Data
 - Define Multiple Response Sets...

Figure 7-6
Define Multiple Response Sets dialog box



- ▶ Select two or more variables. If your variables are coded as dichotomies, indicate which value you want to have counted.
- ▶ Enter a unique name for each multiple response set. The name can be up to 63 bytes long. A dollar sign is automatically added to the beginning of the set name.
- ▶ Enter a descriptive label for the set. (This is optional.)
- ▶ Click Add to add the multiple response set to the list of defined sets.

Dichotomies

A multiple dichotomy set typically consists of multiple dichotomous variables: variables with only two possible values of a yes/no, present/absent, checked/not checked nature. Although the variables may not be strictly dichotomous, all of the variables in the set are coded the same way, and the Counted Value represents the positive/present/checked condition.

For example, a survey asks the question, “Which of the following sources do you rely on for news?” and provides five possible responses. The respondent can indicate multiple choices by checking a box next to each choice. The five responses become five variables in the data file, coded 0 for *No* (not checked) and 1 for *Yes* (checked). In the multiple dichotomy set, the Counted Value is 1.

The sample data file *survey_sample.sav* already has three defined multiple response sets. *\$mltnews* is a multiple dichotomy set.

- ▶ Select (click) *\$mltnews* in the Mult. Response Sets list.

This displays the variables and settings used to define this multiple response set.

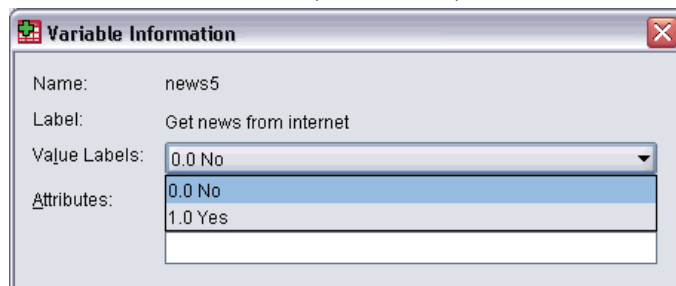
- The Variables in Set list displays the five variables used to construct the multiple response set.
- The Variable Coding group indicates that the variables are dichotomous.
- The Counted Value is 1.

- ▶ Select (click) one of the variables in the Variables in Set list.

- ▶ Right-click the variable and select Variable Information from the pop-up context menu.

- ▶ In the Variable Information window, click the arrow on the Value Labels drop-down list to display the entire list of defined value labels.

Figure 7-7
Variable information for multiple dichotomy source variable



The value labels indicate that the variable is a dichotomy with values of 0 and 1, representing *No* and *Yes*, respectively. All five variables in the list are coded the same way, and the value of 1 (the code for *Yes*) is the counted value for the multiple dichotomy set.

Categories

A multiple category set consists of multiple variables, all coded the same way, often with many possible response categories. For example, a survey item states, “Name up to three nationalities that best describe your ethnic heritage.” There may be hundreds of possible responses, but for coding purposes the list is limited to the 40 most common nationalities, with everything else relegated to an “other” category. In the data file, the three choices become three variables, each with 41 categories (40 coded nationalities and one “other” category).

In the sample data file, *\$ethmult* and *\$mltcars* are multiple category sets.

Category Label Source

For multiple dichotomies, you can control how sets are labeled.

- **Variable labels.** Uses the defined variable labels (or variable names for variables without defined variable labels) as the set category labels. For example, if all of the variables in the set have the same value label (or no defined value labels) for the counted value (for example, *Yes*), then you should use the variable labels as the set category labels.
- **Labels of counted values.** Uses the defined value labels of the counted values as set category labels. Select this option only if all variables have a defined value label for the counted value and the value label for the counted value is different for each variable.
- **Use variable label as set label.** If you select Label of counted values, you can also use the variable label for the first variable in the set with a defined variable label as the set label. If none of the variables in the set have defined variable labels, the name of the first variable in the set is used as the set label.

Copying Data Properties

The Copy Data Properties Wizard provides the ability to use an external SPSS Statistics data file as a template for defining file and variable properties in the active dataset. You can also use variables in the active dataset as templates for other variables in the active dataset. You can:

- Copy selected file properties from an external data file or open dataset to the active dataset. File properties include documents, file labels, multiple response sets, variable sets, and weighting.
- Copy selected variable properties from an external data file or open dataset to matching variables in the active dataset. Variable properties include value labels, missing values, level of measurement, variable labels, print and write formats, alignment, and column width (in the Data Editor).
- Copy selected variable properties from one variable in either an external data file, open dataset, or the active dataset to many variables in the active dataset.
- Create new variables in the active dataset based on selected variables in an external data file or open dataset.

When copying data properties, the following general rules apply:

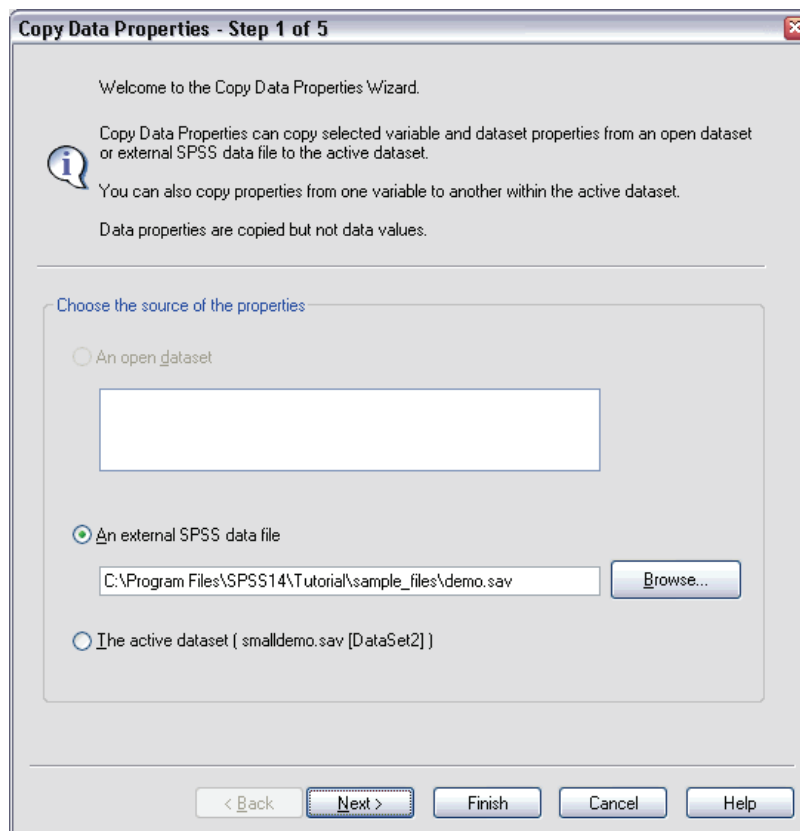
- If you use an external data file as the source data file, it must be a data file in SPSS Statistics format.
- If you use the active dataset as the source data file, it must contain at least one variable. You cannot use a completely blank active dataset as the source data file.
- Undefined (empty) properties in the source dataset do not overwrite defined properties in the active dataset.
- Variable properties are copied from the source variable only to target variables of a matching type—string (alphanumeric) or numeric (including numeric, date, and currency).

Note: Copy Data Properties replaces Apply Data Dictionary, formerly available on the File menu.

To Copy Data Properties

- ▶ From the menus in the Data Editor window choose:
Data
Copy Data Properties...

Figure 7-8
Copy Data Properties Wizard: Step 1

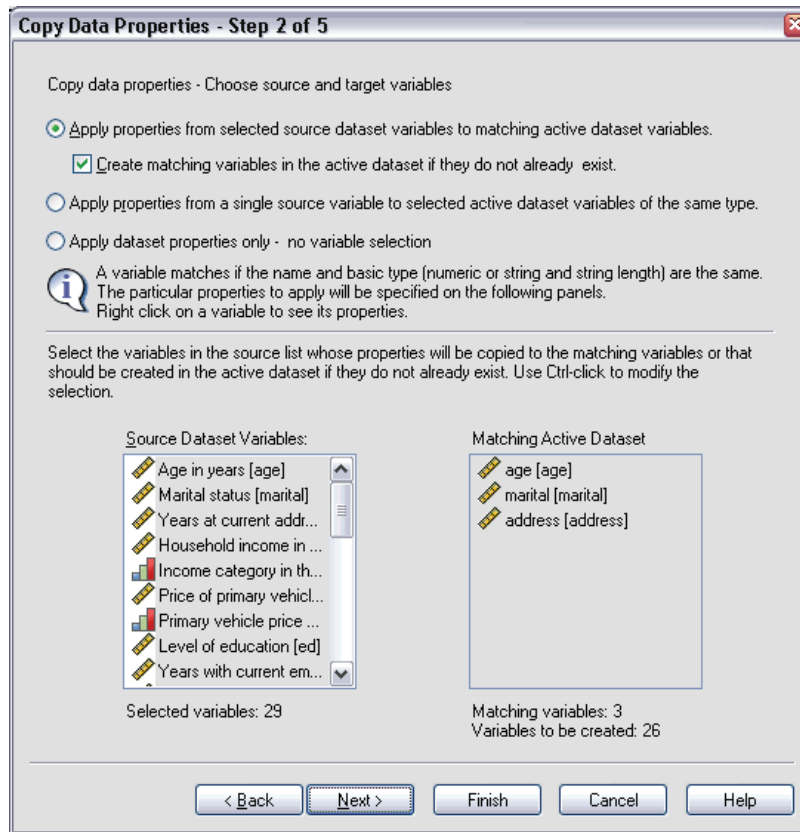


- ▶ Select the data file with the file and/or variable properties that you want to copy. This can be a currently open dataset, an external SPSS Statistics data file, or the active dataset.
- ▶ Follow the step-by-step instructions in the Copy Data Properties Wizard.

Selecting Source and Target Variables

In this step, you can specify the source variables containing the variable properties that you want to copy and the target variables that will receive those variable properties.

Figure 7-9
Copy Data Properties Wizard: Step 2



Apply properties from selected source dataset variables to matching active dataset variables.

Variable properties are copied from one or more selected source variables to matching variables in the active dataset. Variables “match” if both the variable name and type (string or numeric) are the same. For string variables, the defined length must also be the same. By default, only matching variables are displayed in the two variable lists.

- **Create matching variables in the active dataset if they do not already exist.** This updates the source list to display all variables in the source data file. If you select source variables that do not exist in the active dataset (based on variable name), new variables will be created in the active dataset with the variable names and properties from the source data file.

If the active dataset contains no variables (a blank, new dataset), all variables in the source data file are displayed and new variables based on the selected source variables are automatically created in the active dataset.

Apply properties from a single source variable to selected active dataset variables of the same type.

Variable properties from a single selected variable in the source list can be applied to one or more selected variables in the active dataset list. Only variables of the same type (numeric or string) as the selected variable in the source list are displayed in the active dataset list. For string variables, only strings of the same defined length as the source variable are displayed. This option is not available if the active dataset contains no variables.

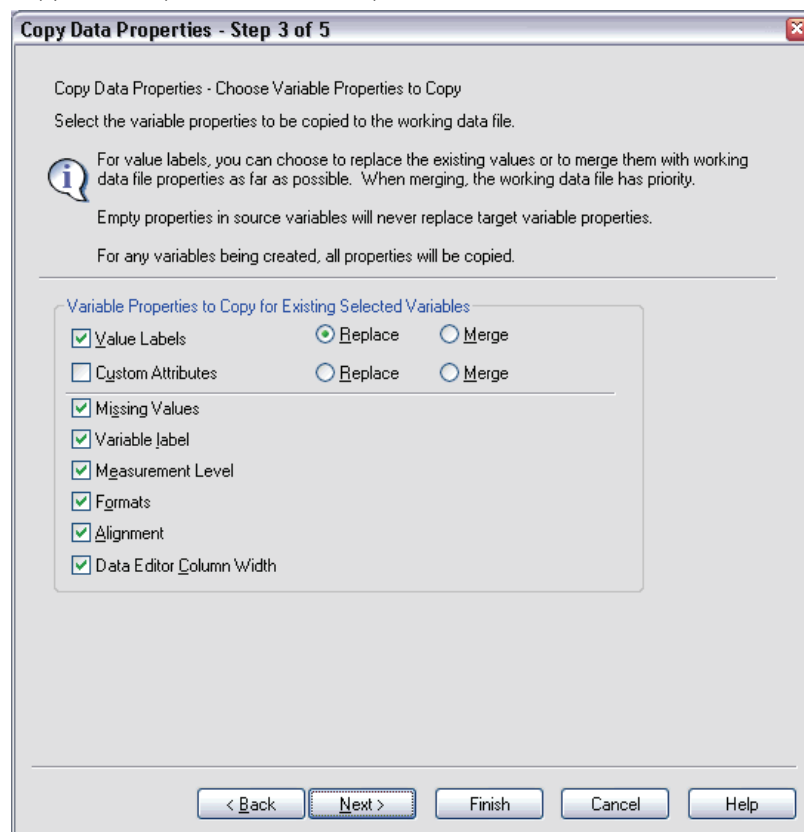
Note: You cannot create new variables in the active dataset with this option.

Apply dataset properties only—no variable selection. Only file properties (for example, documents, file label, weight) will be applied to the active dataset. No variable properties will be applied. This option is not available if the active dataset is also the source data file.

Choosing Variable Properties to Copy

You can copy selected variable properties from the source variables to the target variables. Undefined (empty) properties in the source variables do not overwrite defined properties in the target variables.

Figure 7-10
Copy Data Properties Wizard: Step 3



Value Labels. Value labels are descriptive labels associated with data values. Value labels are often used when numeric data values are used to represent non-numeric categories (for example, codes of 1 and 2 for *Male* and *Female*). You can replace or merge value labels in the target variables.

- Replace deletes any defined value labels for the target variable and replaces them with the defined value labels from the source variable.
- Merge merges the defined value labels from the source variable with any existing defined value label for the target variable. If the same value has a defined value label in both the source and target variables, the value label in the target variable is unchanged.

Custom Attributes. User-defined custom variable attributes. For more information, see [Custom Variable Attributes in Chapter 5 on p. 80](#).

- Replace deletes any custom attributes for the target variable and replaces them with the defined attributes from the source variable.
- Merge merges the defined attributes from the source variable with any existing defined attributes for the target variable.

Missing Values. Missing values are values identified as representing missing data (for example, 98 for *Do not know* and 99 for *Not applicable*). Typically, these values also have defined value labels that describe what the missing value codes stand for. Any existing defined missing values for the target variable are deleted and replaced with the defined missing values from the source variable.

Variable Label. Descriptive variable labels can contain spaces and reserved characters not allowed in variable names. If you're copying variable properties from a single source variable to multiple target variables, you might want to think twice before selecting this option.

Measurement Level. The measurement level can be nominal, ordinal, or scale. For those procedures that distinguish between different measurement levels, nominal and ordinal are both considered **categorical**.

Formats. For numeric variables, this controls numeric type (such as numeric, date, or currency), width (total number of displayed characters, including leading and trailing characters and decimal indicator), and number of decimal places displayed. This option is ignored for string variables.

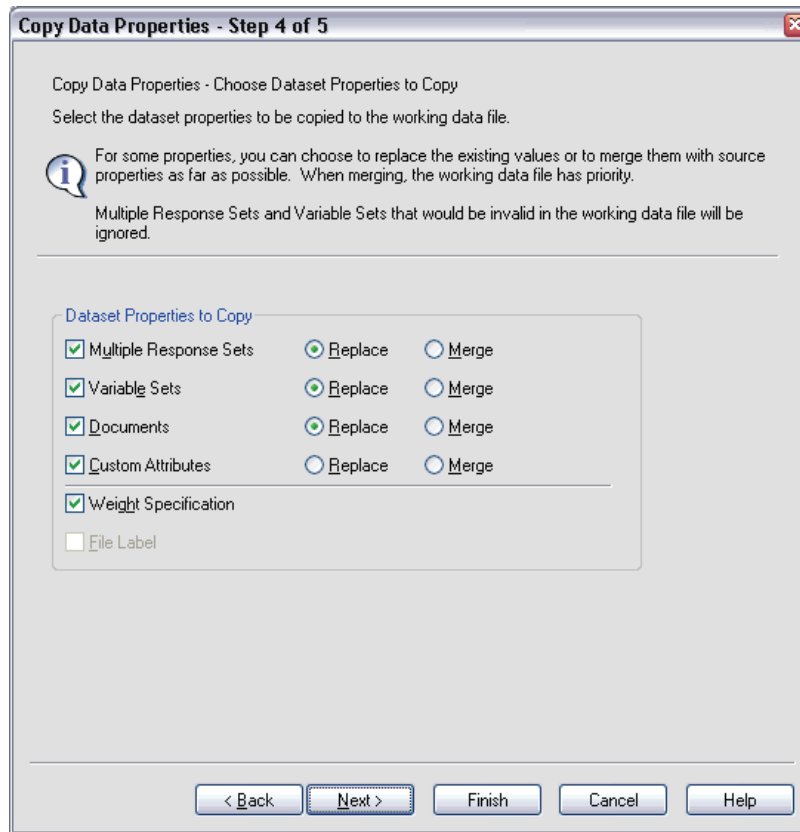
Alignment. This affects only alignment (left, right, center) in Data View in the Data Editor.

Data Editor Column Width. This affects only column width in Data View in the Data Editor.

Copying Dataset (File) Properties

You can apply selected, global dataset properties from the source data file to the active dataset. (This is not available if the active dataset is the source data file.)

Figure 7-11
Copy Data Properties Wizard: Step 4



Multiple Response Sets. Applies multiple response set definitions from the source data file to the active dataset.

- Multiple response sets in the source data file that contain variables that do not exist in the active dataset are ignored unless those variables will be created based on specifications in step 2 (Selecting Source and Target Variables) in the Copy Data Properties Wizard.
- Replace deletes all multiple response sets in the active dataset and replaces them with the multiple response sets from the source data file.
- Merge adds multiple response sets from the source data file to the collection of multiple response sets in the active dataset. If a set with the same name exists in both files, the existing set in the active dataset is unchanged.

Variable Sets. Variable sets are used to control the list of variables that are displayed in dialog boxes. Variable sets are defined by selecting Define Sets from the Utilities menu.

- Sets in the source data file that contain variables that do not exist in the active dataset are ignored unless those variables will be created based on specifications in step 2 (Selecting Source and Target Variables) in the Copy Data Properties Wizard.

- **Replace** deletes any existing variable sets in the active dataset, replacing them with variable sets from the source data file.
- **Merge** adds variable sets from the source data file to the collection of variable sets in the active dataset. If a set with the same name exists in both files, the existing set in the active dataset is unchanged.

Documents. Notes appended to the data file via the `DOCUMENT` command.

- **Replace** deletes any existing documents in the active dataset, replacing them with the documents from the source data file.
- **Merge** combines documents from the source and active dataset. Unique documents in the source file that do not exist in the active dataset are added to the active dataset. All documents are then sorted by date.

Custom Attributes. Custom data file attributes, typically created with the `DATAFILE ATTRIBUTE` command in command syntax.

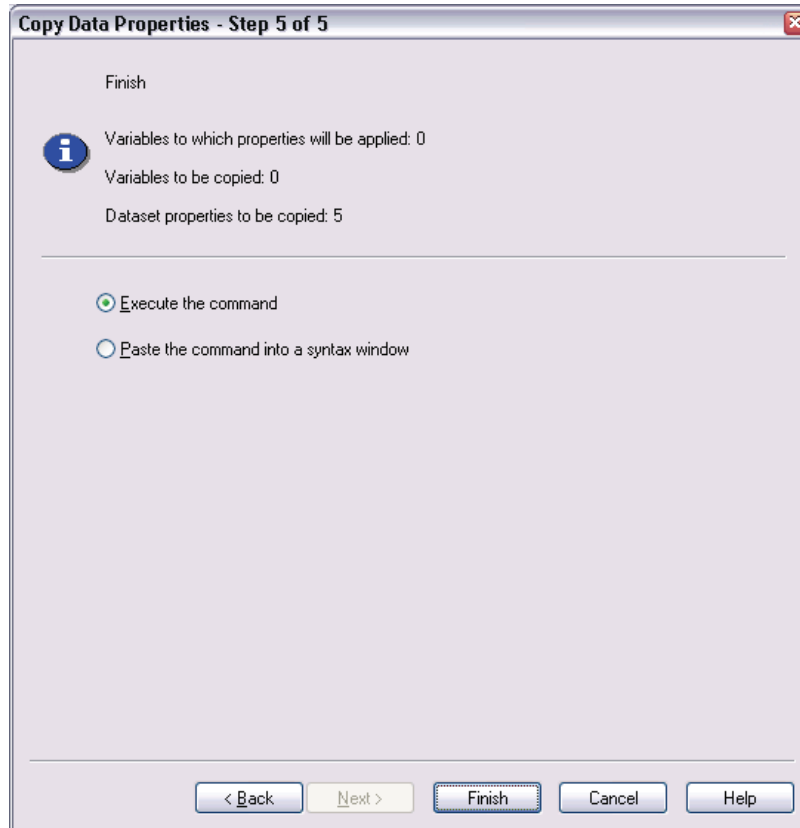
- **Replace** deletes any existing custom data file attributes in the active dataset, replacing them with the data file attributes from the source data file.
- **Merge** combines data file attributes from the source and active dataset. Unique attribute names in the source file that do not exist in the active dataset are added to the active dataset. If the same attribute name exists in both data files, the named attribute in the active dataset is unchanged.

Weight Specification. Weights cases by the current weight variable in the source data file if there is a matching variable in the active dataset. This overrides any weighting currently in effect in the active dataset.

File Label. Descriptive label applied to a data file with the `FILE LABEL` command.

Results

Figure 7-12
Copy Data Properties Wizard: Step 5



The last step in the Copy Data Properties Wizard provides information on the number of variables for which variable properties will be copied from the source data file, the number of new variables that will be created, and the number of dataset (file) properties that will be copied.

You can also choose to paste the generated command syntax into a syntax window and save the syntax for later use.

Identifying Duplicate Cases

“Duplicate” cases may occur in your data for many reasons, including:

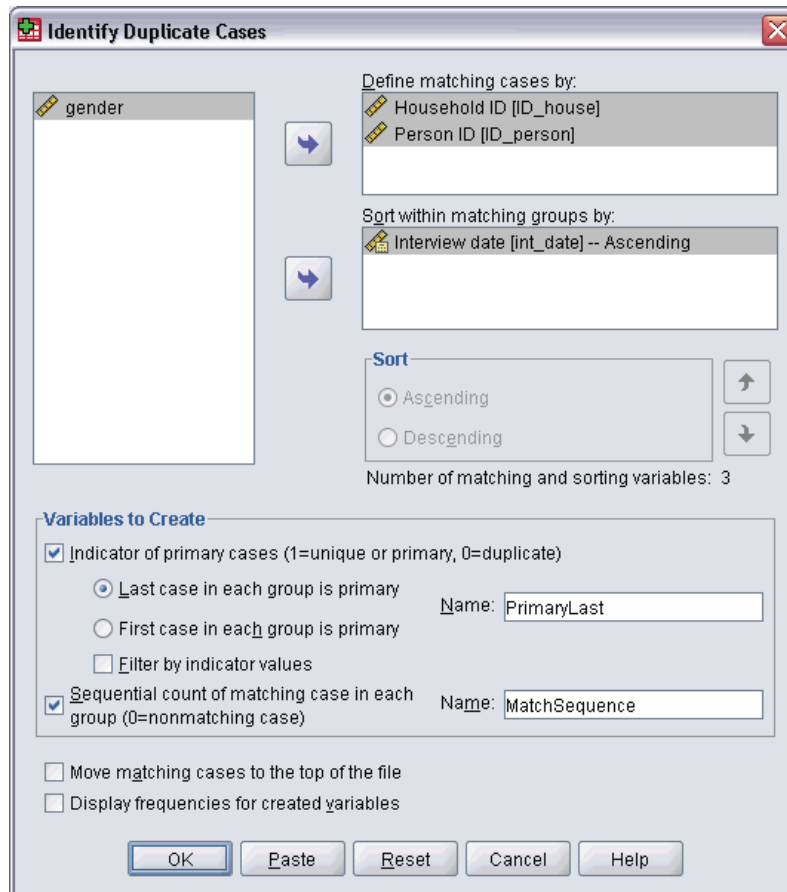
- Data entry errors in which the same case is accidentally entered more than once.
- Multiple cases share a common primary ID value but have different secondary ID values, such as family members who all live in the same house.
- Multiple cases represent the same case but with different values for variables other than those that identify the case, such as multiple purchases made by the same person or company for different products or at different times.

Identify Duplicate Cases allows you to define *duplicate* almost any way that you want and provides some control over the automatic determination of primary versus duplicate cases.

To Identify and Flag Duplicate Cases

- ▶ From the menus choose:
Data
Identify Duplicate Cases...
 - ▶ Select one or more variables that identify matching cases.
 - ▶ Select one or more of the options in the Variables to Create group.
- Optionally, you can:
- ▶ Select one or more variables to sort cases within groups defined by the selected matching cases variables. The sort order defined by these variables determines the “first” and “last” case in each group. Otherwise, the original file order is used.
 - ▶ Automatically filter duplicate cases so that they won’t be included in reports, charts, or calculations of statistics.

Figure 7-13
Identify Duplicate Cases dialog box



Define matching cases by. Cases are considered duplicates if their values match for *all* selected variables. If you want to identify only cases that are a 100% match in all respects, select all of the variables.

Sort within matching groups by. Cases are automatically sorted by the variables that define matching cases. You can select additional sorting variables that will determine the sequential order of cases in each matching group.

- For each sort variable, you can sort in ascending or descending order.
- If you select multiple sort variables, cases are sorted by each variable within categories of the preceding variable in the list. For example, if you select *date* as the first sorting variable and *amount* as the second sorting variable, cases will be sorted by amount within each date.
- Use the up and down arrow buttons to the right of the list to change the sort order of the variables.
- The sort order determines the “first” and “last” case within each matching group, which determines the value of the optional primary indicator variable. For example, if you want to filter out all but the most recent case in each matching group, you could sort cases within the group in ascending order of a date variable, which would make the most recent date the last date in the group.

Indicator of primary cases. Creates a variable with a value of 1 for all unique cases and the case identified as the primary case in each group of matching cases and a value of 0 for the nonprimary duplicates in each group.

- The primary case can be either the last or first case in each matching group, as determined by the sort order within the matching group. If you don’t specify any sort variables, the original file order determines the order of cases within each group.
- You can use the indicator variable as a **filter variable** to exclude nonprimary duplicates from reports and analyses without deleting those cases from the data file.

Sequential count of matching cases in each group. Creates a variable with a sequential value from 1 to *n* for cases in each matching group. The sequence is based on the current order of cases in each group, which is either the original file order or the order determined by any specified sort variables.

Move matching cases to the top. Sorts the data file so that all groups of matching cases are at the top of the data file, making it easy to visually inspect the matching cases in the Data Editor.

Display frequencies for created variables. Frequency tables containing counts for each value of the created variables. For example, for the primary indicator variable, the table would show the number of cases with a value 0 for that variable, which indicates the number of duplicates, and the number of cases with a value of 1 for that variable, which indicates the number of unique and primary cases.

Missing Values. For numeric variables, the system-missing value is treated like any other value—cases with the system-missing value for an identifier variable are treated as having matching values for that variable. For string variables, cases with no value for an identifier variable are treated as having matching values for that variable.

Visual Binning

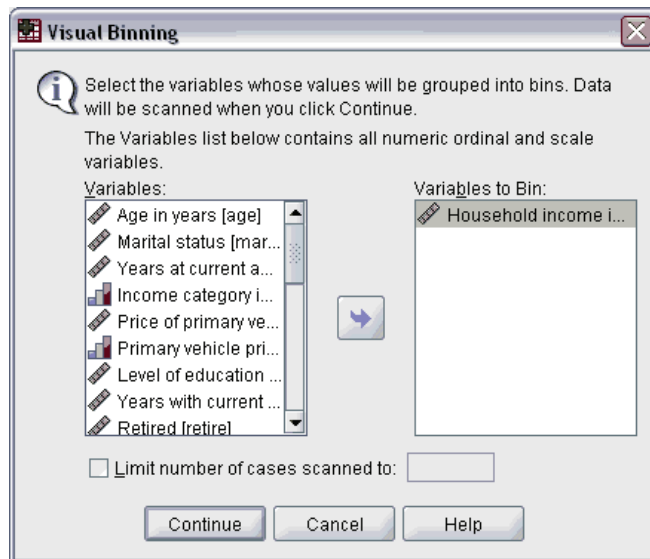
Visual Binning is designed to assist you in the process of creating new variables based on grouping contiguous values of existing variables into a limited number of distinct categories. You can use Visual Binning to:

- Create categorical variables from continuous scale variables. For example, you could use a scale income variable to create a new categorical variable that contains income ranges.
- Collapse a large number of ordinal categories into a smaller set of categories. For example, you could collapse a rating scale of nine down to three categories representing low, medium, and high.

In the first step, you:

- ▶ Select the numeric scale and/or ordinal variables for which you want to create new categorical (binned) variables.

Figure 7-14
Initial dialog box for selecting variables to bin



Optionally, you can limit the number of cases to scan. For data files with a large number of cases, limiting the number of cases scanned can save time, but you should avoid this if possible because it will affect the distribution of values used in subsequent calculations in Visual Binning.

Note: String variables and nominal numeric variables are not displayed in the source variable list. Visual Binning requires numeric variables, measured on either a scale or ordinal level, since it assumes that the data values represent some logical order that can be used to group values in a meaningful fashion. You can change the defined measurement level of a variable in Variable View in the Data Editor. [For more information, see Variable Measurement Level in Chapter 5 on p. 73.](#)

To Bin Variables

- ▶ From the menus in the Data Editor window choose:
Transform
Visual Binning...
- ▶ Select the numeric scale and/or ordinal variables for which you want to create new categorical (binned) variables.
- ▶ Select a variable in the Scanned Variable List.
- ▶ Enter a name for the new binned variable. Variable names must be unique and must follow variable naming rules. [For more information, see Variable Names in Chapter 5 on p. 72.](#)
- ▶ Define the binning criteria for the new variable. [For more information, see Binning Variables on p. 117.](#)
- ▶ Click OK.

Binning Variables

Figure 7-15
Visual Binning, main dialog box

Visual Binning

Scanned Variable List:

L	Variable
1	Household income in thousands

Current Variable: income Name: Household income in thousands

Binned Variable: inccat2 Label: Household income in thousands (Binned)

Minimum: 9.00 Nonmissing Values Maximum: 1116.00

Grid:

	Value	Label
1	25.00	<25.00
2	50.00	25.00 - 49.00
3	75.00	50.00 - 74.00
4	HIGH	75.00+
5		

Upper Endpoints

Included (<=)

Excluded (<)

Make Cutpoints...

Make Labels

Reverse scale

OK Paste Reset Cancel Help

The Visual Binning main dialog box provides the following information for the scanned variables:

Scanned Variable List. Displays the variables you selected in the initial dialog box. You can sort the list by measurement level (scale or ordinal) or by variable label or name by clicking on the column headings.

Cases Scanned. Indicates the number of cases scanned. All scanned cases without user-missing or system-missing values for the selected variable are used to generate the distribution of values used in calculations in Visual Binning, including the histogram displayed in the main dialog box and cutpoints based on percentiles or standard deviation units.

Missing Values. Indicates the number of scanned cases with user-missing or system-missing values. Missing values are not included in any of the binned categories. [For more information, see User-Missing Values in Visual Binning on p. 123.](#)

Current Variable. The name and variable label (if any) for the currently selected variable that will be used as the basis for the new, binned variable.

Binned Variable. Name and optional variable label for the new, binned variable.

- **Name.** You must enter a name for the new variable. Variable names must be unique and must follow variable naming rules. [For more information, see Variable Names in Chapter 5 on p. 72.](#)
- **Label.** You can enter a descriptive variable label up to 255 characters long. The default variable label is the variable label (if any) or variable name of the source variable with (*Binned*) appended to the end of the label.

Minimum and Maximum. Minimum and maximum values for the currently selected variable, based on the scanned cases and not including values defined as user-missing.

Nonmissing Values. The histogram displays the distribution of nonmissing values for the currently selected variable, based on the scanned cases.

- After you define bins for the new variable, vertical lines on the histogram are displayed to indicate the cutpoints that define bins.
- You can click and drag the cutpoint lines to different locations on the histogram, changing the bin ranges.
- You can remove bins by dragging cutpoint lines off the histogram.

Note: The histogram (displaying nonmissing values), the minimum, and the maximum are based on the scanned values. If you do not include all cases in the scan, the true distribution may not be accurately reflected, particularly if the data file has been sorted by the selected variable. If you scan zero cases, no information about the distribution of values is available.

Grid. Displays the values that define the upper endpoints of each bin and optional value labels for each bin.

- **Value.** The values that define the upper endpoints of each bin. You can enter values or use *Make Cutpoints* to automatically create bins based on selected criteria. By default, a cutpoint with a value of *HIGH* is automatically included. This bin will contain any nonmissing values above the other cutpoints. The bin defined by the lowest cutpoint will include all nonmissing values lower than or equal to that value (or simply lower than that value, depending on how you define upper endpoints).
- **Label.** Optional, descriptive labels for the values of the new, binned variable. Since the values of the new variable will simply be sequential integers from 1 to *n*, labels that describe what the values represent can be very useful. You can enter labels or use *Make Labels* to automatically create value labels.

To Delete a Bin from the Grid

- ▶ Right-click on the either the *Value* or *Label* cell for the bin.
- ▶ From the pop-up context menu, select Delete Row.

Note: If you delete the *HIGH* bin, any cases with values higher than the last specified cutpoint value will be assigned the system-missing value for the new variable.

To Delete All Labels or Delete All Defined Bins

- ▶ Right-click anywhere in the grid.
- ▶ From the pop-up context menu select either Delete All Labels or Delete All Cutpoints.

Upper Endpoints. Controls treatment of upper endpoint values entered in the *Value* column of the grid.

- **Included (<=).** Cases with the value specified in the *Value* cell are included in the binned category. For example, if you specify values of 25, 50, and 75, cases with a value of exactly 25 will go in the first bin, since this will include all cases with values less than or equal to 25.
- **Excluded (<).** Cases with the value specified in the *Value* cell are not included in the binned category. Instead, they are included in the next bin. For example, if you specify values of 25, 50, and 75, cases with a value of exactly 25 will go in the second bin rather than the first, since the first bin will contain only cases with values less than 25.

Make Cutpoints. Generates binned categories automatically for equal width intervals, intervals with the same number of cases, or intervals based on standard deviations. This is not available if you scanned zero cases. [For more information, see Automatically Generating Binned Categories on p. 119.](#)

Make Labels. Generates descriptive labels for the sequential integer values of the new, binned variable, based on the values in the grid and the specified treatment of upper endpoints (included or excluded).

Reverse scale. By default, values of the new, binned variable are ascending sequential integers from 1 to n . Reversing the scale makes the values descending sequential integers from n to 1.

Copy Bins. You can copy the binning specifications from another variable to the currently selected variable or from the selected variable to multiple other variables. [For more information, see Copying Binned Categories on p. 121.](#)

Automatically Generating Binned Categories

The Make Cutpoints dialog box allows you to auto-generate binned categories based on selected criteria.

To Use the Make Cutpoints Dialog Box

- ▶ Select (click) a variable in the Scanned Variable List.
- ▶ Click Make Cutpoints.

- ▶ Select the criteria for generating cutpoints that will define the binned categories.
- ▶ Click Apply.

Figure 7-16
Make Cutpoints dialog box

Note: The Make Cutpoints dialog box is not available if you scanned zero cases.

Equal Width Intervals. Generates binned categories of equal width (for example, 1–10, 11–20, and 21–30) based on any two of the following three criteria:

- **First Cutpoint Location.** The value that defines the upper end of the lowest binned category (for example, a value of 10 indicates a range that includes all values up to 10).
- **Number of Cutpoints.** The number of binned categories is the number of cutpoints *plus one*. For example, 9 cutpoints generate 10 binned categories.
- **Width.** The width of each interval. For example, a value of 10 would bin *age in years* into 10-year intervals.

Equal Percentiles Based on Scanned Cases. Generates binned categories with an equal number of cases in each bin (using the **aempirical algorithm** for percentiles), based on either of the following criteria:

- **Number of Cutpoints.** The number of binned categories is the number of cutpoints *plus one*. For example, three cutpoints generate four percentile bins (quartiles), each containing 25% of the cases.
- **Width (%).** Width of each interval, expressed as a percentage of the total number of cases. For example, a value of 33.3 would produce three binned categories (two cutpoints), each containing 33.3% of the cases.

If the source variable contains a relatively small number of distinct values or a large number of cases with the same value, you may get fewer bins than requested. If there are multiple identical values at a cutpoint, they will all go into the same interval; so the actual percentages may not always be exactly equal.

Cutpoints at Mean and Selected Standard Deviations Based on Scanned Cases. Generates binned categories based on the values of the mean and standard deviation of the distribution of the variable.

- If you don't select any of the standard deviation intervals, two binned categories will be created, with the mean as the cutpoint dividing the bins.
- You can select any combination of standard deviation intervals based on one, two, and/or three standard deviations. For example, selecting all three would result in eight binned categories—six bins in one standard deviation intervals and two bins for cases more than three standard deviations above and below the mean.

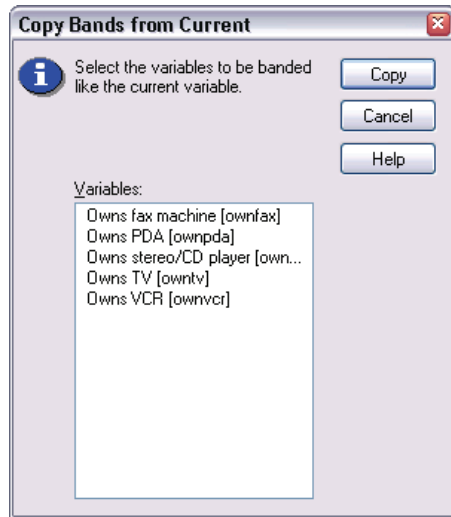
In a normal distribution, 68% of the cases fall within one standard deviation of the mean; 95%, within two standard deviations; and 99%, within three standard deviations. Creating binned categories based on standard deviations may result in some defined bins outside of the actual data range and even outside of the range of possible data values (for example, a negative salary range).

Note: Calculations of percentiles and standard deviations are based on the scanned cases. If you limit the number of cases scanned, the resulting bins may not contain the proportion of cases that you wanted in those bins, particularly if the data file is sorted by the source variable. For example, if you limit the scan to the first 100 cases of a data file with 1000 cases and the data file is sorted in ascending order of age of respondent, instead of four percentile age bins each containing 25% of the cases, you may find that the first three bins each contain only about 3.3% of the cases, and the last bin contains 90% of the cases.

Copying Binned Categories

When creating binned categories for one or more variables, you can copy the binning specifications from another variable to the currently selected variable or from the selected variable to multiple other variables.

Figure 7-17
Copying bins to or from the current variable



To Copy Binning Specifications

- ▶ Define binned categories for at least one variable—but do *not* click OK or Paste.
 - ▶ Select (click) a variable in the Scanned Variable List for which you have defined binned categories.
 - ▶ Click To Other Variables.
 - ▶ Select the variables for which you want to create new variables with the same binned categories.
 - ▶ Click Copy.
- or*
- ▶ Select (click) a variable in the Scanned Variable List to which you want to copy defined binned categories.
 - ▶ Click From Another Variable.
 - ▶ Select the variable with the defined binned categories that you want to copy.
 - ▶ Click Copy.

If you have specified value labels for the variable from which you are copying the binning specifications, those are also copied.

Note: Once you click OK in the Visual Binning main dialog box to create new binned variables (or close the dialog box in any other way), you cannot use Visual Binning to copy those binned categories to other variables.

User-Missing Values in Visual Binning

Values defined as user-missing (values identified as codes for missing data) for the source variable are not included in the binned categories for the new variable. User-missing values for the source variables are copied as user-missing values for the new variable, and any defined value labels for missing value codes are also copied.

If a missing value code conflicts with one of the binned category values for the new variable, the missing value code for the new variable is recoded to a nonconflicting value by adding 100 to the highest binned category value. For example, if a value of 1 is defined as user-missing for the source variable and the new variable will have six binned categories, any cases with a value of 1 for the source variable will have a value of 106 for the new variable, and 106 will be defined as user-missing. If the user-missing value for the source variable had a defined value label, that label will be retained as the value label for the recoded value of the new variable.

Note: If the source variable has a defined range of user-missing values of the form $LO-n$, where n is a positive number, the corresponding user-missing values for the new variable will be negative numbers.

Data Transformations

In an ideal situation, your raw data are perfectly suitable for the type of analysis you want to perform, and any relationships between variables are either conveniently linear or neatly orthogonal. Unfortunately, this is rarely the case. Preliminary analysis may reveal inconvenient coding schemes or coding errors, or data transformations may be required in order to expose the true relationship between variables.

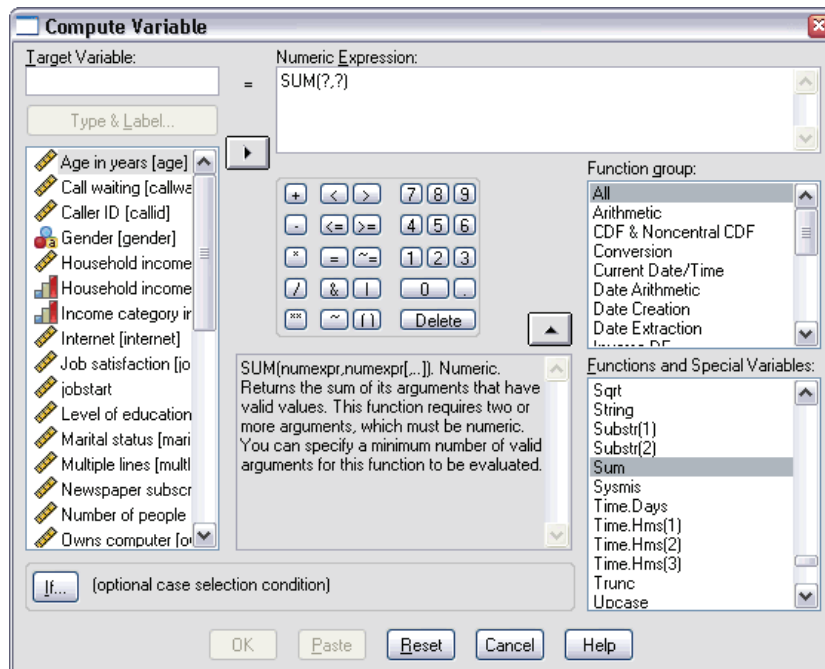
You can perform data transformations ranging from simple tasks, such as collapsing categories for analysis, to more advanced tasks, such as creating new variables based on complex equations and conditional statements.

Computing Variables

Use the Compute dialog box to compute values for a variable based on numeric transformations of other variables.

- You can compute values for numeric or string (alphanumeric) variables.
- You can create new variables or replace the values of existing variables. For new variables, you can also specify the variable type and label.
- You can compute values selectively for subsets of data based on logical conditions.
- You can use more than 70 built-in functions, including arithmetic functions, statistical functions, distribution functions, and string functions.

Figure 8-1
Compute Variable dialog box



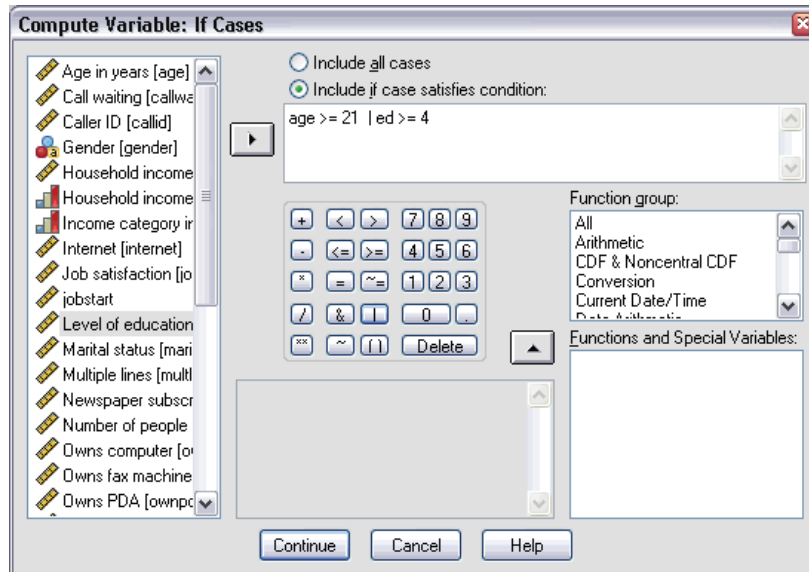
To Compute Variables

- ▶ From the menus choose:
 - Transform
 - Compute Variable...
- ▶ Type the name of a single target variable. It can be an existing variable or a new variable to be added to the active dataset.
- ▶ To build an expression, either paste components into the Expression field or type directly in the Expression field.
 - You can paste functions or commonly used system variables by selecting a group from the Function group list and double-clicking the function or variable in the Functions and Special Variables list (or select the function or variable and click the arrow adjacent to the Function group list). Fill in any parameters indicated by question marks (only applies to functions). The function group labeled All provides a listing of all available functions and system variables. A brief description of the currently selected function or variable is displayed in a reserved area in the dialog box.
 - String constants must be enclosed in quotation marks or apostrophes.
 - If values contain decimals, a period (.) must be used as the decimal indicator.
 - For new string variables, you must also select Type & Label to specify the data type.

Compute Variable: If Cases

The If Cases dialog box allows you to apply data transformations to selected subsets of cases, using conditional expressions. A conditional expression returns a value of *true*, *false*, or *missing* for each case.

Figure 8-2
Compute Variable If Cases dialog box



- If the result of a conditional expression is *true*, the case is included in the selected subset.
- If the result of a conditional expression is *false* or *missing*, the case is not included in the selected subset.
- Most conditional expressions use one or more of the six relational operators (<, >, <=, >=, =, and ~=) on the calculator pad.
- Conditional expressions can include variable names, constants, arithmetic operators, numeric (and other) functions, logical variables, and relational operators.

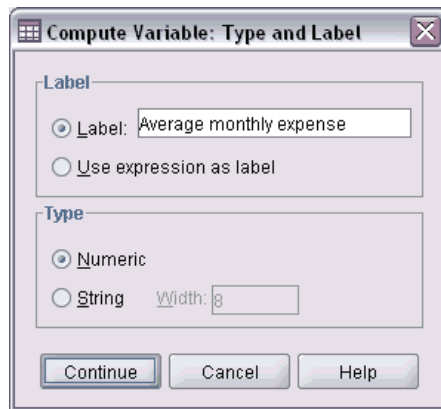
Compute Variable: Type and Label

By default, new computed variables are numeric. To compute a new string variable, you must specify the data type and width.

Label. Optional, descriptive variable label up to 255 bytes long. You can enter a label or use the first 110 characters of the compute expression as the label.

Type. Computed variables can be numeric or string (alphanumeric). String variables cannot be used in calculations.

Figure 8-3
Type and Label dialog box



Functions

Many types of functions are supported, including:

- Arithmetic functions
- Statistical functions
- String functions
- Date and time functions
- Distribution functions
- Random variable functions
- Missing value functions
- Scoring functions (SPSS Statistics Server only)

For more information and a detailed description of each function, type functions on the Index tab of the Help system.

■

Missing Values in Functions

Functions and simple arithmetic expressions treat missing values in different ways. In the expression:

```
(var1+var2+var3)/3
```

the result is missing if a case has a missing value for any of the three variables.

In the expression:

```
MEAN(var1, var2, var3)
```

the result is missing only if the case has missing values for all three variables.

For statistical functions, you can specify the minimum number of arguments that must have nonmissing values. To do so, type a period and the minimum number after the function name, as in:

```
MEAN.2(var1, var2, var3)
```

Random Number Generators

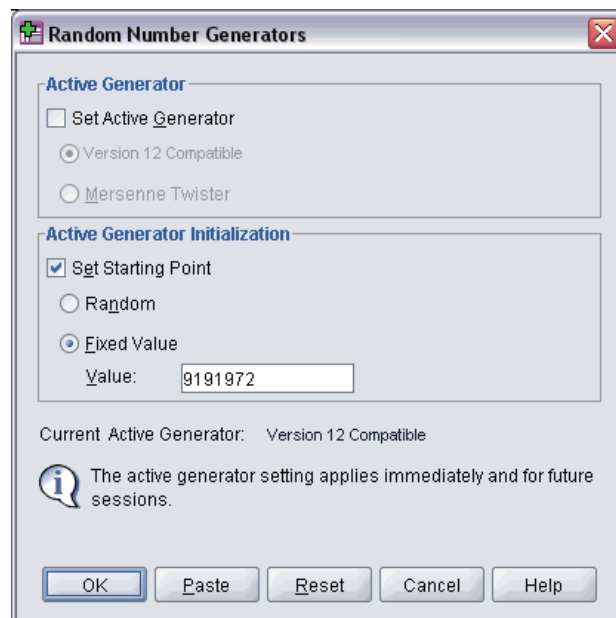
The Random Number Generators dialog box allows you to select the random number generator and set the starting sequence value so you can reproduce a sequence of random numbers.

Active Generator. Two different random number generators are available:

- **Version 12 Compatible.** The random number generator used in version 12 and previous releases. If you need to reproduce randomized results generated in previous releases based on a specified seed value, use this random number generator.
- **Mersenne Twister.** A newer random number generator that is more reliable for simulation purposes. If reproducing randomized results from version 12 or earlier is not an issue, use this random number generator.

Active Generator Initialization. The random number seed changes each time a random number is generated for use in transformations (such as random distribution functions), random sampling, or case weighting. To replicate a sequence of random numbers, set the initialization starting point value prior to each analysis that uses the random numbers. The value must be a positive integer.

Figure 8-4
Random Number Generators dialog box



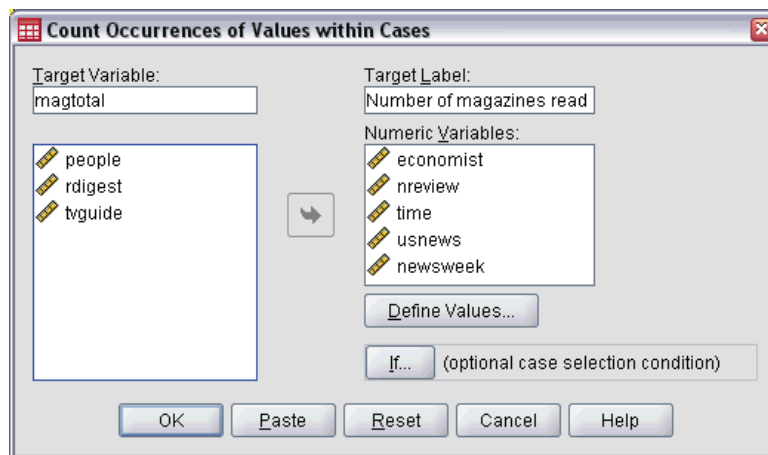
To select the random number generator and/or set the initialization value:

- ▶ From the menus choose:
 - Transform
 - Random Number Generators

Count Occurrences of Values within Cases

This dialog box creates a variable that counts the occurrences of the same value(s) in a list of variables for each case. For example, a survey might contain a list of magazines with *yes/no* check boxes to indicate which magazines each respondent reads. You could count the number of *yes* responses for each respondent to create a new variable that contains the total number of magazines read.

Figure 8-5
Count Occurrences of Values within Cases dialog box



To Count Occurrences of Values within Cases

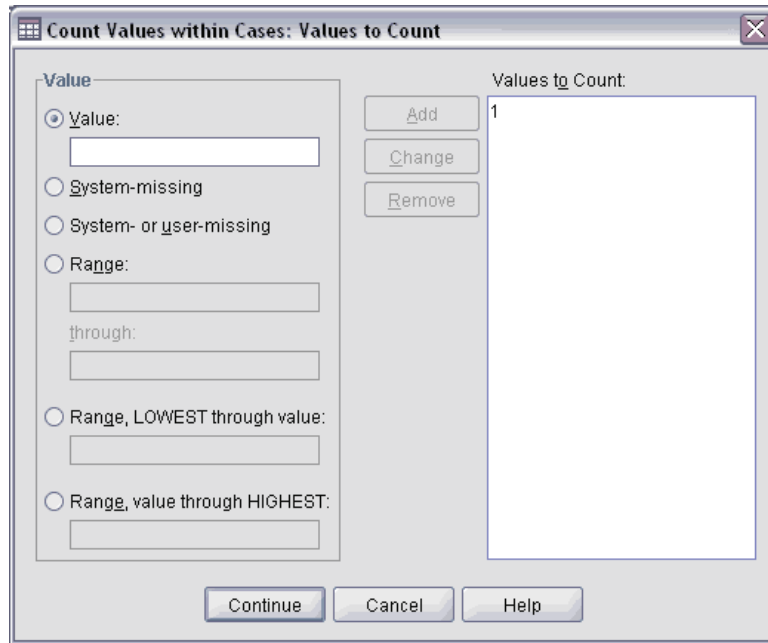
- ▶ From the menus choose:
 - Transform
 - Count Values within Cases...
- ▶ Enter a target variable name.
- ▶ Select two or more variables of the same type (numeric or string).
- ▶ Click Define Values and specify which value or values should be counted.
 - Optionally, you can define a subset of cases for which to count occurrences of values.

Count Values within Cases: Values to Count

The value of the target variable (on the main dialog box) is incremented by 1 each time one of the selected variables matches a specification in the Values to Count list here. If a case matches several specifications for any variable, the target variable is incremented several times for that variable.

Value specifications can include individual values, missing or system-missing values, and ranges. Ranges include their endpoints and any user-missing values that fall within the range.

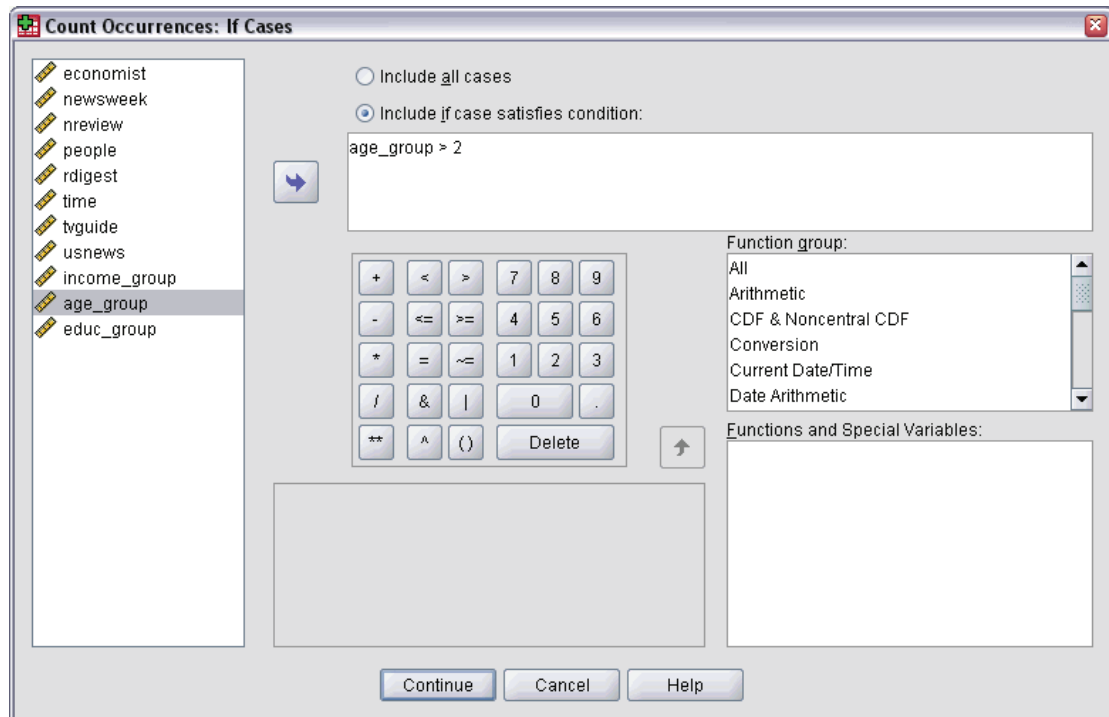
Figure 8-6
Values to Count dialog box



Count Occurrences: If Cases

The If Cases dialog box allows you to count occurrences of values for a selected subset of cases, using conditional expressions. A conditional expression returns a value of *true*, *false*, or *missing* for each case.

Figure 8-7
Count Occurrences If Cases dialog box



For general considerations on using the If Cases dialog box, see Compute Variable: If Cases on p. 126.

Shift Values

Shift Values creates new variables that contain the values of existing variables from preceding or subsequent cases.

Name. Name for the new variable. This must be a name that does not already exist in the active dataset.

Get value from earlier case (lag). Get the value from a previous case in the active dataset. For example, with the default number of cases value of 1, each case for the new variable has the value of the original variable from the case that immediately precedes it.

Get value from following case (lead). Get the value from a subsequent case in the active dataset. For example, with the default number of cases value of 1, each case for the new variable has the value of the original variable from the next case.

Number of cases to shift. Get the value from the n th preceding or subsequent case, where n is the value specified. The value must be a non-negative integer.

- If split file processing is on, the scope of the shift is limited to each split group. A shift value cannot be obtained from a case in a preceding or subsequent split group.

- Filter status is ignored.
- The value of the result variable is set to system-missing for the first or last n cases in the dataset or split group, where n is the value specified for **Number of cases to shift**. For example, using the Lag method with a value of 1 would set the result variable to system-missing for the first case in the dataset (or first case in each split group).
- User-missing values are preserved.
- Dictionary information from the original variable, including defined value labels and user-missing value assignments, is applied to the new variable. (Note: Custom variable attributes are not included.)
- A variable label is automatically generated for the new variable that describes the shift operation that created the variable.

To Create a New Variable with Shifted Values

- ▶ From the menus, choose:
 - Transform
 - Shift Values
- ▶ Select the variable to use as the source of values for the new variable.
- ▶ Enter a name for the new variable.
- ▶ Select the shift method (lag or lead) and the number of cases to shift.
- ▶ Click Change.
- ▶ Repeat for each new variable you want to create.

Recoding Values

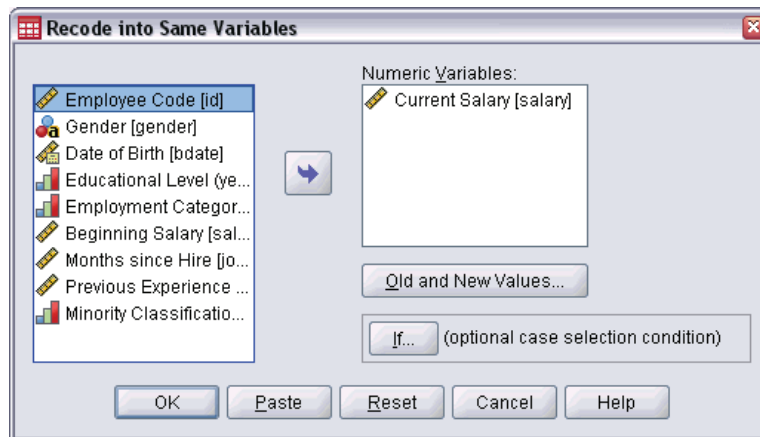
You can modify data values by recoding them. This is particularly useful for collapsing or combining categories. You can recode the values within existing variables, or you can create new variables based on the recoded values of existing variables.

Recode into Same Variables

The Recode into Same Variables dialog box allows you to reassign the values of existing variables or collapse ranges of existing values into new values. For example, you could collapse salaries into salary range categories.

You can recode numeric and string variables. If you select multiple variables, they must all be the same type. You cannot recode numeric and string variables together.

Figure 8-8
 Recode into Same Variables dialog box



To Recode Values of a Variable

- ▶ From the menus choose:
 Transform
 Recode into Same Variables...
- ▶ Select the variables you want to recode. If you select multiple variables, they must be the same type (numeric or string).
- ▶ Click Old and New Values and specify how to recode values.

Optionally, you can define a subset of cases to recode. The If Cases dialog box for doing this is the same as the one described for Count Occurrences.

Recode into Same Variables: Old and New Values

You can define values to recode in this dialog box. All value specifications must be the same data type (numeric or string) as the variables selected in the main dialog box.

Old Value. The value(s) to be recoded. You can recode single values, ranges of values, and missing values. System-missing values and ranges cannot be selected for string variables because neither concept applies to string variables. Ranges include their endpoints and any user-missing values that fall within the range.

- **Value.** Individual old value to be recoded into a new value. The value must be the same data type (numeric or string) as the variable(s) being recoded.
- **System-missing.** Values assigned by the program when values in your data are undefined according to the format type you have specified, when a numeric field is blank, or when a value resulting from a transformation command is undefined. Numeric system-missing values are displayed as periods. String variables cannot have system-missing values, since any character is legal in a string variable.

- **System- or user-missing.** Observations with values that either have been defined as user-missing values or are unknown and have been assigned the system-missing value, which is indicated with a period (.).
- **Range.** Inclusive range of values. Not available for string variables. Any user-missing values within the range are included.
- **All other values.** Any remaining values not included in one of the specifications on the Old-New list. This appears as ELSE on the Old-New list.

New Value. The single value into which each old value or range of values is recoded. You can enter a value or assign the system-missing value.

- **Value.** Value into which one or more old values will be recoded. The value must be the same data type (numeric or string) as the old value.
- **System-missing.** Recodes specified old values into the system-missing value. The system-missing value is not used in calculations, and cases with the system-missing value are excluded from many procedures. Not available for string variables.

Old→New. The list of specifications that will be used to recode the variable(s). You can add, change, and remove specifications from the list. The list is automatically sorted, based on the old value specification, using the following order: single values, missing values, ranges, and all other values. If you change a recode specification on the list, the procedure automatically re-sorts the list, if necessary, to maintain this order.

Figure 8-9
Old and New Values dialog box

Recode into Same Variables: Old and New Values

Old Value

Value:

System-missing

System- or user-missing

Range:

Range, LOWEST through value:

Range, value through HIGHEST:

All other values

New Value

Value:

System-missing

Old --> New:

Add Change Remove

Lowest thru 24999 --> 1
25000 thru 49999 --> 2
50000 thru 74999 --> 3
75000 thru Highest --> 4

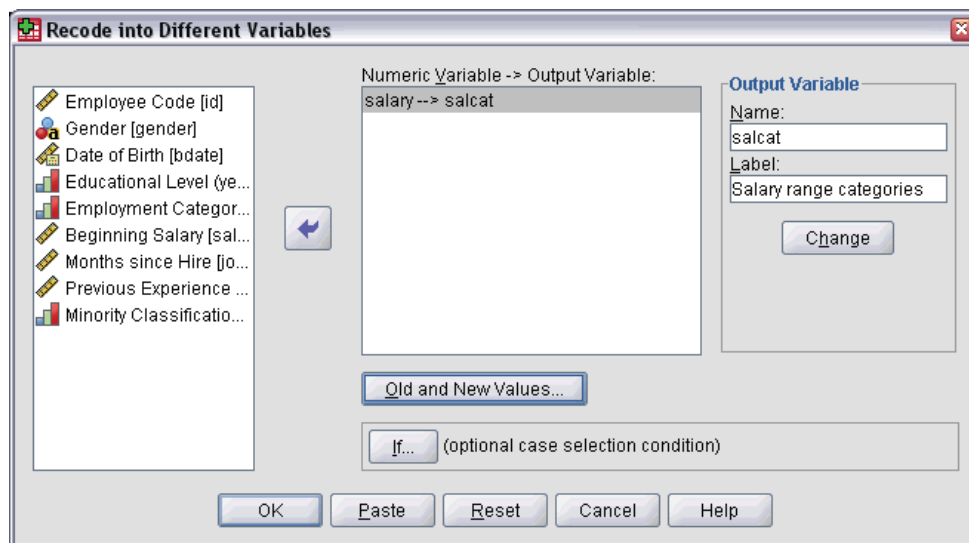
Continue Cancel Help

Recode into Different Variables

The Recode into Different Variables dialog box allows you to reassign the values of existing variables or collapse ranges of existing values into new values for a new variable. For example, you could collapse salaries into a new variable containing salary-range categories.

- You can recode numeric and string variables.
- You can recode numeric variables into string variables and vice versa.
- If you select multiple variables, they must all be the same type. You cannot recode numeric and string variables together.

Figure 8-10
Recode into Different Variables dialog box



To Recode Values of a Variable into a New Variable

- ▶ From the menus choose:
Transform
Recode into Different Variables...
- ▶ Select the variables you want to recode. If you select multiple variables, they must be the same type (numeric or string).
- ▶ Enter an output (new) variable name for each new variable and click Change.
- ▶ Click Old and New Values and specify how to recode values.

Optionally, you can define a subset of cases to recode. The If Cases dialog box for doing this is the same as the one described for Count Occurrences.

Recode into Different Variables: Old and New Values

You can define values to recode in this dialog box.

Old Value. The value(s) to be recoded. You can recode single values, ranges of values, and missing values. System-missing values and ranges cannot be selected for string variables because neither concept applies to string variables. Old values must be the same data type (numeric or string) as the original variable. Ranges include their endpoints and any user-missing values that fall within the range.

- **Value.** Individual old value to be recoded into a new value. The value must be the same data type (numeric or string) as the variable(s) being recoded.
- **System-missing.** Values assigned by the program when values in your data are undefined according to the format type you have specified, when a numeric field is blank, or when a value resulting from a transformation command is undefined. Numeric system-missing values are displayed as periods. String variables cannot have system-missing values, since any character is legal in a string variable.
- **System- or user-missing.** Observations with values that either have been defined as user-missing values or are unknown and have been assigned the system-missing value, which is indicated with a period (.).
- **Range.** Inclusive range of values. Not available for string variables. Any user-missing values within the range are included.
- **All other values.** Any remaining values not included in one of the specifications on the Old-New list. This appears as ELSE on the Old-New list.

New Value. The single value into which each old value or range of values is recoded. New values can be numeric or string.

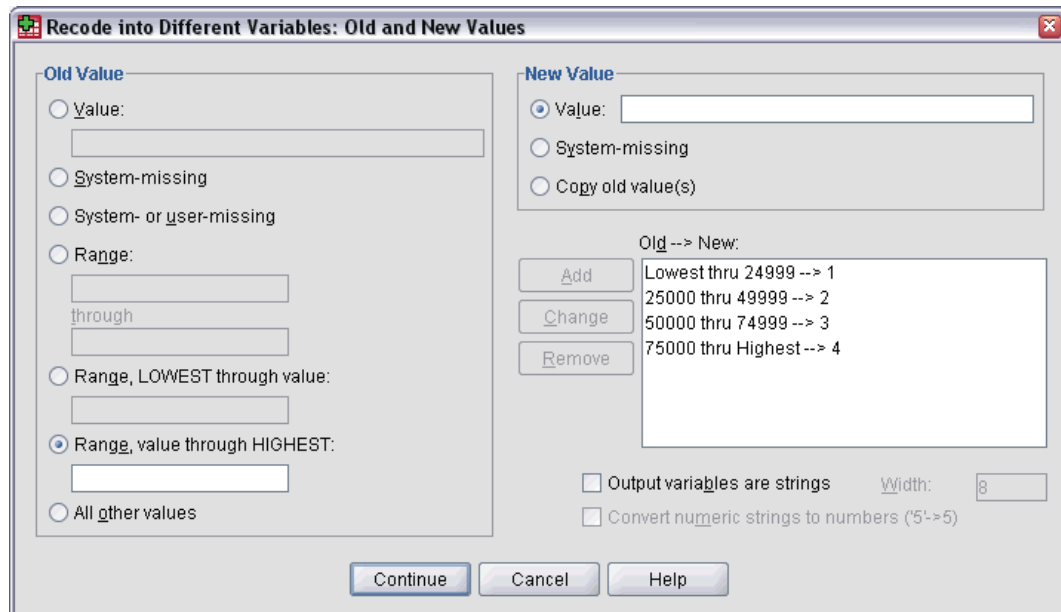
- **Value.** Value into which one or more old values will be recoded. The value must be the same data type (numeric or string) as the old value.
- **System-missing.** Recodes specified old values into the system-missing value. The system-missing value is not used in calculations, and cases with the system-missing value are excluded from many procedures. Not available for string variables.
- **Copy old values.** Retains the old value. If some values don't require recoding, use this to include the old values. Any old values that are not specified are not included in the new variable, and cases with those values will be assigned the system-missing value for the new variable.

Output variables are strings. Defines the new, recoded variable as a string (alphanumeric) variable. The old variable can be numeric or string.

Convert numeric strings to numbers. Converts string values containing numbers to numeric values. Strings containing anything other than numbers and an optional sign (+ or -) are assigned the system-missing value.

Old→New. The list of specifications that will be used to recode the variable(s). You can add, change, and remove specifications from the list. The list is automatically sorted, based on the old value specification, using the following order: single values, missing values, ranges, and all other values. If you change a recode specification on the list, the procedure automatically re-sorts the list, if necessary, to maintain this order.

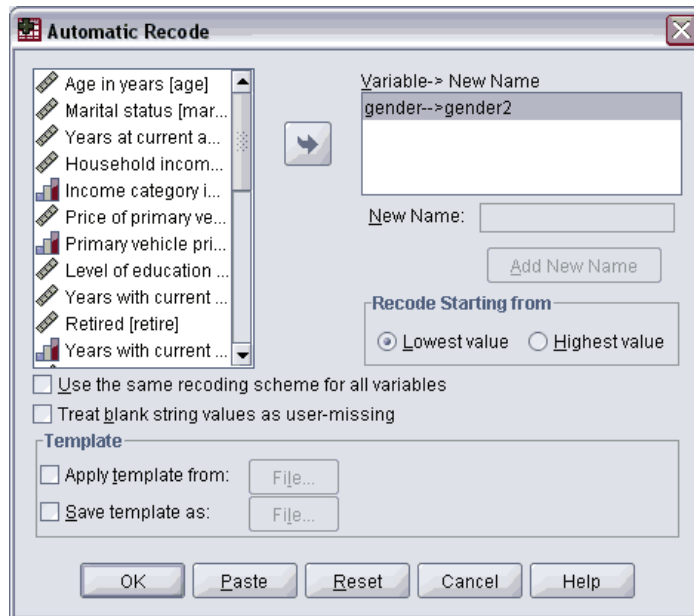
Figure 8-11
Old and New Values dialog box



Automatic Recode

The Automatic Recode dialog box allows you to convert string and numeric values into consecutive integers. When category codes are not sequential, the resulting empty cells reduce performance and increase memory requirements for many procedures. Additionally, some procedures cannot use string variables, and some require consecutive integer values for factor levels.

Figure 8-12
Automatic Recode dialog box



- The new variable(s) created by Automatic Recode retain any defined variable and value labels from the old variable. For any values without a defined value label, the original value is used as the label for the recoded value. A table displays the old and new values and value labels.
- String values are recoded in alphabetical order, with uppercase letters preceding their lowercase counterparts.
- Missing values are recoded into missing values higher than any nonmissing values, with their order preserved. For example, if the original variable has 10 nonmissing values, the lowest missing value would be recoded to 11, and the value 11 would be a missing value for the new variable.

Use the same recoding scheme for all variables. This option allows you to apply a single autorecoding scheme to all the selected variables, yielding a consistent coding scheme for all the new variables.

If you select this option, the following rules and limitations apply:

- All variables must be of the same type (numeric or string).
- All observed values for all selected variables are used to create a sorted order of values to recode into sequential integers.
- User-missing values for the new variables are based on the first variable in the list with defined user-missing values. All other values from other original variables, except for system-missing, are treated as valid.

Treat blank string values as user-missing. For string variables, blank or null values are not treated as **system-missing**. This option will autorecode blank strings into a **user-missing** value higher than the highest nonmissing value.

Templates

You can save the autorecoding scheme in a template file and then apply it to other variables and other data files.

For example, you may have a large number of alphanumeric product codes that you autorecode into integers every month, but some months new product codes are added that change the original autorecoding scheme. If you save the original scheme in a template and then apply it to the new data that contain the new set of codes, any new codes encountered in the data are autorecoded into values higher than the last value in the template, preserving the original autorecode scheme of the original product codes.

Save template as. Saves the autorecode scheme for the selected variables in an external template file.

- The template contains information that maps the original nonmissing values to the recoded values.
- Only information for nonmissing values is saved in the template. User-missing value information is not retained.
- If you have selected multiple variables for recoding but you have not selected to use the same autorecoding scheme for all variables or you are not applying an existing template as part of the autorecoding, the template will be based on the first variable in the list.
- If you have selected multiple variables for recoding and you have also selected Use the same recoding scheme for all variables and/or you have selected Apply template, then the template will contain the combined autorecoding scheme for all variables.

Apply template from. Applies a previously saved autorecode template to variables selected for recoding, appending any additional values found in the variables to the end of the scheme and preserving the relationship between the original and autorecoded values stored in the saved scheme.

- All variables selected for recoding must be the same type (numeric or string), and that type must match the type defined in the template.
- Templates do not contain any information on user-missing values. User-missing values for the target variables are based on the first variable in the original variable list with defined user-missing values. All other values from other original variables, except for system-missing, are treated as valid.
- Value mappings from the template are applied first. All remaining values are recoded into values higher than the last value in the template, with user-missing values (based on the first variable in the list with defined user-missing values) recoded into values higher than the last valid value.
- If you have selected multiple variables for autorecoding, the template is applied first, followed by a common, combined autorecoding for all additional values found in the selected variables, resulting in a single, common autorecoding scheme for all selected variables.

To Recode String or Numeric Values into Consecutive Integers

- ▶ From the menu choose:
 - Transform
 - Automatic Recode...
- ▶ Select one or more variables to recode.
- ▶ For each selected variable, enter a name for the new variable and click New Name.

Rank Cases

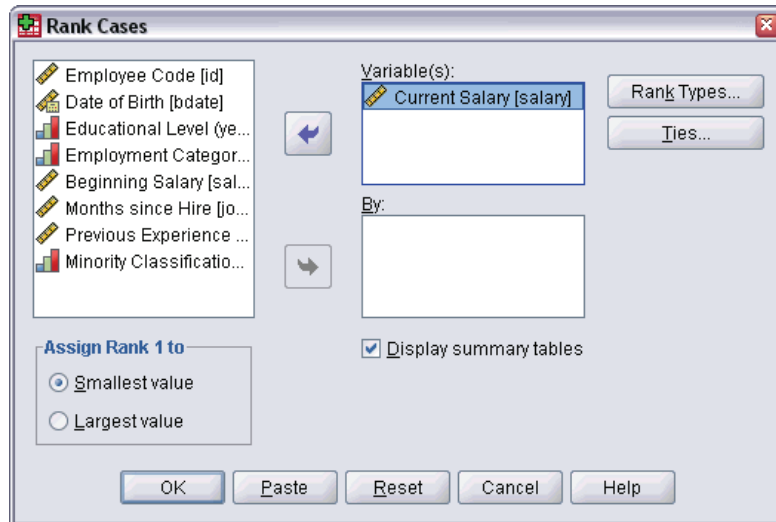
The Rank Cases dialog box allows you to create new variables containing ranks, normal and Savage scores, and percentile values for numeric variables.

New variable names and descriptive variable labels are automatically generated, based on the original variable name and the selected measure(s). A summary table lists the original variables, the new variables, and the variable labels. (Note: The automatically generated new variable names are limited to a maximum length of 8 bytes.)

Optionally, you can:

- Rank cases in ascending or descending order.
- Organize rankings into subgroups by selecting one or more grouping variables for the By list. Ranks are computed within each group. Groups are defined by the combination of values of the grouping variables. For example, if you select *gender* and *minority* as grouping variables, ranks are computed for each combination of *gender* and *minority*.

Figure 8-13
Rank Cases dialog box



To Rank Cases

- ▶ From the menu choose:
 - Transform
 - Rank Cases...

- ▶ Select one or more variables to rank. You can rank only numeric variables.
Optionally, you can rank cases in ascending or descending order and organize ranks into subgroups.

Rank Cases: Types

You can select multiple ranking methods. A separate ranking variable is created for each method. Ranking methods include simple ranks, Savage scores, fractional ranks, and percentiles. You can also create rankings based on proportion estimates and normal scores.

Rank. Simple rank. The value of the new variable equals its rank.

Savage score. The new variable contains Savage scores based on an exponential distribution.

Fractional rank. The value of the new variable equals rank divided by the sum of the weights of the nonmissing cases.

Fractional rank as percent. Each rank is divided by the number of cases with valid values and multiplied by 100.

Sum of case weights. The value of the new variable equals the sum of case weights. The new variable is a constant for all cases in the same group.

Ntiles. Ranks are based on percentile groups, with each group containing approximately the same number of cases. For example, 4 Ntiles would assign a rank of 1 to cases below the 25th percentile, 2 to cases between the 25th and 50th percentile, 3 to cases between the 50th and 75th percentile, and 4 to cases above the 75th percentile.

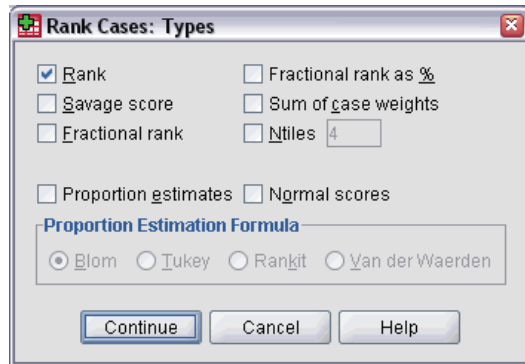
Proportion estimates. Estimates of the cumulative proportion of the distribution corresponding to a particular rank.

Normal scores. The z scores corresponding to the estimated cumulative proportion.

Proportion Estimation Formula. For proportion estimates and normal scores, you can select the proportion estimation formula: Blom, Tukey, Rankit, or Van der Waerden.

- **Blom.** Creates new ranking variable based on proportion estimates that uses the formula $(r-3/8) / (w+1/4)$, where w is the sum of the case weights and r is the rank.
- **Tukey.** Uses the formula $(r-1/3) / (w+1/3)$, where r is the rank and w is the sum of the case weights.
- **Rankit.** Uses the formula $(r-1/2) / w$, where w is the number of observations and r is the rank, ranging from 1 to w .
- **Van der Waerden.** Van der Waerden's transformation, defined by the formula $r/(w+1)$, where w is the sum of the case weights and r is the rank, ranging from 1 to w .


Figure 8-14
Rank Cases Types dialog box



Rank Cases: Ties

This dialog box controls the method for assigning rankings to cases with the same value on the original variable.

Figure 8-15
Rank Cases Ties dialog box



The following table shows how the different methods assign ranks to tied values:

Value	Mean	Low	High	Sequential
10	1	1	1	1
15	3	2	4	2
15	3	2	4	2
15	3	2	4	2
16	5	5	5	3
20	6	6	6	4

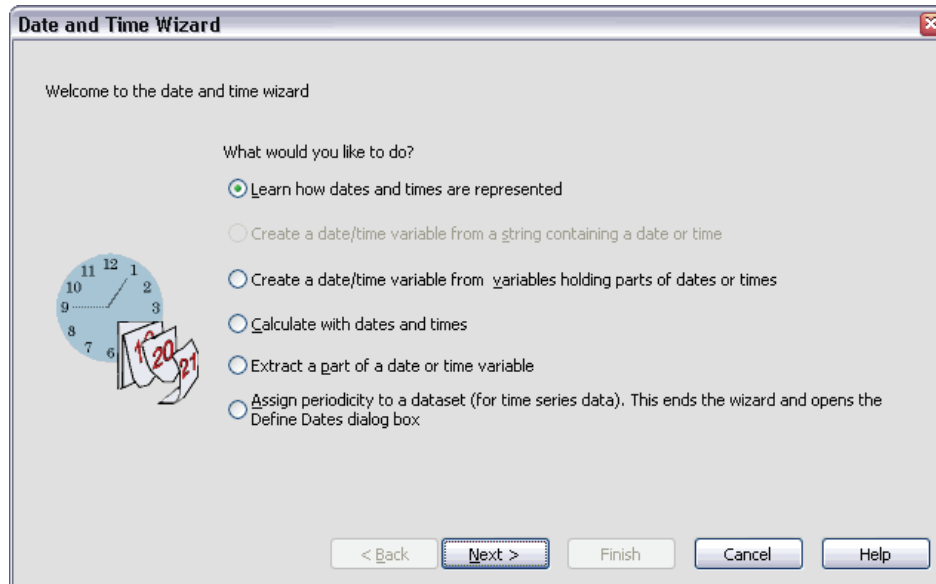
Date and Time Wizard

The Date and Time Wizard simplifies a number of common tasks associated with date and time variables.

To Use the Date and Time Wizard

- ▶ From the menus choose:
Transform
Date and Time Wizard...
- ▶ Select the task you wish to accomplish and follow the steps to define the task.

Figure 8-16
Date and Time Wizard introduction screen



- **Learn how dates and times are represented.** This choice leads to a screen that provides a brief overview of date/time variables in SPSS Statistics. By clicking on the Help button, it also provides a link to more detailed information.
- **Create a date/time variable from a string containing a date or time.** Use this option to create a date/time variable from a string variable. For example, you have a string variable representing dates in the form mm/dd/yyyy and want to create a date/time variable from this.
- **Create a date/time variable from variables holding parts of dates or times.** This choice allows you to construct a date/time variable from a set of existing variables. For example, you have a variable that represents the month (as an integer), a second that represents the day of the month, and a third that represents the year. You can combine these three variables into a single date/time variable.
- **Calculate with dates and times.** Use this option to add or subtract values from date/time variables. For example, you can calculate the duration of a process by subtracting a variable representing the start time of the process from another variable representing the end time of the process.

- **Extract a part of a date or time variable.** This option allows you to extract part of a date/time variable, such as the day of the month from a date/time variable, which has the form mm/dd/yyyy.
- **Assign periodicity to a dataset.** This choice takes you to the Define Dates dialog box, used to create date/time variables that consist of a set of sequential dates. This feature is typically used to associate dates with time series data.

Note: Tasks are disabled when the dataset lacks the types of variables required to accomplish the task. For instance, if the dataset contains no string variables, then the task to create a date/time variable from a string does not apply and is disabled.

Dates and Times in SPSS Statistics

Variables that represent dates and times in SPSS Statistics have a variable type of numeric, with display formats that correspond to the specific date/time formats. These variables are generally referred to as date/time variables. Date/time variables that actually represent dates are distinguished from those that represent a time duration that is independent of any date, such as 20 hours, 10 minutes, and 15 seconds. The latter are referred to as duration variables and the former as date or date/time variables. For a complete list of display formats, see “Date and Time” in the “Universals” section of the *Command Syntax Reference*.

Date and date/time variables. Date variables have a format representing a date, such as mm/dd/yyyy. Date/time variables have a format representing a date and time, such as dd-mmm-yyyy hh:mm:ss. Internally, date and date/time variables are stored as the number of seconds from October 14, 1582. Date and date/time variables are sometimes referred to as date-format variables.

- Both two- and four-digit year specifications are recognized. By default, two-digit years represent a range beginning 69 years prior to the current date and ending 30 years after the current date. This range is determined by your Options settings and is configurable (from the Edit menu, choose Options and click the Data tab).
- Dashes, periods, commas, slashes, or blanks can be used as delimiters in day-month-year formats.
- Months can be represented in digits, Roman numerals, or three-character abbreviations, and they can be fully spelled out. Three-letter abbreviations and fully spelled-out month names must be in English; month names in other languages are not recognized.

Duration variables. Duration variables have a format representing a time duration, such as hh:mm. They are stored internally as seconds without reference to a particular date.

- In time specifications (applies to date/time and duration variables), colons can be used as delimiters between hours, minutes, and seconds. Hours and minutes are required, but seconds are optional. A period is required to separate seconds from fractional seconds. Hours can be of unlimited magnitude, but the maximum value for minutes is 59 and for seconds is 59.999....

Current date and time. The system variable *\$TIME* holds the current date and time. It represents the number of seconds from October 14, 1582, to the date and time when the transformation command that uses it is executed.

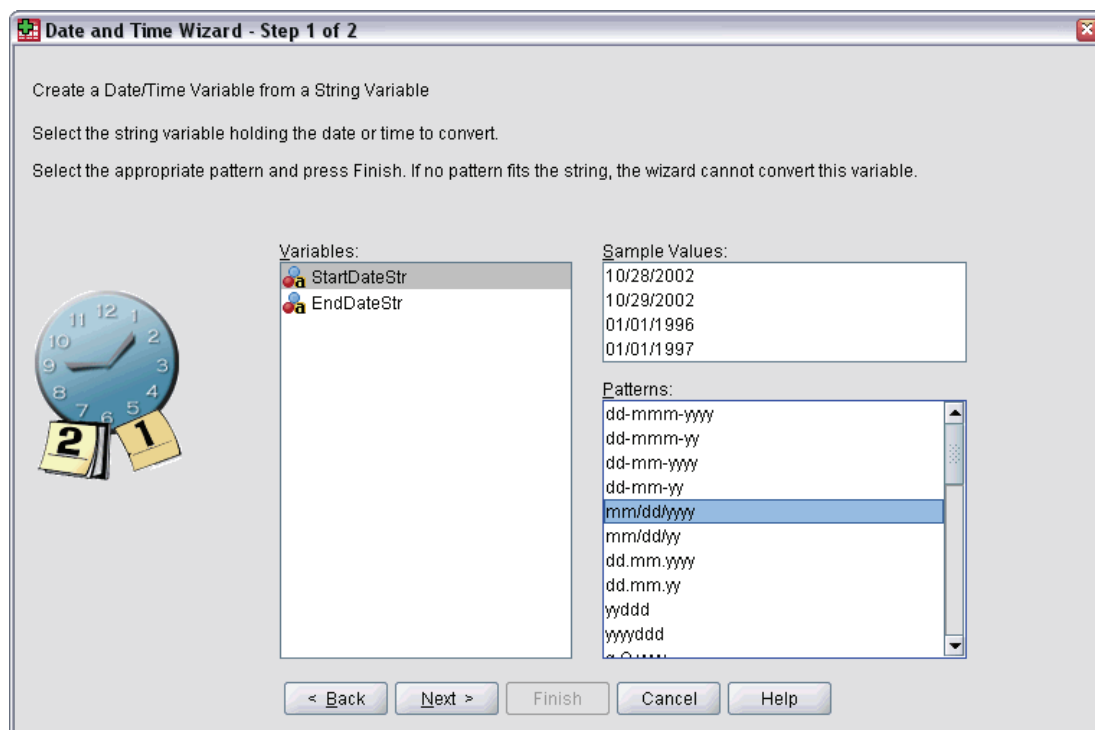
Create a Date/Time Variable from a String

To create a date/time variable from a string variable:

- ▶ Select Create a date/time variable from a string containing a date or time on the introduction screen of the Date and Time Wizard.

Select String Variable to Convert to Date/Time Variable

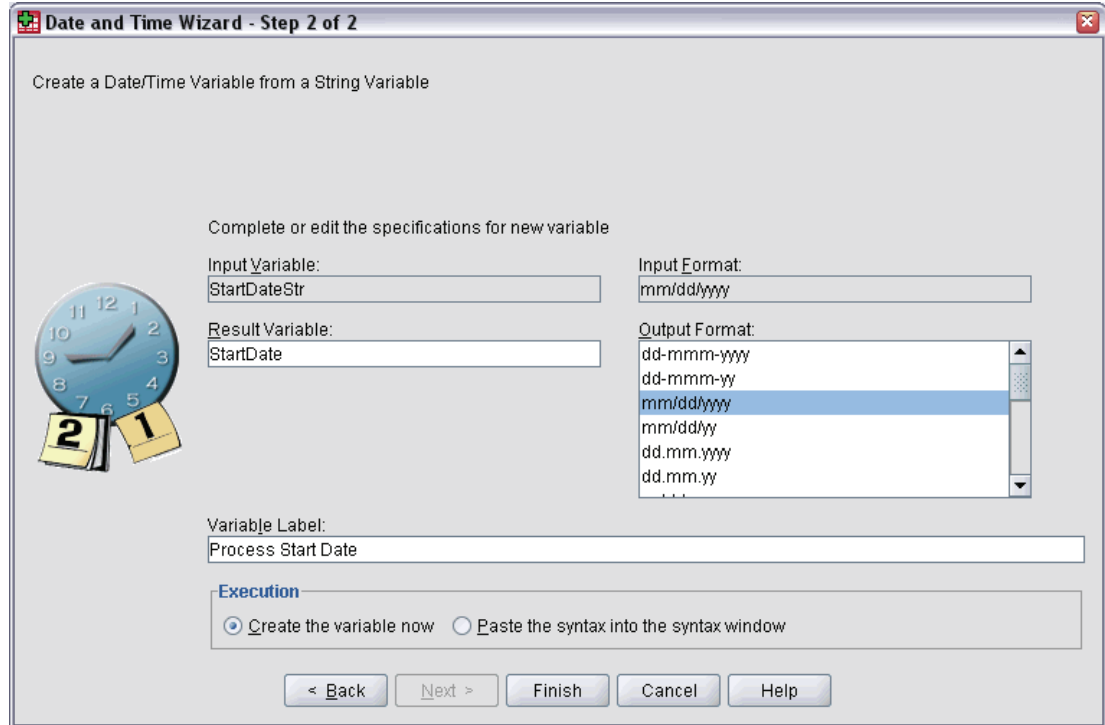
Figure 8-17
Create date/time variable from string variable, step 1



- ▶ Select the string variable to convert in the Variables list. Note that the list displays only string variables.
- ▶ Select the pattern from the Patterns list that matches how dates are represented by the string variable. The Sample Values list displays actual values of the selected variable in the data file. Values of the string variable that do not fit the selected pattern result in a value of system-missing for the new variable.

Specify Result of Converting String Variable to Date/Time Variable

Figure 8-18
Create date/time variable from string variable, step 2



- ▶ Enter a name for the Result Variable. This cannot be the name of an existing variable.

Optionally, you can:

- Select a date/time format for the new variable from the Output Format list.
- Assign a descriptive variable label to the new variable.

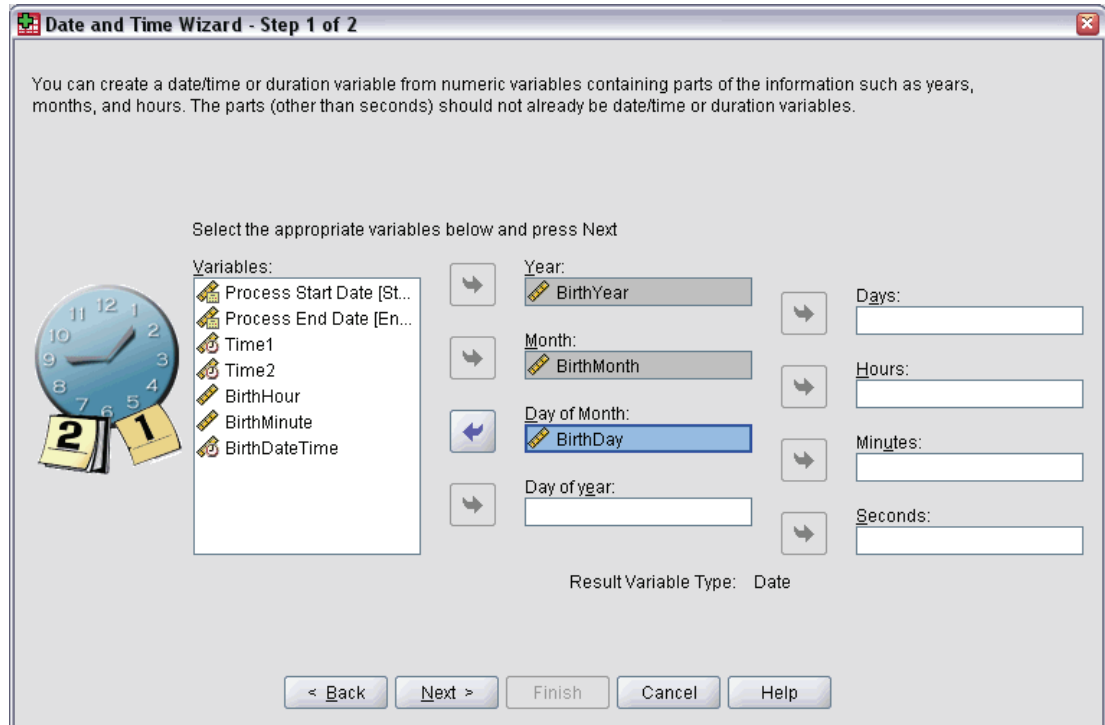
Create a Date/Time Variable from a Set of Variables

To merge a set of existing variables into a single date/time variable:

- ▶ Select Create a date/time variable from variables holding parts of dates or times on the introduction screen of the Date and Time Wizard.

Select Variables to Merge into Single Date/Time Variable

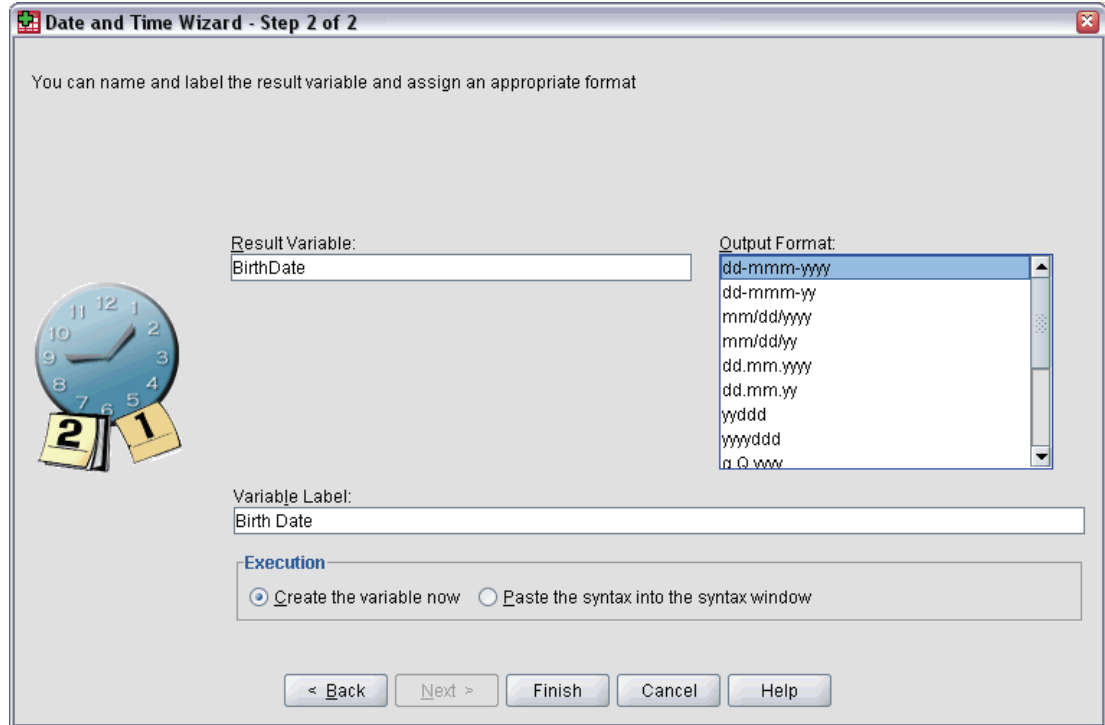
Figure 8-19
Create date/time variable from variable set, step 1



- ▶ Select the variables that represent the different parts of the date/time.
 - Some combinations of selections are not allowed. For instance, creating a date/time variable from Year and Day of Month is invalid because once Year is chosen, a full date is required.
 - You cannot use an existing date/time variable as one of the parts of the final date/time variable you're creating. Variables that make up the parts of the new date/time variable must be integers. The exception is the allowed use of an existing date/time variable as the Seconds part of the new variable. Since fractional seconds are allowed, the variable used for Seconds is not required to be an integer.
 - Values, for any part of the new variable, that are not within the allowed range result in a value of system-missing for the new variable. For instance, if you inadvertently use a variable representing day of month for Month, any cases with day of month values in the range 14–31 will be assigned the system-missing value for the new variable since the valid range for months in SPSS Statistics is 1–13.

Specify Date/Time Variable Created by Merging Variables

Figure 8-20
Create date/time variable from variable set, step 2



- ▶ Enter a name for the Result Variable. This cannot be the name of an existing variable.
- ▶ Select a date/time format from the Output Format list.

Optionally, you can:

- Assign a descriptive variable label to the new variable.

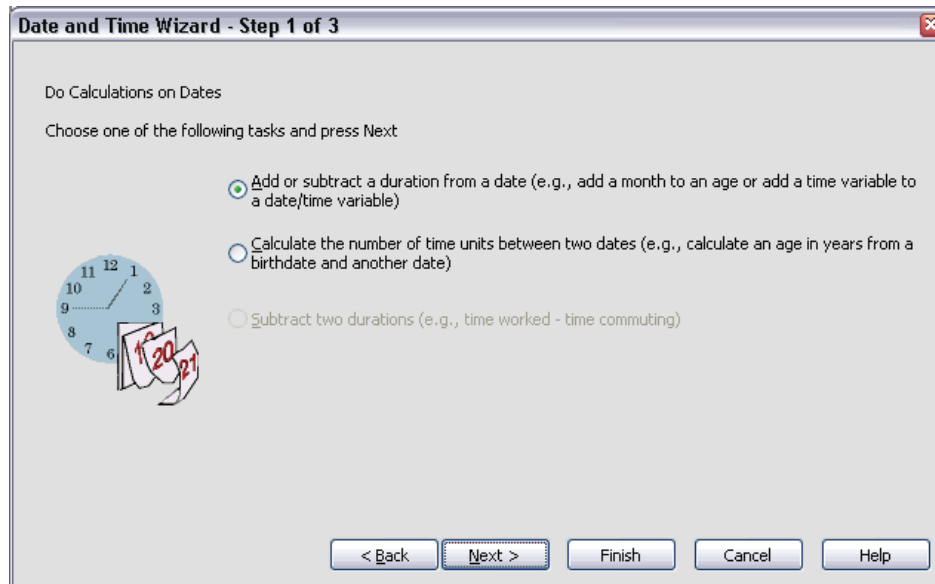
Add or Subtract Values from Date/Time Variables

To add or subtract values from date/time variables:

- ▶ Select Calculate with dates and times on the introduction screen of the Date and Time Wizard.

Select Type of Calculation to Perform with Date/Time Variables

Figure 8-21
Add or subtract values from date/time variables, step 1



- **Add or subtract a duration from a date.** Use this option to add to or subtract from a date-format variable. You can add or subtract durations that are fixed values, such as 10 days, or the values from a numeric variable, such as a variable that represents years.
- **Calculate the number of time units between two dates.** Use this option to obtain the difference between two dates as measured in a chosen unit. For example, you can obtain the number of years or the number of days separating two dates.
- **Subtract two durations.** Use this option to obtain the difference between two variables that have formats of durations, such as hh:mm or hh:mm:ss.

Note: Tasks are disabled when the dataset lacks the types of variables required to accomplish the task. For instance, if the dataset lacks two variables with formats of durations, then the task to subtract two durations does not apply and is disabled.

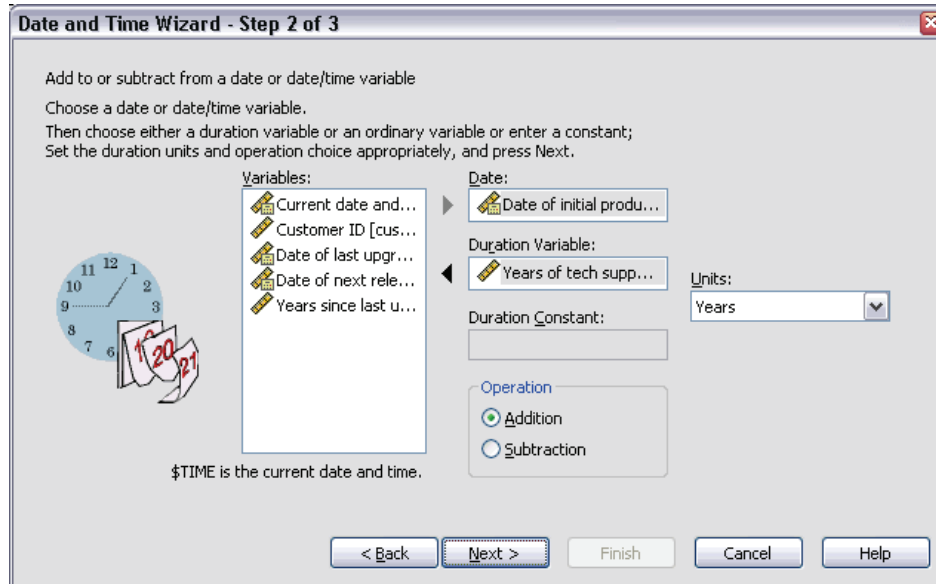
Add or Subtract a Duration from a Date

To add or subtract a duration from a date-format variable:

- ▶ Select Add or subtract a duration from a date on the screen of the Date and Time Wizard labeled *Do Calculations on Dates*.

Select Date/Time Variable and Duration to Add or Subtract

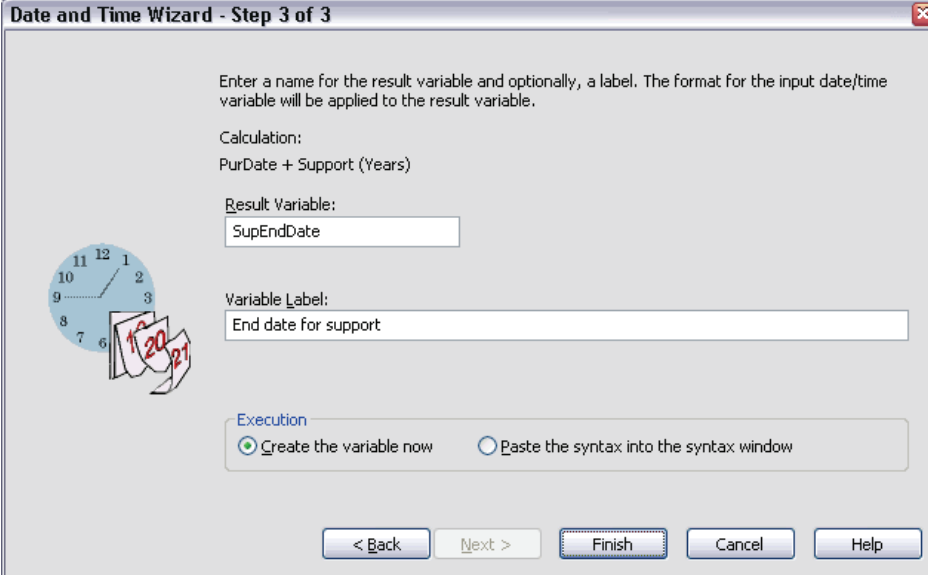
Figure 8-22
Add or subtract duration, step 2



- ▶ Select a date (or time) variable.
- ▶ Select a duration variable or enter a value for Duration Constant. Variables used for durations cannot be date or date/time variables. They can be duration variables or simple numeric variables.
- ▶ Select the unit that the duration represents from the drop-down list. Select Duration if using a variable and the variable is in the form of a duration, such as hh:mm or hh:mm:ss.

Specify Result of Adding or Subtracting a Duration from a Date/Time Variable

Figure 8-23
Add or subtract duration, step 3



Date and Time Wizard - Step 3 of 3

Enter a name for the result variable and optionally, a label. The format for the input date/time variable will be applied to the result variable.

Calculation:
PurDate + Support (Years)

Result Variable:
SupEndDate

Variable Label:
End date for support

Execution
 Create the variable now Paste the syntax into the syntax window

< Back Next > Finish Cancel Help

- ▶ Enter a name for Result Variable. This cannot be the name of an existing variable.
Optionally, you can:
 - Assign a descriptive variable label to the new variable.

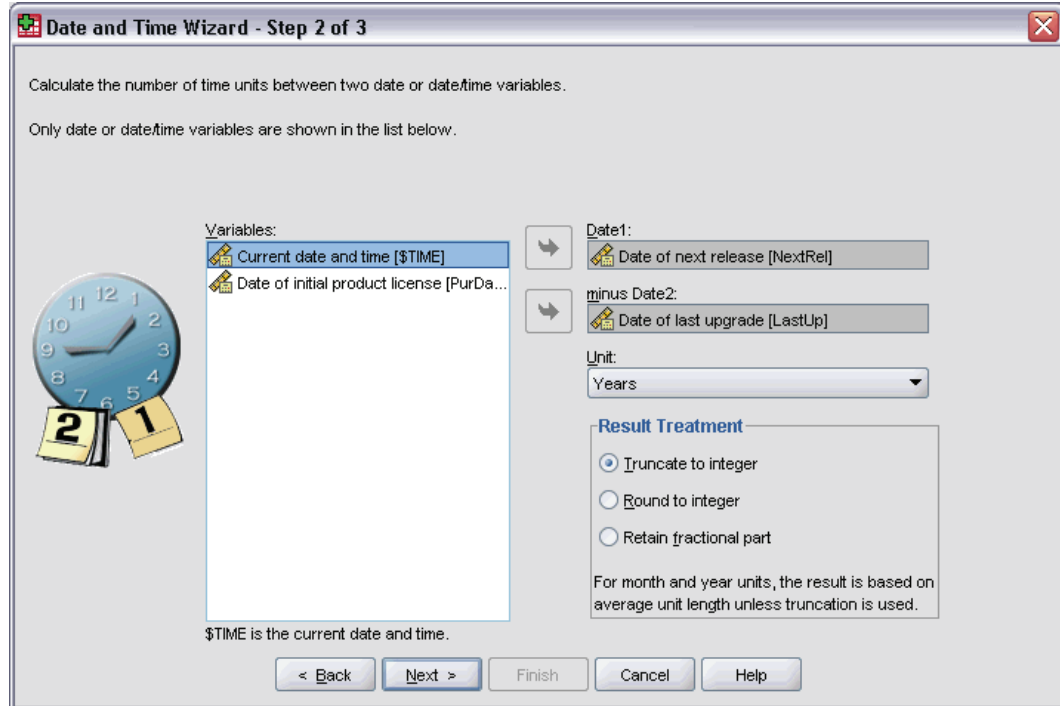
Subtract Date-Format Variables

To subtract two date-format variables:

- ▶ Select Calculate the number of time units between two dates on the screen of the Date and Time Wizard labeled *Do Calculations on Dates*.

Select Date-Format Variables to Subtract

Figure 8-24
Subtract dates, step 2



- ▶ Select the variables to subtract.
- ▶ Select the unit for the result from the drop-down list.
- ▶ Select how the result should be calculated (Result Treatment).

Result Treatment

The following options are available for how the result is calculated:

- **Truncate to integer.** Any fractional portion of the result is ignored. For example, subtracting 10/28/2006 from 10/21/2007 returns a result of 0 for years and 11 for months.
- **Round to integer.** The result is rounded to the closest integer. For example, subtracting 10/28/2006 from 10/21/2007 returns a result of 1 for years and 12 for months.
- **Retain fractional part.** The complete value is retained; no rounding or truncation is applied. For example, subtracting 10/28/2006 from 10/21/2007 returns a result of 0.98 for years and 11.76 for months.

For rounding and fractional retention, the result for years is based on average number of days in a year (365.25), and the result for months is based on the average number of days in a month (30.4375). For example, subtracting 2/1/2007 from 3/1/2007 (m/d/y format) returns a fractional result of 0.92 months, whereas subtracting 3/1/2007 from 2/1/2007 returns a fractional difference of 1.02 months. This also affects values calculated on time spans that include leap years. For

example, subtracting 2/1/2008 from 3/1/2008 returns a fractional difference of 0.95 months, compared to 0.92 for the same time span in a non-leap year.

Date 1	Date 2	Years			Months		
		Truncate	Round	Fraction	Truncate	Round	Fraction
10/21/2006	10/28/2007	1	1	1.02	12	12	12.22
10/28/2006	10/21/2007	0	1	.98	11	12	11.76
2/1/2007	3/1/2007	0	0	.08	1	1	.92
2/1/2008	3/1/2008	0	0	.08	1	1	.95
3/1/2007	4/1/2007	0	0	.08	1	1	1.02
4/1/2007	5/1/2007	0	0	.08	1	1	.99

Specify Result of Subtracting Two Date-Format Variables

Figure 8-25
Subtract dates, step 3

Calculation: NextRel - LastUp

Result Variable: YearsLastUp Units: Years

Variable Label: Years since last upgrade

Execution
 Create the variable now Paste the syntax into the syntax window

< Back Next > Finish Cancel Help

- ▶ Enter a name for Result Variable. This cannot be the name of an existing variable.

Optionally, you can:

- Assign a descriptive variable label to the new variable.

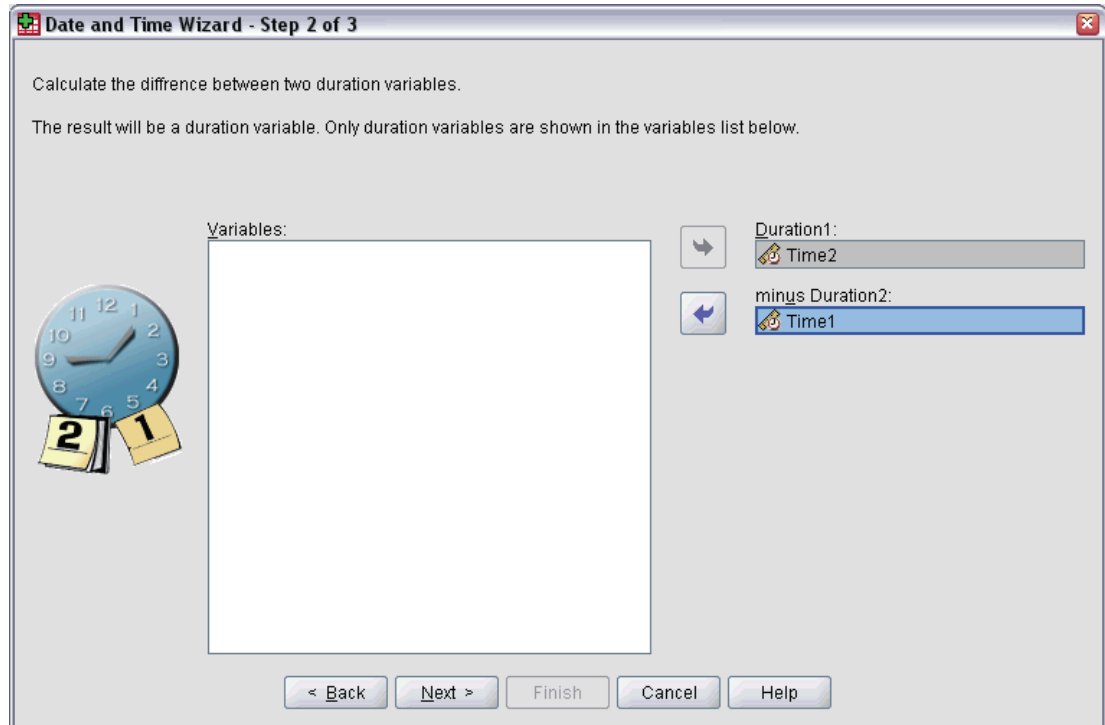
Subtract Duration Variables

To subtract two duration variables:

- ▶ Select Subtract two durations on the screen of the Date and Time Wizard labeled *Do Calculations on Dates*.

Select Duration Variables to Subtract

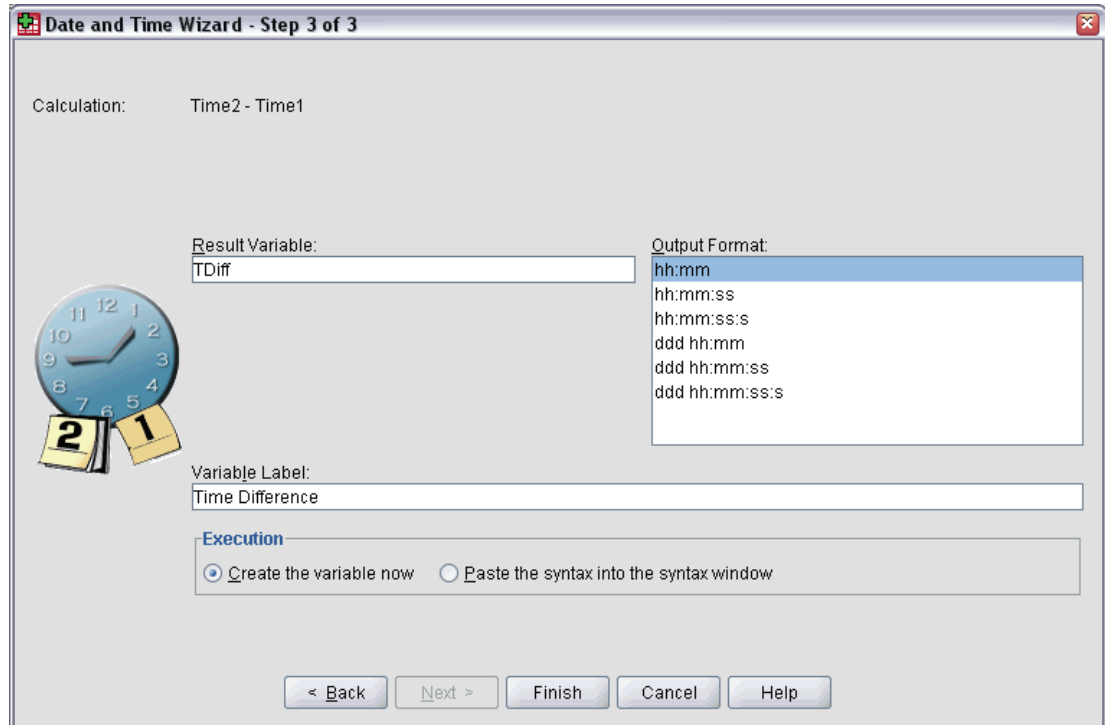
Figure 8-26
Subtract two durations, step 2



- ▶ Select the variables to subtract.

Specify Result of Subtracting Two Duration Variables

Figure 8-27
Subtract two durations, step 3



- ▶ Enter a name for Result Variable. This cannot be the name of an existing variable.
- ▶ Select a duration format from the Output Format list.

Optionally, you can:

- Assign a descriptive variable label to the new variable.

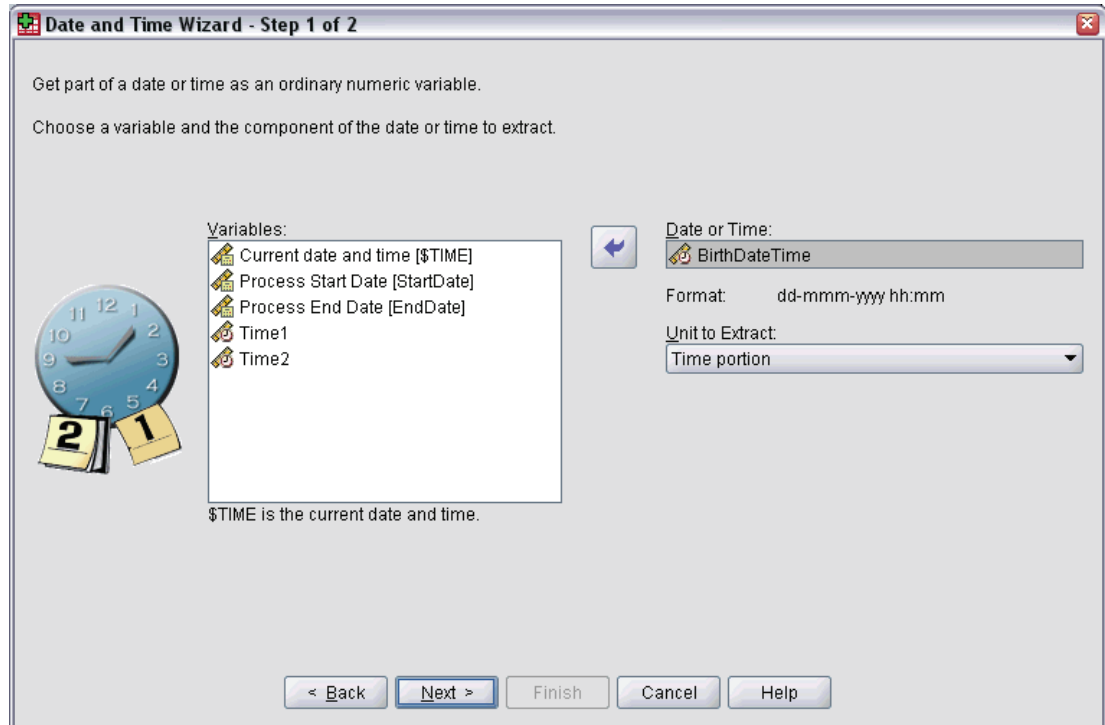
Extract Part of a Date/Time Variable

To extract a component—such as the year—from a date/time variable:

- ▶ Select Extract a part of a date or time variable on the introduction screen of the Date and Time Wizard.

Select Component to Extract from Date/Time Variable

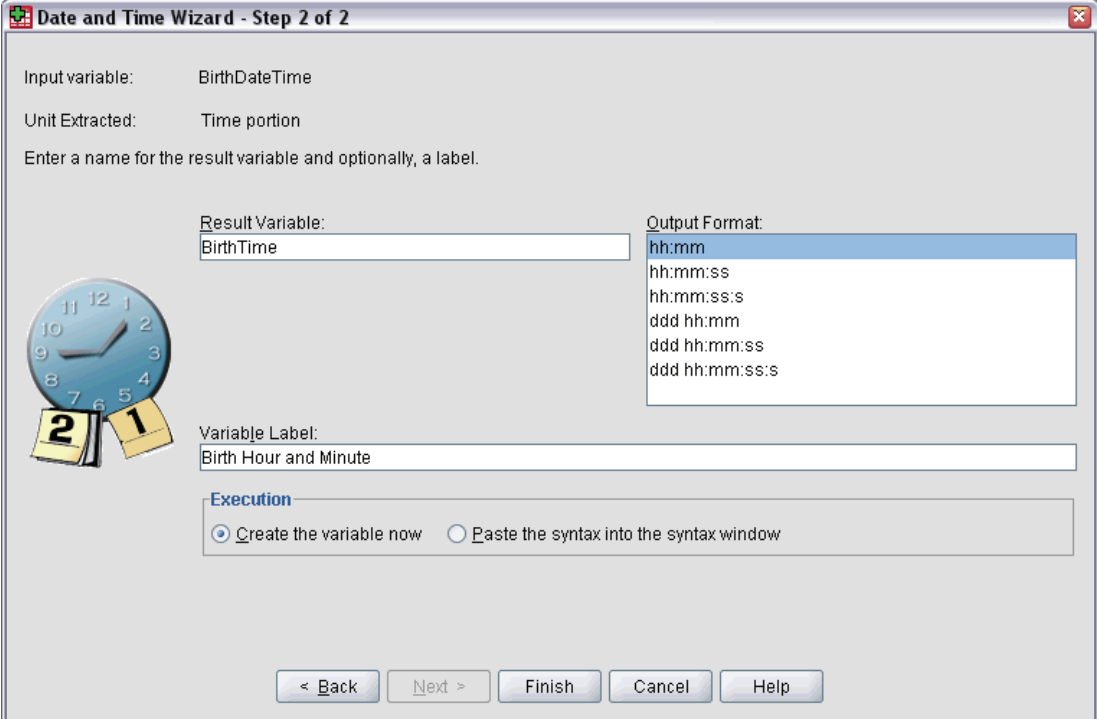
Figure 8-28
Get part of a date/time variable, step 1



- ▶ Select the variable containing the date or time part to extract.
- ▶ Select the part of the variable to extract, from the drop-down list. You can extract information from dates that is not explicitly part of the display date, such as day of the week.

Specify Result of Extracting Component from Date/Time Variable

Figure 8-29
Get part of a date/time variable, step 2



Date and Time Wizard - Step 2 of 2

Input variable: BirthDateTime
Unit Extracted: Time portion
Enter a name for the result variable and optionally, a label.

Result Variable: BirthTime
Output Format: hh:mm
hh:mm:ss
hh:mm:ss:s
ddd hh:mm
ddd hh:mm:ss
ddd hh:mm:ss:s

Variable Label: Birth Hour and Minute

Execution
 Create the variable now Paste the syntax into the syntax window

< Back Next > Finish Cancel Help

- ▶ Enter a name for Result Variable. This cannot be the name of an existing variable.
- ▶ If you're extracting the date or time part of a date/time variable, then you must select a format from the Output Format list. In cases where the output format is not required, the Output Format list will be disabled.

Optionally, you can:

- Assign a descriptive variable label to the new variable.

Time Series Data Transformations

Several data transformations that are useful in time series analysis are provided:

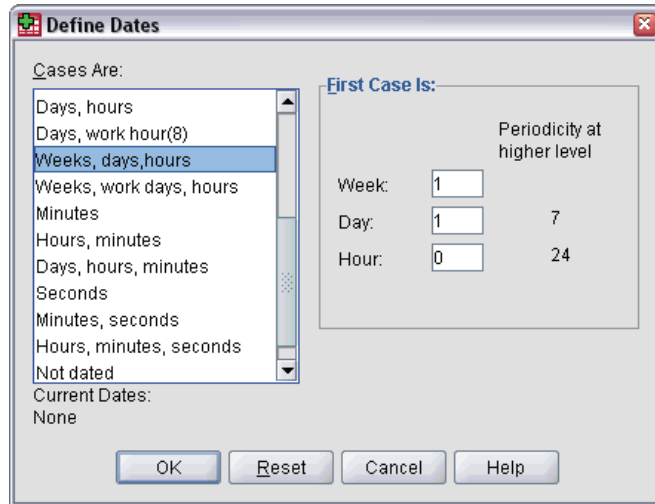
- Generate date variables to establish periodicity and to distinguish between historical, validation, and forecasting periods.
- Create new time series variables as functions of existing time series variables.
- Replace system- and user-missing values with estimates based on one of several methods.

A **time series** is obtained by measuring a variable (or set of variables) regularly over a period of time. Time series data transformations assume a data file structure in which each case (row) represents a set of observations at a different time, and the length of time between cases is uniform.

Define Dates

The Define Dates dialog box allows you to generate date variables that can be used to establish the periodicity of a **time series** and to label output from time series analysis.

Figure 8-30
Define Dates dialog box



Cases Are. Defines the time interval used to generate dates.

- Not dated removes any previously defined date variables. Any variables with the following names are deleted: *year_*, *quarter_*, *month_*, *week_*, *day_*, *hour_*, *minute_*, *second_*, and *date_*.
- Custom indicates the presence of custom date variables created with command syntax (for example, a four-day workweek). This item merely reflects the current state of the active dataset. Selecting it from the list has no effect.

First Case Is. Defines the starting date value, which is assigned to the first case. Sequential values, based on the time interval, are assigned to subsequent cases.

Periodicity at higher level. Indicates the repetitive cyclical variation, such as the number of months in a year or the number of days in a week. The value displayed indicates the maximum value you can enter.

A new numeric variable is created for each component that is used to define the date. The new variable names end with an underscore. A descriptive string variable, *date_*, is also created from the components. For example, if you selected *Weeks, days, hours*, four new variables are created: *week_*, *day_*, *hour_*, and *date_*.

If date variables have already been defined, they are replaced when you define new date variables that will have the same names as the existing date variables.

To Define Dates for Time Series Data

- From the menus choose:
 - Data
 - Define Dates...

- ▶ Select a time interval from the Cases Are list.
- ▶ Enter the value(s) that define the starting date for First Case Is, which determines the date assigned to the first case.

Date Variables versus Date Format Variables

Date variables created with Define Dates should not be confused with date format variables defined in the Variable View of the Data Editor. Date variables are used to establish periodicity for time series data. Date format variables represent dates and/or times displayed in various date/time formats. Date variables are simple integers representing the number of days, weeks, hours, and so on, from a user-specified starting point. Internally, most date format variables are stored as the number of seconds from October 14, 1582.

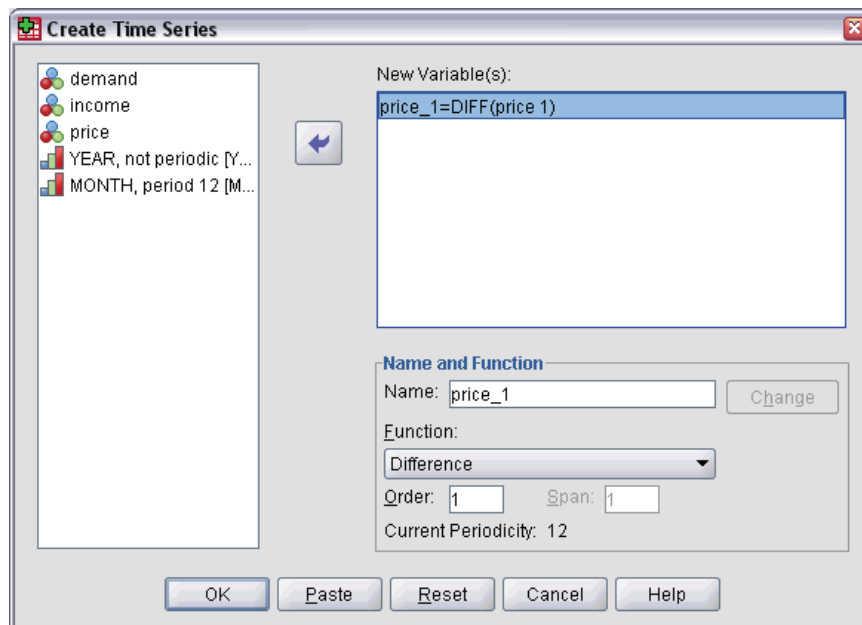
Create Time Series

The Create Time Series dialog box allows you to create new variables based on functions of existing numeric **time series** variables. These transformed values are useful in many time series analysis procedures.

Default new variable names are the first six characters of the existing variable used to create it, followed by an underscore and a sequential number. For example, for the variable *price*, the new variable name would be *price_1*. The new variables retain any defined value labels from the original variables.

Available functions for creating time series variables include differences, moving averages, running medians, lag, and lead functions.

Figure 8-31
Create Time Series dialog box



To Create New Time Series Variables

- ▶ From the menus choose:
 - Transform
 - Create Time Series...
- ▶ Select the time series function that you want to use to transform the original variable(s).
- ▶ Select the variable(s) from which you want to create new time series variables. Only numeric variables can be used.

Optionally, you can:

- Enter variable names to override the default new variable names.
- Change the function for a selected variable.

Time Series Transformation Functions

Difference. Nonseasonal difference between successive values in the series. The order is the number of previous values used to calculate the difference. Because one observation is lost for each order of difference, system-missing values appear at the beginning of the series. For example, if the difference order is 2, the first two cases will have the system-missing value for the new variable.

Seasonal difference. Difference between series values a constant span apart. The span is based on the currently defined periodicity. To compute seasonal differences, you must have defined date variables (Data menu, Define Dates) that include a periodic component (such as months of the year). The order is the number of seasonal periods used to compute the difference. The number of cases with the system-missing value at the beginning of the series is equal to the periodicity multiplied by the order. For example, if the current periodicity is 12 and the order is 2, the first 24 cases will have the system-missing value for the new variable.

Centered moving average. Average of a span of series values surrounding and including the current value. The span is the number of series values used to compute the average. If the span is even, the moving average is computed by averaging each pair of uncentered means. The number of cases with the system-missing value at the beginning and at the end of the series for a span of n is equal to $n/2$ for even span values and $(n-1)/2$ for odd span values. For example, if the span is 5, the number of cases with the system-missing value at the beginning and at the end of the series is 2.

Prior moving average. Average of the span of series values preceding the current value. The span is the number of preceding series values used to compute the average. The number of cases with the system-missing value at the beginning of the series is equal to the span value.

Running medians. Median of a span of series values surrounding and including the current value. The span is the number of series values used to compute the median. If the span is even, the median is computed by averaging each pair of uncentered medians. The number of cases with the system-missing value at the beginning and at the end of the series for a span of n is equal to $n/2$ for even span values and $(n-1)/2$ for odd span values. For example, if the span is 5, the number of cases with the system-missing value at the beginning and at the end of the series is 2.

Cumulative sum. Cumulative sum of series values up to and including the current value.

Lag. Value of a previous case, based on the specified lag order. The order is the number of cases prior to the current case from which the value is obtained. The number of cases with the system-missing value at the beginning of the series is equal to the order value.

Lead. Value of a subsequent case, based on the specified lead order. The order is the number of cases after the current case from which the value is obtained. The number of cases with the system-missing value at the end of the series is equal to the order value.

Smoothing. New series values based on a compound data smoother. The smoother starts with a running median of 4, which is centered by a running median of 2. It then resmooths these values by applying a running median of 5, a running median of 3, and hanning (running weighted averages). Residuals are computed by subtracting the smoothed series from the original series. This whole process is then repeated on the computed residuals. Finally, the smoothed residuals are computed by subtracting the smoothed values obtained the first time through the process. This is sometimes referred to as **T4253H smoothing**.

Replace Missing Values

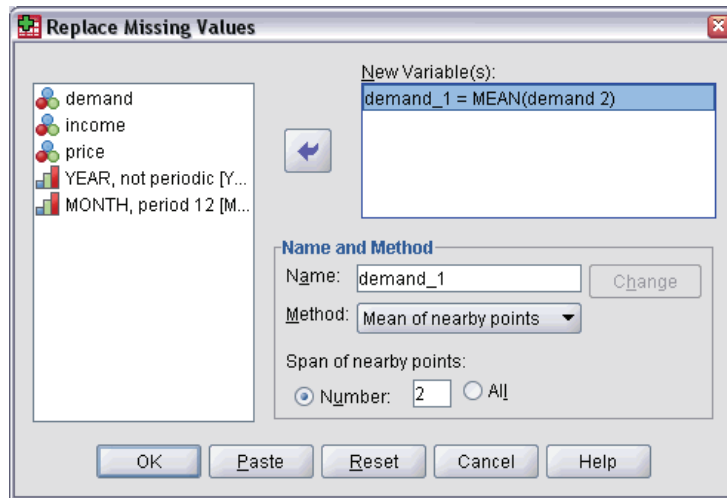
Missing observations can be problematic in analysis, and some time series measures cannot be computed if there are missing values in the series. Sometimes the value for a particular observation is simply not known. In addition, missing data can result from any of the following:

- Each degree of differencing reduces the length of a series by 1.
- Each degree of seasonal differencing reduces the length of a series by one season.
- If you create new series that contain forecasts beyond the end of the existing series (by clicking a Save button and making suitable choices), the original series and the generated residual series will have missing data for the new observations.
- Some transformations (for example, the log transformation) produce missing data for certain values of the original series.

Missing data at the beginning or end of a series pose no particular problem; they simply shorten the useful length of the series. Gaps in the middle of a series (*embedded* missing data) can be a much more serious problem. The extent of the problem depends on the analytical procedure you are using.

The Replace Missing Values dialog box allows you to create new **time series** variables from existing ones, replacing missing values with estimates computed with one of several methods. Default new variable names are the first six characters of the existing variable used to create it, followed by an underscore and a sequential number. For example, for the variable *price*, the new variable name would be *price_1*. The new variables retain any defined value labels from the original variables.

Figure 8-32
 Replace Missing Values dialog box



To Replace Missing Values for Time Series Variables

- ▶ From the menus choose:
 Transform
 Replace Missing Values...
 - ▶ Select the estimation method you want to use to replace missing values.
 - ▶ Select the variable(s) for which you want to replace missing values.
- Optionally, you can:
- Enter variable names to override the default new variable names.
 - Change the estimation method for a selected variable.

Estimation Methods for Replacing Missing Values

Series mean. Replaces missing values with the mean for the entire series.

Mean of nearby points. Replaces missing values with the mean of valid surrounding values. The span of nearby points is the number of valid values above and below the missing value used to compute the mean.

Median of nearby points. Replaces missing values with the median of valid surrounding values. The span of nearby points is the number of valid values above and below the missing value used to compute the median.

Linear interpolation. Replaces missing values using a linear interpolation. The last valid value before the missing value and the first valid value after the missing value are used for the interpolation. If the first or last case in the series has a missing value, the missing value is not replaced.

Linear trend at point. Replaces missing values with the linear trend for that point. The existing series is regressed on an index variable scaled 1 to n . Missing values are replaced with their predicted values.

Scoring Data with Predictive Models

The process of applying a predictive model to a set of data is referred to as **scoring** the data. SPSS Statistics, Clementine, and AnswerTree have procedures for building predictive models such as regression, clustering, tree, and neural network models. Once a model has been built, the model specifications can be saved as an XML file containing all of the information necessary to reconstruct the model. The SPSS Statistics Server product then provides the means to read an XML model file and apply the model to a dataset.

Example. A credit application is rated for risk based on various aspects of the applicant and the loan in question. The credit score obtained from the risk model is used to accept or reject the loan application.

Scoring is treated as a transformation of the data. The model is expressed internally as a set of numeric transformations to be applied to a given set of variables—the predictor variables specified in the model—in order to obtain a predicted result. In this sense, the process of scoring data with a given model is inherently the same as applying any function, such as a square root function, to a set of data.

Scoring is only available with SPSS Statistics Server and can be done interactively by users working in distributed analysis mode. For scoring large data files, you'll probably want to make use of the Batch Facility, a separate executable provided with SPSS Statistics Server. For information about using the Batch Facility, see the *Batch Facility User's Guide*, provided as a PDF file on the SPSS Statistics Server product CD.

The scoring process consists of the following:

- ▶ **Loading** a model from a file in XML (PMML) format.
- ▶ **Computing** your scores as a new variable, using the `ApplyModel` or `StrApplyModel` function available in the **Compute Variable** dialog box.

For details concerning the `ApplyModel` or `StrApplyModel` function, see Scoring Expressions in the Transformation Expressions section of the *Command Syntax Reference*.

The following table lists the procedures that support the export of model specifications to XML. The exported models can be used with SPSS Statistics Server to score new data as described above. The full list of model types that can be scored with SPSS Statistics Server is provided in the description of the `ApplyModel` function.

Procedure name	Command name	Option
Discriminant	DISCRIMINANT	Base
Linear Regression	REGRESSION	Base
TwoStep Cluster	TWOSTEP CLUSTER	Base
Generalized Linear Models	GENLIN	Advanced Statistics
Complex Samples General Linear Model	CSGLM	Complex Samples

Procedure name	Command name	Option
Complex Samples Logistic Regression	CSLOGISTIC	Complex Samples
Complex Samples Ordinal Regression	CSORDINAL	Complex Samples
Logistic Regression	LOGISTIC REGRESSION	Regression
Multinomial Logistic Regression	NOMREG	Regression
Decision Tree	TREE	Decision Tree

Loading a Saved Model

The Load Model dialog box allows you to load predictive models saved in XML (PMML) format and is only available when working in distributed analysis mode. Loading models is a necessary first step to scoring data with them.

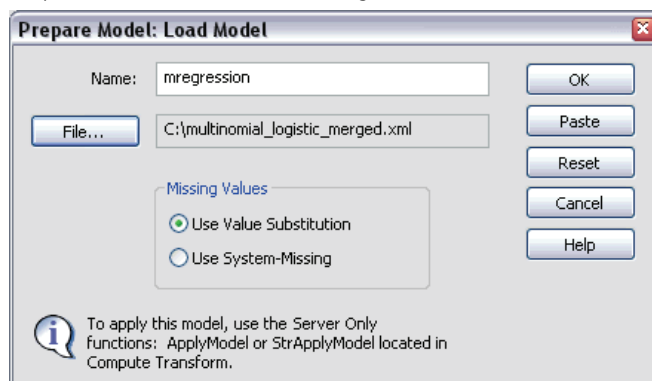
Figure 8-33
Load Model output

SPSS Variables	Model Variables					
	Name	Label	Type	Width	Role	Measurement
Age_Group	Age_Group	Age group	Numeric	8	Factor	Nominal
Frequency	Frequency	Number of purchases	Numeric	7	Covariate	Continuous
Gender	Gender	Gender	String	6	Factor	Nominal
Has_Broadband	Has_Broadband	Has home broadband connection	Numeric	10	Factor	Nominal
Has_Child	Has_Child	Has children	Numeric	10	Factor	Nominal
Income_Group	Income_Group	Income group	Numeric	10	Factor	Nominal
Log_Amount	Log_Amount	Log-transformed amount	Numeric	8	Covariate	Continuous
Recency	Recency	Days since last purchase	Numeric	10	Covariate	Continuous
Response	Response	Response	Numeric	10	Target	Nominal

To Load a Model

- From the menus choose:
Transform
Prepare Model
Load Model...

Figure 8-34
Prepare Model Load Model dialog box



- ▶ Enter a name to associate with this model. Each loaded model must have a unique name.
- ▶ Click File and select a model file. The resulting Open File dialog box displays the files that are available in distributed analysis mode. This includes files on the machine where SPSS Statistics Server is installed and files on your local computer that reside in shared folders or on shared drives.

Note: When you score data, the model will be applied to variables in the active dataset with the same names as the variables from the model file. You can map variables in the original model to different variables in the active dataset with the use of command syntax (see the `MODEL HANDLE` command).

Name. A name used to identify this model. The rules for valid model names are the same as for [variable names](#) (see Variable Names in Chapter 5 on p. 72), with the addition of the \$ character as an allowed first character. You will use this name to specify the model when scoring data with the `ApplyModel` or `StrApplyModel` functions.

File. The XML (PMML) file containing the model specifications.

Missing Values

This group of options controls the treatment of missing values, encountered during the scoring process, for the predictor variables defined in the model. A missing value in the context of scoring refers to one of the following:

- A predictor variable contains no value. For numeric variables, this means the system-missing value. For string variables, this means a null string.
- The value has been defined as user-missing, in the model, for the given predictor. Values defined as user-missing in the active dataset, but not in the model, are not treated as missing values in the scoring process.
- The predictor variable is categorical and the value is not one of the categories defined in the model.

Use Value Substitution. Attempt to use value substitution when scoring cases with missing values. The method for determining a value to substitute for a missing value depends on the type of predictive model.

- **SPSS Statistics Models.** For independent variables in linear regression and discriminant models, if mean value substitution for missing values was specified when building and saving the model, then this mean value is used in place of the missing value in the scoring computation, and scoring proceeds. If the mean value is not available, then the system-missing value is returned.
- **AnswerTree models & TREE command models.** For the CHAID and Exhaustive CHAID models, the biggest child node is selected for a missing split variable. The biggest child node is the one with the largest population among the child nodes using learning sample cases. For C&RT and QUEST models, surrogate split variables (if any) are used first. (Surrogate splits are splits that attempt to match the original split as closely as possible using alternate predictors.) If no surrogate splits are specified or all surrogate split variables are missing, the biggest child node is used.

- **Clementine models.** Linear regression models are handled as described under SPSS Statistics models. Logistic regression models are handled as described under Logistic Regression models. C&RT Tree models are handled as described for C&RT models under AnswerTree models.
- **Logistic Regression models.** For covariates in logistic regression models, if a mean value of the predictor was included as part of the saved model, then this mean value is used in place of the missing value in the scoring computation, and scoring proceeds. If the predictor is categorical (for example, a factor in a logistic regression model), or if the mean value is not available, then the system-missing value is returned.

Use System-Missing. Return the system-missing value when scoring a case with a missing value.

Displaying a List of Loaded Models

You can obtain a list of the currently loaded models. From the menus (only available in distributed analysis mode) choose:

```
Transform
  Prepare Model
    List Model(s)
```

This will generate a Model Handles table. The table contains a list of all currently loaded models and includes the name (referred to as the model handle) assigned to the model, the type of model, the path to the model file, and the method for handling missing values.

Figure 8-35
Listing of loaded models

Model Handles				
Name	Type	Input File Name	Validation handling	Missing handling
mregression	NOMREG	C:\models\mregression.xml	ON	SUBSTITUTE
tree	TREE	C:\models\tree.xml	ON	SUBSTITUTE

Additional Features Available with Command Syntax

From the Load Model dialog box, you can paste your selections into a syntax window and edit the resulting MODEL HANDLE command syntax. This allows you to:

- Map variables in the original model to different variables in the active dataset (with the MAP subcommand). By default, the model is applied to variables in the active dataset with the same names as the variables from the model file.

See the *Command Syntax Reference* for complete syntax information.

File Handling and File Transformations

Data files are not always organized in the ideal form for your specific needs. You may want to combine data files, sort the data in a different order, select a subset of cases, or change the unit of analysis by grouping cases together. A wide range of file transformation capabilities is available, including the ability to:

Sort data. You can sort cases based on the value of one or more variables.

Transpose cases and variables. The SPSS Statistics data file format reads rows as cases and columns as variables. For data files in which this order is reversed, you can switch the rows and columns and read the data in the correct format.

Merge files. You can merge two or more data files. You can combine files with the same variables but different cases or the same cases but different variables.

Select subsets of cases. You can restrict your analysis to a subset of cases or perform simultaneous analyses on different subsets.

Aggregate data. You can change the unit of analysis by aggregating cases based on the value of one or more grouping variables.

Weight data. You can weight cases for analysis based on the value of a weight variable.

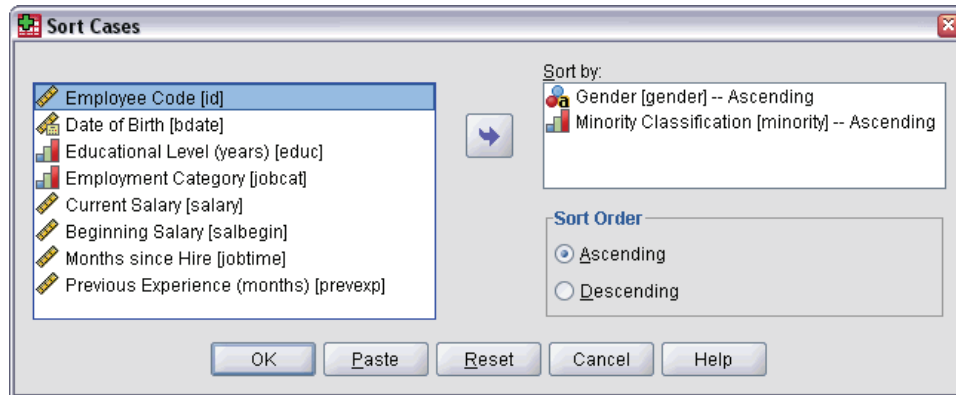
Restructure data. You can restructure data to create a single case (record) from multiple cases or create multiple cases from a single case.

Sort Cases

This dialog box sorts cases (rows) of the data file based on the values of one or more sorting variables. You can sort cases in ascending or descending order.

- If you select multiple sort variables, cases are sorted by each variable within categories of the preceding variable on the Sort list. For example, if you select *gender* as the first sorting variable and *minority* as the second sorting variable, cases will be sorted by minority classification within each gender category.
- The sort sequence is based on the locale-defined order (and is not necessarily the same as the numerical order of the character codes). The default locale is the operating system locale. You can control the locale with the Language setting on the General tab of the Options dialog box (Edit menu).

Figure 9-1
Sort Cases dialog box



To Sort Cases

- ▶ From the menus choose:
 - Data
 - Sort Cases...
- ▶ Select one or more sorting variables.

Sort Variables

You can sort the variables in the active dataset based on the values of any of the variable attributes (e.g., variable name, data type, measurement level), including custom variable attributes.

- Values can be sorted in ascending or descending order.
- You can save the original (pre-sorted) variable order in a custom variable attribute.
- Sorting by values of custom variable attributes is limited to custom variable attributes that are currently visible in Variable View.

For more information on custom variable attributes, see [Custom Variable Attributes](#).

To Sort Variables

In Variable View of the Data Editor:

- ▶ Right-click on the attribute column heading and from the context menu choose Sort Ascending or Sort Descending.
- or
- ▶ From the menus in Variable View or Data View, choose:
 - Data
 - Sort Variables
- ▶ Select the attribute you want to use to sort variables.
- ▶ Select the sort order (ascending or descending).

Figure 9-2
Sort Variables dialog



- The list of variable attributes matches the attribute column names displayed in Variable View of the Data Editor.
- You can save the original (pre-sorted) variable order in a custom variable attribute. For each variable, the value of the attribute is an integer value indicating its position prior to sorting; so by sorting variables based on the value of that custom attribute you can restore the original variable order.

Transpose

Transpose creates a new data file in which the rows and columns in the original data file are transposed so that cases (rows) become variables and variables (columns) become cases. Transpose automatically creates new variable names and displays a list of the new variable names.

- A new string variable that contains the original variable name, *case_lbl*, is automatically created.
- If the active dataset contains an ID or name variable with unique values, you can use it as the name variable, and its values will be used as variable names in the transposed data file. If it is a numeric variable, the variable names start with the letter *V*, followed by the numeric value.
- User-missing values are converted to the system-missing value in the transposed data file. To retain any of these values, change the definition of missing values in the Variable View in the Data Editor.

To Transpose Variables and Cases

- ▶ From the menus choose:
 - Data
 - Transpose...
- ▶ Select one or more variables to transpose into cases.

Merging Data Files

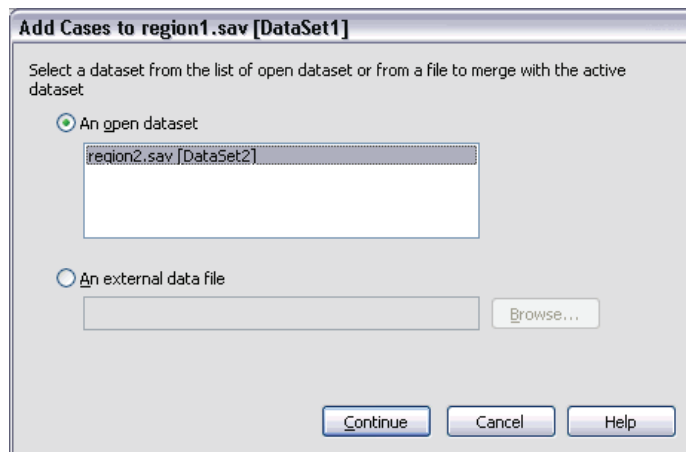
You can merge data from two files in two different ways. You can:

- Merge the active dataset with another open dataset or SPSS Statistics data file containing the same variables but different cases.
- Merge the active dataset with another open dataset or SPSS Statistics data file containing the same cases but different variables.

To Merge Files

- ▶ From the menus choose:
 - Data
 - Merge Files
- ▶ Select Add Cases or Add Variables.

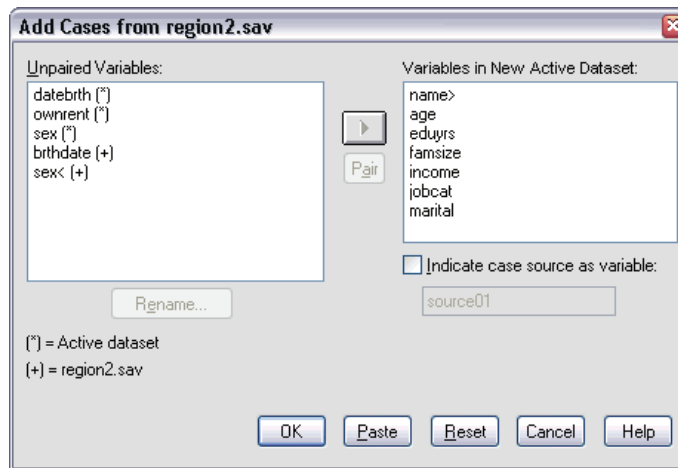
Figure 9-3
Selecting files to merge



Add Cases

Add Cases merges the active dataset with a second dataset or external SPSS Statistics data file that contains the same variables (columns) but different cases (rows). For example, you might record the same information for customers in two different sales regions and maintain the data for each region in separate files. The second dataset can be an external SPSS Statistics data file or a dataset available in the current session.

Figure 9-4
Add Cases dialog box



Unpaired Variables. Variables to be excluded from the new, merged data file. Variables from the active dataset are identified with an asterisk (*). Variables from the other dataset are identified with a plus sign (+). By default, this list contains:

- Variables from either data file that do not match a variable name in the other file. You can create pairs from unpaired variables and include them in the new, merged file.
- Variables defined as numeric data in one file and string data in the other file. Numeric variables cannot be merged with string variables.
- String variables of unequal width. The defined width of a string variable must be the same in both data files.

Variables in New Active Dataset. Variables to be included in the new, merged data file. By default, all of the variables that match both the name and the data type (numeric or string) are included on the list.

- You can remove variables from the list if you do not want them to be included in the merged file.
- Any unpaired variables included in the merged file will contain missing data for cases from the file that does not contain that variable.

Indicate case source as variable. Indicates the source data file for each case. This variable has a value of 0 for cases from the active dataset and a value of 1 for cases from the external data file.

To Merge Data Files with the Same Variables and Different Cases

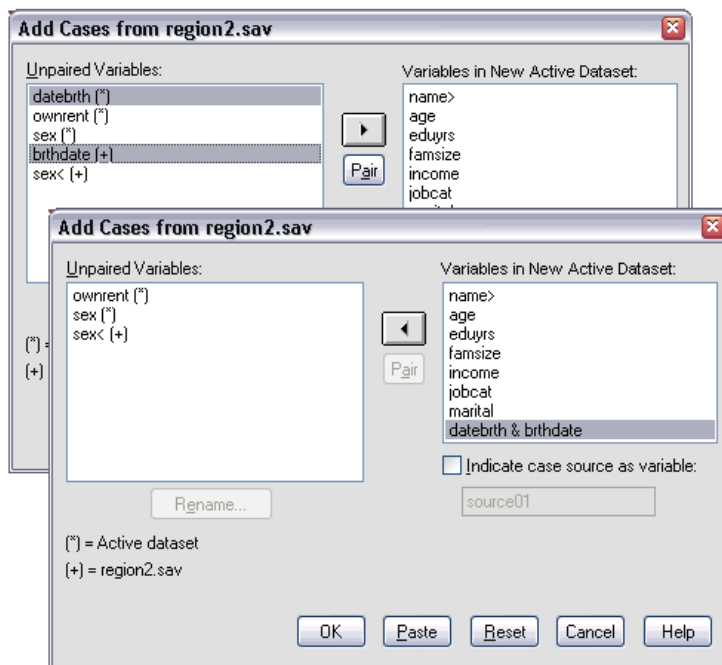
- ▶ Open at least one of the data files that you want to merge. If you have multiple datasets open, make one of the datasets that you want to merge the active dataset. The cases from this file will appear first in the new, merged data file.
- ▶ From the menus choose:
 - Data
 - Merge Files
 - Add Cases...

- ▶ Select the dataset or external SPSS Statistics data file to merge with the active dataset.
- ▶ Remove any variables that you do not want from the Variables in New Active Dataset list.
- ▶ Add any variable pairs from the Unpaired Variables list that represent the same information recorded under different variable names in the two files. For example, date of birth might have the variable name *brthdate* in one file and *datebrth* in the other file.

To Select a Pair of Unpaired Variables

- ▶ Click one of the variables on the Unpaired Variables list.
- ▶ Ctrl-click the other variable on the list. (Press the Ctrl key and click the left mouse button at the same time.)
- ▶ Click Pair to move the variable pair to the Variables in New Active Dataset list. (The variable name from the active dataset is used as the variable name in the merged file.)

Figure 9-5
Selecting pairs of variables with Ctrl-click



Add Cases: Rename

You can rename variables from either the active dataset or the other dataset before moving them from the unpaired list to the list of variables to be included in the merged data file. Renaming variables enables you to:

- Use the variable name from the other dataset rather than the name from the active dataset for variable pairs.
- Include two variables with the same name but of unmatched types or different string widths. For example, to include both the numeric variable *sex* from the active dataset and the string variable *sex* from the other dataset, one of them must be renamed first.

Add Cases: Dictionary Information

Any existing dictionary information (variable and value labels, user-missing values, display formats) in the active dataset is applied to the merged data file.

- If any dictionary information for a variable is undefined in the active dataset, dictionary information from the other dataset is used.
- If the active dataset contains any defined value labels or user-missing values for a variable, any additional value labels or user-missing values for that variable in the other dataset are ignored.

Merging More Than Two Data Sources

Using command syntax, you can merge up to 50 datasets and/or data files. For more information, see the `ADD FILES` command in the *Command Syntax Reference* (available from the Help menu).

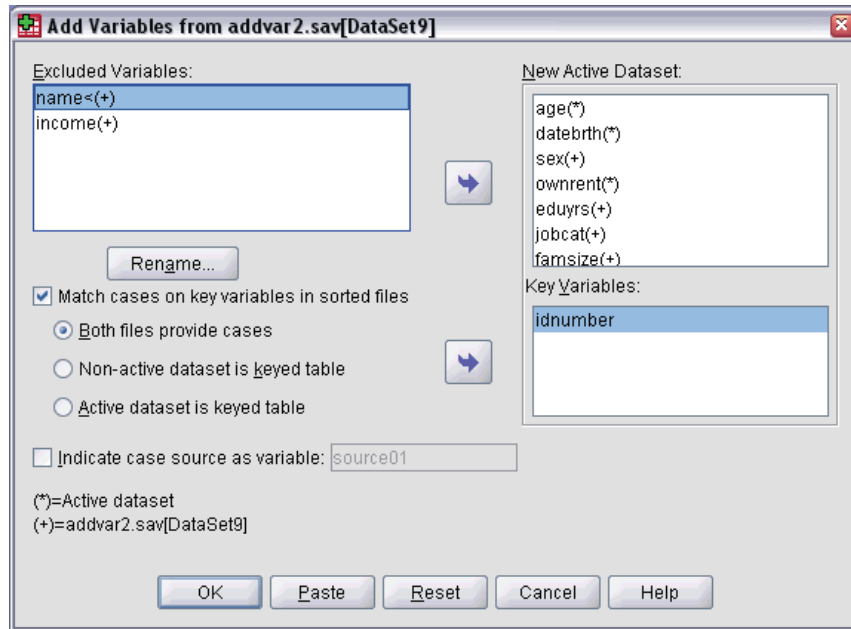
Add Variables

Add Variables merges the active dataset with another open dataset or external SPSS Statistics data file that contains the same cases (rows) but different variables (columns). For example, you might want to merge a data file that contains pre-test results with one that contains post-test results.

- Cases must be sorted in the same order in both datasets.
- If one or more key variables are used to match cases, the two datasets must be sorted by ascending order of the key variable(s).
- Variable names in the second data file that duplicate variable names in the active dataset are excluded by default because Add Variables assumes that these variables contain duplicate information.

Indicate case source as variable. Indicates the source data file for each case. This variable has a value of 0 for cases from the active dataset and a value of 1 for cases from the external data file.

Figure 9-6
Add Variables dialog box



Excluded Variables. Variables to be excluded from the new, merged data file. By default, this list contains any variable names from the other dataset that duplicate variable names in the active dataset. Variables from the active dataset are identified with an asterisk (*). Variables from the other dataset are identified with a plus sign (+). If you want to include an excluded variable with a duplicate name in the merged file, you can rename it and add it to the list of variables to be included.

New Active Dataset. Variables to be included in the new, merged dataset. By default, all unique variable names in both datasets are included on the list.

Key Variables. If some cases in one dataset do not have matching cases in the other dataset (that is, some cases are missing in one dataset), use key variables to identify and correctly match cases from the two datasets. You can also use key variables with table lookup files.

- The key variables must have the same names in both datasets.
- Both datasets must be sorted by ascending order of the key variables, and the order of variables on the Key Variables list must be the same as their sort sequence.
- Cases that do not match on the key variables are included in the merged file but are not merged with cases from the other file. Unmatched cases contain values for only the variables in the file from which they are taken; variables from the other file contain the system-missing value.

Non-active or active dataset is keyed table. A keyed table, or **table lookup file**, is a file in which data for each “case” can be applied to multiple cases in the other data file. For example, if one file contains information on individual family members (such as sex, age, education) and the other file contains overall family information (such as total income, family size, location), you can use the file of family data as a table lookup file and apply the common family data to each individual family member in the merged data file.

To Merge Files with the Same Cases but Different Variables

- ▶ Open at least one of the data files that you want to merge. If you have multiple datasets open, make one of the datasets that you want to merge the active dataset.
- ▶ From the menus choose:
 - Data
 - Merge Files
 - Add Variables...
- ▶ Select the dataset or external SPSS Statistics data file to merge with the active dataset.

To Select Key Variables

- ▶ Select the variables from the external file variables (+) on the Excluded Variables list.
- ▶ Select Match cases on key variables in sorted files.
- ▶ Add the variables to the Key Variables list.

The key variables must exist in both the active dataset and the other dataset. Both datasets must be sorted by ascending order of the key variables, and the order of variables on the Key Variables list must be the same as their sort sequence.

Add Variables: Rename

You can rename variables from either the active dataset or the other dataset before moving them to the list of variables to be included in the merged data file. This is primarily useful if you want to include two variables with the same name that contain different information in the two files.

Merging More Than Two Data Sources

Using command syntax, you can merge up to 50 datasets and/or data files. For more information, see the `MATCH FILES` command in the *Command Syntax Reference* (available from the Help menu).

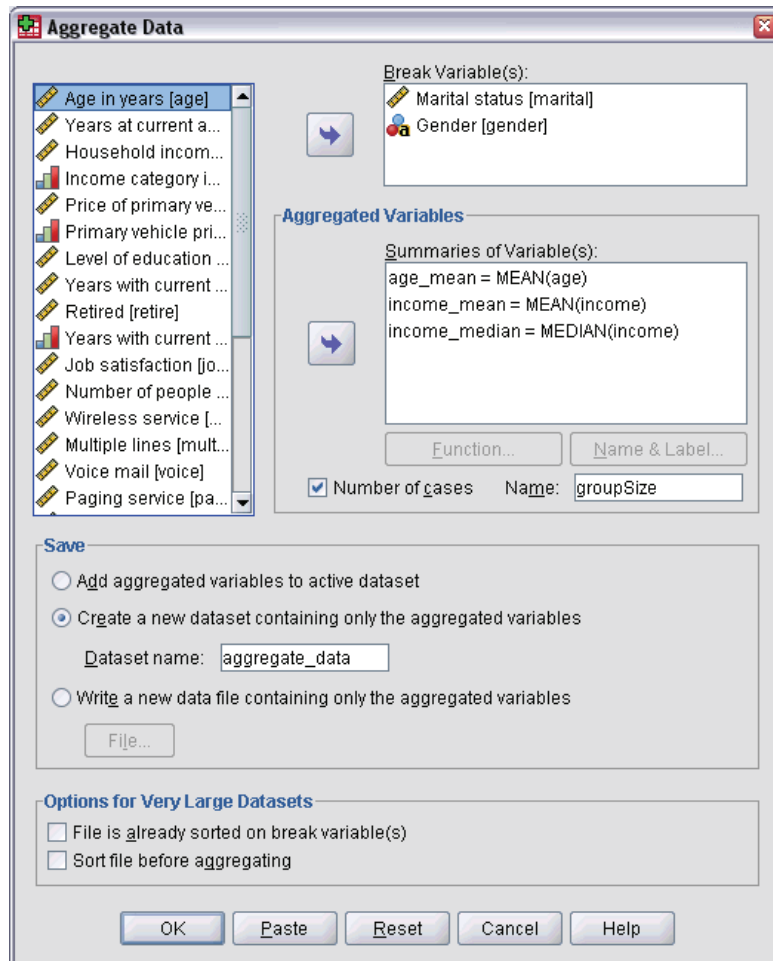
Aggregate Data

Aggregate Data aggregates groups of cases in the active dataset into single cases and creates a new, aggregated file or creates new variables in the active dataset that contain aggregated data. Cases are aggregated based on the value of *zero* or more break (grouping) variables. If no break variables are specified, then the entire dataset is a single break group.

- If you create a new, aggregated data file, the new data file contains one case for each group defined by the break variables. For example, if there is one break variable with two values, the new data file will contain only two cases. If no break variable is specified, the new data file will contain one case.
- If you add aggregate variables to the active dataset, the data file itself is not aggregated. Each case with the same value(s) of the break variable(s) receives the same values for the new aggregate variables. For example, if *gender* is the only break variable, all males would

receive the same value for a new aggregate variable that represents average age. If no break variable is specified, all cases would receive the same value for a new aggregate variable that represents average age.

Figure 9-7
Aggregate Data dialog box



Break Variable(s). Cases are grouped together based on the values of the break variables. Each unique combination of break variable values defines a group. When creating a new, aggregated data file, all break variables are saved in the new file with their existing names and dictionary information. The break variable, if specified, can be either numeric or string.

Aggregated Variables. Source variables are used with aggregate functions to create new aggregate variables. The aggregate variable name is followed by an optional variable label, the name of the aggregate function, and the source variable name in parentheses.

You can override the default aggregate variable names with new variable names, provide descriptive variable labels, and change the functions used to compute the aggregated data values. You can also create a variable that contains the number of cases in each break group.

To Aggregate a Data File

- ▶ From the menus choose:
 - Data
 - Aggregate...
- ▶ Optionally select break variables that define how cases are grouped to create aggregated data. If no break variables are specified, then the entire dataset is a single break group.
- ▶ Select one or more aggregate variables.
- ▶ Select an aggregate function for each aggregate variable.

Saving Aggregated Results

You can add aggregate variables to the active dataset or create a new, aggregated data file.

- **Add aggregated variables to active dataset.** New variables based on aggregate functions are added to the active dataset. The data file itself is not aggregated. Each case with the same value(s) of the break variable(s) receives the same values for the new aggregate variables.
- **Create a new dataset containing only the aggregated variables.** Saves aggregated data to a new dataset in the current session. The dataset includes the break variables that define the aggregated cases and all aggregate variables defined by aggregate functions. The active dataset is unaffected.
- **Write a new data file containing only the aggregated variables.** Saves aggregated data to an external data file. The file includes the break variables that define the aggregated cases and all aggregate variables defined by aggregate functions. The active dataset is unaffected.

Sorting Options for Large Data Files

For very large data files, it may be more efficient to aggregate presorted data.

File is already sorted on break variable(s). If the data have already been sorted by values of the break variables, this option enables the procedure to run more quickly and use less memory. Use this option with caution.

- Data must be sorted by values of the break variables in the same order as the break variables specified for the Aggregate Data procedure.
- If you are adding variables to the active dataset, select this option only if the data are sorted by ascending values of the break variables.

Sort file before aggregating. In very rare instances with large data files, you may find it necessary to sort the data file by values of the break variables prior to aggregating. This option is not recommended unless you encounter memory or performance problems.

Aggregate Data: Aggregate Function

This dialog box specifies the function to use to calculate aggregated data values for selected variables on the Aggregate Variables list in the Aggregate Data dialog box. Aggregate functions include:

- Summary functions for numeric variables, including mean, median, standard deviation, and sum
- Number of cases, including unweighted, weighted, nonmissing, and missing
- Percentage or fraction of values above or below a specified value
- Percentage or fraction of values inside or outside of a specified range

Figure 9-8
Aggregate Function dialog box

Aggregate Data: Aggregate Function

Summary Statistics	Specific Values	Number of cases
<input type="radio"/> Mean	<input type="radio"/> First	<input type="radio"/> Weighted
<input type="radio"/> Median	<input type="radio"/> Last	<input type="radio"/> Weighted missing
<input type="radio"/> Sum	<input type="radio"/> Minimum	<input type="radio"/> Unweighted
<input type="radio"/> Standard Deviation	<input type="radio"/> Maximum	<input type="radio"/> Unweighted missing

Percentages

Above Value:

Below

Inside Low: High:

Outside

Fractions

Above Value:

Below

Inside Low: High:

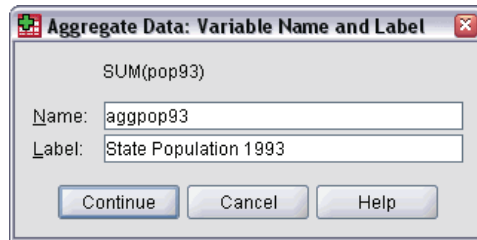
Outside

Continue Cancel Help

Aggregate Data: Variable Name and Label

Aggregate Data assigns default variable names for the aggregated variables in the new data file. This dialog box enables you to change the variable name for the selected variable on the Aggregate Variables list and provide a descriptive variable label. [For more information, see Variable Names in Chapter 5 on p. 72.](#)

Figure 9-9
Variable Name and Label dialog box

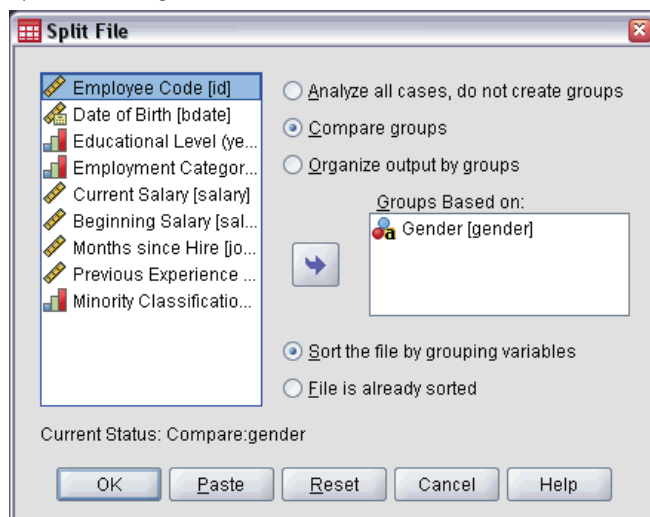


Split File

Split File splits the data file into separate groups for analysis based on the values of one or more grouping variables. If you select multiple grouping variables, cases are grouped by each variable within categories of the preceding variable on the Groups Based On list. For example, if you select *gender* as the first grouping variable and *minority* as the second grouping variable, cases will be grouped by minority classification within each gender category.

- You can specify up to eight grouping variables.
- Each eight bytes of a long string variable (string variables longer than eight bytes) counts as a variable toward the limit of eight grouping variables.
- Cases should be sorted by values of the grouping variables and in the same order that variables are listed in the Groups Based On list. If the data file isn't already sorted, select Sort the file by grouping variables.

Figure 9-10
Split File dialog box



Compare groups. Split-file groups are presented together for comparison purposes. For pivot tables, a single pivot table is created and each split-file variable can be moved between table dimensions. For charts, a separate chart is created for each split-file group and the charts are displayed together in the Viewer.

Organize output by groups. All results from each procedure are displayed separately for each split-file group.

To Split a Data File for Analysis

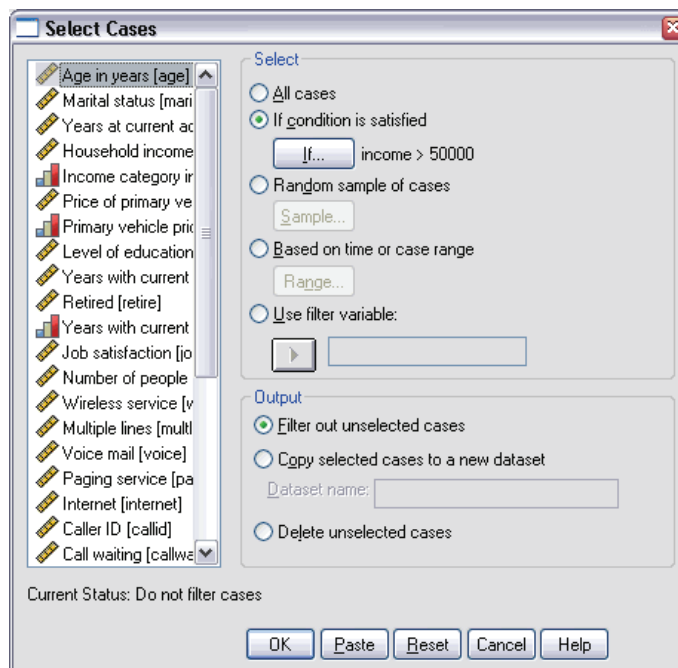
- ▶ From the menus choose:
Data
Split File...
- ▶ Select Compare groups or Organize output by groups.
- ▶ Select one or more grouping variables.

Select Cases

Select Cases provides several methods for selecting a subgroup of cases based on criteria that include variables and complex expressions. You can also select a random sample of cases. The criteria used to define a subgroup can include:

- Variable values and ranges
- Date and time ranges
- Case (row) numbers
- Arithmetic expressions
- Logical expressions
- Functions

Figure 9-11
Select Cases dialog box



All cases. Turns case filtering off and uses all cases.

If condition is satisfied. Use a conditional expression to select cases. If the result of the conditional expression is true, the case is selected. If the result is false or missing, the case is not selected.

Random sample of cases. Selects a random sample based on an approximate percentage or an exact number of cases.

Based on time or case range. Selects cases based on a range of case numbers or a range of dates/times.

Use filter variable. Use the selected numeric variable from the data file as the filter variable. Cases with any value other than 0 or missing for the filter variable are selected.

Output

This section controls the treatment of unselected cases. You can choose one of the following alternatives for the treatment of unselected cases:

- **Filter out unselected cases.** Unselected cases are not included in the analysis but remain in the dataset. You can use the unselected cases later in the session if you turn filtering off. If you select a random sample or if you select cases based on a conditional expression, this generates a variable named *filter_\$* with a value of 1 for selected cases and a value of 0 for unselected cases.
- **Copy selected cases to a new dataset.** Selected cases are copied to a new dataset, leaving the original dataset unaffected. Unselected cases are not included in the new dataset and are left in their original state in the original dataset.
- **Delete unselected cases.** Unselected cases are deleted from the dataset. Deleted cases can be recovered only by exiting from the file without saving any changes and then reopening the file. The deletion of cases is permanent if you save the changes to the data file.

Note: If you delete unselected cases and save the file, the cases cannot be recovered.

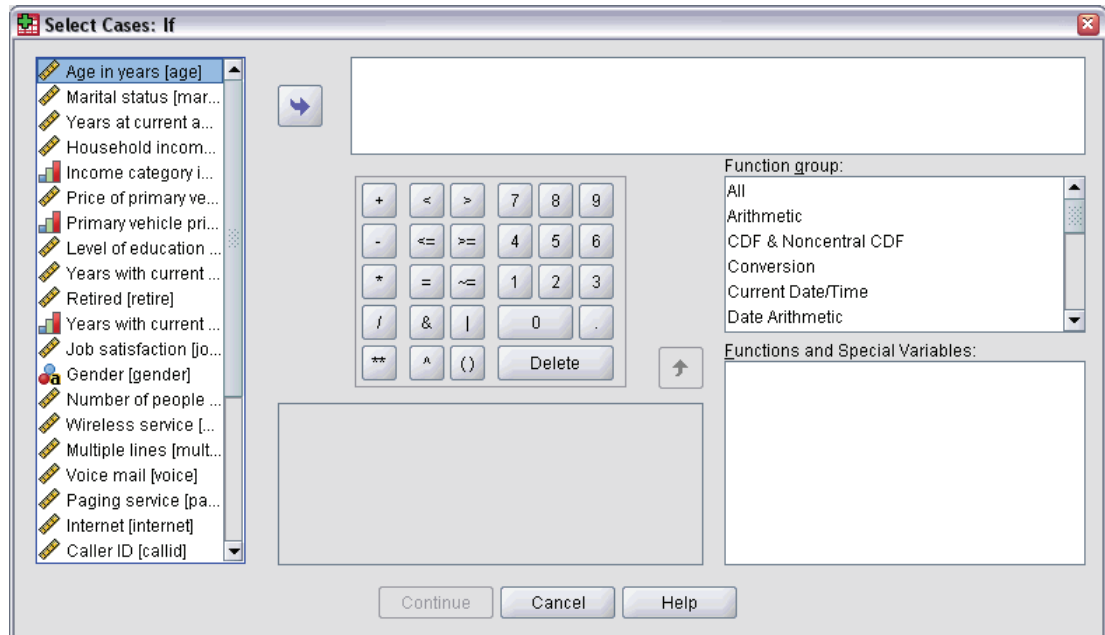
To Select a Subset of Cases

- ▶ From the menus choose:
 - Data
 - Select Cases...
- ▶ Select one of the methods for selecting cases.
- ▶ Specify the criteria for selecting cases.

Select Cases: If

This dialog box allows you to select subsets of cases using conditional expressions. A conditional expression returns a value of *true*, *false*, or *missing* for each case.

Figure 9-12
Select Cases If dialog box




- If the result of a conditional expression is *true*, the case is included in the selected subset.
- If the result of a conditional expression is *false* or *missing*, the case is not included in the selected subset.
- Most conditional expressions use one or more of the six relational operators (<, >, <=, >=, =, and ~=) on the calculator pad.
- Conditional expressions can include variable names, constants, arithmetic operators, numeric (and other) functions, logical variables, and relational operators.

Select Cases: Random Sample

This dialog box allows you to select a random sample based on an approximate percentage or an exact number of cases. Sampling is performed without replacement; so, the same case cannot be selected more than once.

Figure 9-13
Select Cases Random Sample dialog box



Approximately. Generates a random sample of approximately the specified percentage of cases. Since this routine makes an independent pseudo-random decision for each case, the percentage of cases selected can only approximate the specified percentage. The more cases there are in the data file, the closer the percentage of cases selected is to the specified percentage.

Exactly. A user-specified number of cases. You must also specify the number of cases from which to generate the sample. This second number should be less than or equal to the total number of cases in the data file. If the number exceeds the total number of cases in the data file, the sample will contain proportionally fewer cases than the requested number.

Select Cases: Range

This dialog box selects cases based on a range of case numbers or a range of dates or times.

- Case ranges are based on row number as displayed in the Data Editor.
- Date and time ranges are available only for **time series data** with defined date variables (Data menu, Define Dates).

Figure 9-14

Select Cases Range dialog box for range of cases (no defined date variables)





Figure 9-15

Select Cases Range dialog box for time series data with defined date variables



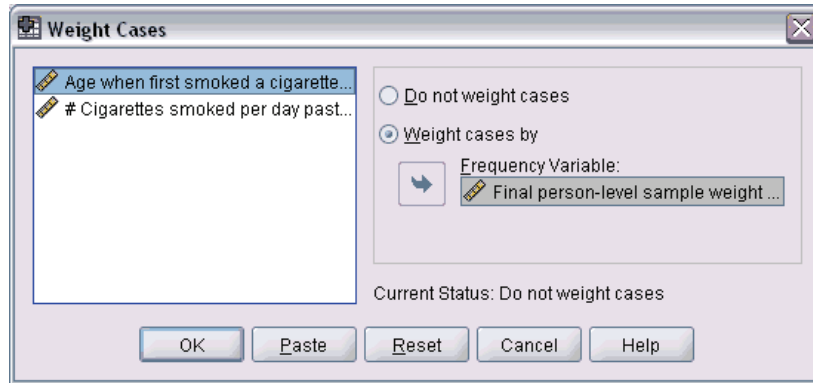
Weight Cases

Weight Cases gives cases different weights (by simulated replication) for statistical analysis.

- The values of the weighting variable should indicate the number of observations represented by single cases in your data file.
- Cases with zero, negative, or missing values for the weighting variable are excluded from analysis.
- Fractional values are valid and some procedures, such as Frequencies, Crosstabs, and Custom Tables, will use fractional weight values. However, most procedures treat the weighting variable as a replication weight and will simply round fractional weights to the nearest integer.

Some procedures ignore the weighting variable completely, and this limitation is noted in the procedure-specific documentation.

Figure 9-16
Weight Cases dialog box



Once you apply a weight variable, it remains in effect until you select another weight variable or turn off weighting. If you save a weighted data file, weighting information is saved with the data file. You can turn off weighting at any time, even after the file has been saved in weighted form.

Weights in Crosstabs. The Crosstabs procedure has several options for handling case weights. For more information, see [Crosstabs Cell Display in Chapter 18 on p. 291](#).

Weights in scatterplots and histograms. Scatterplots and histograms have an option for turning case weights on and off, but this does not affect cases with a zero, negative, or missing value for the weight variable. These cases remain excluded from the chart even if you turn weighting off from within the chart.

To Weight Cases

- ▶ From the menus choose:
Data
Weight Cases...
- ▶ Select Weight cases by.
- ▶ Select a frequency variable.

The values of the frequency variable are used as case weights. For example, a case with a value of 3 for the frequency variable will represent three cases in the weighted data file.

Restructuring Data

Use the Restructure Data Wizard to restructure your data for the procedure that you want to use. The wizard replaces the current file with a new, restructured file. The wizard can:

- Restructure selected variables into cases
- Restructure selected cases into variables
- Transpose all data

To Restructure Data

- ▶ From the menus choose:
 - Data
 - Restructure...
- ▶ Select the type of restructuring that you want to do.
- ▶ Select the data to restructure.

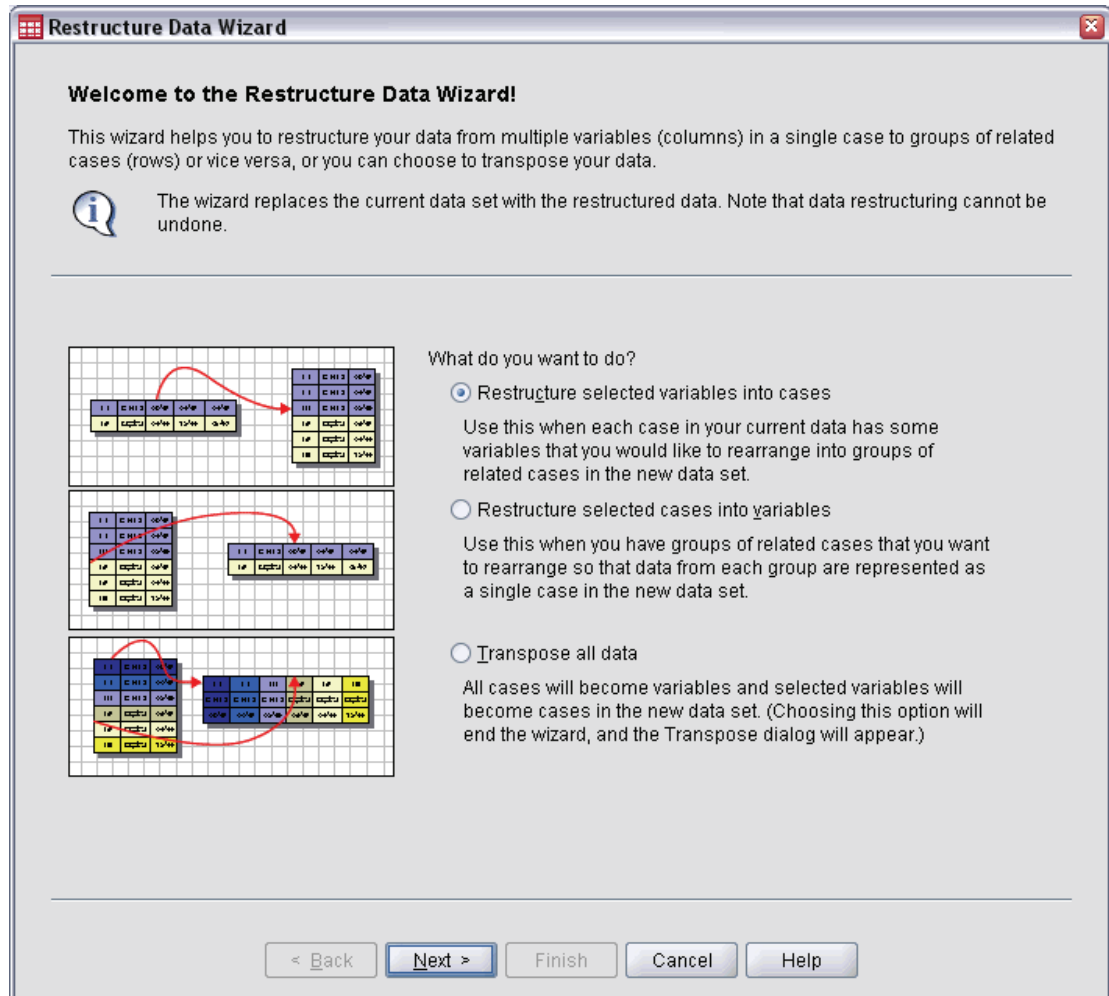
Optionally, you can:

- Create identification variables, which allow you to trace a value in the new file back to a value in the original file
- Sort the data prior to restructuring
- Define options for the new file
- Paste the command syntax into a syntax window

Restructure Data Wizard: Select Type

Use the Restructure Data Wizard to restructure your data. In the first dialog box, select the type of restructuring that you want to do.

Figure 9-17
Restructure Data Wizard



- **Restructure selected variables into cases.** Choose this when you have groups of related columns in your data and you want them to appear in groups of rows in the new data file. If you choose this, the wizard will display the steps for **Variables to Cases**.
- **Restructure selected cases into variables.** Choose this when you have groups of related rows in your data and you want them to appear in groups of columns in the new data file. If you choose this, the wizard will display the steps for **Cases to Variables**.
- **Transpose all data.** Choose this when you want to transpose your data. All rows will become columns and all columns will become rows in the new data. This choice closes the Restructure Data Wizard and opens the Transpose Data dialog box.

Deciding How to Restructure the Data

A **variable** contains information that you want to analyze—for example, a measurement or a score. A **case** is an observation—for example, an individual. In a *simple* data structure, each variable is a single column in your data and each case is a single row. So, for example, if you were

measuring test scores for all students in a class, all score values would appear in only one column, and there would be a row for each student.

When you analyze data, you are often analyzing how a variable varies according to some condition. The condition can be a specific experimental treatment, a demographic, a point in time, or something else. In data analysis, conditions of interest are often referred to as **factors**. When you analyze factors, you have a *complex* data structure. You may have information about a variable in more than one column in your data (for example, a column for each level of a factor), or you may have information about a case in more than one row (for example, a row for each level of a factor). The Restructure Data Wizard helps you to restructure files with a complex data structure.

The structure of the current file and the structure that you want in the new file determine the choices that you make in the wizard.

How are the data arranged in the current file? The current data may be arranged so that factors are recorded in a *separate* variable (in groups of cases) or *with* the variable (in groups of variables).

- **Groups of cases.** Does the current file have variables and conditions recorded in separate columns? For example:

var	factor
8	1
9	1
3	2
1	2

In this example, the first two rows are a **case group** because they are related. They contain data for the same factor level. In SPSS Statistics data analysis, the factor is often referred to as a **grouping variable** when the data are structured this way.

- **Groups of columns.** Does the current file have variables and conditions recorded in the same column? For example:

var_1	var_2
8	3
9	1

In this example, the two columns are a **variable group** because they are related. They contain data for the same variable—*var_1* for factor level 1 and *var_2* for factor level 2. In SPSS Statistics data analysis, the factor is often referred to as a **repeated measure** when the data are structured this way.

How should the data be arranged in the new file? This is usually determined by the procedure that you want to use to analyze your data.

- **Procedures that require groups of cases.** Your data must be structured in case groups to do analyses that require a grouping variable. Examples are *univariate*, *multivariate*, and *variance components* with General Linear Model, Mixed Models, and OLAP Cubes and *independent*

samples with T Test or Nonparametric Tests. If your current data structure is variable groups and you want to do these analyses, select Restructure selected variables into cases.

- **Procedures that require groups of variables.** Your data must be structured in variable groups to analyze repeated measures. Examples are *repeated measures* with General Linear Model, *time-dependent covariate* analysis with Cox Regression Analysis, *paired samples* with T Test, or *related samples* with Nonparametric Tests. If your current data structure is case groups and you want to do these analyses, select Restructure selected cases into variables.

Example of Variables to Cases

In this example, test scores are recorded in separate columns for each factor, *A* and *B*.

Figure 9-18
Current data for variables to cases

	score_a	score_b
1	1014.00	864.00
2	684.00	636.00
3	810.00	638.00

You want to do an independent-samples *t* test. You have a column group consisting of *score_a* and *score_b*, but you don't have the **grouping variable** that the procedure requires. Select Restructure selected variables into cases in the Restructure Data Wizard, restructure one variable group into a new variable named *score*, and create an index named *group*. The new data file is shown in the following figure.

Figure 9-19
New, restructured data for variables to cases

	group	score
1	SCORE_A	1014.00
2	SCORE_B	864.00
3	SCORE_A	684.00
4	SCORE_B	636.00
5	SCORE_A	810.00
6	SCORE_B	638.00

When you run the independent-samples *t* test, you can now use *group* as the grouping variable.

Example of Cases to Variables

In this example, test scores are recorded twice for each subject—before and after a treatment.

Figure 9-20
Current data for cases to variables

	id	scor	time
1	1	1014.00	bef
2	1	864.00	aft
3	2	684.00	bef
4	2	636.00	aft

You want to do a paired-samples t test. Your data structure is case groups, but you don't have the **repeated measures** for the paired variables that the procedure requires. Select Restructure selected cases into variables in the Restructure Data Wizard, use *id* to identify the row groups in the current data, and use *time* to create the variable group in the new file.

Figure 9-21

New, restructured data for cases to variables

	id	aft	bef
1	1	864.00	1014.00
2	2	636.00	684.00

When you run the paired-samples t test, you can now use *bef* and *aft* as the variable pair.

Restructure Data Wizard (Variables to Cases): Number of Variable Groups

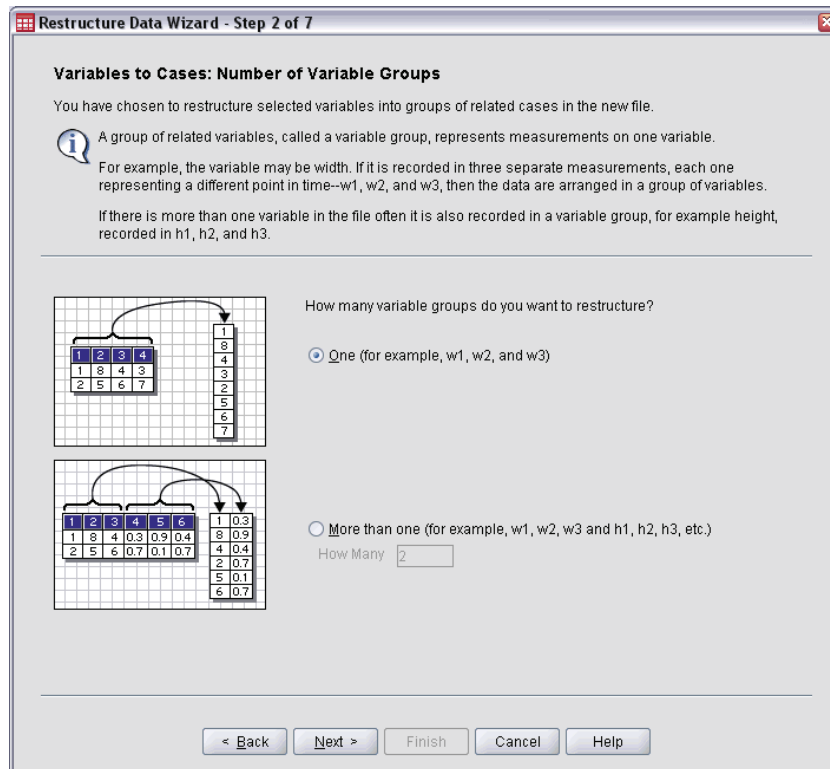
Note: The wizard presents this step if you choose to restructure variable groups into rows.

In this step, decide how many variable groups in the current file that you want to restructure in the new file.

How many variable groups are in the current file? Think about how many variable groups exist in the current data. A group of related columns, called a **variable group**, records repeated measures of the same variable in separate columns. For example, if you have three columns in the current data—*w1*, *w2*, and *w3*—that record **width**, you have one variable group. If you have an additional three columns—*h1*, *h2*, and *h3*—that record **height**, you have two variable groups.

How many variable groups should be in the new file? Consider how many variable groups you want to have represented in the new data file. You do not have to restructure all variable groups into the new file.

Figure 9-22
Restructure Data Wizard: Number of Variable Groups, Step 2



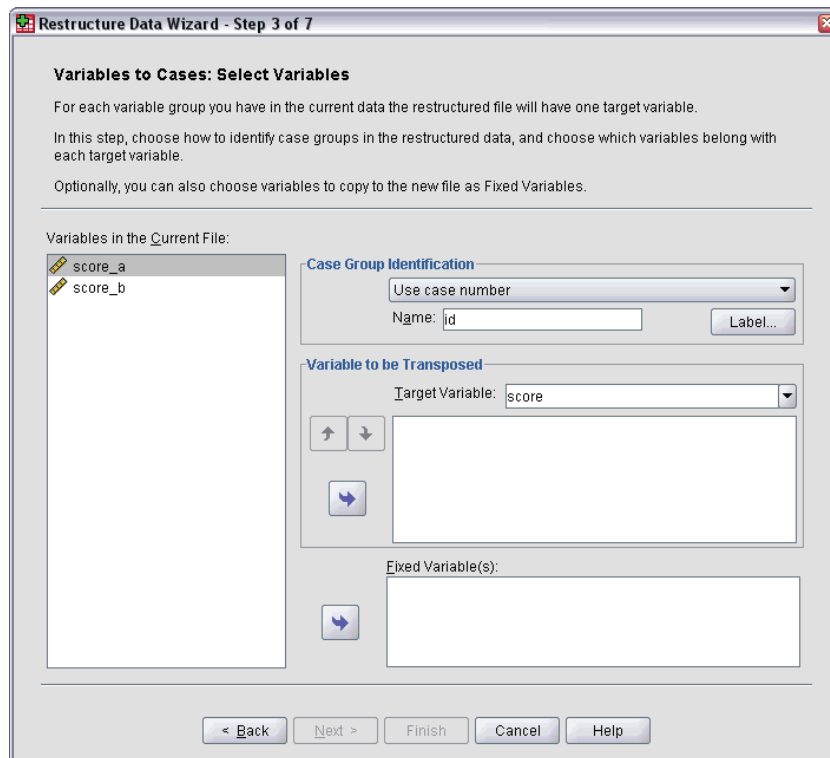
- **One.** The wizard will create a single restructured variable in the new file from one variable group in the current file.
- **More than one.** The wizard will create multiple restructured variables in the new file. The number that you specify affects the next step, in which the wizard automatically creates the specified number of new variables.

Restructure Data Wizard (Variables to Cases): Select Variables

Note: The wizard presents this step if you choose to restructure variable groups into rows.

In this step, provide information about how the variables in the current file should be used in the new file. You can also create a variable that identifies the rows in the new file.

Figure 9-23
Restructure Data Wizard: Select Variables, Step 3



How should the new rows be identified? You can create a variable in the new data file that identifies the row in the current data file that was used to create a group of new rows. The identifier can be a sequential case number or it can be the values of the variable. Use the controls in Case Group Identification to define the identification variable in the new file. Click a cell to change the default variable name and provide a descriptive variable label for the identification variable.

What should be restructured in the new file? In the previous step, you told the wizard how many variable groups you want to restructure. The wizard created one new variable for each group. The values for the variable group will appear in that variable in the new file. Use the controls in Variable to be Transposed to define the restructured variable in the new file.

To Specify One Restructured Variable

- ▶ Put the variables that make up the variable group that you want to transform into the Variable to be Transposed list. All of the variables in the group must be of the same type (numeric or string).

You can include the same variable more than once in the variable group (variables are copied rather than moved from the source variable list); its values are repeated in the new file.

To Specify Multiple Restructured Variables

- ▶ Select the first target variable that you want to define from the Target Variable drop-down list.

- ▶ Put the variables that make up the variable group that you want to transform into the Variable to be Transposed list. All of the variables in the group must be of the same type (numeric or string). You can include the same variable more than once in the variable group. (A variable is copied rather than moved from the source variable list, and its values are repeated in the new file.)
- ▶ Select the next target variable that you want to define, and repeat the variable selection process for all available target variables.
 - Although you can include the same variable more than once in the same target variable group, you cannot include the same variable in more than one target variable group.
 - Each target variable group list must contain the same number of variables. (Variables that are listed more than once are included in the count.)
 - The number of target variable groups is determined by the number of variable groups that you specified in the previous step. You can change the default variable names here, but you must return to the previous step to change the number of variable groups to restructure.
 - You must define variable groups (by selecting variables in the source list) for all available target variables before you can proceed to the next step.

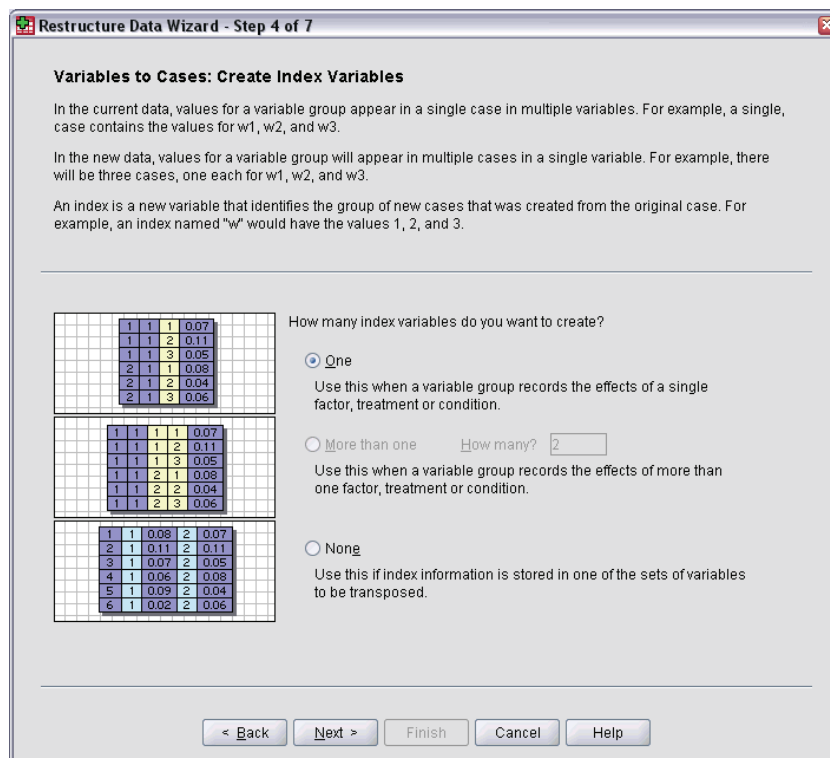
What should be copied into the new file? Variables that aren't restructured can be copied into the new file. Their values will be propagated in the new rows. Move variables that you want to copy into the new file into the Fixed Variable(s) list.

Restructure Data Wizard (Variables to Cases): Create Index Variables

Note: The wizard presents this step if you choose to restructure variable groups into rows.

In this step, decide whether to create index variables. An index is a new variable that sequentially identifies a row group based on the original variable from which the new row was created.

Figure 9-24
Restructure Data Wizard: Create Index Variables, Step 4



How many index variables should be in the new file? Index variables can be used as grouping variables in procedures. In most cases, a single index variable is sufficient; however, if the variable groups in your current file reflect multiple factor levels, multiple indices may be appropriate.

- **One.** The wizard will create a single index variable.
- **More than one.** The wizard will create multiple indices and enter the number of indices that you want to create. The number that you specify affects the next step, in which the wizard automatically creates the specified number of indices.
- **None.** Select this if you do not want to create index variables in the new file.

Example of One Index for Variables to Cases

In the current data, there is one variable group, *width*, and one factor, *time*. Width was measured three times and recorded in *w1*, *w2*, and *w3*.

Figure 9-25
Current data for one index

	subject	w1	w2	w3
1	1	6.70	4.30	5.70
2	2	7.10	5.90	5.60

We'll restructure the variable group into a single variable, *width*, and create a single numeric index. The new data are shown in the following table.

Figure 9-26
New, restructured data with one index

	subject	index	width
1	1	1	6.70
2	1	2	4.30
3	1	3	5.70
4	2	1	7.10
5	2	2	5.90
6	2	3	5.60

Index starts with 1 and increments for each variable in the group. It restarts each time a new row is encountered in the original file. We can now use *index* in procedures that require a grouping variable.

Example of Two Indices for Variables to Cases

When a variable group records more than one factor, you can create more than one index; however, the current data must be arranged so that the levels of the first factor are a primary index within which the levels of subsequent factors cycle. In the current data, there is one variable group, *width*, and two factors, *A* and *B*. The data are arranged so that levels of factor *B* cycle within levels of factor *A*.

Figure 9-27
Current data for two indices

	subject	w_a1b1	w_a1b2	w_a2b1	w_a2b2
1	1	5.50	6.40	5.80	5.90
2	2	7.40	7.10	5.60	6.70

We'll restructure the variable group into a single variable, *width*, and create two indices. The new data are shown in the following table.

Figure 9-28
New, restructured data with two indices

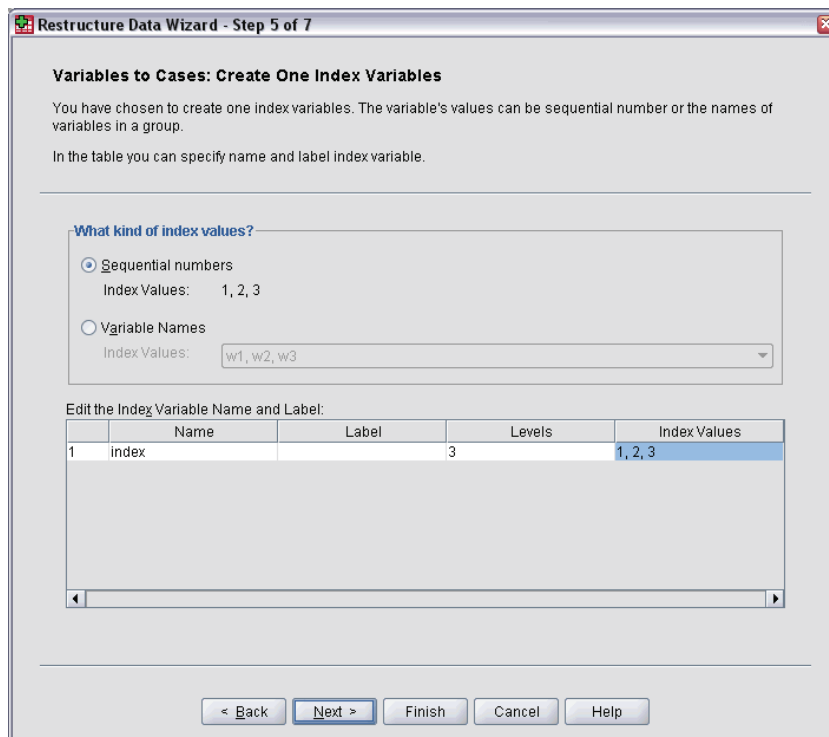
	subject	index_a	index_b	width
1	1	1	1	5.50
2	1	1	2	6.40
3	1	2	1	5.80
4	1	2	2	5.90
5	2	1	1	7.40
6	2	1	2	7.10
7	2	2	1	5.60
8	2	2	2	6.70

Restructure Data Wizard (Variables to Cases): Create One Index Variable

Note: The wizard presents this step if you choose to restructure variable groups into rows and create one index variable.

In this step, decide what values you want for the index variable. The values can be sequential numbers or the names of the variables in an original variable group. You can also specify a name and a label for the new index variable.

Figure 9-29
Restructure Data Wizard: Create One Index Variable, Step 5



For more information, see [Example of One Index for Variables to Cases](#) on p. 193.

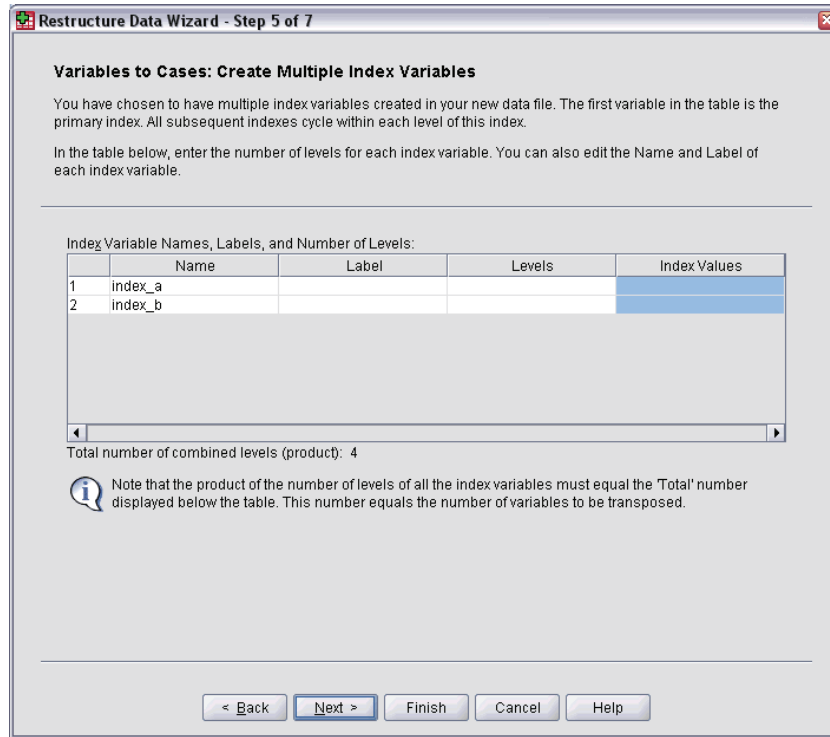
- **Sequential numbers.** The wizard will automatically assign sequential numbers as index values.
- **Variable names.** The wizard will use the names of the selected variable group as index values. Choose a variable group from the list.
- **Names and labels.** Click a cell to change the default variable name and provide a descriptive variable label for the index variable.

Restructure Data Wizard (Variables to Cases): Create Multiple Index Variables

Note: The wizard presents this step if you choose to restructure variable groups into rows and create multiple index variables.

In this step, specify the number of levels for each index variable. You can also specify a name and a label for the new index variable.

Figure 9-30
Restructure Data Wizard: Create Multiple Index Variables, Step 5



For more information, see [Example of Two Indices for Variables to Cases](#) on p. 194.

How many levels are recorded in the current file? Consider how many factor levels are recorded in the current data. A **level** defines a group of cases that experienced identical conditions. If there are multiple factors, the current data must be arranged so that the levels of the first factor are a primary index within which the levels of subsequent factors cycle.

How many levels should be in the new file? Enter the number of levels for each index. The values for multiple index variables are always sequential numbers. The values start at 1 and increment for each level. The first index increments the slowest, and the last index increments the fastest.

Total combined levels. You cannot create more levels than exist in the current data. Because the restructured data will contain one row for each combination of treatments, the wizard checks the number of levels that you create. It will compare the product of the levels that you create to the number of variables in your variable groups. They must match.

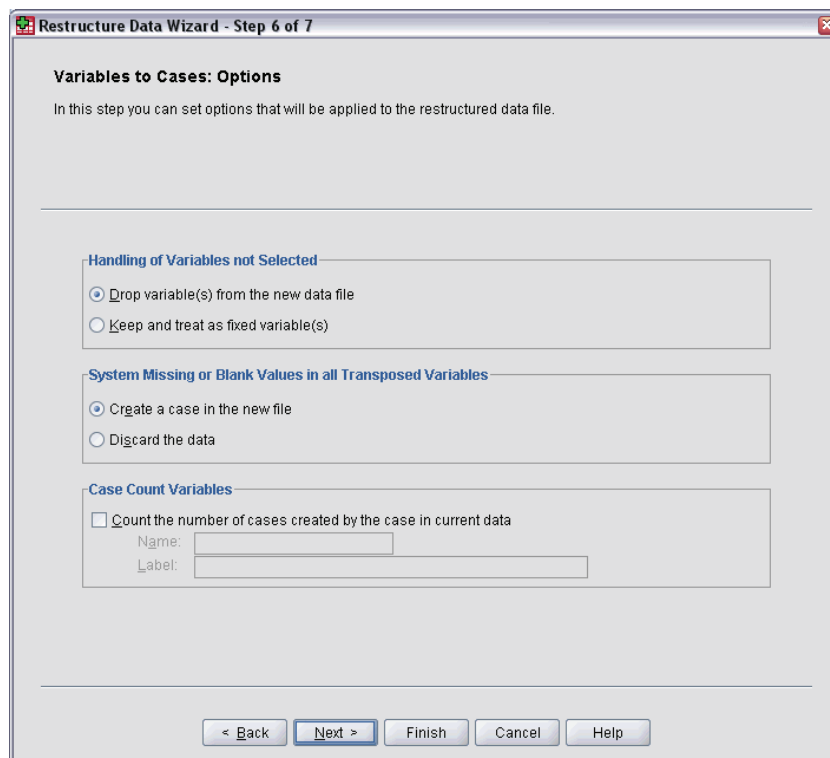
Names and labels. Click a cell to change the default variable name and provide a descriptive variable label for the index variables.

Restructure Data Wizard (Variables to Cases): Options

Note: The wizard presents this step if you choose to restructure variable groups into rows.

In this step, specify options for the new, restructured file.

Figure 9-31
Restructure Data Wizard: Options, Step 6



Drop unselected variables? In the Select Variables step (step 3), you selected variable groups to be restructured, variables to be copied, and an identification variable from the current data. The data from the selected variables will appear in the new file. If there are other variables in the current data, you can choose to discard or keep them.

Keep missing data? The wizard checks each potential new row for null values. A **null value** is a system-missing or blank value. You can choose to keep or discard rows that contain only null values.

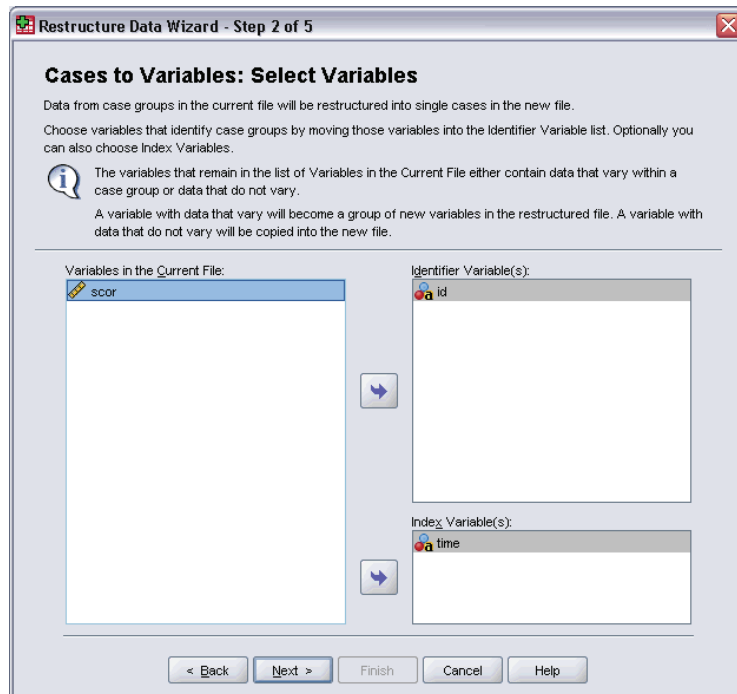
Create a count variable? The wizard can create a **count variable** in the new file. It contains the number of new rows generated by a row in the current data. A count variable may be useful if you choose to discard null values from the new file because that makes it possible to generate a different number of new rows for a given row in the current data. Click a cell to change the default variable name and provide a descriptive variable label for the count variable.

Restructure Data Wizard (Cases to Variables): Select Variables

Note: The wizard presents this step if you choose to restructure case groups into columns.

In this step, provide information about how the variables in the current file should be used in the new file.

Figure 9-32
Restructure Data Wizard: Select Variables, Step 2



What identifies case groups in the current data? A **case group** is a group of rows that are related because they measure the same observational unit—for example, an individual or an institution. The wizard needs to know which variables in the current file identify the case groups so that it can consolidate each group into a single row in the new file. Move variables that identify case groups in the current file into the Identifier Variable(s) list. Variables that are used to split the current data file are automatically used to identify case groups. Each time a new combination of identification values is encountered, the wizard will create a new row, so cases in the current file should be sorted by values of the identification variables in the same order that variables are listed in the Identifier Variable(s) list. If the current data file isn't already sorted, you can sort it in the next step.

How should the new variable groups be created in the new file? In the original data, a variable appears in a single column. In the new data file, that variable will appear in multiple new columns. **Index variables** are variables in the current data that the wizard should use to create the new columns. The restructured data will contain one new variable for each unique value in these columns. Move the variables that should be used to form the new variable groups to the Index Variable(s) list. When the wizard presents options, you can also choose to order the new columns by index.

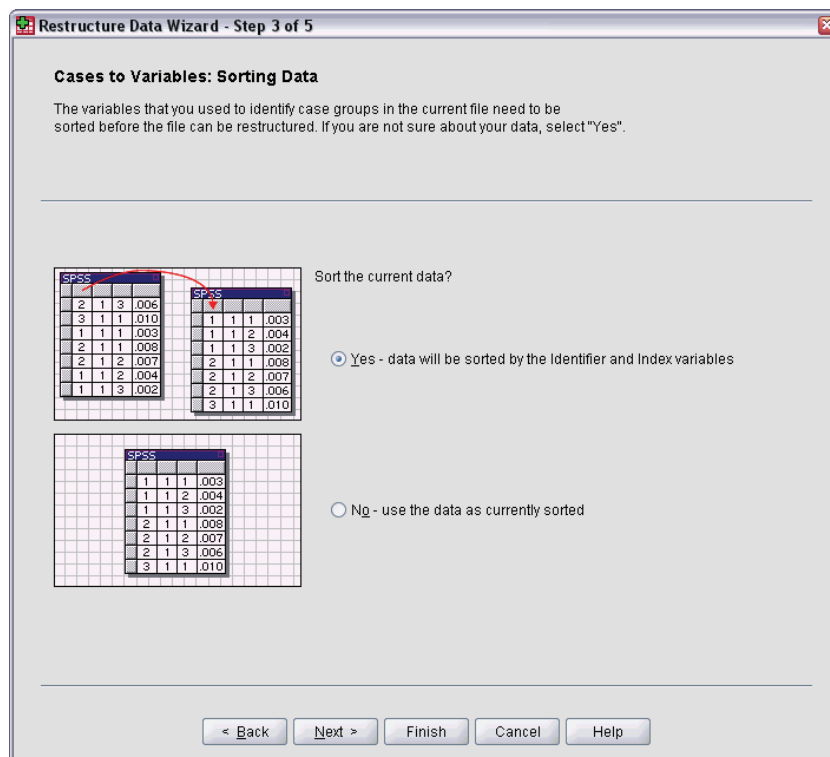
What happens to the other columns? The wizard automatically decides what to do with the variables that remain in the Current File list. It checks each variable to see if the data values vary within a case group. If they do, the wizard restructures the values into a variable group in the new file. If they don't, the wizard copies the values into the new file.

Restructure Data Wizard (Cases to Variables): Sort Data

Note: The wizard presents this step if you choose to restructure case groups into columns.

In this step, decide whether to sort the current file before restructuring it. Each time the wizard encounters a new combination of identification values, a new row is created, so it is important that the data are sorted by the variables that identify case groups.

Figure 9-33
Restructure Data Wizard: Sorting Data, Step 3



How are the rows ordered in the current file? Consider how the current data are sorted and which variables you are using to identify case groups (specified in the previous step).

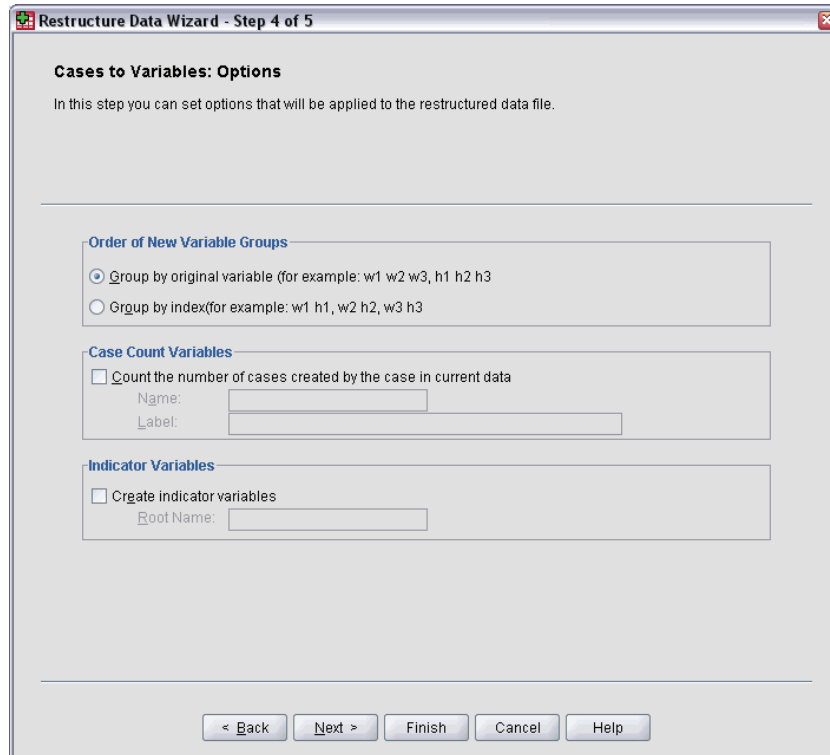
- **Yes.** The wizard will automatically sort the current data by the identification variables in the same order that variables are listed in the Identifier Variable(s) list in the previous step. Choose this when the data aren't sorted by the identification variables or when you aren't sure. This choice requires a separate pass of the data, but it guarantees that the rows are correctly ordered for restructuring.
- **No.** The wizard will not sort the current data. Choose this when you are sure that the current data are sorted by the variables that identify case groups.

Restructure Data Wizard (Cases to Variables): Options

Note: The wizard presents this step if you choose to restructure case groups into columns.

In this step, specify options for the new, restructured file.

Figure 9-34
Restructure Data Wizard: Options, Step 4



How should the new variable groups be ordered in the new file?

- **By variable.** The wizard groups the new variables created from an original variable together.
- **By index.** The wizard groups the variables according to the values of the index variables.

Example. The variables to be restructured are *w* and *h*, and the index is *month*:

w h month

Grouping by variable results in:

w.jan w.feb h.jan

Grouping by index results in:

w.jan h.jan w.feb

Create a count variable? The wizard can create a count variable in the new file. It contains the number of rows in the current data that were used to create a row in the new data file.

Create indicator variables? The wizard can use the index variables to create **indicator variables** in the new data file. It creates one new variable for each unique value of the index variable. The indicator variables signal the presence or absence of a value for a case. An indicator variable has the value of 1 if the case has a value; otherwise, it is 0.

Example. The index variable is *product*. It records the products that a customer purchased. The original data are:

customer	product
1	chick
1	eggs
2	eggs
3	chick

Creating an indicator variable results in one new variable for each unique value of *product*. The restructured data are:

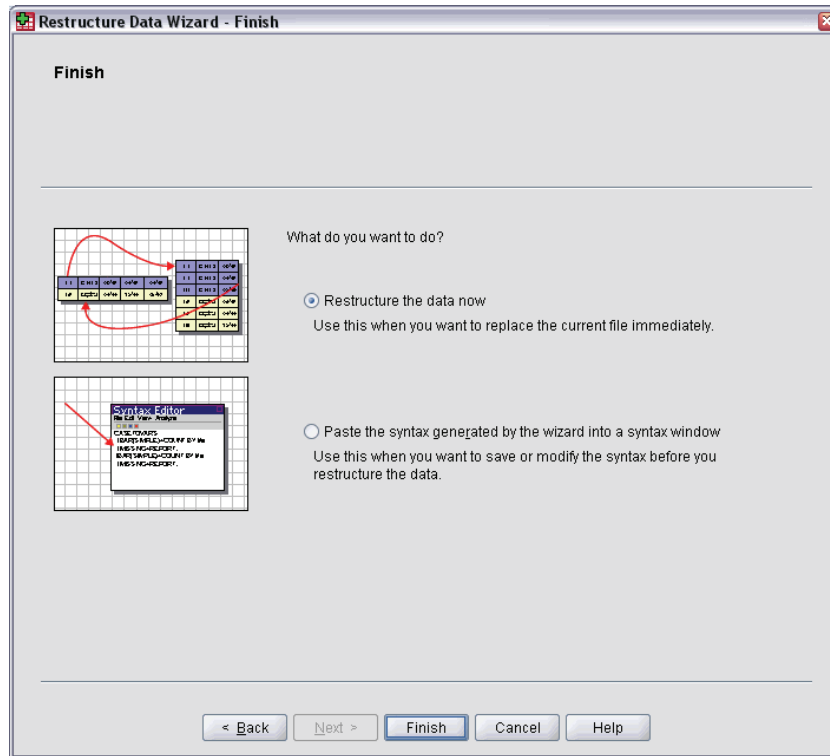
customer	indchick	indeggs
1	1	1
2	0	1
3	1	0

In this example, the restructured data could be used to get frequency counts of the products that customers buy.

Restructure Data Wizard: Finish

This is the final step of the Restructure Data Wizard. Decide what to do with your specifications.

Figure 9-35
Restructure Data Wizard: Finish



- **Restructure now.** The wizard will create the new, restructured file. Choose this if you want to replace the current file immediately.

Note: If original data are weighted, the new data will be weighted unless the variable that is used as the weight is restructured or dropped from the new file.

- **Paste syntax.** The wizard will paste the syntax it generates into a syntax window. Choose this when you are not ready to replace the current file, when you want to modify the syntax, or when you want to save it for future use.

Working with Output

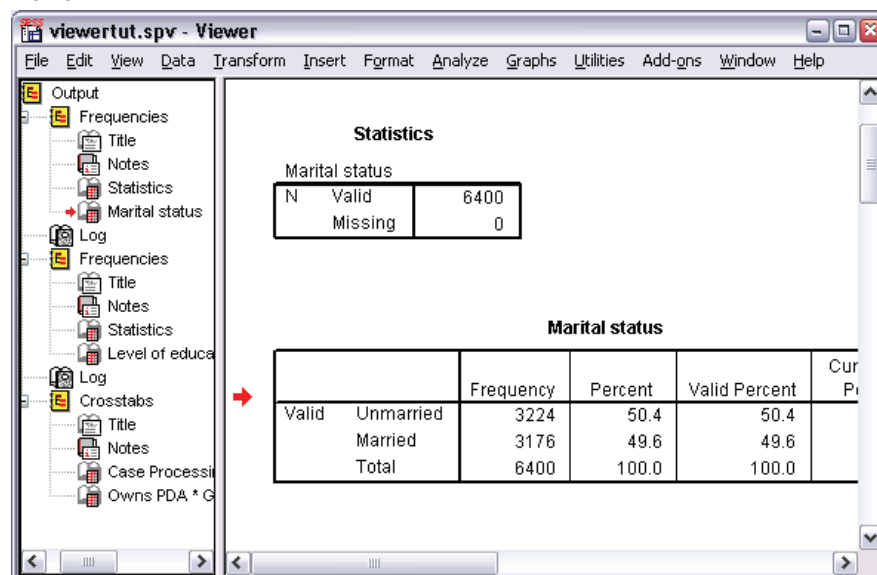
When you run a procedure, the results are displayed in a window called the Viewer. In this window, you can easily navigate to the output that you want to see. You can also manipulate the output and create a document that contains precisely the output that you want.

Viewer

Results are displayed in the Viewer. You can use the Viewer to:

- Browse results
- Show or hide selected tables and charts
- Change the display order of results by moving selected items
- Move items between the Viewer and other applications

Figure 10-1
Viewer



The Viewer is divided into two panes:

- The left pane contains an outline view of the contents.
- The right pane contains statistical tables, charts, and text output.

You can click an item in the outline to go directly to the corresponding table or chart. You can click and drag the right border of the outline pane to change the width of the outline pane.

Showing and Hiding Results

In the Viewer, you can selectively show and hide individual tables or results from an entire procedure. This process is useful when you want to shorten the amount of visible output in the contents pane.

To Hide Tables and Charts

- ▶ Double-click the item's book icon in the outline pane of the Viewer.

or

- ▶ Click the item to select it.
- ▶ From the menus choose:

View
Hide

or

- ▶ Click the closed book (Hide) icon on the Outlining toolbar.

The open book (Show) icon becomes the active icon, indicating that the item is now hidden.

To Hide Procedure Results

- ▶ Click the box to the left of the procedure name in the outline pane.

This hides all results from the procedure and collapses the outline view.

Moving, Deleting, and Copying Output

You can rearrange the results by copying, moving, or deleting an item or a group of items.

To Move Output in the Viewer

- ▶ Select the items in the outline or contents pane.
- ▶ Drag and drop the selected items into a different location.

To Delete Output in the Viewer

- ▶ Select the items in the outline or contents pane.
- ▶ Press the Delete key.

or

- ▶ From the menus choose:

Edit
Delete

Changing Initial Alignment

By default, all results are initially left-aligned. To change the initial alignment of new output items:

- ▶ From the menus choose:
 - Edit
 - Options
- ▶ Click the Viewer tab.
- ▶ In the Initial Output State group, select the item type (for example, pivot table, chart, text output).
- ▶ Select the alignment option you want.

Changing Alignment of Output Items

- ▶ In the outline or contents pane, select the items that you want to align.
- ▶ From the menus choose:
 - Format
 - Align Left

or

 - Format
 - Center

or

 - Format
 - Align Right

Viewer Outline

The outline pane provides a table of contents of the Viewer document. You can use the outline pane to navigate through your results and control the display. Most actions in the outline pane have a corresponding effect on the contents pane.

- Selecting an item in the outline pane displays the corresponding item in the contents pane.
- Moving an item in the outline pane moves the corresponding item in the contents pane.
- Collapsing the outline view hides the results from all items in the collapsed levels.

Controlling the outline display. To control the outline display, you can:

- Expand and collapse the outline view
- Change the outline level for selected items
- Change the size of items in the outline display
- Change the font that is used in the outline display

To Collapse and Expand the Outline View

- ▶ Click the box to the left of the outline item that you want to collapse or expand.

or

- ▶ Click the item in the outline.
- ▶ From the menus choose:
 - View
 - Collapse

or

- View
- Expand

To Change the Outline Level

- ▶ Click the item in the outline pane.
- ▶ Click the left arrow on the Outlining toolbar to promote the item (move the item to the left).

or

Click the right arrow on the Outlining toolbar to demote the item (move the item to the right).

or

- ▶ From the menus choose:
 - Edit
 - Outline
 - Promote

or

- Edit
- Outline
- Demote

Changing the outline level is particularly useful after you move items in the outline level. Moving items can change the outline level of the items, and you can use the left- and right-arrow buttons on the Outlining toolbar to restore the original outline level.

To Change the Size of Outline Items

- ▶ From the menus choose:
 - View
 - Outline Size
- ▶ Select the outline size (Small, Medium, or Large).

To Change the Font in the Outline

- ▶ From the menus choose:
 - View
 - Outline Font...
- ▶ Select a font.

Adding Items to the Viewer

In the Viewer, you can add items such as titles, new text, charts, or material from other applications.

To Add a Title or Text

Text items that are not connected to a table or chart can be added to the Viewer.

- ▶ Click the table, chart, or other object that will precede the title or text.
- ▶ From the menus choose:
 - Insert
 - New Title
- or*
- Insert
 - New Text
- ▶ Double-click the new object.
- ▶ Enter the text.

To Add a Text File

- ▶ In the outline pane or contents pane of the Viewer, click the table, chart, or other object that will precede the text.
- ▶ From the menus choose:
 - Insert
 - Text File...
- ▶ Select a text file.
 - To edit the text, double-click it.

Pasting Objects into the Viewer

Objects from other applications can be pasted into the Viewer. You can use either Paste After or Paste Special. Either type of pasting puts the new object after the currently selected object in the Viewer. Use Paste Special when you want to choose the format of the pasted object.

Finding and Replacing Information in the Viewer

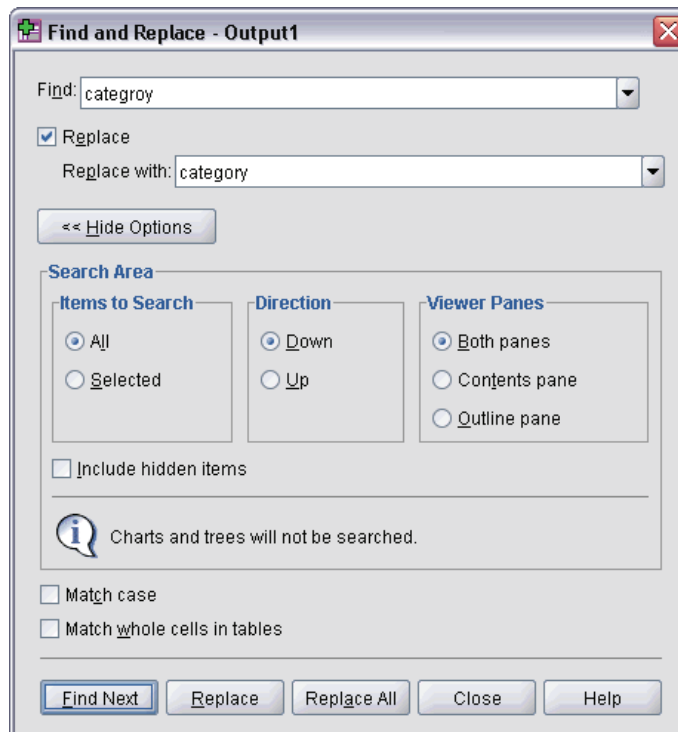
- ▶ To find or replace information in the Viewer, from the menus choose:

Edit
Find

or

Edit
Replace

Figure 10-2
Find and Replace dialog box



You can use Find and Replace to:

- Search the entire document or just the selected items.
- Search down or up from the current location.
- Search both panes or restrict the search to the contents or outline pane.
- Search for hidden items. These include any items hidden in the contents pane (for example, Notes tables, which are hidden by default) and hidden rows and columns in pivot tables.
- Restrict the search criteria to case-sensitive matches.
- Restrict the search criteria in pivot tables to matches of the entire cell contents.

Hidden Items and Pivot Table Layers

- Layers beneath the currently visible layer of a multidimensional pivot table are not considered hidden and will be included in the search area even when hidden items are not included in the search.
- Hidden items include hidden items in the contents pane (items with closed book icons in the outline pane or included within collapsed blocks of the outline pane) and rows and columns in pivot tables either hidden by default (for example, empty rows and columns are hidden by default) or manually hidden by editing the table and selectively hiding specific rows or columns. Hidden items are only included in the search if you explicitly select Include hidden items.
- In both cases, the hidden or nonvisible element that contains the search text or value is displayed when it is found, but the item is returned to its original state afterward.

Copying Output into Other Applications

Output objects can be copied and pasted into other applications, such as a word-processing program or a spreadsheet. You can paste output in a variety of formats. Depending on the target application, some or all of the following formats may be available:

Picture (metafile). Pivot tables and charts can be pasted as metafile pictures. The picture format can be resized in the other application, and sometimes a limited amount of editing can be done with the facilities of the other application. Pivot tables that are pasted as pictures retain all borders and font characteristics. For pivot tables, select Files from the Paste Special menu. This format is available only on Windows operating systems.

RTF (rich text format). Pivot tables can be pasted into other applications in RTF format. In most applications, this means that the pivot table is pasted as a table that can then be edited in the other application.

Note: Microsoft Word may not display extremely wide tables properly.

Bitmap. Charts and other graphics can be pasted into other applications as bitmaps.

BIFF. The contents of a table can be pasted into a spreadsheet and retain numeric precision. This format is available only on Windows operating systems.

Text. The contents of a table can be copied and pasted as text. This process can be useful for applications such as e-mail, where the application can accept or transmit only text.

If the target application supports multiple available formats, it may have a Paste Special menu item that allows you to select the format, or it may automatically display a list of available formats.

To Copy and Paste Output Items into Another Application

- ▶ Select the objects in either the outline or contents pane of the Viewer.
- ▶ From the Viewer menus choose:
 - Edit
 - Copy

- ▶ From the menus in the target application choose:

- Edit
 - Paste

or

- Edit
 - Paste Special...

- **Paste.** Output is copied to the Clipboard in a number of formats. Each application determines the “best” format to use for Paste.
- **Paste Special.** Results are copied to the Clipboard in multiple formats. Paste Special allows you to select the format that you want from the list of formats that are available to the target application.

Copying and pasting model views. All model views, including tables, are pasted as graphic images.

Copying and pasting pivot tables. When pivot tables are pasted in Word/RTF format, tables that are too wide for the document width will either be wrapped, scaled down to fit the document width, or left unchanged, depending on the pivot table options settings. [For more information, see Pivot Table Options in Chapter 48 on p. 505.](#)

Export Output

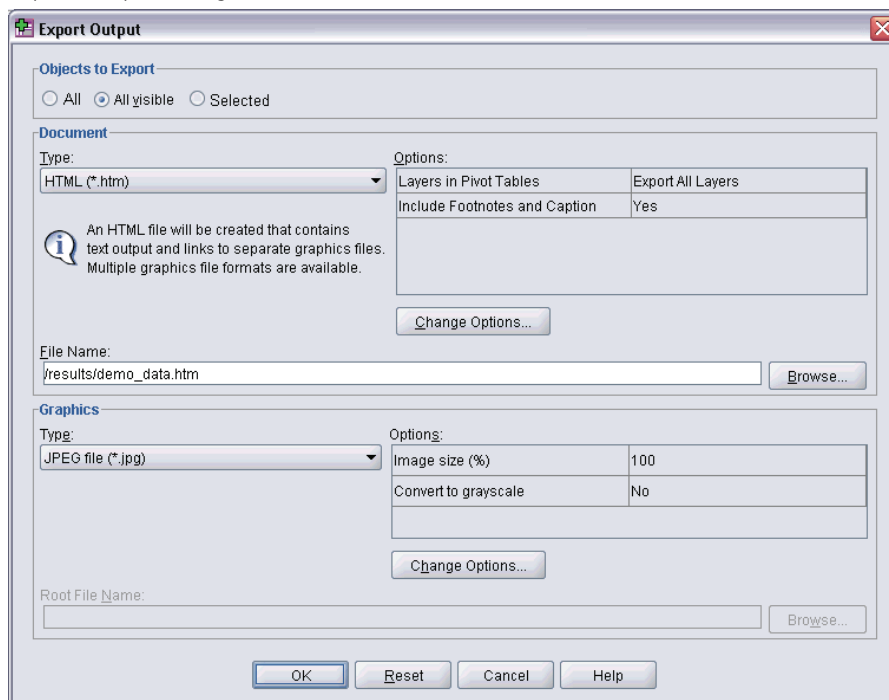
Export Output saves Viewer output in HTML, text, Word/RTF, Excel, PowerPoint (requires PowerPoint 97 or later), and PDF formats. Charts can also be exported in a number of different graphics formats.

Note: Export to PowerPoint is available only on Windows operating systems and is not available in the Student Version.

To Export Output

- ▶ Make the Viewer the active window (click anywhere in the window).
- ▶ From the menus choose:
 - File
 - Export...
- ▶ Enter a filename (or prefix for charts) and select an export format.

Figure 10-3
Export Output dialog box



Objects to Export. You can export all objects in the Viewer, all visible objects, or only selected objects.

Document Type. The available options are:

- **Word/RTF (*.doc).** Pivot tables are exported as Word tables with all formatting attributes intact—for example, cell borders, font styles, and background colors. Text output is exported as formatted RTF. Charts, tree diagrams, and model views are included in PNG format.

Note: Microsoft Word may not display extremely wide tables properly.

- **Excel (*.xls).** Pivot table rows, columns, and cells are exported as Excel rows, columns, and cells, with all formatting attributes intact—for example, cell borders, font styles, and background colors. Text output is exported with all font attributes intact. Each line in the text output is a row in the Excel file, with the entire contents of the line contained in a single cell. Charts, tree diagrams, and model views are included in PNG format.
- **HTML (*.htm).** Pivot tables are exported as HTML tables. Text output is exported as preformatted HTML. Charts, tree diagrams, and model views are embedded by reference, and you should export charts in a suitable format for inclusion in HTML documents (for example, PNG and JPEG).
- **Portable Document Format (*.pdf).** All output is exported as it appears in Print Preview, with all formatting attributes intact.
- **PowerPoint file (*.ppt).** Pivot tables are exported as Word tables and are embedded on separate slides in the PowerPoint file, with one slide for each pivot table. All formatting attributes of the pivot table are retained—for example, cell borders, font styles, and background colors. Charts, tree diagrams, and model views are exported in TIFF format. Text output is not included.

Export to PowerPoint is available only on Windows operating systems.

- **Text (*.txt).** Text output formats include plain text, UTF-8, and UTF-16. Pivot tables can be exported in tab-separated or space-separated format. All text output is exported in space-separated format. For charts, tree diagrams, and model views, a line is inserted in the text file for each graphic, indicating the image filename.
- **None (Graphics Only).** Available export formats include: EPS, JPEG, TIFF, PNG, and BMP. On Windows operating systems, EMF (enhanced metafile) format is also available.

Output Management System. You can also automatically export all output or user-specified types of output as Word, Excel, PDF, HTML, text or SPSS Statistics-format data files. [For more information, see Output Management System in Chapter 52 on p. 552.](#)

HTML Options

The following options are available for exporting output in HTML format:

Layers in pivot tables. By default, inclusion or exclusion of pivot table layers is controlled by the table properties for each pivot table. You can override this setting and include all layers or exclude all but the currently visible layer. [For more information, see Table Properties: Printing in Chapter 11 on p. 240.](#)

Include footnotes and captions. Controls the inclusion or exclusion of all pivot table footnotes and captions.

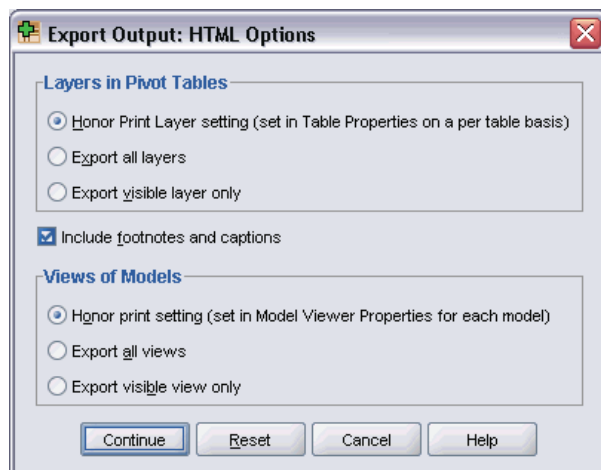
Views of Models. By default, inclusion or exclusion of model views is controlled by the model properties for each model. You can override this setting and include all views or exclude all but the currently visible view. [For more information, see Model Properties in Chapter 12 on p. 251.](#) (Note: all model views, including tables, are exported as graphics.)

Note: For HTML, you can also control the image file format for exported charts. [For more information, see Graphics Format Options on p. 221.](#)

To Set HTML Export Options

- ▶ Select HTML as the export format.
- ▶ Click Change Options.

Figure 10-4
HTML Export Output Options



Word/RTF Options

The following options are available for exporting output in Word format:

Layers in pivot tables. By default, inclusion or exclusion of pivot table layers is controlled by the table properties for each pivot table. You can override this setting and include all layers or exclude all but the currently visible layer. [For more information, see Table Properties: Printing in Chapter 11 on p. 240.](#)

Wide Pivot Tables. Controls the treatment of tables that are too wide for the defined document width. By default, the table is wrapped to fit. The table is divided into sections, and row labels are repeated for each section of the table. Alternatively, you can shrink wide tables or make no changes to wide tables and allow them to extend beyond the defined document width.

Include footnotes and captions. Controls the inclusion or exclusion of all pivot table footnotes and captions.

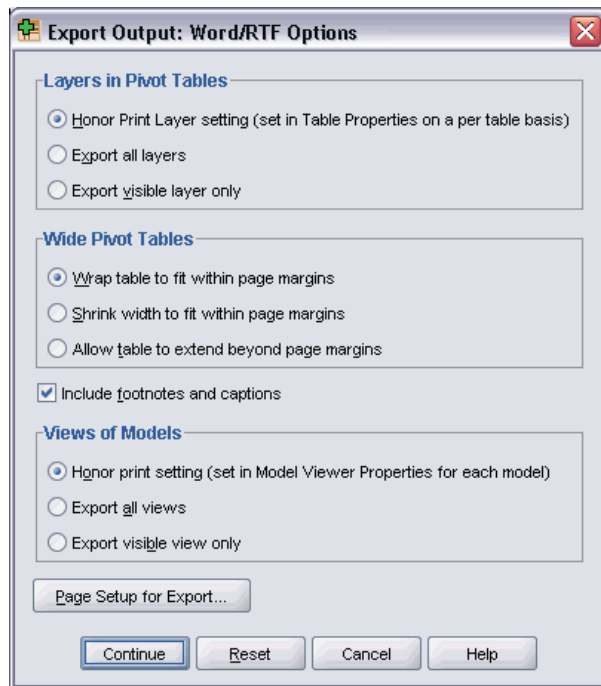
Views of Models. By default, inclusion or exclusion of model views is controlled by the model properties for each model. You can override this setting and include all views or exclude all but the currently visible view. [For more information, see Model Properties in Chapter 12 on p. 251.](#) (Note: all model views, including tables, are exported as graphics.)

Page Setup for Export. This opens a dialog where you can define the page size and margins for the exported document. The document width used to determine wrapping and shrinking behavior is the page width minus the left and right margins.

To Set Word Export Options

- ▶ Select Word/RTF as the export format.
- ▶ Click Change Options.

Figure 10-5
Word Export Output Options



Excel Options

The following options are available for exporting output in Excel format:

Create a worksheet or workbook or modify an existing worksheet. By default, a new workbook is created. If a file with the specified name already exists, it will be overwritten. If you select the option to create a worksheet, if a worksheet with the specified name already exists in the specified file, it will be overwritten. If you select the option to modify an existing worksheet, you must also specify the worksheet name. (This is optional for creating a worksheet.) Worksheet names cannot exceed 31 characters and cannot contain forward or back slashes, square brackets, question marks, or asterisks.

If you modify an existing worksheet, charts, model views, and tree diagrams are not included in the exported output.

Location in worksheet. Controls the location within the worksheet for the exported output. By default, exported output will be added after the last column that has any content, starting in the first row, without modifying any existing contents. This is a good choice for adding new columns to an existing worksheet. Adding exported output after the last row is a good choice for adding new rows to an existing worksheet. Adding exported output starting at a specific cell location will overwrite any existing content in the area where the exported output is added.

Layers in pivot tables. By default, inclusion or exclusion of pivot table layers is controlled by the table properties for each pivot table. You can override this setting and include all layers or exclude all but the currently visible layer. [For more information, see Table Properties: Printing in Chapter 11 on p. 240.](#)

Include footnotes and captions. Controls the inclusion or exclusion of all pivot table footnotes and captions.

Views of Models. By default, inclusion or exclusion of model views is controlled by the model properties for each model. You can override this setting and include all views or exclude all but the currently visible view. [For more information, see Model Properties in Chapter 12 on p. 251.](#) (Note: all model views, including tables, are exported as graphics.)

To Set Excel Export Options

- ▶ Select Excel as the export format.
- ▶ Click Change Options.

Figure 10-6
Excel Export Output Options

Export Output: Excel Options

What Do You Want to Do?

Create a workbook
If you specify the name of an existing workbook, it will be completely overwritten.

Create a worksheet
If you specify the name of an existing worksheet, it will be completely overwritten.

Modify an existing worksheet

Worksheet name:

Location in Worksheet

After last column

After last row

Starting in a specific cell
Starting cell:

Existing data at the specified location will be overwritten.

Layers in Pivot Tables

Honor Print Layer setting (set in Table Properties on a per table basis)

Export all layers

Export visible layer only

Include footnotes and captions

Views of Models

Honor print setting (set in Model Viewer Properties for each model)

Export all views

Export visible view only

Continue Reset Cancel Help

PowerPoint Options

The following options are available for PowerPoint:

Layers in pivot tables. By default, inclusion or exclusion of pivot table layers is controlled by the table properties for each pivot table. You can override this setting and include all layers or exclude all but the currently visible layer. [For more information, see Table Properties: Printing in Chapter 11 on p. 240.](#)

Wide Pivot Tables. Controls the treatment of tables that are too wide for the defined document width. By default, the table is wrapped to fit. The table is divided into sections, and row labels are repeated for each section of the table. Alternatively, you can shrink wide tables or make no changes to wide tables and allow them to extend beyond the defined document width.

Include footnotes and captions. Controls the inclusion or exclusion of all pivot table footnotes and captions.

Use Viewer outline entries as slide titles. Includes a title on each slide that is created by the export. Each slide contains a single item that is exported from the Viewer. The title is formed from the outline entry for the item in the outline pane of the Viewer.

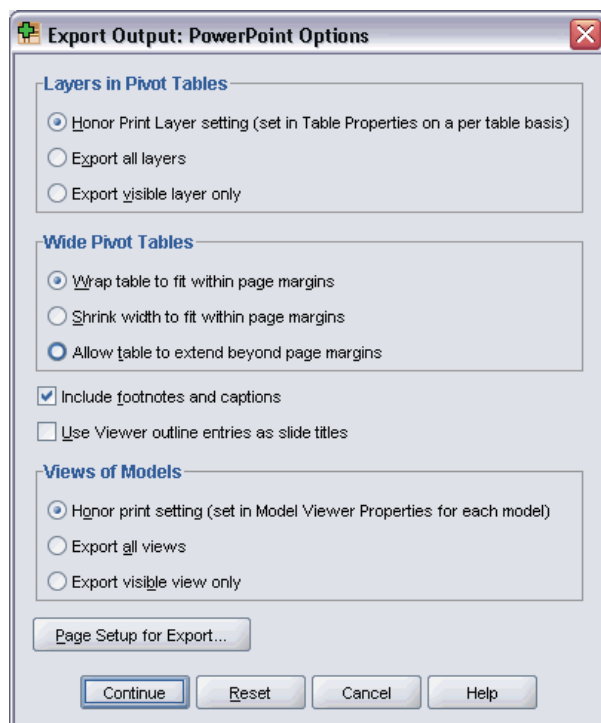
Views of Models. By default, inclusion or exclusion of model views is controlled by the model properties for each model. You can override this setting and include all views or exclude all but the currently visible view. [For more information, see Model Properties in Chapter 12 on p. 251.](#) (Note: all model views, including tables, are exported as graphics.)

Page Setup for Export. This opens a dialog where you can define the page size and margins for the exported document. The document width used to determine wrapping and shrinking behavior is the page width minus the left and right margins.

To Set PowerPoint Export Options

- ▶ Select PowerPoint as the export format.
- ▶ Click Change Options.

Figure 10-7
Powerpoint Export Options dialog



Note: Export to PowerPoint is available only on Windows operating systems.

PDF Options

The following options are available for PDF:

Embed bookmarks. This option includes bookmarks in the PDF document that correspond to the Viewer outline entries. Like the Viewer outline pane, bookmarks can make it much easier to navigate documents with a large number of output objects.

Embed fonts. Embedding fonts ensures that the PDF document will look the same on all computers. Otherwise, if some fonts used in the document are not available on the computer being used to view (or print) the PDF document, font substitution may yield suboptimal results.

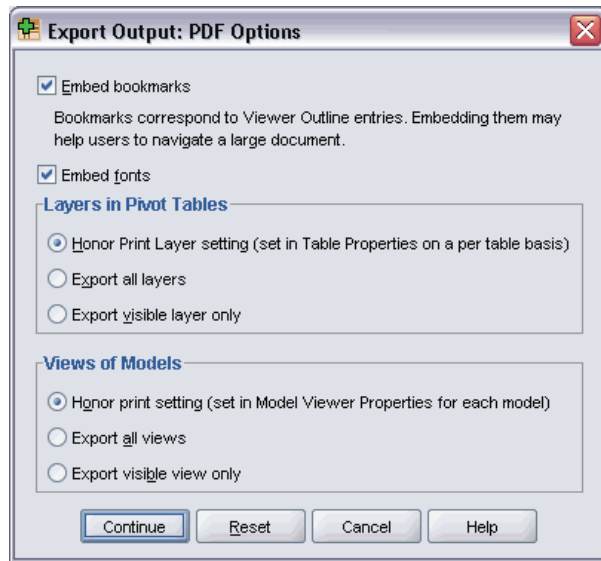
Layers in pivot tables. By default, inclusion or exclusion of pivot table layers is controlled by the table properties for each pivot table. You can override this setting and include all layers or exclude all but the currently visible layer. [For more information, see Table Properties: Printing in Chapter 11 on p. 240.](#)

Views of Models. By default, inclusion or exclusion of model views is controlled by the model properties for each model. You can override this setting and include all views or exclude all but the currently visible view. [For more information, see Model Properties in Chapter 12 on p. 251.](#) (Note: all model views, including tables, are exported as graphics.)

To Set PDF Export Options

- ▶ Select Portable Document Format as the export format.
- ▶ Click Change Options.

Figure 10-8
PDF Options dialog box



Other Settings That Affect PDF Output

Page Setup/Page Attributes. Page size, orientation, margins, content and display of page headers and footers, and printed chart size in PDF documents are controlled by page setup and page attribute options.

Table Properties/TableLooks. Scaling of wide and/or long tables and printing of table layers are controlled by table properties for each table. These properties can also be saved in TableLooks. For more information, see [Table Properties: Printing in Chapter 11 on p. 240](#).

Default/Current Printer. The resolution (DPI) of the PDF document is the current resolution setting for the default or currently selected printer (which can be changed using Page Setup). The maximum resolution is 1200 DPI. If the printer setting is higher, the PDF document resolution will be 1200 DPI.

Note: High-resolution documents may yield poor results when printed on lower-resolution printers.

Text Options

The following options are available for text export:

Pivot Table Format. Pivot tables can be exported in tab-separated or space-separated format. For space-separated format, you can also control:

- **Column Width.** Autofit does not wrap any column contents, and each column is as wide as the widest label or value in that column. Custom sets a maximum column width that is applied to all columns in the table, and values that exceed that width wrap onto the next line in that column.
- **Row/Column Border Character.** Controls the characters used to create row and column borders. To suppress display of row and column borders, enter blank spaces for the values.

Layers in pivot tables. By default, inclusion or exclusion of pivot table layers is controlled by the table properties for each pivot table. You can override this setting and include all layers or exclude all but the currently visible layer. [For more information, see Table Properties: Printing in Chapter 11 on p. 240.](#)

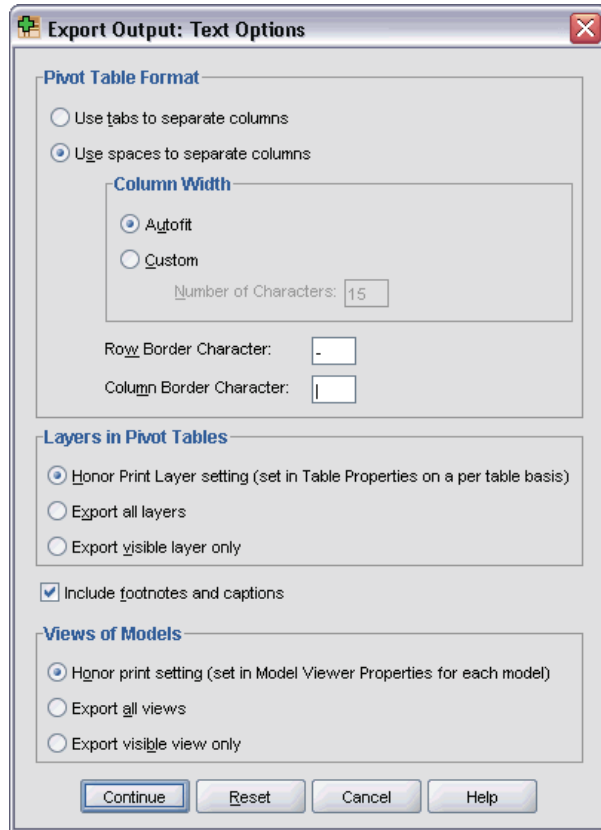
Include footnotes and captions. Controls the inclusion or exclusion of all pivot table footnotes and captions.

Views of Models. By default, inclusion or exclusion of model views is controlled by the model properties for each model. You can override this setting and include all views or exclude all but the currently visible view. [For more information, see Model Properties in Chapter 12 on p. 251.](#) (Note: all model views, including tables, are exported as graphics.)

To Set Text Export Options

- ▶ Select Text as the export format.
- ▶ Click Change Options.

Figure 10-9
Text Options dialog box

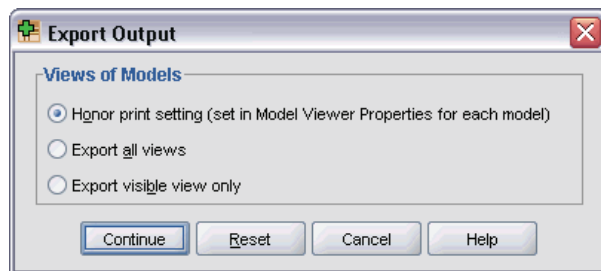


Graphics Only Options

The following options are available for exporting graphics only:

Views of Models. By default, inclusion or exclusion of model views is controlled by the model properties for each model. You can override this setting and include all views or exclude all but the currently visible view. [For more information, see Model Properties in Chapter 12 on p. 251.](#) (Note: all model views, including tables, are exported as graphics.)

Figure 10-10
Graphics Only Options



Graphics Format Options

For HTML and text documents and for exporting charts only, you can select the graphic format, and for each graphic format you can control various optional settings.

To select the graphic format and options for exported charts:

- ▶ Select HTML, Text, or None (Graphics only) as the document type.
- ▶ Select the graphic file format from the drop-down list.
- ▶ Click Change Options to change the options for the selected graphic file format.

JPEG Chart Export Options

- **Image size.** Percentage of original chart size, up to 200 percent.
- **Convert to grayscale.** Converts colors to shades of gray.

BMP Chart Export Options

- **Image size.** Percentage of original chart size, up to 200 percent.
- **Compress image to reduce file size.** A lossless compression technique that creates smaller files without affecting image quality.

PNG Chart Export Options

Image size. Percentage of original chart size, up to 200 percent.

Color Depth. Determines the number of colors in the exported chart. A chart that is saved under any depth will have a minimum of the number of colors that are actually used and a maximum of the number of colors that are allowed by the depth. For example, if the chart contains three colors—red, white, and black—and you save it as 16 colors, the chart will remain as three colors.

- If the number of colors in the chart exceeds the number of colors for that depth, the colors will be dithered to replicate the colors in the chart.
- Current screen depth is the number of colors currently displayed on your computer monitor.

EMF and TIFF Chart Export Options

Image size. Percentage of original chart size, up to 200 percent.

Note: EMF (enhanced metafile) format is available only on Windows operating systems.

EPS Chart Export Options

Image size. You can specify the size as a percentage of the original image size (up to 200 percent), or you can specify an image width in pixels (with height determined by the width value and the aspect ratio). The exported image is always proportional to the original.

Include TIFF preview image. Saves a preview with the EPS image in TIFF format for display in applications that cannot display EPS images on screen.

Fonts. Controls the treatment of fonts in EPS images.

- **Use font references.** If the fonts that are used in the chart are available on the output device, the fonts are used. Otherwise, the output device uses alternate fonts.
- **Replace fonts with curves.** Turns fonts into PostScript curve data. The text itself is no longer editable as text in applications that can edit EPS graphics. This option is useful if the fonts that are used in the chart are not available on the output device.

Viewer Printing

There are two options for printing the contents of the Viewer window:

All visible output. Prints only items that are currently displayed in the contents pane. Hidden items (items with a closed book icon in the outline pane or hidden in collapsed outline layers) are not printed.

Selection. Prints only items that are currently selected in the outline and/or contents panes.

To Print Output and Charts

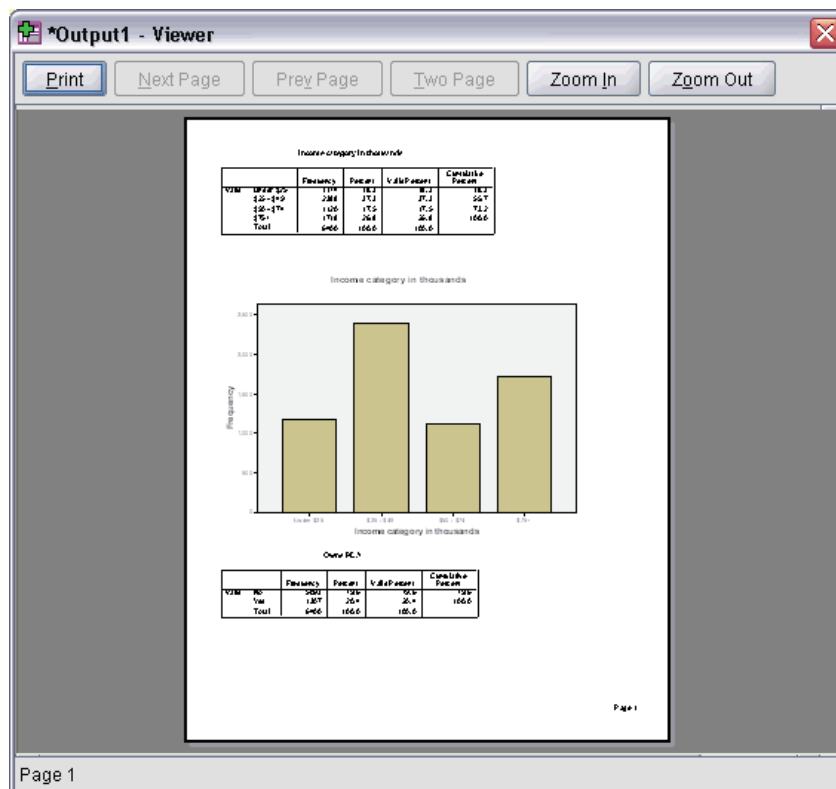
- ▶ Make the Viewer the active window (click anywhere in the window).
- ▶ From the menus choose:
 - File
 - Print...
- ▶ Select the print settings that you want.
- ▶ Click OK to print.

Print Preview

Print Preview shows you what will print on each page for Viewer documents. It is a good idea to check Print Preview before actually printing a Viewer document, because Print Preview shows you items that may not be visible by looking at the contents pane of the Viewer, including:

- Page breaks
- Hidden layers of pivot tables
- Breaks in wide tables
- Headers and footers that are printed on each page

Figure 10-11
Print Preview



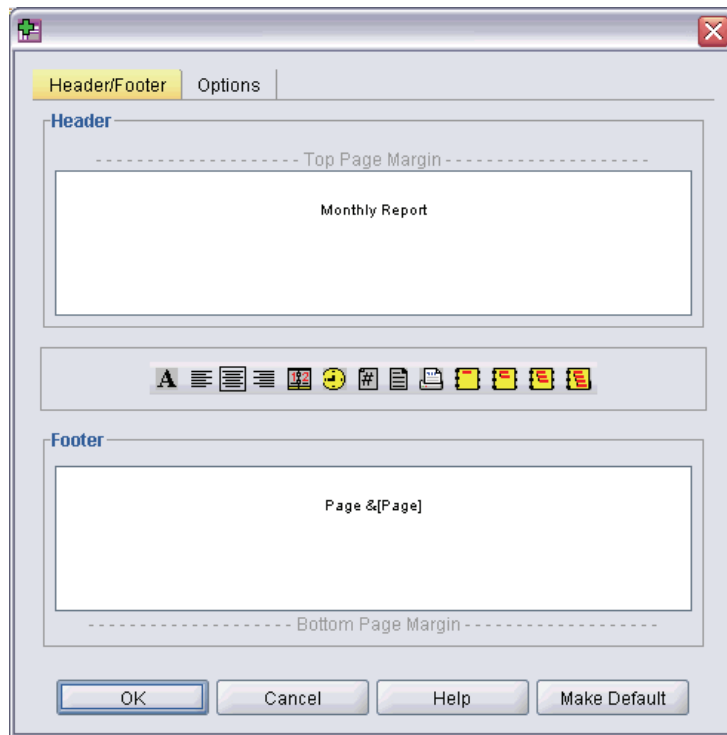
If any output is currently selected in the Viewer, the preview displays only the selected output. To view a preview for all output, make sure nothing is selected in the Viewer.

Page Attributes: Headers and Footers

Headers and footers are the information that is printed at the top and bottom of each page. You can enter any text that you want to use as headers and footers. You can also use the toolbar in the middle of the dialog box to insert:

- Date and time
- Page numbers
- Viewer filename
- Outline heading labels
- Page titles and subtitles

Figure 10-12
Page Attributes dialog box, Header/Footer tab



- Make Default uses the settings specified here as the default settings for new Viewer documents. (Note: this makes the current settings on both the Header/Footer tab and the Options tab the default settings.)
- Outline heading labels indicate the first-, second-, third-, and/or fourth-level outline heading for the first item on each page.
- Page titles and subtitles print the current page titles and subtitles. These can be created with New Page Title on the Viewer Insert menu or with the TITLE and SUBTITLE commands. If you have not specified any page titles or subtitles, this setting is ignored.

Note: Font characteristics for new page titles and subtitles are controlled on the Viewer tab of the Options dialog box (accessed by choosing Options on the Edit menu). Font characteristics for existing page titles and subtitles can be changed by editing the titles in the Viewer.

To see how your headers and footers will look on the printed page, choose Print Preview from the File menu.

To Insert Page Headers and Footers

- ▶ Make the Viewer the active window (click anywhere in the window).
- ▶ From the menus choose:
 - File
 - Page Attributes...
- ▶ Click the Header/Footer tab.

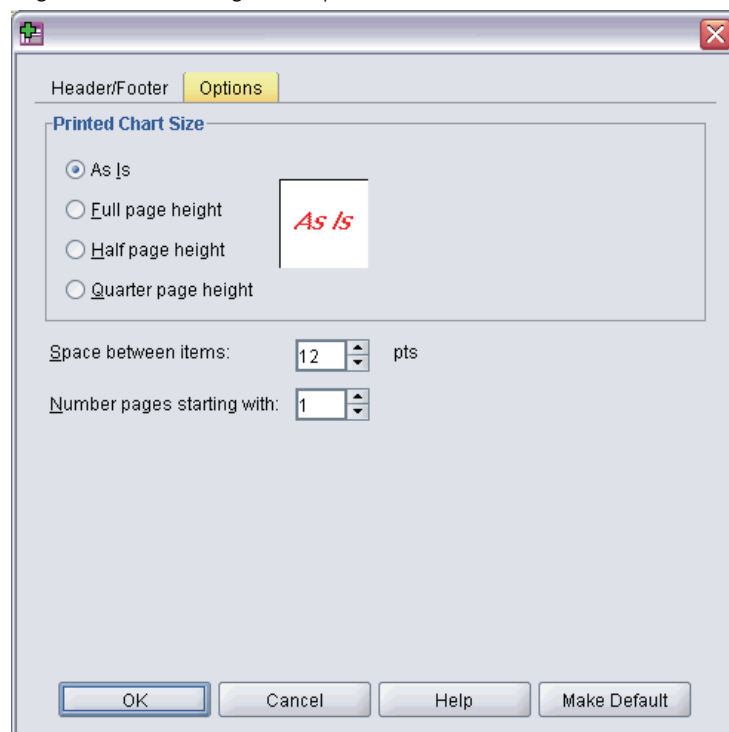
- ▶ Enter the header and/or footer that you want to appear on each page.

Page Attributes: Options

This dialog box controls the printed chart size, the space between printed output items, and page numbering.

- **Printed Chart Size.** Controls the size of the printed chart relative to the defined page size. The chart's aspect ratio (width-to-height ratio) is not affected by the printed chart size. The overall printed size of a chart is limited by both its height and width. When the outer borders of a chart reach the left and right borders of the page, the chart size cannot increase further to fill additional page height.
- **Space between items.** Controls the space between printed items. Each pivot table, chart, and text object is a separate item. This setting does not affect the display of items in the Viewer.
- **Number pages starting with.** Numbers pages sequentially, starting with the specified number.
- **Make Default.** This option uses the settings specified here as the default settings for new Viewer documents. (Note: this makes the current settings on both the Header/Footer tab and the Options tab the default settings.)

Figure 10-13
Page Attributes dialog box, Options tab



To Change Printed Chart Size, Page Numbering, and Space between Printed Items

- ▶ Make the Viewer the active window (click anywhere in the window).

- ▶ From the menus choose:
 - File
 - Page Attributes...
- ▶ Click the Options tab.
- ▶ Change the settings and click OK.

Saving Output

The contents of the Viewer can be saved to a Viewer document. The saved document includes both panes of the Viewer window (the outline and the contents).

To Save a Viewer Document

- ▶ From the Viewer window menus choose:
 - File
 - Save
- ▶ Enter the name of the document, and then click Save.

Optionally, you can lock files to prevent editing in Smartreader (a separate product for working with Viewer documents). If a Viewer document is locked, you can manipulate pivot tables (swap rows and columns, change the displayed layer, etc.) but you cannot edit any output or save any changes to the Viewer document in Smartreader. This setting has no effect on Viewer documents opened in SPSS Statistics.

To save results in external formats (for example, HTML or text), use Export on the File menu.

Pivot Tables

Many results are presented in tables that can be pivoted interactively. That is, you can rearrange the rows, columns, and layers.

Manipulating a Pivot Table

Options for manipulating a pivot table include:

- Transposing rows and columns
- Moving rows and columns
- Creating multidimensional layers
- Grouping and ungrouping rows and columns
- Showing and hiding rows, columns, and other information
- Rotating row and column labels
- Finding definitions of terms

Activating a Pivot Table

Before you can manipulate or modify a pivot table, you need to **activate** the table. To activate a table:

- ▶ Double-click the table.

or

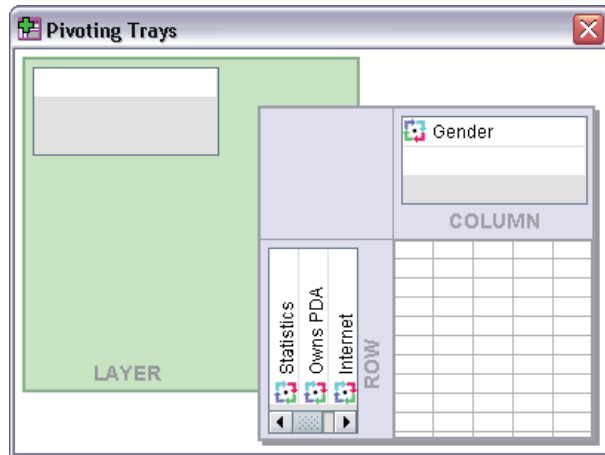
- ▶ Right-click the table and from the context menu choose Edit Content.
- ▶ From the sub-menu choose either In Viewer or In Separate Window.
 - By default, activating the table by double-clicking will activate all but very large tables in the Viewer window. [For more information, see Pivot Table Options in Chapter 48 on p. 505.](#)
 - If you want to have more than one pivot table activated at the same time, you must activate the tables in separate windows.

Pivoting a Table

- ▶ Activate the pivot table.

- ▶ From the menus choose:
 - Pivot
 - Pivoting Trays

Figure 11-1
Pivoting trays



A table has three dimensions: rows, columns, and layers. A dimension can contain multiple elements (or none at all). You can change the organization of the table by moving elements between or within dimensions. To move an element, just drag and drop it where you want it.

Changing Display Order of Elements within a Dimension

To change the display order of elements within a table dimension (row, column, or layer):

- ▶ If pivoting trays are not already on, from the Pivot Table menu choose:
 - Pivot
 - Pivoting Trays
- ▶ Drag and drop the elements within the dimension in the pivoting tray.

Moving Rows and Columns within a Dimension Element

- ▶ In the table itself (not the pivoting trays), click the label for the row or column you want to move.
- ▶ Drag the label to the new position.
- ▶ From the context menu choose Insert Before or Swap.

Note: Make sure that Drag to Copy on the Edit menu is *not* enabled (checked). If Drag to Copy is enabled, deselect it.

Transposing Rows and Columns

If you just want to flip the rows and columns, there's a simple alternative to using the pivoting trays:

- ▶ From the menus choose:
Pivot
Transpose Rows and Columns

This has the same effect as dragging all of the row elements into the column dimension and dragging all of the column elements into the row dimension.

Grouping Rows or Columns

- ▶ Select the labels for the rows or columns that you want to group together (click and drag or Shift+click to select multiple labels).
- ▶ From the menus choose:
Edit
Group

A group label is automatically inserted. Double-click the group label to edit the label text.

Figure 11-2
Row and column groups and labels

		Column Group Label		Total
		Female	Male	
Row Group Label	Clerical	206	157	363
	Custodial		27	27
	Manager	10	74	84

Note: To add rows or columns to an existing group, you must first ungroup the items that are currently in the group. Then you can create a new group that includes the additional items.

Ungrouping Rows or Columns

- ▶ Click anywhere in the group label for the rows or columns that you want to ungroup.
- ▶ From the menus choose:
Edit
Ungroup

Ungrouping automatically deletes the group label.

Rotating Row or Column Labels

You can rotate labels between horizontal and vertical display for the innermost column labels and the outermost row labels in a table.

- ▶ From the menus choose:
Format
Rotate Inner Column Labels

or

Format
Rotate Outer Row Labels

Figure 11-3
Rotated column labels

	Frequency	Percent	Valid Percent	Cumulative Percent
Did not complete high school	1390	21.7	21.7	21.7
High school degree	1936	30.3	30.3	52.0
Some college				
College degree				
Post-undergraduate degree				
Total				

	Frequency	Percent	Valid Percent	Cumulative Percent
Did not complete high school	1390	21.7	21.7	21.7
High school degree	1936	30.3	30.3	52.0
Some college	1360	21.3	21.3	73.2
College degree	1355	21.2	21.2	94.4
Post-undergraduate degree	359	5.6	5.6	100.0
Total	6400	100.0	100.0	

Only the innermost column labels and the outermost row labels can be rotated.

Working with Layers

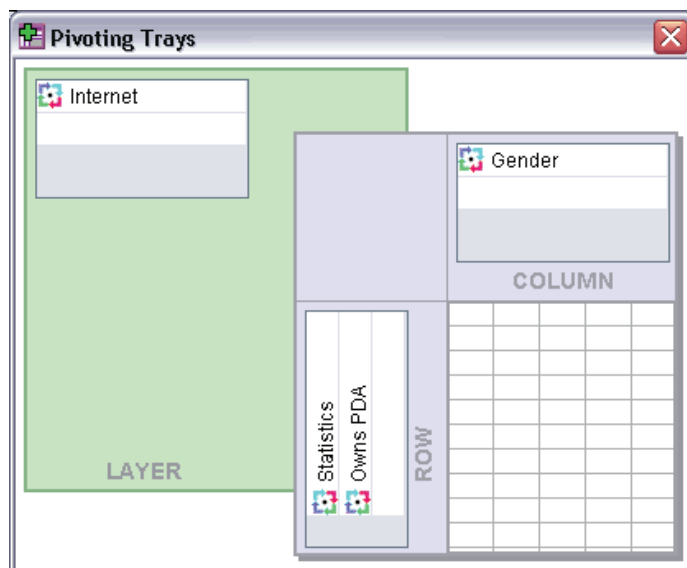
You can display a separate two-dimensional table for each category or combination of categories. The table can be thought of as stacked in layers, with only the top layer visible.

Creating and Displaying Layers

To create layers:

- ▶ Activate the pivot table.
- ▶ If pivoting trays are not already on, from the Pivot Table menu choose:
 - Pivot
 - Pivoting Trays
- ▶ Drag an element from the row or column dimension into the layer dimension.

Figure 11-4
Moving categories into layers



Moving elements to the layer dimension creates a multidimensional table, but only a single two-dimensional “slice” is displayed. The visible table is the table for the top layer. For example, if a yes/no categorical variable is in the layer dimension, then the multidimensional table has two layers: one for the *yes* category and one for the *no* category.

Figure 11-5
Categories in separate layers

Internet: Yes			Gender		Total
Owns	No	Count	Female	Male	
PDA		Expected Count	512.5	512.5	1025.0

Internet: No			Gender		Total
Owns	No	Count	Female	Male	
PDA		Expected Count	1903.4	1955.6	3859.0
		% within Owns PDA	49.2%	50.8%	100.0%
	Yes	Count	327	323	650
		Expected Count	320.6	329.4	650.0
		% within Owns PDA	50.3%	49.7%	100.0%
Total		Count	2224	2285	4509
		Expected Count	2224.0	2285.0	4509.0
		% within Owns PDA	49.3%	50.7%	100.0%

Changing the Displayed Layer

- Choose a category from the drop-down list of layers (in the pivot table itself, not the pivoting tray).

Figure 11-6
Selecting layers from drop-down lists

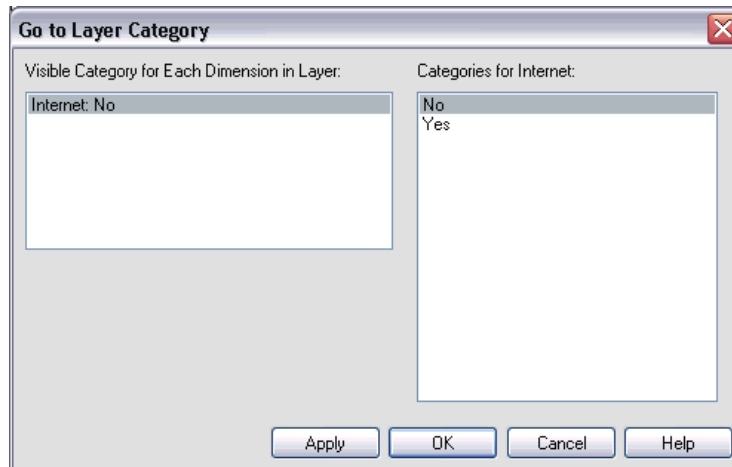
Internet		No			
		No	Gender		Total
		Yes	Female	Male	
Owns PDA	No	Count	1897	1962	3859
		Expected Count	1903.4	1955.6	3859.0
		% within Owns PDA	49.2%	50.8%	100.0%
	Yes	Count	327	323	650
		Expected Count	320.6	329.4	650.0
		% within Owns PDA	50.3%	49.7%	100.0%
Total	Count	2224	2285	4509	
	Expected Count	2224.0	2285.0	4509.0	
	% within Owns PDA	49.3%	50.7%	100.0%	

Go to Layer Category

Go to Layer Category allows you to change layers in a pivot table. This dialog box is particularly useful when there are many layers or the selected layer has many categories.

- ▶ From the menus choose:
Pivot
Go to Layer...

Figure 11-7
Go to Layer Category dialog box



- ▶ In the Visible Category list, select a layer dimension. The Categories list will display all categories for the selected dimension.
- ▶ In the Categories list, select the category that you want, and then click OK or Apply.

Showing and Hiding Items

Many types of cells can be hidden, including:

- Dimension labels

- Categories, including the label cell and data cells in a row or column
- Category labels (without hiding the data cells)
- Footnotes, titles, and captions

Hiding Rows and Columns in a Table

- ▶ Right-click the category label for the row or column you want to hide.
 - ▶ From the context menu choose:
 - Select
 - Data and Label Cells
 - ▶ Right-click the category label again and from the context menu choose Hide Category.
- or*
- ▶ From the View menu choose Hide.

Showing Hidden Rows and Columns in a Table

- ▶ Right-click another row or column label in the same dimension as the hidden row or column.
- ▶ From the context menu choose:
 - Select
 - Data and Label Cells
- ▶ From the menus choose:
 - View
 - Show All Categories in [dimension name]

or

- ▶ To display all hidden rows and columns in an activated pivot table, from the menus choose:
 - View
 - Show All

This displays all hidden rows and columns in the table. (If Hide empty rows and columns is selected in Table Properties for this table, a completely empty row or column remains hidden.)

Hiding and Showing Dimension Labels

- ▶ Select the dimension label or any category label within the dimension.
- ▶ From the View menu or the context menu choose Hide Dimension Label or Show Dimension Label.

Hiding and Showing Table Titles

To hide a title:

- ▶ Select the title.
- ▶ From the View menu choose Hide.

To show hidden titles:

- ▶ From the View menu choose Show All.

TableLooks

A TableLook is a set of properties that define the appearance of a table. You can select a previously defined TableLook or create your own TableLook.

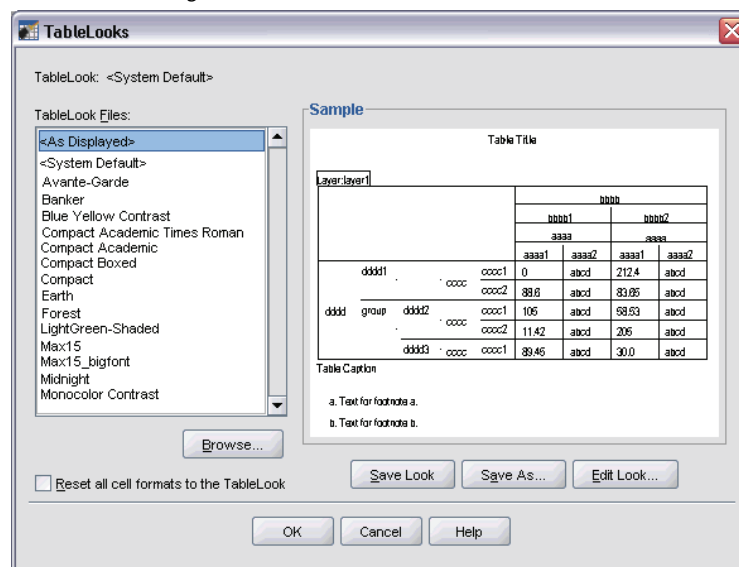
- Before or after a TableLook is applied, you can change cell formats for individual cells or groups of cells by using cell properties. The edited cell formats will remain intact, even when you apply a new TableLook. [For more information, see Cell Properties on p. 241.](#)
- Optionally, you can reset all cells to the cell formats that are defined by the current TableLook. This resets any cells that have been edited. If As Displayed is selected in the TableLook Files list, any edited cells are reset to the current table properties.

Note: TableLooks created in earlier versions of SPSS Statistics cannot be used in version 16.0 or later.

To Apply or Save a TableLook

- ▶ Activate a pivot table.
- ▶ From the menus choose:
Format
TableLooks...

Figure 11-8
TableLooks dialog box



- ▶ Select a TableLook from the list of files. To select a file from another directory, click Browse.
- ▶ Click OK to apply the TableLook to the selected pivot table.

To Edit or Create a TableLook

- ▶ In the TableLooks dialog box, select a TableLook from the list of files.
- ▶ Click Edit Look.
- ▶ Adjust the table properties for the attributes that you want, and then click OK.
- ▶ Click Save Look to save the edited TableLook, or click Save As to save it as a new TableLook.

Editing a TableLook affects only the selected pivot table. An edited TableLook is not applied to any other tables that uses that TableLook unless you select those tables and reapply the TableLook.

Table Properties

Table Properties allows you to set general properties of a table, set cell styles for various parts of a table, and save a set of those properties as a TableLook. You can:

- Control general properties, such as hiding empty rows or columns and adjusting printing properties.
- Control the format and position of footnote markers.
- Determine specific formats for cells in the data area, for row and column labels, and for other areas of the table.
- Control the width and color of the lines that form the borders of each area of the table.

To Change Pivot Table Properties

- ▶ Activate the pivot table.
- ▶ From the menus choose:
 - Format
 - Table Properties...
- ▶ Select a tab (General, Footnotes, Cell Formats, Borders, or Printing).
- ▶ Select the options that you want.
- ▶ Click OK or Apply.

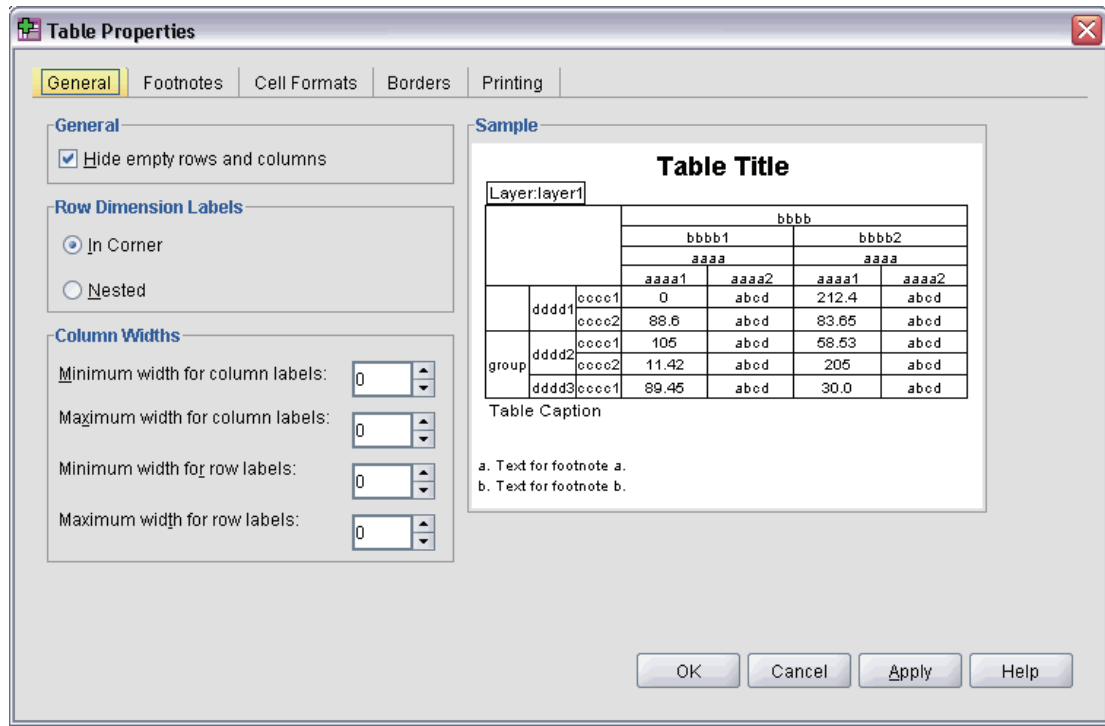
The new properties are applied to the selected pivot table. To apply new table properties to a TableLook instead of just the selected table, edit the TableLook (Format menu, TableLooks).

Table Properties: General

Several properties apply to the table as a whole. You can:

- Show or hide empty rows and columns. (An empty row or column has nothing in any of the data cells.)
- Control the placement of row labels, which can be in the upper left corner or nested.
- Control maximum and minimum column width (expressed in points).

Figure 11-9
Table Properties dialog box, General tab



To change general table properties:

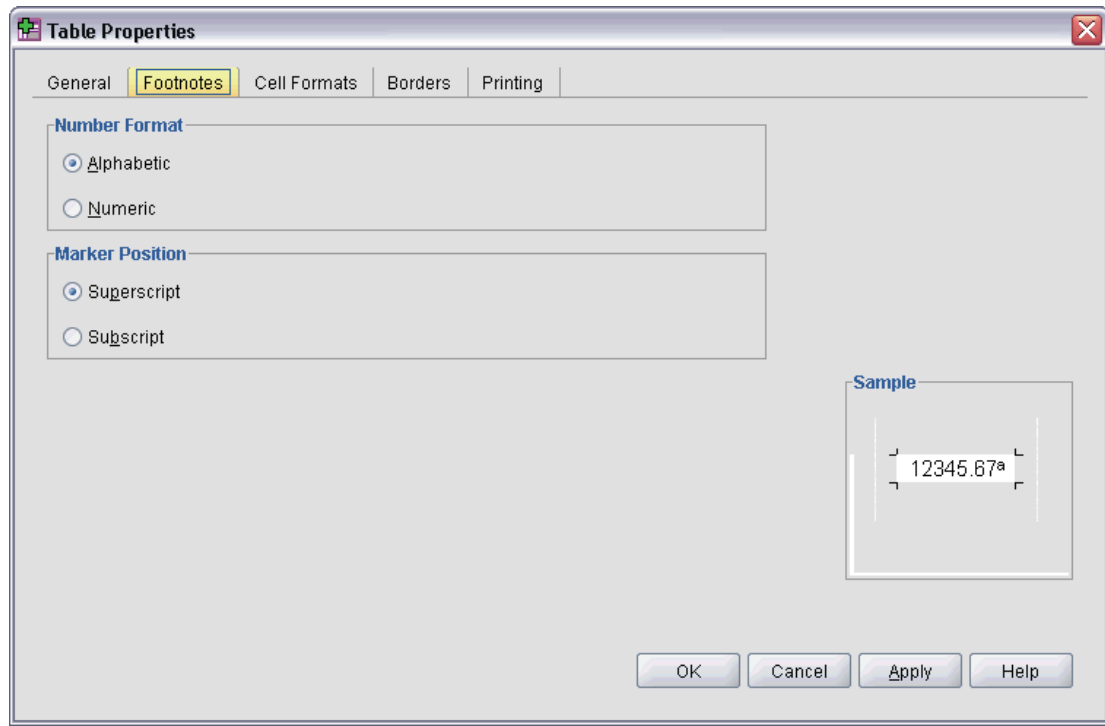
- ▶ Click the General tab.
- ▶ Select the options that you want.
- ▶ Click OK or Apply.

Table Properties: Footnotes

The properties of footnote markers include style and position in relation to text.

- The style of footnote markers is either numbers (1, 2, 3, ...) or letters (a, b, c, ...).
- The footnote markers can be attached to text as superscripts or subscripts.

Figure 11-10
Table Properties dialog box, Footnotes tab



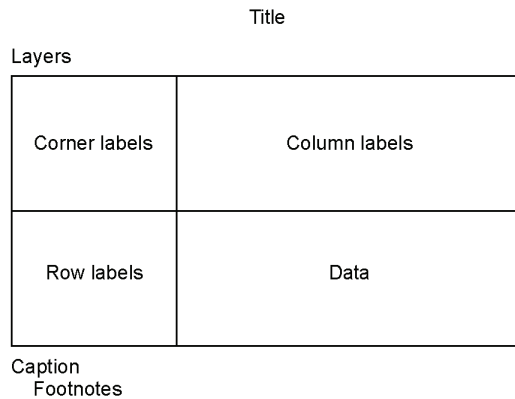
To change footnote marker properties:

- ▶ Click the Footnotes tab.
- ▶ Select a footnote number format.
- ▶ Select a marker position.
- ▶ Click OK or Apply.

Table Properties: Cell Formats

For formatting, a table is divided into areas: title, layers, corner labels, row labels, column labels, data, caption, and footnotes. For each area of a table, you can modify the associated cell formats. Cell formats include text characteristics (such as font, size, color, and style), horizontal and vertical alignment, background colors, and inner cell margins.

Figure 11-11
Areas of a table

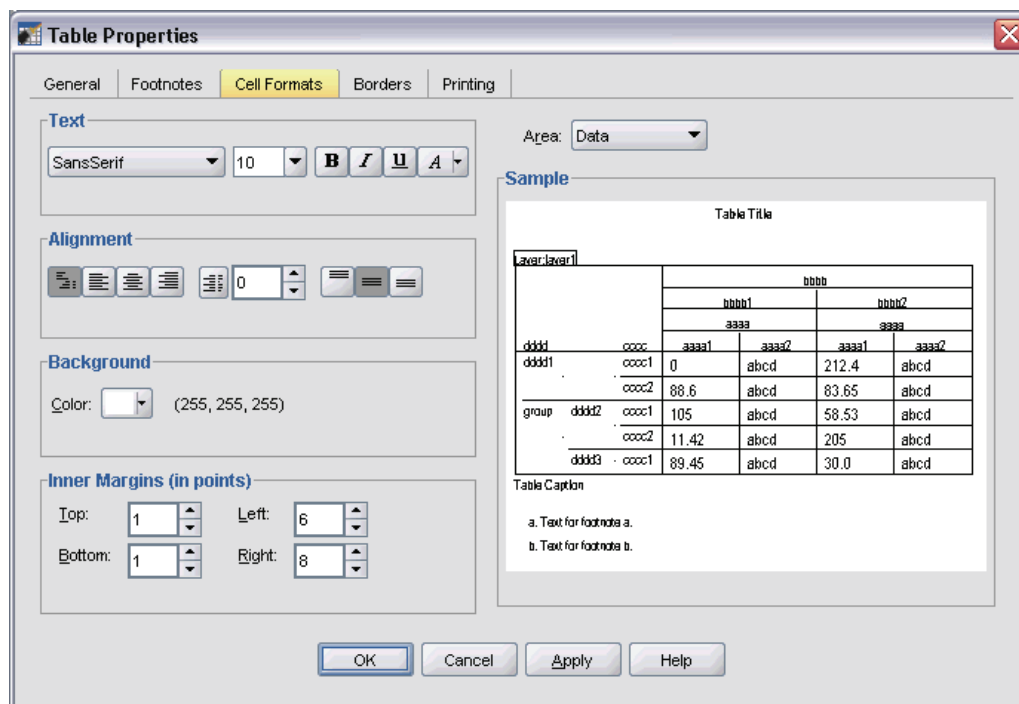


Cell formats are applied to areas (categories of information). They are not characteristics of individual cells. This distinction is an important consideration when pivoting a table.

For example,

- If you specify a bold font as a cell format of column labels, the column labels will appear bold no matter what information is currently displayed in the column dimension. If you move an item from the column dimension to another dimension, it does not retain the bold characteristic of the column labels.
- If you make column labels bold simply by highlighting the cells in an activated pivot table and clicking the Bold button on the toolbar, the contents of those cells will remain bold no matter what dimension you move them to, and the column labels will not retain the bold characteristic for other items moved into the column dimension.

Figure 11-12
Table Properties dialog box, Cell Formats tab



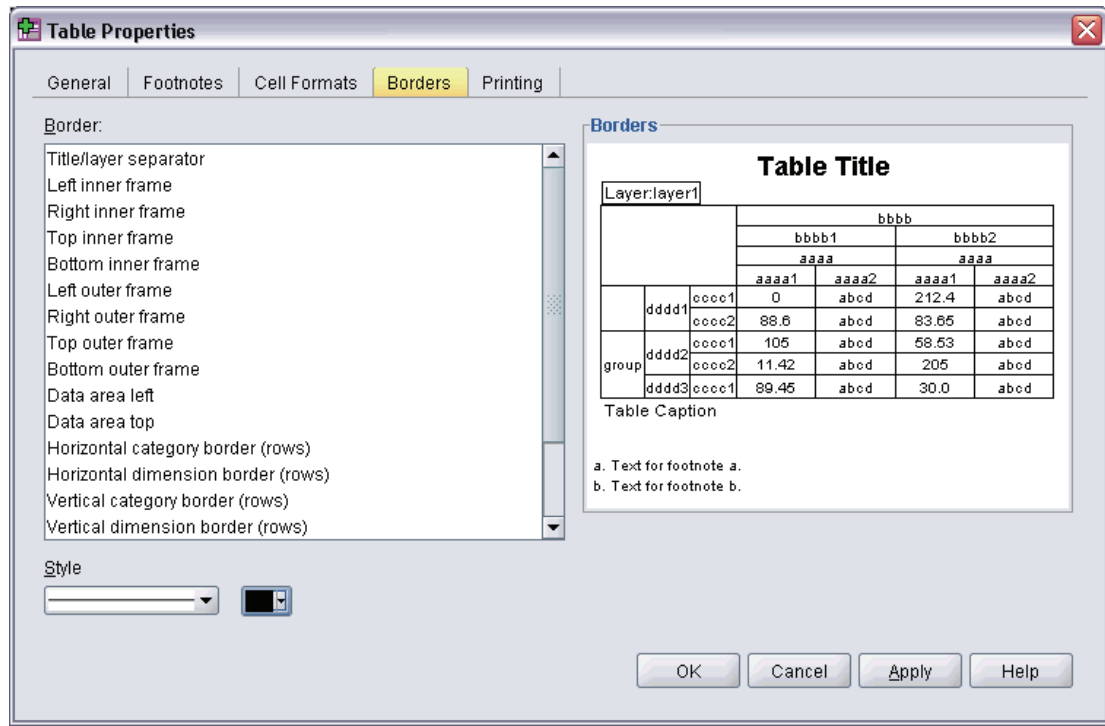
To change cell formats:

- ▶ Select the Cell Formats tab.
- ▶ Select an area from the drop-down list or click an area of the sample.
- ▶ Select characteristics for the area. Your selections are reflected in the sample.
- ▶ Click OK or Apply.

Table Properties: Borders

For each border location in a table, you can select a line style and a color. If you select None as the style, there will be no line at the selected location.

Figure 11-13
Table Properties dialog box, Borders tab



To change table borders:

- ▶ Click the Borders tab.
- ▶ Select a border location, either by clicking its name in the list or by clicking a line in the Sample area.
- ▶ Select a line style or select None.
- ▶ Select a color.
- ▶ Click OK or Apply.

Table Properties: Printing

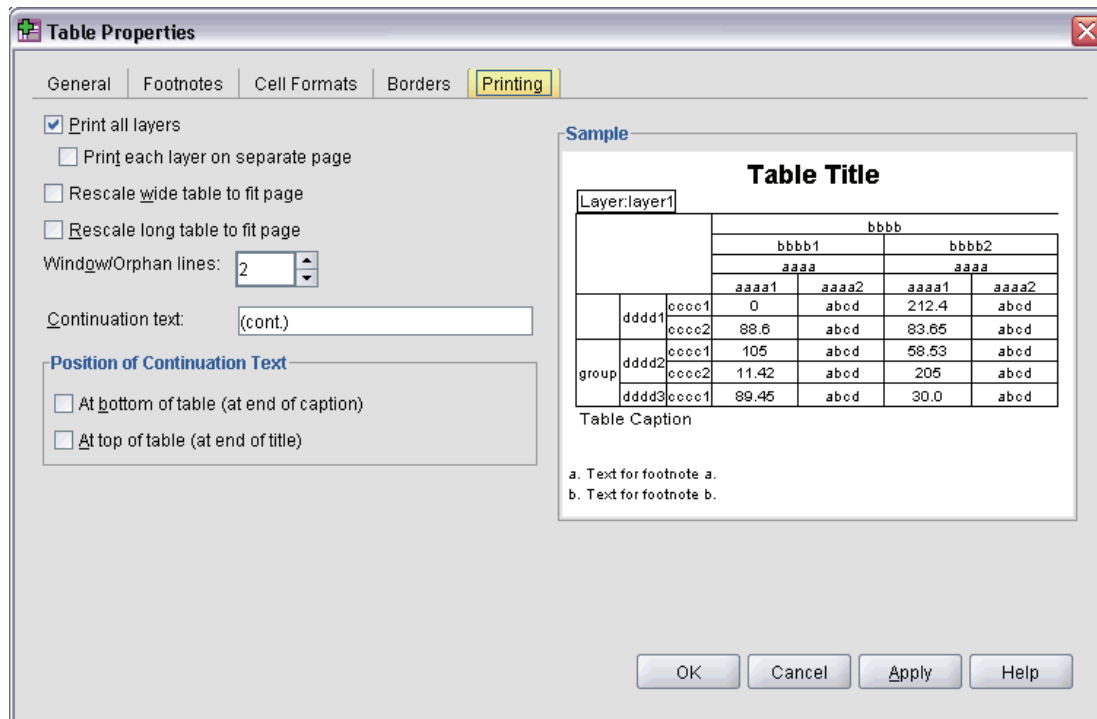
You can control the following properties for printed pivot tables:

- Print all layers or only the top layer of the table, and print each layer on a separate page.
- Shrink a table horizontally or vertically to fit the page for printing.
- Control widow/orphan lines by controlling the minimum number of rows and columns that will be contained in any printed section of a table if the table is too wide and/or too long for the defined page size.

Note: If a table is too long to fit on the current page because there is other output above it, but it will fit within the defined page length, the table is automatically printed on a new page, regardless of the widow/orphan setting.

- Include continuation text for tables that don't fit on a single page. You can display continuation text at the bottom of each page and at the top of each page. If neither option is selected, the continuation text will not be displayed.

Figure 11-14
Table Properties dialog box, Printing tab



To control pivot table printing properties:

- ▶ Click the Printing tab.
- ▶ Select the printing options that you want.
- ▶ Click OK or Apply.

Cell Properties

Cell properties are applied to a selected cell. You can change the font, value format, alignment, margins, and colors. Cell properties override table properties; therefore, if you change table properties, you do not change any individually applied cell properties.

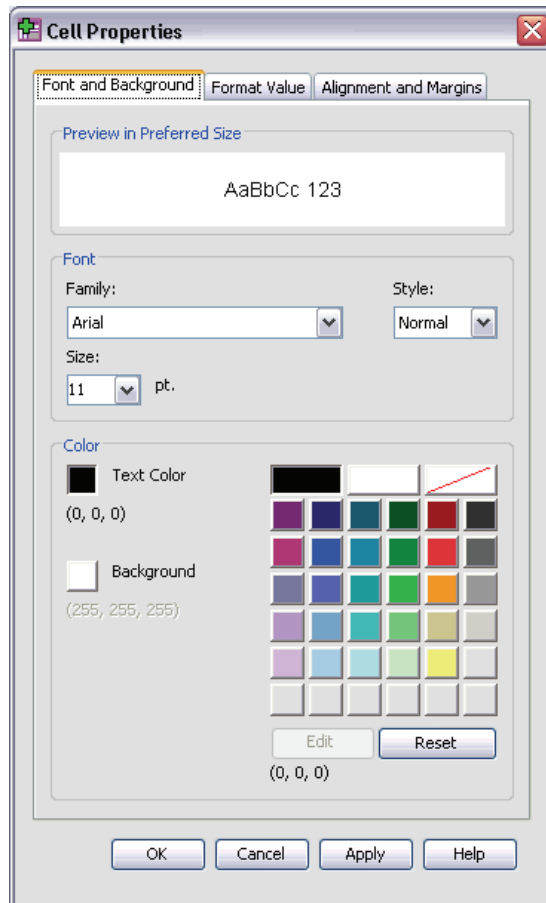
To change cell properties:

- ▶ Activate a table and select the cell(s) in the table.
- ▶ From the Format menu or the context menu choose Cell Properties.

Font and Background

The Font and Background tab controls the font style and color and background color for the selected cells in the table.

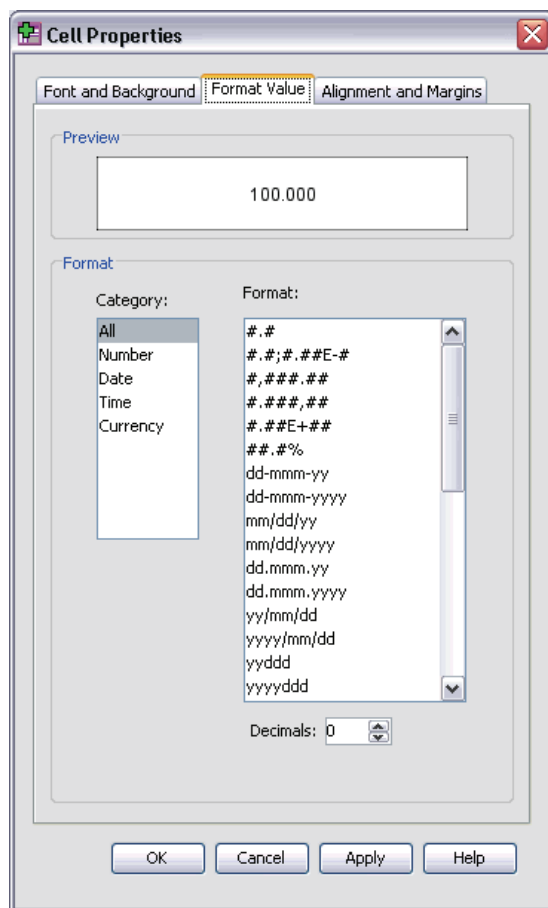
Figure 11-15
Cell Properties dialog box, Font and Background tab



Format Value

The Format Value tab controls value formats for the selected cells. You can select formats for numbers, dates, time, or currencies, and you can adjust the number of decimal digits that are displayed.

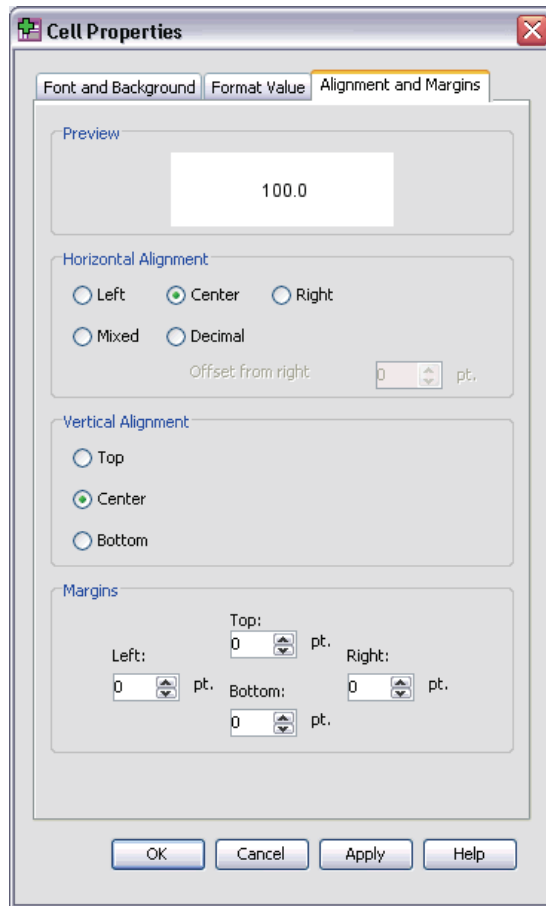
Figure 11-16
Cell Properties dialog box, Format Value tab



Alignment and Margins

The Alignment and Margins tab controls horizontal and vertical alignment of values and top, bottom, left, and right margins for the selected cells. Mixed horizontal alignment aligns the content of each cell according to its type. For example, dates are right-aligned and text values are left-aligned.

Figure 11-17
Cell Properties dialog box, Alignment and Margins tab



Footnotes and Captions

You can add footnotes and captions to a table. You can also hide footnotes or captions, change footnote markers, and renumber footnotes.

Some footnote attributes are controlled by table properties. [For more information, see Table Properties: Footnotes on p. 236.](#)

Adding Footnotes and Captions

To add a caption to a table:

- ▶ From the Insert menu choose Caption.

A footnote can be attached to any item in a table. To add a footnote:

- ▶ Click a title, cell, or caption within an activated pivot table.
- ▶ From the Insert menu choose Footnote.

To Hide or Show a Caption

To hide a caption:

- ▶ Select the caption.
- ▶ From the View menu choose Hide.

To show hidden captions:

- ▶ From the View menu choose Show All.

To Hide or Show a Footnote in a Table

To hide a footnote:

- ▶ Select the footnote.
- ▶ From the View menu choose Hide or from the context menu choose Hide Footnote.

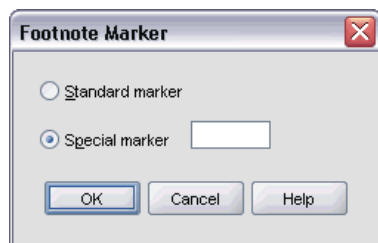
To show hidden footnotes:

- ▶ From the View menu choose Show All Footnotes.

Footnote Marker

Footnote Marker changes the character(s) that can be used to mark a footnote.

Figure 11-18
Footnote Marker dialog box



To change footnote markers:

- ▶ Select a footnote.
- ▶ From the Format menu choose Footnote Marker.
- ▶ Enter one or two characters.

Renumbering Footnotes


When you have pivoted a table by switching rows, columns, and layers, the footnotes may be out of order. To renumber the footnotes:

- ▶ From the Format menu choose Renumber Footnotes.

Data Cell Widths

Set Data Cell Width is used to set all data cells to the same width.

Figure 11-19
Set Data Cell Width dialog box



To set the width for all data cells:

- ▶ From the menus choose:
Format
Set Data Cell Widths...
- ▶ Enter a value for the cell width.

Changing Column Width

- ▶ Click and drag the column border.

Displaying Hidden Borders in a Pivot Table

For tables without many visible borders, you can display the hidden borders. This can simplify tasks like changing column widths.

- ▶ From the View menu choose Gridlines.

Figure 11-20
Gridlines displayed for hidden borders

Internet				Gender		Total		
				Female	Male			
No	Owns PDA	No	Count	1897	1962	3859		
			Expected Count	1903.4	1955.6	3859.0		
			% within Owns PDA	49.2%	50.8%	100.0%		
	Yes	Count	327	323	650			
Total				Gender		Total		
Yes	Owns PDA	No	Count	1897	1962	3859		
			Expected Count	1903.4	1955.6	3859.0		
			% within Owns PDA	49.2%	50.8%	100.0%		
		Yes	Count	327	323	650		
			Expected Count	320.6	329.4	650.0		
			% within Owns PDA	50.3%	49.7%	100.0%		
		Total			Count	2224	2285	4509
					Expected Count	2224.0	2285.0	4509.0
					% within Owns PDA	49.3%	50.7%	100.0%
Total	Yes	Owns PDA	No	Count	513	512	1025	
				Expected Count	512.5	512.5	1025.0	
				% within Owns PDA	50.0%	50.0%	100.0%	
	Yes	Count	305	306	611			
		Expected Count	305.5	305.5	611.0			
		% within Owns PDA	49.9%	50.1%	100.0%			
Total			Count	818	818	1636		
			Expected Count	818.0	818.0	1636.0		
			% within Owns PDA	50.0%	50.0%	100.0%		

Selecting Rows and Columns in a Pivot Table

In pivot tables, there are some constraints on how you select entire rows and columns, and the visual highlight that indicates the selected row or column may span noncontiguous areas of the table. To select an entire row or column:

- ▶ Click a row or column label.

- ▶ From the menus choose:

Edit
Select
Data and Label Cells

or

- ▶ Right-click the category label for the row or column.

- ▶ From the context menu choose:

Select
Data and Label Cells

or

- ▶ Ctrl+Alt+click the row or column label.

Printing Pivot Tables

Several factors can affect the way that printed pivot tables look, and these factors can be controlled by changing pivot table attributes.

- For multidimensional pivot tables (tables with layers), you can either print all layers or print only the top (visible) layer. [For more information, see Table Properties: Printing on p. 240.](#)
- For long or wide pivot tables, you can automatically resize the table to fit the page or control the location of table breaks and page breaks. [For more information, see Table Properties: Printing on p. 240.](#)
- For tables that are too wide or too long for a single page, you can control the location of table breaks between pages.

Use Print Preview on the File menu to see how printed pivot tables will look.

Controlling Table Breaks for Wide and Long Tables

Pivot tables that are either too wide or too long to print within the defined page size are automatically split and printed in multiple sections. (For wide tables, multiple sections will print on the same page if there is room.) You can:

- Control the row and column locations where large tables are split.
- Specify rows and columns that should be kept together when tables are split.
- Rescale large tables to fit the defined page size.

To Specify Row and Column Breaks for Printed Pivot Tables

- ▶ Click the column label to the left of where you want to insert the break, or click the row label above where you want to insert the break.
- ▶ From the menus choose:
 - Format
 - Break Here

To Specify Rows or Columns to Keep Together

- ▶ Select the labels of the rows or columns that you want to keep together. (Click and drag or Shift+click to select multiple row or column labels.)
- ▶ From the menus choose:
 - Format
 - Keep Together

Creating a Chart from a Pivot Table

- ▶ Double-click the pivot table to activate it.
- ▶ Select the rows, columns, or cells you want to display in the chart.
- ▶ Right-click anywhere in the selected area.

- ▶ Choose Create Graph from the context menu and select a chart type.

Models

Some results are presented as models, which appear in the output Viewer as a special type of visualization. The visualization displayed in the output Viewer is not the only view of the model that is available. A single model contains many different views. You can activate the model in the Model Viewer and interact with the model directly to display the available model views. You can also choose to print and export all the views in the model.

Interacting with a Model

To interact with a model, you first activate it:

- ▶ Double-click the model.

or

- ▶ Right-click the model and from the context menu choose Edit Content.
- ▶ From the submenu choose In Separate Window.

Activating the model displays the model in the Model Viewer. [For more information, see Working with the Model Viewer on p. 250.](#)

Working with the Model Viewer

The Model Viewer is an interactive tool for displaying the available model views and editing the look of the model views. (For information about displaying the Model Viewer, see [Interacting with a Model](#) on p. 250.) The Model Viewer is split into two parts:

- **Main view.** The main view appears in the left part of the Model Viewer. The main view displays some general visualization (for example, a network graph) for the model. The main view itself may have more than one model view. The drop-down list below the main view allows you to choose from the available main views.
- **Auxiliary view.** The auxiliary view appears in the right part of the Model Viewer. The auxiliary view typically displays a more detailed visualization (including tables) of the model compared to the general visualization in the main view. Like the main view, the auxiliary view may have more than one model view. The drop-down list below the auxiliary view allows you to choose from the available main views. The auxiliary can also display specific visualizations for elements that are selected in the main view. For example, depending on the type of model, you may be able to select a variable node in the main view to display a table for that variable in the auxiliary view.

The specific visualizations that are displayed depend on the procedure that created the model. For information about working with specific models, refer to the documentation for the procedure that created the model.

Model View Tables

Tables displayed in the Model Viewer are not pivot tables. You cannot manipulate these tables as you can manipulate pivot tables.

Editing and Exploring Model Views

By default, the Model Viewer is activated in **Explore mode**. Explore mode allows you to interact with the model by dragging elements and selecting them to display related visualization in the auxiliary view.

The Model Viewer also offers an **Edit mode**. Edit mode allows you to edit the styles of the model views. The Model Viewer includes the Graphboard Editor for making these changes. When you close the Model Viewer, any edits that you make are saved to the model.

The View menu allows you to switch between Explore mode and Edit mode.

Setting Model Properties

Within the Model Viewer, you can set specific properties for the model. [For more information, see Model Properties on p. 251.](#)

Copying Model Views

You can also copy individual model views within the Model Viewer. [For more information, see Copying Model Views on p. 251.](#)

Model Properties

From within the Model Viewer, choose from the menus:

- File
- Properties

Each model has associated properties that let you specify which views are printed from the output Viewer. By default, only the view that is visible in the output Viewer is printed. This is always a main view, and only one main view. You can also specify that all available model views are printed. These include all the main views and all the auxiliary views. Note that you can also print individual model views within the Model Viewer itself. [For more information, see Printing a Model on p. 252.](#)

Copying Model Views

From the Edit menu within the Model Viewer, you can copy the currently displayed main view or the currently display auxiliary view. Only one model view is copied. You can paste the model view into the output Viewer, where the individual model view is subsequently rendered as a

visualization that can be edited in the Graphboard Editor. Pasting into the output Viewer allows you to display multiple model views simultaneously. You can also paste into other applications, where the view may appear as an image or a table depending on the target application.

Printing a Model

Printing from the Model Viewer

You can print a single model view within the Model Viewer itself.

- ▶ Activate the model in the Model Viewer. [For more information, see Interacting with a Model on p. 250.](#)
- ▶ From the menus choose:
 - View
 - Edit Mode
- ▶ On the General toolbar palette in the main or auxiliary view (depending on which one you want to print), click the print icon. (If this palette is not displayed, choose Palettes>General from the View menu.)

Printing from the Output Viewer

When you print from the output Viewer, the number of views that are printed for a specific model depend on the model's properties. The model can be set to print only the displayed view or all of the available model views. [For more information, see Model Properties on p. 251.](#)

Exporting a Model

By default, when you export models from the output Viewer, inclusion or exclusion of model views is controlled by the model properties for each model. For more information about model properties, see [Model Properties](#) on p. 251. On export, you can override this setting and include all model views or only the currently visible model view. In the Export Output dialog box, click Change Options... in the Document group. For more information about exporting and this dialog box, see [Export Output](#) on p. 210. Note that all model views, including tables, are exported as graphics.

Working with Command Syntax

The powerful command language allows you to save and automate many common tasks. It also provides some functionality not found in the menus and dialog boxes.

Most commands are accessible from the menus and dialog boxes. However, some commands and options are available only by using the command language. The command language also allows you to save your jobs in a syntax file so that you can repeat your analysis at a later date or run it in an automated job with the a production job.

A syntax file is simply a text file that contains commands. While it is possible to open a syntax window and type in commands, it is often easier if you let the software help you build your syntax file using one of the following methods:

- Pasting command syntax from dialog boxes
- Copying syntax from the output log
- Copying syntax from the journal file

Detailed command syntax reference information is available in two forms: integrated into the overall Help system and as a separate PDF file, called the *Command Syntax Reference*, also available from the Help menu.

Context-sensitive Help for the current command in a syntax window is available by pressing the F1 key.

Syntax Rules

When you run commands from a command syntax window during a session, you are running commands in **interactive** mode.

The following rules apply to command specifications in interactive mode:

- Each command must start on a new line. Commands can begin in any column of a command line and continue for as many lines as needed. The exception is the `END DATA` command, which must begin in the first column of the first line after the end of data.
- Each command should end with a period as a command terminator. It is best to omit the terminator on `BEGIN DATA`, however, so that inline data are treated as one continuous specification.
- The command terminator must be the last nonblank character in a command.
- In the absence of a period as the command terminator, a blank line is interpreted as a command terminator.

Note: For compatibility with other modes of command execution (including command files run with `INSERT` or `INCLUDE` commands in an interactive session), each line of command syntax should not exceed 256 bytes.

- Most subcommands are separated by slashes (/). The slash before the first subcommand on a command is usually optional.
- Variable names must be spelled out fully.
- Text included within apostrophes or quotation marks must be contained on a single line.
- A period (.) must be used to indicate decimals, regardless of your regional or locale settings.
- Variable names ending in a period can cause errors in commands created by the dialog boxes. You cannot create such variable names in the dialog boxes, and you should generally avoid them.

Command syntax is case insensitive, and three- or four-letter abbreviations can be used for many command specifications. You can use as many lines as you want to specify a single command. You can add space or break lines at almost any point where a single blank is allowed, such as around slashes, parentheses, arithmetic operators, or between variable names. For example,

```
FREQUENCIES
VARIABLES=JOB CAT GENDER
/PERCENTILES=25 50 75
/BARCHART.
```

and

```
freq var=jobcat gender /percent=25 50 75 /bar.
```

are both acceptable alternatives that generate the same results.

INCLUDE Files

For command files run via the `INCLUDE` command, batch mode syntax rules apply.

The following rules apply to command specifications in batch mode:

- All commands in the command file must begin in column 1. You can use plus (+) or minus (–) signs in the first column if you want to indent the command specification to make the command file more readable.
- If multiple lines are used for a command, column 1 of each continuation line must be blank.
- Command terminators are optional.
- A line cannot exceed 256 bytes; any additional characters are truncated.

Unless you have existing command files that already use the `INCLUDE` command, you should probably use the `INSERT` command instead, since it can accommodate command files that conform to either set of rules. If you generate command syntax by pasting dialog box choices into a syntax window, the format of the commands is suitable for any mode of operation. See the *Command Syntax Reference* (available in PDF format from the Help menu) for more information.

Pasting Syntax from Dialog Boxes

The easiest way to build a command syntax file is to make selections in dialog boxes and paste the syntax for the selections into a syntax window. By pasting the syntax at each step of a lengthy analysis, you can build a job file that allows you to repeat the analysis at a later date or run an automated job with the Production Facility.

In the syntax window, you can run the pasted syntax, edit it, and save it in a syntax file.

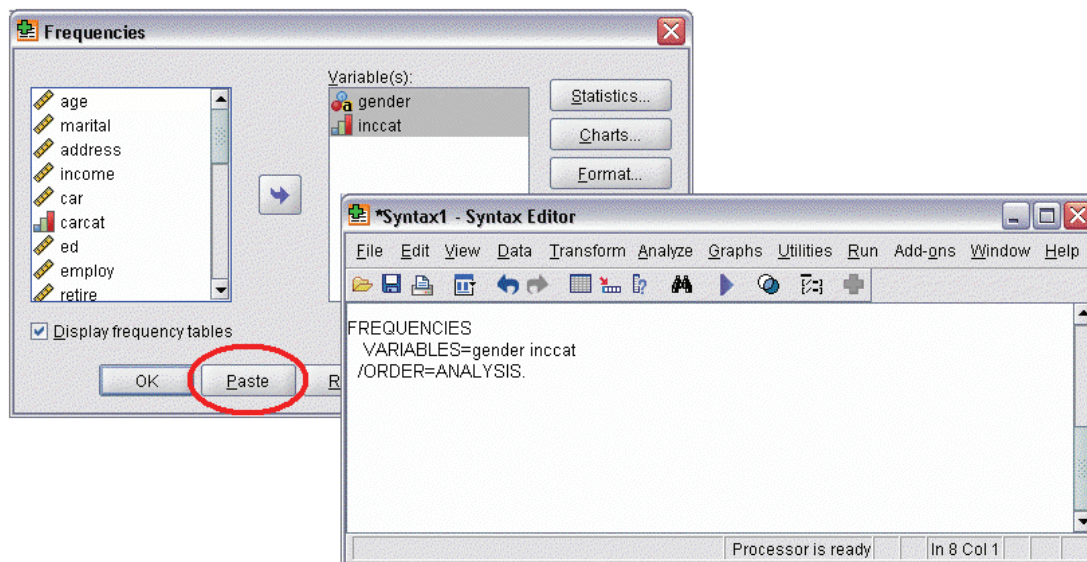
To Paste Syntax from Dialog Boxes

- ▶ Open the dialog box and make the selections that you want.
- ▶ Click Paste.

The command syntax is pasted to the designated syntax window. If you do not have an open syntax window, a new syntax window opens automatically, and the syntax is pasted there.

Figure 13-1

Command syntax pasted from a dialog box

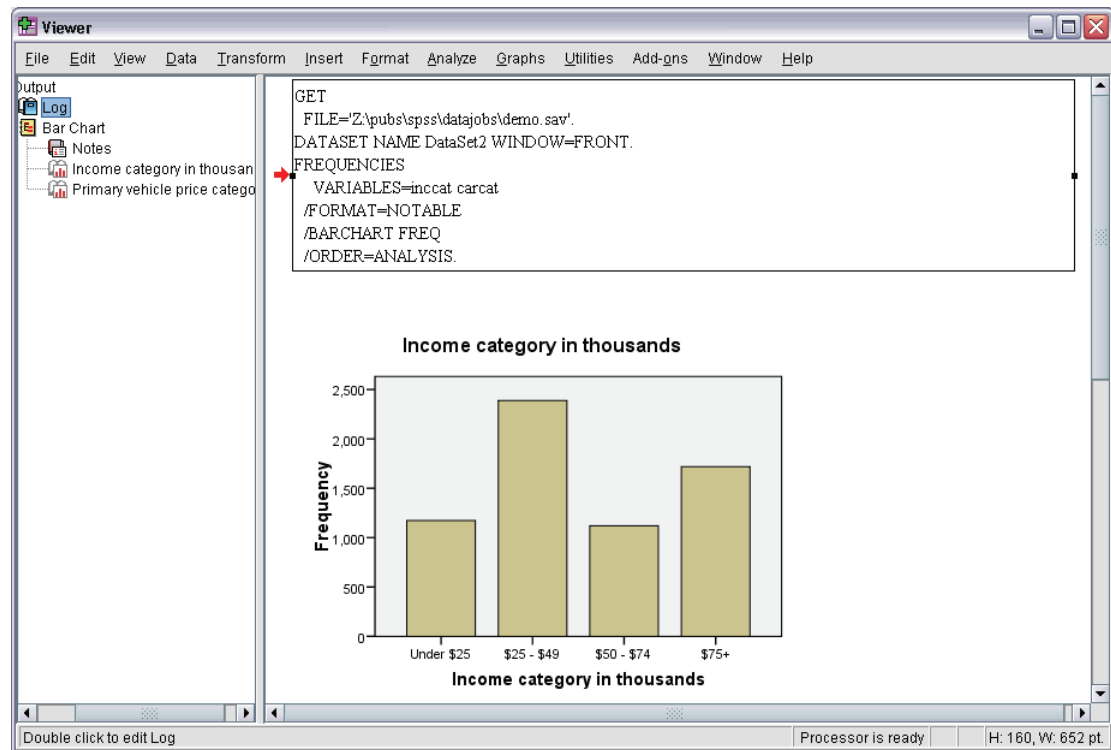


Copying Syntax from the Output Log

You can build a syntax file by copying command syntax from the log that appears in the Viewer. To use this method, you must select Display commands in the log in the Viewer settings (Edit menu, Options, Viewer tab) before running the analysis. Each command will then appear in the Viewer along with the output from the analysis.

In the syntax window, you can run the pasted syntax, edit it, and save it in a syntax file.

Figure 13-2
Command syntax in the log



To Copy Syntax from the Output Log

- ▶ Before running the analysis, from the menus choose:
Edit
Options...
- ▶ On the Viewer tab, select Display commands in the log.
As you run analyses, the commands for your dialog box selections are recorded in the log.
- ▶ Open a previously saved syntax file or create a new one. To create a new syntax file, from the menus choose:
File
New
Syntax
- ▶ In the Viewer, double-click a log item to activate it.
- ▶ Select the text that you want to copy.
- ▶ From the Viewer menus choose:
Edit
Copy

- ▶ In a syntax window, from the menus choose:
 - Edit
 - Paste

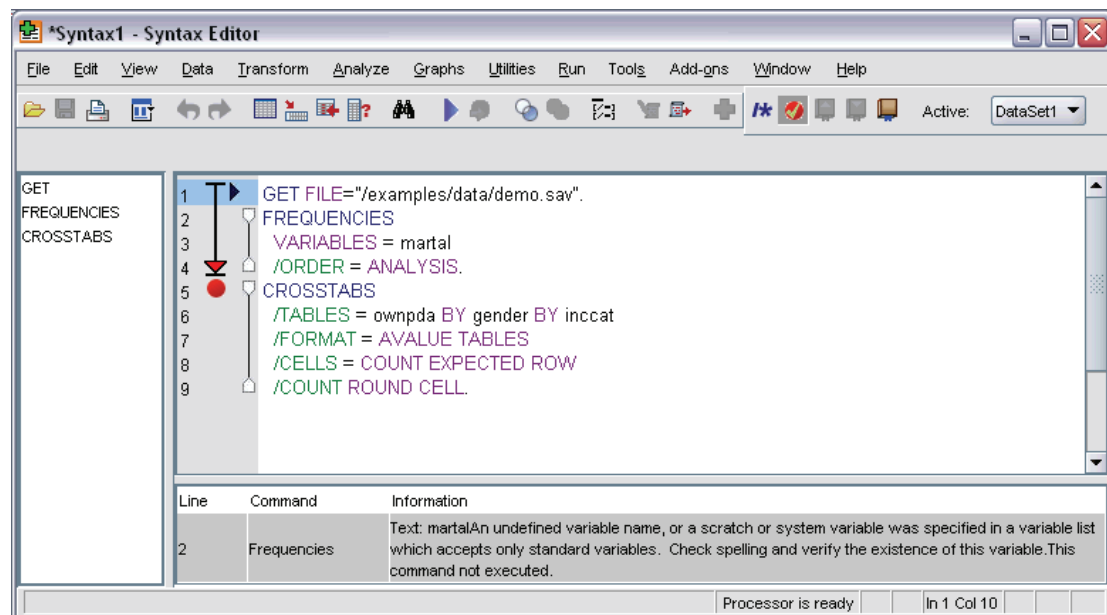
Using the Syntax Editor

The Syntax Editor provides an environment specifically designed for creating, editing, and running command syntax. The Syntax Editor features:

- **Auto-Completion.** As you type, you can select commands, subcommands, keywords, and keyword values from a context-sensitive list. You can choose to be prompted automatically with the list or display the list on demand.
- **Color Coding.** Recognized elements of command syntax (commands, subcommands, keywords, and keyword values) are color coded so, at a glance, you can spot unrecognized terms. Also, a number of common syntactical errors—such as unmatched quotes—are color coded for quick identification.
- **Breakpoints.** You can stop execution of command syntax at specified points, allowing you to inspect the data or output before proceeding.
- **Bookmarks.** You can set bookmarks that allow you to quickly navigate large command syntax files.
- **Step Through.** You can step through command syntax one command at a time, advancing to the next command with a single click.

Syntax Editor Window

Figure 13-3
Syntax Editor



The Syntax Editor window is divided into four areas:

- The editor pane is the main part of the Syntax Editor window and is where you enter and edit command syntax.
- The gutter is adjacent to the editor pane and displays information such as line numbers and breakpoint positions.
- The navigation pane is to the left of the gutter and editor pane and displays a list of all commands in the Syntax Editor window and provides single click navigation to any command.
- The error pane is below the editor pane and displays runtime errors.

Gutter Contents

Line numbers, breakpoints, bookmarks, command spans, and a progress indicator are displayed in the gutter to the left of the editor pane in the syntax window.

- Line numbers do not account for any external files referenced in `INSERT` and `INCLUDE` commands. You can show or hide line numbers by choosing `View > Show Line Numbers` from the menus.
- Breakpoints stop execution at specified points and are represented as a red circle adjacent to the command on which the breakpoint is set.
- Bookmarks mark specific lines in a command syntax file and are represented as a square enclosing the number (1-9) assigned to the bookmark. Hovering over the icon for a bookmark displays the number of the bookmark and the name, if any, assigned to the bookmark.
- Command spans are icons that provide visual indicators of the start and end of a command. You can show or hide command spans by choosing `View > Show Command Spans` from the menus.
- The progress of a given syntax run is indicated with a downward pointing arrow in the gutter, stretching from the first command run to the last command run. This is most useful when running command syntax containing breakpoints and when stepping through command syntax. [For more information, see *Running Command Syntax* on p. 263.](#)

Navigation Pane

The navigation pane contains a list of all recognized commands in the syntax window, displayed in the order in which they occur in the window. Clicking on a command in the navigation pane positions the cursor at the start of the command.

- You can use the Up and Down arrow keys to move through the list of commands or click on a command to navigate to it. A double click will select the command.
- Command names for commands containing certain types of syntactical errors—such as unmatched quotes—are colored red and in bold text by default. [For more information, see *Color Coding* on p. 260.](#)
- The first word of each line of unrecognized text is shown in gray.
- You can show or hide the navigation pane by choosing `View > Show Navigation Pane` from the menus.

Error Pane

The error pane displays runtime errors from the most previous run.

- The information for each error contains the line number on which the error occurred.
- You can use the Up and Down arrow keys to move through the list of errors.
- Clicking on an entry in the list will position the cursor on the line that generated the error.
- You can show or hide the error pane by choosing View > Show Error Pane from the menus.

Terminology

Commands. The basic unit of syntax is the command. Each command begins with the command name, which consists of one, two, or three words—for example, `DESCRIPTIVES`, `SORT CASES`, or `ADD VALUE LABELS`.

Subcommands. Most commands contain subcommands. Subcommands provide for additional specifications and begin with a forward slash followed by the name of the subcommand.

Keywords. Keywords are fixed terms that are typically used within a subcommand to specify options available for the subcommand.

Keyword Values. Keywords can have values such as a fixed term that specifies an option or a numeric value.

Example

```
CODEBOOK gender jobcat salary  
  /VARINFO VALUELABELS MISSING  
  /OPTIONS VARORDER=MEASURE.
```

- The name of the command is `CODEBOOK`.
- `VARINFO` and `OPTIONS` are subcommands.
- `VALUELABELS`, `MISSING`, and `VARORDER` are keywords.
- `MEASURE` is a keyword value associated with `VARORDER`.

Auto-Completion

The Syntax Editor provides assistance in the form of auto-completion of commands, subcommands, keywords, and keyword values. By default, you are prompted with a context-sensitive list of available terms. You can display the list on demand by pressing Ctrl+Spacebar and you can close the list by pressing the Esc key.

The Auto Complete menu item on the Tools menu toggles the automatic display of the auto-complete list on or off. You can also enable or disable automatic display of the list from the Syntax Editor tab in the Options dialog box. Toggling the Auto Complete menu item overrides the setting on the Options dialog but does not persist across sessions.

Note: The auto-completion list will close if a space is entered. For commands consisting of multiple words—such as `ADD FILES`—select the desired command before entering any spaces.

Color Coding

The Syntax Editor color codes recognized elements of command syntax, such as commands and subcommands, as well as various syntactical errors like unmatched quotes or parentheses. Unrecognized text is not color coded.

Commands. By default, recognized commands are colored blue and in bold text. If, however, there is a recognized syntactical error within the command—such as a missing parenthesis—the command name is colored red and in bold text by default.

Note: Abbreviations of command names—such as `FREQ` for `FREQUENCIES`—are not colored, but such abbreviations are valid.

Subcommands. Recognized subcommands are colored green by default. If, however, the subcommand is missing a required equals sign or an invalid equals sign follows it, the subcommand name is colored red by default.

Keywords. Recognized keywords are colored maroon by default. If, however, the keyword is missing a required equals sign or an invalid equals sign follows it, the keyword is colored red by default.

Keyword values. Recognized keyword values are colored pink by default. User-specified values of keywords such as integers, real numbers, and quoted strings are not color coded.

Comments. Text within a comment is colored gray.

Syntactical Errors. Text associated with the following syntactical errors is colored red by default.

- **Unmatched Parentheses, Brackets, and Quotes.** Unmatched parentheses and brackets within comments and quoted strings are not detected. Unmatched single or double quotes within quoted strings are syntactically valid.

Certain commands contain blocks of text that are not command syntax—such as `BEGIN DATA-END DATA`, `BEGIN GPL-END GPL`, and `BEGIN PROGRAM-END PROGRAM`. Unmatched values are not detected within such blocks.

- **Long lines.** Long lines are lines containing more than 251 characters.
- **End statements.** Several commands require either an `END` statement prior to the command terminator (for example, `BEGIN DATA-END DATA`) or require a matching `END` command at some point later in the command stream (for example, `LOOP-END LOOP`). In both cases, the command will be colored red, by default, until the required `END` statement is added.

From the Syntax Editor tab in the Options dialog box, you can change default colors and text styles and you can turn color coding off or on. You can also turn color coding of commands, subcommands, keywords, and keyword values off or on by choosing `Tools > Color Coding` from the menus. You can turn color coding of syntactical errors off or on by choosing `Tools > Validation`. Choices made on the Tools menu override settings in the Options dialog box but do not persist across sessions.

Note: Color coding of command syntax within macros is not supported.

Breakpoints

Breakpoints allow you to stop execution of command syntax at specified points within the syntax window and continue execution when ready.

- Breakpoints are set at the level of a command and stop execution prior to running the command.
- Breakpoints cannot occur within `LOOP-END LOOP`, `DO IF-END IF`, `DO REPEAT-END REPEAT`, `INPUT PROGRAM-END INPUT PROGRAM`, and `MATRIX-END MATRIX` blocks. They can, however, be set at the beginning of such blocks and will stop execution prior to running the block.
- Breakpoints cannot be set on lines containing non-SPSS Statistics command syntax, such as occur within `BEGIN PROGRAM-END PROGRAM`, `BEGIN DATA-END DATA`, and `BEGIN GPL-END GPL` blocks.
- Breakpoints are not saved with the command syntax file and are not included in copied text.
- By default, breakpoints are honored during execution. You can toggle whether breakpoints are honored or not from **Tools > Honor Breakpoints**.

To Insert a Breakpoint

- ▶ Click anywhere in the gutter to the left of the command text.

or

- ▶ Position the cursor within the desired command.
- ▶ From the menus choose:
 - Tools
 - Toggle Breakpoint

The breakpoint is represented as a red circle in the gutter to the left of the command text and on the same line as the command name.

Clearing Breakpoints

To clear a single breakpoint:

- ▶ Click on the icon representing the breakpoint in the gutter to the left of the command text.

or

- ▶ Position the cursor within the desired command.
- ▶ From the menus choose:
 - Tools
 - Toggle Breakpoint

To clear all breakpoints:

- ▶ From the menus choose:
 - Tools
 - Clear All Breakpoints

See [Running Command Syntax](#) on p. 263 for information about the run-time behavior in the presence of breakpoints.

Bookmarks

Bookmarks allow you to quickly navigate to specified positions in a command syntax file. You can have up to 9 bookmarks in a given file. Bookmarks are saved with the file, but are not included when copying text.

To Insert a Bookmark

- ▶ Position the cursor on the line where you want to insert the bookmark.
- ▶ From the menus choose:
 - Tools
 - Toggle Bookmark

The new bookmark is assigned the next available number, from 1 to 9. It is represented as a square enclosing the assigned number and displayed in the gutter to the left of the command text.

Clearing Bookmarks

To clear a single bookmark:

- ▶ Position the cursor on the line containing the bookmark.
- ▶ From the menus choose:
 - Tools
 - Toggle Bookmark

To clear all bookmarks:

- ▶ From the menus choose:
 - Tools
 - Clear All Bookmarks

Renaming a Bookmark

You can associate a name with a bookmark. This is in addition to the number (1-9) assigned to the bookmark when it was created.

- ▶ From the menus choose:
 - Tools
 - Rename Bookmark

- ▶ Enter a name for the bookmark and click OK.

The specified name replaces any existing name for the bookmark.

Navigating with Bookmarks

To navigate to the next or previous bookmark:

- ▶ From the menus choose:
 - Tools
 - Next Bookmark

or

- Tools
- Previous Bookmark

To navigate to a specific bookmark:

- ▶ From the menus choose:
 - Tools
 - Go To Bookmark
- ▶ Select the desired bookmark.

Commenting Out Text

You can comment out entire commands as well as text that is not recognized as command syntax.

- ▶ Select the desired text. When selecting multiple commands, a command will be commented out if any part of it is selected.
- ▶ From the menus choose:
 - Tools
 - Comment Selection

You can comment out a single command by positioning the cursor anywhere within the command and clicking the Comment Selection button on the toolbar.

Running Command Syntax

- ▶ Highlight the commands that you want to run in the syntax window.
- ▶ Click the Run button (the right-pointing triangle) on the Syntax Editor toolbar. It runs the selected commands or the command where the cursor is located if there is no selection.

or

- ▶ Choose one of the items from the Run menu.
 - **All.** Runs all commands in the syntax window, honoring any breakpoints.
 - **Selection.** Runs the currently selected commands, honoring any breakpoints. This includes any partially highlighted commands. If there is no selection, the command where the cursor is positioned is run.

- **To End.** Runs all commands starting from the first command in the current selection to the last command in the syntax window, honoring any breakpoints. If nothing is selected, the run starts from the command where the cursor is positioned.
- **Step Through.** Runs the command syntax one command at a time starting from the first command in the syntax window (Step Through From Start) or from the command where the cursor is positioned (Step Through From Current). If there is selected text, the run starts from the first command in the selection. After a given command has run, the cursor advances to the next command and you continue the step through sequence by choosing Continue.

LOOP-END LOOP, DO IF-END IF, DO REPEAT-END REPEAT, INPUT PROGRAM-END INPUT PROGRAM, and MATRIX-END MATRIX blocks are treated as single commands when using Step Through. You can not step into one of these blocks.
- **Continue.** Continues a run stopped by a breakpoint or Step Through.

Progress Indicator

The progress of a given syntax run is indicated with a downward pointing arrow in the gutter, spanning the last set of commands run. For instance, you choose to run all commands in a syntax window that contains breakpoints. At the first breakpoint, the arrow will span the region from the first command in the window to the command prior to the one containing the breakpoint. At the second breakpoint, the arrow will stretch from the command containing the first breakpoint to the command prior to the one containing the second breakpoint.

Run-time Behavior with Breakpoints

- When running command syntax containing breakpoints, execution stops at each breakpoint. Specifically, the block of command syntax from a given breakpoint (or beginning of the run) to the next breakpoint (or end of the run) is submitted for execution exactly as if you had selected that syntax and chosen Run > Selection.
- You can work with multiple syntax windows, each with its own set of breakpoints, but there is only one queue for executing command syntax. Once a block of command syntax has been submitted—such as the block of command syntax up to the first breakpoint—no other block of command syntax will be executed until the previous block has completed, regardless of whether the blocks are in the same or different syntax windows.
- With execution stopped at a breakpoint, you can run command syntax in other syntax windows, and inspect Data Editor or Viewer windows. However, modifying the contents of the syntax window containing the breakpoint or changing the cursor position in that window will cancel the run.

Unicode Syntax Files

In Unicode mode, the default format for saving command syntax files created or modified during the session is also Unicode (UTF-8). Unicode-format command syntax files cannot be read by versions of SPSS Statistics prior to 16.0. For more information on Unicode mode, see [General Options on p. 494](#).

To save a syntax file in a format compatible with earlier releases:

- ▶ From the syntax window menus, choose:
 - File
 - Save As
- ▶ In the Save As dialog, from the Encoding drop-down list, choose Local Encoding. The local encoding is determined by the current locale.

Multiple Execute Commands

Syntax pasted from dialog boxes or copied from the log or the journal may contain EXECUTE commands. When you run commands from a syntax window, EXECUTE commands are generally unnecessary and may slow performance, particularly with larger data files, because each EXECUTE command reads the entire data file. For more information, see the EXECUTE command in the *Command Syntax Reference* (available from the Help menu in any SPSS Statistics window).

Lag Functions

One notable exception is transformation commands that contain lag functions. In a series of transformation commands without any intervening EXECUTE commands or other commands that read the data, lag functions are calculated after all other transformations, regardless of command order. For example,

```
COMPUTE lagvar=LAG(var1).  
COMPUTE var1=var1*2.
```

and

```
COMPUTE lagvar=LAG(var1).  
EXECUTE.  
COMPUTE var1=var1*2.
```

yield very different results for the value of *lagvar*, since the former uses the transformed value of *var1* while the latter uses the original value.

Codebook

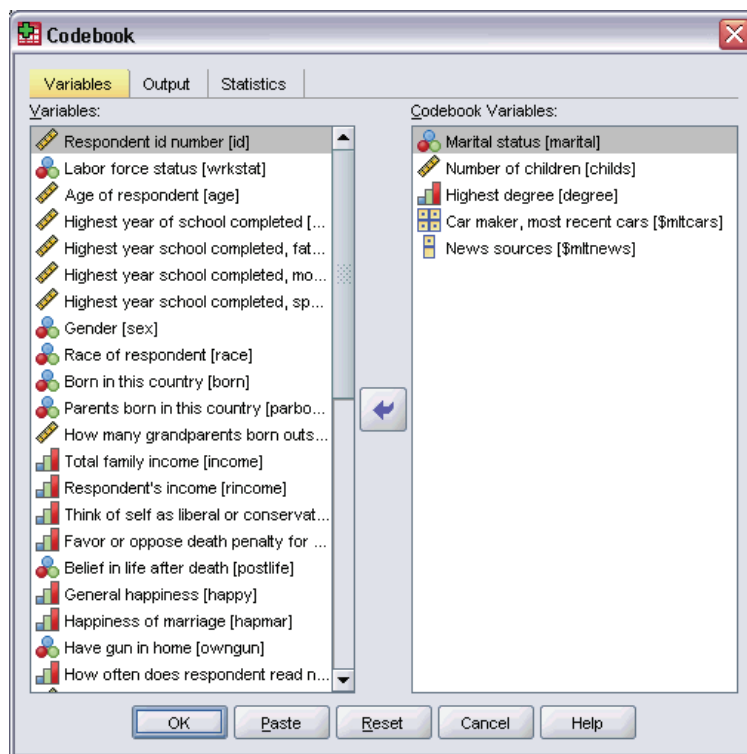
Codebook reports the dictionary information — such as variable names, variable labels, value labels, missing values — and summary statistics for all or specified variables and multiple response sets in the active dataset. For nominal and ordinal variables and multiple response sets, summary statistics include counts and percents. For scale variables, summary statistics include mean, standard deviation, and quartiles.

Note: Codebook ignores split file status. This includes split-file groups created for multiple imputation of missing values (available in the Missing Values add-on option). For more information on split-file processing, see [Split File](#).

To Obtain a Codebook

- ▶ From the menus choose:
 - Analyze
 - Reports
 - Codebook
- ▶ Click the Variables tab.

Figure 14-1
Codebook dialog, Variables tab



- ▶ Select one or more variables and/or multiple response sets.

Optionally, you can:

- Control the variable information that is displayed.
- Control the statistics that are displayed (or exclude all summary statistics).
- Control the order in which variables and multiple response sets are displayed.
- Change the measurement level for any variable in the source list in order to change the summary statistics displayed. [For more information, see Codebook Statistics Tab on p. 270.](#)

Changing Measurement Level

You can temporarily change the measurement level for variables. (You cannot change the measurement level for multiple response sets. They are always treated as nominal.)

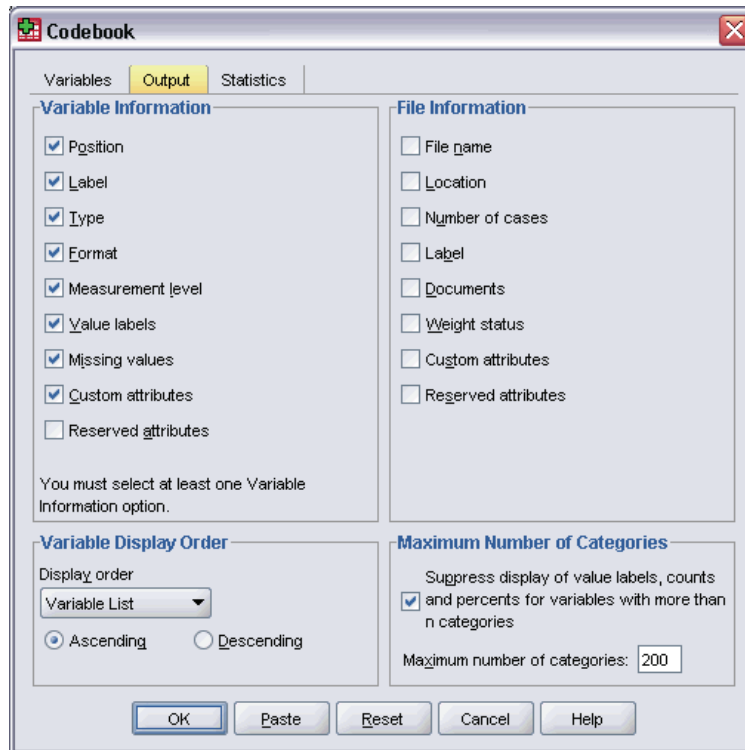
- ▶ Right-click a variable in the source list.
- ▶ Select a measurement level from the pop-up context menu.

This changes the measurement level temporarily. In practical terms, this is only useful for numeric variables. The measurement level for string variables is restricted to nominal or ordinal, which are both treated the same by the Codebook procedure.

Codebook Output Tab

The Output tab controls the variable information included for each variable and multiple response set, the order in which the variables and multiple response sets are displayed, and the contents of the optional file information table.

Figure 14-2
Codebook dialog, Output tab



Variable Information

This controls the dictionary information displayed for each variable.

Position. An integer that represents the position of the variable in file order. This is not available for multiple response sets.

Label. The descriptive label associated with the variable or multiple response set.

Type. Fundamental data type. This is either *Numeric*, *String*, or *Multiple Response Set*.

Format. The display format for the variable, such as *A4*, *F8.2*, or *DATE11*. This is not available for multiple response sets.

Measurement level. The possible values are *Nominal*, *Ordinal*, *Scale*, and *Unknown*. The value displayed is the measurement level stored in the dictionary and is not affected by any temporary measurement level override specified by changing the measurement level in the source variable list on the Variables tab. This is not available for multiple response sets.

Note: The measurement level for numeric variables may be “unknown” prior to the first data pass when the measurement level has not been explicitly set, such as data read from an external source or newly created variables. [For more information, see Variable Measurement Level in Chapter 5 on p. 73.](#)

Value labels. Descriptive labels associated with specific data values. [For more information, see Value Labels in Chapter 5 on p. 76.](#)

- If Count or Percent is selected on the Statistics tab, defined value labels are included in the output even if you don’t select Value labels here.
- For multiple dichotomy sets, “value labels” are either the variable labels for the elementary variables in the set or the labels of counted values, depending on how the set is defined. [For more information, see Multiple Response Sets in Chapter 7 on p. 103.](#)

Missing values. User-defined missing values. [For more information, see Missing Values in Chapter 5 on p. 77.](#) If Count or Percent is selected on the Statistics tab, defined value labels are included in the output even if you don’t select Missing values here. This is not available for multiple response sets.

Custom attributes. User-defined custom variable attributes. Output includes both the names and values for any custom variable attributes associated with each variable. [For more information, see Creating Custom Variable Attributes in Chapter 5 on p. 80.](#) This is not available for multiple response sets.

Reserved attributes. Reserved system variable attributes. You can display system attributes, but you should not alter them. System attribute names start with a dollar sign (\$) . Non-display attributes, with names that begin with either “@” or “\$@”, are not included. Output includes both the names and values for any system attributes associated with each variable. This is not available for multiple response sets.

File Information

The optional file information table can include any of the following file attributes:

File name. Name of the SPSS Statistics data file. If the dataset has never been saved in SPSS Statistics format, then there is no data file name. (If there is no file name displayed in the title bar of the Data Editor window, then the active dataset does not have a file name.)

Location. Directory (folder) location of the SPSS Statistics data file. If the dataset has never been saved in SPSS Statistics format, then there is no location.

Number of cases. Number of cases in the active dataset. This is the total number of cases, including any cases that may be excluded from summary statistics due to filter conditions.

Label. This is the file label (if any) defined by the `FILE LABEL` command.

Documents. Data file document text. [For more information, see Data File Comments in Chapter 47 on p. 489.](#)

Weight status. If weighting is on, the name of the weight variable is displayed. [For more information, see Weight Cases in Chapter 9 on p. 183.](#)

Custom attributes. User-defined custom data file attributes. Data file attributes defined with the `DATAFILE ATTRIBUTE` command.

Reserved attributes. Reserved system data file attributes. You can display system attributes, but you should not alter them. System attribute names start with a dollar sign (\$) . Non-display attributes, with names that begin with either “@” or “\$@”, are not included. Output includes both the names and values for any system data file attributes.

Variable Display Order

The following alternatives are available for controlling the order in which variables and multiple response sets are displayed.

Alphabetical. Alphabetic order by variable name.

File. The order in which variables appear in the dataset (the order in which they are displayed in the Data Editor). In ascending order, multiple response sets are displayed last, after all selected variables.

Measurement level. Sort by measurement level. This creates four sorting groups: nominal, ordinal, scale, and unknown. Multiple response sets are treated as nominal.

Note: The measurement level for numeric variables may be “unknown” prior to the first data pass when the measurement level has not been explicitly set, such as data read from an external source or newly created variables.

Variable list. The order in which variables and multiple response sets appear in the selected variables list on the Variables tab.

Custom attribute name. The list of sort order options also includes the names of any user-defined custom variable attributes. In ascending order, variables that don’t have the attribute sort to the top, followed by variables that have the attribute but no defined value for the attribute, followed by variables with defined values for the attribute in alphabetic order of the values.

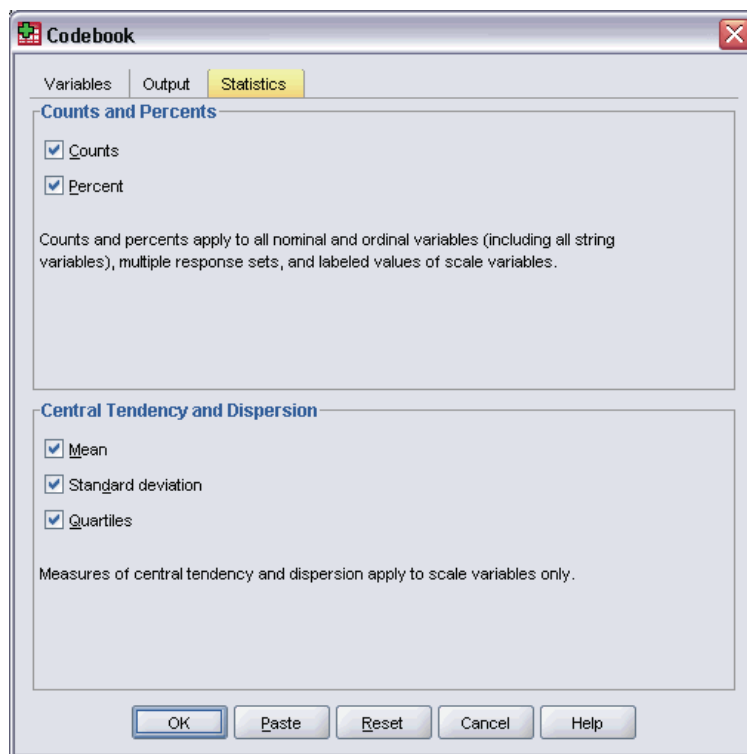
Maximum Number of Categories

If the output includes value labels, counts, or percents for each unique value, you can suppress this information from the table if the number of values exceeds the specified value. By default, this information is suppressed if the number of unique values for the variable exceeds 200.

Codebook Statistics Tab

The Statistics tab allows you to control the summary statistics that are included in the output, or suppress the display of summary statistics entirely.

Figure 14-3
Codebook dialog, Statistics tab



Counts and Percents

For nominal and ordinal variables, multiple response sets, and labeled values or scale variables, the available statistics are:

Count. The count or number of cases having each value (or range of values) of a variable.

Percent. The percentage of cases having a particular value.

Central Tendency and Dispersion

For scale variables, the available statistics are:

Mean. A measure of central tendency. The arithmetic average, the sum divided by the number of cases.

Standard Deviation. A measure of dispersion around the mean. In a normal distribution, 68% of cases fall within one standard deviation of the mean and 95% of cases fall within two standard deviations. For example, if the mean age is 45, with a standard deviation of 10, 95% of the cases would be between 25 and 65 in a normal distribution.

Quartiles. Displays values corresponding to the 25th, 50th, and 75th percentiles.

Note: You can temporarily change the measurement level associated with a variable (and thereby change the summary statistics displayed for that variable) in the source variable list on the Variables tab.

Frequencies

The Frequencies procedure provides statistics and graphical displays that are useful for describing many types of variables. The Frequencies procedure is a good place to start looking at your data.

For a frequency report and bar chart, you can arrange the distinct values in ascending or descending order, or you can order the categories by their frequencies. The frequencies report can be suppressed when a variable has many distinct values. You can label charts with frequencies (the default) or percentages.

Example. What is the distribution of a company's customers by industry type? From the output, you might learn that 37.5% of your customers are in government agencies, 24.9% are in corporations, 28.1% are in academic institutions, and 9.4% are in the healthcare industry. For continuous, quantitative data, such as sales revenue, you might learn that the average product sale is \$3,576, with a standard deviation of \$1,078.

Statistics and plots. Frequency counts, percentages, cumulative percentages, mean, median, mode, sum, standard deviation, variance, range, minimum and maximum values, standard error of the mean, skewness and kurtosis (both with standard errors), quartiles, user-specified percentiles, bar charts, pie charts, and histograms.

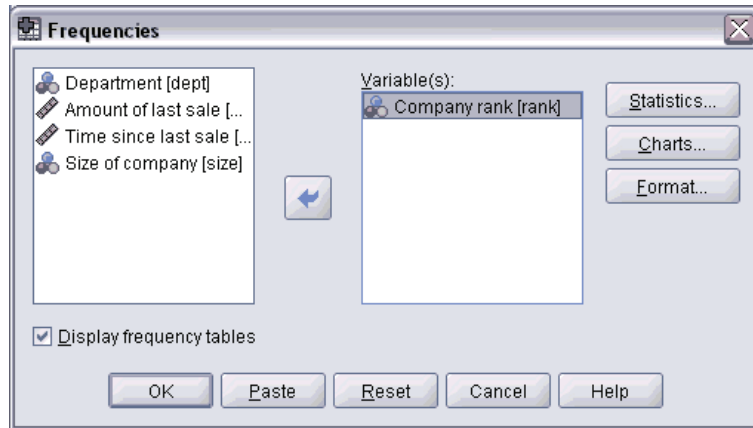
Data. Use numeric codes or strings to code categorical variables (nominal or ordinal level measurements).

Assumptions. The tabulations and percentages provide a useful description for data from any distribution, especially for variables with ordered or unordered categories. Most of the optional summary statistics, such as the mean and standard deviation, are based on normal theory and are appropriate for quantitative variables with symmetric distributions. Robust statistics, such as the median, quartiles, and percentiles, are appropriate for quantitative variables that may or may not meet the assumption of normality.

To Obtain Frequency Tables

- ▶ From the menu choose:
 - Analyze
 - Descriptive Statistics
 - Frequencies...

Figure 15-1
Frequencies main dialog box



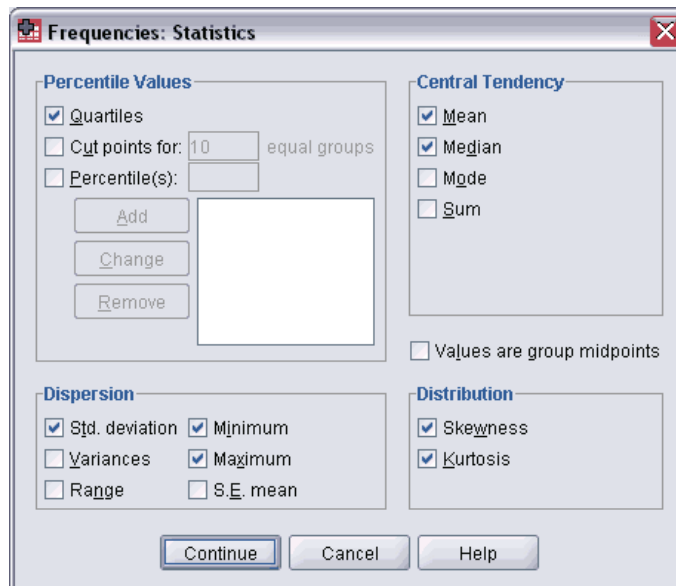
- ▶ Select one or more categorical or quantitative variables.

Optionally, you can:

- Click Statistics for descriptive statistics for quantitative variables.
- Click Charts for bar charts, pie charts, and histograms.
- Click Format for the order in which results are displayed.

Frequencies Statistics

Figure 15-2
Frequencies Statistics dialog box



Percentile Values. Values of a quantitative variable that divide the ordered data into groups so that a certain percentage is above and another percentage is below. Quartiles (the 25th, 50th, and 75th percentiles) divide the observations into four groups of equal size. If you want an equal number

of groups other than four, select Cut points for n equal groups. You can also specify individual percentiles (for example, the 95th percentile, the value below which 95% of the observations fall).

Central Tendency. Statistics that describe the location of the distribution include the mean, median, mode, and sum of all the values.

- **Mean.** A measure of central tendency. The arithmetic average, the sum divided by the number of cases.
- **Median.** The value above and below which half of the cases fall, the 50th percentile. If there is an even number of cases, the median is the average of the two middle cases when they are sorted in ascending or descending order. The median is a measure of central tendency not sensitive to outlying values (unlike the mean, which can be affected by a few extremely high or low values).
- **Mode.** The most frequently occurring value. If several values share the greatest frequency of occurrence, each of them is a mode. The Frequencies procedure reports only the smallest of such multiple modes.
- **Sum.** The sum or total of the values, across all cases with nonmissing values.

Dispersion. Statistics that measure the amount of variation or spread in the data include the standard deviation, variance, range, minimum, maximum, and standard error of the mean.

- **Std. deviation.** A measure of dispersion around the mean. In a normal distribution, 68% of cases fall within one standard deviation of the mean and 95% of cases fall within two standard deviations. For example, if the mean age is 45, with a standard deviation of 10, 95% of the cases would be between 25 and 65 in a normal distribution.
- **Variance.** A measure of dispersion around the mean, equal to the sum of squared deviations from the mean divided by one less than the number of cases. The variance is measured in units that are the square of those of the variable itself.
- **Range.** The difference between the largest and smallest values of a numeric variable, the maximum minus the minimum.
- **Minimum.** The smallest value of a numeric variable.
- **Maximum.** The largest value of a numeric variable.
- **S. E. mean.** A measure of how much the value of the mean may vary from sample to sample taken from the same distribution. It can be used to roughly compare the observed mean to a hypothesized value (that is, you can conclude the two values are different if the ratio of the difference to the standard error is less than -2 or greater than +2).

Distribution. Skewness and kurtosis are statistics that describe the shape and symmetry of the distribution. These statistics are displayed with their standard errors.

- **Skewness.** A measure of the asymmetry of a distribution. The normal distribution is symmetric and has a skewness value of 0. A distribution with a significant positive skewness has a long right tail. A distribution with a significant negative skewness has a long left tail. As a guideline, a skewness value more than twice its standard error is taken to indicate a departure from symmetry.
- **Kurtosis.** A measure of the extent to which observations cluster around a central point. For a normal distribution, the value of the kurtosis statistic is zero. Positive kurtosis indicates that the observations cluster more and have longer tails than those in the normal distribution, and negative kurtosis indicates that the observations cluster less and have shorter tails.

Values are group midpoints. If the values in your data are midpoints of groups (for example, ages of all people in their thirties are coded as 35), select this option to estimate the median and percentiles for the original, ungrouped data.

Frequencies Charts

Figure 15-3
Frequencies Charts dialog box

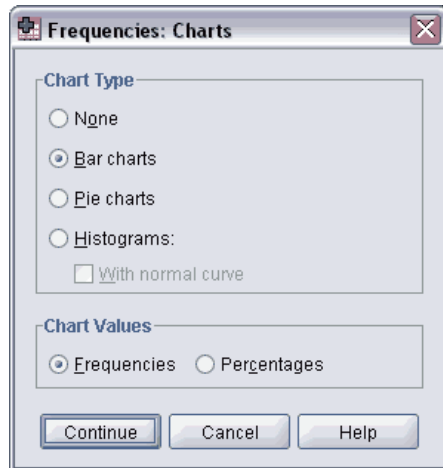
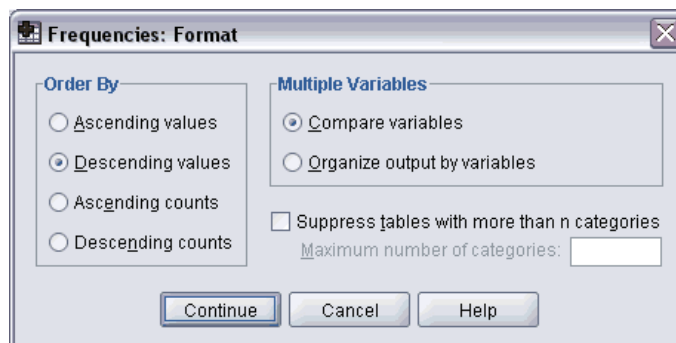


Chart Type. A pie chart displays the contribution of parts to a whole. Each slice of a pie chart corresponds to a group that is defined by a single grouping variable. A bar chart displays the count for each distinct value or category as a separate bar, allowing you to compare categories visually. A histogram also has bars, but they are plotted along an equal interval scale. The height of each bar is the count of values of a quantitative variable falling within the interval. A histogram shows the shape, center, and spread of the distribution. A normal curve superimposed on a histogram helps you judge whether the data are normally distributed.

Chart Values. For bar charts, the scale axis can be labeled by frequency counts or percentages.

Frequencies Format

Figure 15-4
Frequencies Format dialog box



Order by. The frequency table can be arranged according to the actual values in the data or according to the count (frequency of occurrence) of those values, and the table can be arranged in either ascending or descending order. However, if you request a histogram or percentiles, Frequencies assumes that the variable is quantitative and displays its values in ascending order.

Multiple Variables. If you produce statistics tables for multiple variables, you can either display all variables in a single table (Compare variables) or display a separate statistics table for each variable (Organize output by variables).

Suppress tables with more than n categories. This option prevents the display of tables with more than the specified number of values.

Descriptives

The Descriptives procedure displays univariate summary statistics for several variables in a single table and calculates standardized values (z scores). Variables can be ordered by the size of their means (in ascending or descending order), alphabetically, or by the order in which you select the variables (the default).

When z scores are saved, they are added to the data in the Data Editor and are available for charts, data listings, and analyses. When variables are recorded in different units (for example, gross domestic product per capita and percentage literate), a z -score transformation places variables on a common scale for easier visual comparison.

Example. If each case in your data contains the daily sales totals for each member of the sales staff (for example, one entry for Bob, one entry for Kim, and one entry for Brian) collected each day for several months, the Descriptives procedure can compute the average daily sales for each staff member and can order the results from highest average sales to lowest average sales.

Statistics. Sample size, mean, minimum, maximum, standard deviation, variance, range, sum, standard error of the mean, and kurtosis and skewness with their standard errors.

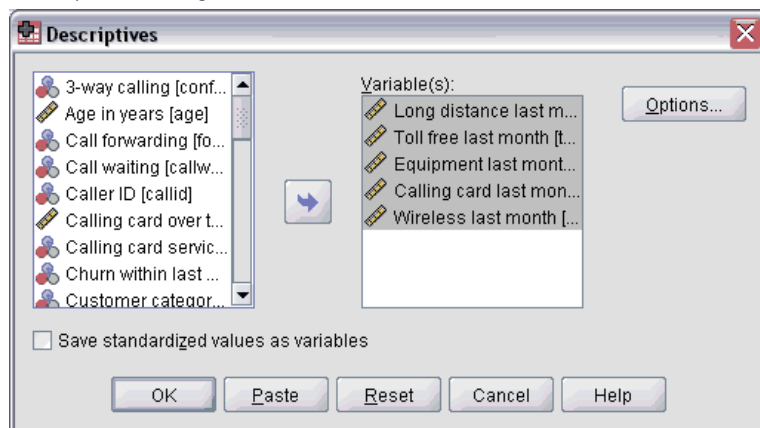
Data. Use numeric variables after you have screened them graphically for recording errors, outliers, and distributional anomalies. The Descriptives procedure is very efficient for large files (thousands of cases).

Assumptions. Most of the available statistics (including z scores) are based on normal theory and are appropriate for quantitative variables (interval- or ratio-level measurements) with symmetric distributions. Avoid variables with unordered categories or skewed distributions. The distribution of z scores has the same shape as that of the original data; therefore, calculating z scores is not a remedy for problem data.

To Obtain Descriptive Statistics

- ▶ From the menus choose:
 - Analyze
 - Descriptive Statistics
 - Descriptives...

Figure 16-1
Descriptives dialog box



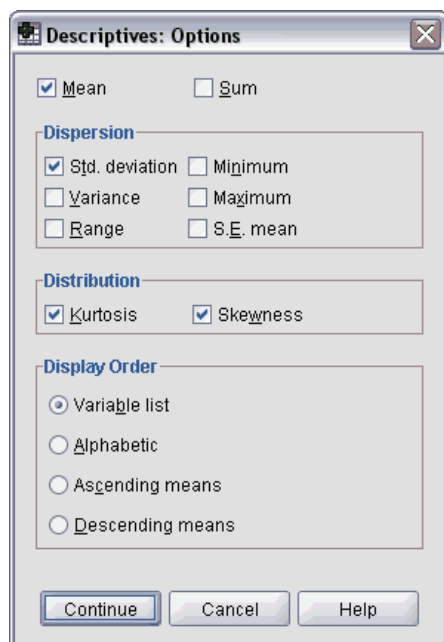
- ▶ Select one or more variables.

Optionally, you can:

- Select Save standardized values as variables to save z scores as new variables.
- Click Options for optional statistics and display order.

Descriptives Options

Figure 16-2
Descriptives Options dialog box



Mean and Sum. The mean, or arithmetic average, is displayed by default.

Dispersion. Statistics that measure the spread or variation in the data include the standard deviation, variance, range, minimum, maximum, and standard error of the mean.

- **Std. deviation.** A measure of dispersion around the mean. In a normal distribution, 68% of cases fall within one standard deviation of the mean and 95% of cases fall within two standard deviations. For example, if the mean age is 45, with a standard deviation of 10, 95% of the cases would be between 25 and 65 in a normal distribution.
- **Variance.** A measure of dispersion around the mean, equal to the sum of squared deviations from the mean divided by one less than the number of cases. The variance is measured in units that are the square of those of the variable itself.
- **Range.** The difference between the largest and smallest values of a numeric variable, the maximum minus the minimum.
- **Minimum.** The smallest value of a numeric variable.
- **Maximum.** The largest value of a numeric variable.
- **S.E. mean.** A measure of how much the value of the mean may vary from sample to sample taken from the same distribution. It can be used to roughly compare the observed mean to a hypothesized value (that is, you can conclude the two values are different if the ratio of the difference to the standard error is less than -2 or greater than +2).

Distribution. Kurtosis and skewness are statistics that characterize the shape and symmetry of the distribution. These statistics are displayed with their standard errors.

- **Kurtosis.** A measure of the extent to which observations cluster around a central point. For a normal distribution, the value of the kurtosis statistic is zero. Positive kurtosis indicates that the observations cluster more and have longer tails than those in the normal distribution, and negative kurtosis indicates that the observations cluster less and have shorter tails.
- **Skewness.** A measure of the asymmetry of a distribution. The normal distribution is symmetric and has a skewness value of 0. A distribution with a significant positive skewness has a long right tail. A distribution with a significant negative skewness has a long left tail. As a guideline, a skewness value more than twice its standard error is taken to indicate a departure from symmetry.

Display Order. By default, the variables are displayed in the order in which you selected them. Optionally, you can display variables alphabetically, by ascending means, or by descending means.

DESCRIPTIVES Command Additional Features

The command syntax language also allows you to:

- Save standardized scores (z scores) for some but not all variables (with the `VARIABLES` subcommand).
- Specify names for new variables that contain standardized scores (with the `VARIABLES` subcommand).
- Exclude from the analysis cases with missing values for any variable (with the `MISSING` subcommand).
- Sort the variables in the display by the value of any statistic, not just the mean (with the `SORT` subcommand).

See the *Command Syntax Reference* for complete syntax information.

Explore

The Explore procedure produces summary statistics and graphical displays, either for all of your cases or separately for groups of cases. There are many reasons for using the Explore procedure—data screening, outlier identification, description, assumption checking, and characterizing differences among subpopulations (groups of cases). Data screening may show that you have unusual values, extreme values, gaps in the data, or other peculiarities. Exploring the data can help to determine whether the statistical techniques that you are considering for data analysis are appropriate. The exploration may indicate that you need to transform the data if the technique requires a normal distribution. Or you may decide that you need nonparametric tests.

Example. Look at the distribution of maze-learning times for rats under four different reinforcement schedules. For each of the four groups, you can see if the distribution of times is approximately normal and whether the four variances are equal. You can also identify the cases with the five largest and five smallest times. The boxplots and stem-and-leaf plots graphically summarize the distribution of learning times for each of the groups.

Statistics and plots. Mean, median, 5% trimmed mean, standard error, variance, standard deviation, minimum, maximum, range, interquartile range, skewness and kurtosis and their standard errors, confidence interval for the mean (and specified confidence level), percentiles, Huber's M-estimator, Andrews' wave estimator, Hampel's redescending M-estimator, Tukey's biweight estimator, the five largest and five smallest values, the Kolmogorov-Smirnov statistic with a Lilliefors significance level for testing normality, and the Shapiro-Wilk statistic. Boxplots, stem-and-leaf plots, histograms, normality plots, and spread-versus-level plots with Levene tests and transformations.

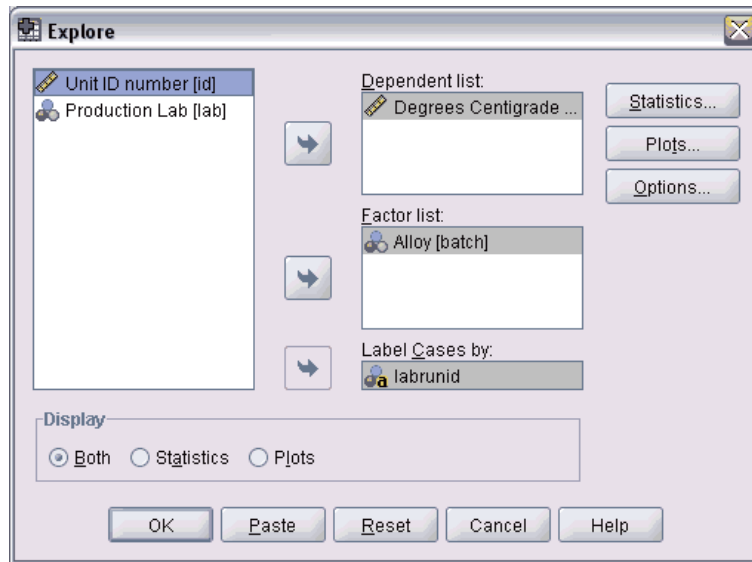
Data. The Explore procedure can be used for quantitative variables (interval- or ratio-level measurements). A factor variable (used to break the data into groups of cases) should have a reasonable number of distinct values (categories). These values may be short string or numeric. The case label variable, used to label outliers in boxplots, can be short string, long string (first 15 bytes), or numeric.

Assumptions. The distribution of your data does not have to be symmetric or normal.

To Explore Your Data

- ▶ From the menus choose:
 - Analyze
 - Descriptive Statistics
 - Explore...

Figure 17-1
Explore dialog box



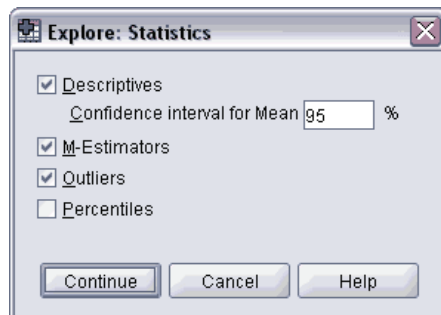
- ▶ Select one or more dependent variables.

Optionally, you can:

- Select one or more factor variables, whose values will define groups of cases.
- Select an identification variable to label cases.
- Click Statistics for robust estimators, outliers, percentiles, and frequency tables.
- Click Plots for histograms, normal probability plots and tests, and spread-versus-level plots with Levene's statistics.
- Click Options for the treatment of missing values.

Explore Statistics

Figure 17-2
Explore Statistics dialog box



Descriptives. These measures of central tendency and dispersion are displayed by default. Measures of central tendency indicate the location of the distribution; they include the mean, median, and 5% trimmed mean. Measures of dispersion show the dissimilarity of the values; these

include standard error, variance, standard deviation, minimum, maximum, range, and interquartile range. The descriptive statistics also include measures of the shape of the distribution; skewness and kurtosis are displayed with their standard errors. The 95% level confidence interval for the mean is also displayed; you can specify a different confidence level.

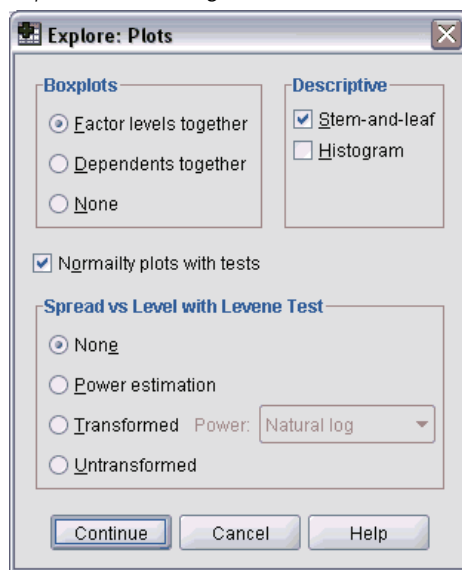
M-estimators. Robust alternatives to the sample mean and median for estimating the location. The estimators calculated differ in the weights they apply to cases. Huber's M-estimator, Andrews' wave estimator, Hampel's redescending M-estimator, and Tukey's biweight estimator are displayed.

Outliers. Displays the five largest and five smallest values with case labels.

Percentiles. Displays the values for the 5th, 10th, 25th, 50th, 75th, 90th, and 95th percentiles.

Explore Plots

Figure 17-3
Explore Plots dialog box



Boxplots. These alternatives control the display of boxplots when you have more than one dependent variable. Factor levels together generates a separate display for each dependent variable. Within a display, boxplots are shown for each of the groups defined by a factor variable. Dependents together generates a separate display for each group defined by a factor variable. Within a display, boxplots are shown side by side for each dependent variable. This display is particularly useful when the different variables represent a single characteristic measured at different times.

Descriptive. The Descriptive group allows you to choose stem-and-leaf plots and histograms.

Normality plots with tests. Displays normal probability and detrended normal probability plots. The Kolmogorov-Smirnov statistic, with a Lilliefors significance level for testing normality, is displayed. If non-integer weights are specified, the Shapiro-Wilk statistic is calculated when the

weighted sample size lies between 3 and 50. For no weights or integer weights, the statistic is calculated when the weighted sample size lies between 3 and 5,000.

Spread vs. Level with Levene Test. Controls data transformation for spread-versus-level plots. For all spread-versus-level plots, the slope of the regression line and Levene's robust tests for homogeneity of variance are displayed. If you select a transformation, Levene's tests are based on the transformed data. If no factor variable is selected, spread-versus-level plots are not produced. Power estimation produces a plot of the natural logs of the interquartile ranges against the natural logs of the medians for all cells, as well as an estimate of the power transformation for achieving equal variances in the cells. A spread-versus-level plot helps to determine the power for a transformation to stabilize (make more equal) variances across groups. Transformed allows you to select one of the power alternatives, perhaps following the recommendation from power estimation, and produces plots of transformed data. The interquartile range and median of the transformed data are plotted. Untransformed produces plots of the raw data. This is equivalent to a transformation with a power of 1.

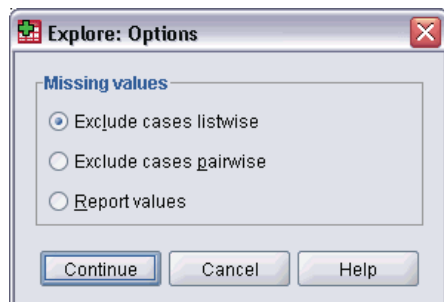
Explore Power Transformations

These are the power transformations for spread-versus-level plots. To transform data, you must select a power for the transformation. You can choose one of the following alternatives:

- **Natural log.** Natural log transformation. This is the default.
- **1/square root.** For each data value, the reciprocal of the square root is calculated.
- **Reciprocal.** The reciprocal of each data value is calculated.
- **Square root.** The square root of each data value is calculated.
- **Square.** Each data value is squared.
- **Cube.** Each data value is cubed.

Explore Options

Figure 17-4
Explore Options dialog box



Missing Values. Controls the treatment of missing values.

- **Exclude cases listwise.** Cases with missing values for any dependent or factor variable are excluded from all analyses. This is the default.

- **Exclude cases pairwise.** Cases with no missing values for variables in a group (cell) are included in the analysis of that group. The case may have missing values for variables used in other groups.
- **Report values.** Missing values for factor variables are treated as a separate category. All output is produced for this additional category. Frequency tables include categories for missing values. Missing values for a factor variable are included but labeled as missing.

EXAMINE Command Additional Features

The Explore procedure uses `EXAMINE` command syntax. The command syntax language also allows you to:

- Request total output and plots in addition to output and plots for groups defined by the factor variables (with the `TOTAL` subcommand).
- Specify a common scale for a group of boxplots (with the `SCALE` subcommand).
- Specify interactions of the factor variables (with the `VARIABLES` subcommand).
- Specify percentiles other than the defaults (with the `PERCENTILES` subcommand).
- Calculate percentiles according to any of five methods (with the `PERCENTILES` subcommand).
- Specify any power transformation for spread-versus-level plots (with the `PLOT` subcommand).
- Specify the number of extreme values to be displayed (with the `STATISTICS` subcommand).
- Specify parameters for the M-estimators, robust estimators of location (with the `MESTIMATORS` subcommand).

See the *Command Syntax Reference* for complete syntax information.

Crosstabs

The Crosstabs procedure forms two-way and multiway tables and provides a variety of tests and measures of association for two-way tables. The structure of the table and whether categories are ordered determine what test or measure to use.

Crosstabs' statistics and measures of association are computed for two-way tables only. If you specify a row, a column, and a layer factor (control variable), the Crosstabs procedure forms one panel of associated statistics and measures for each value of the layer factor (or a combination of values for two or more control variables). For example, if *gender* is a layer factor for a table of *married* (yes, no) against *life* (is life exciting, routine, or dull), the results for a two-way table for the females are computed separately from those for the males and printed as panels following one another.

Example. Are customers from small companies more likely to be profitable in sales of services (for example, training and consulting) than those from larger companies? From a crosstabulation, you might learn that the majority of small companies (fewer than 500 employees) yield high service profits, while the majority of large companies (more than 2,500 employees) yield low service profits.

Statistics and measures of association. Pearson chi-square, likelihood-ratio chi-square, linear-by-linear association test, Fisher's exact test, Yates' corrected chi-square, Pearson's r , Spearman's rho, contingency coefficient, phi, Cramér's V , symmetric and asymmetric lambdas, Goodman and Kruskal's tau, uncertainty coefficient, gamma, Somers' d , Kendall's tau- b , Kendall's tau- c , eta coefficient, Cohen's kappa, relative risk estimate, odds ratio, McNemar test, and Cochran's and Mantel-Haenszel statistics.

Data. To define the categories of each table variable, use values of a numeric or string (eight or fewer bytes) variable. For example, for *gender*, you could code the data as 1 and 2 or as *male* and *female*.

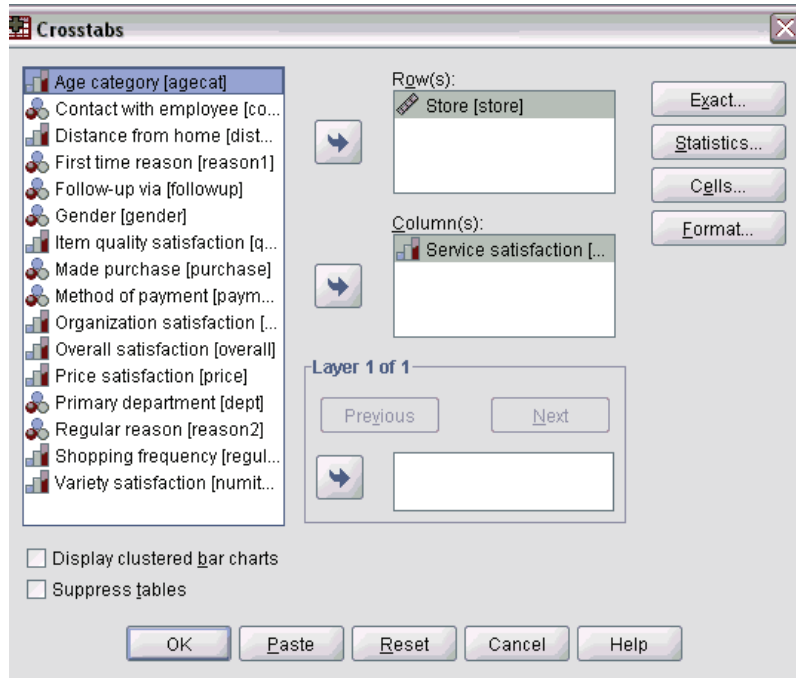
Assumptions. Some statistics and measures assume ordered categories (ordinal data) or quantitative values (interval or ratio data), as discussed in the section on statistics. Others are valid when the table variables have unordered categories (nominal data). For the chi-square-based statistics (phi, Cramér's V , and contingency coefficient), the data should be a random sample from a multinomial distribution.

Note: Ordinal variables can be either numeric codes that represent categories (for example, 1 = *low*, 2 = *medium*, 3 = *high*) or string values. However, the alphabetic order of string values is assumed to reflect the true order of the categories. For example, for a string variable with the values of *low*, *medium*, *high*, the order of the categories is interpreted as *high*, *low*, *medium*—which is not the correct order. In general, it is more reliable to use numeric codes to represent ordinal data.

To Obtain Crosstabulations

- ▶ From the menu choose:
 - Analyze
 - Descriptive Statistics
 - Crosstabs...

Figure 18-1
Crosstabs dialog box



- ▶ Select one or more row variables and one or more column variables.

Optionally, you can:

- Select one or more control variables.
- Click Statistics for tests and measures of association for two-way tables or subtables.
- Click Cells for observed and expected values, percentages, and residuals.
- Click Format for controlling the order of categories.

Crosstabs Layers

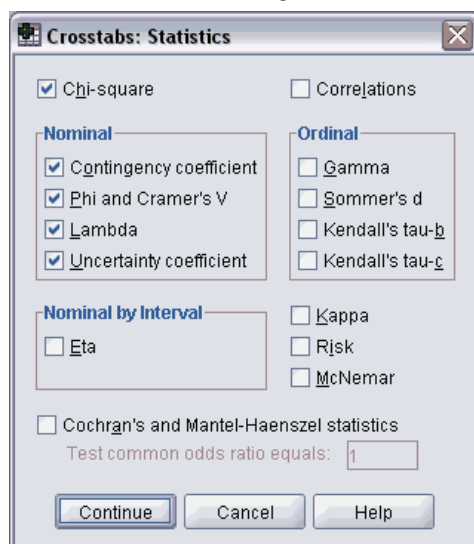
If you select one or more layer variables, a separate crosstabulation is produced for each category of each layer variable (control variable). For example, if you have one row variable, one column variable, and one layer variable with two categories, you get a two-way table for each category of the layer variable. To make another layer of control variables, click Next. Subtables are produced for each combination of categories for each first-layer variable, each second-layer variable, and so on. If statistics and measures of association are requested, they apply to two-way subtables only.

Crosstabs Clustered Bar Charts

Display clustered bar charts. A clustered bar chart helps summarize your data for groups of cases. There is one cluster of bars for each value of the variable you specified under Rows. The variable that defines the bars within each cluster is the variable you specified under Columns. There is one set of differently colored or patterned bars for each value of this variable. If you specify more than one variable under Columns or Rows, a clustered bar chart is produced for each combination of two variables.

Crosstabs Statistics

Figure 18-2
Crosstabs Statistics dialog box



Chi-square. For tables with two rows and two columns, select Chi-square to calculate the Pearson chi-square, the likelihood-ratio chi-square, Fisher's exact test, and Yates' corrected chi-square (continuity correction). For 2×2 tables, Fisher's exact test is computed when a table that does not result from missing rows or columns in a larger table has a cell with an expected frequency of less than 5. Yates' corrected chi-square is computed for all other 2×2 tables. For tables with any number of rows and columns, select Chi-square to calculate the Pearson chi-square and the likelihood-ratio chi-square. When both table variables are quantitative, Chi-square yields the linear-by-linear association test.

Correlations. For tables in which both rows and columns contain ordered values, Correlations yields Spearman's correlation coefficient, rho (numeric data only). Spearman's rho is a measure of association between rank orders. When both table variables (factors) are quantitative, Correlations yields the Pearson correlation coefficient, r , a measure of linear association between the variables.

Nominal. For nominal data (no intrinsic order, such as Catholic, Protestant, and Jewish), you can select Contingency coefficient, Phi (coefficient) and Cramér's V, Lambda (symmetric and asymmetric lambdas and Goodman and Kruskal's tau), and Uncertainty coefficient.

- **Contingency coefficient.** A measure of association based on chi-square. The value ranges between 0 and 1, with 0 indicating no association between the row and column variables and values close to 1 indicating a high degree of association between the variables. The maximum value possible depends on the number of rows and columns in a table.
- **Phi and Cramer's V.** Phi is a chi-square-based measure of association that involves dividing the chi-square statistic by the sample size and taking the square root of the result. Cramer's V is a measure of association based on chi-square.
- **Lambda.** A measure of association that reflects the proportional reduction in error when values of the independent variable are used to predict values of the dependent variable. A value of 1 means that the independent variable perfectly predicts the dependent variable. A value of 0 means that the independent variable is no help in predicting the dependent variable.
- **Uncertainty coefficient.** A measure of association that indicates the proportional reduction in error when values of one variable are used to predict values of the other variable. For example, a value of 0.83 indicates that knowledge of one variable reduces error in predicting values of the other variable by 83%. The program calculates both symmetric and asymmetric versions of the uncertainty coefficient.

Ordinal. For tables in which both rows and columns contain ordered values, select Gamma (zero-order for 2-way tables and conditional for 3-way to 10-way tables), Kendall's tau-b, and Kendall's tau-c. For predicting column categories from row categories, select Somers' d.

- **Gamma.** A symmetric measure of association between two ordinal variables that ranges between -1 and 1. Values close to an absolute value of 1 indicate a strong relationship between the two variables. Values close to 0 indicate little or no relationship. For 2-way tables, zero-order gammas are displayed. For 3-way to n-way tables, conditional gammas are displayed.
- **Somers' d.** A measure of association between two ordinal variables that ranges from -1 to 1. Values close to an absolute value of 1 indicate a strong relationship between the two variables, and values close to 0 indicate little or no relationship between the variables. Somers' d is an asymmetric extension of gamma that differs only in the inclusion of the number of pairs not tied on the independent variable. A symmetric version of this statistic is also calculated.
- **Kendall's tau-b.** A nonparametric measure of correlation for ordinal or ranked variables that take ties into account. The sign of the coefficient indicates the direction of the relationship, and its absolute value indicates the strength, with larger absolute values indicating stronger relationships. Possible values range from -1 to 1, but a value of -1 or +1 can be obtained only from square tables.
- **Kendall's tau-c.** A nonparametric measure of association for ordinal variables that ignores ties. The sign of the coefficient indicates the direction of the relationship, and its absolute value indicates the strength, with larger absolute values indicating stronger relationships. Possible values range from -1 to 1, but a value of -1 or +1 can be obtained only from square tables.

Nominal by Interval. When one variable is categorical and the other is quantitative, select Eta. The categorical variable must be coded numerically.

- **Eta.** A measure of association that ranges from 0 to 1, with 0 indicating no association between the row and column variables and values close to 1 indicating a high degree of association. Eta is appropriate for a dependent variable measured on an interval scale (for example, income) and an independent variable with a limited number of categories (for example,

gender). Two eta values are computed: one treats the row variable as the interval variable, and the other treats the column variable as the interval variable.

Kappa. Cohen's kappa measures the agreement between the evaluations of two raters when both are rating the same object. A value of 1 indicates perfect agreement. A value of 0 indicates that agreement is no better than chance. Kappa is available only for tables in which both variables use the same category values and both variables have the same number of categories.

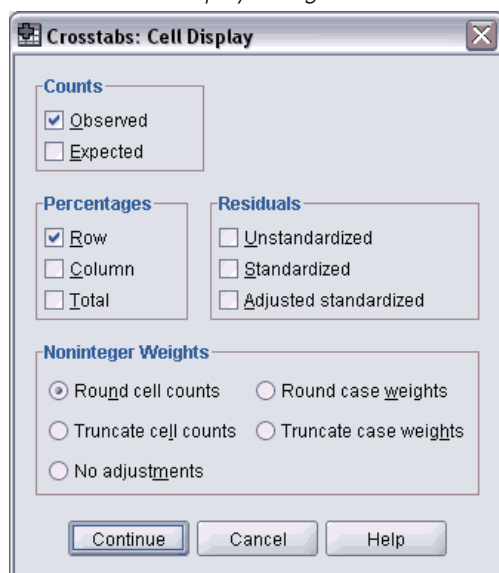
Risk. For 2 x 2 tables, a measure of the strength of the association between the presence of a factor and the occurrence of an event. If the confidence interval for the statistic includes a value of 1, you cannot assume that the factor is associated with the event. The odds ratio can be used as an estimate or relative risk when the occurrence of the factor is rare.

McNemar. A nonparametric test for two related dichotomous variables. Tests for changes in responses using the chi-square distribution. Useful for detecting changes in responses due to experimental intervention in "before-and-after" designs. For larger square tables, the McNemar-Bowker test of symmetry is reported.

Cochran's and Mantel-Haenszel statistics. Cochran's and Mantel-Haenszel statistics can be used to test for independence between a dichotomous factor variable and a dichotomous response variable, conditional upon covariate patterns defined by one or more layer (control) variables. Note that while other statistics are computed layer by layer, the Cochran's and Mantel-Haenszel statistics are computed once for all layers.

Crosstabs Cell Display

Figure 18-3
Crosstabs Cell Display dialog box



To help you uncover patterns in the data that contribute to a significant chi-square test, the Crosstabs procedure displays expected frequencies and three types of residuals (deviates) that measure the difference between observed and expected frequencies. Each cell of the table can contain any combination of counts, percentages, and residuals selected.

Counts. The number of cases actually observed and the number of cases expected if the row and column variables are independent of each other.

Percentages. The percentages can add up across the rows or down the columns. The percentages of the total number of cases represented in the table (one layer) are also available.

Residuals. Raw unstandardized residuals give the difference between the observed and expected values. Standardized and adjusted standardized residuals are also available.

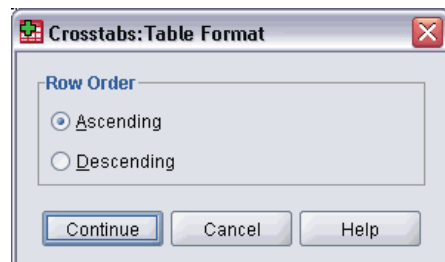
- **Unstandardized.** The difference between an observed value and the expected value. The expected value is the number of cases you would expect in the cell if there were no relationship between the two variables. A positive residual indicates that there are more cases in the cell than there would be if the row and column variables were independent.
- **Standardized.** The residual divided by an estimate of its standard deviation. Standardized residuals, which are also known as Pearson residuals, have a mean of 0 and a standard deviation of 1.
- **Adjusted standardized.** The residual for a cell (observed minus expected value) divided by an estimate of its standard error. The resulting standardized residual is expressed in standard deviation units above or below the mean.

Noninteger Weights. Cell counts are normally integer values, since they represent the number of cases in each cell. But if the data file is currently weighted by a weight variable with fractional values (for example, 1.25), cell counts can also be fractional values. You can truncate or round either before or after calculating the cell counts or use fractional cell counts for both table display and statistical calculations.

- **Round cell counts.** Case weights are used as is but the accumulated weights in the cells are rounded before computing any statistics.
- **Truncate cell counts.** Case weights are used as is but the accumulated weights in the cells are truncated before computing any statistics.
- **Round case weights.** Case weights are rounded before use.
- **Truncate case weights.** Case weights are truncated before use.
- **No adjustments.** Case weights are used as is and fractional cell counts are used. However, when Exact Statistics (available only with the Exact Tests option) are requested, the accumulated weights in the cells are either truncated or rounded before computing the Exact test statistics.

Crosstabs Table Format

Figure 18-4
Crosstabs Table Format dialog box



You can arrange rows in ascending or descending order of the values of the row variable.

Summarize

The Summarize procedure calculates subgroup statistics for variables within categories of one or more grouping variables. All levels of the grouping variable are crosstabulated. You can choose the order in which the statistics are displayed. Summary statistics for each variable across all categories are also displayed. Data values in each category can be listed or suppressed. With large datasets, you can choose to list only the first n cases.

Example. What is the average product sales amount by region and customer industry? You might discover that the average sales amount is slightly higher in the western region than in other regions, with corporate customers in the western region yielding the highest average sales amount.

Statistics. Sum, number of cases, mean, median, grouped median, standard error of the mean, minimum, maximum, range, variable value of the first category of the grouping variable, variable value of the last category of the grouping variable, standard deviation, variance, kurtosis, standard error of kurtosis, skewness, standard error of skewness, percentage of total sum, percentage of total N , percentage of sum in, percentage of N in, geometric mean, and harmonic mean.

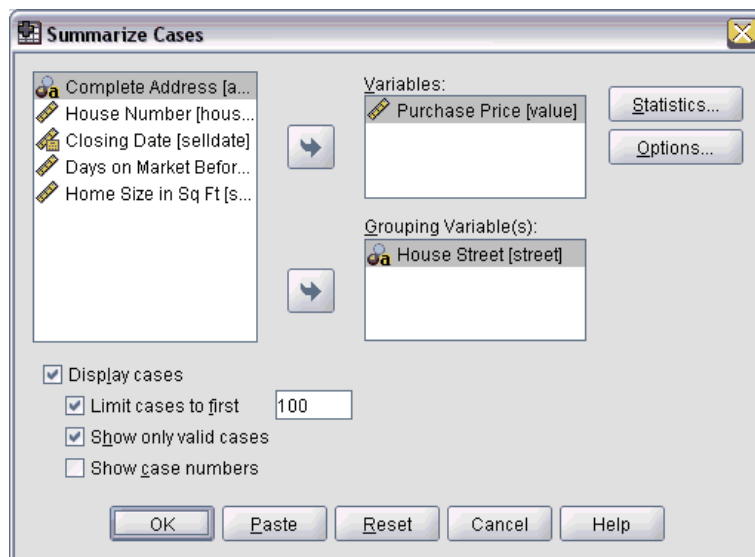
Data. Grouping variables are categorical variables whose values can be numeric or string. The number of categories should be reasonably small. The other variables should be able to be ranked.

Assumptions. Some of the optional subgroup statistics, such as the mean and standard deviation, are based on normal theory and are appropriate for quantitative variables with symmetric distributions. Robust statistics, such as the median and the range, are appropriate for quantitative variables that may or may not meet the assumption of normality.

To Obtain Case Summaries

- ▶ From the menus choose:
 - Analyze
 - Reports
 - Case Summaries...

Figure 19-1
Summarize Cases dialog box



- ▶ Select one or more variables.

Optionally, you can:

- Select one or more grouping variables to divide your data into subgroups.
- Click Options to change the output title, add a caption below the output, or exclude cases with missing values.
- Click Statistics for optional statistics.
- Select Display cases to list the cases in each subgroup. By default, the system lists only the first 100 cases in your file. You can raise or lower the value for Limit cases to first n or deselect that item to list all cases.

Summarize Options

Figure 19-2
Options dialog box

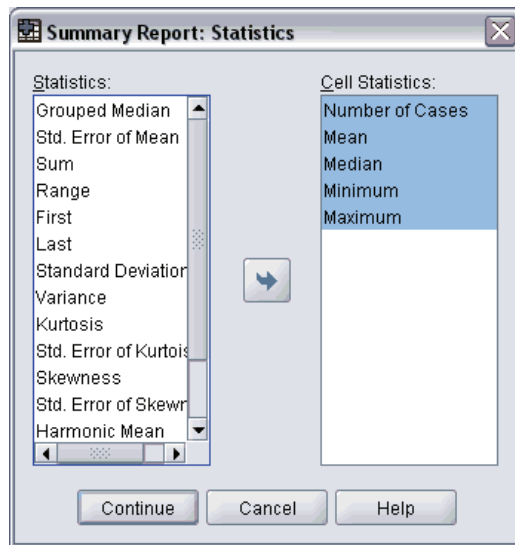


Summarize allows you to change the title of your output or add a caption that will appear below the output table. You can control line wrapping in titles and captions by typing `\n` wherever you want to insert a line break in the text.

You can also choose to display or suppress subheadings for totals and to include or exclude cases with missing values for any of the variables used in any of the analyses. Often it is desirable to denote missing cases in output with a period or an asterisk. Enter a character, phrase, or code that you would like to have appear when a value is missing; otherwise, no special treatment is applied to missing cases in the output.

Summarize Statistics

Figure 19-3
Summary Report Statistics dialog box



You can choose one or more of the following subgroup statistics for the variables within each category of each grouping variable: sum, number of cases, mean, median, grouped median, standard error of the mean, minimum, maximum, range, variable value of the first category of the grouping variable, variable value of the last category of the grouping variable, standard deviation, variance, kurtosis, standard error of kurtosis, skewness, standard error of skewness, percentage of total sum, percentage of total N , percentage of sum in, percentage of N in, geometric mean, harmonic mean. The order in which the statistics appear in the Cell Statistics list is the order in which they will be displayed in the output. Summary statistics are also displayed for each variable across all categories.

First. Displays the first data value encountered in the data file.

Geometric Mean. The n th root of the product of the data values, where n represents the number of cases.

Grouped Median. Median that is calculated for data that is coded into groups. For example, with age data, if each value in the 30s is coded 35, each value in the 40s is coded 45, and so on, the grouped median is the median calculated from the coded data.

Harmonic Mean. Used to estimate an average group size when the sample sizes in the groups are not equal. The harmonic mean is the total number of samples divided by the sum of the reciprocals of the sample sizes.

Kurtosis. A measure of the extent to which observations cluster around a central point. For a normal distribution, the value of the kurtosis statistic is zero. Positive kurtosis indicates that the observations cluster more and have longer tails than those in the normal distribution, and negative kurtosis indicates that the observations cluster less and have shorter tails.

Last. Displays the last data value encountered in the data file.

Maximum. The largest value of a numeric variable.

Mean. A measure of central tendency. The arithmetic average, the sum divided by the number of cases.

Median. The value above and below which half of the cases fall, the 50th percentile. If there is an even number of cases, the median is the average of the two middle cases when they are sorted in ascending or descending order. The median is a measure of central tendency not sensitive to outlying values (unlike the mean, which can be affected by a few extremely high or low values).

Minimum. The smallest value of a numeric variable.

N. The number of cases (observations or records).

Percent of Total N. Percentage of the total number of cases in each category.

Percent of Total Sum. Percentage of the total sum in each category.

Range. The difference between the largest and smallest values of a numeric variable, the maximum minus the minimum.

Skewness. A measure of the asymmetry of a distribution. The normal distribution is symmetric and has a skewness value of 0. A distribution with a significant positive skewness has a long right tail. A distribution with a significant negative skewness has a long left tail. As a guideline, a skewness value more than twice its standard error is taken to indicate a departure from symmetry.

Standard Error of Kurtosis. The ratio of kurtosis to its standard error can be used as a test of normality (that is, you can reject normality if the ratio is less than -2 or greater than +2). A large positive value for kurtosis indicates that the tails of the distribution are longer than those of a normal distribution; a negative value for kurtosis indicates shorter tails (becoming like those of a box-shaped uniform distribution).

Standard Error of Skewness. The ratio of skewness to its standard error can be used as a test of normality (that is, you can reject normality if the ratio is less than -2 or greater than +2). A large positive value for skewness indicates a long right tail; an extreme negative value indicates a long left tail.

Sum. The sum or total of the values, across all cases with nonmissing values.

Variance. A measure of dispersion around the mean, equal to the sum of squared deviations from the mean divided by one less than the number of cases. The variance is measured in units that are the square of those of the variable itself.

Means

The Means procedure calculates subgroup means and related univariate statistics for dependent variables within categories of one or more independent variables. Optionally, you can obtain a one-way analysis of variance, eta, and tests for linearity.

Example. Measure the average amount of fat absorbed by three different types of cooking oil, and perform a one-way analysis of variance to see whether the means differ.

Statistics. Sum, number of cases, mean, median, grouped median, standard error of the mean, minimum, maximum, range, variable value of the first category of the grouping variable, variable value of the last category of the grouping variable, standard deviation, variance, kurtosis, standard error of kurtosis, skewness, standard error of skewness, percentage of total sum, percentage of total N , percentage of sum in, percentage of N in, geometric mean, and harmonic mean. Options include analysis of variance, eta, eta squared, and tests for linearity R and R^2 .

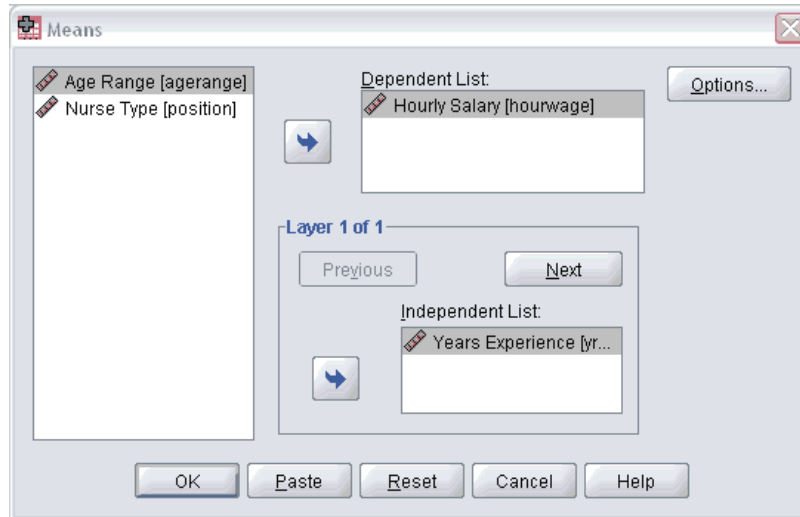
Data. The dependent variables are quantitative, and the independent variables are categorical. The values of categorical variables can be numeric or string.

Assumptions. Some of the optional subgroup statistics, such as the mean and standard deviation, are based on normal theory and are appropriate for quantitative variables with symmetric distributions. Robust statistics, such as the median, are appropriate for quantitative variables that may or may not meet the assumption of normality. Analysis of variance is robust to departures from normality, but the data in each cell should be symmetric. Analysis of variance also assumes that the groups come from populations with equal variances. To test this assumption, use Levene's homogeneity-of-variance test, available in the One-Way ANOVA procedure.

To Obtain Subgroup Means

- ▶ From the menus choose:
 - Analyze
 - Compare Means
 - Means...

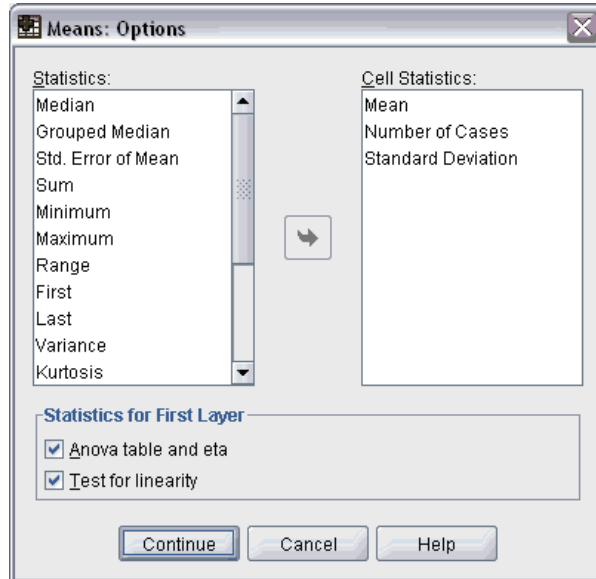
Figure 20-1
Means dialog box



- ▶ Select one or more dependent variables.
- ▶ Use one of the following methods to select categorical independent variables:
 - Select one or more independent variables. Separate results are displayed for each independent variable.
 - Select one or more layers of independent variables. Each layer further subdivides the sample. If you have one independent variable in Layer 1 and one independent variable in Layer 2, the results are displayed in one crossed table, as opposed to separate tables for each independent variable.
- ▶ Optionally, click Options for optional statistics, an analysis of variance table, eta, eta squared, R , and R^2 .

Means Options

Figure 20-2
Means Options dialog box



You can choose one or more of the following subgroup statistics for the variables within each category of each grouping variable: sum, number of cases, mean, median, grouped median, standard error of the mean, minimum, maximum, range, variable value of the first category of the grouping variable, variable value of the last category of the grouping variable, standard deviation, variance, kurtosis, standard error of kurtosis, skewness, standard error of skewness, percentage of total sum, percentage of total N , percentage of sum in, percentage of N in, geometric mean, and harmonic mean. You can change the order in which the subgroup statistics appear. The order in which the statistics appear in the Cell Statistics list is the order in which they are displayed in the output. Summary statistics are also displayed for each variable across all categories.

First. Displays the first data value encountered in the data file.

Geometric Mean. The n th root of the product of the data values, where n represents the number of cases.

Grouped Median. Median that is calculated for data that is coded into groups. For example, with age data, if each value in the 30s is coded 35, each value in the 40s is coded 45, and so on, the grouped median is the median calculated from the coded data.

Harmonic Mean. Used to estimate an average group size when the sample sizes in the groups are not equal. The harmonic mean is the total number of samples divided by the sum of the reciprocals of the sample sizes.

Kurtosis. A measure of the extent to which observations cluster around a central point. For a normal distribution, the value of the kurtosis statistic is zero. Positive kurtosis indicates that the observations cluster more and have longer tails than those in the normal distribution, and negative kurtosis indicates that the observations cluster less and have shorter tails.

Last. Displays the last data value encountered in the data file.

Maximum. The largest value of a numeric variable.

Mean. A measure of central tendency. The arithmetic average, the sum divided by the number of cases.

Median. The value above and below which half of the cases fall, the 50th percentile. If there is an even number of cases, the median is the average of the two middle cases when they are sorted in ascending or descending order. The median is a measure of central tendency not sensitive to outlying values (unlike the mean, which can be affected by a few extremely high or low values).

Minimum. The smallest value of a numeric variable.

N. The number of cases (observations or records).

Percent of total N. Percentage of the total number of cases in each category.

Percent of total sum. Percentage of the total sum in each category.

Range. The difference between the largest and smallest values of a numeric variable, the maximum minus the minimum.

Skewness. A measure of the asymmetry of a distribution. The normal distribution is symmetric and has a skewness value of 0. A distribution with a significant positive skewness has a long right tail. A distribution with a significant negative skewness has a long left tail. As a guideline, a skewness value more than twice its standard error is taken to indicate a departure from symmetry.

Standard Error of Kurtosis. The ratio of kurtosis to its standard error can be used as a test of normality (that is, you can reject normality if the ratio is less than -2 or greater than +2). A large positive value for kurtosis indicates that the tails of the distribution are longer than those of a normal distribution; a negative value for kurtosis indicates shorter tails (becoming like those of a box-shaped uniform distribution).

Standard Error of Skewness. The ratio of skewness to its standard error can be used as a test of normality (that is, you can reject normality if the ratio is less than -2 or greater than +2). A large positive value for skewness indicates a long right tail; an extreme negative value indicates a long left tail.

Sum. The sum or total of the values, across all cases with nonmissing values.

Variance. A measure of dispersion around the mean, equal to the sum of squared deviations from the mean divided by one less than the number of cases. The variance is measured in units that are the square of those of the variable itself.

Statistics for First Layer

Anova table and eta. Displays a one-way analysis-of-variance table and calculates eta and eta-squared (measures of association) for each independent variable in the first layer.

Test for linearity. Calculates the sum of squares, degrees of freedom, and mean square associated with linear and nonlinear components, as well as the F ratio, R and R-squared. Linearity is not calculated if the independent variable is a short string.

OLAP Cubes

The OLAP (Online Analytical Processing) Cubes procedure calculates totals, means, and other univariate statistics for continuous summary variables within categories of one or more categorical grouping variables. A separate layer in the table is created for each category of each grouping variable.

Example. Total and average sales for different regions and product lines within regions.

Statistics. Sum, number of cases, mean, median, grouped median, standard error of the mean, minimum, maximum, range, variable value of the first category of the grouping variable, variable value of the last category of the grouping variable, standard deviation, variance, kurtosis, standard error of kurtosis, skewness, standard error of skewness, percentage of total cases, percentage of total sum, percentage of total cases within grouping variables, percentage of total sum within grouping variables, geometric mean, and harmonic mean.

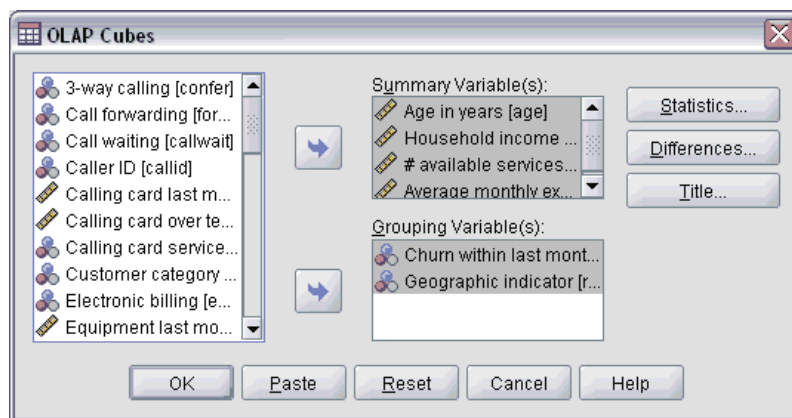
Data. The summary variables are quantitative (continuous variables measured on an interval or ratio scale), and the grouping variables are categorical. The values of categorical variables can be numeric or string.

Assumptions. Some of the optional subgroup statistics, such as the mean and standard deviation, are based on normal theory and are appropriate for quantitative variables with symmetric distributions. Robust statistics, such as the median and range, are appropriate for quantitative variables that may or may not meet the assumption of normality.

To Obtain OLAP Cubes

- ▶ From the menus choose:
 - Analyze
 - Reports
 - OLAP Cubes...

Figure 21-1
OLAP Cubes dialog box



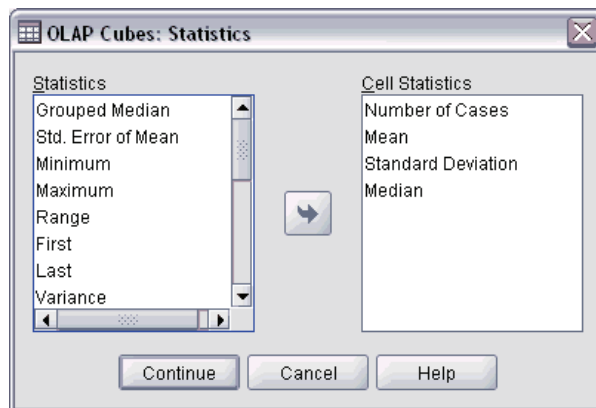
- ▶ Select one or more continuous summary variables.
- ▶ Select one or more categorical grouping variables.

Optionally:

- Select different summary statistics (click Statistics). You must select one or more grouping variables before you can select summary statistics.
- Calculate differences between pairs of variables and pairs of groups that are defined by a grouping variable (click Differences).
- Create custom table titles (click Title).

OLAP Cubes Statistics

Figure 21-2
OLAP Cubes Statistics dialog box



You can choose one or more of the following subgroup statistics for the summary variables within each category of each grouping variable: sum, number of cases, mean, median, grouped median, standard error of the mean, minimum, maximum, range, variable value of the first category of the grouping variable, variable value of the last category of the grouping variable, standard deviation, variance, kurtosis, standard error of kurtosis, skewness, standard error of skewness, percentage

of total cases, percentage of total sum, percentage of total cases within grouping variables, percentage of total sum within grouping variables, geometric mean, and harmonic mean.

You can change the order in which the subgroup statistics appear. The order in which the statistics appear in the Cell Statistics list is the order in which they are displayed in the output. Summary statistics are also displayed for each variable across all categories.

First. Displays the first data value encountered in the data file.

Geometric Mean. The n th root of the product of the data values, where n represents the number of cases.

Grouped Median. Median that is calculated for data that is coded into groups. For example, with age data, if each value in the 30s is coded 35, each value in the 40s is coded 45, and so on, the grouped median is the median calculated from the coded data.

Harmonic Mean. Used to estimate an average group size when the sample sizes in the groups are not equal. The harmonic mean is the total number of samples divided by the sum of the reciprocals of the sample sizes.

Kurtosis. A measure of the extent to which observations cluster around a central point. For a normal distribution, the value of the kurtosis statistic is zero. Positive kurtosis indicates that the observations cluster more and have longer tails than those in the normal distribution, and negative kurtosis indicates that the observations cluster less and have shorter tails.

Last. Displays the last data value encountered in the data file.

Maximum. The largest value of a numeric variable.

Mean. A measure of central tendency. The arithmetic average, the sum divided by the number of cases.

Median. The value above and below which half of the cases fall, the 50th percentile. If there is an even number of cases, the median is the average of the two middle cases when they are sorted in ascending or descending order. The median is a measure of central tendency not sensitive to outlying values (unlike the mean, which can be affected by a few extremely high or low values).

Minimum. The smallest value of a numeric variable.

N. The number of cases (observations or records).

Percent of N in. Percentage of the number of cases for the specified grouping variable within categories of other grouping variables. If you only have one grouping variable, this value is identical to percentage of total number of cases.

Percent of Sum in. Percentage of the sum for the specified grouping variable within categories of other grouping variables. If you only have one grouping variable, this value is identical to percentage of total sum.

Percent of Total N. Percentage of the total number of cases in each category.

Percent of Total Sum. Percentage of the total sum in each category.

Range. The difference between the largest and smallest values of a numeric variable, the maximum minus the minimum.

Skewness. A measure of the asymmetry of a distribution. The normal distribution is symmetric and has a skewness value of 0. A distribution with a significant positive skewness has a long right tail. A distribution with a significant negative skewness has a long left tail. As a guideline, a skewness value more than twice its standard error is taken to indicate a departure from symmetry.

Standard Error of Kurtosis. The ratio of kurtosis to its standard error can be used as a test of normality (that is, you can reject normality if the ratio is less than -2 or greater than +2). A large positive value for kurtosis indicates that the tails of the distribution are longer than those of a normal distribution; a negative value for kurtosis indicates shorter tails (becoming like those of a box-shaped uniform distribution).

Standard Error of Skewness. The ratio of skewness to its standard error can be used as a test of normality (that is, you can reject normality if the ratio is less than -2 or greater than +2). A large positive value for skewness indicates a long right tail; an extreme negative value indicates a long left tail.

Sum. The sum or total of the values, across all cases with nonmissing values.

Variance. A measure of dispersion around the mean, equal to the sum of squared deviations from the mean divided by one less than the number of cases. The variance is measured in units that are the square of those of the variable itself.

OLAP Cubes Differences

Figure 21-3
OLAP Cubes Differences dialog box

OLAP Cubes: Differences

Differences for Summary Statistics

None

Differences between variables

Differences between groups

Type of Difference

Percentage difference

Arithmetic difference

Differences between Variables

Variable:

Minus variable:

Percent label:

Arithmetic label:

Pairs:

Delete Pair

Differences between Groups of Cases

Grouping variable:

Category:

Minus category:

Percent label:

Arithmetic label:

Delete Pair

Continue Cancel Help

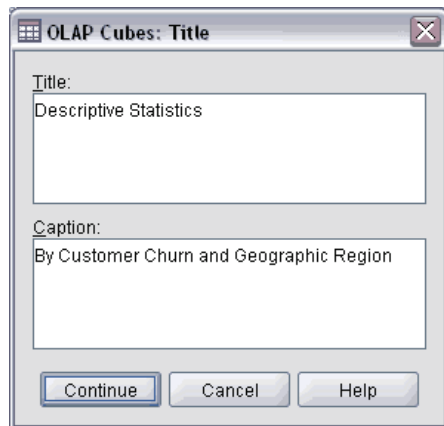
This dialog box allows you to calculate percentage and arithmetic differences between summary variables or between groups that are defined by a grouping variable. Differences are calculated for all measures that are selected in the OLAP Cubes Statistics dialog box.

Differences between Variables. Calculates differences between pairs of variables. Summary statistics values for the second variable (the Minus variable) in each pair are subtracted from summary statistics values for the first variable in the pair. For percentage differences, the value of the summary variable for the Minus variable is used as the denominator. You must select at least two summary variables in the main dialog box before you can specify differences between variables.

Differences between Groups of Cases. Calculates differences between pairs of groups defined by a grouping variable. Summary statistics values for the second category in each pair (the Minus category) are subtracted from summary statistics values for the first category in the pair. Percentage differences use the value of the summary statistic for the Minus category as the denominator. You must select one or more grouping variables in the main dialog box before you can specify differences between groups.

OLAP Cubes Title

Figure 21-4
OLAP Cubes Title dialog box



You can change the title of your output or add a caption that will appear below the output table. You can also control line wrapping of titles and captions by typing `\n` wherever you want to insert a line break in the text.

T Tests

Three types of t tests are available:

Independent-samples t test (two-sample t test). Compares the means of one variable for two groups of cases. Descriptive statistics for each group and Levene's test for equality of variances are provided, as well as both equal- and unequal-variance t values and a 95% confidence interval for the difference in means.

Paired-samples t test (dependent t test). Compares the means of two variables for a single group. This test is also for matched pairs or case-control study designs. The output includes descriptive statistics for the test variables, the correlation between the variables, descriptive statistics for the paired differences, the t test, and a 95% confidence interval.

One-sample t test. Compares the mean of one variable with a known or hypothesized value. Descriptive statistics for the test variables are displayed along with the t test. A 95% confidence interval for the difference between the mean of the test variable and the hypothesized test value is part of the default output.

Independent-Samples T Test

The Independent-Samples T Test procedure compares means for two groups of cases. Ideally, for this test, the subjects should be randomly assigned to two groups, so that any difference in response is due to the treatment (or lack of treatment) and not to other factors. This is not the case if you compare average income for males and females. A person is not randomly assigned to be a male or female. In such situations, you should ensure that differences in other factors are not masking or enhancing a significant difference in means. Differences in average income may be influenced by factors such as education (and not by sex alone).

Example. Patients with high blood pressure are randomly assigned to a placebo group and a treatment group. The placebo subjects receive an inactive pill, and the treatment subjects receive a new drug that is expected to lower blood pressure. After the subjects are treated for two months, the two-sample t test is used to compare the average blood pressures for the placebo group and the treatment group. Each patient is measured once and belongs to one group.

Statistics. For each variable: sample size, mean, standard deviation, and standard error of the mean. For the difference in means: mean, standard error, and confidence interval (you can specify the confidence level). Tests: Levene's test for equality of variances and both pooled-variances and separate-variances t tests for equality of means.

Data. The values of the quantitative variable of interest are in a single column in the data file. The procedure uses a grouping variable with two values to separate the cases into two groups. The grouping variable can be numeric (values such as 1 and 2 or 6.25 and 12.5) or short string

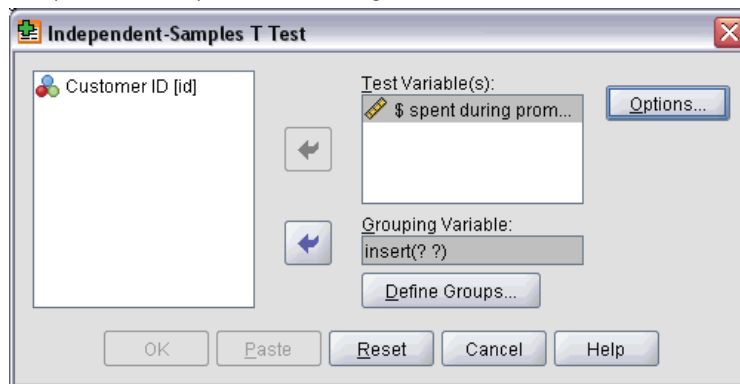
(such as *yes* and *no*). As an alternative, you can use a quantitative variable, such as *age*, to split the cases into two groups by specifying a cutpoint (cutpoint 21 splits *age* into an under-21 group and a 21-and-over group).

Assumptions. For the equal-variance *t* test, the observations should be independent, random samples from normal distributions with the same population variance. For the unequal-variance *t* test, the observations should be independent, random samples from normal distributions. The two-sample *t* test is fairly robust to departures from normality. When checking distributions graphically, look to see that they are symmetric and have no outliers.

To Obtain an Independent-Samples T Test

- ▶ From the menus choose:
Analyze
 Compare Means
 Independent-Samples T Test...

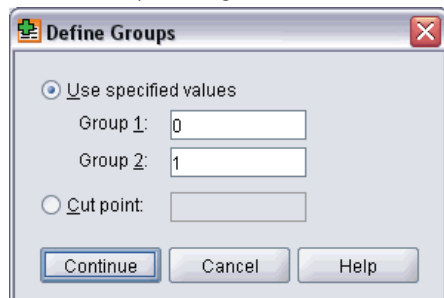
Figure 22-1
Independent-Samples T Test dialog box



- ▶ Select one or more quantitative test variables. A separate *t* test is computed for each variable.
- ▶ Select a single grouping variable, and then click Define Groups to specify two codes for the groups that you want to compare.
- ▶ Optionally, click Options to control the treatment of missing data and the level of the confidence interval.

Independent-Samples T Test Define Groups


Figure 22-2
Define Groups dialog box for numeric variables



For numeric grouping variables, define the two groups for the t test by specifying two values or a cutpoint:

- **Use specified values.** Enter a value for Group 1 and another value for Group 2. Cases with any other values are excluded from the analysis. Numbers need not be integers (for example, 6.25 and 12.5 are valid).
- **Cutpoint.** Enter a number that splits the values of the grouping variable into two sets. All cases with values that are less than the cutpoint form one group, and cases with values that are greater than or equal to the cutpoint form the other group.

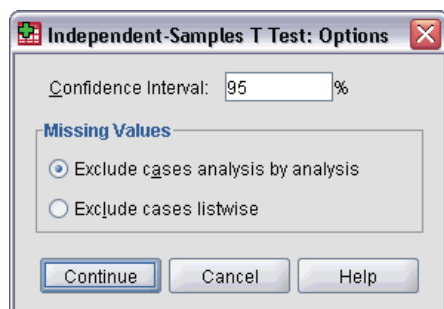
Figure 22-3
Define Groups dialog box for string variables



For string grouping variables, enter a string for Group 1 and another value for Group 2, such as *yes* and *no*. Cases with other strings are excluded from the analysis.

Independent-Samples T Test Options

Figure 22-4
Independent-Samples T Test Options dialog box



Confidence Interval. By default, a 95% confidence interval for the difference in means is displayed. Enter a value between 1 and 99 to request a different confidence level.

Missing Values. When you test several variables, and data are missing for one or more variables, you can tell the procedure which cases to include (or exclude).

- **Exclude cases analysis by analysis.** Each t test uses all cases that have valid data for the tested variables. Sample sizes may vary from test to test.
- **Exclude cases listwise.** Each t test uses only cases that have valid data for all variables that are used in the requested t tests. The sample size is constant across tests.

Paired-Samples T Test

The Paired-Samples T Test procedure compares the means of two variables for a single group. The procedure computes the differences between values of the two variables for each case and tests whether the average differs from 0.

Example. In a study on high blood pressure, all patients are measured at the beginning of the study, given a treatment, and measured again. Thus, each subject has two measures, often called *before* and *after* measures. An alternative design for which this test is used is a matched-pairs or case-control study, in which each record in the data file contains the response for the patient and also for his or her matched control subject. In a blood pressure study, patients and controls might be matched by age (a 75-year-old patient with a 75-year-old control group member).

Statistics. For each variable: mean, sample size, standard deviation, and standard error of the mean. For each pair of variables: correlation, average difference in means, t test, and confidence interval for mean difference (you can specify the confidence level). Standard deviation and standard error of the mean difference.

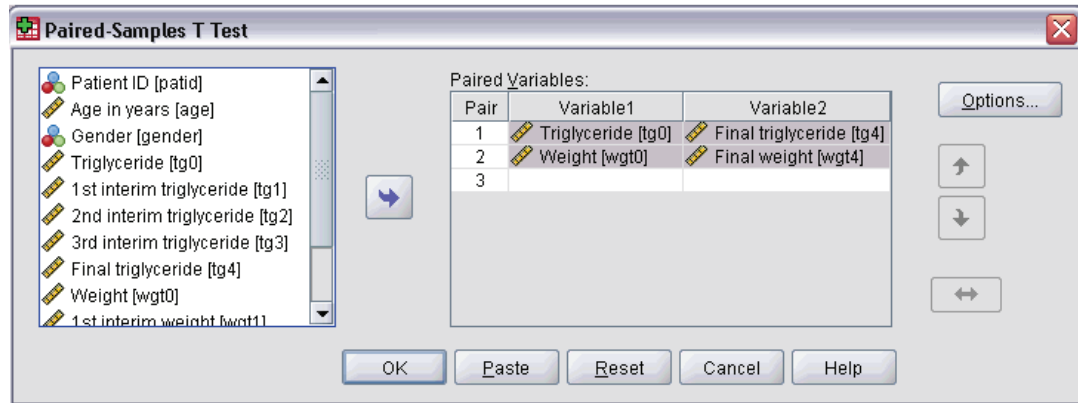
Data. For each paired test, specify two quantitative variables (interval level of measurement or ratio level of measurement). For a matched-pairs or case-control study, the response for each test subject and its matched control subject must be in the same case in the data file.

Assumptions. Observations for each pair should be made under the same conditions. The mean differences should be normally distributed. Variances of each variable can be equal or unequal.

To Obtain a Paired-Samples T Test

- From the menus choose:
 - Analyze
 - Compare Means
 - Paired-Samples T Test...

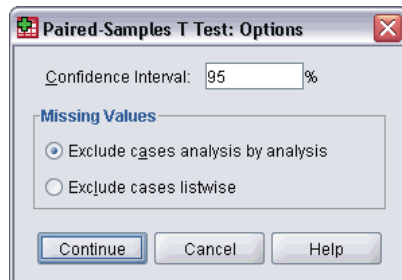
Figure 22-5
Paired-Samples T Test dialog box



- ▶ Select one or more pairs of variables
- ▶ Optionally, click Options to control the treatment of missing data and the level of the confidence interval.

Paired-Samples T Test Options

Figure 22-6
Paired-Samples T Test Options dialog box



Confidence Interval. By default, a 95% confidence interval for the difference in means is displayed. Enter a value between 1 and 99 to request a different confidence level.

Missing Values. When you test several variables, and data are missing for one or more variables, you can tell the procedure which cases to include (or exclude):

- **Exclude cases analysis by analysis.** Each t test uses all cases that have valid data for the tested pair of variables. Sample sizes may vary from test to test.
- **Exclude cases listwise.** Each t test uses only cases that have valid data for all pairs of tested variables. The sample size is constant across tests.

One-Sample T Test

The One-Sample T Test procedure tests whether the mean of a single variable differs from a specified constant.

Examples. A researcher might want to test whether the average IQ score for a group of students differs from 100. Or a cereal manufacturer can take a sample of boxes from the production line and check whether the mean weight of the samples differs from 1.3 pounds at the 95% confidence level.

Statistics. For each test variable: mean, standard deviation, and standard error of the mean. The average difference between each data value and the hypothesized test value, a t test that tests that this difference is 0, and a confidence interval for this difference (you can specify the confidence level).

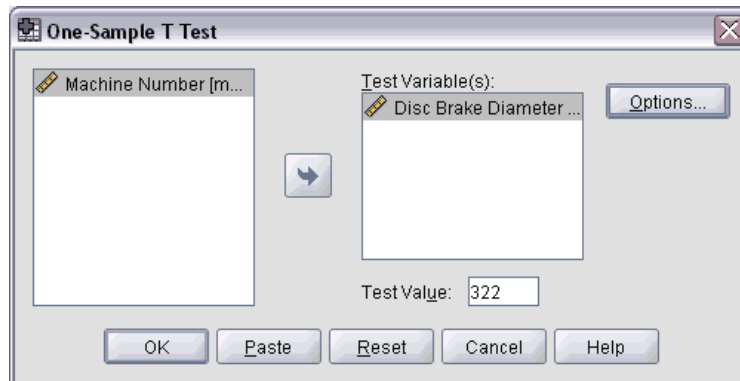
Data. To test the values of a quantitative variable against a hypothesized test value, choose a quantitative variable and enter a hypothesized test value.

Assumptions. This test assumes that the data are normally distributed; however, this test is fairly robust to departures from normality.

To Obtain a One-Sample T Test

- ▶ From the menu choose:
 - Analyze
 - Compare Means
 - One-Sample T Test...

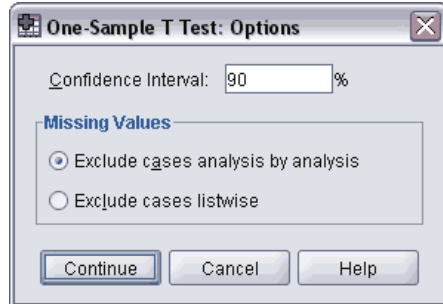
Figure 22-7
One-Sample T Test dialog box



- ▶ Select one or more variables to be tested against the same hypothesized value.
- ▶ Enter a numeric test value against which each sample mean is compared.
- ▶ Optionally, click Options to control the treatment of missing data and the level of the confidence interval.

One-Sample T Test Options

Figure 22-8
One-Sample T Test Options dialog box



Confidence Interval. By default, a 95% confidence interval for the difference between the mean and the hypothesized test value is displayed. Enter a value between 1 and 99 to request a different confidence level.

Missing Values. When you test several variables, and data are missing for one or more variables, you can tell the procedure which cases to include (or exclude).

- **Exclude cases analysis by analysis.** Each t test uses all cases that have valid data for the tested variable. Sample sizes may vary from test to test.
- **Exclude cases listwise.** Each t test uses only cases that have valid data for all variables that are used in any of the requested t tests. The sample size is constant across tests.

T-TEST Command Additional Features

The command syntax language also allows you to:

- Produce both one-sample and independent-samples t tests by running a single command.
- Test a variable against each variable on a list in a paired t test (with the `PAIRS` subcommand).

See the *Command Syntax Reference* for complete syntax information.

One-Way ANOVA

The One-Way ANOVA procedure produces a one-way analysis of variance for a quantitative dependent variable by a single factor (independent) variable. Analysis of variance is used to test the hypothesis that several means are equal. This technique is an extension of the two-sample t test.

In addition to determining that differences exist among the means, you may want to know which means differ. There are two types of tests for comparing means: a priori contrasts and post hoc tests. Contrasts are tests set up *before* running the experiment, and post hoc tests are run *after* the experiment has been conducted. You can also test for trends across categories.

Example. Doughnuts absorb fat in various amounts when they are cooked. An experiment is set up involving three types of fat: peanut oil, corn oil, and lard. Peanut oil and corn oil are unsaturated fats, and lard is a saturated fat. Along with determining whether the amount of fat absorbed depends on the type of fat used, you could set up an a priori contrast to determine whether the amount of fat absorption differs for saturated and unsaturated fats.

Statistics. For each group: number of cases, mean, standard deviation, standard error of the mean, minimum, maximum, and 95% confidence interval for the mean. Levene's test for homogeneity of variance, analysis-of-variance table and robust tests of the equality of means for each dependent variable, user-specified a priori contrasts, and post hoc range tests and multiple comparisons: Bonferroni, Sidak, Tukey's honestly significant difference, Hochberg's GT2, Gabriel, Dunnett, Ryan-Einot-Gabriel-Welsch F test (R-E-G-W F), Ryan-Einot-Gabriel-Welsch range test (R-E-G-W Q), Tamhane's T2, Dunnett's T3, Games-Howell, Dunnett's C , Duncan's multiple range test, Student-Newman-Keuls (S-N-K), Tukey's b , Waller-Duncan, Scheffé, and least-significant difference.

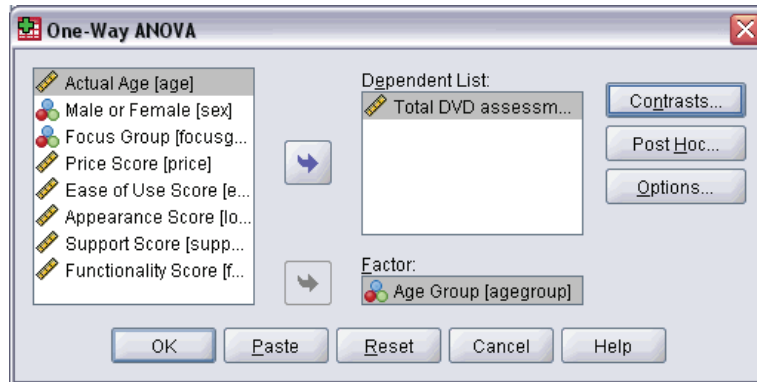
Data. Factor variable values should be integers, and the dependent variable should be quantitative (interval level of measurement).

Assumptions. Each group is an independent random sample from a normal population. Analysis of variance is robust to departures from normality, although the data should be symmetric. The groups should come from populations with equal variances. To test this assumption, use Levene's homogeneity-of-variance test.

To Obtain a One-Way Analysis of Variance

- ▶ From the menus choose:
 - Analyze
 - Compare Means
 - One-Way ANOVA...

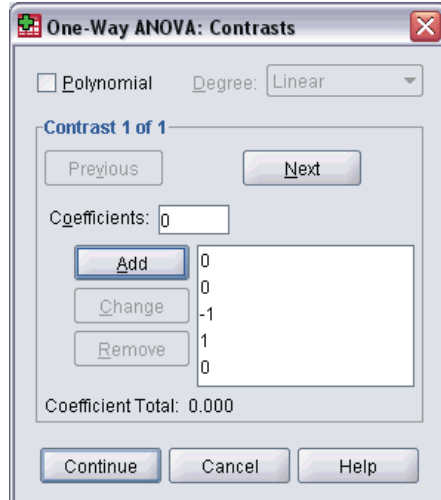
Figure 23-1
One-Way ANOVA dialog box



- ▶ Select one or more dependent variables.
- ▶ Select a single independent factor variable.

One-Way ANOVA Contrasts

Figure 23-2
One-Way ANOVA Contrasts dialog box



You can partition the between-groups sums of squares into trend components or specify a priori contrasts.

Polynomial. Partitions the between-groups sums of squares into trend components. You can test for a trend of the dependent variable across the ordered levels of the factor variable. For example, you could test for a linear trend (increasing or decreasing) in salary across the ordered levels of highest degree earned.

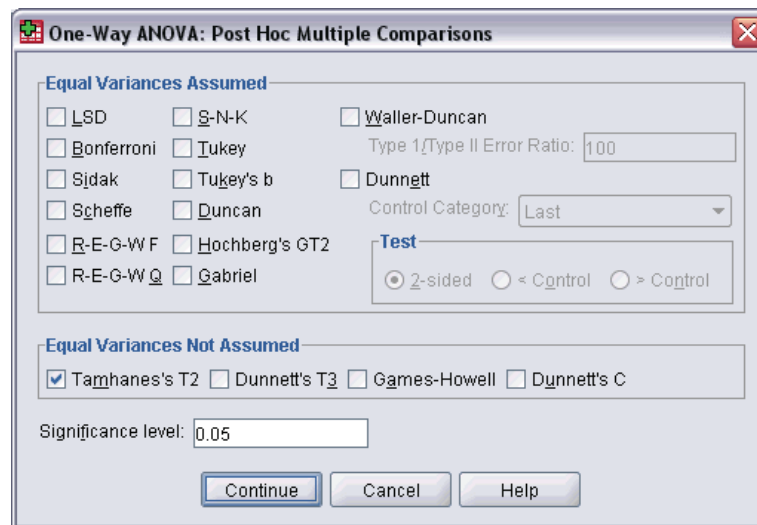
- **Degree.** You can choose a 1st, 2nd, 3rd, 4th, or 5th degree polynomial.

Coefficients. User-specified a priori contrasts to be tested by the t statistic. Enter a coefficient for each group (category) of the factor variable and click Add after each entry. Each new value is added to the bottom of the coefficient list. To specify additional sets of contrasts, click Next. Use Next and Previous to move between sets of contrasts.

The order of the coefficients is important because it corresponds to the ascending order of the category values of the factor variable. The first coefficient on the list corresponds to the lowest group value of the factor variable, and the last coefficient corresponds to the highest value. For example, if there are six categories of the factor variable, the coefficients $-1, 0, 0, 0, 0.5,$ and 0.5 contrast the first group with the fifth and sixth groups. For most applications, the coefficients should sum to 0. Sets that do not sum to 0 can also be used, but a warning message is displayed.

One-Way ANOVA Post Hoc Tests

Figure 23-3
One-way ANOVA Post Hoc Multiple Comparisons dialog box



Once you have determined that differences exist among the means, post hoc range tests and pairwise multiple comparisons can determine which means differ. Range tests identify homogeneous subsets of means that are not different from each other. Pairwise multiple comparisons test the difference between each pair of means and yield a matrix where asterisks indicate significantly different group means at an alpha level of 0.05.

Equal Variances Assumed

Tukey's honestly significant difference test, Hochberg's GT2, Gabriel, and Scheffé are multiple comparison tests and range tests. Other available range tests are Tukey's b , S-N-K (Student-Newman-Keuls), Duncan, R-E-G-W F (Ryan-Einot-Gabriel-Welsch F test), R-E-G-W Q (Ryan-Einot-Gabriel-Welsch range test), and Waller-Duncan. Available multiple comparison tests are Bonferroni, Tukey's honestly significant difference test, Sidak, Gabriel, Hochberg, Dunnett, Scheffé, and LSD (least significant difference).

- **LSD.** Uses t tests to perform all pairwise comparisons between group means. No adjustment is made to the error rate for multiple comparisons.
- **Bonferroni.** Uses t tests to perform pairwise comparisons between group means, but controls overall error rate by setting the error rate for each test to the experimentwise error rate divided by the total number of tests. Hence, the observed significance level is adjusted for the fact that multiple comparisons are being made.
- **Sidak.** Pairwise multiple comparison test based on a t statistic. Sidak adjusts the significance level for multiple comparisons and provides tighter bounds than Bonferroni.
- **Scheffe.** Performs simultaneous joint pairwise comparisons for all possible pairwise combinations of means. Uses the F sampling distribution. Can be used to examine all possible linear combinations of group means, not just pairwise comparisons.
- **R-E-G-W F.** Ryan-Einot-Gabriel-Welsch multiple stepdown procedure based on an F test.
- **R-E-G-W Q.** Ryan-Einot-Gabriel-Welsch multiple stepdown procedure based on the Studentized range.
- **S-N-K.** Makes all pairwise comparisons between means using the Studentized range distribution. With equal sample sizes, it also compares pairs of means within homogeneous subsets, using a stepwise procedure. Means are ordered from highest to lowest, and extreme differences are tested first.
- **Tukey.** Uses the Studentized range statistic to make all of the pairwise comparisons between groups. Sets the experimentwise error rate at the error rate for the collection for all pairwise comparisons.
- **Tukey's b.** Uses the Studentized range distribution to make pairwise comparisons between groups. The critical value is the average of the corresponding value for the Tukey's honestly significant difference test and the Student-Newman-Keuls.
- **Duncan.** Makes pairwise comparisons using a stepwise order of comparisons identical to the order used by the Student-Newman-Keuls test, but sets a protection level for the error rate for the collection of tests, rather than an error rate for individual tests. Uses the Studentized range statistic.
- **Hochberg's GT2.** Multiple comparison and range test that uses the Studentized maximum modulus. Similar to Tukey's honestly significant difference test.
- **Gabriel.** Pairwise comparison test that used the Studentized maximum modulus and is generally more powerful than Hochberg's GT2 when the cell sizes are unequal. Gabriel's test may become liberal when the cell sizes vary greatly.
- **Waller-Duncan.** Multiple comparison test based on a t statistic; uses a Bayesian approach.
- **Dunnett.** Pairwise multiple comparison t test that compares a set of treatments against a single control mean. The last category is the default control category. Alternatively, you can choose the first category. 2-sided tests that the mean at any level (except the control category) of the factor is not equal to that of the control category. < Control tests if the mean at any level of the factor is smaller than that of the control category. > Control tests if the mean at any level of the factor is greater than that of the control category.

Equal Variances Not Assumed

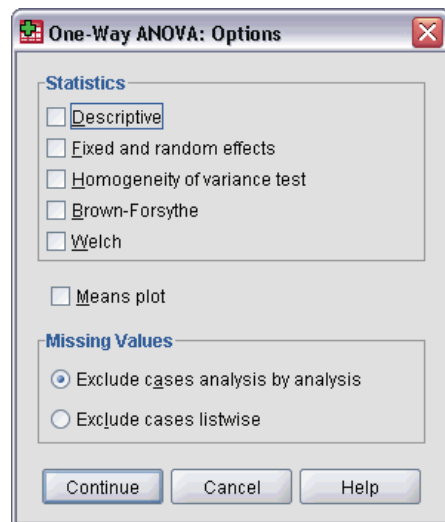
Multiple comparison tests that do not assume equal variances are Tamhane's T2, Dunnett's T3, Games-Howell, and Dunnett's C.

- **Tamhane's T2.** Conservative pairwise comparisons test based on a t test. This test is appropriate when the variances are unequal.
- **Dunnett's T3.** Pairwise comparison test based on the Studentized maximum modulus. This test is appropriate when the variances are unequal.
- **Games-Howell.** Pairwise comparison test that is sometimes liberal. This test is appropriate when the variances are unequal.
- **Dunnett's C.** Pairwise comparison test based on the Studentized range. This test is appropriate when the variances are unequal.

Note: You may find it easier to interpret the output from post hoc tests if you deselect Hide empty rows and columns in the Table Properties dialog box (in an activated pivot table, choose Table Properties from the Format menu).

One-Way ANOVA Options

Figure 23-4
One-Way ANOVA Options dialog box



Statistics. Choose one or more of the following:

- **Descriptive.** Calculates the number of cases, mean, standard deviation, standard error of the mean, minimum, maximum, and 95% confidence intervals for each dependent variable for each group.
- **Fixed and random effects.** Displays the standard deviation, standard error, and 95% confidence interval for the fixed-effects model, and the standard error, 95% confidence interval, and estimate of between-components variance for the random-effects model.
- **Homogeneity of variance test.** Calculates the Levene statistic to test for the equality of group variances. This test is not dependent on the assumption of normality.

- **Brown-Forsythe.** Calculates the Brown-Forsythe statistic to test for the equality of group means. This statistic is preferable to the F statistic when the assumption of equal variances does not hold.
- **Welch.** Calculates the Welch statistic to test for the equality of group means. This statistic is preferable to the F statistic when the assumption of equal variances does not hold.

Means plot. Displays a chart that plots the subgroup means (the means for each group defined by values of the factor variable).

Missing Values. Controls the treatment of missing values.

- **Exclude cases analysis by analysis.** A case with a missing value for either the dependent or the factor variable for a given analysis is not used in that analysis. Also, a case outside the range specified for the factor variable is not used.
- **Exclude cases listwise.** Cases with missing values for the factor variable or for any dependent variable included on the dependent list in the main dialog box are excluded from all analyses. If you have not specified multiple dependent variables, this has no effect.

ONEWAY Command Additional Features

The command syntax language also allows you to:

- Obtain fixed- and random-effects statistics. Standard deviation, standard error of the mean, and 95% confidence intervals for the fixed-effects model. Standard error, 95% confidence intervals, and estimate of between-components variance for random-effects model (using `STATISTICS=EFFECTS`).
- Specify alpha levels for the least significance difference, Bonferroni, Duncan, and Scheffé multiple comparison tests (with the `RANGES` subcommand).
- Write a matrix of means, standard deviations, and frequencies, or read a matrix of means, frequencies, pooled variances, and degrees of freedom for the pooled variances. These matrices can be used in place of raw data to obtain a one-way analysis of variance (with the `MATRIX` subcommand).

See the *Command Syntax Reference* for complete syntax information.

GLM Univariate Analysis

The GLM Univariate procedure provides regression analysis and analysis of variance for one dependent variable by one or more factors and/or variables. The factor variables divide the population into groups. Using this General Linear Model procedure, you can test null hypotheses about the effects of other variables on the means of various groupings of a single dependent variable. You can investigate interactions between factors as well as the effects of individual factors, some of which may be random. In addition, the effects of covariates and covariate interactions with factors can be included. For regression analysis, the independent (predictor) variables are specified as covariates.

Both balanced and unbalanced models can be tested. A design is balanced if each cell in the model contains the same number of cases. In addition to testing hypotheses, GLM Univariate produces estimates of parameters.

Commonly used a priori contrasts are available to perform hypothesis testing. Additionally, after an overall F test has shown significance, you can use post hoc tests to evaluate differences among specific means. Estimated marginal means give estimates of predicted mean values for the cells in the model, and profile plots (interaction plots) of these means allow you to easily visualize some of the relationships.

Residuals, predicted values, Cook's distance, and leverage values can be saved as new variables in your data file for checking assumptions.

WLS Weight allows you to specify a variable used to give observations different weights for a weighted least-squares (WLS) analysis, perhaps to compensate for a different precision of measurement.

Example. Data are gathered for individual runners in the Chicago marathon for several years. The time in which each runner finishes is the dependent variable. Other factors include weather (cold, pleasant, or hot), number of months of training, number of previous marathons, and gender. Age is considered a covariate. You might find that gender is a significant effect and that the interaction of gender with weather is significant.

Methods. Type I, Type II, Type III, and Type IV sums of squares can be used to evaluate different hypotheses. Type III is the default.

Statistics. Post hoc range tests and multiple comparisons: least significant difference, Bonferroni, Sidak, Scheffé, Ryan-Einot-Gabriel-Welsch multiple F , Ryan-Einot-Gabriel-Welsch multiple range, Student-Newman-Keuls, Tukey's honestly significant difference, Tukey's b , Duncan, Hochberg's GT2, Gabriel, Waller-Duncan t test, Dunnett (one-sided and two-sided), Tamhane's T2, Dunnett's T3, Games-Howell, and Dunnett's C . Descriptive statistics: observed means, standard deviations, and counts for all of the dependent variables in all cells. The Levene test for homogeneity of variance.

Plots. Spread-versus-level, residual, and profile (interaction).

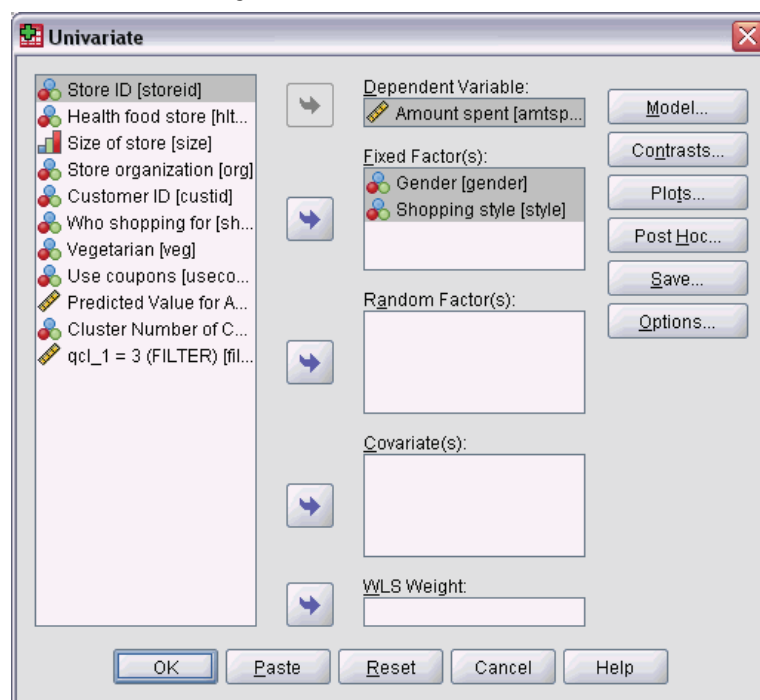
Data. The dependent variable is quantitative. Factors are categorical. They can have numeric values or string values of up to eight characters. Covariates are quantitative variables that are related to the dependent variable.

Assumptions. The data are a random sample from a normal population; in the population, all cell variances are the same. Analysis of variance is robust to departures from normality, although the data should be symmetric. To check assumptions, you can use homogeneity of variances tests and spread-versus-level plots. You can also examine residuals and residual plots.

To Obtain GLM Univariate Tables

- ▶ From the menus choose:
Analyze
 General Linear Model
 Univariate...

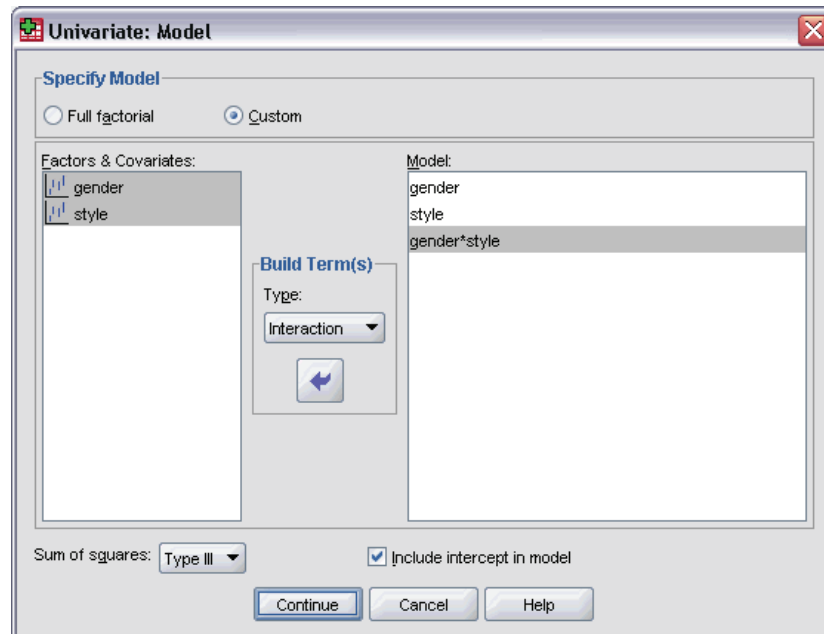
Figure 24-1
GLM Univariate dialog box



- ▶ Select a dependent variable.
- ▶ Select variables for Fixed Factor(s), Random Factor(s), and Covariate(s), as appropriate for your data.
- ▶ Optionally, you can use WLS Weight to specify a weight variable for weighted least-squares analysis. If the value of the weighting variable is zero, negative, or missing, the case is excluded from the analysis. A variable already used in the model cannot be used as a weighting variable.

GLM Model

Figure 24-2
Univariate Model dialog box



Specify Model. A full factorial model contains all factor main effects, all covariate main effects, and all factor-by-factor interactions. It does not contain covariate interactions. Select Custom to specify only a subset of interactions or to specify factor-by-covariate interactions. You must indicate all of the terms to be included in the model.

Factors and Covariates. The factors and covariates are listed.

Model. The model depends on the nature of your data. After selecting Custom, you can select the main effects and interactions that are of interest in your analysis.

Sum of squares. The method of calculating the sums of squares. For balanced or unbalanced models with no missing cells, the Type III sum-of-squares method is most commonly used.

Include intercept in model. The intercept is usually included in the model. If you can assume that the data pass through the origin, you can exclude the intercept.

Build Terms

For the selected factors and covariates:

Interaction. Creates the highest-level interaction term of all selected variables. This is the default.

Main effects. Creates a main-effects term for each variable selected.

All 2-way. Creates all possible two-way interactions of the selected variables.

All 3-way. Creates all possible three-way interactions of the selected variables.

All 4-way. Creates all possible four-way interactions of the selected variables.

All 5-way. Creates all possible five-way interactions of the selected variables.

Sum of Squares

For the model, you can choose a type of sums of squares. Type III is the most commonly used and is the default.

Type I. This method is also known as the hierarchical decomposition of the sum-of-squares method. Each term is adjusted for only the term that precedes it in the model. Type I sums of squares are commonly used for:

- A balanced ANOVA model in which any main effects are specified before any first-order interaction effects, any first-order interaction effects are specified before any second-order interaction effects, and so on.
- A polynomial regression model in which any lower-order terms are specified before any higher-order terms.
- A purely nested model in which the first-specified effect is nested within the second-specified effect, the second-specified effect is nested within the third, and so on. (This form of nesting can be specified only by using syntax.)

Type II. This method calculates the sums of squares of an effect in the model adjusted for all other “appropriate” effects. An appropriate effect is one that corresponds to all effects that do not contain the effect being examined. The Type II sum-of-squares method is commonly used for:

- A balanced ANOVA model.
- Any model that has main factor effects only.
- Any regression model.
- A purely nested design. (This form of nesting can be specified by using syntax.)

Type III. The default. This method calculates the sums of squares of an effect in the design as the sums of squares adjusted for any other effects that do not contain it and orthogonal to any effects (if any) that contain it. The Type III sums of squares have one major advantage in that they are invariant with respect to the cell frequencies as long as the general form of estimability remains constant. Hence, this type of sums of squares is often considered useful for an unbalanced model with no missing cells. In a factorial design with no missing cells, this method is equivalent to the Yates’ weighted-squares-of-means technique. The Type III sum-of-squares method is commonly used for:

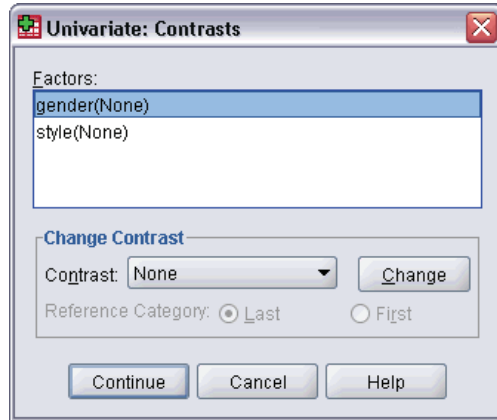
- Any models listed in Type I and Type II.
- Any balanced or unbalanced model with no empty cells.

Type IV. This method is designed for a situation in which there are missing cells. For any effect F in the design, if F is not contained in any other effect, then Type IV = Type III = Type II. When F is contained in other effects, Type IV distributes the contrasts being made among the parameters in F to all higher-level effects equitably. The Type IV sum-of-squares method is commonly used for:

- Any models listed in Type I and Type II.
- Any balanced model or unbalanced model with empty cells.

GLM Contrasts

Figure 24-3
Univariate Contrasts dialog box



Contrasts are used to test for differences among the levels of a factor. You can specify a contrast for each factor in the model (in a repeated measures model, for each between-subjects factor). Contrasts represent linear combinations of the parameters.

Hypothesis testing is based on the null hypothesis $\mathbf{LB} = 0$, where \mathbf{L} is the contrast coefficients matrix and \mathbf{B} is the parameter vector. When a contrast is specified, an \mathbf{L} matrix is created. The columns of the \mathbf{L} matrix corresponding to the factor match the contrast. The remaining columns are adjusted so that the \mathbf{L} matrix is estimable.

The output includes an F statistic for each set of contrasts. Also displayed for the contrast differences are Bonferroni-type simultaneous confidence intervals based on Student's t distribution.

Available Contrasts

Available contrasts are deviation, simple, difference, Helmert, repeated, and polynomial. For deviation contrasts and simple contrasts, you can choose whether the reference category is the last or first category.

Contrast Types

Deviation. Compares the mean of each level (except a reference category) to the mean of all of the levels (grand mean). The levels of the factor can be in any order.

Simple. Compares the mean of each level to the mean of a specified level. This type of contrast is useful when there is a control group. You can choose the first or last category as the reference.

Difference. Compares the mean of each level (except the first) to the mean of previous levels. (Sometimes called reverse Helmert contrasts.)

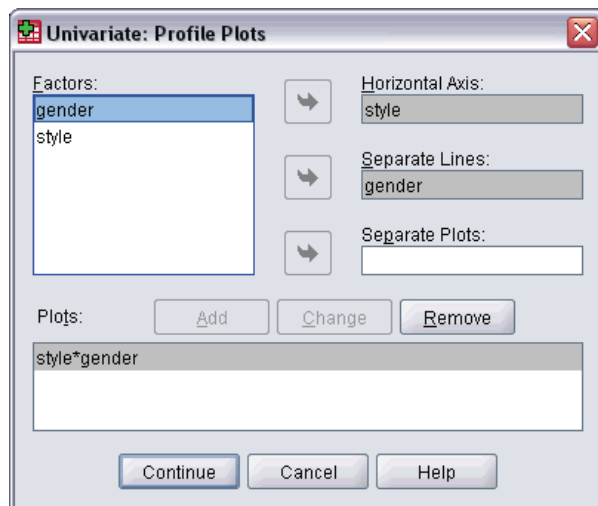
Helmert. Compares the mean of each level of the factor (except the last) to the mean of subsequent levels.

Repeated. Compares the mean of each level (except the last) to the mean of the subsequent level.

Polynomial. Compares the linear effect, quadratic effect, cubic effect, and so on. The first degree of freedom contains the linear effect across all categories; the second degree of freedom, the quadratic effect; and so on. These contrasts are often used to estimate polynomial trends.

GLM Profile Plots


Figure 24-4
Univariate Profile Plots dialog box



Profile plots (interaction plots) are useful for comparing marginal means in your model. A profile plot is a line plot in which each point indicates the estimated marginal mean of a dependent variable (adjusted for any covariates) at one level of a factor. The levels of a second factor can be used to make separate lines. Each level in a third factor can be used to create a separate plot. All fixed and random factors, if any, are available for plots. For multivariate analyses, profile plots are created for each dependent variable. In a repeated measures analysis, both between-subjects factors and within-subjects factors can be used in profile plots. GLM Multivariate and GLM Repeated Measures are available only if you have the Advanced Statistics option installed.

A profile plot of one factor shows whether the estimated marginal means are increasing or decreasing across levels. For two or more factors, parallel lines indicate that there is no interaction between factors, which means that you can investigate the levels of only one factor. Nonparallel lines indicate an interaction.

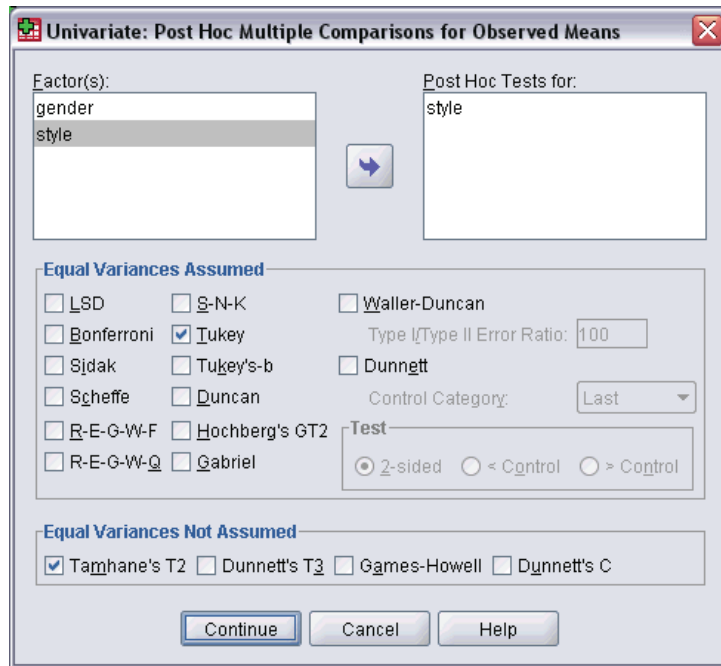
Figure 24-5
Nonparallel plot (left) and parallel plot (right)



After a plot is specified by selecting factors for the horizontal axis and, optionally, factors for separate lines and separate plots, the plot must be added to the Plots list.

GLM Post Hoc Comparisons

Figure 24-6
Post Hoc dialog box



Post hoc multiple comparison tests. Once you have determined that differences exist among the means, post hoc range tests and pairwise multiple comparisons can determine which means differ. Comparisons are made on unadjusted values. These tests are used for fixed between-subjects factors only. In GLM Repeated Measures, these tests are not available if there are no between-subjects factors, and the post hoc multiple comparison tests are performed for the average across the levels of the within-subjects factors. For GLM Multivariate, the post hoc tests are performed for each dependent variable separately. GLM Multivariate and GLM Repeated Measures are available only if you have the Advanced Statistics option installed.

The Bonferroni and Tukey's honestly significant difference tests are commonly used multiple comparison tests. The **Bonferroni test**, based on Student's t statistic, adjusts the observed significance level for the fact that multiple comparisons are made. **Sidak's t test** also adjusts the significance level and provides tighter bounds than the Bonferroni test. **Tukey's honestly significant difference test** uses the Studentized range statistic to make all pairwise comparisons between groups and sets the experimentwise error rate to the error rate for the collection for all pairwise comparisons. When testing a large number of pairs of means, Tukey's honestly significant difference test is more powerful than the Bonferroni test. For a small number of pairs, Bonferroni is more powerful.

Hochberg's GT2 is similar to Tukey's honestly significant difference test, but the Studentized maximum modulus is used. Usually, Tukey's test is more powerful. **Gabriel's pairwise comparisons test** also uses the Studentized maximum modulus and is generally more powerful than Hochberg's GT2 when the cell sizes are unequal. Gabriel's test may become liberal when the cell sizes vary greatly.

Dunnnett's pairwise multiple comparison t test compares a set of treatments against a single control mean. The last category is the default control category. Alternatively, you can choose the first category. You can also choose a two-sided or one-sided test. To test that the mean at any level (except the control category) of the factor is not equal to that of the control category, use a two-sided test. To test whether the mean at any level of the factor is smaller than that of the control category, select < Control. Likewise, to test whether the mean at any level of the factor is larger than that of the control category, select > Control.

Ryan, Einot, Gabriel, and Welsch (R-E-G-W) developed two multiple step-down range tests. Multiple step-down procedures first test whether all means are equal. If all means are not equal, subsets of means are tested for equality. **R-E-G-W F** is based on an F test and **R-E-G-W Q** is based on the Studentized range. These tests are more powerful than Duncan's multiple range test and Student-Newman-Keuls (which are also multiple step-down procedures), but they are not recommended for unequal cell sizes.

When the variances are unequal, use **Tamhane's T2** (conservative pairwise comparisons test based on a t test), **Dunnnett's T3** (pairwise comparison test based on the Studentized maximum modulus), **Games-Howell pairwise comparison test** (sometimes liberal), or **Dunnnett's C** (pairwise comparison test based on the Studentized range). Note that these tests are not valid and will not be produced if there are multiple factors in the model.

Duncan's multiple range test, Student-Newman-Keuls (**S-N-K**), and **Tukey's b** are range tests that rank group means and compute a range value. These tests are not used as frequently as the tests previously discussed.

The **Waller-Duncan t test** uses a Bayesian approach. This range test uses the harmonic mean of the sample size when the sample sizes are unequal.

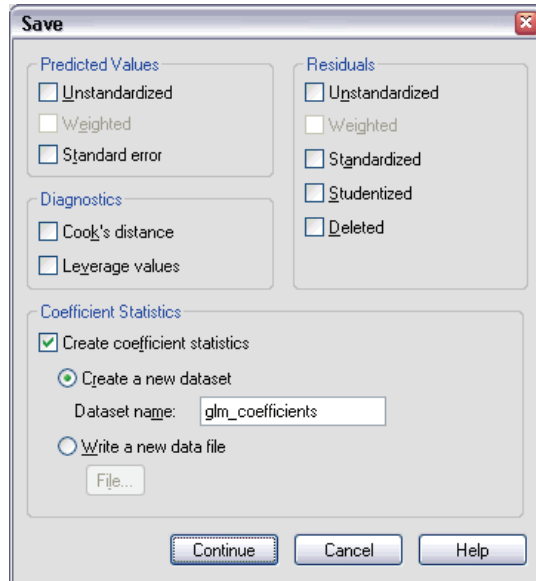
The significance level of the **Scheffé** test is designed to allow all possible linear combinations of group means to be tested, not just pairwise comparisons available in this feature. The result is that the Scheffé test is often more conservative than other tests, which means that a larger difference between means is required for significance.

The least significant difference (**LSD**) pairwise multiple comparison test is equivalent to multiple individual t tests between all pairs of groups. The disadvantage of this test is that no attempt is made to adjust the observed significance level for multiple comparisons.

Tests displayed. Pairwise comparisons are provided for LSD, Sidak, Bonferroni, Games-Howell, Tamhane's T2 and T3, Dunnnett's C, and Dunnnett's T3. Homogeneous subsets for range tests are provided for S-N-K, Tukey's b , Duncan, R-E-G-W F , R-E-G-W Q , and Waller. Tukey's honestly significant difference test, Hochberg's GT2, Gabriel's test, and Scheffé's test are both multiple comparison tests and range tests.

GLM Save

Figure 24-7
Save dialog box



You can save values predicted by the model, residuals, and related measures as new variables in the Data Editor. Many of these variables can be used for examining assumptions about the data. To save the values for use in another SPSS Statistics session, you must save the current data file.

Predicted Values. The values that the model predicts for each case.

- **Unstandardized.** The value the model predicts for the dependent variable.
- **Weighted.** Weighted unstandardized predicted values. Available only if a WLS variable was previously selected.
- **Standard error.** An estimate of the standard deviation of the average value of the dependent variable for cases that have the same values of the independent variables.

Diagnostics. Measures to identify cases with unusual combinations of values for the independent variables and cases that may have a large impact on the model.

- **Cook's distance.** A measure of how much the residuals of all cases would change if a particular case were excluded from the calculation of the regression coefficients. A large Cook's D indicates that excluding a case from computation of the regression statistics changes the coefficients substantially.
- **Leverage values.** Uncentered leverage values. The relative influence of each observation on the model's fit.

Residuals. An unstandardized residual is the actual value of the dependent variable minus the value predicted by the model. Standardized, Studentized, and deleted residuals are also available. If a WLS variable was chosen, weighted unstandardized residuals are available.

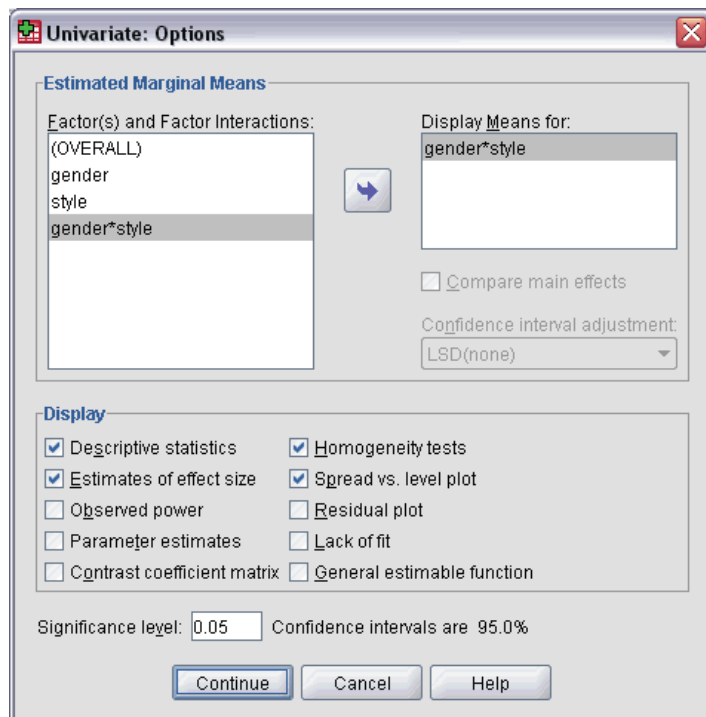
- **Unstandardized.** The difference between an observed value and the value predicted by the model.

- **Weighted.** Weighted unstandardized residuals. Available only if a WLS variable was previously selected.
- **Standardized.** The residual divided by an estimate of its standard deviation. Standardized residuals, which are also known as Pearson residuals, have a mean of 0 and a standard deviation of 1.
- **Studentized.** The residual divided by an estimate of its standard deviation that varies from case to case, depending on the distance of each case's values on the independent variables from the means of the independent variables.
- **Deleted.** The residual for a case when that case is excluded from the calculation of the regression coefficients. It is the difference between the value of the dependent variable and the adjusted predicted value.

Coefficient Statistics. Writes a variance-covariance matrix of the parameter estimates in the model to a new dataset in the current session or an external SPSS Statistics data file. Also, for each dependent variable, there will be a row of parameter estimates, a row of significance values for the t statistics corresponding to the parameter estimates, and a row of residual degrees of freedom. For a multivariate model, there are similar rows for each dependent variable. You can use this matrix file in other procedures that read matrix files.

GLM Options

Figure 24-8
Options dialog box



Optional statistics are available from this dialog box. Statistics are calculated using a fixed-effects model.

Estimated Marginal Means. Select the factors and interactions for which you want estimates of the population marginal means in the cells. These means are adjusted for the covariates, if any.

- **Compare main effects.** Provides uncorrected pairwise comparisons among estimated marginal means for any main effect in the model, for both between- and within-subjects factors. This item is available only if main effects are selected under the Display Means For list.
- **Confidence interval adjustment.** Select least significant difference (LSD), Bonferroni, or Sidak adjustment to the confidence intervals and significance. This item is available only if Compare main effects is selected.

Display. Select Descriptive statistics to produce observed means, standard deviations, and counts for all of the dependent variables in all cells. Estimates of effect size gives a partial eta-squared value for each effect and each parameter estimate. The eta-squared statistic describes the proportion of total variability attributable to a factor. Select Observed power to obtain the power of the test when the alternative hypothesis is set based on the observed value. Select Parameter estimates to produce the parameter estimates, standard errors, *t* tests, confidence intervals, and the observed power for each test. Select Contrast coefficient matrix to obtain the **L** matrix.

Homogeneity tests produces the Levene test of the homogeneity of variance for each dependent variable across all level combinations of the between-subjects factors, for between-subjects factors only. The spread-versus-level and residual plots options are useful for checking assumptions about the data. This item is disabled if there are no factors. Select Residual plot to produce an observed-by-predicted-by-standardized residual plot for each dependent variable. These plots are useful for investigating the assumption of equal variance. Select Lack of fit to check if the relationship between the dependent variable and the independent variables can be adequately described by the model. General estimable function allows you to construct custom hypothesis tests based on the general estimable function. Rows in any contrast coefficient matrix are linear combinations of the general estimable function.

Significance level. You might want to adjust the significance level used in post hoc tests and the confidence level used for constructing confidence intervals. The specified value is also used to calculate the observed power for the test. When you specify a significance level, the associated level of the confidence intervals is displayed in the dialog box.

UNIANOVA Command Additional Features

The command syntax language also allows you to:

- Specify nested effects in the design (using the `DESIGN` subcommand).
- Specify tests of effects versus a linear combination of effects or a value (using the `TEST` subcommand).
- Specify multiple contrasts (using the `CONTRAST` subcommand).
- Include user-missing values (using the `MISSING` subcommand).
- Specify EPS criteria (using the `CRITERIA` subcommand).
- Construct a custom **L** matrix, **M** matrix, or **K** matrix (using the `LMATRIX`, `MMATRIX`, and `KMATRIX` subcommands).
- For deviation or simple contrasts, specify an intermediate reference category (using the `CONTRAST` subcommand).

- Specify metrics for polynomial contrasts (using the `CONTRAST` subcommand).
- Specify error terms for post hoc comparisons (using the `POSTHOC` subcommand).
- Compute estimated marginal means for any factor or factor interaction among the factors in the factor list (using the `EMMEANS` subcommand).
- Specify names for temporary variables (using the `SAVE` subcommand).
- Construct a correlation matrix data file (using the `OUTFILE` subcommand).
- Construct a matrix data file that contains statistics from the between-subjects ANOVA table (using the `OUTFILE` subcommand).
- Save the design matrix to a new data file (using the `OUTFILE` subcommand).

See the *Command Syntax Reference* for complete syntax information.

Bivariate Correlations

The Bivariate Correlations procedure computes Pearson's correlation coefficient, Spearman's rho, and Kendall's tau-*b* with their significance levels. Correlations measure how variables or rank orders are related. Before calculating a correlation coefficient, screen your data for outliers (which can cause misleading results) and evidence of a linear relationship. Pearson's correlation coefficient is a measure of linear association. Two variables can be perfectly related, but if the relationship is not linear, Pearson's correlation coefficient is not an appropriate statistic for measuring their association.

Example. Is the number of games won by a basketball team correlated with the average number of points scored per game? A scatterplot indicates that there is a linear relationship. Analyzing data from the 1994–1995 NBA season yields that Pearson's correlation coefficient (0.581) is significant at the 0.01 level. You might suspect that the more games won per season, the fewer points the opponents scored. These variables are negatively correlated (–0.401), and the correlation is significant at the 0.05 level.

Statistics. For each variable: number of cases with nonmissing values, mean, and standard deviation. For each pair of variables: Pearson's correlation coefficient, Spearman's rho, Kendall's tau-*b*, cross-product of deviations, and covariance.

Data. Use symmetric quantitative variables for Pearson's correlation coefficient and quantitative variables or variables with ordered categories for Spearman's rho and Kendall's tau-*b*.

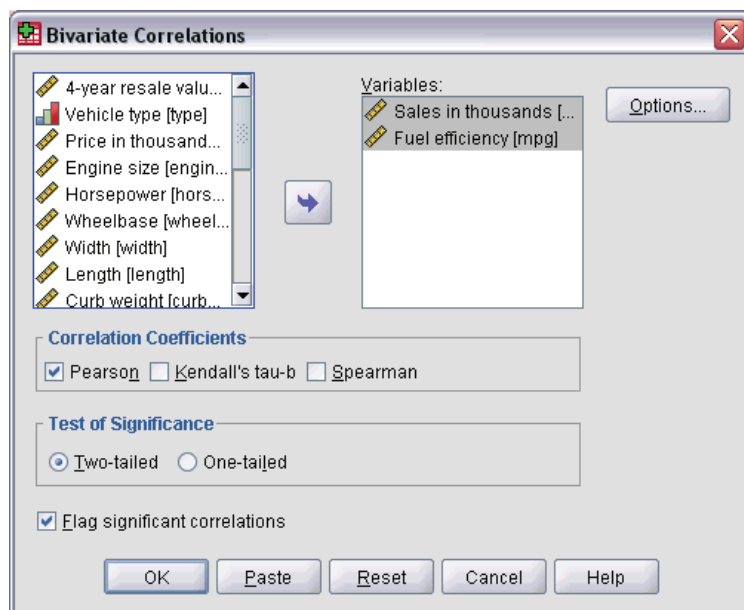
Assumptions. Pearson's correlation coefficient assumes that each pair of variables is bivariate normal.

To Obtain Bivariate Correlations

From the menus choose:

Analyze
 Correlate
 Bivariate...

Figure 25-1
Bivariate Correlations dialog box



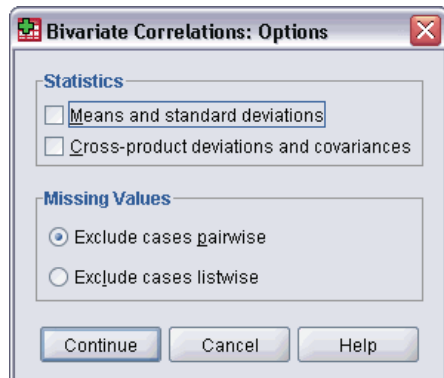
- ▶ Select two or more numeric variables.

The following options are also available:

- **Correlation Coefficients.** For quantitative, normally distributed variables, choose the Pearson correlation coefficient. If your data are not normally distributed or have ordered categories, choose Kendall's tau-b or Spearman, which measure the association between rank orders. Correlation coefficients range in value from -1 (a perfect negative relationship) and $+1$ (a perfect positive relationship). A value of 0 indicates no linear relationship. When interpreting your results, be careful not to draw any cause-and-effect conclusions due to a significant correlation.
- **Test of Significance.** You can select two-tailed or one-tailed probabilities. If the direction of association is known in advance, select One-tailed. Otherwise, select Two-tailed.
- **Flag significant correlations.** Correlation coefficients significant at the 0.05 level are identified with a single asterisk, and those significant at the 0.01 level are identified with two asterisks.

Bivariate Correlations Options

Figure 25-2
Bivariate Correlations Options dialog box



Statistics. For Pearson correlations, you can choose one or both of the following:

- **Means and standard deviations.** Displayed for each variable. The number of cases with nonmissing values is also shown. Missing values are handled on a variable-by-variable basis regardless of your missing values setting.
- **Cross-product deviations and covariances.** Displayed for each pair of variables. The cross-product of deviations is equal to the sum of the products of mean-corrected variables. This is the numerator of the Pearson correlation coefficient. The covariance is an unstandardized measure of the relationship between two variables, equal to the cross-product deviation divided by $N-1$.

Missing Values. You can choose one of the following:

- **Exclude cases pairwise.** Cases with missing values for one or both of a pair of variables for a correlation coefficient are excluded from the analysis. Since each coefficient is based on all cases that have valid codes on that particular pair of variables, the maximum information available is used in every calculation. This can result in a set of coefficients based on a varying number of cases.
- **Exclude cases listwise.** Cases with missing values for any variable are excluded from all correlations.

CORRELATIONS and NONPAR CORR Command Additional Features

The command syntax language also allows you to:

- Write a correlation matrix for Pearson correlations that can be used in place of raw data to obtain other analyses such as factor analysis (with the `MATRIX` subcommand).
- Obtain correlations of each variable on a list with each variable on a second list (using the keyword `WITH` on the `VARIABLES` subcommand).

See the *Command Syntax Reference* for complete syntax information.

Partial Correlations

The Partial Correlations procedure computes partial correlation coefficients that describe the linear relationship between two variables while controlling for the effects of one or more additional variables. Correlations are measures of linear association. Two variables can be perfectly related, but if the relationship is not linear, a correlation coefficient is not an appropriate statistic for measuring their association.

Example. Is there a relationship between healthcare funding and disease rates? Although you might expect any such relationship to be a negative one, a study reports a significant *positive* correlation: as healthcare funding increases, disease rates appear to increase. Controlling for the rate of visits to healthcare providers, however, virtually eliminates the observed positive correlation. Healthcare funding and disease rates only appear to be positively related because more people have access to healthcare when funding increases, which leads to more reported diseases by doctors and hospitals.

Statistics. For each variable: number of cases with nonmissing values, mean, and standard deviation. Partial and zero-order correlation matrices, with degrees of freedom and significance levels.

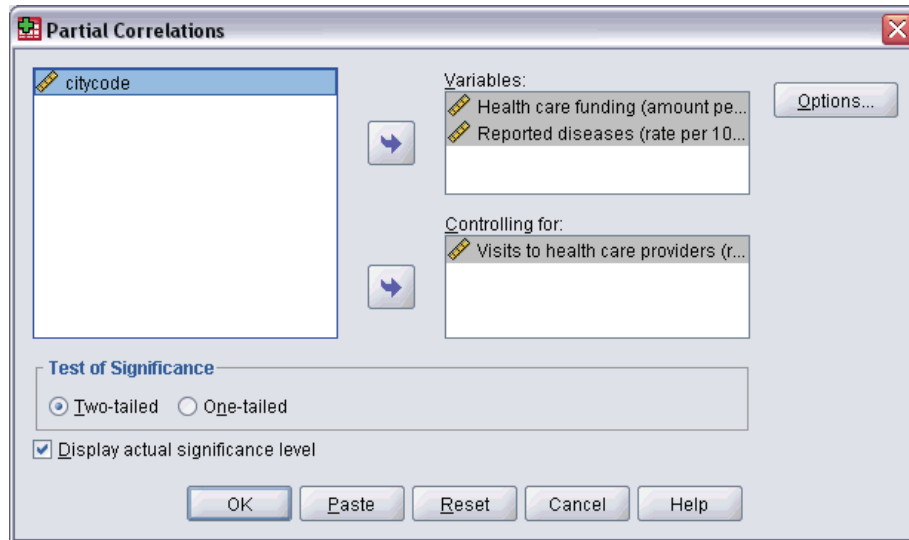
Data. Use symmetric, quantitative variables.

Assumptions. The Partial Correlations procedure assumes that each pair of variables is bivariate normal.

To Obtain Partial Correlations

- ▶ From the menus choose:
 - Analyze
 - Correlate
 - Partial...

Figure 26-1
Partial Correlations dialog box



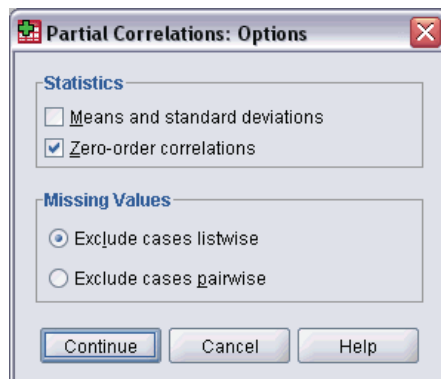
- ▶ Select two or more numeric variables for which partial correlations are to be computed.
- ▶ Select one or more numeric control variables.

The following options are also available:

- **Test of Significance.** You can select two-tailed or one-tailed probabilities. If the direction of association is known in advance, select One-tailed. Otherwise, select Two-tailed.
- **Display actual significance level.** By default, the probability and degrees of freedom are shown for each correlation coefficient. If you deselect this item, coefficients significant at the 0.05 level are identified with a single asterisk, coefficients significant at the 0.01 level are identified with a double asterisk, and degrees of freedom are suppressed. This setting affects both partial and zero-order correlation matrices.

Partial Correlations Options

Figure 26-2
Partial Correlations Options dialog box



Statistics. You can choose one or both of the following:

- **Means and standard deviations.** Displayed for each variable. The number of cases with nonmissing values is also shown.
- **Zero-order correlations.** A matrix of simple correlations between all variables, including control variables, is displayed.

Missing Values. You can choose one of the following alternatives:

- **Exclude cases listwise.** Cases having missing values for any variable, including a control variable, are excluded from all computations.
- **Exclude cases pairwise.** For computation of the zero-order correlations on which the partial correlations are based, a case having missing values for both or one of a pair of variables is not used. Pairwise deletion uses as much of the data as possible. However, the number of cases may differ across coefficients. When pairwise deletion is in effect, the degrees of freedom for a particular partial coefficient are based on the smallest number of cases used in the calculation of any of the zero-order correlations.

PARTIAL CORR Command Additional Features

The command syntax language also allows you to:

- Read a zero-order correlation matrix or write a partial correlation matrix (with the `MATRIX` subcommand).
- Obtain partial correlations between two lists of variables (using the keyword `WITH` on the `VARIABLES` subcommand).
- Obtain multiple analyses (with multiple `VARIABLES` subcommands).
- Specify order values to request (for example, both first- and second-order partial correlations) when you have two control variables (with the `VARIABLES` subcommand).
- Suppress redundant coefficients (with the `FORMAT` subcommand).
- Display a matrix of simple correlations when some coefficients cannot be computed (with the `STATISTICS` subcommand).

See the *Command Syntax Reference* for complete syntax information.

Distances

This procedure calculates any of a wide variety of statistics measuring either similarities or dissimilarities (distances), either between pairs of variables or between pairs of cases. These similarity or distance measures can then be used with other procedures, such as factor analysis, cluster analysis, or multidimensional scaling, to help analyze complex datasets.

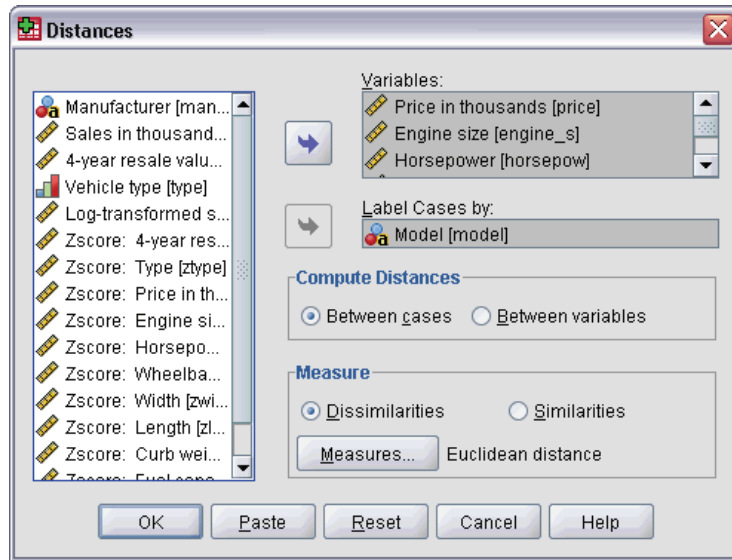
Example. Is it possible to measure similarities between pairs of automobiles based on certain characteristics, such as engine size, MPG, and horsepower? By computing similarities between autos, you can gain a sense of which autos are similar to each other and which are different from each other. For a more formal analysis, you might consider applying a hierarchical cluster analysis or multidimensional scaling to the similarities to explore the underlying structure.

Statistics. Dissimilarity (distance) measures for interval data are Euclidean distance, squared Euclidean distance, Chebychev, block, Minkowski, or customized; for count data, chi-square or phi-square; for binary data, Euclidean distance, squared Euclidean distance, size difference, pattern difference, variance, shape, or Lance and Williams. Similarity measures for interval data are Pearson correlation or cosine; for binary data, Russel and Rao, simple matching, Jaccard, dice, Rogers and Tanimoto, Sokal and Sneath 1, Sokal and Sneath 2, Sokal and Sneath 3, Kulczynski 1, Kulczynski 2, Sokal and Sneath 4, Hamann, Lambda, Anderberg's *D*, Yule's *Y*, Yule's *Q*, Ochiai, Sokal and Sneath 5, phi 4-point correlation, or dispersion.

To Obtain Distance Matrices

- ▶ From the menus choose:
 - Analyze
 - Correlate
 - Distances...

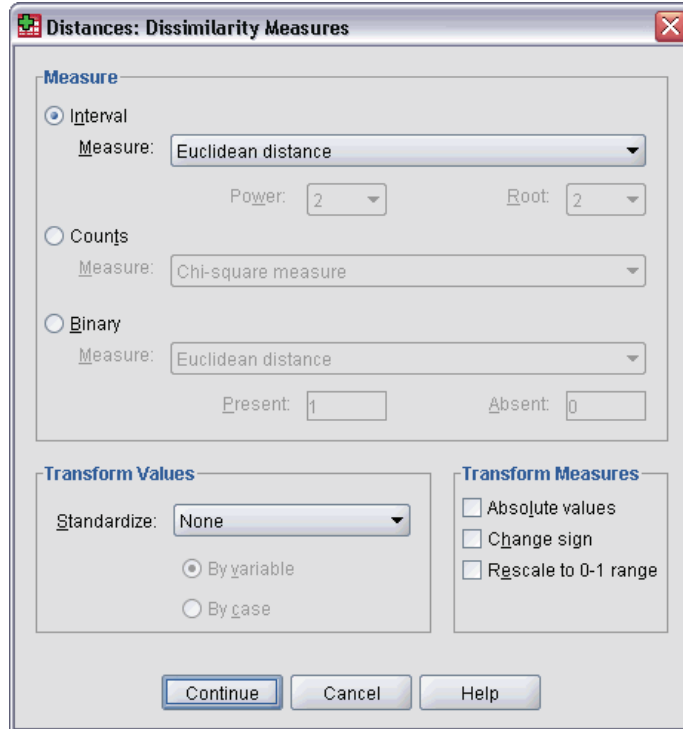
Figure 27-1
Distances dialog box



- ▶ Select at least one numeric variable to compute distances between cases, or select at least two numeric variables to compute distances between variables.
- ▶ Select an alternative in the Compute Distances group to calculate proximities either between cases or between variables.

Distances Dissimilarity Measures

Figure 27-2
Distances Dissimilarity Measures dialog box



From the Measure group, select the alternative that corresponds to your type of data (interval, count, or binary); then, from the drop-down list, select one of the measures that corresponds to that type of data. Available measures, by data type, are:

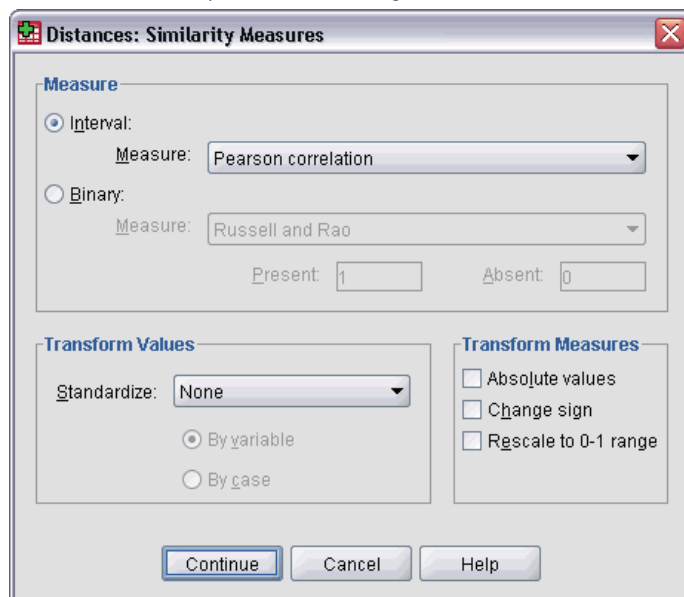
- **Interval data.** Euclidean distance, squared Euclidean distance, Chebychev, block, Minkowski, or customized.
- **Count data.** Chi-square measure or phi-square measure.
- **Binary data.** Euclidean distance, squared Euclidean distance, size difference, pattern difference, variance, shape, or Lance and Williams. (Enter values for Present and Absent to specify which two values are meaningful; Distances will ignore all other values.)

The Transform Values group allows you to standardize data values for either cases or variables *before* computing proximities. These transformations are not applicable to binary data. Available standardization methods are *z* scores, range -1 to 1 , range 0 to 1 , maximum magnitude of 1 , mean of 1 , or standard deviation of 1 .

The Transform Measures group allows you to transform the values generated by the distance measure. They are applied after the distance measure has been computed. Available options are absolute values, change sign, and rescale to $0-1$ range.

Distances Similarity Measures

Figure 27-3
Distances Similarity Measures dialog box



From the Measure group, select the alternative that corresponds to your type of data (interval or binary); then, from the drop-down list, select one of the measures that corresponds to that type of data. Available measures, by data type, are:

- **Interval data.** Pearson correlation or cosine.
- **Binary data.** Russell and Rao, simple matching, Jaccard, Dice, Rogers and Tanimoto, Sokal and Sneath 1, Sokal and Sneath 2, Sokal and Sneath 3, Kulczynski 1, Kulczynski 2, Sokal and Sneath 4, Hamann, Lambda, Anderberg's *D*, Yule's *Y*, Yule's *Q*, Ochiai, Sokal and Sneath 5, phi 4-point correlation, or dispersion. (Enter values for Present and Absent to specify which two values are meaningful; Distances will ignore all other values.)

The Transform Values group allows you to standardize data values for either cases or variables before computing proximities. These transformations are not applicable to binary data. Available standardization methods are *z* scores, range -1 to 1 , range 0 to 1 , maximum magnitude of 1 , mean of 1 , and standard deviation of 1 .

The Transform Measures group allows you to transform the values generated by the distance measure. They are applied after the distance measure has been computed. Available options are absolute values, change sign, and rescale to $0-1$ range.

PROXIMITIES Command Additional Features

The Distances procedure uses PROXIMITIES command syntax. The command syntax language also allows you to:

- Specify any integer as the power for the Minkowski distance measure.
- Specify any integers as the power and root for a customized distance measure.

See the *Command Syntax Reference* for complete syntax information.

Linear Regression

Linear Regression estimates the coefficients of the linear equation, involving one or more independent variables, that best predict the value of the dependent variable. For example, you can try to predict a salesperson's total yearly sales (the dependent variable) from independent variables such as age, education, and years of experience.

Example. Is the number of games won by a basketball team in a season related to the average number of points the team scores per game? A scatterplot indicates that these variables are linearly related. The number of games won and the average number of points scored by the opponent are also linearly related. These variables have a negative relationship. As the number of games won increases, the average number of points scored by the opponent decreases. With linear regression, you can model the relationship of these variables. A good model can be used to predict how many games teams will win.

Statistics. For each variable: number of valid cases, mean, and standard deviation. For each model: regression coefficients, correlation matrix, part and partial correlations, multiple R , R^2 , adjusted R^2 , change in R^2 , standard error of the estimate, analysis-of-variance table, predicted values, and residuals. Also, 95%-confidence intervals for each regression coefficient, variance-covariance matrix, variance inflation factor, tolerance, Durbin-Watson test, distance measures (Mahalanobis, Cook, and leverage values), DfBeta, DfFit, prediction intervals, and casewise diagnostics. Plots: scatterplots, partial plots, histograms, and normal probability plots.

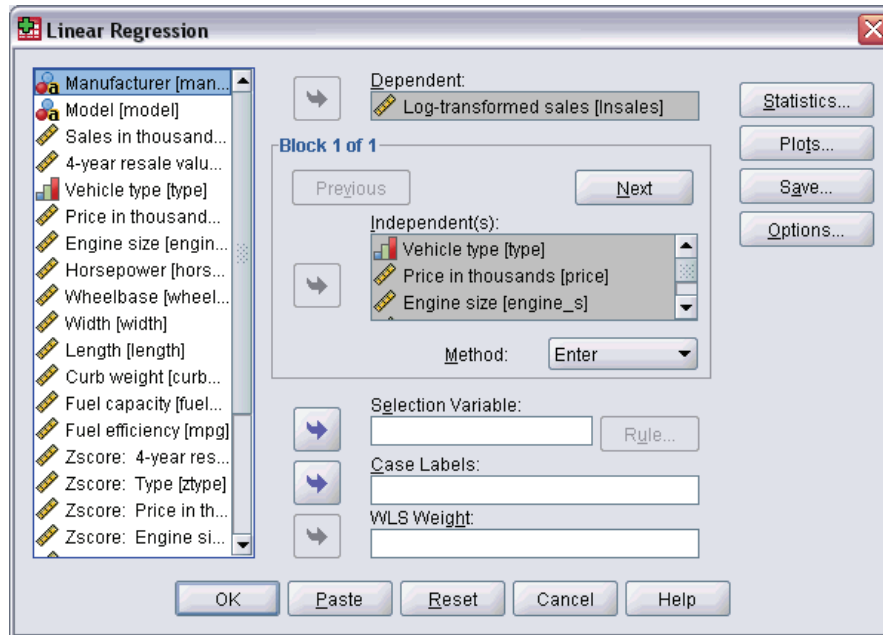
Data. The dependent and independent variables should be quantitative. Categorical variables, such as religion, major field of study, or region of residence, need to be recoded to binary (dummy) variables or other types of contrast variables.

Assumptions. For each value of the independent variable, the distribution of the dependent variable must be normal. The variance of the distribution of the dependent variable should be constant for all values of the independent variable. The relationship between the dependent variable and each independent variable should be linear, and all observations should be independent.

To Obtain a Linear Regression Analysis

- ▶ From the menu choose:
 - Analyze
 - Regression
 - Linear...

Figure 28-1
Linear Regression dialog box



- ▶ In the Linear Regression dialog box, select a numeric dependent variable.
- ▶ Select one or more numeric independent variables.

Optionally, you can:

- Group independent variables into blocks and specify different entry methods for different subsets of variables.
- Choose a selection variable to limit the analysis to a subset of cases having a particular value(s) for this variable.
- Select a case identification variable for identifying points on plots.
- Select a numeric WLS Weight variable for a weighted least squares analysis.

WLS. Allows you to obtain a weighted least-squares model. Data points are weighted by the reciprocal of their variances. This means that observations with large variances have less impact on the analysis than observations associated with small variances. If the value of the weighting variable is zero, negative, or missing, the case is excluded from the analysis.

Linear Regression Variable Selection Methods

Method selection allows you to specify how independent variables are entered into the analysis. Using different methods, you can construct a variety of regression models from the same set of variables.

- **Enter (Regression).** A procedure for variable selection in which all variables in a block are entered in a single step.

- **Stepwise.** At each step, the independent variable not in the equation that has the smallest probability of F is entered, if that probability is sufficiently small. Variables already in the regression equation are removed if their probability of F becomes sufficiently large. The method terminates when no more variables are eligible for inclusion or removal.
- **Remove.** A procedure for variable selection in which all variables in a block are removed in a single step.
- **Backward Elimination.** A variable selection procedure in which all variables are entered into the equation and then sequentially removed. The variable with the smallest partial correlation with the dependent variable is considered first for removal. If it meets the criterion for elimination, it is removed. After the first variable is removed, the variable remaining in the equation with the smallest partial correlation is considered next. The procedure stops when there are no variables in the equation that satisfy the removal criteria.
- **Forward Selection.** A stepwise variable selection procedure in which variables are sequentially entered into the model. The first variable considered for entry into the equation is the one with the largest positive or negative correlation with the dependent variable. This variable is entered into the equation only if it satisfies the criterion for entry. If the first variable is entered, the independent variable not in the equation that has the largest partial correlation is considered next. The procedure stops when there are no variables that meet the entry criterion.

The significance values in your output are based on fitting a single model. Therefore, the significance values are generally invalid when a stepwise method (stepwise, forward, or backward) is used.

All variables must pass the tolerance criterion to be entered in the equation, regardless of the entry method specified. The default tolerance level is 0.0001. Also, a variable is not entered if it would cause the tolerance of another variable already in the model to drop below the tolerance criterion.

All independent variables selected are added to a single regression model. However, you can specify different entry methods for different subsets of variables. For example, you can enter one block of variables into the regression model using stepwise selection and a second block using forward selection. To add a second block of variables to the regression model, click Next.

Linear Regression Set Rule

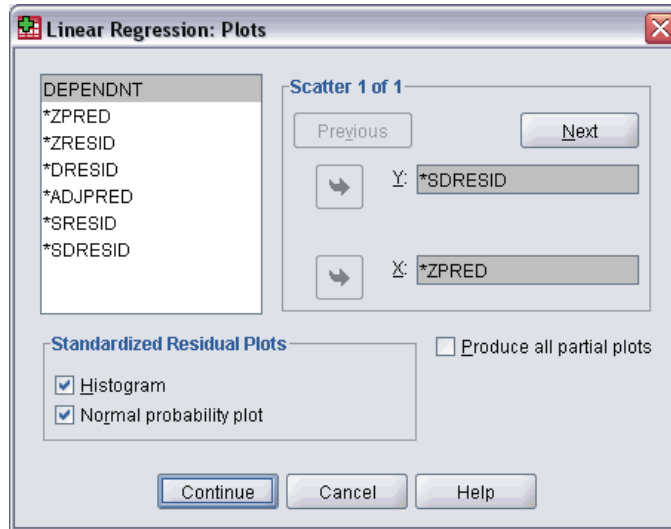
Figure 28-2
Linear Regression Set Rule dialog box



Cases defined by the selection rule are included in the analysis. For example, if you select a variable, choose equals, and type 5 for the value, then only cases for which the selected variable has a value equal to 5 are included in the analysis. A string value is also permitted.

Linear Regression Plots

Figure 28-3
Linear Regression Plots dialog box



Plots can aid in the validation of the assumptions of normality, linearity, and equality of variances. Plots are also useful for detecting outliers, unusual observations, and influential cases. After saving them as new variables, predicted values, residuals, and other diagnostics are available in the Data Editor for constructing plots with the independent variables. The following plots are available:

Scatterplots. You can plot any two of the following: the dependent variable, standardized predicted values, standardized residuals, deleted residuals, adjusted predicted values, Studentized residuals, or Studentized deleted residuals. Plot the standardized residuals against the standardized predicted values to check for linearity and equality of variances.

Source variable list. Lists the dependent variable (DEPENDNT) and the following predicted and residual variables: Standardized predicted values (*ZPRED), Standardized residuals (*ZRESID), Deleted residuals (*DRESID), Adjusted predicted values (*ADJPRED), Studentized residuals (*SRESID), Studentized deleted residuals (*SDRESID).

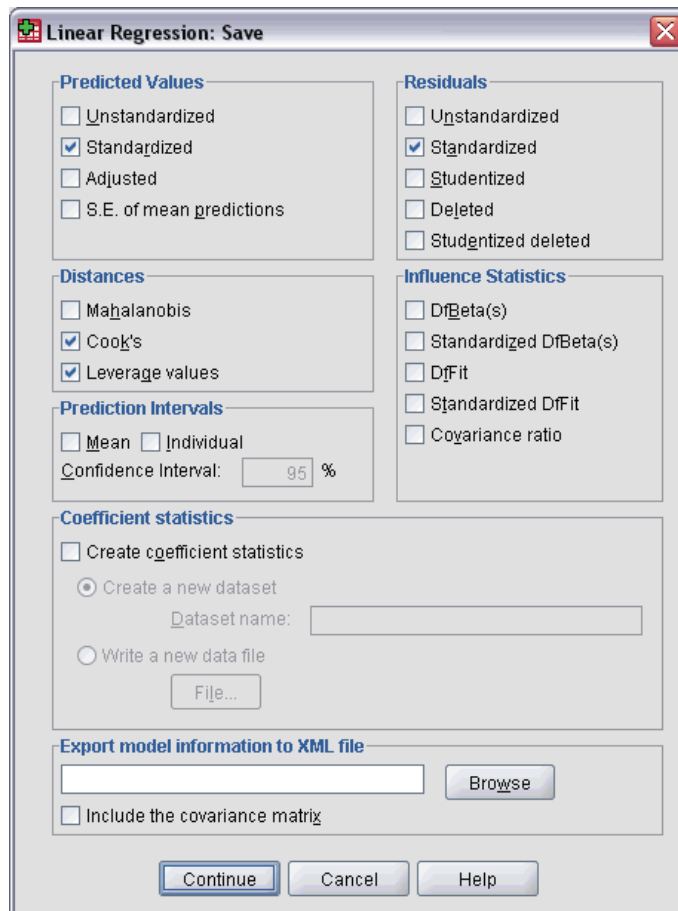
Produce all partial plots. Displays scatterplots of residuals of each independent variable and the residuals of the dependent variable when both variables are regressed separately on the rest of the independent variables. At least two independent variables must be in the equation for a partial plot to be produced.

Standardized Residual Plots. You can obtain histograms of standardized residuals and normal probability plots comparing the distribution of standardized residuals to a normal distribution.

If any plots are requested, summary statistics are displayed for standardized predicted values and standardized residuals (*ZPRED and *ZRESID).

Linear Regression: Saving New Variables

Figure 28-4
Linear Regression Save dialog box



You can save predicted values, residuals, and other statistics useful for diagnostics. Each selection adds one or more new variables to your active data file.

Predicted Values. Values that the regression model predicts for each case.

- **Unstandardized.** The value the model predicts for the dependent variable.
- **Standardized.** A transformation of each predicted value into its standardized form. That is, the mean predicted value is subtracted from the predicted value, and the difference is divided by the standard deviation of the predicted values. Standardized predicted values have a mean of 0 and a standard deviation of 1.
- **Adjusted.** The predicted value for a case when that case is excluded from the calculation of the regression coefficients.
- **S.E. of mean predictions.** Standard errors of the predicted values. An estimate of the standard deviation of the average value of the dependent variable for cases that have the same values of the independent variables.

Distances. Measures to identify cases with unusual combinations of values for the independent variables and cases that may have a large impact on the regression model.

- **Mahalanobis.** A measure of how much a case's values on the independent variables differ from the average of all cases. A large Mahalanobis distance identifies a case as having extreme values on one or more of the independent variables.
- **Cook's.** A measure of how much the residuals of all cases would change if a particular case were excluded from the calculation of the regression coefficients. A large Cook's D indicates that excluding a case from computation of the regression statistics changes the coefficients substantially.
- **Leverage values.** Measures the influence of a point on the fit of the regression. The centered leverage ranges from 0 (no influence on the fit) to $(N-1)/N$.

Prediction Intervals. The upper and lower bounds for both mean and individual prediction intervals.

- **Mean.** Lower and upper bounds (two variables) for the prediction interval of the mean predicted response.
- **Individual.** Lower and upper bounds (two variables) for the prediction interval of the dependent variable for a single case.
- **Confidence Interval.** Enter a value between 1 and 99.99 to specify the confidence level for the two Prediction Intervals. Mean or Individual must be selected before entering this value. Typical confidence interval values are 90, 95, and 99.

Residuals. The actual value of the dependent variable minus the value predicted by the regression equation.

- **Unstandardized.** The difference between an observed value and the value predicted by the model.
- **Standardized.** The residual divided by an estimate of its standard deviation. Standardized residuals, which are also known as Pearson residuals, have a mean of 0 and a standard deviation of 1.
- **Studentized.** The residual divided by an estimate of its standard deviation that varies from case to case, depending on the distance of each case's values on the independent variables from the means of the independent variables.
- **Deleted.** The residual for a case when that case is excluded from the calculation of the regression coefficients. It is the difference between the value of the dependent variable and the adjusted predicted value.
- **Studentized deleted.** The deleted residual for a case divided by its standard error. The difference between a Studentized deleted residual and its associated Studentized residual indicates how much difference eliminating a case makes on its own prediction.

Influence Statistics. The change in the regression coefficients (DfBeta[s]) and predicted values (DfFit) that results from the exclusion of a particular case. Standardized DfBetas and DfFit values are also available along with the covariance ratio.

- **DfBeta(s).** The difference in beta value is the change in the regression coefficient that results from the exclusion of a particular case. A value is computed for each term in the model, including the constant.

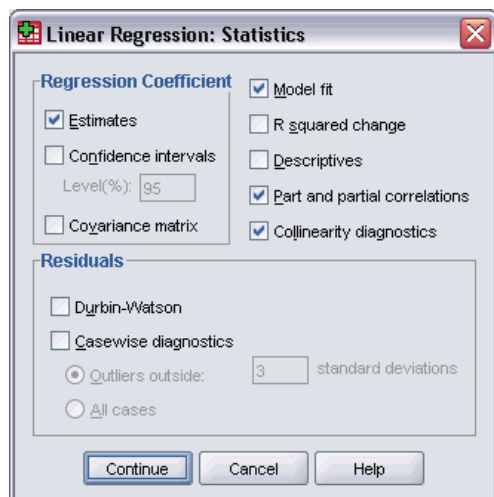
- **Standardized DfBeta.** Standardized difference in beta value. The change in the regression coefficient that results from the exclusion of a particular case. You may want to examine cases with absolute values greater than 2 divided by the square root of N , where N is the number of cases. A value is computed for each term in the model, including the constant.
- **DfFit.** The difference in fit value is the change in the predicted value that results from the exclusion of a particular case.
- **Standardized DfFit.** Standardized difference in fit value. The change in the predicted value that results from the exclusion of a particular case. You may want to examine standardized values which in absolute value exceed 2 times the square root of p/N , where p is the number of parameters in the model and N is the number of cases.
- **Covariance ratio.** The ratio of the determinant of the covariance matrix with a particular case excluded from the calculation of the regression coefficients to the determinant of the covariance matrix with all cases included. If the ratio is close to 1, the case does not significantly alter the covariance matrix.

Coefficient Statistics. Saves regression coefficients to a dataset or a data file. Datasets are available for subsequent use in the same session but are not saved as files unless explicitly saved prior to the end of the session. Dataset names must conform to variable naming rules. [For more information, see Variable Names in Chapter 5 on p. 72.](#)

Export model information to XML file. Parameter estimates and (optionally) their covariances are exported to the specified file in XML (PMML) format. SmartScore and SPSS Statistics Server (a separate product) can use this model file to apply the model information to other data files for scoring purposes.

Linear Regression Statistics

Figure 28-5
Statistics dialog box



The following statistics are available:

Regression Coefficients. Estimates displays Regression coefficient B , standard error of B , standardized coefficient beta, t value for B , and two-tailed significance level of t . Confidence intervals displays confidence intervals with the specified level of confidence for each regression coefficient or a covariance matrix. Covariance matrix displays a variance-covariance matrix of regression coefficients with covariances off the diagonal and variances on the diagonal. A correlation matrix is also displayed.

Model fit. The variables entered and removed from the model are listed, and the following goodness-of-fit statistics are displayed: multiple R , R^2 and adjusted R^2 , standard error of the estimate, and an analysis-of-variance table.

R squared change. The change in the R^2 statistic that is produced by adding or deleting an independent variable. If the R^2 change associated with a variable is large, that means that the variable is a good predictor of the dependent variable.

Descriptives. Provides the number of valid cases, the mean, and the standard deviation for each variable in the analysis. A correlation matrix with a one-tailed significance level and the number of cases for each correlation are also displayed.

Partial Correlation. The correlation that remains between two variables after removing the correlation that is due to their mutual association with the other variables. The correlation between the dependent variable and an independent variable when the linear effects of the other independent variables in the model have been removed from both.

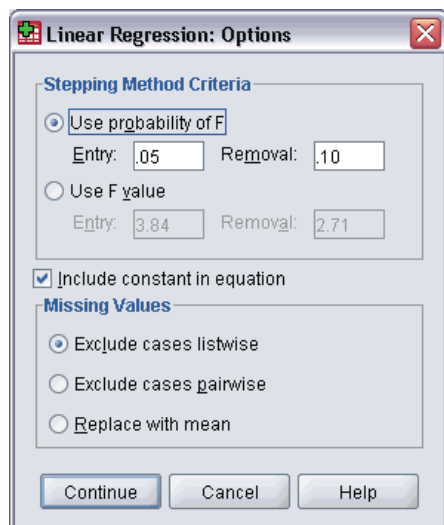
Part Correlation. The correlation between the dependent variable and an independent variable when the linear effects of the other independent variables in the model have been removed from the independent variable. It is related to the change in R-squared when a variable is added to an equation. Sometimes called the semipartial correlation.

Collinearity diagnostics. Collinearity (or multicollinearity) is the undesirable situation when one independent variable is a linear function of other independent variables. Eigenvalues of the scaled and uncentered cross-products matrix, condition indices, and variance-decomposition proportions are displayed along with variance inflation factors (VIF) and tolerances for individual variables.

Residuals. Displays the Durbin-Watson test for serial correlation of the residuals and casewise diagnostics for the cases meeting the selection criterion (outliers above n standard deviations).

Linear Regression Options

Figure 28-6
Linear Regression Options dialog box



The following options are available:

Stepping Method Criteria. These options apply when either the forward, backward, or stepwise variable selection method has been specified. Variables can be entered or removed from the model depending on either the significance (probability) of the F value or the F value itself.

- **Use Probability of F.** A variable is entered into the model if the significance level of its F value is less than the Entry value and is removed if the significance level is greater than the Removal value. Entry must be less than Removal, and both values must be positive. To enter more variables into the model, increase the Entry value. To remove more variables from the model, lower the Removal value.
- **Use F Value.** A variable is entered into the model if its F value is greater than the Entry value and is removed if the F value is less than the Removal value. Entry must be greater than Removal, and both values must be positive. To enter more variables into the model, lower the Entry value. To remove more variables from the model, increase the Removal value.

Include constant in equation. By default, the regression model includes a constant term. Deselecting this option forces regression through the origin, which is rarely done. Some results of regression through the origin are not comparable to results of regression that do include a constant. For example, R^2 cannot be interpreted in the usual way.

Missing Values. You can choose one of the following:

- **Exclude cases listwise.** Only cases with valid values for all variables are included in the analyses.
- **Exclude cases pairwise.** Cases with complete data for the pair of variables being correlated are used to compute the correlation coefficient on which the regression analysis is based. Degrees of freedom are based on the minimum pairwise N .
- **Replace with mean.** All cases are used for computations, with the mean of the variable substituted for missing observations.

REGRESSION Command Additional Features

The command syntax language also allows you to:

- Write a correlation matrix or read a matrix in place of raw data to obtain your regression analysis (with the `MATRIX` subcommand).
- Specify tolerance levels (with the `CRITERIA` subcommand).
- Obtain multiple models for the same or different dependent variables (with the `METHOD` and `DEPENDENT` subcommands).
- Obtain additional statistics (with the `DESCRIPTIVES` and `STATISTICS` subcommands).

See the *Command Syntax Reference* for complete syntax information.

Ordinal Regression

Ordinal Regression allows you to model the dependence of a polytomous ordinal response on a set of predictors, which can be factors or covariates. The design of Ordinal Regression is based on the methodology of McCullagh (1980, 1998), and the procedure is referred to as `PLUM` in the syntax.

Standard linear regression analysis involves minimizing the sum-of-squared differences between a response (dependent) variable and a weighted combination of predictor (independent) variables. The estimated coefficients reflect how changes in the predictors affect the response. The response is assumed to be numerical, in the sense that changes in the level of the response are equivalent throughout the range of the response. For example, the difference in height between a person who is 150 cm tall and a person who is 140 cm tall is 10 cm, which has the same meaning as the difference in height between a person who is 210 cm tall and a person who is 200 cm tall. These relationships do not necessarily hold for ordinal variables, in which the choice and number of response categories can be quite arbitrary.

Example. Ordinal Regression could be used to study patient reaction to drug dosage. The possible reactions may be classified as *none*, *mild*, *moderate*, or *severe*. The difference between a mild and moderate reaction is difficult or impossible to quantify and is based on perception. Moreover, the difference between a mild and moderate response may be greater or less than the difference between a moderate and severe response.

Statistics and plots. Observed and expected frequencies and cumulative frequencies, Pearson residuals for frequencies and cumulative frequencies, observed and expected probabilities, observed and expected cumulative probabilities of each response category by covariate pattern, asymptotic correlation and covariance matrices of parameter estimates, Pearson's chi-square and likelihood-ratio chi-square, goodness-of-fit statistics, iteration history, test of parallel lines assumption, parameter estimates, standard errors, confidence intervals, and Cox and Snell's, Nagelkerke's, and McFadden's R^2 statistics.

Data. The dependent variable is assumed to be ordinal and can be numeric or string. The ordering is determined by sorting the values of the dependent variable in ascending order. The lowest value defines the first category. Factor variables are assumed to be categorical. Covariate variables must be numeric. Note that using more than one continuous covariate can easily result in the creation of a very large cell probabilities table.

Assumptions. Only one response variable is allowed, and it must be specified. Also, for each distinct pattern of values across the independent variables, the responses are assumed to be independent multinomial variables.

Related procedures. Nominal logistic regression uses similar models for nominal dependent variables.

Obtaining an Ordinal Regression

- ▶ From the menus choose:
 - Analyze
 - Regression
 - Ordinal...

Figure 29-1
Ordinal Regression dialog box

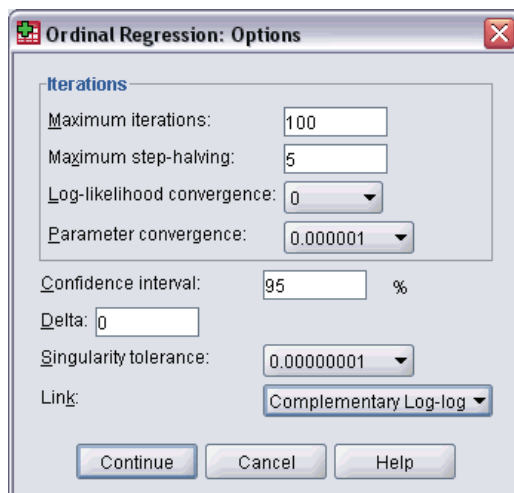


- ▶ Select one dependent variable.
- ▶ Click OK.

Ordinal Regression Options

The Options dialog box allows you to adjust parameters used in the iterative estimation algorithm, choose a level of confidence for your parameter estimates, and select a link function.

Figure 29-2
Ordinal Regression Options dialog box



Iterations. You can customize the iterative algorithm.

- **Maximum iterations.** Specify a non-negative integer. If 0 is specified, the procedure returns the initial estimates.
- **Maximum step-halving.** Specify a positive integer.
- **Log-likelihood convergence.** The algorithm stops if the absolute or relative change in the log-likelihood is less than this value. The criterion is not used if 0 is specified.
- **Parameter convergence.** The algorithm stops if the absolute or relative change in each of the parameter estimates is less than this value. The criterion is not used if 0 is specified.

Confidence interval. Specify a value greater than or equal to 0 and less than 100.

Delta. The value added to zero cell frequencies. Specify a non-negative value less than 1.

Singularity tolerance. Used for checking for highly dependent predictors. Select a value from the list of options.

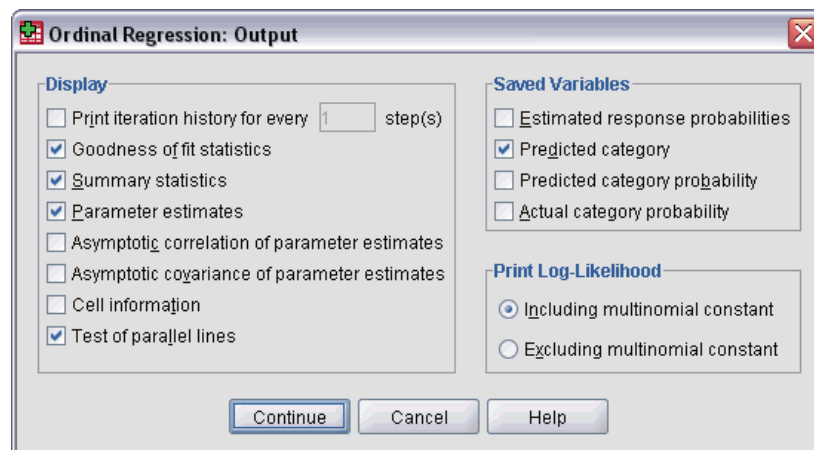
Link function. The link function is a transformation of the cumulative probabilities that allows estimation of the model. Five link functions are available, summarized in the following table.

Function	Form	Typical application
Logit	$\log(\xi / (1-\xi))$	Evenly distributed categories
Complementary log-log	$\log(-\log(1-\xi))$	Higher categories more probable
Negative log-log	$-\log(-\log(\xi))$	Lower categories more probable
Probit	$\Phi^{-1}(\xi)$	Latent variable is normally distributed
Cauchit (inverse Cauchy)	$\tan(\pi(\xi-0.5))$	Latent variable has many extreme values

Ordinal Regression Output

The Output dialog box allows you to produce tables for display in the Viewer and save variables to the working file.

Figure 29-3
Ordinal Regression Output dialog box



Display. Produces tables for:

- **Print iteration history.** The log-likelihood and parameter estimates are printed for the print iteration frequency specified. The first and last iterations are always printed.
- **Goodness-of-fit statistics.** The Pearson and likelihood-ratio chi-square statistics. They are computed based on the classification specified in the variable list.
- **Summary statistics.** Cox and Snell's, Nagelkerke's, and McFadden's R^2 statistics.
- **Parameter estimates.** Parameter estimates, standard errors, and confidence intervals.
- **Asymptotic correlation of parameter estimates.** Matrix of parameter estimate correlations.
- **Asymptotic covariance of parameter estimates.** Matrix of parameter estimate covariances.
- **Cell information.** Observed and expected frequencies and cumulative frequencies, Pearson residuals for frequencies and cumulative frequencies, observed and expected probabilities, and observed and expected cumulative probabilities of each response category by covariate pattern. Note that for models with many covariate patterns (for example, models with continuous covariates), this option can generate a very large, unwieldy table.
- **Test of parallel lines.** Test of the hypothesis that the location parameters are equivalent across the levels of the dependent variable. This is available only for the location-only model.

Saved variables. Saves the following variables to the working file:

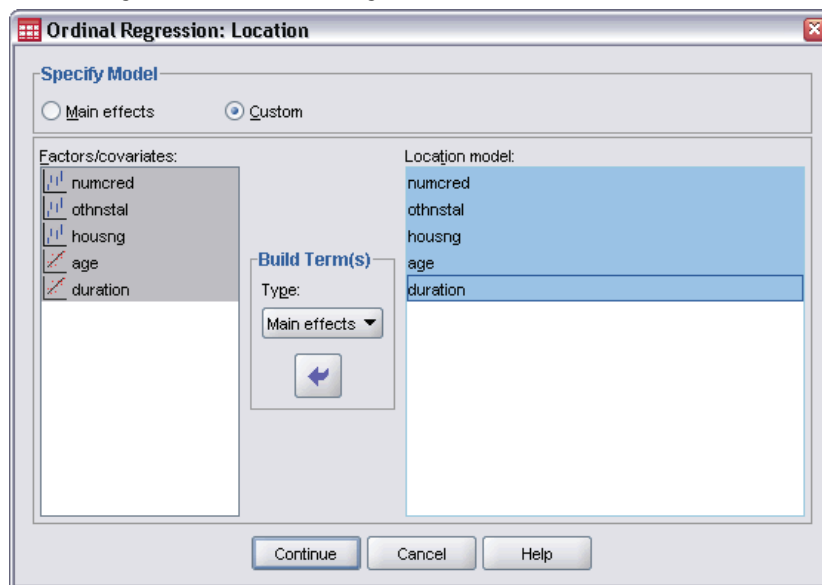
- **Estimated response probabilities.** Model-estimated probabilities of classifying a factor/covariate pattern into the response categories. There are as many probabilities as the number of response categories.
- **Predicted category.** The response category that has the maximum estimated probability for a factor/covariate pattern.
- **Predicted category probability.** Estimated probability of classifying a factor/covariate pattern into the predicted category. This probability is also the maximum of the estimated probabilities of the factor/covariate pattern.
- **Actual category probability.** Estimated probability of classifying a factor/covariate pattern into the actual category.

Print log-likelihood. Controls the display of the log-likelihood. Including multinomial constant gives you the full value of the likelihood. To compare your results across products that do not include the constant, you can choose to exclude it.

Ordinal Regression Location Model

The Location dialog box allows you to specify the location model for your analysis.

Figure 29-4
Ordinal Regression Location dialog box



Specify model. A main-effects model contains the covariate and factor main effects but no interaction effects. You can create a custom model to specify subsets of factor interactions or covariate interactions.

Factors/covariates. The factors and covariates are listed.

Location model. The model depends on the main effects and interaction effects that you select.

Build Terms

For the selected factors and covariates:

Interaction. Creates the highest-level interaction term of all selected variables. This is the default.

Main effects. Creates a main-effects term for each variable selected.

All 2-way. Creates all possible two-way interactions of the selected variables.

All 3-way. Creates all possible three-way interactions of the selected variables.

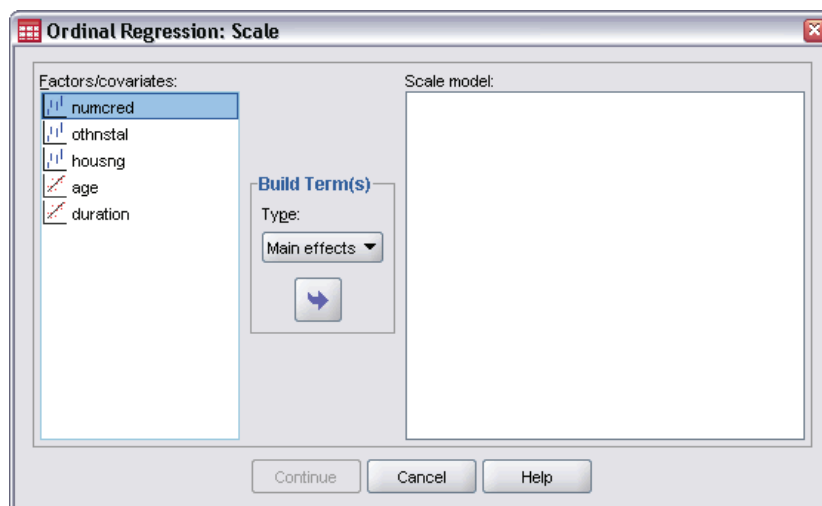
All 4-way. Creates all possible four-way interactions of the selected variables.

All 5-way. Creates all possible five-way interactions of the selected variables.

Ordinal Regression Scale Model

The Scale dialog box allows you to specify the scale model for your analysis.

Figure 29-5
Ordinal Regression Scale dialog box



Factors/covariates. The factors and covariates are listed.

Scale model. The model depends on the main and interaction effects that you select.

Build Terms

For the selected factors and covariates:

Interaction. Creates the highest-level interaction term of all selected variables. This is the default.

Main effects. Creates a main-effects term for each variable selected.

All 2-way. Creates all possible two-way interactions of the selected variables.

All 3-way. Creates all possible three-way interactions of the selected variables.

All 4-way. Creates all possible four-way interactions of the selected variables.

All 5-way. Creates all possible five-way interactions of the selected variables.

PLUM Command Additional Features

You can customize your Ordinal Regression if you paste your selections into a syntax window and edit the resulting PLUM command syntax. The command syntax language also allows you to:

- Create customized hypothesis tests by specifying null hypotheses as linear combinations of parameters.

See the *Command Syntax Reference* for complete syntax information.

Curve Estimation

The Curve Estimation procedure produces curve estimation regression statistics and related plots for 11 different curve estimation regression models. A separate model is produced for each dependent variable. You can also save predicted values, residuals, and prediction intervals as new variables.

Example. An Internet service provider tracks the percentage of virus-infected e-mail traffic on its networks over time. A scatterplot reveals that the relationship is nonlinear. You might fit a quadratic or cubic model to the data and check the validity of assumptions and the goodness of fit of the model.

Statistics. For each model: regression coefficients, multiple R , R^2 , adjusted R^2 , standard error of the estimate, analysis-of-variance table, predicted values, residuals, and prediction intervals. Models: linear, logarithmic, inverse, quadratic, cubic, power, compound, S-curve, logistic, growth, and exponential.

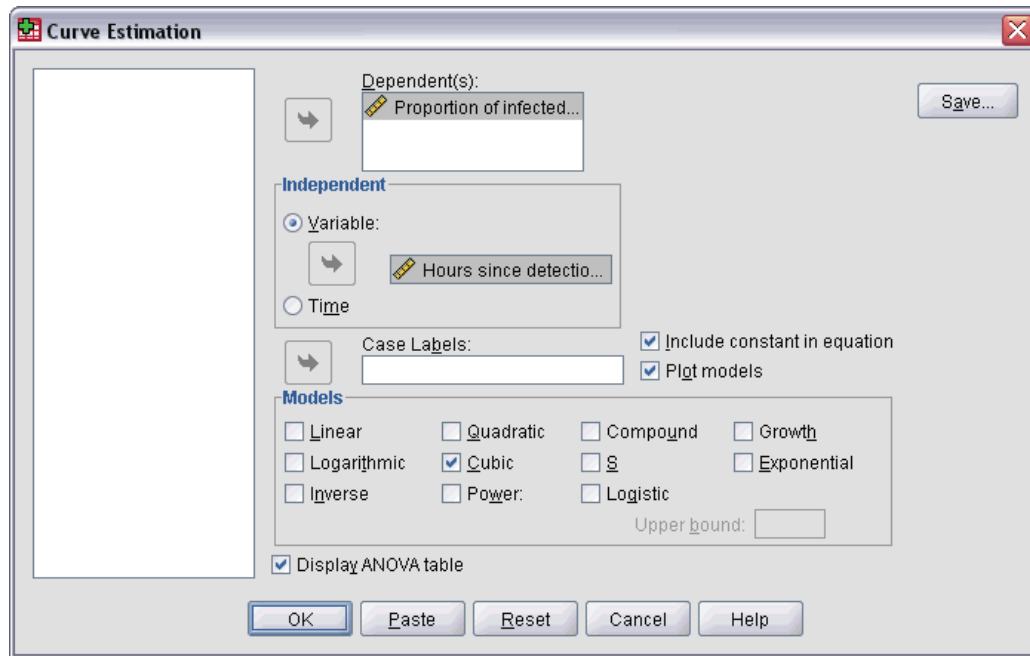
Data. The dependent and independent variables should be quantitative. If you select Time from the active dataset as the independent variable (instead of selecting a variable), the Curve Estimation procedure generates a time variable where the length of time between cases is uniform. If Time is selected, the dependent variable should be a time-series measure. Time-series analysis requires a data file structure in which each case (row) represents a set of observations at a different time and the length of time between cases is uniform.

Assumptions. Screen your data graphically to determine how the independent and dependent variables are related (linearly, exponentially, etc.). The residuals of a good model should be randomly distributed and normal. If a linear model is used, the following assumptions should be met: For each value of the independent variable, the distribution of the dependent variable must be normal. The variance of the distribution of the dependent variable should be constant for all values of the independent variable. The relationship between the dependent variable and the independent variable should be linear, and all observations should be independent.

To Obtain a Curve Estimation

- ▶ From the menu choose:
 - Analyze
 - Regression
 - Curve Estimation...

Figure 30-1
Curve Estimation dialog box



- ▶ Select one or more dependent variables. A separate model is produced for each dependent variable.
- ▶ Select an independent variable (either select a variable in the active dataset or select Time).
- ▶ Optionally:
 - Select a variable for labeling cases in scatterplots. For each point in the scatterplot, you can use the Point Selection tool to display the value of the Case Label variable.
 - Click Save to save predicted values, residuals, and prediction intervals as new variables.

The following options are also available:

- **Include constant in equation.** Estimates a constant term in the regression equation. The constant is included by default.
- **Plot models.** Plots the values of the dependent variable and each selected model against the independent variable. A separate chart is produced for each dependent variable.
- **Display ANOVA table.** Displays a summary analysis-of-variance table for each selected model.

Curve Estimation Models

You can choose one or more curve estimation regression models. To determine which model to use, plot your data. If your variables appear to be related linearly, use a simple linear regression model. When your variables are not linearly related, try transforming your data. When a transformation does not help, you may need a more complicated model. View a scatterplot of your data; if the plot resembles a mathematical function you recognize, fit your data to that type of model. For example, if your data resemble an exponential function, use an exponential model.

Linear. Model whose equation is $Y = b_0 + (b_1 * t)$. The series values are modeled as a linear function of time.

Logarithmic. Model whose equation is $Y = b_0 + (b_1 * \ln(t))$.

Inverse. Model whose equation is $Y = b_0 + (b_1 / t)$.

Quadratic. Model whose equation is $Y = b_0 + (b_1 * t) + (b_2 * t^{**2})$. The quadratic model can be used to model a series that "takes off" or a series that dampens.

Cubic. Model that is defined by the equation $Y = b_0 + (b_1 * t) + (b_2 * t^{**2}) + (b_3 * t^{**3})$.

Power. Model whose equation is $Y = b_0 * (t^{**b_1})$ or $\ln(Y) = \ln(b_0) + (b_1 * \ln(t))$.

Compound. Model whose equation is $Y = b_0 * (b_1^{**t})$ or $\ln(Y) = \ln(b_0) + (\ln(b_1) * t)$.

S-curve. Model whose equation is $Y = e^{**}(b_0 + (b_1/t))$ or $\ln(Y) = b_0 + (b_1/t)$.

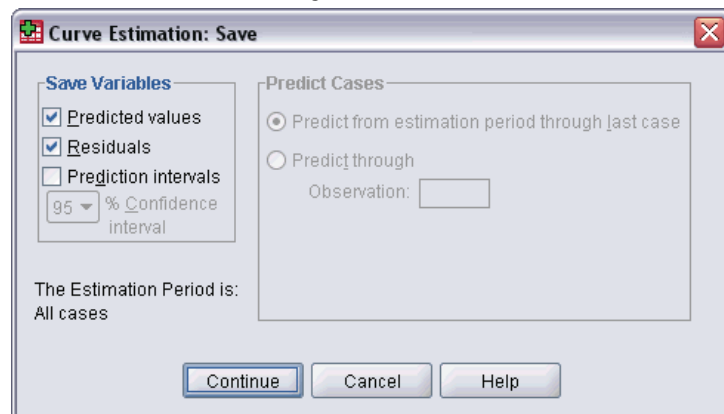
Logistic. Model whose equation is $Y = 1 / (1/u + (b_0 * (b_1^{**t})))$ or $\ln(1/y-1/u) = \ln(b_0) + (\ln(b_1) * t)$ where u is the upper boundary value. After selecting Logistic, specify the upper boundary value to use in the regression equation. The value must be a positive number that is greater than the largest dependent variable value.

Growth. Model whose equation is $Y = e^{**}(b_0 + (b_1 * t))$ or $\ln(Y) = b_0 + (b_1 * t)$.

Exponential. Model whose equation is $Y = b_0 * (e^{**}(b_1 * t))$ or $\ln(Y) = \ln(b_0) + (b_1 * t)$.

Curve Estimation Save

Figure 30-2
Curve Estimation Save dialog box



Save Variables. For each selected model, you can save predicted values, residuals (observed value of the dependent variable minus the model predicted value), and prediction intervals (upper and lower bounds). The new variable names and descriptive labels are displayed in a table in the output window.

Predict Cases. In the active dataset, if you select Time instead of a variable as the independent variable, you can specify a forecast period beyond the end of the time series. You can choose one of the following alternatives:

- **Predict from estimation period through last case.** Predicts values for all cases in the file, based on the cases in the estimation period. The estimation period, displayed at the bottom of the dialog box, is defined with the Range subdialog box of the Select Cases option on the Data menu. If no estimation period has been defined, all cases are used to predict values.
- **Predict through.** Predicts values through the specified date, time, or observation number, based on the cases in the estimation period. This feature can be used to forecast values beyond the last case in the time series. The currently defined date variables determine what text boxes are available for specifying the end of the prediction period. If there are no defined date variables, you can specify the ending observation (case) number.

Use the Define Dates option on the Data menu to create date variables.

Partial Least Squares Regression

The Partial Least Squares Regression procedure estimates partial least squares (PLS, also known as “projection to latent structure”) regression models. PLS is a predictive technique that is an alternative to ordinary least squares (OLS) regression, canonical correlation, or structural equation modeling, and it is particularly useful when predictor variables are highly correlated or when the number of predictors exceeds the number of cases.

PLS combines features of principal components analysis and multiple regression. It first extracts a set of latent factors that explain as much of the covariance as possible between the independent and dependent variables. Then a regression step predicts values of the dependent variables using the decomposition of the independent variables.

Availability. PLS is an extension command that requires the Python Extension Module to be installed on the system where you plan to run PLS. The PLS Extension Module must be installed separately, and the installer can be downloaded from <http://www.spss.com/devcentral>.

Note: The PLS Extension Module is dependent upon Python software. SPSS Inc. is not the owner or licensor of the Python software. Any user of Python must agree to the terms of the Python license agreement located on the Python Web site. SPSS Inc. is not making any statement about the quality of the Python program. SPSS Inc. fully disclaims all liability associated with your use of the Python program.

Tables. Proportion of variance explained (by latent factor), latent factor weights, latent factor loadings, independent variable importance in projection (VIP), and regression parameter estimates (by dependent variable) are all produced by default.

Charts. Variable importance in projection (VIP), factor scores, factor weights for the first three latent factors, and distance to the model are all produced from the [Options](#) tab.

Measurement level. The dependent and independent (predictor) variables can be scale, nominal, or ordinal. The procedure assumes that the appropriate measurement level has been assigned to all variables, although you can temporarily change the measurement level for a variable by right-clicking the variable in the source variable list and selecting a measurement level from the context menu. Categorical (nominal or ordinal) variables are treated equivalently by the procedure.

Categorical variable coding. The procedure temporarily recodes categorical dependent variables using one-of- c coding for the duration of the procedure. If there are c categories of a variable, then the variable is stored as c vectors, with the first category denoted $(1,0,\dots,0)$, the next category $(0,1,0,\dots,0)$, ..., and the final category $(0,0,\dots,0,1)$. Categorical dependent variables are represented using dummy coding; that is, simply omit the indicator corresponding to the reference category.

Frequency weights. Weight values are rounded to the nearest whole number before use. Cases with missing weights or weights less than 0.5 are not used in the analyses.

Missing values. User- and system-missing values are treated as invalid.

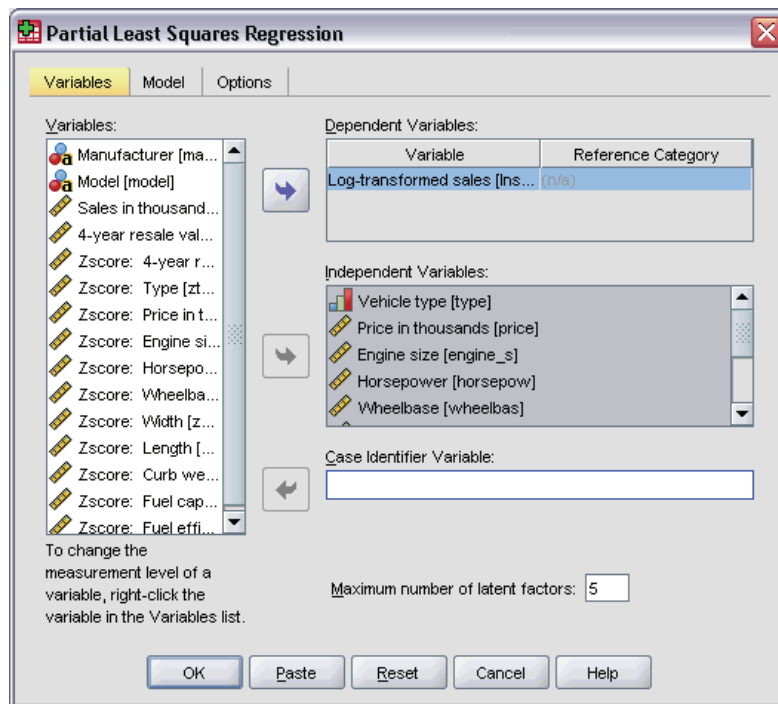
Rescaling. All model variables are centered and standardized, including indicator variables representing categorical variables.

To Obtain Partial Least Squares Regression

From the menus choose:

Analyze
 Regression
 Partial Least Squares...

Figure 31-1
 Partial Least Squares Regression Variables tab



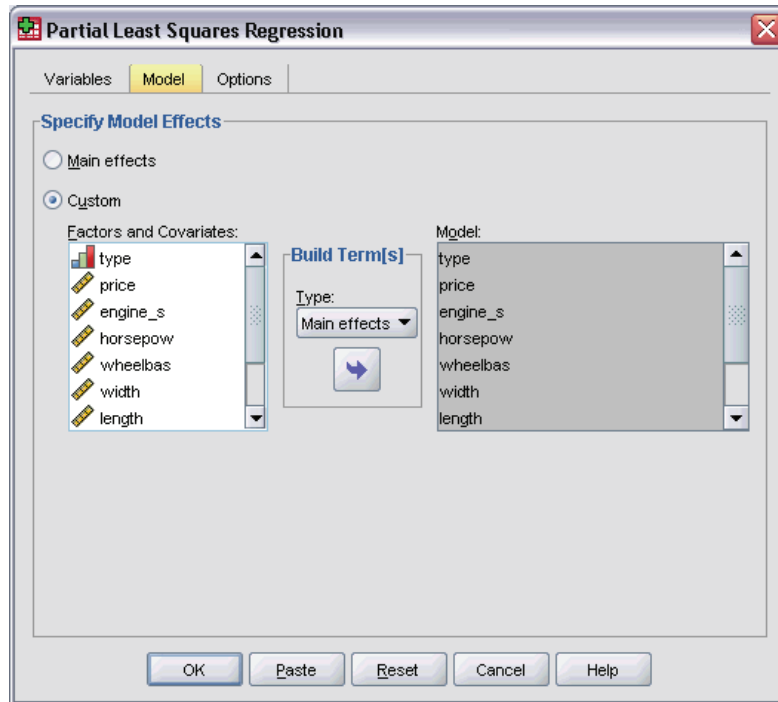
- ▶ Select at least one dependent variable.
- ▶ Select at least one independent variable.

Optionally, you can:

- Specify a reference category for categorical (nominal or ordinal) dependent variables.
- Specify a variable to be used as a unique identifier for casewise output and saved datasets.
- Specify an upper limit on the number of latent factors to be extracted.

Model

Figure 31-2
Partial Least Squares Regression Model tab



Specify Model Effects. A main-effects model contains all factor and covariate main effects. Select Custom to specify interactions. You must indicate all of the terms to be included in the model.

Factors and Covariates. The factors and covariates are listed.

Model. The model depends on the nature of your data. After selecting Custom, you can select the main effects and interactions that are of interest in your analysis.

Build Terms

For the selected factors and covariates:

Interaction. Creates the highest-level interaction term of all selected variables. This is the default.

Main effects. Creates a main-effects term for each variable selected.

All 2-way. Creates all possible two-way interactions of the selected variables.

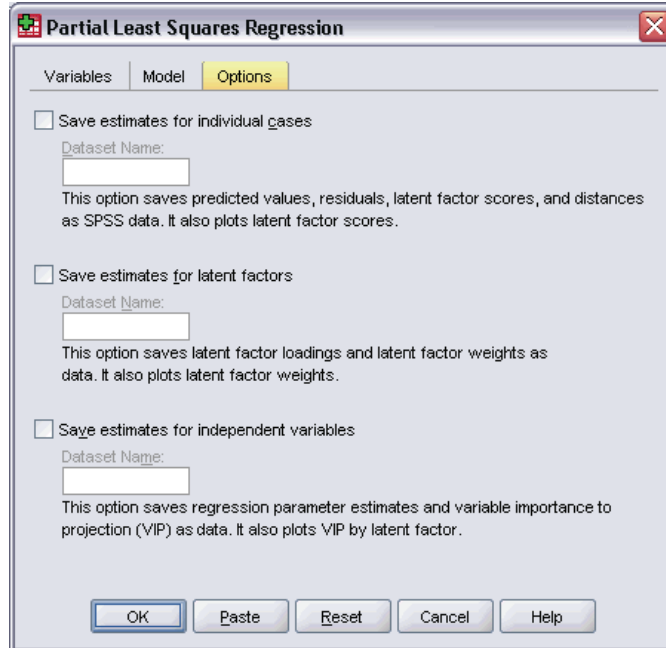
All 3-way. Creates all possible three-way interactions of the selected variables.

All 4-way. Creates all possible four-way interactions of the selected variables.

All 5-way. Creates all possible five-way interactions of the selected variables.

Options

Figure 31-3
Partial Least Squares Regression Options tab



The Options tab allows the user to save and plot model estimates for individual cases, latent factors, and predictors.

For each type of data, specify the name of a dataset. The dataset names must be unique. If you specify the name of an existing dataset, its contents are replaced; otherwise, a new dataset is created.

- **Save estimates for individual cases.** Saves the following casewise model estimates: predicted values, residuals, distance to latent factor model, and latent factor scores. It also plots latent factor scores.
- **Save estimates for latent factors.** Saves latent factor loadings and latent factor weights. It also plots latent factor weights.
- **Save estimates for independent variables.** Saves regression parameter estimates and variable importance to projection (VIP). It also plots VIP by latent factor.


Nearest Neighbor Analysis

Nearest Neighbor Analysis is a method for classifying cases based on their similarity to other cases. In machine learning, it was developed as a way to recognize patterns of data without requiring an exact match to any stored patterns, or cases. Similar cases are near each other and dissimilar cases are distant from each other. Thus, the distance between two cases is a measure of their dissimilarity.

Cases that are near each other are said to be “neighbors.” When a new case (holdout) is presented, its distance from each of the cases in the model is computed. The classifications of the most similar cases – the nearest neighbors – are tallied and the new case is placed into the category that contains the greatest number of nearest neighbors.

You can specify the number of nearest neighbors to examine; this value is called k . The pictures show how a new case would be classified using two different values of k . When $k = 5$, the new case is placed in category 1 because a majority of the nearest neighbors belong to category 1. However, when $k = 9$, the new case is placed in category 0 because a majority of the nearest neighbors belong to category 0.

Figure 32-1
The effects of changing k on classification



Nearest neighbor analysis can also be used to compute values for a continuous target. In this situation, the average or median target value of the nearest neighbors is used to obtain the predicted value for the new case.












Target and Features. The target and features can be:

- **Nominal.** A variable can be treated as nominal when its values represent categories with no intrinsic ranking (for example, the department of the company in which an employee works). Examples of nominal variables include region, zip code, and religious affiliation.

- **Ordinal.** A variable can be treated as ordinal when its values represent categories with some intrinsic ranking (for example, levels of service satisfaction from highly dissatisfied to highly satisfied). Examples of ordinal variables include attitude scores representing degree of satisfaction or confidence and preference rating scores.
- **Scale.** A variable can be treated as scale when its values represent ordered categories with a meaningful metric, so that distance comparisons between values are appropriate. Examples of scale variables include age in years and income in thousands of dollars.

Nominal and Ordinal variables are treated equivalently by Nearest Neighbor Analysis. The procedure assumes that the appropriate measurement level has been assigned to each variable; however, you can temporarily change the measurement level for a variable by right-clicking the variable in the source variable list and selecting a measurement level from the context menu. To permanently change the level of measurement for a variable, see [Variable Measurement Level](#).

An icon next to each variable in the variable list identifies the measurement level and data type:

Measurement Level	Data Type			
	Numeric	String	Date	Time
Scale		n/a		
Ordinal				
Nominal				

Categorical variable coding. The procedure temporarily recodes categorical predictors and dependent variables using one-of- c coding for the duration of the procedure. If there are c categories of a variable, then the variable is stored as c vectors, with the first category denoted $(1,0,\dots,0)$, the next category $(0,1,0,\dots,0)$, ..., and the final category $(0,0,\dots,0,1)$.

This coding scheme increases the dimensionality of the feature space. In particular, the total number of dimensions is the number of scale predictors plus the number of categories across all categorical predictors. As a result, this coding scheme can lead to slower training. If your nearest neighbors training is proceeding very slowly, you might try reducing the number of categories in your categorical predictors by combining similar categories or dropping cases that have extremely rare categories before running the procedure.

All one-of- c coding is based on the training data, even if a holdout sample is defined (see [Partitions](#)). Thus, if the holdout sample contains cases with predictor categories that are not present in the training data, then those cases are not scored. If the holdout sample contains cases with dependent variable categories that are not present in the training data, then those cases are scored.

Rescaling. Scale features are normalized by default. All rescaling is performed based on the training data, even if a holdout sample is defined (see [Partitions](#) on p. 373). If you specify a variable to define partitions, it is important that the features have similar distributions across

the training and holdout samples. Use, for example, the [Explore](#) procedure to examine the distributions across partitions.

Frequency weights. Frequency weights are ignored by this procedure.

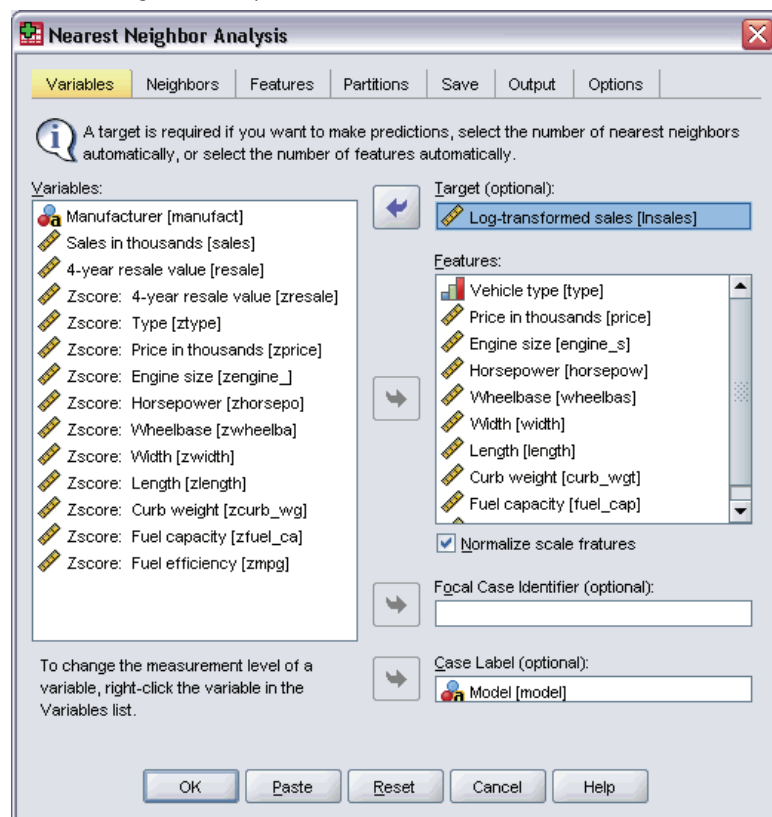
Replicating results. The procedure uses random number generation during random assignment of partitions and cross-validation folds. If you want to replicate your results exactly, in addition to using the same procedure settings, set a seed for the Mersenne Twister (see [Partitions](#) on p. 373), or use variables to define partitions and cross-validation folds.

To Obtain a Nearest Neighbor Analysis

From the menus choose:

```
Analyze
  Classify
    Nearest Neighbor...
```

Figure 32-2
Nearest Neighbor Analysis Variables tab



- Specify one or more features, which can be thought of independent variables or predictors if there is a target.

Target (optional). If no target (dependent variable or response) is specified, then the procedure finds the k nearest neighbors only – no classification or prediction is done.

Normalize scale features. Normalized features have the same range of values, which can improve the performance of the estimation algorithm. Adjusted normalization, $[2*(x-\min)/(\max-\min)]-1$, is used. Adjusted normalized values fall between -1 and 1 .

Focal case identifier (optional). This allows you to mark cases of particular interest. For example, a researcher wants to determine whether the test scores from one school district – the focal case – are comparable to those from similar school districts. He uses nearest neighbor analysis to find the school districts that are most similar with respect to a given set of features. Then he compares the test scores from the focal school district to those from the nearest neighbors.

Focal cases could also be used in clinical studies to select control cases that are similar to clinical cases. Focal cases are displayed in the k nearest neighbors and distances table, feature space chart, peers chart, and quadrant map. Information on focal cases is saved to the files specified on the Output tab.

Cases with a positive value on the specified variable are treated as focal cases. It is invalid to specify a variable with no positive values.

Case label (optional). Cases are labeled using these values in the feature space chart, peers chart, and quadrant map.

Neighbors

Figure 32-3
Nearest Neighbor Analysis Neighbors tab

Number of Nearest Neighbors (k). Specify the number of nearest neighbors. Note that using a greater number of neighbors will not necessarily result in a more accurate model.

If a target is specified on the Variables tab, you can alternatively specify a range of values and allow the procedure to choose the “best” number of neighbors within that range. The method for determining the number of nearest neighbors depends upon whether feature selection is requested on the Features tab.

- If feature selection is in effect, then feature selection is performed for each value of k in the requested range, and the k , and accompanying feature set, with the lowest error rate (or the lowest sum-of-squares error if the target is scale) is selected.
- If feature selection is not in effect, then V -fold cross-validation is used to select the “best” number of neighbors. See the Partition tab for control over assignment of folds.

Distance Computation. This is the metric used to specify the distance metric used to measure the similarity of cases.

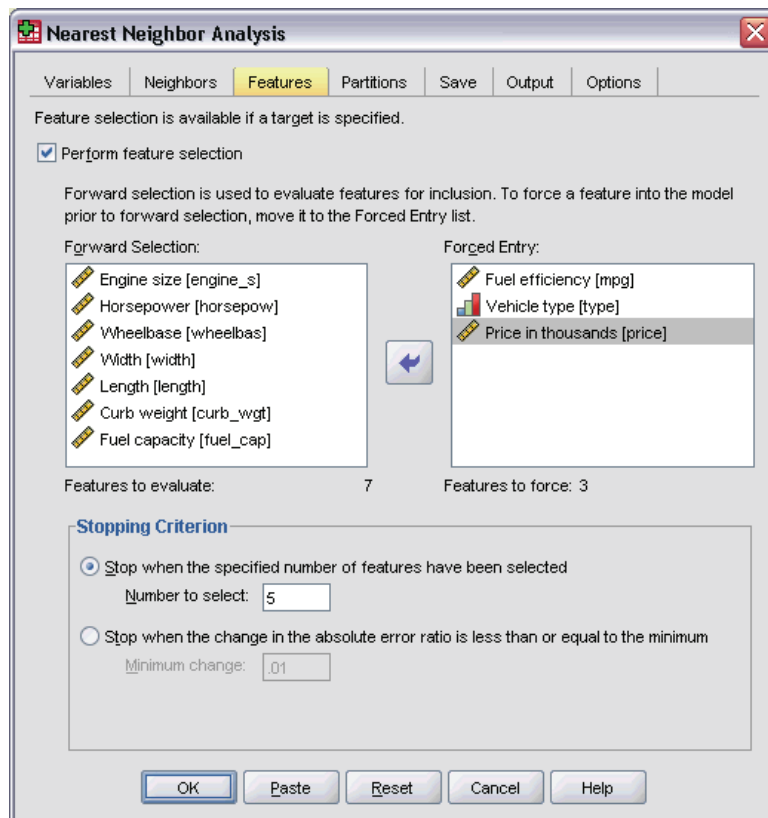
- **Euclidean metric.** The distance between two cases, x and y , is the square root of the sum, over all dimensions, of the squared differences between the values for the cases.
- **City block metric.** The distance between two cases is the sum, over all dimensions, of the absolute differences between the values for the cases. Also called Manhattan distance.

Optionally, if a target is specified on the Variables tab, you can choose to weight features by their normalized importance when computing distances. Feature importance for a predictor is calculated by the ratio of the error rate or sum-of-squares error of the model with the predictor removed from the model to the error rate or sum-of-squares error for the full model. Normalized importance is calculated by reweighting the feature importance values so that they sum to 1.

Predictions for Scale Target. If a scale target is specified on the Variables tab, this specifies whether the predicted value is computed based upon the mean or the median value of the nearest neighbors.

Features

Figure 32-4
Nearest Neighbor Analysis Features tab



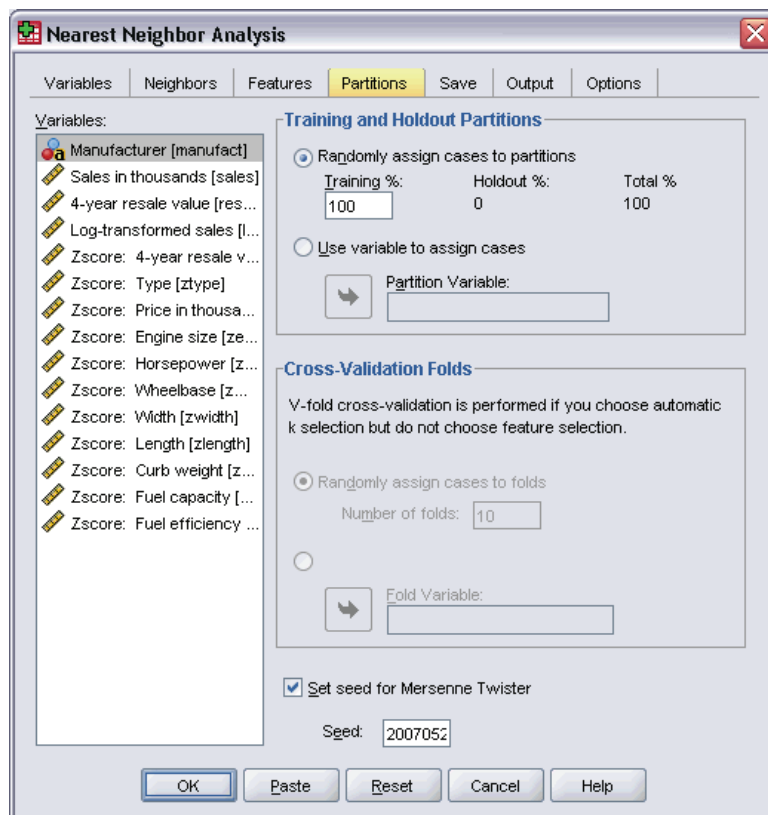
The Features tab allows you to request and specify options for feature selection when a target is specified on the Variables tab. By default, all features are considered for feature selection, but you can optionally select a subset of features to force into the model.

Stopping Criterion. At each step, the feature whose addition to the model results in the smallest error (computed as the error rate for a categorical target and sum of squares error for a scale target) is considered for inclusion in the model set. Forward selection continues until the specified condition is met.

- **Specified Number of Features.** The algorithm adds a fixed number of features in addition to those forced into the model. Specify a positive integer. Decreasing values of the number to select creates a more parsimonious model, at the risk of missing important features. Increasing values of the number to select will capture all the important features, at the risk of eventually adding features that actually increase the model error.
- **Minimum Change in Absolute Error Ratio.** The algorithm stops when the change in the absolute error ratio indicates that the model cannot be further improved by adding more features. Specify a positive number. Decreasing values of the minimum change will tend to include more features, at the risk of including features that don't add much value to the model. Increasing the value of the minimum change will tend to exclude more features, at the risk of losing features that are important to the model. The "optimal" value of the minimum change will depend upon your data and application. See the Feature Selection Error Log in the output to help you assess which features are most important. [For more information, see Feature Selection Error Log on p. 383.](#)

Partitions

Figure 32-5
Nearest Neighbor Analysis Partitions tab



The Partitions tab allows you to divide the dataset into training and holdout sets and, when applicable, assign cases into cross-validation folds

Training and Holdout Partitions. This group specifies the method of partitioning the active dataset into training and holdout samples. The **training sample** comprises the data records used to train the nearest neighbor model; some percentage of cases in the dataset must be assigned to the training sample in order to obtain a model. The **holdout sample** is an independent set of data records used to assess the final model; the error for the holdout sample gives an “honest” estimate of the predictive ability of the model because the holdout cases were not used to build the model.

- **Randomly assign cases to partitions.** Specify the percentage of cases to assign to the training sample. The rest are assigned to the holdout sample.
- **Use variable to assign cases.** Specify a numeric variable that assigns each case in the active dataset to the training or holdout sample. Cases with a positive value on the variable are assigned to the training sample, cases with a value of 0 or a negative value, to the holdout sample. Cases with a system-missing value are excluded from the analysis. Any user-missing values for the partition variable are always treated as valid.

Cross-Validation Folds. V -fold cross-validation is used to determine the “best” number of neighbors. It is not available in conjunction with feature selection for performance reasons.

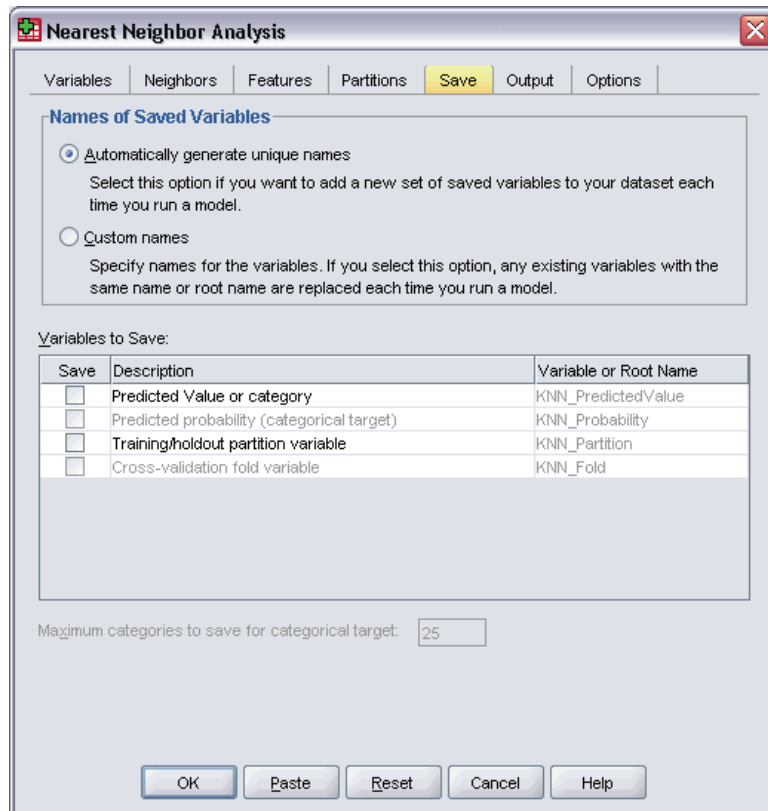
Cross-validation divides the sample into a number of subsamples, or folds. Nearest neighbor models are then generated, excluding the data from each subsample in turn. The first model is based on all of the cases except those in the first sample fold, the second model is based on all of the cases except those in the second sample fold, and so on. For each model, the error is estimated by applying the model to the subsample excluded in generating it. The “best” number of nearest neighbors is the one which produces the lowest error across folds.

- **Randomly assign cases to folds.** Specify the number of folds that should be used for cross-validation. The procedure randomly assigns cases to folds, numbered from 1 to V , the number of folds.
- **Use variable to assign cases.** Specify a numeric variable that assigns each case in the active dataset to a fold. The variable must be numeric and take values from 1 to V . If any values in this range are missing, and on any splits if split files are in effect, this will cause an error. [For more information, see Split File in Chapter 9 on p. 179.](#)

Set seed for Mersenne Twister. Setting a seed allows you to replicate analyses. This control is essentially a shortcut to setting the Mersenne Twister as the active generator and specifying a fixed starting point on the Random Number Generators dialog. [For more information, see Random Number Generators in Chapter 8 on p. 128.](#)

Save

Figure 32-6
Nearest Neighbor Analysis Save tab



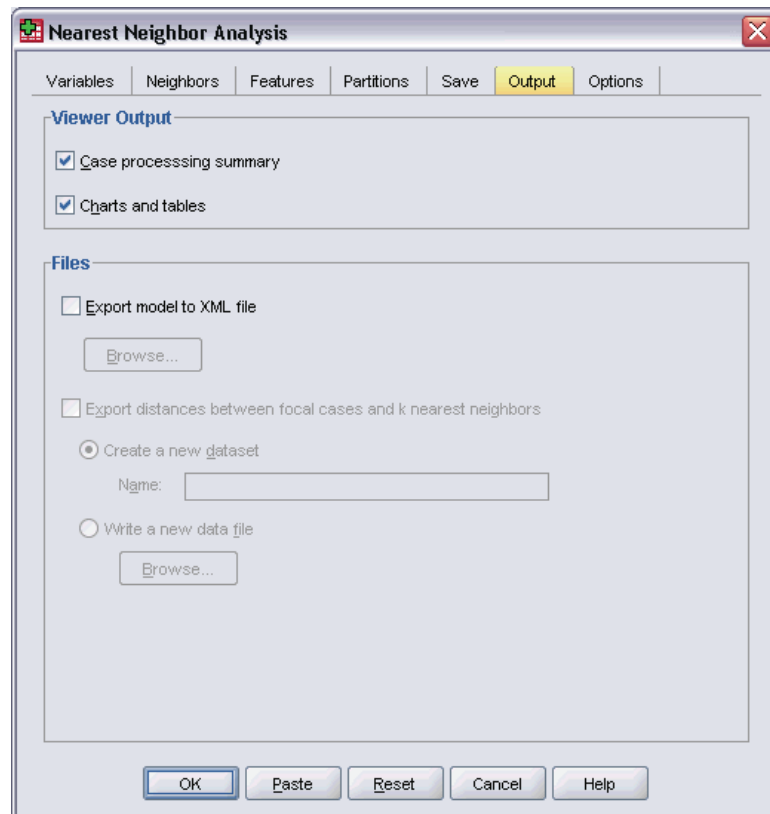
Names of Saved Variables. Automatic name generation ensures that you keep all of your work. Custom names allow you to discard/replace results from previous runs without first deleting the saved variables in the Data Editor.

Variables to Save

- **Predicted value or category.** This saves the predicted value for a scale target or the predicted category for a categorical target.
- **Predicted probability.** This saves the predicted probabilities for a categorical target. A separate variable is saved for each of the first n categories, where n is specified in the Maximum categories to save for categorical target control.
- **Training/holdout partition variables.** If cases are randomly assigned to the training and holdout samples on the Partitions tab, this saves the value of the partition (training or holdout) to which the case was assigned.
- **Cross-validation fold variable.** If cases are randomly assigned to cross-validation folds on the Partitions tab, this saves the value of the fold to which the case was assigned.

Output

Figure 32-7
Nearest Neighbor Analysis Output tab



Viewer Output

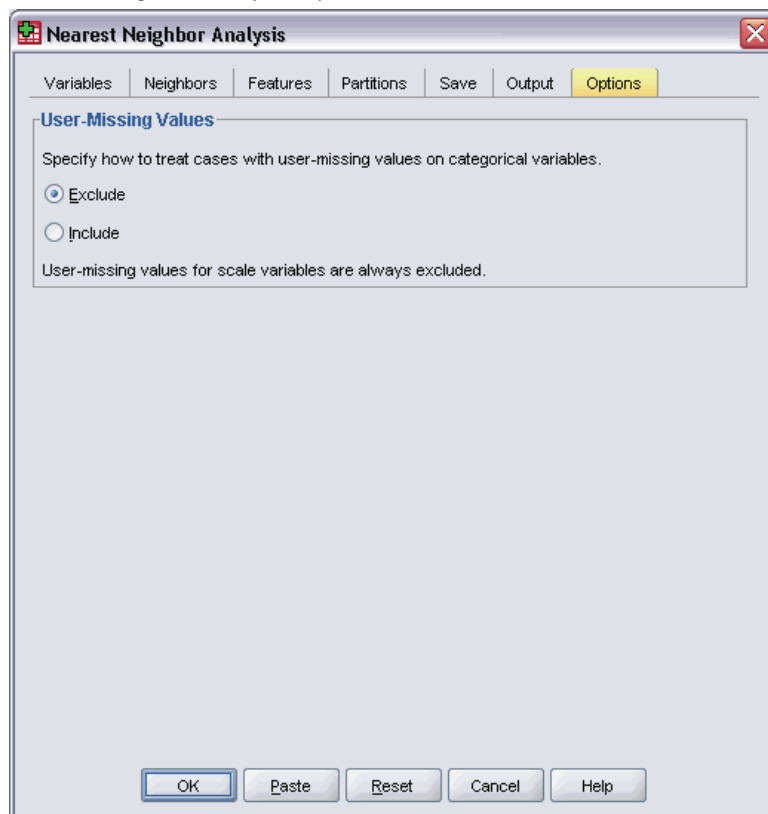
- **Case processing summary.** Displays the case processing summary table, which summarizes the number of cases included and excluded in the analysis, in total and by training and holdout samples.
- **Charts and tables.** Displays model-related output, including tables and charts. Tables in the model view include k nearest neighbors and distances for focal cases, classification of categorical response variables, and an error summary. Graphical output in the model view includes a selection error log, feature importance chart, feature space chart, peers chart, and quadrant map. [For more information, see Model View on p. 378.](#)

Files

- **Export model to XML.** SmartScore and SPSS Statistics Server (a separate product) can use this model file to apply the model information to other data files for scoring purposes. This option is not available if split files have been defined.
- **Export distances between focal cases and k nearest neighbors.** For each focal case, a separate variable is created for each of the focal case's k nearest neighbors (from the training sample) and the corresponding k nearest distances.

Options

Figure 32-8
Nearest Neighbor Analysis Options tab

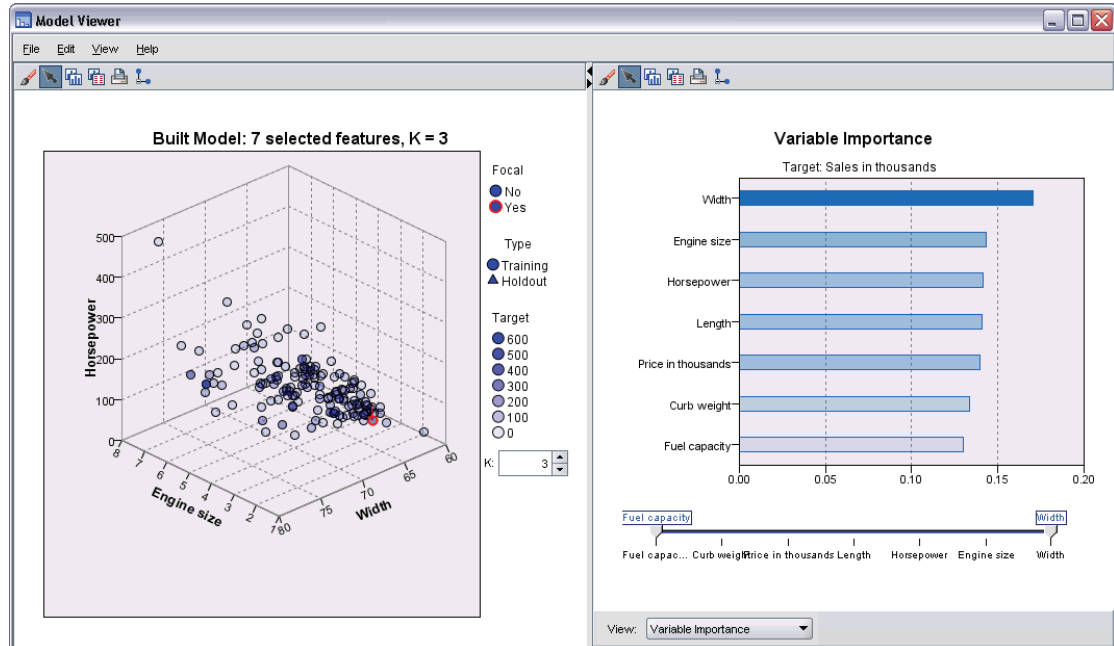


User-Missing Values. Categorical variables must have valid values for a case to be included in the analysis. These controls allow you to decide whether user-missing values are treated as valid among categorical variables.

System-missing values and missing values for scale variables are always treated as invalid.

Model View

Figure 32-9
Nearest Neighbor Analysis Model View



When you select Charts and tables in the Output tab, the procedure creates a Nearest Neighbor Model object in the Viewer. By activating (double-clicking) this object, you gain an interactive view of the model. The model view has a 2-panel window:

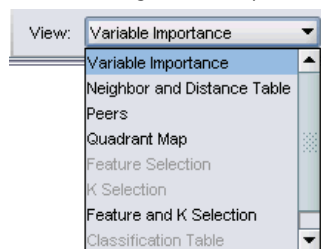
- The first panel displays an overview of the model called the main view.
- The second panel displays one of two types of views:

An auxiliary model view shows more information about the model, but is not focused on the model itself.

A linked view is a view that shows details about one feature of the model when the user drills down on part of the main view.

By default, the first panel shows the feature space and the second panel shows the variable importance chart. If the variable importance chart is not available; that is, when Weight features by importance was not selected on the Features tab, the first available view in the View dropdown is shown.

Figure 32-10
Nearest Neighbor Analysis Model View dropdown

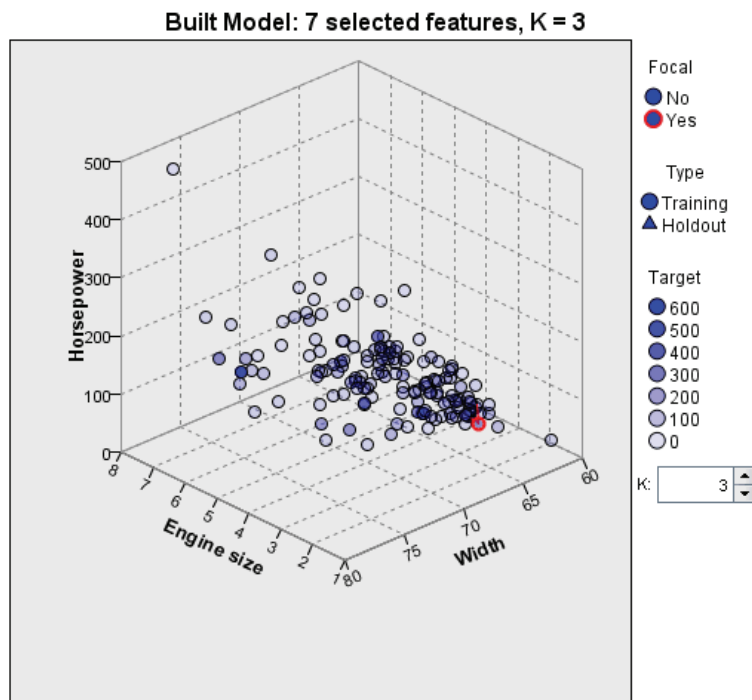


When a view has no available information, its item text in the View dropdown is disabled.

For general information on Model objects, see [Models](#) on p. 250.

Feature Space

Figure 32-11
Feature Space



The feature space chart is an interactive graph of the feature space (or a subspace, if there are more than 3 features). Each axis represents a feature in the model, and the location of points in the chart show the values of these features for cases in the training and holdout partitions.

Keys. In addition to the feature values, points in the plot convey other information.

- Shape indicates the partition to which a point belongs, either Training or Holdout.

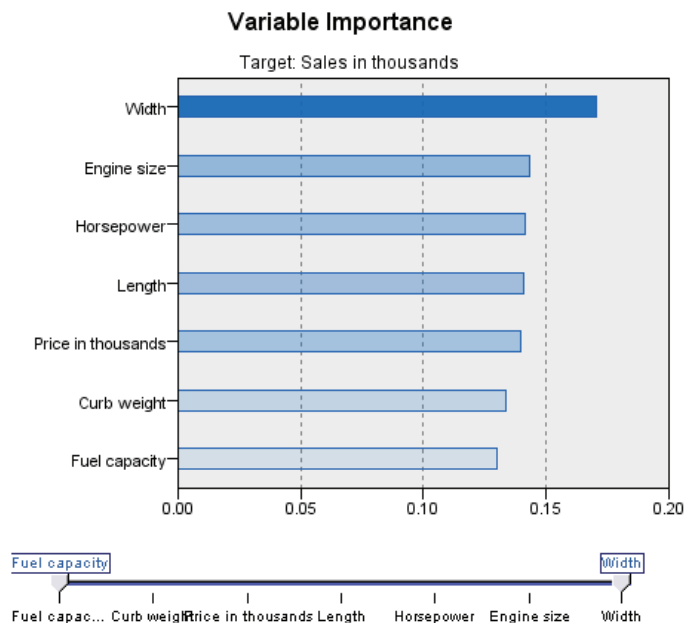
- The color/shading of a point indicates the value of the target for that case; with distinct color values equal to the categories of a categorical target, and shades indicating the range of values of a continuous target. The indicated value for the training partition is the observed value; for the holdout partition, it is the predicted value. If no target is specified, this key is not shown.
- Heavier outlines indicate a case is focal. Focal cases are shown linked to their k nearest neighbors.

Controls and Interactivity. A number of controls in the chart allow you explore the Feature Space.

- You can choose which subset of features to show in the chart and change which features are represented on the dimensions.
- “Focal cases” are simply points selected in the Feature Space chart. If you specified a focal case variable, the points representing the focal cases will initially be selected. However, any point can temporarily become a focal case if you select it. The “usual” controls for point selection apply; clicking on a point selects that point and deselects all others; Control-clicking on a point adds it to the set of selected points. Linked views, such as the Peers Chart, will automatically update based upon the cases selected in the Feature Space.
- You can change the number of nearest neighbors (k) to display for focal cases.
- Hovering over a point in the chart displays a tooltip with the value of the case label, or case number if case labels are not defined, and the observed and predicted target values.
- A “Reset” button allows you to return the Feature Space to its original state.

Variable Importance

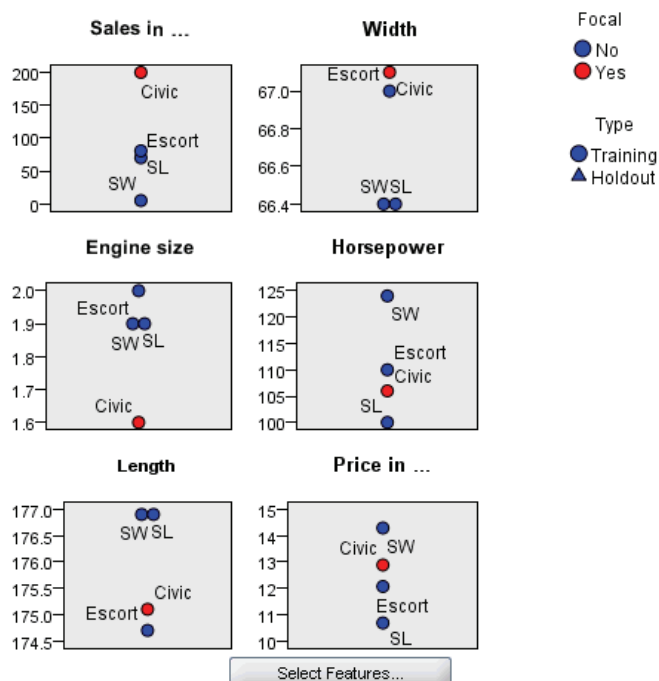
Figure 32-12
Variable Importance



Typically, you will want to focus your modeling efforts on the variables that matter most and consider dropping or ignoring those that matter least. The variable importance chart helps you do this by indicating the relative importance of each variable in estimating the model. Since the values are relative, the sum of the values for all variables on the display is 1.0. Variable importance does not relate to model accuracy. It just relates to the importance of each variable in making a prediction, not whether or not the prediction is accurate.

Peers

Figure 32-13
Peers Chart



This chart displays the focal cases and their k nearest neighbors on each feature and on the target. It is available if a focal case is selected in the Feature Space.

Linking Behavior. The Peers chart is linked to the Feature Space in two ways.

- Cases selected (focal) in the Feature Space are displayed in the Peers chart, along with their k nearest neighbors.
- The value of k selected in the Feature Space is used in the Peers chart.

Nearest Neighbor Distances

Figure 32-14
Nearest Neighbor Distances

Focal Case	Nearest Neighbors			Nearest Distances		
	1	2	3	1	2	3
Civic	SL	Escort	SW	0.053	0.059	0.064


This table displays the k nearest neighbors and distances for focal cases only. It is available if a focal case identifier is specified on the Variables tab, and only displays focal cases identified by this variable.

Each row of:

- The Focal Case column contains the value of the case labeling variable for the focal case; if case labels are not defined, this column contains the case number of the focal case.
- The i th column under the Nearest Neighbors group contains the value of the case labeling variable for the i th nearest neighbor of the focal case; if case labels are not defined, this column contains the case number of the i th nearest neighbor of the focal case.
- The i th column under the Nearest Distances group contains the distance of the i th nearest neighbor to the focal case

Quadrant Map

Figure 32-15
Quadrant Map

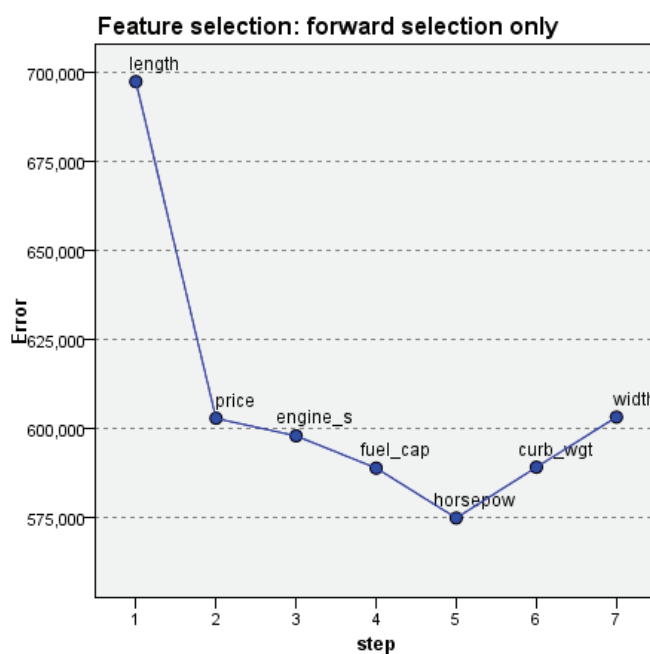


This chart displays the focal cases and their k nearest neighbors on a scatterplot (or dotplot, depending upon the measurement level of the target) with the target on the y -axis and a scale feature on the x -axis, paneled by features. It is available if there is a target and if a focal case is selected in the Feature Space.

- Reference lines are drawn for continuous variables, at the variable means in the training partition.

Feature Selection Error Log

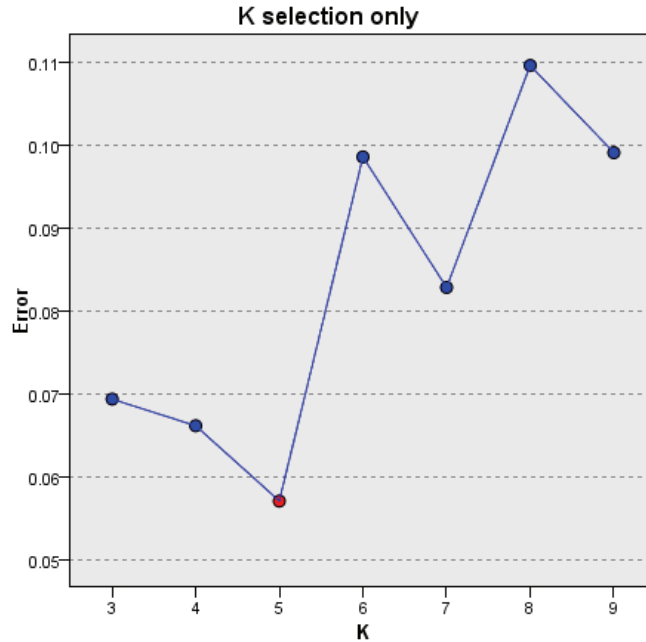
Figure 32-16
Feature Selection



Points on the chart display the error (either the error rate or sum-of-squares error, depending upon the measurement level of the target) on the y -axis for the model with the feature listed on the x -axis (plus all features to the left on the x -axis). This chart is available if there is a target and feature selection is in effect.

***k* Selection Error Log**

Figure 32-17
k Selection




Points on the chart display the error (either the error rate or sum-of-squares error, depending upon the measurement level of the target) on the y -axis for the model with the number of nearest neighbors (k) on the x -axis. This chart is available if there is a target and k selection is in effect.

***k** and Feature Selection Error Log*

Figure 32-18
k and Feature Selection

K selection and Feature selection (forward selection)



These are feature selection charts (see [Feature Selection Error Log on p. 383](#)), paneled by k . This chart is available if there is a target and k and feature selection are both in effect.

Classification Table

Figure 32-19
Classification Table

Partition		Predicted		
		0	1	Percent Correct
Training	0	111	1	99.11%
	1	7	33	82.50%
	Overall Percent	77.64%	22.37%	94.74%

This table displays the cross-classification of observed versus predicted values of the target, by partition. It is available if there is a target and it is categorical.

- The (Missing) row in the Holdout partition contains holdout cases with missing values on the target. These cases contribute to the Holdout Sample: Overall Percent values but not to the Percent Correct values.

Error Summary

Figure 32-20
Error Summary

Partition	Sum-of-Squares Error
Training	622043

This table is available if there is a target variable. It displays the error associated with the model; sum-of-squares for a continuous target and the error rate (100% – overall percent correct) for a categorical target.

Discriminant Analysis

Discriminant analysis builds a predictive model for group membership. The model is composed of a discriminant function (or, for more than two groups, a set of discriminant functions) based on linear combinations of the predictor variables that provide the best discrimination between the groups. The functions are generated from a sample of cases for which group membership is known; the functions can then be applied to new cases that have measurements for the predictor variables but have unknown group membership.

Note: The grouping variable can have more than two values. The codes for the grouping variable must be integers, however, and you need to specify their minimum and maximum values. Cases with values outside of these bounds are excluded from the analysis.

Example. On average, people in temperate zone countries consume more calories per day than people in the tropics, and a greater proportion of the people in the temperate zones are city dwellers. A researcher wants to combine this information into a function to determine how well an individual can discriminate between the two groups of countries. The researcher thinks that population size and economic information may also be important. Discriminant analysis allows you to estimate coefficients of the linear discriminant function, which looks like the right side of a multiple linear regression equation. That is, using coefficients a , b , c , and d , the function is:

$$D = a * \text{climate} + b * \text{urban} + c * \text{population} + d * \text{gross domestic product per capita}$$

If these variables are useful for discriminating between the two climate zones, the values of D will differ for the temperate and tropic countries. If you use a stepwise variable selection method, you may find that you do not need to include all four variables in the function.

Statistics. For each variable: means, standard deviations, univariate ANOVA. For each analysis: Box's M , within-groups correlation matrix, within-groups covariance matrix, separate-groups covariance matrix, total covariance matrix. For each canonical discriminant function: eigenvalue, percentage of variance, canonical correlation, Wilks' lambda, chi-square. For each step: prior probabilities, Fisher's function coefficients, unstandardized function coefficients, Wilks' lambda for each canonical function.

Data. The grouping variable must have a limited number of distinct categories, coded as integers. Independent variables that are nominal must be recoded to dummy or contrast variables.

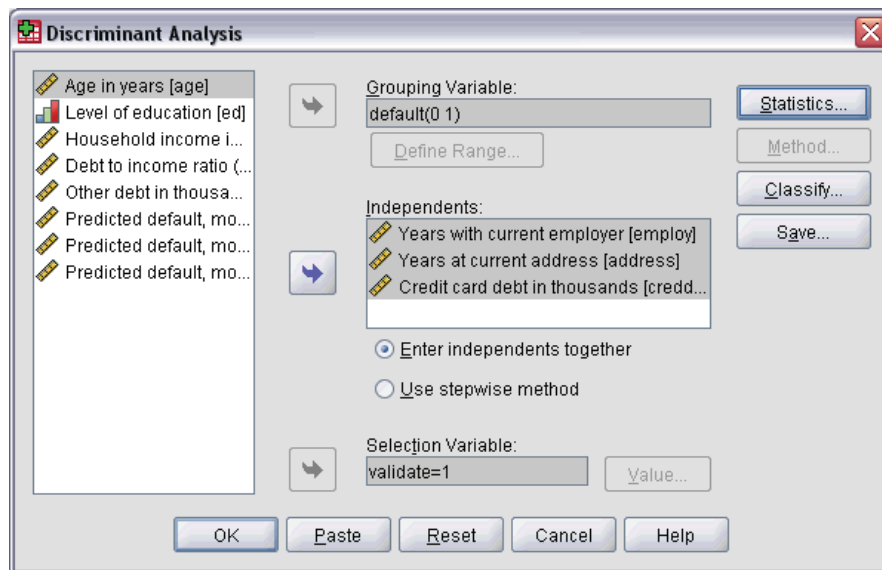
Assumptions. Cases should be independent. Predictor variables should have a multivariate normal distribution, and within-group variance-covariance matrices should be equal across groups. Group membership is assumed to be mutually exclusive (that is, no case belongs to more than one group) and collectively exhaustive (that is, all cases are members of a group). The procedure is most effective when group membership is a truly categorical variable; if group membership is

based on values of a continuous variable (for example, high IQ versus low IQ), consider using linear regression to take advantage of the richer information that is offered by the continuous variable itself.

To Obtain a Discriminant Analysis

- ▶ From the menus choose:
 - Analyze
 - Classify
 - Discriminant...


Figure 33-1
Discriminant Analysis dialog box



- ▶ Select an integer-valued grouping variable and click Define Range to specify the categories of interest.
- ▶ Select the independent, or predictor, variables. (If your grouping variable does not have integer values, Automatic Recode on the Transform menu will create a variable that does.)
- ▶ Select the method for entering the independent variables.
 - **Enter independents together.** Simultaneously enters all independent variables that satisfy tolerance criteria.
 - **Use stepwise method.** Uses stepwise analysis to control variable entry and removal.
- ▶ Optionally, select cases with a selection variable.

Discriminant Analysis Define Range


Figure 33-2
Discriminant Analysis Define Range dialog box



Specify the minimum and maximum value of the grouping variable for the analysis. Cases with values outside of this range are not used in the discriminant analysis but are classified into one of the existing groups based on the results of the analysis. The minimum and maximum values must be integers.

Discriminant Analysis Select Cases

Figure 33-3
Discriminant Analysis Set Value dialog box



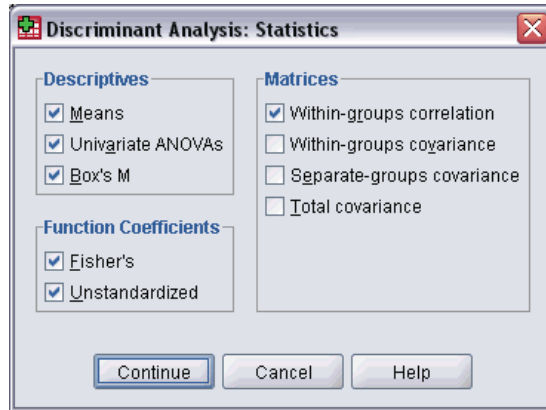
To select cases for your analysis:

- ▶ In the Discriminant Analysis dialog box, choose a selection variable.
- ▶ Click Value to enter an integer as the selection value.

Only cases with the specified value for the selection variable are used to derive the discriminant functions. Statistics and classification results are generated for both selected and unselected cases. This process provides a mechanism for classifying new cases based on previously existing data or for partitioning your data into training and testing subsets to perform validation on the model generated.

Discriminant Analysis Statistics

Figure 33-4
Discriminant Analysis Statistics dialog box



Descriptives. Available options are means (including standard deviations), univariate ANOVAs, and Box's *M* test.

- **Means.** Displays total and group means, as well as standard deviations for the independent variables.
- **Univariate ANOVAs.** Performs a one-way analysis-of-variance test for equality of group means for each independent variable.
- **Box's *M*.** A test for the equality of the group covariance matrices. For sufficiently large samples, a nonsignificant *p* value means there is insufficient evidence that the matrices differ. The test is sensitive to departures from multivariate normality.

Function Coefficients. Available options are Fisher's classification coefficients and unstandardized coefficients.

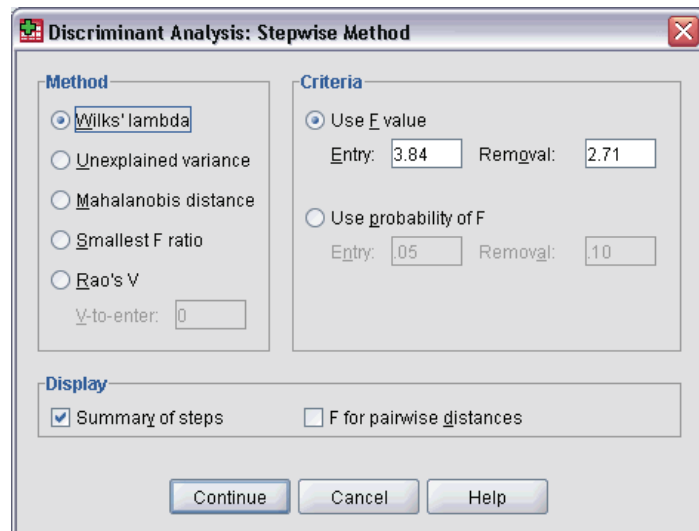
- **Fisher's.** Displays Fisher's classification function coefficients that can be used directly for classification. A separate set of classification function coefficients is obtained for each group, and a case is assigned to the group for which it has the largest discriminant score (classification function value).
- **Unstandardized.** Displays the unstandardized discriminant function coefficients.

Matrices. Available matrices of coefficients for independent variables are within-groups correlation matrix, within-groups covariance matrix, separate-groups covariance matrix, and total covariance matrix.

- **Within-groups correlation.** Displays a pooled within-groups correlation matrix that is obtained by averaging the separate covariance matrices for all groups before computing the correlations.
- **Within-groups covariance.** Displays a pooled within-groups covariance matrix, which may differ from the total covariance matrix. The matrix is obtained by averaging the separate covariance matrices for all groups.
- **Separate-groups covariance.** Displays separate covariance matrices for each group.
- **Total covariance.** Displays a covariance matrix from all cases as if they were from a single sample.

Discriminant Analysis Stepwise Method

Figure 33-5
Discriminant Analysis Stepwise Method dialog box



Method. Select the statistic to be used for entering or removing new variables. Available alternatives are Wilks' lambda, unexplained variance, Mahalanobis distance, smallest F ratio, and Rao's V . With Rao's V , you can specify the minimum increase in V for a variable to enter.

- **Wilks' lambda.** A variable selection method for stepwise discriminant analysis that chooses variables for entry into the equation on the basis of how much they lower Wilks' lambda. At each step, the variable that minimizes the overall Wilks' lambda is entered.
- **Unexplained variance.** At each step, the variable that minimizes the sum of the unexplained variation between groups is entered.
- **Mahalanobis distance.** A measure of how much a case's values on the independent variables differ from the average of all cases. A large Mahalanobis distance identifies a case as having extreme values on one or more of the independent variables.
- **Smallest F ratio.** A method of variable selection in stepwise analysis based on maximizing an F ratio computed from the Mahalanobis distance between groups.
- **Rao's V .** A measure of the differences between group means. Also called the Lawley-Hotelling trace. At each step, the variable that maximizes the increase in Rao's V is entered. After selecting this option, enter the minimum value a variable must have to enter the analysis.

Criteria. Available alternatives are Use F value and Use probability of F . Enter values for entering and removing variables.

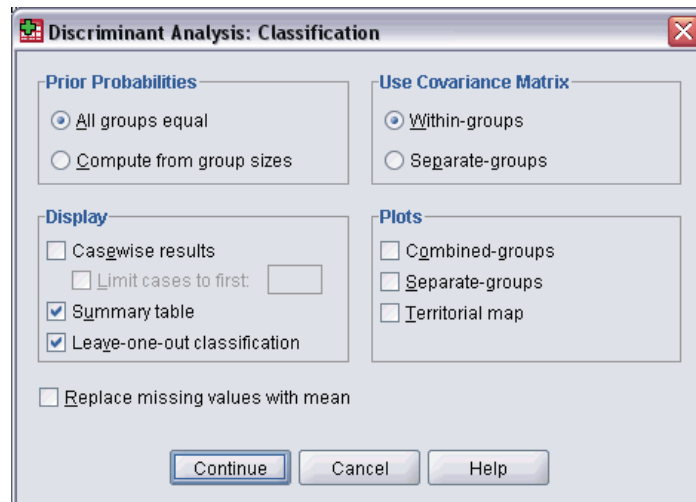
- **Use F value.** A variable is entered into the model if its F value is greater than the Entry value and is removed if the F value is less than the Removal value. Entry must be greater than Removal, and both values must be positive. To enter more variables into the model, lower the Entry value. To remove more variables from the model, increase the Removal value.
- **Use probability of F .** A variable is entered into the model if the significance level of its F value is less than the Entry value and is removed if the significance level is greater than the Removal value. Entry must be less than Removal, and both values must be positive. To enter

more variables into the model, increase the Entry value. To remove more variables from the model, lower the Removal value.

Display. Summary of steps displays statistics for all variables after each step; F for pairwise distances displays a matrix of pairwise F ratios for each pair of groups.

Discriminant Analysis Classification

Figure 33-6
Discriminant Analysis Classify dialog box



Prior Probabilities. This option determines whether the classification coefficients are adjusted for a priori knowledge of group membership.

- **All groups equal.** Equal prior probabilities are assumed for all groups; this has no effect on the coefficients.
- **Compute from group sizes.** The observed group sizes in your sample determine the prior probabilities of group membership. For example, if 50% of the observations included in the analysis fall into the first group, 25% in the second, and 25% in the third, the classification coefficients are adjusted to increase the likelihood of membership in the first group relative to the other two.

Display. Available display options are casewise results, summary table, and leave-one-out classification.

- **Casewise results.** Codes for actual group, predicted group, posterior probabilities, and discriminant scores are displayed for each case.
- **Summary table.** The number of cases correctly and incorrectly assigned to each of the groups based on the discriminant analysis. Sometimes called the "Confusion Matrix."
- **Leave-one-out classification.** Each case in the analysis is classified by the functions derived from all cases other than that case. It is also known as the "U-method."

Replace missing values with mean. Select this option to substitute the mean of an independent variable for a missing value during the classification phase only.

Use Covariance Matrix. You can choose to classify cases using a within-groups covariance matrix or a separate-groups covariance matrix.

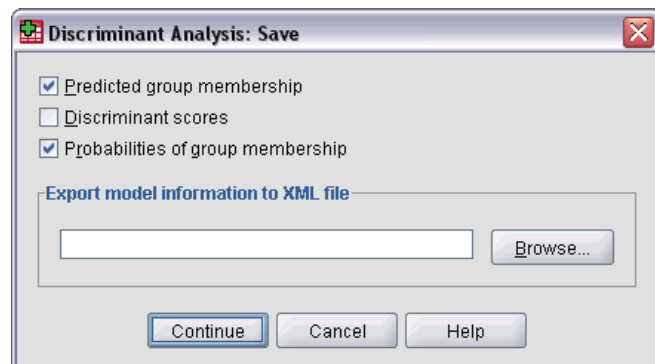
- **Within-groups.** The pooled within-groups covariance matrix is used to classify cases.
- **Separate-groups.** Separate-groups covariance matrices are used for classification. Because classification is based on the discriminant functions (not based on the original variables), this option is not always equivalent to quadratic discrimination.

Plots. Available plot options are combined-groups, separate-groups, and territorial map.

- **Combined-groups.** Creates an all-groups scatterplot of the first two discriminant function values. If there is only one function, a histogram is displayed instead.
- **Separate-groups.** Creates separate-group scatterplots of the first two discriminant function values. If there is only one function, histograms are displayed instead.
- **Territorial map.** A plot of the boundaries used to classify cases into groups based on function values. The numbers correspond to groups into which cases are classified. The mean for each group is indicated by an asterisk within its boundaries. The map is not displayed if there is only one discriminant function.

Discriminant Analysis Save

Figure 33-7
Discriminant Analysis Save dialog box



You can add new variables to your active data file. Available options are predicted group membership (a single variable), discriminant scores (one variable for each discriminant function in the solution), and probabilities of group membership given the discriminant scores (one variable for each group).

You can also export model information to the specified file in XML (PMML) format. SmartScore and SPSS Statistics Server (a separate product) can use this model file to apply the model information to other data files for scoring purposes.

DISCRIMINANT Command Additional Features

The command syntax language also allows you to:

- Perform multiple discriminant analyses (with one command) and control the order in which variables are entered (with the `ANALYSIS` subcommand).

- Specify prior probabilities for classification (with the `PRIORS` subcommand).
- Display rotated pattern and structure matrices (with the `ROTATE` subcommand).
- Limit the number of extracted discriminant functions (with the `FUNCTIONS` subcommand).
- Restrict classification to the cases that are selected (or unselected) for the analysis (with the `SELECT` subcommand).
- Read and analyze a correlation matrix (with the `MATRIX` subcommand).
- Write a correlation matrix for later analysis (with the `MATRIX` subcommand).

See the *Command Syntax Reference* for complete syntax information.

Factor Analysis

Factor analysis attempts to identify underlying variables, or **factors**, that explain the pattern of correlations within a set of observed variables. Factor analysis is often used in data reduction to identify a small number of factors that explain most of the variance that is observed in a much larger number of manifest variables. Factor analysis can also be used to generate hypotheses regarding causal mechanisms or to screen variables for subsequent analysis (for example, to identify collinearity prior to performing a linear regression analysis).

The factor analysis procedure offers a high degree of flexibility:

- Seven methods of factor extraction are available.
- Five methods of rotation are available, including direct oblimin and promax for nonorthogonal rotations.
- Three methods of computing factor scores are available, and scores can be saved as variables for further analysis.

Example. What underlying attitudes lead people to respond to the questions on a political survey as they do? Examining the correlations among the survey items reveals that there is significant overlap among various subgroups of items—questions about taxes tend to correlate with each other, questions about military issues correlate with each other, and so on. With factor analysis, you can investigate the number of underlying factors and, in many cases, identify what the factors represent conceptually. Additionally, you can compute factor scores for each respondent, which can then be used in subsequent analyses. For example, you might build a logistic regression model to predict voting behavior based on factor scores.

Statistics. For each variable: number of valid cases, mean, and standard deviation. For each factor analysis: correlation matrix of variables, including significance levels, determinant, and inverse; reproduced correlation matrix, including anti-image; initial solution (communalities, eigenvalues, and percentage of variance explained); Kaiser-Meyer-Olkin measure of sampling adequacy and Bartlett's test of sphericity; unrotated solution, including factor loadings, communalities, and eigenvalues; and rotated solution, including rotated pattern matrix and transformation matrix. For oblique rotations: rotated pattern and structure matrices; factor score coefficient matrix and factor covariance matrix. Plots: scree plot of eigenvalues and loading plot of first two or three factors.

Data. The variables should be quantitative at the **interval** or **ratio** level. Categorical data (such as religion or country of origin) are not suitable for factor analysis. Data for which Pearson correlation coefficients can sensibly be calculated should be suitable for factor analysis.

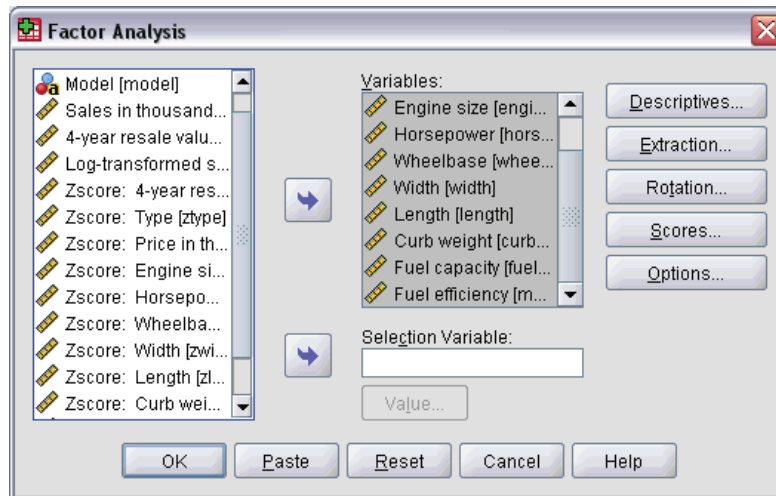
Assumptions. The data should have a bivariate normal distribution for each pair of variables, and observations should be independent. The factor analysis model specifies that variables are determined by common factors (the factors estimated by the model) and unique factors (which do

not overlap between observed variables); the computed estimates are based on the assumption that all unique factors are uncorrelated with each other and with the common factors.

To Obtain a Factor Analysis


- ▶ From the menus choose:
 - Analyze
 - Dimension Reduction
 - Factor...
- ▶ Select the variables for the factor analysis.

Figure 34-1
Factor Analysis dialog box



Factor Analysis Select Cases

Figure 34-2
Factor Analysis Set Value dialog box



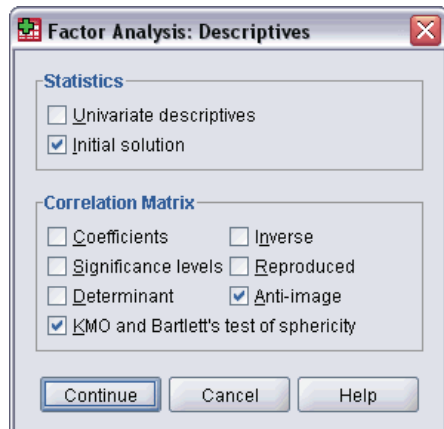
To select cases for your analysis:

- ▶ Choose a selection variable.
- ▶ Click Value to enter an integer as the selection value.

Only cases with that value for the selection variable are used in the factor analysis.

Factor Analysis Descriptives

Figure 34-3
Factor Analysis Descriptives dialog box



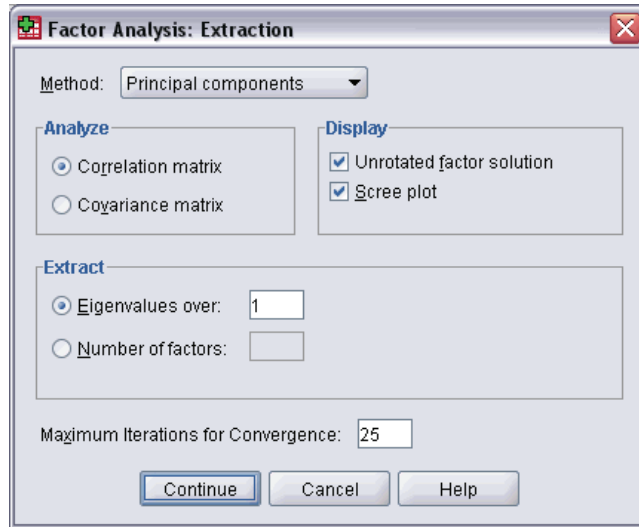
Statistics. Univariate descriptives includes the mean, standard deviation, and number of valid cases for each variable. Initial solution displays initial communalities, eigenvalues, and the percentage of variance explained.

Correlation Matrix. The available options are coefficients, significance levels, determinant, KMO and Bartlett's test of sphericity, inverse, reproduced, and anti-image.

- **KMO and Bartlett's Test of Sphericity.** The Kaiser-Meyer-Olkin measure of sampling adequacy tests whether the partial correlations among variables are small. Bartlett's test of sphericity tests whether the correlation matrix is an identity matrix, which would indicate that the factor model is inappropriate.
- **Reproduced.** The estimated correlation matrix from the factor solution. Residuals (difference between estimated and observed correlations) are also displayed.
- **Anti-image.** The anti-image correlation matrix contains the negatives of the partial correlation coefficients, and the anti-image covariance matrix contains the negatives of the partial covariances. In a good factor model, most of the off-diagonal elements will be small. The measure of sampling adequacy for a variable is displayed on the diagonal of the anti-image correlation matrix.

Factor Analysis Extraction

Figure 34-4
Factor Analysis Extraction dialog box



Method. Allows you to specify the method of factor extraction. Available methods are principal components, unweighted least squares, generalized least squares, maximum likelihood, principal axis factoring, alpha factoring, and image factoring.

- **Principal Components Analysis.** A factor extraction method used to form uncorrelated linear combinations of the observed variables. The first component has maximum variance. Successive components explain progressively smaller portions of the variance and are all uncorrelated with each other. Principal components analysis is used to obtain the initial factor solution. It can be used when a correlation matrix is singular.
- **Unweighted Least-Squares Method.** A factor extraction method that minimizes the sum of the squared differences between the observed and reproduced correlation matrices (ignoring the diagonals).
- **Generalized Least-Squares Method.** A factor extraction method that minimizes the sum of the squared differences between the observed and reproduced correlation matrices. Correlations are weighted by the inverse of their uniqueness, so that variables with high uniqueness are given less weight than those with low uniqueness.
- **Maximum-Likelihood Method.** A factor extraction method that produces parameter estimates that are most likely to have produced the observed correlation matrix if the sample is from a multivariate normal distribution. The correlations are weighted by the inverse of the uniqueness of the variables, and an iterative algorithm is employed.
- **Principal Axis Factoring.** A method of extracting factors from the original correlation matrix, with squared multiple correlation coefficients placed in the diagonal as initial estimates of the communalities. These factor loadings are used to estimate new communalities that replace the old communality estimates in the diagonal. Iterations continue until the changes in the communalities from one iteration to the next satisfy the convergence criterion for extraction.

- **Alpha.** A factor extraction method that considers the variables in the analysis to be a sample from the universe of potential variables. This method maximizes the alpha reliability of the factors.
- **Image Factoring.** A factor extraction method developed by Guttman and based on image theory. The common part of the variable, called the partial image, is defined as its linear regression on remaining variables, rather than a function of hypothetical factors.

Analyze. Allows you to specify either a correlation matrix or a covariance matrix.

- **Correlation matrix.** Useful if variables in your analysis are measured on different scales.
- **Covariance matrix.** Useful when you want to apply your factor analysis to multiple groups with different variances for each variable.

Extract. You can either retain all factors whose eigenvalues exceed a specified value, or you can retain a specific number of factors.

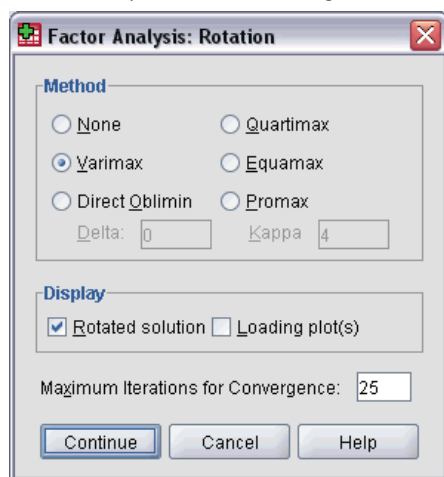
Display. Allows you to request the unrotated factor solution and a scree plot of the eigenvalues.

- **Unrotated Factor Solution.** Displays unrotated factor loadings (factor pattern matrix), communalities, and eigenvalues for the factor solution.
- **Scree plot.** A plot of the variance that is associated with each factor. This plot is used to determine how many factors should be kept. Typically the plot shows a distinct break between the steep slope of the large factors and the gradual trailing of the rest (the scree).

Maximum Iterations for Convergence. Allows you to specify the maximum number of steps that the algorithm can take to estimate the solution.

Factor Analysis Rotation

Figure 34-5
Factor Analysis Rotation dialog box



Method. Allows you to select the method of factor rotation. Available methods are varimax, direct oblimin, quartimax, equamax, or promax.

- **Varimax Method.** An orthogonal rotation method that minimizes the number of variables that have high loadings on each factor. This method simplifies the interpretation of the factors.

- **Direct Oblimin Method.** A method for oblique (nonorthogonal) rotation. When delta equals 0 (the default), solutions are most oblique. As delta becomes more negative, the factors become less oblique. To override the default delta of 0, enter a number less than or equal to 0.8.
- **Quartimax Method.** A rotation method that minimizes the number of factors needed to explain each variable. This method simplifies the interpretation of the observed variables.
- **Equamax Method.** A rotation method that is a combination of the varimax method, which simplifies the factors, and the quartimax method, which simplifies the variables. The number of variables that load highly on a factor and the number of factors needed to explain a variable are minimized.
- **Promax Rotation.** An oblique rotation, which allows factors to be correlated. This rotation can be calculated more quickly than a direct oblimin rotation, so it is useful for large datasets.

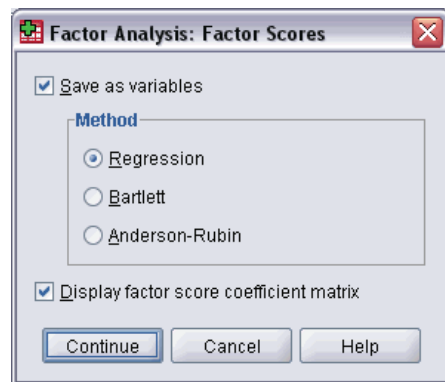
Display. Allows you to include output on the rotated solution, as well as loading plots for the first two or three factors.

- **Rotated Solution.** A rotation method must be selected to obtain a rotated solution. For orthogonal rotations, the rotated pattern matrix and factor transformation matrix are displayed. For oblique rotations, the pattern, structure, and factor correlation matrices are displayed.
- **Factor Loading Plot.** Three-dimensional factor loading plot of the first three factors. For a two-factor solution, a two-dimensional plot is shown. The plot is not displayed if only one factor is extracted. Plots display rotated solutions if rotation is requested.

Maximum Iterations for Convergence. Allows you to specify the maximum number of steps that the algorithm can take to perform the rotation.

Factor Analysis Scores

Figure 34-6
Factor Analysis Factor Scores dialog box



Save as variables. Creates one new variable for each factor in the final solution.

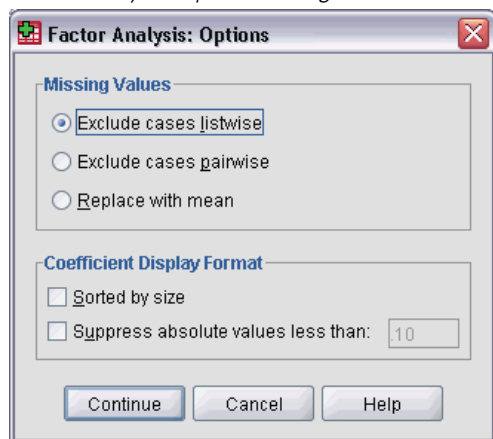
Method. The alternative methods for calculating factor scores are regression, Bartlett, and Anderson-Rubin.

- **Regression Method.** A method for estimating factor score coefficients. The scores that are produced have a mean of 0 and a variance equal to the squared multiple correlation between the estimated factor scores and the true factor values. The scores may be correlated even when factors are orthogonal.
- **Bartlett Scores.** A method of estimating factor score coefficients. The scores that are produced have a mean of 0. The sum of squares of the unique factors over the range of variables is minimized.
- **Anderson-Rubin Method.** A method of estimating factor score coefficients; a modification of the Bartlett method which ensures orthogonality of the estimated factors. The scores that are produced have a mean of 0, have a standard deviation of 1, and are uncorrelated.

Display factor score coefficient matrix. Shows the coefficients by which variables are multiplied to obtain factor scores. Also shows the correlations between factor scores.

Factor Analysis Options

Figure 34-7
Factor Analysis Options dialog box



Missing Values. Allows you to specify how missing values are handled. The available choices are to exclude cases **listwise**, exclude cases **pairwise**, or replace with mean.

Coefficient Display Format. Allows you to control aspects of the output matrices. You sort coefficients by size and suppress coefficients with absolute values that are less than the specified value.

FACTOR Command Additional Features

The command syntax language also allows you to:

- Specify convergence criteria for iteration during extraction and rotation.
- Specify individual rotated-factor plots.
- Specify how many factor scores to save.
- Specify diagonal values for the principal axis factoring method.

- Write correlation matrices or factor-loading matrices to disk for later analysis.
- Read and analyze correlation matrices or factor-loading matrices.

See the *Command Syntax Reference* for complete syntax information.

Choosing a Procedure for Clustering

Cluster analyses can be performed using the TwoStep, Hierarchical, or K-Means Cluster Analysis procedure. Each procedure employs a different algorithm for creating clusters, and each has options not available in the others.

TwoStep Cluster Analysis. For many applications, the TwoStep Cluster Analysis procedure will be the method of choice. It provides the following unique features:

- Automatic selection of the best number of clusters, in addition to measures for choosing between cluster models.
- Ability to create cluster models simultaneously based on categorical and continuous variables.
- Ability to save the cluster model to an external XML file and then read that file and update the cluster model using newer data.

Additionally, the TwoStep Cluster Analysis procedure can analyze large data files.

Hierarchical Cluster Analysis. The Hierarchical Cluster Analysis procedure is limited to smaller data files (hundreds of objects to be clustered) but has the following unique features:

- Ability to cluster cases or variables.
- Ability to compute a range of possible solutions and save cluster memberships for each of those solutions.
- Several methods for cluster formation, variable transformation, and measuring the dissimilarity between clusters.

As long as all the variables are of the same type, the Hierarchical Cluster Analysis procedure can analyze interval (continuous), count, or binary variables.

K-Means Cluster Analysis. The K-Means Cluster Analysis procedure is limited to continuous data and requires you to specify the number of clusters in advance, but it has the following unique features:

- Ability to save distances from cluster centers for each object.
- Ability to read initial cluster centers from and save final cluster centers to an external SPSS Statistics file.

Additionally, the K-Means Cluster Analysis procedure can analyze large data files.

TwoStep Cluster Analysis

The TwoStep Cluster Analysis procedure is an exploratory tool designed to reveal natural groupings (or clusters) within a dataset that would otherwise not be apparent. The algorithm employed by this procedure has several desirable features that differentiate it from traditional clustering techniques:

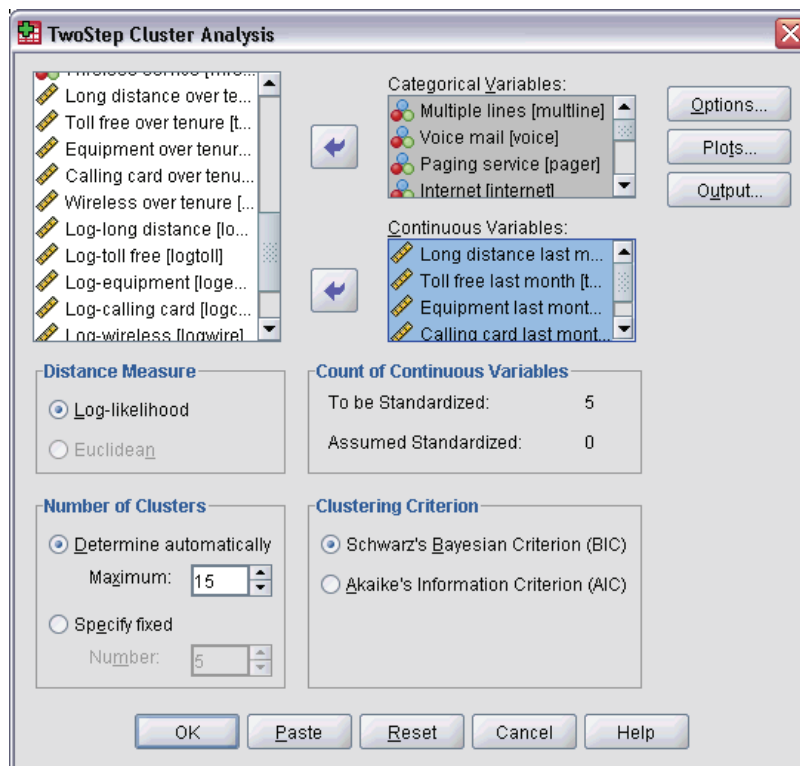
- **Handling of categorical and continuous variables.** By assuming variables to be independent, a joint multinomial-normal distribution can be placed on categorical and continuous variables.
- **Automatic selection of number of clusters.** By comparing the values of a model-choice criterion across different clustering solutions, the procedure can automatically determine the optimal number of clusters.
- **Scalability.** By constructing a cluster features (CF) tree that summarizes the records, the TwoStep algorithm allows you to analyze large data files.

Example. Retail and consumer product companies regularly apply clustering techniques to data that describe their customers' buying habits, gender, age, income level, etc. These companies tailor their marketing and product development strategies to each consumer group to increase sales and build brand loyalty.

Statistics. The procedure produces information criteria (AIC or BIC) by number of clusters in the solution, cluster frequencies for the final clustering, and descriptive statistics by cluster for the final clustering.

Plots. The procedure produces bar charts of cluster frequencies, pie charts of cluster frequencies, and variable importance charts.

Figure 36-1
TwoStep Cluster Analysis dialog box



Distance Measure. This selection determines how the similarity between two clusters is computed.

- **Log-likelihood.** The likelihood measure places a probability distribution on the variables. Continuous variables are assumed to be normally distributed, while categorical variables are assumed to be multinomial. All variables are assumed to be independent.
- **Euclidean.** The Euclidean measure is the “straight line” distance between two clusters. It can be used only when all of the variables are continuous.

Number of Clusters. This selection allows you to specify how the number of clusters is to be determined.

- **Determine automatically.** The procedure will automatically determine the “best” number of clusters, using the criterion specified in the Clustering Criterion group. Optionally, enter a positive integer specifying the maximum number of clusters that the procedure should consider.
- **Specify fixed.** Allows you to fix the number of clusters in the solution. Enter a positive integer.

Count of Continuous Variables. This group provides a summary of the continuous variable standardization specifications made in the Options dialog box. [For more information, see TwoStep Cluster Analysis Options on p. 407.](#)

Clustering Criterion. This selection determines how the automatic clustering algorithm determines the number of clusters. Either the Bayesian Information Criterion (BIC) or the Akaike Information Criterion (AIC) can be specified.

Data. This procedure works with both continuous and categorical variables. Cases represent objects to be clustered, and the variables represent attributes upon which the clustering is based.

Case Order. Note that the cluster features tree and the final solution may depend on the order of cases. To minimize order effects, randomly order the cases. You may want to obtain several different solutions with cases sorted in different random orders to verify the stability of a given solution. In situations where this is difficult due to extremely large file sizes, multiple runs with a sample of cases sorted in different random orders might be substituted.

Assumptions. The likelihood distance measure assumes that variables in the cluster model are independent. Further, each continuous variable is assumed to have a normal (Gaussian) distribution, and each categorical variable is assumed to have a multinomial distribution. Empirical internal testing indicates that the procedure is fairly robust to violations of both the assumption of independence and the distributional assumptions, but you should try to be aware of how well these assumptions are met.

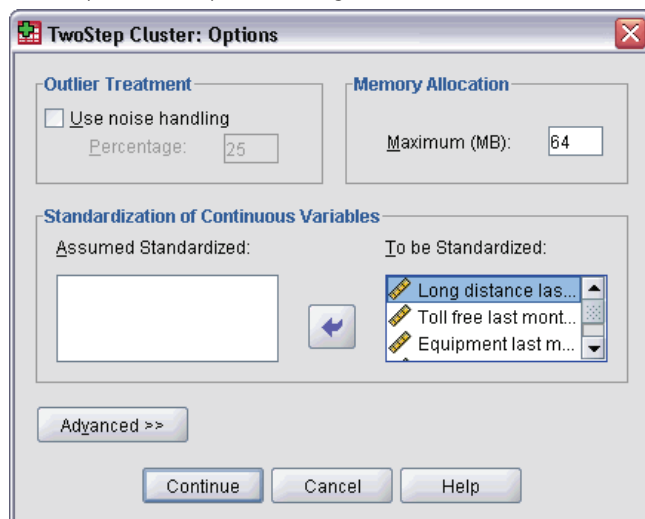
Use the [Bivariate Correlations](#) procedure to test the independence of two continuous variables. Use the [Crosstabs](#) procedure to test the independence of two categorical variables. Use the [Means](#) procedure to test the independence between a continuous variable and categorical variable. Use the [Explore](#) procedure to test the normality of a continuous variable. Use the [Chi-Square Test](#) procedure to test whether a categorical variable has a specified multinomial distribution.

To Obtain a TwoStep Cluster Analysis

- ▶ From the menus choose:
 - Analyze
 - Classify
 - TwoStep Cluster...
- ▶ Select one or more categorical or continuous variables.
 - Optionally, you can:
 - Adjust the criteria by which clusters are constructed.
 - Select settings for noise handling, memory allocation, variable standardization, and cluster model input.
 - Request optional tables and plots.
 - Save model results to the working file or to an external XML file.

TwoStep Cluster Analysis Options

Figure 36-2
TwoStep Cluster Options dialog box



Outlier Treatment. This group allows you to treat outliers specially during clustering if the cluster features (CF) tree fills. The CF tree is full if it cannot accept any more cases in a leaf node and no leaf node can be split.

- If you select noise handling and the CF tree fills, it will be regrown after placing cases in sparse leaves into a “noise” leaf. A leaf is considered sparse if it contains fewer than the specified percentage of cases of the maximum leaf size. After the tree is regrown, the outliers will be placed in the CF tree if possible. If not, the outliers are discarded.
- If you do not select noise handling and the CF tree fills, it will be regrown using a larger distance change threshold. After final clustering, values that cannot be assigned to a cluster are labeled outliers. The outlier cluster is given an identification number of -1 and is not included in the count of the number of clusters.

Memory Allocation. This group allows you to specify the maximum amount of memory in megabytes (MB) that the cluster algorithm should use. If the procedure exceeds this maximum, it will use the disk to store information that will not fit in memory. Specify a number greater than or equal to 4.

- Consult your system administrator for the largest value that you can specify on your system.
- The algorithm may fail to find the correct or desired number of clusters if this value is too low.

Variable standardization. The clustering algorithm works with standardized continuous variables. Any continuous variables that are not standardized should be left as variables in the To be Standardized list. To save some time and computational effort, you can select any continuous variables that you have already standardized as variables in the Assumed Standardized list.

Advanced Options

CF Tree Tuning Criteria. The following clustering algorithm settings apply specifically to the cluster features (CF) tree and should be changed with care:

- **Initial Distance Change Threshold.** This is the initial threshold used to grow the CF tree. If inserting a given case into a leaf of the CF tree would yield tightness less than the threshold, the leaf is not split. If the tightness exceeds the threshold, the leaf is split.
- **Maximum Branches (per leaf node).** The maximum number of child nodes that a leaf node can have.
- **Maximum Tree Depth.** The maximum number of levels that the CF tree can have.
- **Maximum Number of Nodes Possible.** This indicates the maximum number of CF tree nodes that could potentially be generated by the procedure, based on the function $(b^{d+1} - 1) / (b - 1)$, where b is the maximum branches and d is the maximum tree depth. Be aware that an overly large CF tree can be a drain on system resources and can adversely affect the performance of the procedure. At a minimum, each node requires 16 bytes.

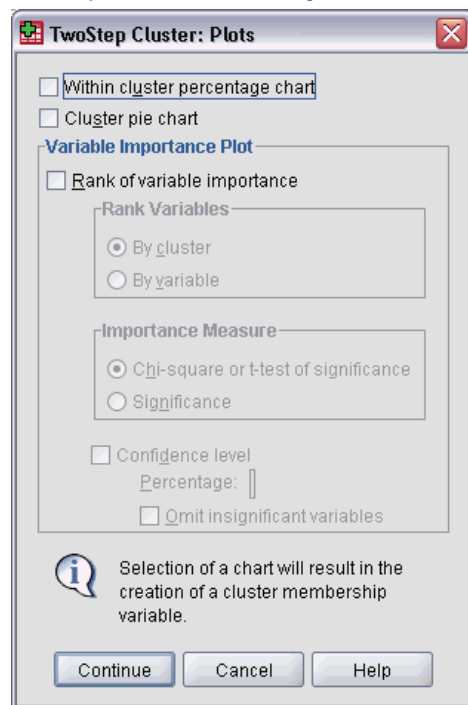
Cluster Model Update. This group allows you to import and update a cluster model generated in a prior analysis. The input file contains the CF tree in XML format. The model will then be updated with the data in the active file. You must select the variable names in the main dialog box in the same order in which they were specified in the prior analysis. The XML file remains unaltered, unless you specifically write the new model information to the same filename. [For more information, see TwoStep Cluster Analysis Output on p. 410.](#)

If a cluster model update is specified, the options pertaining to generation of the CF tree that were specified for the original model are used. More specifically, the distance measure, noise handling, memory allocation, or CF tree tuning criteria settings for the saved model are used, and any settings for these options in the dialog boxes are ignored.

Note: When performing a cluster model update, the procedure assumes that none of the selected cases in the active dataset were used to create the original cluster model. The procedure also assumes that the cases used in the model update come from the same population as the cases used to create the original model; that is, the means and variances of continuous variables and levels of categorical variables are assumed to be the same across both sets of cases. If your “new” and “old” sets of cases come from heterogeneous populations, you should run the TwoStep Cluster Analysis procedure on the combined sets of cases for the best results.

TwoStep Cluster Analysis Plots

Figure 36-3
TwoStep Cluster Plots dialog box



Within cluster percentage chart. Displays charts showing the within-cluster variation of each variable. For each categorical variable, a clustered bar chart is produced, showing the category frequency by cluster ID. For each continuous variable, an error bar chart is produced, showing error bars by cluster ID.

Cluster pie chart. Displays a pie chart showing the percentage and counts of observations within each cluster.

Variable Importance Plot. Displays several different charts showing the importance of each variable within each cluster. The output is sorted by the importance rank of each variable.

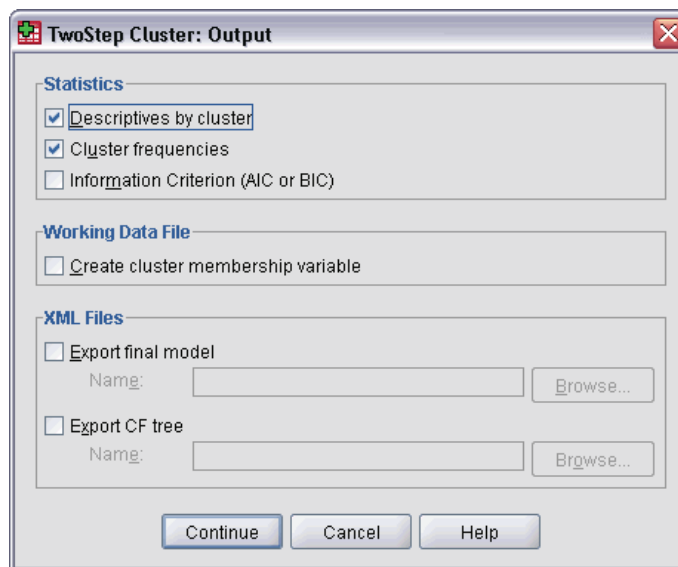
- **Rank Variables.** This option determines whether plots will be created for each cluster (By variable) or for each variable (By cluster).
- **Importance Measure.** This option allows you to select which measure of variable importance to plot. Chi-square or t-test of significance reports a Pearson chi-square statistic as the importance of a categorical variable and a t statistic as the importance of a continuous variable. Significance reports one minus the p value for the test of equality of means for a continuous variable and the expected frequency with the overall dataset for a categorical variable.
- **Confidence level.** This option allows you to set the confidence level for the test of equality of a variable's distribution within a cluster versus the variable's overall distribution. Specify a number less than 100 and greater than or equal to 50. The value of the confidence level is

shown as a vertical line in the variable importance plots, if the plots are created by variable or if the significance measure is plotted.

- **Omit insignificant variables.** Variables that are not significant at the specified confidence level are not displayed in the variable importance plots.

TwoStep Cluster Analysis Output

Figure 36-4
TwoStep Cluster Output dialog box



Statistics. This group provides options for displaying tables of the clustering results. The descriptive statistics and cluster frequencies are produced for the final cluster model, while the information criterion table displays results for a range of cluster solutions.

- **Descriptives by cluster.** Displays two tables that describe the variables in each cluster. In one table, means and standard deviations are reported for continuous variables by cluster. The other table reports frequencies of categorical variables by cluster.
- **Cluster frequencies.** Displays a table that reports the number of observations in each cluster.
- **Information criterion (AIC or BIC).** Displays a table containing the values of the AIC or BIC, depending on the criterion chosen in the main dialog box, for different numbers of clusters. This table is provided only when the number of clusters is being determined automatically. If the number of clusters is fixed, this setting is ignored, and the table is not provided.

Working Data File. This group allows you to save variables to the active dataset.

- **Create cluster membership variable.** This variable contains a cluster identification number for each case. The name of this variable is *tsc_n*, where *n* is a positive integer indicating the ordinal of the active dataset save operation completed by this procedure in a given session.

XML Files. The final cluster model and CF tree are two types of output files that can be exported in XML format.

- **Export final model.** The final cluster model is exported to the specified file in XML (PMML) format. SmartScore and SPSS Statistics Server (a separate product) can use this model file to apply the model information to other data files for scoring purposes.
- **Export CF tree.** This option allows you to save the current state of the cluster tree and update it later using newer data.

Hierarchical Cluster Analysis

This procedure attempts to identify relatively homogeneous groups of cases (or variables) based on selected characteristics, using an algorithm that starts with each case (or variable) in a separate cluster and combines clusters until only one is left. You can analyze raw variables, or you can choose from a variety of standardizing transformations. Distance or similarity measures are generated by the Proximities procedure. Statistics are displayed at each stage to help you select the best solution.

Example. Are there identifiable groups of television shows that attract similar audiences within each group? With hierarchical cluster analysis, you could cluster television shows (cases) into homogeneous groups based on viewer characteristics. This can be used to identify segments for marketing. Or you can cluster cities (cases) into homogeneous groups so that comparable cities can be selected to test various marketing strategies.

Statistics. Agglomeration schedule, distance (or similarity) matrix, and cluster membership for a single solution or a range of solutions. Plots: dendrograms and icicle plots.

Data. The variables can be quantitative, binary, or count data. Scaling of variables is an important issue—differences in scaling may affect your cluster solution(s). If your variables have large differences in scaling (for example, one variable is measured in dollars and the other is measured in years), you should consider standardizing them (this can be done automatically by the Hierarchical Cluster Analysis procedure).

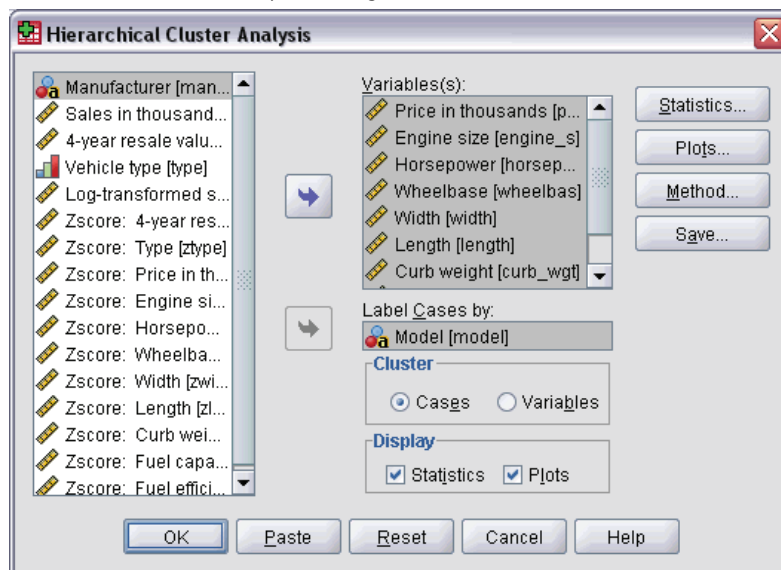
Case order. If tied distances or similarities exist in the input data or occur among updated clusters during joining, the resulting cluster solution may depend on the order of cases in the file. You may want to obtain several different solutions with cases sorted in different random orders to verify the stability of a given solution.

Assumptions. The distance or similarity measures used should be appropriate for the data analyzed (see the Proximities procedure for more information on choices of distance and similarity measures). Also, you should include all relevant variables in your analysis. Omission of influential variables can result in a misleading solution. Because hierarchical cluster analysis is an exploratory method, results should be treated as tentative until they are confirmed with an independent sample.

To Obtain a Hierarchical Cluster Analysis

- ▶ From the menus choose:
 - Analyze
 - Classify
 - Hierarchical Cluster...

Figure 37-1
Hierarchical Cluster Analysis dialog box

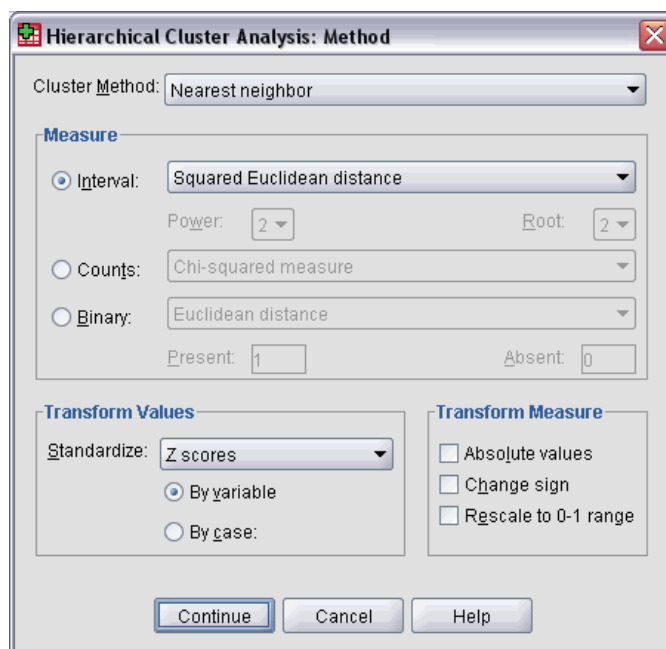


- If you are clustering cases, select at least one numeric variable. If you are clustering variables, select at least three numeric variables.

Optionally, you can select an identification variable to label cases.

Hierarchical Cluster Analysis Method

Figure 37-2
Hierarchical Cluster Analysis Method dialog box



Cluster Method. Available alternatives are between-groups linkage, within-groups linkage, nearest neighbor, furthest neighbor, centroid clustering, median clustering, and Ward's method.

Measure. Allows you to specify the distance or similarity measure to be used in clustering. Select the type of data and the appropriate distance or similarity measure:

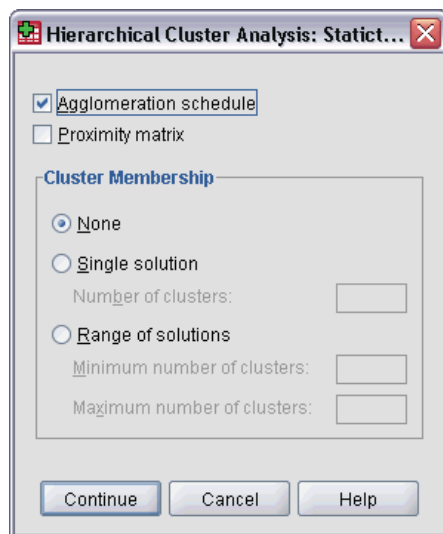
- **Interval.** Available alternatives are Euclidean distance, squared Euclidean distance, cosine, Pearson correlation, Chebychev, block, Minkowski, and customized.
- **Counts.** Available alternatives are chi-square measure and phi-square measure.
- **Binary.** Available alternatives are Euclidean distance, squared Euclidean distance, size difference, pattern difference, variance, dispersion, shape, simple matching, phi 4-point correlation, lambda, Anderberg's *D*, dice, Hamann, Jaccard, Kulczynski 1, Kulczynski 2, Lance and Williams, Ochiai, Rogers and Tanimoto, Russel and Rao, Sokal and Sneath 1, Sokal and Sneath 2, Sokal and Sneath 3, Sokal and Sneath 4, Sokal and Sneath 5, Yule's *Y*, and Yule's *Q*.

Transform Values. Allows you to standardize data values for either cases or values before computing proximities (not available for binary data). Available standardization methods are *z* scores, range -1 to 1, range 0 to 1, maximum magnitude of 1, mean of 1, and standard deviation of 1.

Transform Measures. Allows you to transform the values generated by the distance measure. They are applied after the distance measure has been computed. Available alternatives are absolute values, change sign, and rescale to 0-1 range.

Hierarchical Cluster Analysis Statistics

Figure 37-3
Hierarchical Cluster Analysis Statistics dialog box



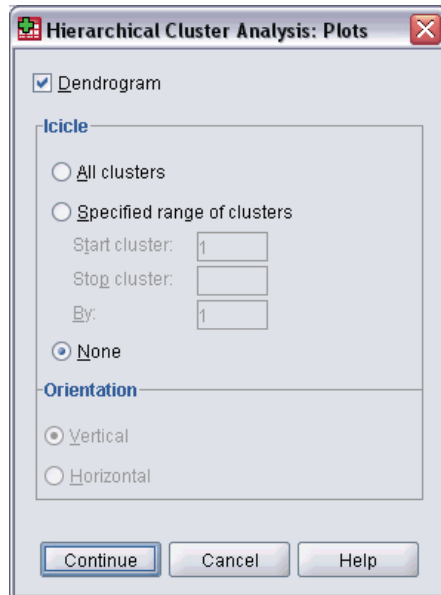
Agglomeration schedule. Displays the cases or clusters combined at each stage, the distances between the cases or clusters being combined, and the last cluster level at which a case (or variable) joined the cluster.

Proximity matrix. Gives the distances or similarities between items.

Cluster Membership. Displays the cluster to which each case is assigned at one or more stages in the combination of clusters. Available options are single solution and range of solutions.

Hierarchical Cluster Analysis Plots

Figure 37-4
Hierarchical Cluster Analysis Plots dialog box



Dendrogram. Displays a **dendrogram**. Dendrograms can be used to assess the cohesiveness of the clusters formed and can provide information about the appropriate number of clusters to keep.

Icicle. Displays an **icicle plot**, including all clusters or a specified range of clusters. Icicle plots display information about how cases are combined into clusters at each iteration of the analysis. Orientation allows you to select a vertical or horizontal plot.

Hierarchical Cluster Analysis Save New Variables

Figure 37-5
Hierarchical Cluster Analysis Save dialog box



Cluster Membership. Allows you to save cluster memberships for a single solution or a range of solutions. Saved variables can then be used in subsequent analyses to explore other differences between groups.

CLUSTER Command Syntax Additional Features

The Hierarchical Cluster procedure uses `CLUSTER` command syntax. The command syntax language also allows you to:

- Use several clustering methods in a single analysis.
- Read and analyze a proximity matrix.
- Write a proximity matrix to disk for later analysis.
- Specify any values for power and root in the customized (Power) distance measure.
- Specify names for saved variables.

See the *Command Syntax Reference* for complete syntax information.

K-Means Cluster Analysis

This procedure attempts to identify relatively homogeneous groups of cases based on selected characteristics, using an algorithm that can handle large numbers of cases. However, the algorithm requires you to specify the number of clusters. You can specify initial cluster centers if you know this information. You can select one of two methods for classifying cases, either updating cluster centers iteratively or classifying only. You can save cluster membership, distance information, and final cluster centers. Optionally, you can specify a variable whose values are used to label casewise output. You can also request analysis of variance F statistics. While these statistics are opportunistic (the procedure tries to form groups that do differ), the relative size of the statistics provides information about each variable's contribution to the separation of the groups.

Example. What are some identifiable groups of television shows that attract similar audiences within each group? With k -means cluster analysis, you could cluster television shows (cases) into k homogeneous groups based on viewer characteristics. This process can be used to identify segments for marketing. Or you can cluster cities (cases) into homogeneous groups so that comparable cities can be selected to test various marketing strategies.

Statistics. Complete solution: initial cluster centers, ANOVA table. Each case: cluster information, distance from cluster center.

Data. Variables should be quantitative at the interval or ratio level. If your variables are binary or counts, use the Hierarchical Cluster Analysis procedure.

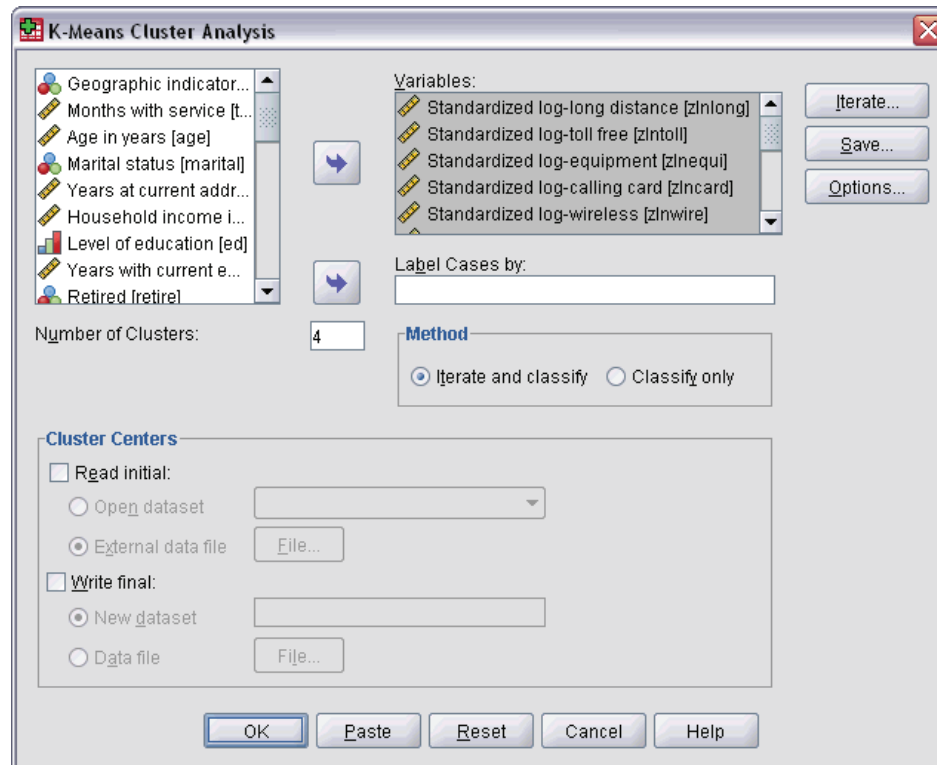
Case and initial cluster center order. The default algorithm for choosing initial cluster centers is not invariant to case ordering. The Use running means option in the Iterate dialog box makes the resulting solution potentially dependent on case order, regardless of how initial cluster centers are chosen. If you are using either of these methods, you may want to obtain several different solutions with cases sorted in different random orders to verify the stability of a given solution. Specifying initial cluster centers and not using the Use running means option will avoid issues related to case order. However, ordering of the initial cluster centers may affect the solution if there are tied distances from cases to cluster centers. To assess the stability of a given solution, you can compare results from analyses with different permutations of the initial center values.

Assumptions. Distances are computed using simple Euclidean distance. If you want to use another distance or similarity measure, use the Hierarchical Cluster Analysis procedure. Scaling of variables is an important consideration. If your variables are measured on different scales (for example, one variable is expressed in dollars and another variable is expressed in years), your results may be misleading. In such cases, you should consider standardizing your variables before you perform the k -means cluster analysis (this task can be done in the Descriptives procedure). The procedure assumes that you have selected the appropriate number of clusters and that you have included all relevant variables. If you have chosen an inappropriate number of clusters or omitted important variables, your results may be misleading.

To Obtain a K-Means Cluster Analysis

- ▶ From the menus choose:
Analyze
 Classify
 K-Means Cluster...

Figure 38-1
K-Means Cluster Analysis dialog box



- ▶ Select the variables to be used in the cluster analysis.
- ▶ Specify the number of clusters. (The number of clusters must be at least 2 and must not be greater than the number of cases in the data file.)
- ▶ Select either Iterate and classify or Classify only.
- ▶ Optionally, select an identification variable to label cases.

K-Means Cluster Analysis Efficiency

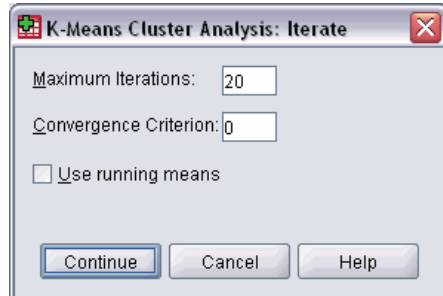
The k -means cluster analysis command is efficient primarily because it does not compute the distances between all pairs of cases, as do many clustering algorithms, including the algorithm that is used by the hierarchical clustering command.

For maximum efficiency, take a sample of cases and select the Iterate and classify method to determine cluster centers. Select Write final as. Then restore the entire data file and select Classify only as the method and select Read initial from to classify the entire file using the centers that

are estimated from the sample. You can write to and read from a file or a dataset. Datasets are available for subsequent use in the same session but are not saved as files unless explicitly saved prior to the end of the session. Dataset names must conform to variable-naming rules. For more information, see [Variable Names in Chapter 5 on p. 72](#).

K-Means Cluster Analysis Iterate

Figure 38-2
K-Means Cluster Analysis Iterate dialog box



Note: These options are available only if you select the Iterate and classify method from the K-Means Cluster Analysis dialog box.

Maximum Iterations. Limits the number of iterations in the k -means algorithm. Iteration stops after this many iterations even if the convergence criterion is not satisfied. This number must be between 1 and 999.

To reproduce the algorithm used by the Quick Cluster command prior to version 5.0, set Maximum Iterations to 1.

Convergence Criterion. Determines when iteration ceases. It represents a proportion of the minimum distance between initial cluster centers, so it must be greater than 0 but not greater than 1. If the criterion equals 0.02, for example, iteration ceases when a complete iteration does not move any of the cluster centers by a distance of more than 2% of the smallest distance between any initial cluster centers.

Use running means. Allows you to request that cluster centers be updated after each case is assigned. If you do not select this option, new cluster centers are calculated after all cases have been assigned.

K-Means Cluster Analysis Save

Figure 38-3
K-Means Cluster Analysis Save New Variables dialog box



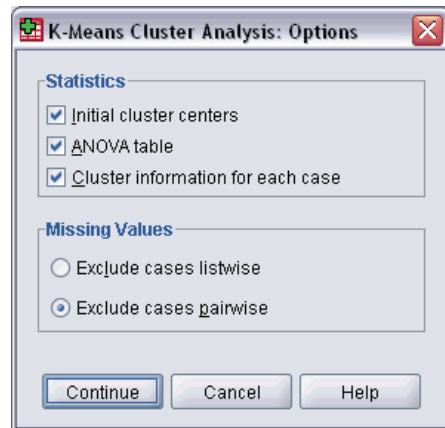
You can save information about the solution as new variables to be used in subsequent analyses:

Cluster membership. Creates a new variable indicating the final cluster membership of each case. Values of the new variable range from 1 to the number of clusters.

Distance from cluster center. Creates a new variable indicating the Euclidean distance between each case and its classification center.

K-Means Cluster Analysis Options

Figure 38-4
K-Means Cluster Analysis Options dialog box



Statistics. You can select the following statistics: initial cluster centers, ANOVA table, and cluster information for each case.

- **Initial cluster centers.** First estimate of the variable means for each of the clusters. By default, a number of well-spaced cases equal to the number of clusters is selected from the data. Initial cluster centers are used for a first round of classification and are then updated.
- **ANOVA table.** Displays an analysis-of-variance table which includes univariate F tests for each clustering variable. The F tests are only descriptive and the resulting probabilities should not be interpreted. The ANOVA table is not displayed if all cases are assigned to a single cluster.
- **Cluster information for each case.** Displays for each case the final cluster assignment and the Euclidean distance between the case and the cluster center used to classify the case. Also displays Euclidean distance between final cluster centers.

Missing Values. Available options are Exclude cases listwise or Exclude cases pairwise.

- **Exclude cases listwise.** Excludes cases with missing values for any clustering variable from the analysis.
- **Exclude cases pairwise.** Assigns cases to clusters based on distances that are computed from all variables with nonmissing values.

QUICK CLUSTER Command Additional Features

The K-Means Cluster procedure uses `QUICK CLUSTER` command syntax. The command syntax language also allows you to:

- Accept the first k cases as initial cluster centers, thereby avoiding the data pass that is normally used to estimate them.
- Specify initial cluster centers directly as a part of the command syntax.
- Specify names for saved variables.

See the *Command Syntax Reference* for complete syntax information.

Nonparametric Tests

The Nonparametric Tests procedure provides several tests that do not require assumptions about the shape of the underlying distribution.

Chi-Square Test. Tabulates a variable into categories and computes a chi-square statistic based on the differences between observed and expected frequencies.

Binomial Test. Compares the observed frequency in each category of a dichotomous variable with expected frequencies from the binomial distribution.

Runs Test. Tests whether the order of occurrence of two values of a variable is random.

One-Sample Kolmogorov-Smirnov Test. Compares the observed cumulative distribution function for a variable with a specified theoretical distribution, which may be normal, uniform, exponential, or Poisson.

Two-Independent-Samples Tests. Compares two groups of cases on one variable. The Mann-Whitney U test, two-sample Kolmogorov-Smirnov test, Moses test of extreme reactions, and Wald-Wolfowitz runs test are available.

Two-Related-Samples Tests. Compares the distributions of two variables. The Wilcoxon signed-rank test, the sign test, and the McNemar test are available.

Tests for Several Independent Samples. Compares two or more groups of cases on one variable. The Kruskal-Wallis test, the Median test, and the Jonckheere-Terpstra test are available.

Tests for Several Related Samples. Compares the distributions of two or more variables. Friedman's test, Kendall's W , and Cochran's Q are available.

Quartiles and the mean, standard deviation, minimum, maximum, and number of nonmissing cases are available for all of the above tests.

Chi-Square Test

The Chi-Square Test procedure tabulates a variable into categories and computes a chi-square statistic. This goodness-of-fit test compares the observed and expected frequencies in each category to test that all categories contain the same proportion of values or test that each category contains a user-specified proportion of values.

Examples. The chi-square test could be used to determine whether a bag of jelly beans contains equal proportions of blue, brown, green, orange, red, and yellow candies. You could also test to see whether a bag of jelly beans contains 5% blue, 30% brown, 10% green, 20% orange, 15% red, and 15% yellow candies.

Statistics. Mean, standard deviation, minimum, maximum, and quartiles. The number and the percentage of nonmissing and missing cases; the number of cases observed and expected for each category; residuals; and the chi-square statistic.

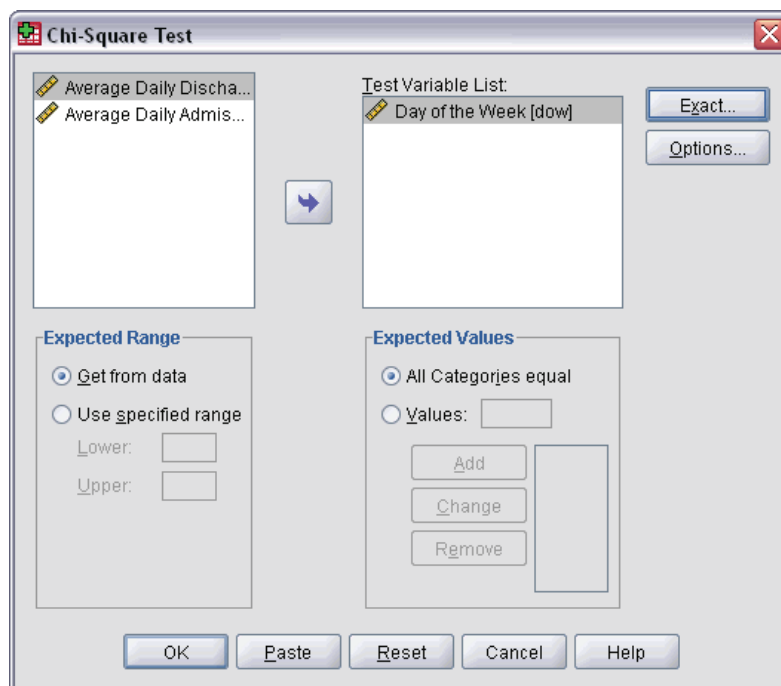
Data. Use ordered or unordered numeric categorical variables (ordinal or nominal levels of measurement). To convert string variables to numeric variables, use the Automatic Recode procedure, which is available on the Transform menu.

Assumptions. Nonparametric tests do not require assumptions about the shape of the underlying distribution. The data are assumed to be a random sample. The expected frequencies for each category should be at least 1. No more than 20% of the categories should have expected frequencies of less than 5.

To Obtain a Chi-Square Test

- ▶ From the menu choose:
 - Analyze
 - Nonparametric Tests
 - Chi-Square...

Figure 39-1
Chi-Square Test dialog box



- ▶ Select one or more test variables. Each variable produces a separate test.
- ▶ Optionally, click Options for descriptive statistics, quartiles, and control of the treatment of missing data.

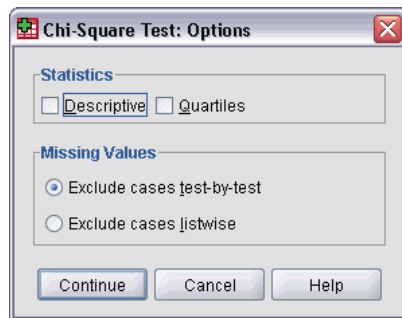
Chi-Square Test Expected Range and Expected Values

Expected Range. By default, each distinct value of the variable is defined as a category. To establish categories within a specific range, select Use specified range and enter integer values for lower and upper bounds. Categories are established for each integer value within the inclusive range, and cases with values outside of the bounds are excluded. For example, if you specify a value of 1 for Lower and a value of 4 for Upper, only the integer values of 1 through 4 are used for the chi-square test.

Expected Values. By default, all categories have equal expected values. Categories can have user-specified expected proportions. Select Values, enter a value that is greater than 0 for each category of the test variable, and then click Add. Each time you add a value, it appears at the bottom of the value list. The order of the values is important; it corresponds to the ascending order of the category values of the test variable. The first value of the list corresponds to the lowest group value of the test variable, and the last value corresponds to the highest value. Elements of the value list are summed, and then each value is divided by this sum to calculate the proportion of cases expected in the corresponding category. For example, a value list of 3, 4, 5, 4 specifies expected proportions of 3/16, 4/16, 5/16, and 4/16.

Chi-Square Test Options

Figure 39-2
Chi-Square Test Options dialog box



Statistics. You can choose one or both summary statistics.

- **Descriptive.** Displays the mean, standard deviation, minimum, maximum, and number of nonmissing cases.
- **Quartiles.** Displays values corresponding to the 25th, 50th, and 75th percentiles.

Missing Values. Controls the treatment of missing values.

- **Exclude cases test-by-test.** When several tests are specified, each test is evaluated separately for missing values.
- **Exclude cases listwise.** Cases with missing values for any variable are excluded from all analyses.

NPARTESTS Command Additional Features (Chi-Square Test)

The command syntax language also allows you to:

- Specify different minimum and maximum values or expected frequencies for different variables (with the `CHISQUARE` subcommand).
- Test the same variable against different expected frequencies or use different ranges (with the `EXPECTED` subcommand).

See the *Command Syntax Reference* for complete syntax information.

Binomial Test

The Binomial Test procedure compares the observed frequencies of the two categories of a dichotomous variable to the frequencies that are expected under a binomial distribution with a specified probability parameter. By default, the probability parameter for both groups is 0.5. To change the probabilities, you can enter a test proportion for the first group. The probability for the second group will be 1 minus the specified probability for the first group.

Example. When you toss a dime, the probability of a head equals 1/2. Based on this hypothesis, a dime is tossed 40 times, and the outcomes are recorded (heads or tails). From the binomial test, you might find that 3/4 of the tosses were heads and that the observed significance level is small (0.0027). These results indicate that it is not likely that the probability of a head equals 1/2; the coin is probably biased.

Statistics. Mean, standard deviation, minimum, maximum, number of nonmissing cases, and quartiles.

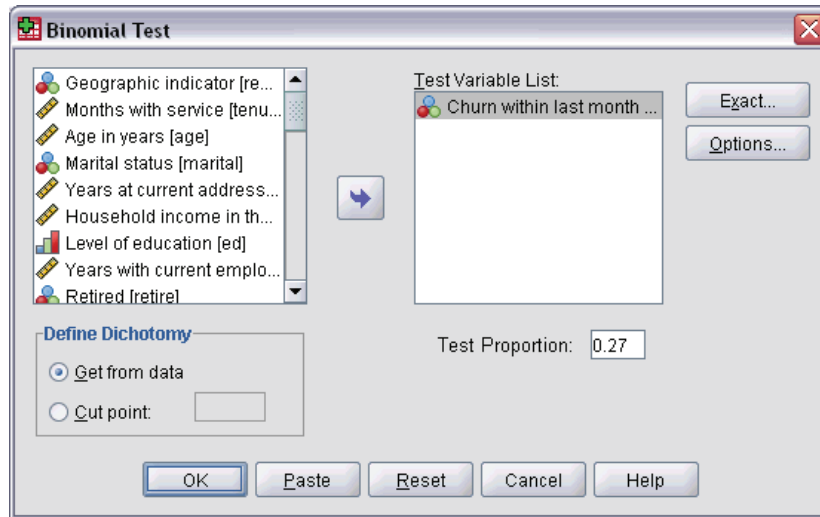
Data. The variables that are tested should be numeric and dichotomous. To convert string variables to numeric variables, use the Automatic Recode procedure, which is available on the Transform menu. A **dichotomous variable** is a variable that can take only two possible values: *yes* or *no*, *true* or *false*, 0 or 1, and so on. The first value encountered in the dataset defines the first group, and the other value defines the second group. If the variables are not dichotomous, you must specify a cut point. The cut point assigns cases with values that are less than or equal to the cut point to the first group and assigns the rest of the cases to the second group.

Assumptions. Nonparametric tests do not require assumptions about the shape of the underlying distribution. The data are assumed to be a random sample.

To Obtain a Binomial Test

- ▶ From the menus choose:
 - Analyze
 - Nonparametric Tests
 - Binomial...

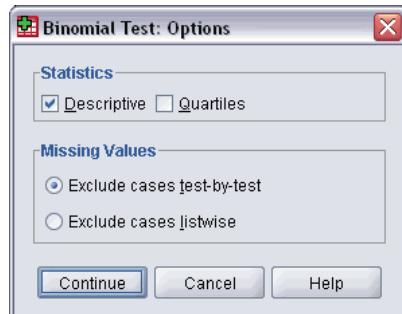
Figure 39-3
Binomial Test dialog box



- ▶ Select one or more numeric test variables.
- ▶ Optionally, click Options for descriptive statistics, quartiles, and control of the treatment of missing data.

Binomial Test Options

Figure 39-4
Binomial Test Options dialog box



Statistics. You can choose one or both summary statistics.

- **Descriptive.** Displays the mean, standard deviation, minimum, maximum, and number of nonmissing cases.
- **Quartiles.** Displays values corresponding to the 25th, 50th, and 75th percentiles.

Missing Values. Controls the treatment of missing values.

- **Exclude cases test-by-test.** When several tests are specified, each test is evaluated separately for missing values.
- **Exclude cases listwise.** Cases with missing values for any variable that is tested are excluded from all analyses.

NPAR TESTS Command Additional Features (Binomial Test)

The command syntax language also allows you to:

- Select specific groups (and exclude other groups) when a variable has more than two categories (with the `BINOMIAL` subcommand).
- Specify different cut points or probabilities for different variables (with the `BINOMIAL` subcommand).
- Test the same variable against different cut points or probabilities (with the `EXPECTED` subcommand).

See the *Command Syntax Reference* for complete syntax information.

Runs Test

The Runs Test procedure tests whether the order of occurrence of two values of a variable is random. A run is a sequence of like observations. A sample with too many or too few runs suggests that the sample is not random.

Examples. Suppose that 20 people are polled to find out whether they would purchase a product. The assumed randomness of the sample would be seriously questioned if all 20 people were of the same gender. The runs test can be used to determine whether the sample was drawn at random.

Statistics. Mean, standard deviation, minimum, maximum, number of nonmissing cases, and quartiles.

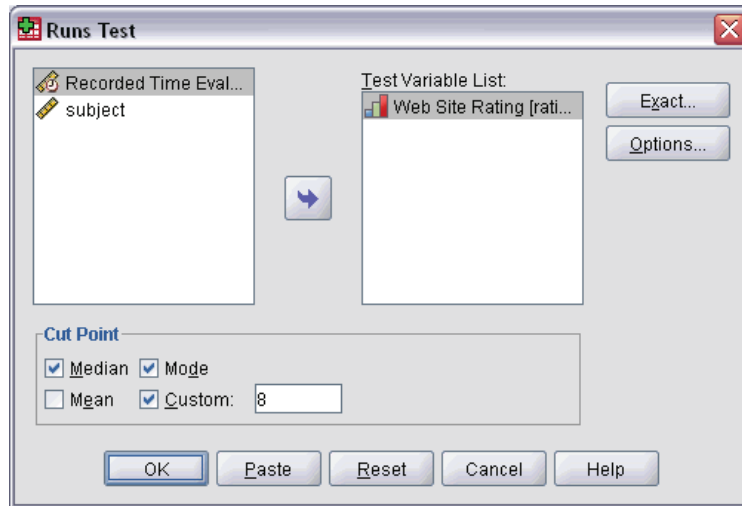
Data. The variables must be numeric. To convert string variables to numeric variables, use the Automatic Recode procedure, which is available on the Transform menu.

Assumptions. Nonparametric tests do not require assumptions about the shape of the underlying distribution. Use samples from continuous probability distributions.

To Obtain a Runs Test

- ▶ From the menus choose:
 - Analyze
 - Nonparametric Tests
 - Runs...

Figure 39-5
Adding a custom cut point



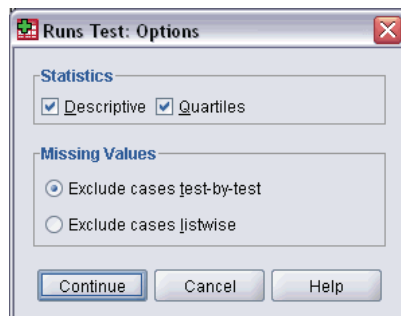
- ▶ Select one or more numeric test variables.
- ▶ Optionally, click Options for descriptive statistics, quartiles, and control of the treatment of missing data.

Runs Test Cut Point

Cut Point. Specifies a cut point to dichotomize the variables that you have chosen. You can use the observed mean, median, or mode, or you can use a specified value as a cut point. Cases with values that are less than the cut point are assigned to one group, and cases with values that are greater than or equal to the cut point are assigned to another group. One test is performed for each chosen cut point.

Runs Test Options

Figure 39-6
Runs Test Options dialog box



Statistics. You can choose one or both summary statistics.

- **Descriptive.** Displays the mean, standard deviation, minimum, maximum, and number of nonmissing cases.
- **Quartiles.** Displays values corresponding to the 25th, 50th, and 75th percentiles.

Missing Values. Controls the treatment of missing values.

- **Exclude cases test-by-test.** When several tests are specified, each test is evaluated separately for missing values.
- **Exclude cases listwise.** Cases with missing values for any variable are excluded from all analyses.

NPAR TESTS Command Additional Features (Runs Test)

The command syntax language also allows you to:

- Specify different cut points for different variables (with the `RUNS` subcommand).
- Test the same variable against different custom cut points (with the `RUNS` subcommand).

See the *Command Syntax Reference* for complete syntax information.

One-Sample Kolmogorov-Smirnov Test

The One-Sample Kolmogorov-Smirnov Test procedure compares the observed cumulative distribution function for a variable with a specified theoretical distribution, which may be normal, uniform, Poisson, or exponential. The Kolmogorov-Smirnov Z is computed from the largest difference (in absolute value) between the observed and theoretical cumulative distribution functions. This goodness-of-fit test tests whether the observations could reasonably have come from the specified distribution.

Example. Many parametric tests require normally distributed variables. The one-sample Kolmogorov-Smirnov test can be used to test that a variable (for example, *income*) is normally distributed.

Statistics. Mean, standard deviation, minimum, maximum, number of nonmissing cases, and quartiles.

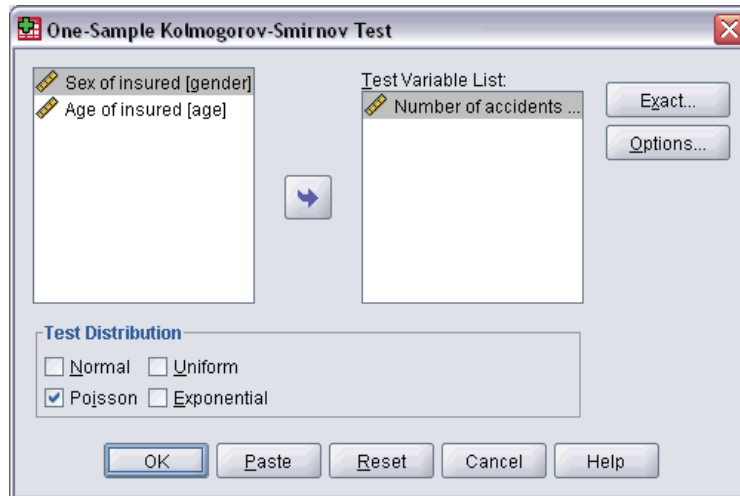
Data. Use quantitative variables (interval or ratio level of measurement).

Assumptions. The Kolmogorov-Smirnov test assumes that the parameters of the test distribution are specified in advance. This procedure estimates the parameters from the sample. The sample mean and sample standard deviation are the parameters for a normal distribution, the sample minimum and maximum values define the range of the uniform distribution, the sample mean is the parameter for the Poisson distribution, and the sample mean is the parameter for the exponential distribution. The power of the test to detect departures from the hypothesized distribution may be seriously diminished. For testing against a normal distribution with estimated parameters, consider the adjusted K-S Lilliefors test (available in the Explore procedure).

To Obtain a One-Sample Kolmogorov-Smirnov Test

- ▶ From the menu choose:
 - Analyze
 - Nonparametric Tests
 - 1-Sample K-S...

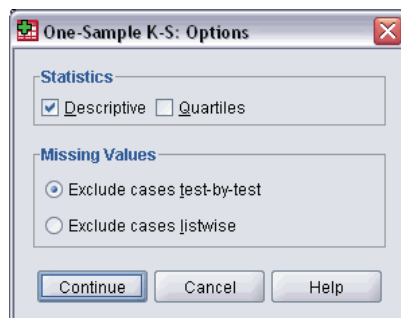
Figure 39-7
One-Sample Kolmogorov-Smirnov Test dialog box



- ▶ Select one or more numeric test variables. Each variable produces a separate test.
- ▶ Optionally, click Options for descriptive statistics, quartiles, and control of the treatment of missing data.

One-Sample Kolmogorov-Smirnov Test Options

Figure 39-8
One-Sample K-S Options dialog box



Statistics. You can choose one or both summary statistics.

- **Descriptive.** Displays the mean, standard deviation, minimum, maximum, and number of nonmissing cases.
- **Quartiles.** Displays values corresponding to the 25th, 50th, and 75th percentiles.

Missing Values. Controls the treatment of missing values.

- **Exclude cases test-by-test.** When several tests are specified, each test is evaluated separately for missing values.
- **Exclude cases listwise.** Cases with missing values for any variable are excluded from all analyses.

NPAR TESTS Command Additional Features (One-Sample Kolmogorov-Smirnov Test)

The command syntax language also allows you to specify the parameters of the test distribution (with the `K-S` subcommand).

See the *Command Syntax Reference* for complete syntax information.

Two-Independent-Samples Tests

The Two-Independent-Samples Tests procedure compares two groups of cases on one variable.

Example. New dental braces have been developed that are intended to be more comfortable, to look better, and to provide more rapid progress in realigning teeth. To find out whether the new braces have to be worn as long as the old braces, 10 children are randomly chosen to wear the old braces, and another 10 children are chosen to wear the new braces. From the Mann-Whitney U test, you might find that, on average, children with the new braces did not have to wear the braces as long as children with the old braces.

Statistics. Mean, standard deviation, minimum, maximum, number of nonmissing cases, and quartiles. Tests: Mann-Whitney U , Moses extreme reactions, Kolmogorov-Smirnov Z , Wald-Wolfowitz runs.

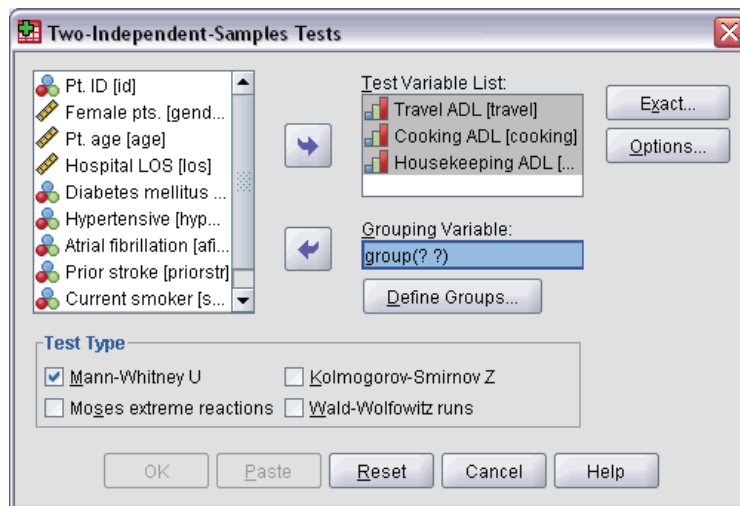
Data. Use numeric variables that can be ordered.

Assumptions. Use independent, random samples. The Mann-Whitney U test tests equality of two distributions. In order to use it to test for differences in location between two distributions, one must assume that the distributions have the same shape.

To Obtain Two-Independent-Samples Tests

- ▶ From the menus choose:
 - Analyze
 - Nonparametric Tests
 - 2 Independent Samples...

Figure 39-9
Two-Independent-Samples Tests dialog box



- ▶ Select one or more numeric variables.
- ▶ Select a grouping variable and click Define Groups to split the file into two groups or samples.

Two-Independent-Samples Test Types

Test Type. Four tests are available to test whether two independent samples (groups) come from the same population.

The **Mann-Whitney U test** is the most popular of the two-independent-samples tests. It is equivalent to the Wilcoxon rank sum test and the Kruskal-Wallis test for two groups. Mann-Whitney tests that two sampled populations are equivalent in location. The observations from both groups are combined and ranked, with the average rank assigned in the case of ties. The number of ties should be small relative to the total number of observations. If the populations are identical in location, the ranks should be randomly mixed between the two samples. The test calculates the number of times that a score from group 1 precedes a score from group 2 and the number of times that a score from group 2 precedes a score from group 1. The Mann-Whitney U statistic is the smaller of these two numbers. The Wilcoxon rank sum W statistic is also displayed. If both samples have the same number of observations, W is the rank sum of the group that is named first in the Two-Independent-Samples Define Groups dialog box.


The **Kolmogorov-Smirnov Z test** and the **Wald-Wolfowitz runs test** are more general tests that detect differences in both the locations and shapes of the distributions. The Kolmogorov-Smirnov test is based on the maximum absolute difference between the observed cumulative distribution functions for both samples. When this difference is significantly large, the two distributions are considered different. The Wald-Wolfowitz runs test combines and ranks the observations from both groups. If the two samples are from the same population, the two groups should be randomly scattered throughout the ranking.

The **Moses extreme reactions test** assumes that the experimental variable will affect some subjects in one direction and other subjects in the opposite direction. The test tests for extreme responses compared to a control group. This test focuses on the span of the control group

and is a measure of how much extreme values in the experimental group influence the span when combined with the control group. The control group is defined by the group 1 value in the Two-Independent-Samples Define Groups dialog box. Observations from both groups are combined and ranked. The span of the control group is computed as the difference between the ranks of the largest and smallest values in the control group plus 1. Because chance outliers can easily distort the range of the span, 5% of the control cases are trimmed automatically from each end.

Two-Independent-Samples Tests Define Groups

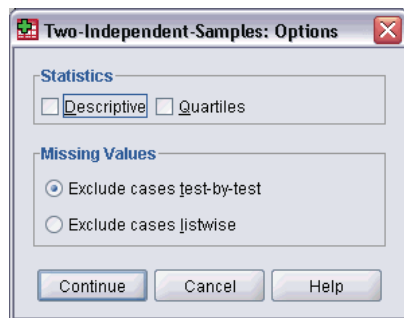
Figure 39-10
Two-Independent-Samples Define Groups dialog box



To split the file into two groups or samples, enter an integer value for Group 1 and another value for Group 2. Cases with other values are excluded from the analysis.

Two-Independent-Samples Tests Options

Figure 39-11
Two-Independent-Samples Options dialog box



Statistics. You can choose one or both summary statistics.

- **Descriptive.** Displays the mean, standard deviation, minimum, maximum, and the number of nonmissing cases.
- **Quartiles.** Displays values corresponding to the 25th, 50th, and 75th percentiles.

Missing Values. Controls the treatment of missing values.

- **Exclude cases test-by-test.** When several tests are specified, each test is evaluated separately for missing values.
- **Exclude cases listwise.** Cases with missing values for any variable are excluded from all analyses.

NPAR TESTS Command Additional Features (Two-Independent-Samples Tests)

The command syntax language also allows you to specify the number of cases to be trimmed for the Moses test (with the `MOSES` subcommand).

See the *Command Syntax Reference* for complete syntax information.

Two-Related-Samples Tests

The Two-Related-Samples Tests procedure compares the distributions of two variables.

Example. In general, do families receive the asking price when they sell their homes? By applying the Wilcoxon signed-rank test to data for 10 homes, you might learn that seven families receive less than the asking price, one family receives more than the asking price, and two families receive the asking price.

Statistics. Mean, standard deviation, minimum, maximum, number of nonmissing cases, and quartiles. Tests: Wilcoxon signed-rank, sign, McNemar. If the Exact Tests option is installed (available only on Windows operating systems), the marginal homogeneity test is also available.

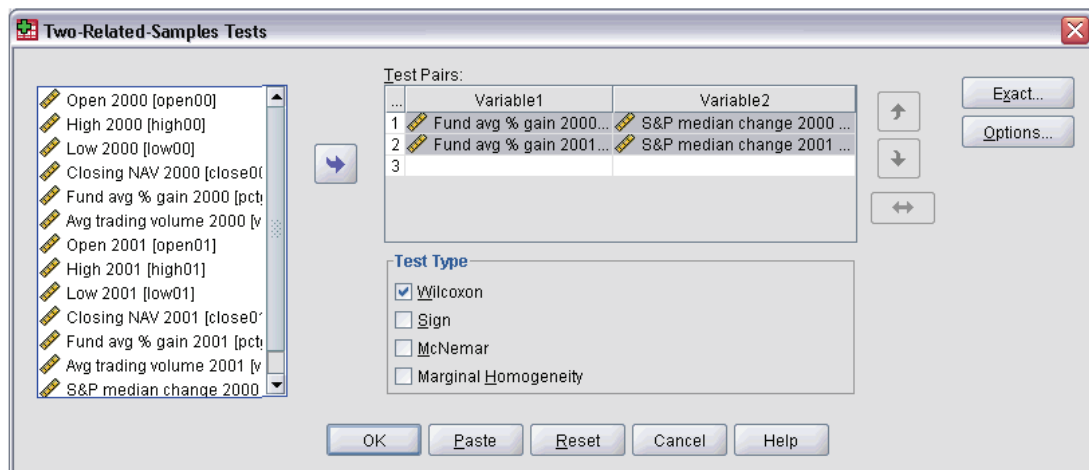
Data. Use numeric variables that can be ordered.

Assumptions. Although no particular distributions are assumed for the two variables, the population distribution of the paired differences is assumed to be symmetric.

To Obtain Two-Related-Samples Tests

- ▶ From the menus choose:
 - Analyze
 - Nonparametric Tests
 - 2 Related Samples...

Figure 39-12
Two-Related-Samples Tests dialog box



- ▶ Select one or more pairs of variables.

Two-Related-Samples Test Types

The tests in this section compare the distributions of two related variables. The appropriate test to use depends on the type of data.

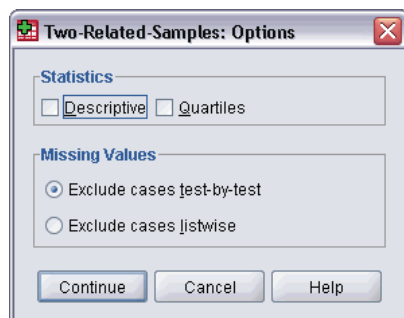
If your data are continuous, use the sign test or the Wilcoxon signed-rank test. The **sign test** computes the differences between the two variables for all cases and classifies the differences as positive, negative, or tied. If the two variables are similarly distributed, the number of positive and negative differences will not differ significantly. The **Wilcoxon signed-rank test** considers information about both the sign of the differences and the magnitude of the differences between pairs. Because the Wilcoxon signed-rank test incorporates more information about the data, it is more powerful than the sign test.

If your data are binary, use the **McNemar test**. This test is typically used in a repeated measures situation, in which each subject's response is elicited twice, once before and once after a specified event occurs. The McNemar test determines whether the initial response rate (before the event) equals the final response rate (after the event). This test is useful for detecting changes in responses due to experimental intervention in before-and-after designs.

If your data are categorical, use the **marginal homogeneity test**. This test is an extension of the McNemar test from binary response to multinomial response. It tests for changes in response (using the chi-square distribution) and is useful for detecting response changes due to experimental intervention in before-and-after designs. The marginal homogeneity test is available only if you have installed Exact Tests.

Two-Related-Samples Tests Options

Figure 39-13
Two-Related-Samples Options dialog box



Statistics. You can choose one or both summary statistics.

- **Descriptive.** Displays the mean, standard deviation, minimum, maximum, and the number of nonmissing cases.
- **Quartiles.** Displays values corresponding to the 25th, 50th, and 75th percentiles.

Missing Values. Controls the treatment of missing values.

- **Exclude cases test-by-test.** When several tests are specified, each test is evaluated separately for missing values.
- **Exclude cases listwise.** Cases with missing values for any variable are excluded from all analyses.

NPART TESTS Command Additional Features (Two Related Samples)

The command syntax language also allows you to test a variable with each variable on a list. See the *Command Syntax Reference* for complete syntax information.

Tests for Several Independent Samples

The Tests for Several Independent Samples procedure compares two or more groups of cases on one variable.

Example. Do three brands of 100-watt lightbulbs differ in the average time that the bulbs will burn? From the Kruskal-Wallis one-way analysis of variance, you might learn that the three brands do differ in average lifetime.

Statistics. Mean, standard deviation, minimum, maximum, number of nonmissing cases, and quartiles. Tests: Kruskal-Wallis H , median.

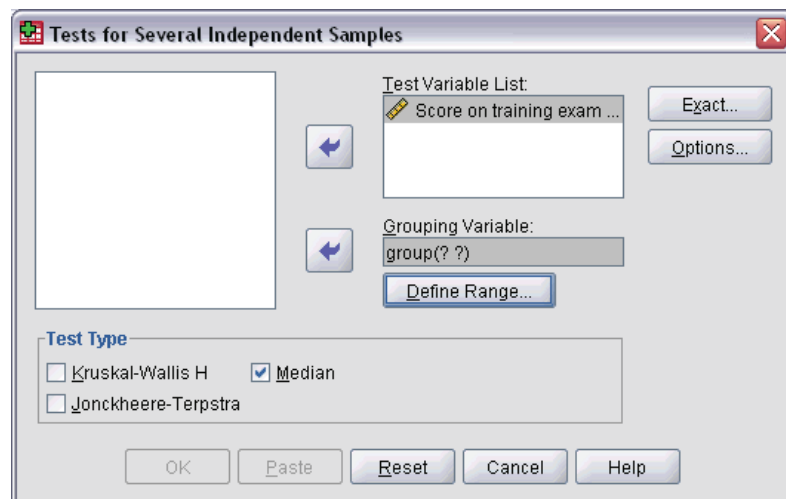
Data. Use numeric variables that can be ordered.

Assumptions. Use independent, random samples. The Kruskal-Wallis H test requires that the tested samples be similar in shape.

To Obtain Tests for Several Independent Samples

- ▶ From the menus choose:
 - Analyze
 - Nonparametric Tests
 - K Independent Samples...

Figure 39-14
Defining the median test



- ▶ Select one or more numeric variables.
- ▶ Select a grouping variable and click Define Range to specify minimum and maximum integer values for the grouping variable.

Tests for Several Independent Samples Test Types


Three tests are available to determine if several independent samples come from the same population. The Kruskal-Wallis H test, the median test, and the Jonckheere-Terpstra test all test whether several independent samples are from the same population.

The **Kruskal-Wallis H test**, an extension of the Mann-Whitney U test, is the nonparametric analog of one-way analysis of variance and detects differences in distribution location. The **median test**, which is a more general test (but not as powerful), detects distributional differences in location and shape. The Kruskal-Wallis H test and the median test assume that there is no *a priori* ordering of the k populations from which the samples are drawn.

When there is a natural *a priori* ordering (ascending or descending) of the k populations, the **Jonckheere-Terpstra test** is more powerful. For example, the k populations might represent k increasing temperatures. The hypothesis that different temperatures produce the same response distribution is tested against the alternative that as the temperature increases, the magnitude of the response increases. Here, the alternative hypothesis is ordered; therefore, Jonckheere-Terpstra is the most appropriate test to use. The Jonckheere-Terpstra test is available only if you have installed the Exact Tests add-on module.

Tests for Several Independent Samples Define Range

Figure 39-15
Several Independent Samples Define Range dialog box



To define the range, enter integer values for Minimum and Maximum that correspond to the lowest and highest categories of the grouping variable. Cases with values outside of the bounds are excluded. For example, if you specify a minimum value of 1 and a maximum value of 3, only the integer values of 1 through 3 are used. The minimum value must be less than the maximum value, and both values must be specified.

Tests for Several Independent Samples Options

Figure 39-16
Several Independent Samples Options dialog box



Statistics. You can choose one or both summary statistics.

- **Descriptive.** Displays the mean, standard deviation, minimum, maximum, and the number of nonmissing cases.
- **Quartiles.** Displays values corresponding to the 25th, 50th, and 75th percentiles.

Missing Values. Controls the treatment of missing values.

- **Exclude cases test-by-test.** When several tests are specified, each test is evaluated separately for missing values.
- **Exclude cases listwise.** Cases with missing values for any variable are excluded from all analyses.

NPAR TESTS Command Additional Features (K Independent Samples)

The command syntax language also allows you to specify a value other than the observed median for the median test (with the `MEDIAN` subcommand).

See the *Command Syntax Reference* for complete syntax information.

Tests for Several Related Samples

The Tests for Several Related Samples procedure compares the distributions of two or more variables.

Example. Does the public associate different amounts of prestige with a doctor, a lawyer, a police officer, and a teacher? Ten people are asked to rank these four occupations in order of prestige. Friedman's test indicates that the public does associate different amounts of prestige with these four professions.

Statistics. Mean, standard deviation, minimum, maximum, number of nonmissing cases, and quartiles. Tests: Friedman, Kendall's W , and Cochran's Q .

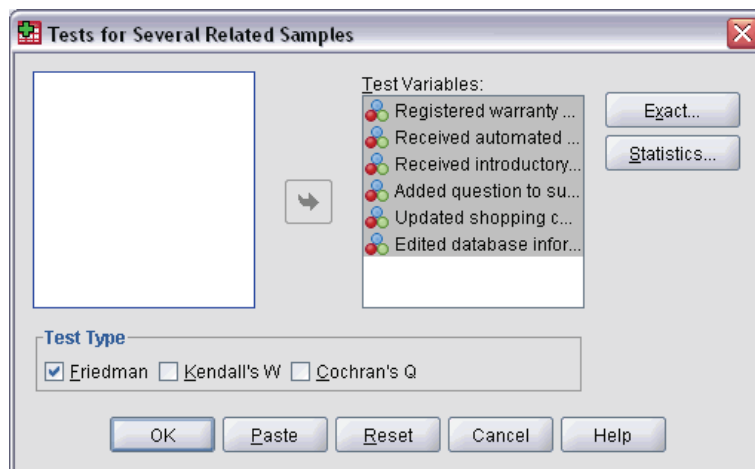
Data. Use numeric variables that can be ordered.

Assumptions. Nonparametric tests do not require assumptions about the shape of the underlying distribution. Use dependent, random samples.

To Obtain Tests for Several Related Samples

- ▶ From the menus choose:
 - Analyze
 - Nonparametric Tests
 - K Related Samples...

Figure 39-17
Selecting Cochran as the test type



- ▶ Select two or more numeric test variables.

Tests for Several Related Samples Test Types

Three tests are available to compare the distributions of several related variables.


The **Friedman test** is the nonparametric equivalent of a one-sample repeated measures design or a two-way analysis of variance with one observation per cell. Friedman tests the null hypothesis that k related variables come from the same population. For each case, the k variables are ranked from 1 to k . The test statistic is based on these ranks.

Kendall's W is a normalization of the Friedman statistic. Kendall's W is interpretable as the coefficient of concordance, which is a measure of agreement among raters. Each case is a judge or rater, and each variable is an item or person being judged. For each variable, the sum of ranks is computed. Kendall's W ranges between 0 (no agreement) and 1 (complete agreement).

Cochran's Q is identical to the Friedman test but is applicable when all responses are binary. This test is an extension of the McNemar test to the k -sample situation. Cochran's Q tests the hypothesis that several related dichotomous variables have the same mean. The variables are measured on the same individual or on matched individuals.

Tests for Several Related Samples Statistics

Figure 39-18
Several Related Samples Statistics dialog box



You can choose statistics.

- **Descriptive.** Displays the mean, standard deviation, minimum, maximum, and the number of nonmissing cases.
- **Quartiles.** Displays values corresponding to the 25th, 50th, and 75th percentiles.

NPAR TESTS Command Additional Features (K Related Samples)

See the *Command Syntax Reference* for complete syntax information.

Multiple Response Analysis

Two procedures are available for analyzing multiple dichotomy and multiple category sets. The Multiple Response Frequencies procedure displays frequency tables. The Multiple Response Crosstabs procedure displays two- and three-dimensional crosstabulations. Before using either procedure, you must define multiple response sets.

Example. This example illustrates the use of multiple response items in a market research survey. The data are fictitious and should not be interpreted as real. An airline might survey passengers flying a particular route to evaluate competing carriers. In this example, American Airlines wants to know about its passengers' use of other airlines on the Chicago-New York route and the relative importance of schedule and service in selecting an airline. The flight attendant hands each passenger a brief questionnaire upon boarding. The first question reads: Circle all airlines you have flown at least once in the last six months on this route—American, United, TWA, USAir, Other. This is a multiple response question, since the passenger can circle more than one response. However, this question cannot be coded directly because a variable can have only one value for each case. You must use several variables to map responses to each question. There are two ways to do this. One is to define a variable corresponding to each of the choices (for example, American, United, TWA, USAir, and Other). If the passenger circles United, the variable *united* is assigned a code of 1, otherwise 0. This is a **multiple dichotomy method** of mapping variables. The other way to map responses is the **multiple category method**, in which you estimate the maximum number of possible responses to the question and set up the same number of variables, with codes used to specify the airline flown. By perusing a sample of questionnaires, you might discover that no user has flown more than three different airlines on this route in the last six months. Further, you find that due to the deregulation of airlines, 10 other airlines are named in the Other category. Using the multiple response method, you would define three variables, each coded as 1 = *american*, 2 = *united*, 3 = *twa*, 4 = *usair*, 5 = *delta*, and so on. If a given passenger circles American and TWA, the first variable has a code of 1, the second has a code of 3, and the third has a missing-value code. Another passenger might have circled American and entered Delta. Thus, the first variable has a code of 1, the second has a code of 5, and the third a missing-value code. If you use the multiple dichotomy method, on the other hand, you end up with 14 separate variables. Although either method of mapping is feasible for this survey, the method you choose depends on the distribution of responses.

Multiple Response Define Sets

The Define Multiple Response Sets procedure groups elementary variables into multiple dichotomy and multiple category sets, for which you can obtain frequency tables and crosstabulations. You can define up to 20 multiple response sets. Each set must have a unique

name. To remove a set, highlight it on the list of multiple response sets and click Remove. To change a set, highlight it on the list, modify any set definition characteristics, and click Change.

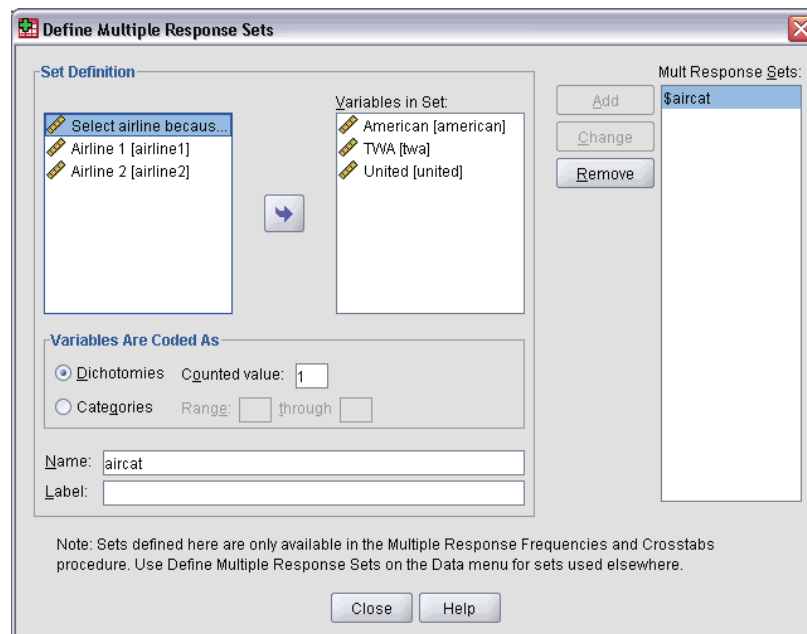
You can code your elementary variables as dichotomies or categories. To use dichotomous variables, select Dichotomies to create a multiple dichotomy set. Enter an integer value for Counted value. Each variable having at least one occurrence of the counted value becomes a category of the multiple dichotomy set. Select Categories to create a multiple category set having the same range of values as the component variables. Enter integer values for the minimum and maximum values of the range for categories of the multiple category set. The procedure totals each distinct integer value in the inclusive range across all component variables. Empty categories are not tabulated.

Each multiple response set must be assigned a unique name of up to seven characters. The procedure prefixes a dollar sign (\$) to the name you assign. You cannot use the following reserved names: *casenum*, *sysmis*, *jdate*, *date*, *time*, *length*, and *width*. The name of the multiple response set exists only for use in multiple response procedures. You cannot refer to multiple response set names in other procedures. Optionally, you can enter a descriptive variable label for the multiple response set. The label can be up to 40 characters long.

To Define Multiple Response Sets

- ▶ From the menus choose:
 - Analyze
 - Multiple Response
 - Define Sets...

Figure 40-1
Define Multiple Response Sets dialog box



- ▶ Select two or more variables.
- ▶ If your variables are coded as dichotomies, indicate which value you want to have counted. If your variables are coded as categories, define the range of the categories.

- ▶ Enter a unique name for each multiple response set.
- ▶ Click Add to add the multiple response set to the list of defined sets.

Multiple Response Frequencies

The Multiple Response Frequencies procedure produces frequency tables for multiple response sets. You must first define one or more multiple response sets (see “Multiple Response Define Sets”).

For multiple dichotomy sets, category names shown in the output come from variable labels defined for elementary variables in the group. If the variable labels are not defined, variable names are used as labels. For multiple category sets, category labels come from the value labels of the first variable in the group. If categories missing for the first variable are present for other variables in the group, define a value label for the missing categories.

Missing Values. Cases with missing values are excluded on a table-by-table basis. Alternatively, you can choose one or both of the following:

- **Exclude cases listwise within dichotomies.** Excludes cases with missing values for any variable from the tabulation of the multiple dichotomy set. This applies only to multiple response sets defined as dichotomy sets. By default, a case is considered missing for a multiple dichotomy set if none of its component variables contains the counted value. Cases with missing values for some (but not all variables) are included in the tabulations of the group if at least one variable contains the counted value.
- **Exclude cases listwise within categories.** Excludes cases with missing values for any variable from tabulation of the multiple category set. This applies only to multiple response sets defined as category sets. By default, a case is considered missing for a multiple category set only if none of its components has valid values within the defined range.

Example. Each variable created from a survey question is an elementary variable. To analyze a multiple response item, you must combine the variables into one of two types of multiple response sets: a multiple dichotomy set or a multiple category set. For example, if an airline survey asked which of three airlines (American, United, TWA) you have flown in the last six months and you used dichotomous variables and defined a **multiple dichotomy set**, each of the three variables in the set would become a category of the group variable. The counts and percentages for the three airlines are displayed in one frequency table. If you discover that no respondent mentioned more than two airlines, you could create two variables, each having three codes, one for each airline. If you define a **multiple category set**, the values are tabulated by adding the same codes in the elementary variables together. The resulting set of values is the same as those for each of the elementary variables. For example, 30 responses for United are the sum of the five United responses for airline 1 and the 25 United responses for airline 2. The counts and percentages for the three airlines are displayed in one frequency table.

Statistics. Frequency tables displaying counts, percentages of responses, percentages of cases, number of valid cases, and number of missing cases.

Data. Use multiple response sets.

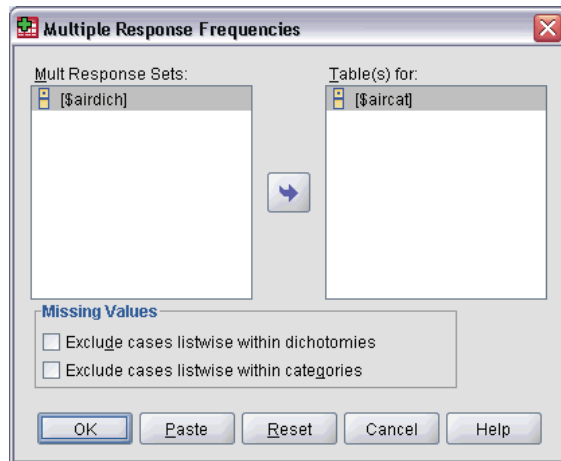
Assumptions. The counts and percentages provide a useful description for data from any distribution.

Related procedures. The Multiple Response Define Sets procedure allows you to define multiple response sets.

To Obtain Multiple Response Frequencies

- ▶ From the menus choose:
 - Analyze
 - Multiple Response
 - Frequencies...

Figure 40-2
Multiple Response Frequencies dialog box



- ▶ Select one or more multiple response sets.

Multiple Response Crosstabs

The Multiple Response Crosstabs procedure crosstabulates defined multiple response sets, elementary variables, or a combination. You can also obtain cell percentages based on cases or responses, modify the handling of missing values, or get paired crosstabulations. You must first define one or more multiple response sets (see “To Define Multiple Response Sets”).

For multiple dichotomy sets, category names shown in the output come from variable labels defined for elementary variables in the group. If the variable labels are not defined, variable names are used as labels. For multiple category sets, category labels come from the value labels of the first variable in the group. If categories missing for the first variable are present for other variables in the group, define a value label for the missing categories. The procedure displays category labels for columns on three lines, with up to eight characters per line. To avoid splitting words, you can reverse row and column items or redefine labels.

Example. Both multiple dichotomy and multiple category sets can be crosstabulated with other variables in this procedure. An airline passenger survey asks passengers for the following information: Circle all of the following airlines you have flown at least once in the last six months (American, United, TWA). Which is more important in selecting a flight—schedule or service? Select only one. After entering the data as dichotomies or multiple categories and combining them into a set, you can crosstabulate the airline choices with the question involving service or schedule.

Statistics. Crosstabulation with cell, row, column, and total counts, and cell, row, column, and total percentages. The cell percentages can be based on cases or responses.

Data. Use multiple response sets or numeric categorical variables.

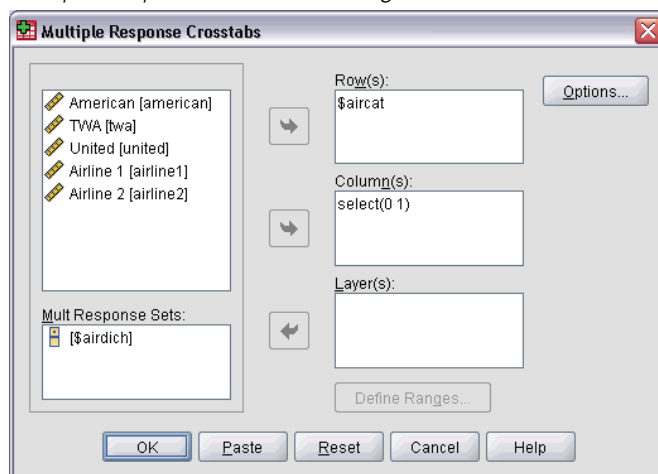
Assumptions. The counts and percentages provide a useful description of data from any distribution.

Related procedures. The Multiple Response Define Sets procedure allows you to define multiple response sets.

To Obtain Multiple Response Crosstabs

- ▶ From the menus choose:
 - Analyze
 - Multiple Response
 - Crosstabs...

Figure 40-3
Multiple Response Crosstabs dialog box




- ▶ Select one or more numeric variables or multiple response sets for each dimension of the crosstabulation.
- ▶ Define the range of each elementary variable.

Optionally, you can obtain a two-way crosstabulation for each category of a control variable or multiple response set. Select one or more items for the Layer(s) list.

Multiple Response Crosstabs Define Ranges

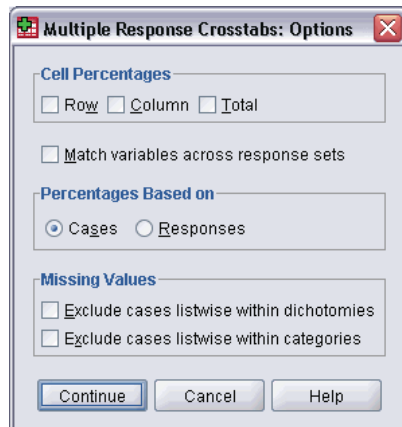
Figure 40-4
Multiple Response Crosstabs Define Variable Range dialog box



Value ranges must be defined for any elementary variable in the crosstabulation. Enter the integer minimum and maximum category values that you want to tabulate. Categories outside the range are excluded from analysis. Values within the inclusive range are assumed to be integers (non-integers are truncated).

Multiple Response Crosstabs Options

Figure 40-5
Multiple Response Crosstabs Options dialog box



Cell Percentages. Cell counts are always displayed. You can choose to display row percentages, column percentages, and two-way table (total) percentages.

Percentages Based on. You can base cell percentages on cases (or respondents). This is not available if you select matching of variables across multiple category sets. You can also base cell percentages on responses. For multiple dichotomy sets, the number of responses is equal to the number of counted values across cases. For multiple category sets, the number of responses is the number of values in the defined range.

Missing Values. You can choose one or both of the following:

- **Exclude cases listwise within dichotomies.** Excludes cases with missing values for any variable from the tabulation of the multiple dichotomy set. This applies only to multiple response sets defined as dichotomy sets. By default, a case is considered missing for a multiple dichotomy set if none of its component variables contains the counted value. Cases with missing values

for some, but not all, variables are included in the tabulations of the group if at least one variable contains the counted value.

- **Exclude cases listwise within categories.** Excludes cases with missing values for any variable from tabulation of the multiple category set. This applies only to multiple response sets defined as category sets. By default, a case is considered missing for a multiple category set only if none of its components has valid values within the defined range.

By default, when crosstabulating two multiple category sets, the procedure tabulates each variable in the first group with each variable in the second group and sums the counts for each cell; therefore, some responses can appear more than once in a table. You can choose the following option:

Match variables across response sets. Pairs the first variable in the first group with the first variable in the second group, and so on. If you select this option, the procedure bases cell percentages on responses rather than respondents. Pairing is not available for multiple dichotomy sets or elementary variables.

MULT RESPONSE Command Additional Features

The command syntax language also allows you to:

- Obtain crosstabulation tables with up to five dimensions (with the `BY` subcommand).
- Change output formatting options, including suppression of value labels (with the `FORMAT` subcommand).

See the *Command Syntax Reference* for complete syntax information.

Reporting Results

Case listings and descriptive statistics are basic tools for studying and presenting data. You can obtain case listings with the Data Editor or the Summarize procedure, frequency counts and descriptive statistics with the Frequencies procedure, and subpopulation statistics with the Means procedure. Each of these uses a format designed to make information clear. If you want to display the information in a different format, Report Summaries in Rows and Report Summaries in Columns give you the control you need over data presentation.

Report Summaries in Rows

Report Summaries in Rows produces reports in which different summary statistics are laid out in rows. Case listings are also available, with or without summary statistics.

Example. A company with a chain of retail stores keeps records of employee information, including salary, job tenure, and the store and division in which each employee works. You could generate a report that provides individual employee information (listing) broken down by store and division (break variables), with summary statistics (for example, mean salary) for each store, division, and division within each store.

Data Columns. Lists the report variables for which you want case listings or summary statistics and controls the display format of data columns.

Break Columns. Lists optional break variables that divide the report into groups and controls the summary statistics and display formats of break columns. For multiple break variables, there will be a separate group for each category of each break variable within categories of the preceding break variable in the list. Break variables should be discrete categorical variables that divide cases into a limited number of meaningful categories. Individual values of each break variable appear, sorted, in a separate column to the left of all data columns.

Report. Controls overall report characteristics, including overall summary statistics, display of missing values, page numbering, and titles.

Display cases. Displays the actual values (or value labels) of the data-column variables for every case. This produces a listing report, which can be much longer than a summary report.

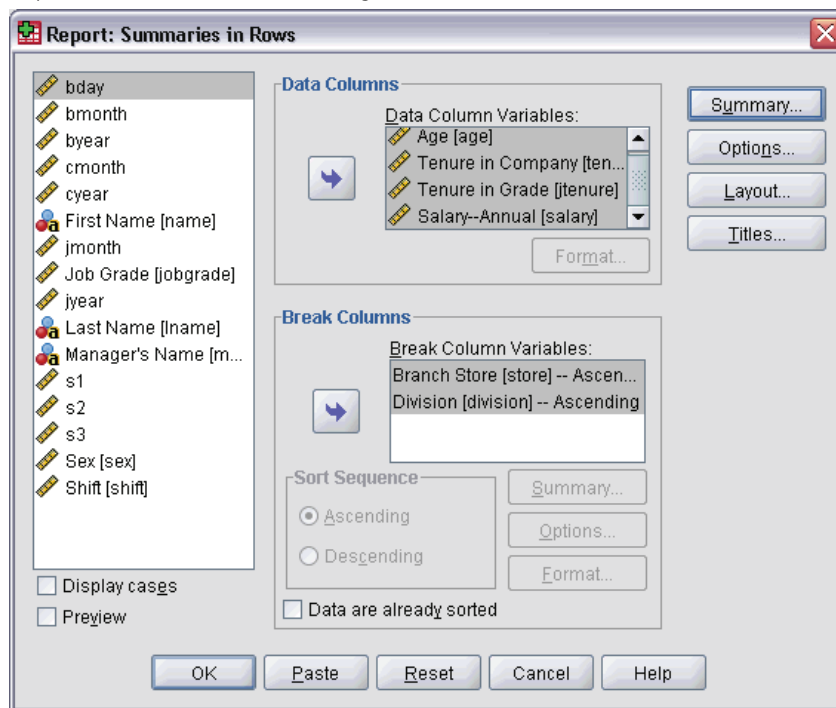
Preview. Displays only the first page of the report. This option is useful for previewing the format of your report without processing the whole report.

Data are already sorted. For reports with break variables, the data file must be sorted by break variable values before generating the report. If your data file is already sorted by values of the break variables, you can save processing time by selecting this option. This option is particularly useful after running a preview report.

To Obtain a Summary Report: Summaries in Rows

- ▶ From the menus choose:
Analyze
Reports
Report Summaries in Rows...
- ▶ Select one or more variables for Data Columns. One column in the report is generated for each variable selected.
- ▶ For reports sorted and displayed by subgroups, select one or more variables for Break Columns.
- ▶ For reports with summary statistics for subgroups defined by break variables, select the break variable in the Break Column Variables list and click Summary in the Break Columns group to specify the summary measure(s).
- ▶ For reports with overall summary statistics, click Summary to specify the summary measure(s).

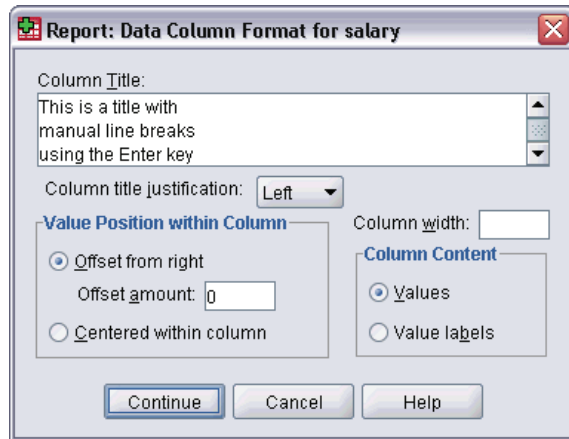
Figure 41-1
Report Summaries in Rows dialog box



Report Data Column/Break Format

The Format dialog boxes control column titles, column width, text alignment, and the display of data values or value labels. Data Column Format controls the format of data columns on the right side of the report page. Break Format controls the format of break columns on the left side.

Figure 41-2
Report Data Column Format dialog box



Column Title. For the selected variable, controls the column title. Long titles are automatically wrapped within the column. Use the Enter key to manually insert line breaks where you want titles to wrap.

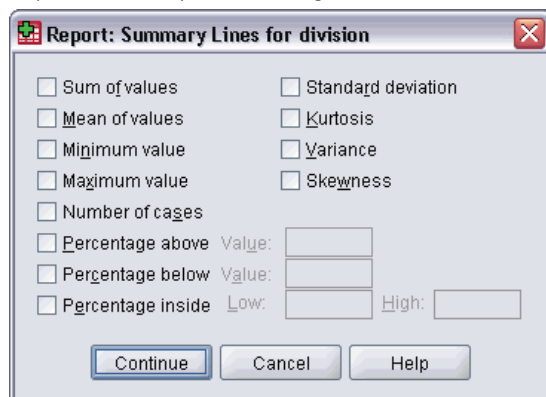
Value Position within Column. For the selected variable, controls the alignment of data values or value labels within the column. Alignment of values or labels does not affect alignment of column headings. You can either indent the column contents by a specified number of characters or center the contents.

Column Content. For the selected variable, controls the display of either data values or defined value labels. Data values are always displayed for any values that do not have defined value labels. (Not available for data columns in column summary reports.)

Report Summary Lines for/Final Summary Lines

The two Summary Lines dialog boxes control the display of summary statistics for break groups and for the entire report. Summary Lines controls subgroup statistics for each category defined by the break variable(s). Final Summary Lines controls overall statistics, displayed at the end of the report.

Figure 41-3
Report Summary Lines dialog box

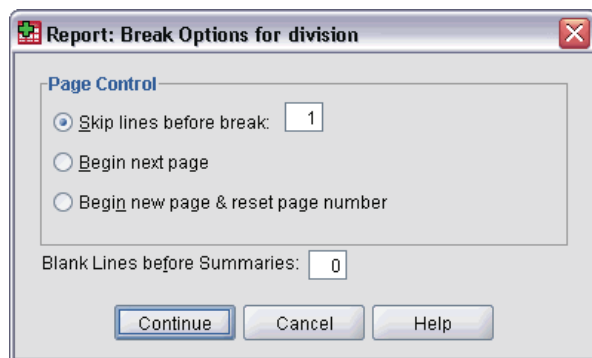


Available summary statistics are sum, mean, minimum, maximum, number of cases, percentage of cases above or below a specified value, percentage of cases within a specified range of values, standard deviation, kurtosis, variance, and skewness.

Report Break Options

Break Options controls spacing and pagination of break category information.

Figure 41-4
Report Break Options dialog box



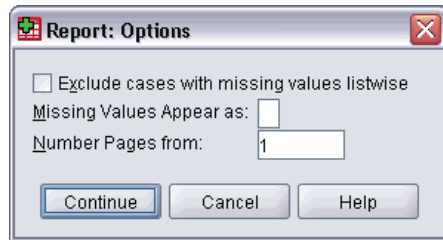
Page Control. Controls spacing and pagination for categories of the selected break variable. You can specify a number of blank lines between break categories or start each break category on a new page.

Blank Lines before Summaries. Controls the number of blank lines between break category labels or data and summary statistics. This is particularly useful for combined reports that include both individual case listings and summary statistics for break categories; in these reports, you can insert space between the case listings and the summary statistics.

Report Options

Report Options controls the treatment and display of missing values and report page numbering.

Figure 41-5
Report Options dialog box



Exclude cases with missing values listwise. Eliminates (from the report) any case with missing values for any of the report variables.

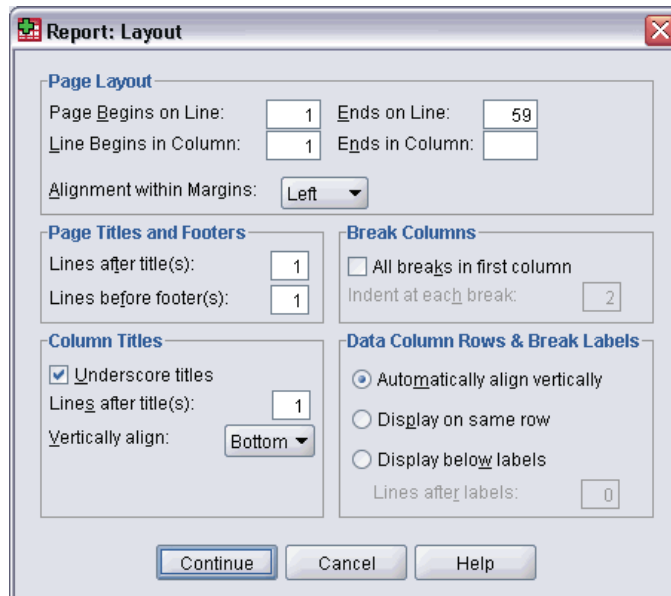
Missing Values Appear as. Allows you to specify the symbol that represents missing values in the data file. The symbol can be only one character and is used to represent both **system-missing** and **user-missing** values.

Number Pages from. Allows you to specify a page number for the first page of the report.

Report Layout

Report Layout controls the width and length of each report page, placement of the report on the page, and the insertion of blank lines and labels.

Figure 41-6
Report Layout dialog box



Page Layout. Controls the page margins expressed in lines (top and bottom) and characters (left and right) and reports alignment within the margins.

Page Titles and Footers. Controls the number of lines that separate page titles and footers from the body of the report.

Break Columns. Controls the display of break columns. If multiple break variables are specified, they can be in separate columns or in the first column. Placing all break variables in the first column produces a narrower report.

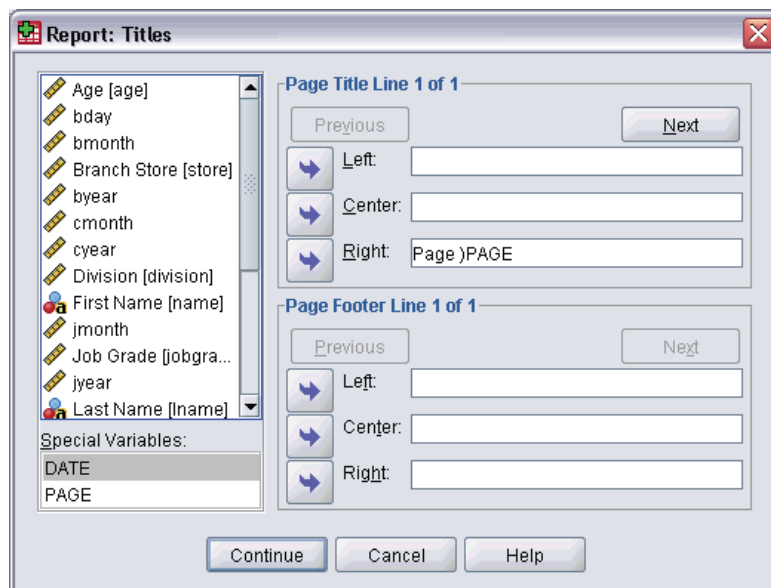
Column Titles. Controls the display of column titles, including title underlining, space between titles and the body of the report, and vertical alignment of column titles.

Data Column Rows and Break Labels. Controls the placement of data column information (data values and/or summary statistics) in relation to the break labels at the start of each break category. The first row of data column information can start either on the same line as the break category label or on a specified number of lines after the break category label. (Not available for column summary reports.)

Report Titles

Report Titles controls the content and placement of report titles and footers. You can specify up to 10 lines of page titles and up to 10 lines of page footers, with left-justified, centered, and right-justified components on each line.

Figure 41-7
Report Titles dialog box



If you insert variables into titles or footers, the current value label or value of the variable is displayed in the title or footer. In titles, the value label corresponding to the value of the variable at the beginning of the page is displayed. In footers, the value label corresponding to the value of the variable at the end of the page is displayed. If there is no value label, the actual value is displayed.

Special Variables. The special variables *DATE* and *PAGE* allow you to insert the current date or the page number into any line of a report header or footer. If your data file contains variables named *DATE* or *PAGE*, you cannot use these variables in report titles or footers.

Report Summaries in Columns

Report Summaries in Columns produces summary reports in which different summary statistics appear in separate columns.

Example. A company with a chain of retail stores keeps records of employee information, including salary, job tenure, and the division in which each employee works. You could generate a report that provides summary salary statistics (for example, mean, minimum, and maximum) for each division.

Data Columns. Lists the report variables for which you want summary statistics and controls the display format and summary statistics displayed for each variable.

Break Columns. Lists optional break variables that divide the report into groups and controls the display formats of break columns. For multiple break variables, there will be a separate group for each category of each break variable within categories of the preceding break variable in the list. Break variables should be discrete categorical variables that divide cases into a limited number of meaningful categories.

Report. Controls overall report characteristics, including display of missing values, page numbering, and titles.

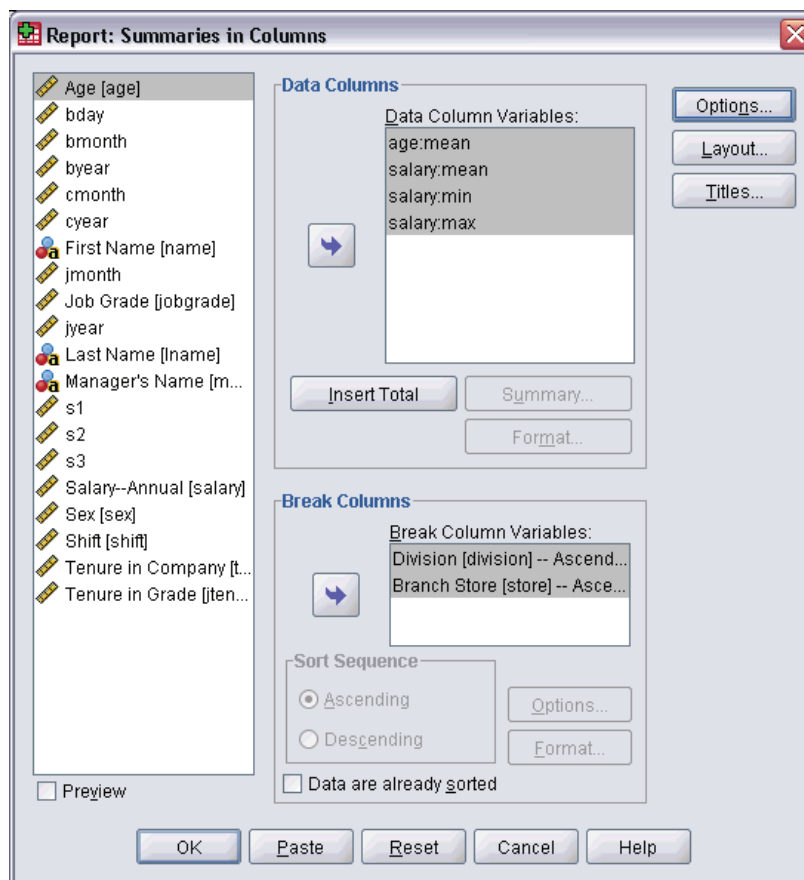
Preview. Displays only the first page of the report. This option is useful for previewing the format of your report without processing the whole report.

Data are already sorted. For reports with break variables, the data file must be sorted by break variable values before generating the report. If your data file is already sorted by values of the break variables, you can save processing time by selecting this option. This option is particularly useful after running a preview report.

To Obtain a Summary Report: Summaries in Columns

- ▶ From the menus choose:
 - Analyze
 - Reports
 - Report Summaries in Columns...
- ▶ Select one or more variables for Data Columns. One column in the report is generated for each variable selected.
- ▶ To change the summary measure for a variable, select the variable in the Data Column Variables list and click Summary.
- ▶ To obtain more than one summary measure for a variable, select the variable in the source list and move it into the Data Column Variables list multiple times, one for each summary measure you want.
- ▶ To display a column containing the sum, mean, ratio, or other function of existing columns, click Insert Total. This places a variable called *total* into the Data Columns list.
- ▶ For reports sorted and displayed by subgroups, select one or more variables for Break Columns.

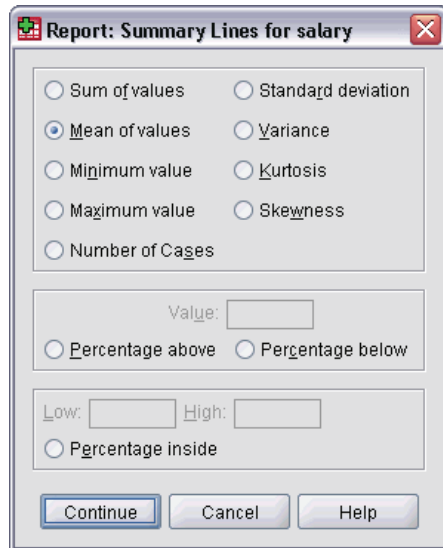
Figure 41-8
Report Summaries in Columns dialog box



Data Columns Summary Function

Summary Lines controls the summary statistic displayed for the selected data column variable.

Figure 41-9
Report Summary Lines dialog box



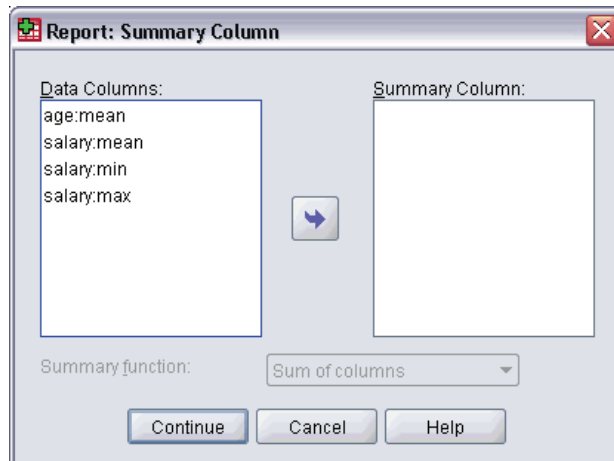
Available summary statistics are sum, mean, minimum, maximum, number of cases, percentage of cases above or below a specified value, percentage of cases within a specified range of values, standard deviation, variance, kurtosis, and skewness.

Data Columns Summary for Total Column

Summary Column controls the total summary statistics that summarize two or more data columns.

Available total summary statistics are sum of columns, mean of columns, minimum, maximum, difference between values in two columns, quotient of values in one column divided by values in another column, and product of columns values multiplied together.

Figure 41-10
Report Summary Column dialog box



Sum of columns. The *total* column is the sum of the columns in the Summary Column list.

Mean of columns. The *total* column is the average of the columns in the Summary Column list.

Minimum of columns. The *total* column is the minimum of the columns in the Summary Column list.

Maximum of columns. The *total* column is the maximum of the columns in the Summary Column list.

1st column – 2nd column. The *total* column is the difference of the columns in the Summary Column list. The Summary Column list must contain exactly two columns.

1st column / 2nd column. The *total* column is the quotient of the columns in the Summary Column list. The Summary Column list must contain exactly two columns.

% 1st column / 2nd column. The *total* column is the first column's percentage of the second column in the Summary Column list. The Summary Column list must contain exactly two columns.

Product of columns. The *total* column is the product of the columns in the Summary Column list.

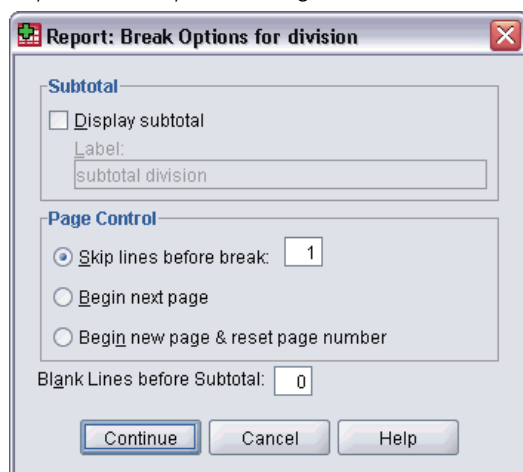
Report Column Format

Data and break column formatting options for Report Summaries in Columns are the same as those described for Report Summaries in Rows.

Report Summaries in Columns Break Options

Break Options controls subtotal display, spacing, and pagination for break categories.

Figure 41-11
Report Break Options dialog box



Subtotal. Controls the display subtotals for break categories.

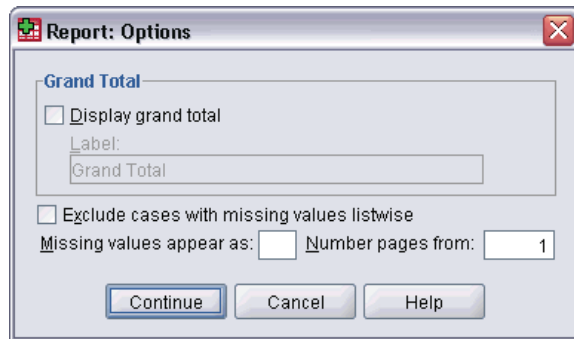
Page Control. Controls spacing and pagination for categories of the selected break variable. You can specify a number of blank lines between break categories or start each break category on a new page.

Blank Lines before Subtotal. Controls the number of blank lines between break category data and subtotals.

Report Summaries in Columns Options

Options controls the display of grand totals, the display of missing values, and pagination in column summary reports.

Figure 41-12
Report Options dialog box



Grand Total. Displays and labels a grand total for each column; displayed at the bottom of the column.

Missing values. You can exclude missing values from the report or select a single character to indicate missing values in the report.

Report Layout for Summaries in Columns

Report layout options for Report Summaries in Columns are the same as those described for Report Summaries in Rows.

REPORT Command Additional Features

The command syntax language also allows you to:

- Display different summary functions in the columns of a single summary line.
- Insert summary lines into data columns for variables other than the data column variable or for various combinations (composite functions) of summary functions.
- Use Median, Mode, Frequency, and Percent as summary functions.
- Control more precisely the display format of summary statistics.
- Insert blank lines at various points in reports.
- Insert blank lines after every *n*th case in listing reports.

Because of the complexity of the `REPORT` syntax, you may find it useful, when building a new report with syntax, to approximate the report generated from the dialog boxes, copy and paste the corresponding syntax, and refine that syntax to yield the exact report that you want.

See the *Command Syntax Reference* for complete syntax information.

Reliability Analysis

Reliability analysis allows you to study the properties of measurement scales and the items that compose the scales. The Reliability Analysis procedure calculates a number of commonly used measures of scale reliability and also provides information about the relationships between individual items in the scale. Intraclass correlation coefficients can be used to compute inter-rater reliability estimates.

Example. Does my questionnaire measure customer satisfaction in a useful way? Using reliability analysis, you can determine the extent to which the items in your questionnaire are related to each other, you can get an overall index of the repeatability or internal consistency of the scale as a whole, and you can identify problem items that should be excluded from the scale.

Statistics. Descriptives for each variable and for the scale, summary statistics across items, inter-item correlations and covariances, reliability estimates, ANOVA table, intraclass correlation coefficients, Hotelling's T^2 , and Tukey's test of additivity.

Models. The following models of reliability are available:

- **Alpha (Cronbach).** This model is a model of internal consistency, based on the average inter-item correlation.
- **Split-half.** This model splits the scale into two parts and examines the correlation between the parts.
- **Guttman.** This model computes Guttman's lower bounds for true reliability.
- **Parallel.** This model assumes that all items have equal variances and equal error variances across replications.
- **Strict parallel.** This model makes the assumptions of the Parallel model and also assumes equal means across items.

Data. Data can be dichotomous, ordinal, or interval, but the data should be coded numerically.

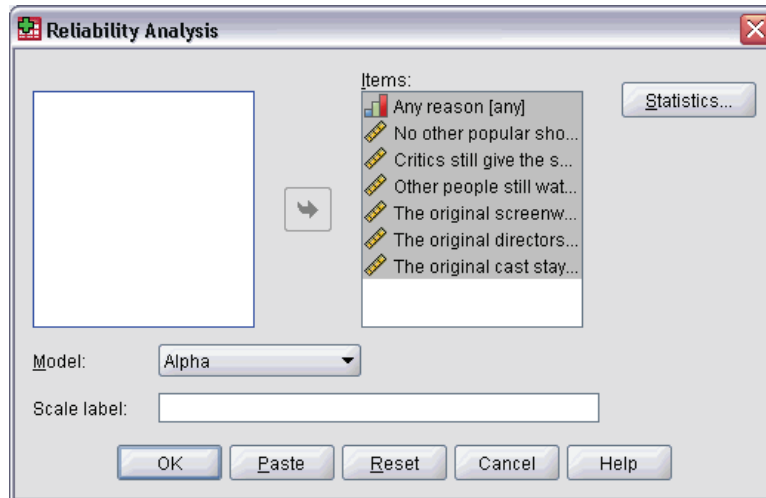
Assumptions. Observations should be independent, and errors should be uncorrelated between items. Each pair of items should have a bivariate normal distribution. Scales should be additive, so that each item is linearly related to the total score.

Related procedures. If you want to explore the dimensionality of your scale items (to see whether more than one construct is needed to account for the pattern of item scores), use factor analysis or multidimensional scaling. To identify homogeneous groups of variables, use hierarchical cluster analysis to cluster variables.

To Obtain a Reliability Analysis

- ▶ From the menus choose:
 - Analyze
 - Scale
 - Reliability Analysis...

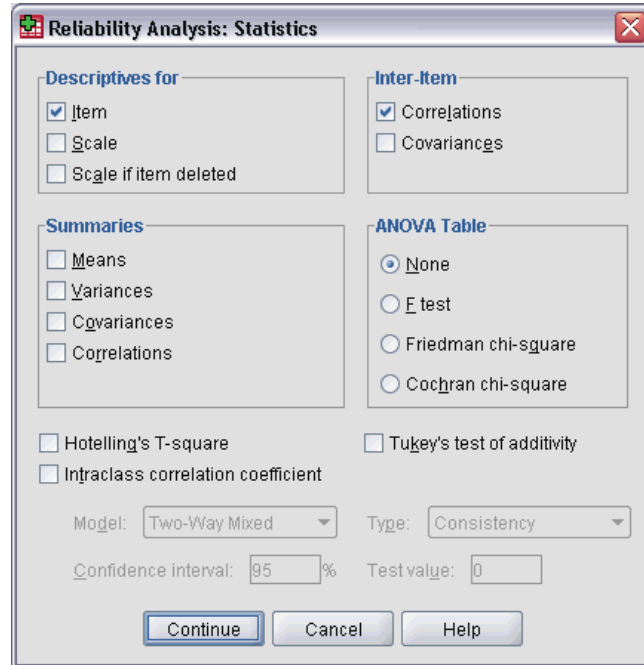
Figure 42-1
Reliability Analysis dialog box



- ▶ Select two or more variables as potential components of an additive scale.
- ▶ Choose a model from the Model drop-down list.

Reliability Analysis Statistics

Figure 42-2
Reliability Analysis Statistics dialog box



You can select various statistics that describe your scale and items. Statistics that are reported by default include the number of cases, the number of items, and reliability estimates as follows:

- **Alpha models.** Coefficient alpha; for dichotomous data, this is equivalent to the Kuder-Richardson 20 (KR20) coefficient.
- **Split-half models.** Correlation between forms, Guttman split-half reliability, Spearman-Brown reliability (equal and unequal length), and coefficient alpha for each half.
- **Guttman models.** Reliability coefficients lambda 1 through lambda 6.
- **Parallel and Strict parallel models.** Test for goodness of fit of model; estimates of error variance, common variance, and true variance; estimated common inter-item correlation; estimated reliability; and unbiased estimate of reliability.

Descriptives for. Produces descriptive statistics for scales or items across cases.

- **Item.** Produces descriptive statistics for items across cases.
- **Scale.** Produces descriptive statistics for scales.
- **Scale if item deleted.** Displays summary statistics comparing each item to the scale that is composed of the other items. Statistics include scale mean and variance if the item were to be deleted from the scale, correlation between the item and the scale that is composed of other items, and Cronbach's alpha if the item were to be deleted from the scale.

Summaries. Provides descriptive statistics of item distributions across all items in the scale.

- **Means.** Summary statistics for item means. The smallest, largest, and average item means, the range and variance of item means, and the ratio of the largest to the smallest item means are displayed.
- **Variiances.** Summary statistics for item variiances. The smallest, largest, and average item variiances, the range and variance of item variiances, and the ratio of the largest to the smallest item variiances are displayed.
- **Covariiances.** Summary statistics for inter-item covariiances. The smallest, largest, and average inter-item covariiances, the range and variance of inter-item covariiances, and the ratio of the largest to the smallest inter-item covariiances are displayed.
- **Correlations.** Summary statistics for inter-item correlations. The smallest, largest, and average inter-item correlations, the range and variance of inter-item correlations, and the ratio of the largest to the smallest inter-item correlations are displayed.

Inter-Item. Produces matrices of correlations or covariiances between items.

ANOVA Table. Produces tests of equal means.

- **F test.** Displays a repeated measures analysis-of-variance table.
- **Friedman chi-square.** Displays Friedman's chi-square and Kendall's coefficient of concordance. This option is appropriate for data that are in the form of ranks. The chi-square test replaces the usual F test in the ANOVA table.
- **Cochran chi-square.** Displays Cochran's Q. This option is appropriate for data that are dichotomous. The Q statistic replaces the usual F statistic in the ANOVA table.

Hotelling's T-square. Produces a multivariate test of the null hypothesis that all items on the scale have the same mean.

Tukey's test of additivity. Produces a test of the assumption that there is no multiplicative interaction among the items.

Intraclass correlation coefficient. Produces measures of consistency or agreement of values within cases.

- **Model.** Select the model for calculating the intraclass correlation coefficient. Available models are Two-Way Mixed, Two-Way Random, and One-Way Random. Select Two-Way Mixed when people effects are random and the item effects are fixed, select Two-Way Random when people effects and the item effects are random, or select One-Way Random when people effects are random.
- **Type.** Select the type of index. Available types are Consistency and Absolute Agreement.
- **Confidence interval.** Specify the level for the confidence interval. The default is 95%.
- **Test value.** Specify the hypothesized value of the coefficient for the hypothesis test. This value is the value to which the observed value is compared. The default value is 0.

RELIABILITY Command Additional Features

The command syntax language also allows you to:

- Read and analyze a correlation matrix.

- Write a correlation matrix for later analysis.
- Specify splits other than equal halves for the split-half method.

See the *Command Syntax Reference* for complete syntax information.

Multidimensional Scaling

Multidimensional scaling attempts to find the structure in a set of distance measures between objects or cases. This task is accomplished by assigning observations to specific locations in a conceptual space (usually two- or three-dimensional) such that the distances between points in the space match the given dissimilarities as closely as possible. In many cases, the dimensions of this conceptual space can be interpreted and used to further understand your data.

If you have objectively measured variables, you can use multidimensional scaling as a data reduction technique (the Multidimensional Scaling procedure will compute distances from multivariate data for you, if necessary). Multidimensional scaling can also be applied to subjective ratings of dissimilarity between objects or concepts. Additionally, the Multidimensional Scaling procedure can handle dissimilarity data from multiple sources, as you might have with multiple raters or questionnaire respondents.

Example. How do people perceive relationships between different cars? If you have data from respondents indicating similarity ratings between different makes and models of cars, multidimensional scaling can be used to identify dimensions that describe consumers' perceptions. You might find, for example, that the price and size of a vehicle define a two-dimensional space, which accounts for the similarities that are reported by your respondents.

Statistics. For each model: data matrix, optimally scaled data matrix, S-stress (Young's), stress (Kruskal's), RSQ, stimulus coordinates, average stress and RSQ for each stimulus (RMDS models). For individual difference (INDSCAL) models: subject weights and weirdness index for each subject. For each matrix in replicated multidimensional scaling models: stress and RSQ for each stimulus. Plots: stimulus coordinates (two- or three-dimensional), scatterplot of disparities versus distances.

Data. If your data are dissimilarity data, all dissimilarities should be quantitative and should be measured in the same metric. If your data are multivariate data, variables can be quantitative, binary, or count data. Scaling of variables is an important issue—differences in scaling may affect your solution. If your variables have large differences in scaling (for example, one variable is measured in dollars and the other variable is measured in years), consider standardizing them (this process can be done automatically by the Multidimensional Scaling procedure).

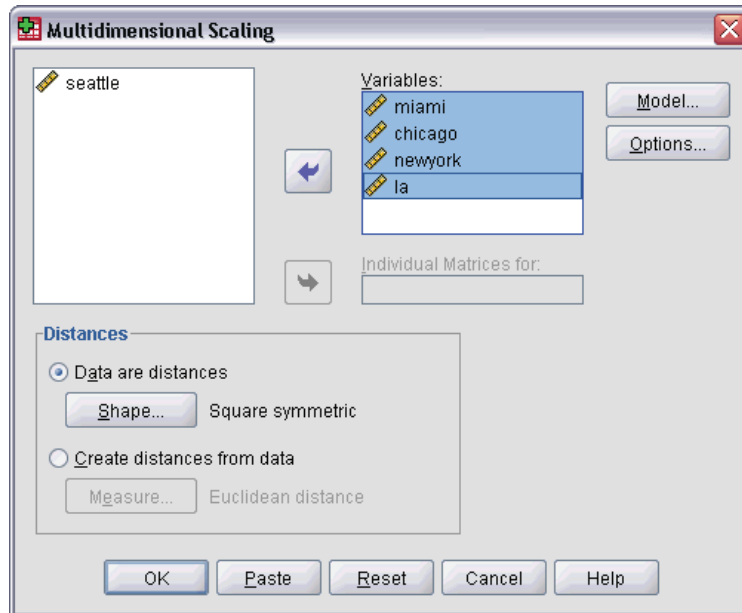
Assumptions. The Multidimensional Scaling procedure is relatively free of distributional assumptions. Be sure to select the appropriate measurement level (ordinal, interval, or ratio) in the Multidimensional Scaling Options dialog box so that the results are computed correctly.

Related procedures. If your goal is data reduction, an alternative method to consider is factor analysis, particularly if your variables are quantitative. If you want to identify groups of similar cases, consider supplementing your multidimensional scaling analysis with a hierarchical or *k*-means cluster analysis.

To Obtain a Multidimensional Scaling Analysis

- ▶ From the menus choose:
Analyze
Scale
Multidimensional Scaling...

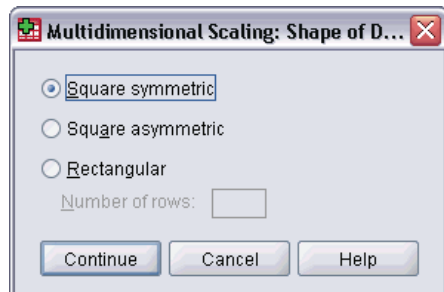
Figure 43-1
Multidimensional Scaling dialog box



- ▶ In Distances, select either Data are distances or Create distances from data.
- ▶ If your data are distances, select at least four numeric variables for analysis. (You can also click Shape to indicate the shape of the distance matrix.)
- ▶ If you want to create the distances before analyzing them, select at least one numeric variable. (You can also click Measure to specify the type of distance measure that you want.) Each grouping variable can be either numeric or string, and you can create separate matrices for each category of a grouping variable by moving that variable into the Individual Matrices For list.

Multidimensional Scaling Shape of Data

Figure 43-2
Multidimensional Scaling Shape dialog box

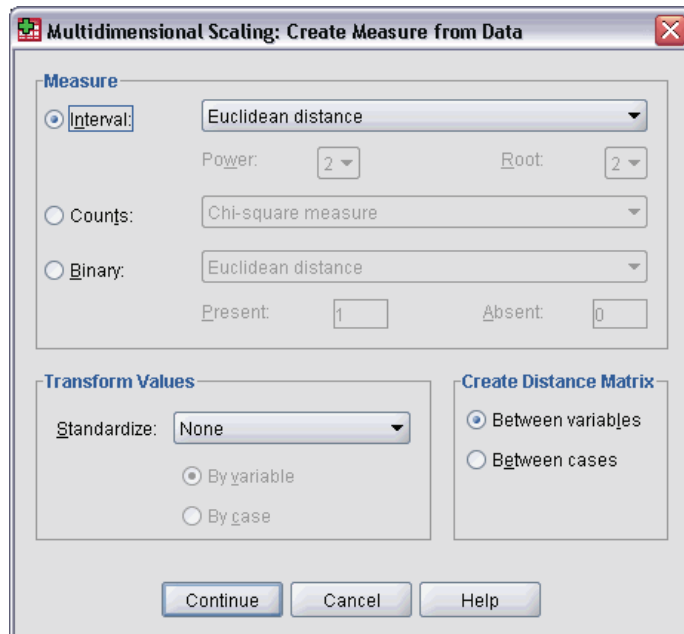


If your active dataset represents distances among a set of objects or represents distances between two sets of objects, specify the shape of your data matrix in order to get the correct results.

Note: You cannot select Square symmetric if the Model dialog box specifies row conditionality.

Multidimensional Scaling Create Measure

Figure 43-3
Multidimensional Scaling Create Measure from Data dialog box



Multidimensional scaling uses dissimilarity data to create a scaling solution. If your data are multivariate data (values of measured variables), you must create dissimilarity data in order to compute a multidimensional scaling solution. You can specify the details of creating dissimilarity measures from your data.

Measure. Allows you to specify the dissimilarity measure for your analysis. Select one alternative from the Measure group corresponding to your type of data, and then choose one of the measures from the drop-down list corresponding to that type of measure. Available alternatives are:

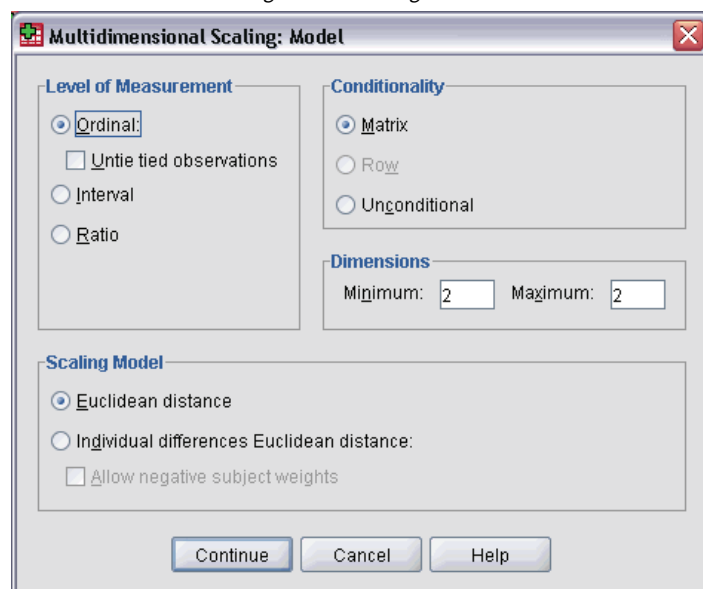
- **Interval.** Euclidean distance, Squared Euclidean distance, Chebychev, Block, Minkowski, or Customized.
- **Counts.** Chi-square measure or Phi-square measure.
- **Binary.** Euclidean distance, Squared Euclidean distance, Size difference, Pattern difference, Variance, or Lance and Williams.

Create Distance Matrix. Allows you to choose the unit of analysis. Alternatives are Between variables or Between cases.

Transform Values. In certain cases, such as when variables are measured on very different scales, you may want to standardize values before computing proximities (not applicable to binary data). Choose a standardization method from the Standardize drop-down list. If no standardization is required, choose None.

Multidimensional Scaling Model

Figure 43-4
Multidimensional Scaling Model dialog box



Correct estimation of a multidimensional scaling model depends on aspects of the data and the model itself.

Level of Measurement. Allows you to specify the level of your data. Alternatives are Ordinal, Interval, or Ratio. If your variables are ordinal, selecting Untie tied observations requests that the variables be treated as continuous variables, so that ties (equal values for different cases) are resolved optimally.

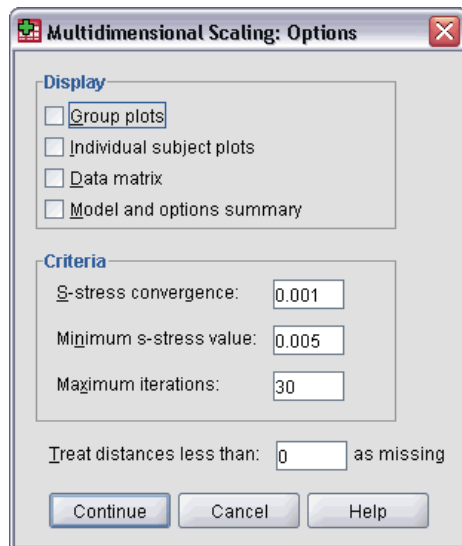
Conditionality. Allows you to specify which comparisons are meaningful. Alternatives are Matrix, Row, or Unconditional.

Dimensions. Allows you to specify the dimensionality of the scaling solution(s). One solution is calculated for each number in the range. Specify integers between 1 and 6; a minimum of 1 is allowed only if you select Euclidean distance as the scaling model. For a single solution, specify the same number for minimum and maximum.

Scaling Model. Allows you to specify the assumptions by which the scaling is performed. Available alternatives are Euclidean distance or Individual differences Euclidean distance (also known as INDSCAL). For the Individual differences Euclidean distance model, you can select Allow negative subject weights, if appropriate for your data.

Multidimensional Scaling Options

Figure 43-5
Multidimensional Scaling Options dialog box



You can specify options for your multidimensional scaling analysis.

Display. Allows you to select various types of output. Available options are Group plots, Individual subject plots, Data matrix, and Model and options summary.

Criteria. Allows you to determine when iteration should stop. To change the defaults, enter values for S-stress convergence, Minimum s-stress value, and Maximum iterations.

Treat distances less than n as missing. Distances that are less than this value are excluded from the analysis.

ALSCAL Command Additional Features

The command syntax language also allows you to:

- Use three additional model types, known as ASCAL, AINDS, and GEMSCAL in the literature about multidimensional scaling.
- Carry out polynomial transformations on interval and ratio data.
- Analyze similarities (rather than distances) with ordinal data.
- Analyze nominal data.
- Save various coordinate and weight matrices into files and read them back in for analysis.
- Constrain multidimensional unfolding.

See the *Command Syntax Reference* for complete syntax information.

Ratio Statistics

The Ratio Statistics procedure provides a comprehensive list of summary statistics for describing the ratio between two scale variables.

You can sort the output by values of a grouping variable in ascending or descending order. The ratio statistics report can be suppressed in the output, and the results can be saved to an external file.

Example. Is there good uniformity in the ratio between the appraisal price and sale price of homes in each of five counties? From the output, you might learn that the distribution of ratios varies considerably from county to county.

Statistics. Median, mean, weighted mean, confidence intervals, coefficient of dispersion (COD), median-centered coefficient of variation, mean-centered coefficient of variation, price-related differential (PRD), standard deviation, average absolute deviation (AAD), range, minimum and maximum values, and the concentration index computed for a user-specified range or percentage within the median ratio.

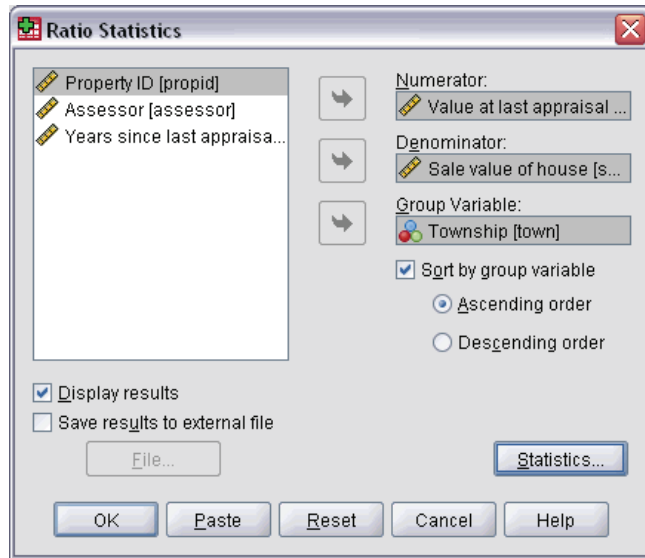
Data. Use numeric codes or strings to code grouping variables (nominal or ordinal level measurements).

Assumptions. The variables that define the numerator and denominator of the ratio should be scale variables that take positive values.

To Obtain Ratio Statistics

- ▶ From the menus choose:
 - Analyze
 - Descriptive Statistics
 - Ratio...

Figure 44-1
Ratio Statistics dialog box



- ▶ Select a numerator variable.
- ▶ Select a denominator variable.

Optionally:

- Select a grouping variable and specify the ordering of the groups in the results.
- Choose whether to display the results in the Viewer.
- Choose whether to save the results to an external file for later use, and specify the name of the file to which the results are saved.

Ratio Statistics

Figure 44-2
Ratio Statistics dialog box

Central Tendency. Measures of central tendency are statistics that describe the distribution of ratios.

- **Median.** The value such that the number of ratios that are less than this value and the number of ratios that are greater than this value are the same.
- **Mean.** The result of summing the ratios and dividing the result by the total number of ratios.
- **Weighted Mean.** The result of dividing the mean of the numerator by the mean of the denominator. Weighted mean is also the mean of the ratios weighted by the denominator.
- **Confidence Intervals.** Displays confidence intervals for the mean, the median, and the weighted mean (if requested). Specify a value that is greater than or equal to 0 and less than 100 as the confidence level.

Dispersion. These statistics measure the amount of variation, or spread, in the observed values.

- **AAD.** The average absolute deviation is the result of summing the absolute deviations of the ratios about the median and dividing the result by the total number of ratios.
- **COD.** The coefficient of dispersion is the result of expressing the average absolute deviation as a percentage of the median.
- **PRD.** The price-related differential, also known as the index of regressivity, is the result of dividing the mean by the weighted mean.
- **Median Centered COV.** The median-centered coefficient of variation is the result of expressing the root mean squares of deviation from the median as a percentage of the median.
- **Mean Centered COV.** The mean-centered coefficient of variation is the result of expressing the standard deviation as a percentage of the mean.

- **Standard deviation.** The standard deviation is the result of summing the squared deviations of the ratios about the mean, dividing the result by the total number of ratios minus one, and taking the positive square root.
- **Range.** The range is the result of subtracting the minimum ratio from the maximum ratio.
- **Minimum.** The minimum is the smallest ratio.
- **Maximum.** The maximum is the largest ratio.

Concentration Index. The coefficient of concentration measures the percentage of ratios that fall within an interval. It can be computed in two different ways:

- **Ratios Between.** Here the interval is defined explicitly by specifying the low and high values of the interval. Enter values for the low proportion and high proportion, and click Add to obtain an interval.
- **Ratios Within.** Here the interval is defined implicitly by specifying the percentage of the median. Enter a value between 0 and 100, and click Add. The lower end of the interval is equal to $(1 - 0.01 \times \text{value}) \times \text{median}$, and the upper end is equal to $(1 + 0.01 \times \text{value}) \times \text{median}$.

ROC Curves

This procedure is a useful way to evaluate the performance of classification schemes in which there is one variable with two categories by which subjects are classified.

Example. It is in a bank's interest to correctly classify customers into those customers who will and will not default on their loans, so special methods are developed for making these decisions. ROC curves can be used to evaluate how well these methods perform.

Statistics. Area under the ROC curve with confidence interval and coordinate points of the ROC curve. Plots: ROC curve.

Methods. The estimate of the area under the ROC curve can be computed either nonparametrically or parametrically using a binegative exponential model.

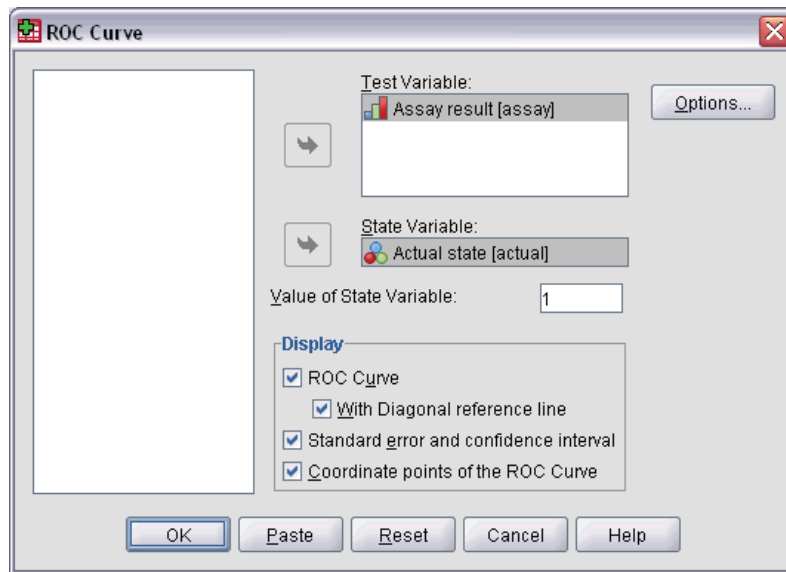
Data. Test variables are quantitative. Test variables are often composed of probabilities from discriminant analysis or logistic regression or composed of scores on an arbitrary scale indicating a rater's "strength of conviction" that a subject falls into one category or another category. The state variable can be of any type and indicates the true category to which a subject belongs. The value of the state variable indicates which category should be considered *positive*.

Assumptions. It is assumed that increasing numbers on the rater scale represent the increasing belief that the subject belongs to one category, while decreasing numbers on the scale represent the increasing belief that the subject belongs to the other category. The user must choose which direction is *positive*. It is also assumed that the *true* category to which each subject belongs is known.

To Obtain an ROC Curve

- ▶ From the menus choose:
 - Analyze
 - ROC Curve...

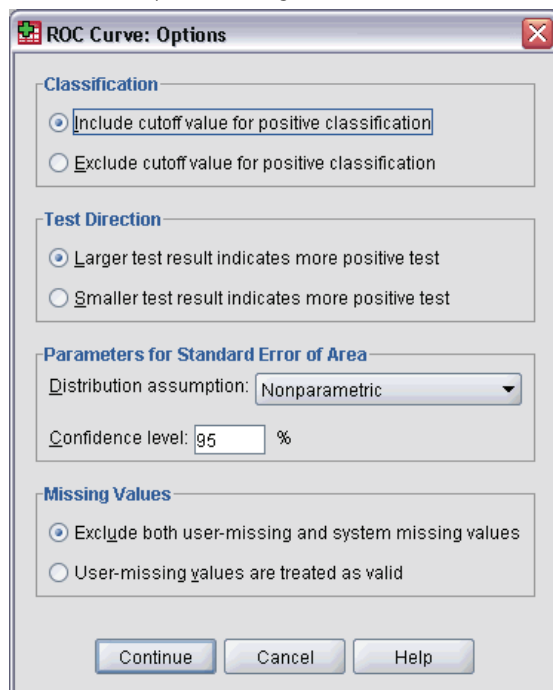
Figure 45-1
ROC Curve dialog box



- ▶ Select one or more test probability variables.
- ▶ Select one state variable.
- ▶ Identify the *positive* value for the state variable.

ROC Curve Options

Figure 45-2
ROC Curve Options dialog box



You can specify the following options for your ROC analysis:

Classification. Allows you to specify whether the cutoff value should be included or excluded when making a *positive* classification. This setting currently has no effect on the output.

Test Direction. Allows you to specify the direction of the scale in relation to the *positive* category.

Parameters for Standard Error of Area. Allows you to specify the method of estimating the standard error of the area under the curve. Available methods are nonparametric and bivariate exponential. Also allows you to set the level for the confidence interval. The available range is 50.1% to 99.9%.

Missing Values. Allows you to specify how missing values are handled.

Overview of the Chart Facility

High-resolution charts and plots are created by the procedures on the Graphs menu and by many of the procedures on the Analyze menu. This chapter provides an overview of the chart facility.

Building and Editing a Chart

Before you can create a chart, you need to have your data in the Data Editor. You can enter the data directly into the Data Editor; open a previously saved data file; or read a spreadsheet, tab-delimited data file, or database file. The Tutorial selection on the Help menu has online examples of creating and modifying a chart, and the online Help system provides information about creating and modifying all chart types.

Building Charts

The Chart Builder allows you to build charts from predefined gallery charts or from the individual parts (for example, axes and bars). You build a chart by dragging and dropping the gallery charts or basic elements onto the **canvas**, which is the large area to the right of the Variables list in the Chart Builder dialog box.

As you are building the chart, the canvas displays a preview of the chart. Although the preview uses defined variable labels and measurement levels, it does not display your actual data. Instead, it uses randomly generated data to provide a rough sketch of how the chart will look.

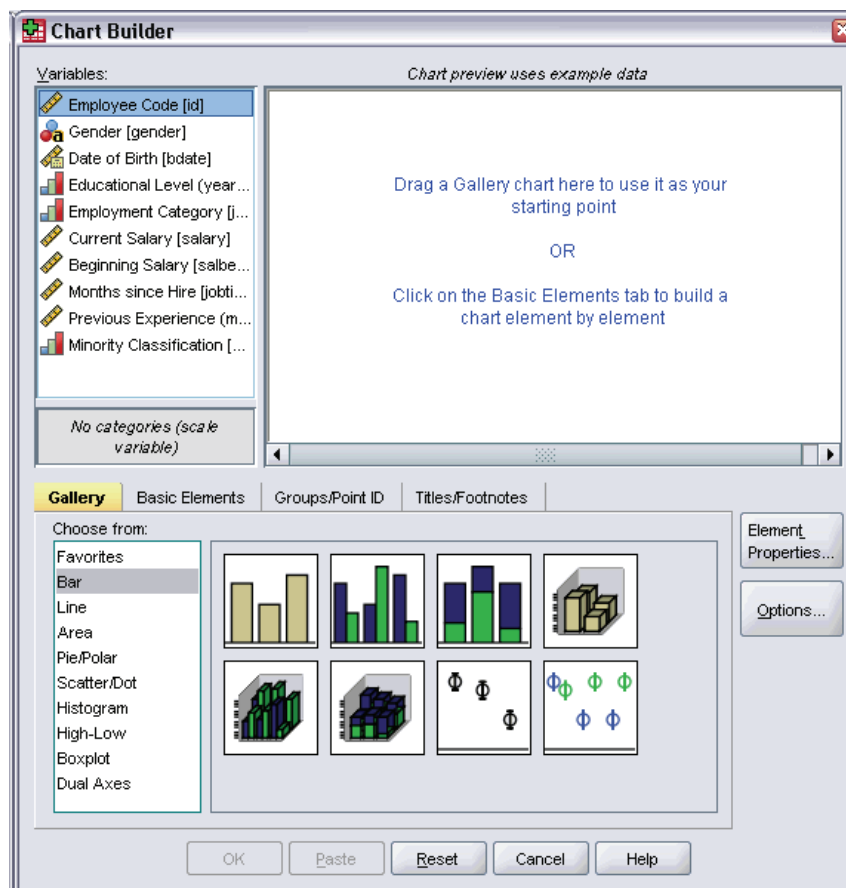
Using the gallery is the preferred method for new users. For information about using the gallery, see [Building a Chart from the Gallery on p. 479](#).

How to Start the Chart Builder

- ▶ From the menus choose:
 - Graphs
 - Chart Builder

This opens the Chart Builder dialog box.

Figure 46-1
Chart Builder dialog box



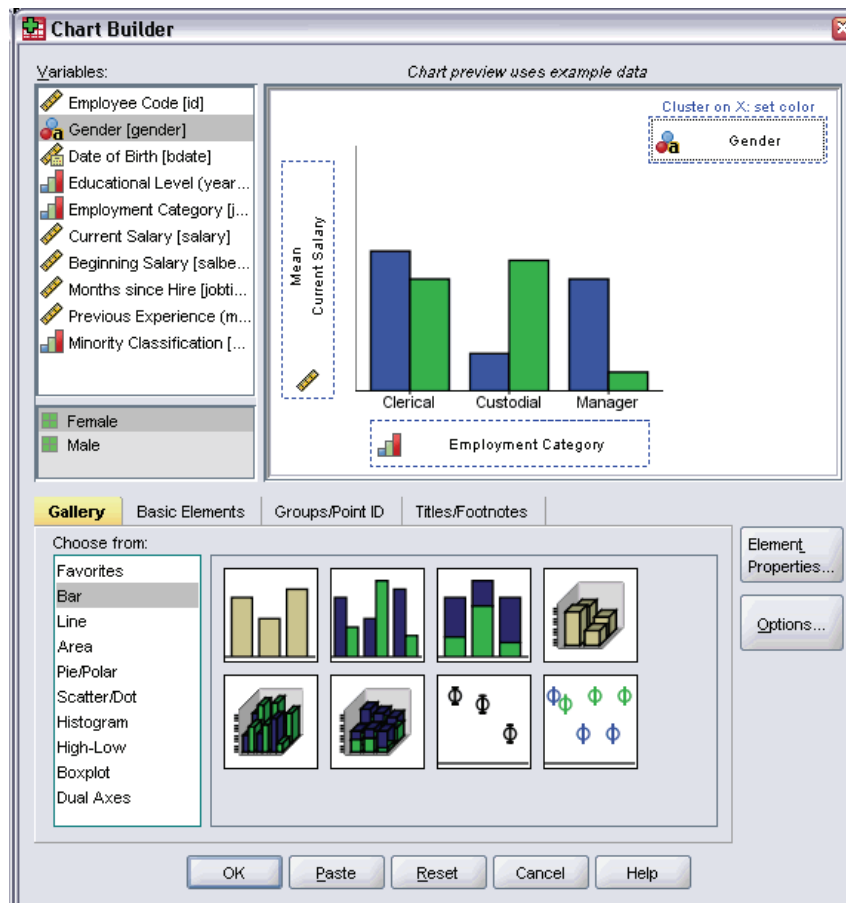
Building a Chart from the Gallery

The easiest method for building charts is to use the gallery. Following are general steps for building a chart from the gallery.

- ▶ Click the Gallery tab if it is not already displayed.
- ▶ In the Choose From list, select a category of charts. Each category offers several types.
- ▶ Drag the picture of the desired chart onto the canvas. You can also double-click the picture. If the canvas already displays a chart, the gallery chart replaces the axis set and graphic elements on the chart.
- ▶ Drag variables from the Variables list and drop them into the axis drop zones and, if available, the grouping drop zone. If an axis drop zone already displays a statistic and you want to use that statistic, you do not have to drag a variable into the drop zone. You need to add a variable to a zone only when the text in the zone is blue. If the text is black, the zone already contains a variable or statistic.

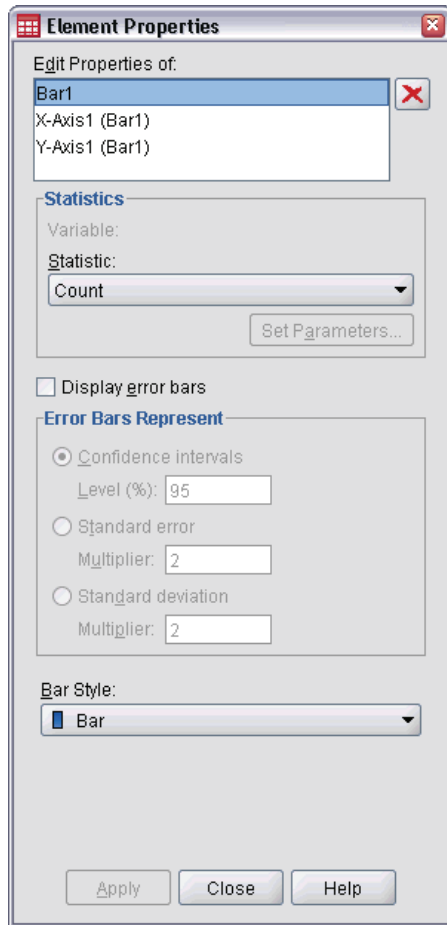
Note: The measurement level of your variables is important. The Chart Builder sets defaults based on the measurement level while you are building the chart. Furthermore, the resulting chart may also look different for different measurement levels. You can temporarily change a variable's measurement level by right-clicking the variable and choosing an option.

Figure 46-2
Chart Builder dialog box with completed drop zones



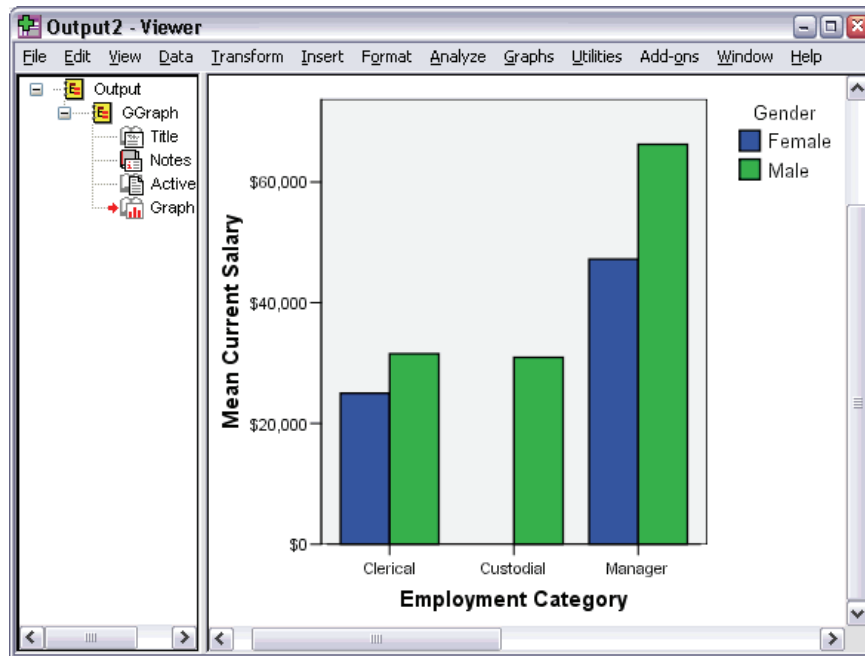
- If you need to change statistics or modify attributes of the axes or legends (such as the scale range), click Element Properties.

Figure 46-3
Element Properties dialog box



- ▶ In the Edit Properties Of list, select the item you want to change. (For information about the specific properties, click Help.)
- ▶ After making any changes, click Apply.
- ▶ If you need to add more variables to the chart (for example, for clustering or paneling), click the Groups/Point ID tab in the Chart Builder dialog box and select one or more options. Then drag categorical variables to the new drop zones that appear on the canvas.
- ▶ If you want to transpose the chart (for example, to make the bars horizontal), click the Basic Elements tab and then click Transpose.
- ▶ Click OK to create the chart. The chart is displayed in the Viewer.

Figure 46-4
Bar chart displayed in Viewer window



Editing Charts

The Chart Editor provides a powerful, easy-to-use environment where you can customize your charts and explore your data. The Chart Editor features:

- **Simple, intuitive user interface.** You can quickly select and edit parts of the chart using menus, context menus, and toolbars. You can also enter text directly on a chart.
- **Wide range of formatting and statistical options.** You can choose from a full range of styles and statistical options.
- **Powerful exploratory tools.** You can explore your data in various ways, such as by labeling, reordering, and rotating it. You can change chart types and the roles of variables in the chart. You can also add distribution curves and fit, interpolation, and reference lines.
- **Flexible templates for consistent look and behavior.** You can create customized templates and use them to easily create charts with the look and options that you want. For example, if you always want a specific orientation for axis labels, you can specify the orientation in a template and apply the template to other charts.

How to View the Chart Editor

- ▶ Create a chart in SPSS Statistics, or open a Viewer file with charts.
- ▶ Double-click a chart in the Viewer.

The chart is displayed in the Chart Editor.

Figure 46-5
Chart displayed in the Chart Editor




Chart Editor Fundamentals

The Chart Editor provides various methods for manipulating charts.

Menus

Many actions that you can perform in the Chart Editor are done with the menus, especially when you are adding an item to the chart. For example, you use the menus to add a fit line to a scatterplot. After adding an item to the chart, you often use the Properties dialog box to specify options for the added item.

Properties Dialog Box

Options for the chart and its chart elements can be found in the Properties dialog box.

To view the Properties dialog box, you can:

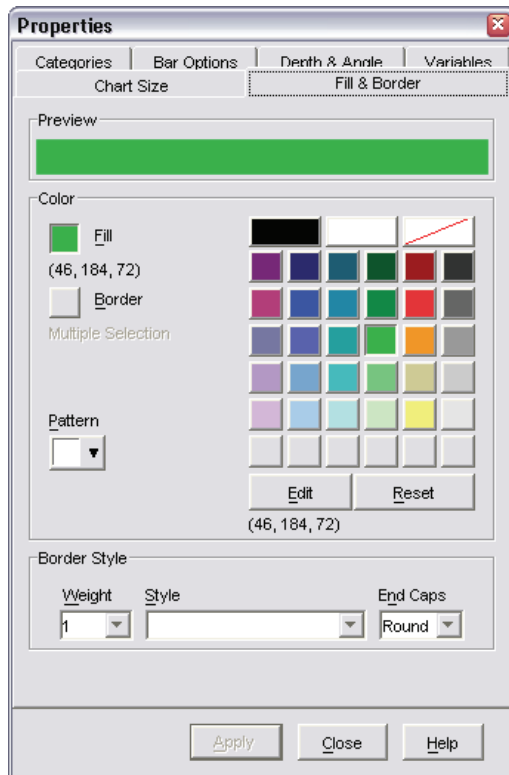
- ▶ Double-click a chart element.

or

- ▶ Select a chart element, and then from the menus choose:
Edit
Properties

Additionally, the Properties dialog box automatically appears when you add an item to the chart.

Figure 46-6
Properties dialog box, Fill & Border tab



The Properties dialog box has tabs that allow you to set the options and make other changes to a chart. The tabs that you see in the Properties dialog box are based on your current selection.

Some tabs include a preview to give you an idea of how the changes will affect the selection when you apply them. However, the chart itself does not reflect your changes until you click Apply. You can make changes on more than one tab before you click Apply. If you have to change the selection to modify a different element on the chart, click Apply before changing the selection. If you do not click Apply before changing the selection, clicking Apply at a later point will apply changes only to the element or elements currently selected.

Depending on your selection, only certain settings will be available. The help for the individual tabs specifies what you need to select to view the tabs. If multiple elements are selected, you can change only those settings that are common to all the elements.

Toolbars

The toolbars provide a shortcut for some of the functionality in the Properties dialog box. For example, instead of using the Text tab in the Properties dialog box, you can use the Edit toolbar to change the font and style of the text.

Saving the Changes

Chart modifications are saved when you close the Chart Editor. The modified chart is subsequently displayed in the Viewer.

Chart Definition Options

When you are defining a chart in the Chart Builder, you can add titles and change options for the chart creation.

Adding and Editing Titles and Footnotes

You can add titles and footnotes to the chart to help a viewer interpret it. The Chart Builder also automatically displays error bar information in the footnotes.

How to Add Titles and Footnotes

- ▶ Click the Titles/Footnotes tab.
- ▶ Select one or more titles and footnotes. The canvas displays some text to indicate that these were added to the chart.
- ▶ Use the Element Properties dialog box to edit the title/footnote text.

How to Remove a Title or Footnote

- ▶ Click the Titles/Footnotes tab.
- ▶ Deselect the title or footnote that you want to remove.

How to Edit the Title or Footnote Text

When you add titles and footnotes, you cannot edit their associated text directly on the chart. As with other items in the Chart Builder, you edit them using the Element Properties dialog box.

- ▶ Click Element Properties if the Element Properties dialog box is not displayed.
- ▶ In the Edit Properties Of list, select a title, subtitle, or footnote (for example, Title 1).
- ▶ In the content box, type the text associated with the title, subtitle, or footnote.
- ▶ Click Apply.

Setting General Options

The Chart Builder offers general options for the chart. These are options that apply to the overall chart, rather than a specific item on the chart. General options include missing value handling, templates, chart size, and panel wrapping.

- ▶ Click Options.

- ▶ Modify the general options. Details about these follow.
- ▶ Click Apply.

User-Missing Values

Break Variables. If there are missing values for the variables used to define categories or subgroups, select Include so that the category or categories of user-missing values (values identified as missing by the user) are included in the chart. These “missing” categories also act as break variables in calculating the statistic. The “missing” category or categories are displayed on the category axis or in the legend, adding, for example, an extra bar or a slice to a pie chart. If there are no missing values, the “missing” categories are not displayed.

If you select this option and want to suppress display after the chart is drawn, open the chart in the Chart Editor and choose Properties from the Edit menu. Use the Categories tab to move the categories that you want to suppress to the Excluded list. Note, however, that the statistics are not recalculated if you hide the “missing” categories. Therefore, something like a percent statistic will still take the “missing” categories into account.

Note: This control does not affect system-missing values. These are always excluded from the chart.

Summary Statistics and Case Values. You can choose one of the following alternatives for exclusion of cases having missing values:

- **Exclude listwise to obtain a consistent base for the chart.** If any of the variables in the chart has a missing value for a given case, the whole case is excluded from the chart.
- **Exclude variable-by-variable to maximize the use of the data.** If a selected variable has any missing values, the cases having those missing values are excluded when the variable is analyzed.

To see the difference between listwise and variable-by-variable exclusion of missing values, consider the following figures, which show a bar chart for each of the two options.

Figure 46-7
Listwise exclusion of missing values





Figure 46-8
Variable-by-variable exclusion of missing values



The data include some system-missing (blank) values in the variables for *Current Salary* and *Employment Category*. In some other cases, the value 0 was entered and defined as missing. For both charts, the option Display groups defined by missing values is selected, which adds the category *Missing* to the other displayed job categories. In each chart, the values of the summary function, *Number of cases*, are displayed in the bar labels.

In both charts, 26 cases have a system-missing value for the job category, and 13 cases have the user-missing value (0). In the listwise chart, the number of cases is the same for both variables in each bar cluster because whenever a value was missing, the case was excluded for all variables. In the variable-by-variable chart, the number of nonmissing cases for each variable in a category is plotted without regard to missing values in other variables.

Templates

A chart template allows you to apply the attributes of one chart to another. When you open a chart in the Chart Editor, you can save it as a template. You can then apply that template by specifying it at creation or later by applying it in the Chart Editor.

Default Template. This is the template specified by the Options. You can access these by choosing Options from the Edit menu in the Data Editor and then clicking the Charts tab. The default template is applied first, which means that the other templates can override it.

Template Files. Click Add to specify one or more templates with the standard file selection dialog box. These are applied in the order in which they appear. Therefore, templates at the end of the list can override templates at the beginning of the list.

Chart Size and Panels

Chart Size. Specify a percentage greater than 100 to enlarge the chart or less than 100 to shrink it. The percentage is relative to the default chart size.

Panels. When there are many panel columns, select Wrap Panels to allow panels to wrap across rows rather than being forced to fit in a specific row. Unless this option is selected, the panels are shrunk to force them to fit in a row.

Utilities

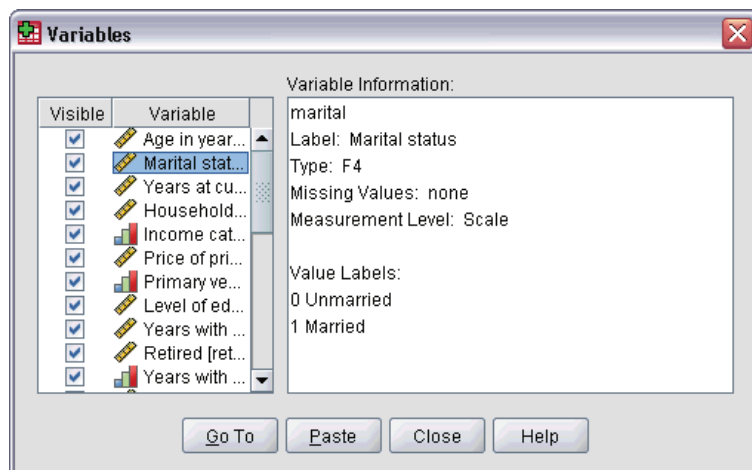
This chapter describes the functions found on the Utilities menu and the ability to reorder target variable lists.

Variable Information

The Variables dialog box displays variable definition information for the currently selected variable, including:

- Variable label
- Data format
- User-missing values
- Value labels
- Measurement level

Figure 47-1
Variables dialog box



Visible. The Visible column in the variable list indicates if the variable is currently visible in the Data Editor and in dialog box variable lists. Visibility is controlled by variable sets. [For more information, see Variable Sets on p. 489.](#)

Go To. Goes to the selected variable in the Data Editor window.

Paste. Pastes the selected variables into the designated syntax window at the cursor location.

To modify variable definitions, use the Variable view in the Data Editor.

To Obtain Variable Information

- ▶ From the menus choose:
 - Utilities
 - Variables...
- ▶ Select the variable for which you want to display variable definition information.

Data File Comments

You can include descriptive comments with a data file. For SPSS Statistics data files, these comments are saved with the data file.

To Add, Modify, Delete, or Display Data File Comments

- ▶ From the menus choose:
 - Utilities
 - Data File Comments...
- ▶ To display the comments in the Viewer, select Display comments in output.

Comments can be any length but are limited to 80 bytes (typically 80 characters in single-byte languages) per line; lines will automatically wrap at 80 characters. Comments are displayed in the same font as text output to accurately reflect how they will appear when displayed in the Viewer.

A date stamp (the current date in parentheses) is automatically appended to the end of the list of comments whenever you add or modify comments. This may lead to some ambiguity concerning the dates associated with comments if you modify an existing comment or insert a new comment between existing comments.

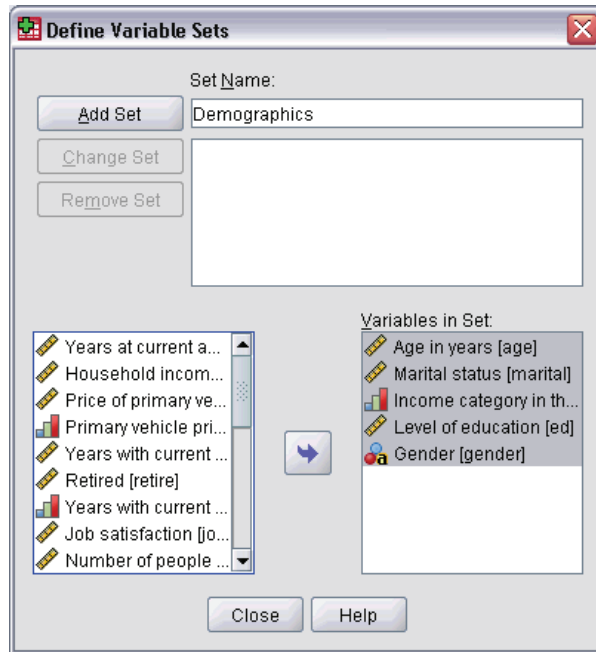
Variable Sets

You can restrict the variables that are displayed in the Data Editor and in dialog box variable lists by defining and using variable sets. This is particularly useful for data files with a large number of variables. Small variable sets make it easier to find and select the variables for your analysis.

Define Variable Sets

Define Variable Sets creates subsets of variables to display in the Data Editor and in dialog box variable lists. Defined variable sets are saved with SPSS Statistics data files.

Figure 47-2
Define Variable Sets dialog box



Set Name. Set names can be up to 64 bytes. Any characters, including blanks, can be used.

Variables in Set. Any combination of numeric and string variables can be included in a set. The order of variables in the set has no effect on the display order of the variables in the Data Editor or in dialog box variable lists. A variable can belong to multiple sets.

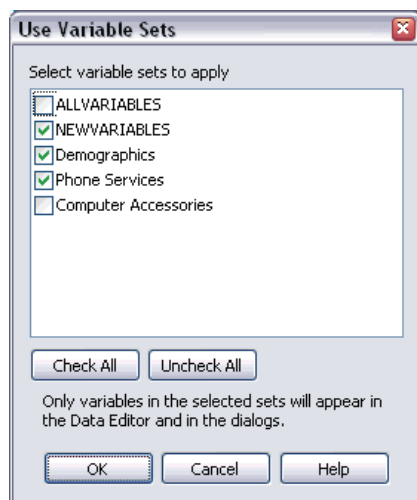
To Define Variable Sets

- ▶ From the menus choose:
Utilities
Define Variable Sets...
- ▶ Select the variables that you want to include in the set.
- ▶ Enter a name for the set (up to 64 bytes).
- ▶ Click Add Set.

Use Variable Sets

Use Variable Sets restricts the variables displayed in the Data Editor and in dialog box variable lists to the variables in the selected (checked) sets.

Figure 47-3
Use Variable Sets dialog box



- The set of variables displayed in the Data Editor and in dialog box variable lists is the union of all selected sets.
- A variable can be included in multiple selected sets.
- The order of variables in the selected sets and the order of selected sets have no effect on the display order of variables in the Data Editor or in dialog box variable lists.
- Although the defined variable sets are saved with SPSS Statistics data files, the list of currently selected sets is reset to the default, built-in sets each time you open the data file.

The list of available variable sets includes any variable sets defined for the active dataset, plus two built-in sets:

- **ALLVARIABLES.** This set contains all variables in the data file, including new variables created during a session.
- **NEWVARIABLES.** This set contains only new variables created during the session.

Note: Even if you save the data file after creating new variables, the new variables are still included in the *NEWVARIABLES* set until you close and reopen the data file.

At least one variable set must be selected. If *ALLVARIABLES* is selected, any other selected sets will not have any visible effect, since this set contains all variables.

To Select Variable Sets to Display

- ▶ From the menus choose:
Utilities
Use Variable Sets...
- ▶ Select the defined variable sets that contain the variables that you want to appear in the Data Editor and in dialog box variable lists.

To Display All Variables

- ▶ From the menus choose:
 - Utilities
 - Show All Variables

Reordering Target Variable Lists

Variables appear on dialog box target lists in the order in which they are selected from the source list. If you want to change the order of variables on a target list—but you don't want to deselect all of the variables and reselect them in the new order—you can move variables up and down on the target list using the Ctrl key (Macintosh: Command key) with the up and down arrow keys. You can move multiple variables simultaneously if they are contiguous (grouped together). You cannot move noncontiguous groups of variables.

Options

Options control a wide variety of settings, including:

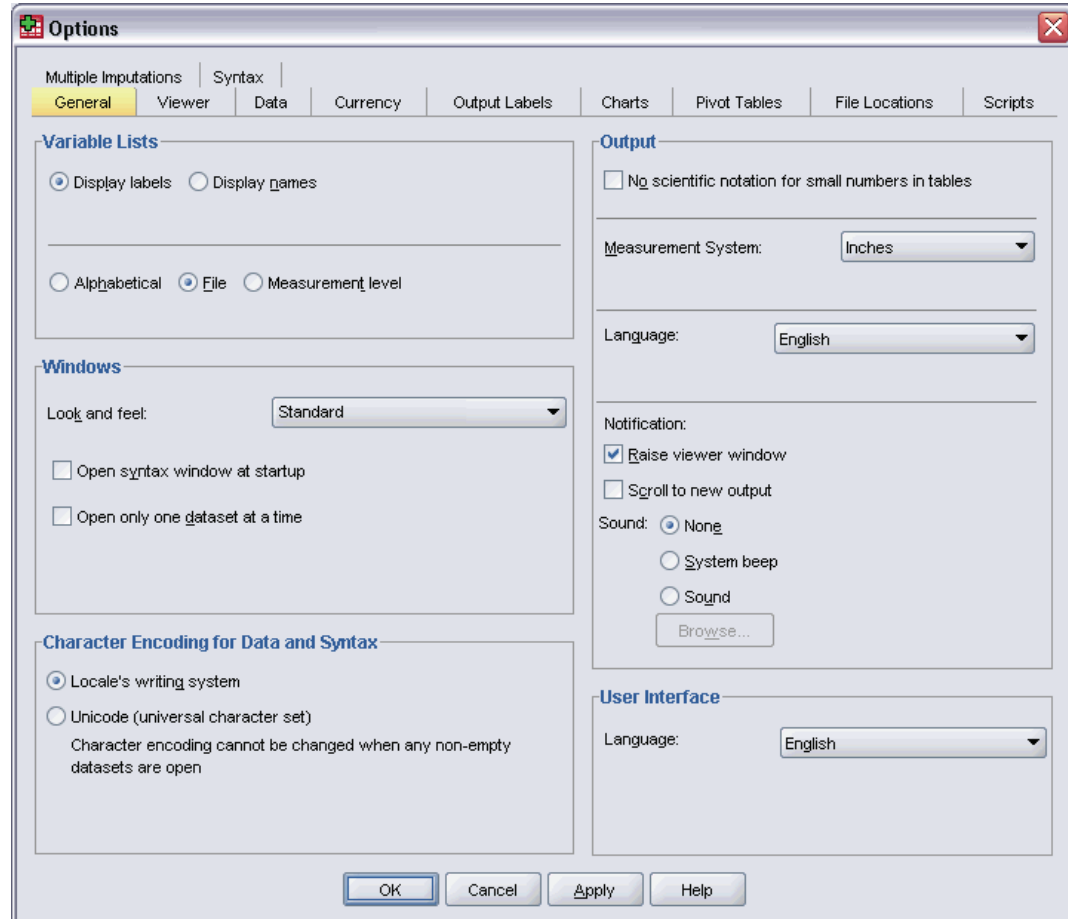
- Session journal, which keeps a record of all commands run in every session
- Display order for variables in dialog box source lists
- Items displayed and hidden in new output results
- TableLook for new pivot tables
- Custom currency formats

To Change Options Settings

- ▶ From the menus choose:
 - Edit
 - Options...
- ▶ Click the tabs for the settings that you want to change.
- ▶ Change the settings.
- ▶ Click OK or Apply.

General Options

Figure 48-1
Options dialog box: General tab



Variable Lists

These settings control the display of variables in dialog box lists. You can display variable names or variable labels. Names or labels can be displayed in alphabetical or file order or grouped by measurement level. Display order affects only source variable lists. Target variable lists always reflect the order in which variables were selected.

Windows

Look and feel. Controls the basic appearance of windows and dialog boxes.

Open syntax window at startup. Syntax windows are text file windows used to enter, edit, and run commands. If you frequently work with command syntax, select this option to automatically open a syntax window at the beginning of each session. This is useful primarily for experienced users who prefer to work with command syntax instead of dialog boxes.

Open only one dataset at a time. Closes the currently open data source each time you open a different data source using the menus and dialog boxes. By default, every time you use the menus and dialog boxes to open a new data source, that data source opens in a new Data Editor window, and any other data sources open in other Data Editor windows remain open and available during the session until explicitly closed.

When you select this option, it takes effect immediately but does not close any datasets that were open at the time the setting was changed. This setting has no effect on data sources opened using command syntax, which relies on `DATASET` commands to control multiple datasets. [For more information, see Working with Multiple Data Sources in Chapter 6 on p. 93.](#)

Character Encoding for Data Files and Syntax Files

This controls the default behavior for determining the encoding for reading and writing data files and syntax files. You can change this setting only when there are no open data sources, and the setting remains in effect for subsequent sessions until explicitly changed.

- **Locale's writing system.** Use the current locale setting to determine the encoding for reading and writing files. This is referred to as **code page mode**.
- **Unicode (universal character set).** Use Unicode encoding (UTF-8) for reading and writing files. This is referred to as **Unicode mode**.

There are a number of important implications regarding Unicode mode and Unicode files:

- SPSS Statistics data files and syntax files saved in Unicode encoding should not be used in releases of SPSS Statistics prior to 16.0. For syntax files, you can specify the encoding when you save the file. For data files, you should open the data file in code page mode and then resave it if you want to read the file with earlier versions.
- When code page data files are read in Unicode mode, the defined width of all string variables is tripled. To automatically set the width of each string variable to the longest observed value for that variable, select Minimize string widths based on observed values in the Open Data dialog box.

Output

No scientific notation for small numbers in tables. Suppresses the display of scientific notation for small decimal values in output. Very small decimal values will be displayed as 0 (or 0.000).

Measurement System. The measurement system used (points, inches, or centimeters) for specifying attributes such as pivot table cell margins, cell widths, and space between tables for printing.

Language. Controls the language used in the output. Does not apply to simple text output. The list of available languages depends on the currently installed language files. (Note: This does *not* affect the user interface language.)

Note: Custom scripts that rely on language-specific text strings in the output may not run correctly when you change the output language. [For more information, see Script Options on p. 507.](#)

Notification. Controls the manner in which the program notifies you that it has finished running a procedure and that the results are available in the Viewer.

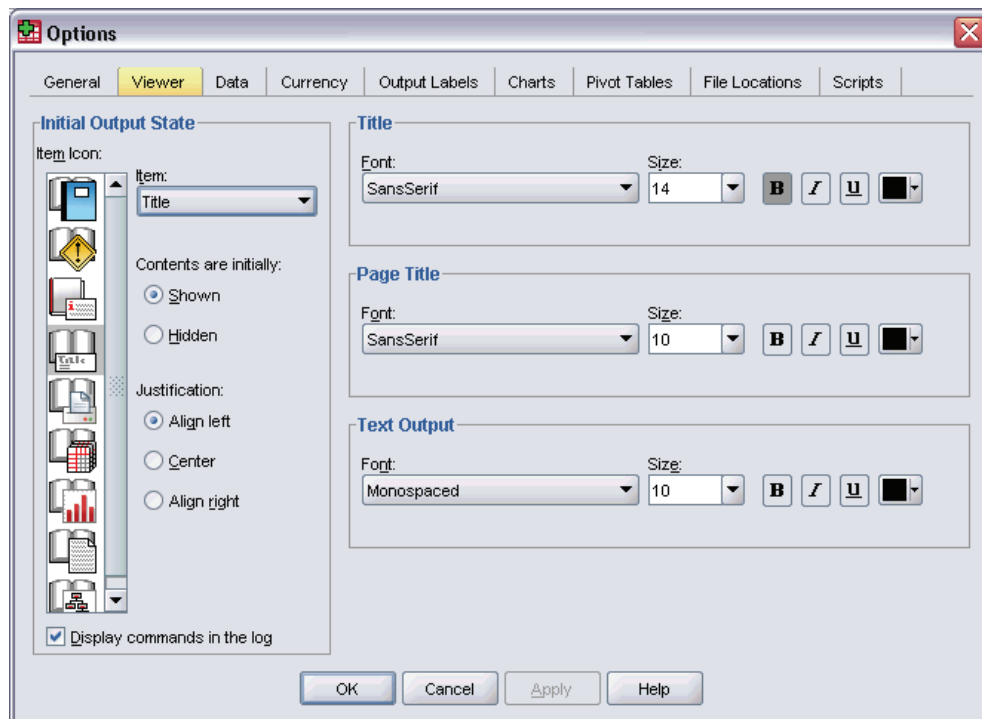
User Interface

This controls the language used in menus, dialogs, and other user interface features. (Note: This does *not* affect the output language.)

Viewer Options

Viewer output display options affect only new output produced after you change the settings. Output already displayed in the Viewer is not affected by changes in these settings.

Figure 48-2
Options dialog box: Viewer tab



Initial Output State. Controls which items are automatically displayed or hidden each time you run a procedure and how items are initially aligned. You can control the display of the following items: log, warnings, notes, titles, pivot tables, charts, tree diagrams, and text output. You can also turn the display of commands in the log on or off. You can copy command syntax from the log and save it in a syntax file.

Note: All output items are displayed left-aligned in the Viewer. Only the alignment of printed output is affected by the justification settings. Centered and right-aligned items are identified by a small symbol above and to the left of the item.

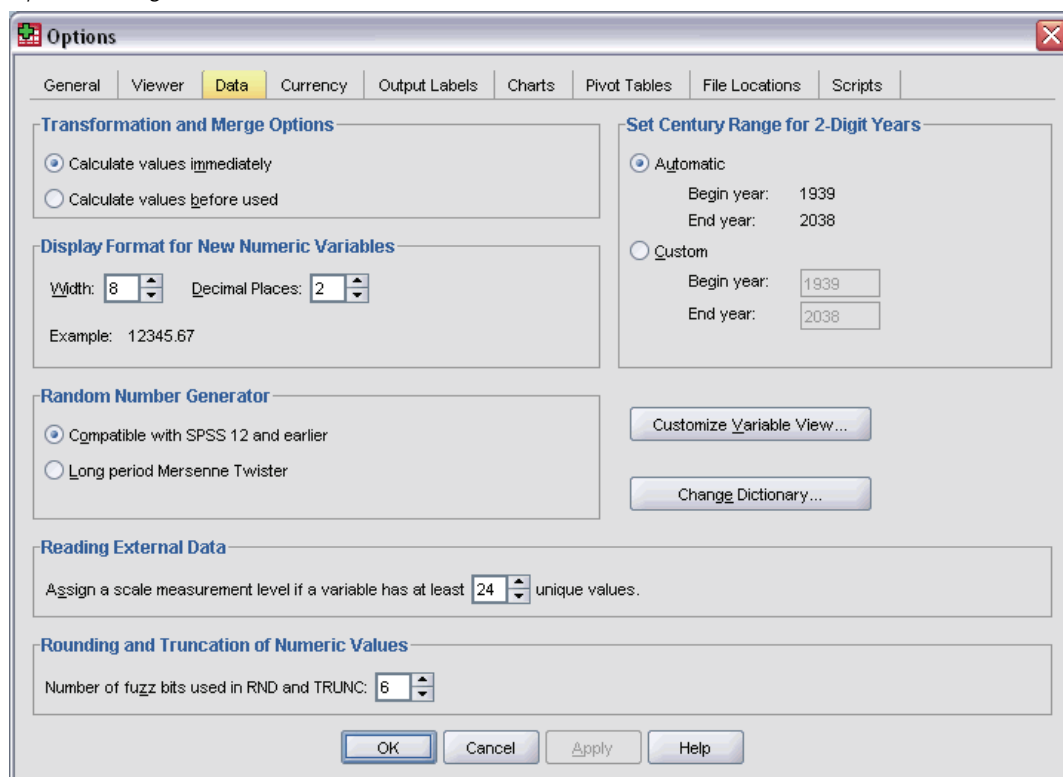
Title. Controls the font style, size, and color for new output titles.

Page Title. Controls the font style, size, and color for new page titles and page titles generated by TITLE and SUBTITLE command syntax or created by New Page Title on the Insert menu.

Text Output Font used for text output. Text output is designed for use with a monospaced (fixed-pitch) font. If you select a proportional font, tabular output will not align properly.

Data Options

Figure 48-3
Options dialog box: Data tab



Transformation and Merge Options. Each time the program executes a command, it reads the data file. Some data transformations (such as Compute and Recode) and file transformations (such as Add Variables and Add Cases) do not require a separate pass of the data, and execution of these commands can be delayed until the program reads the data to execute another command, such as a statistical or charting procedure.

- For large data files, where reading the data can take some time, you may want to select Calculate values before used to delay execution and save processing time. When this option is selected, the results of transformations you make using dialog boxes such as Compute Variable will not appear immediately in the Data Editor; new variables created by transformations will be displayed without any data values; and data values in the Data Editor cannot be changed while there are pending transformations. Any command that reads the data, such as a statistical or charting procedure, will execute the pending transformations and update the data displayed in the Data Editor. Alternatively, you can use Run Pending Transforms on the Transform menu.
- With the default setting of Calculate values immediately, when you paste command syntax from dialogs, an EXECUTE command is pasted after each transformation command. [For more information, see Multiple Execute Commands in Chapter 13 on p. 265.](#)

Display Format for New Numeric Variables. Controls the default display width and number of decimal places for new numeric variables. There is no default display format for new string variables. If a value is too large for the specified display format, first decimal places are rounded and then values are converted to scientific notation. Display formats do not affect internal data values. For example, the value 123456.78 may be rounded to 123457 for display, but the original unrounded value is used in any calculations.

Set Century Range for 2-Digit Years. Defines the range of years for date-format variables entered and/or displayed with a two-digit year (for example, 10/28/86, 29-OCT-87). The automatic range setting is based on the current year, beginning 69 years prior to and ending 30 years after the current year (adding the current year makes a total range of 100 years). For a custom range, the ending year is automatically determined based on the value that you enter for the beginning year.

Random Number Generator. Two different random number generators are available:

- **Version 12 Compatible.** The random number generator used in version 12 and previous releases. If you need to reproduce randomized results generated in previous releases based on a specified seed value, use this random number generator.
- **Mersenne Twister.** A newer random number generator that is more reliable for simulation purposes. If reproducing randomized results from version 12 or earlier is not an issue, use this random number generator.

Reading External Data. For data read from external file formats and older SPSS Statistics data files (prior to release 8.0), you can specify the minimum number of data values for a numeric variable used to classify the variable as scale or nominal. Variables with fewer than the specified number of unique values are classified as nominal.

Rounding and Truncation of Numeric Values. For the `RND` and `TRUNC` functions, this setting controls the default threshold for rounding up values that are very close to a rounding boundary. The setting is specified as a number of bits and is set to 6 at install time, which should be sufficient for most applications. Setting the number of bits to 0 produces the same results as in release 10. Setting the number of bits to 10 produces the same results as in releases 11 and 12.

- For the `RND` function, this setting specifies the number of least-significant bits by which the value to be rounded may fall short of the threshold for rounding up but still be rounded up. For example, when rounding a value between 1.0 and 2.0 to the nearest integer this setting specifies how much the value can fall short of 1.5 (the threshold for rounding up to 2.0) and still be rounded up to 2.0.
- For the `TRUNC` function, this setting specifies the number of least-significant bits by which the value to be truncated may fall short of the nearest rounding boundary and be rounded up before truncating. For example, when truncating a value between 1.0 and 2.0 to the nearest integer this setting specifies how much the value can fall short of 2.0 and be rounded up to 2.0.

Customize Variable View. Controls the default display and order of attributes in Variable View. [For more information, see Changing the Default Variable View on p. 499.](#)

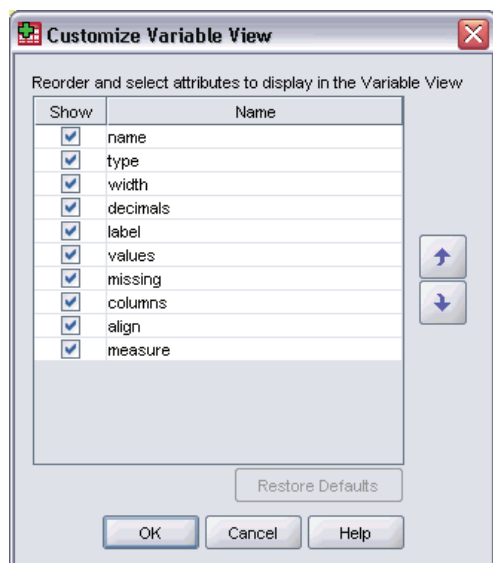
Change Dictionary. Controls the language version of the dictionary used for checking the spelling of items in Variable View. [For more information, see Spell Checking in Chapter 5 on p. 83.](#)

Changing the Default Variable View

You can use Customize Variable View to control which attributes are displayed by default in Variable View (for example, name, type, label) and the order in which they are displayed.

Click Customize Variable View.

Figure 48-4
Customize Variable View (default)



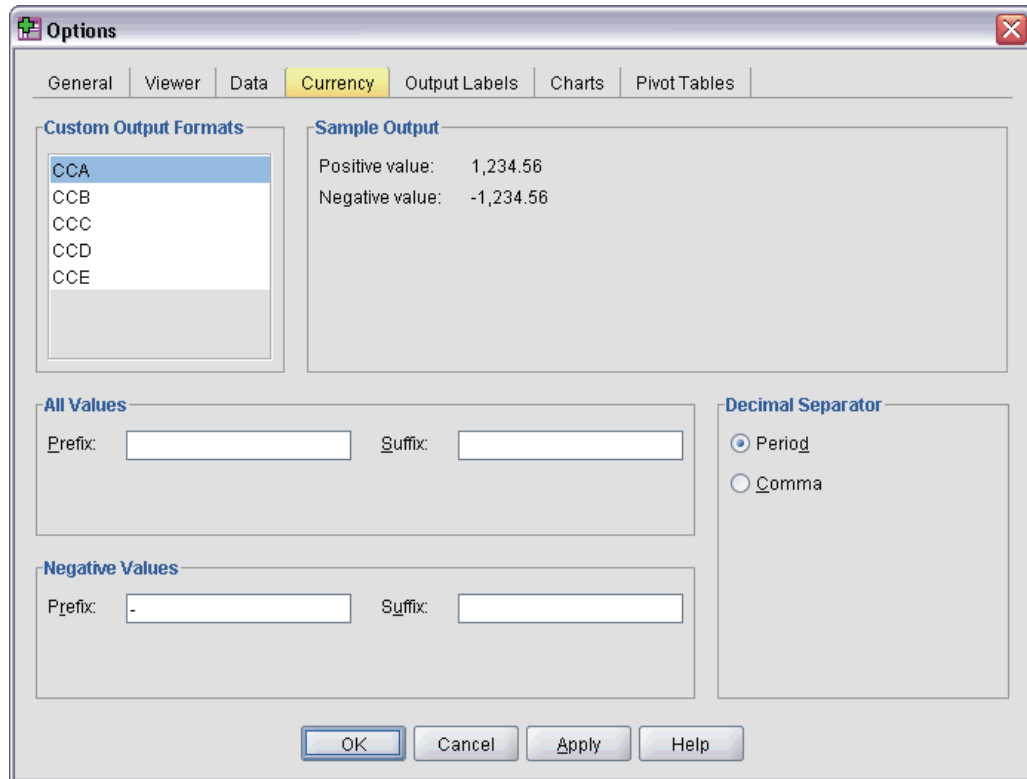
- ▶ Select (check) the variable attributes you want to display.
- ▶ Use the up and down arrow buttons to change the display order of the attributes.

Currency Options

You can create up to five custom currency display formats that can include special prefix and suffix characters and special treatment for negative values.

The five custom currency format names are CCA, CCB, CCC, CCD, and CCE. You cannot change the format names or add new ones. To modify a custom currency format, select the format name from the source list and make the changes that you want.

Figure 48-5
Options dialog box: Currency tab



Prefixes, suffixes, and decimal indicators defined for custom currency formats are for display purposes only. You cannot enter values in the Data Editor using custom currency characters.

To Create Custom Currency Formats

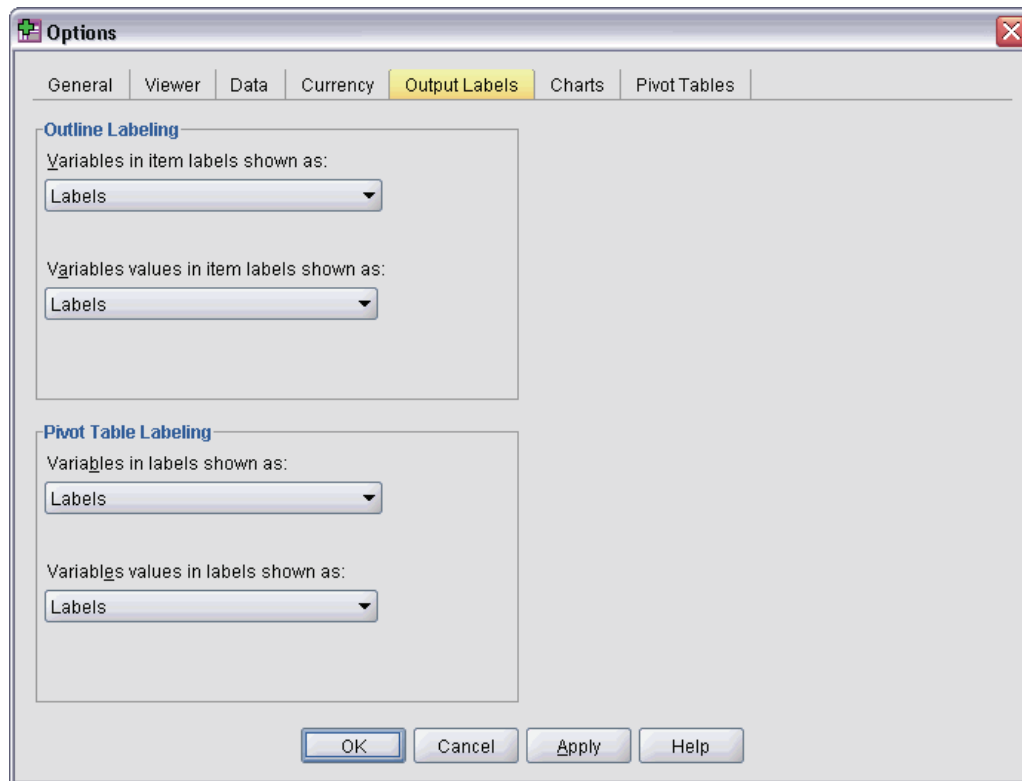
- ▶ Click the Currency tab.
- ▶ Select one of the currency formats from the list (CCA, CCB, CCC, CCD, and CCE).
- ▶ Enter the prefix, suffix, and decimal indicator values.
- ▶ Click OK or Apply.

Output Label Options

Output label options control the display of variable and data value information in the outline and pivot tables. You can display variable names, defined variable labels and actual data values, defined value labels, or a combination.

Descriptive variable and value labels (Variable view in the Data Editor, *Label* and *Values* columns) often make it easier to interpret your results. However, long labels can be awkward in some tables.

Figure 48-6
Options dialog box: Output Labels tab



Output label options affect only new output produced after you change the settings. Output already displayed in the Viewer is not affected by changes in these settings. These settings affect only pivot table output. Text output is not affected by these settings.

Chart Options

Figure 48-7
Options dialog box: Charts tab

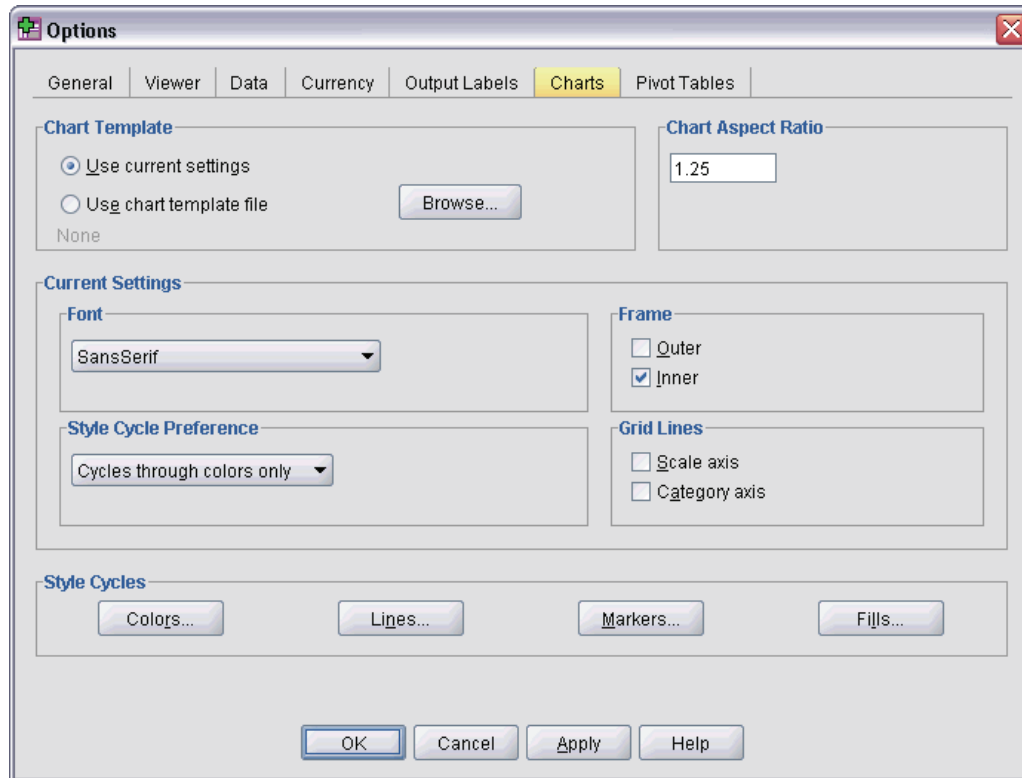


Chart Template. New charts can use either the settings selected here or the settings from a chart template file. Click **Browse** to select a chart template file. To create a chart template file, create a chart with the attributes that you want and save it as a template (choose **Save Chart Template** from the **File** menu).

Chart Aspect Ratio. The width-to-height ratio of the outer frame of new charts. You can specify a width-to-height ratio from 0.1 to 10.0. Values less than 1 make charts that are taller than they are wide. Values greater than 1 make charts that are wider than they are tall. A value of 1 produces a square chart. Once a chart is created, its aspect ratio cannot be changed.

Current Settings. Available settings include:

- **Font.** Font used for all text in new charts.
- **Style Cycle Preference.** The initial assignment of colors and patterns for new charts. Cycle through colors only uses only colors to differentiate chart elements and does not use patterns. Cycle through patterns only uses only line styles, marker symbols, or fill patterns to differentiate chart elements and does not use color.
- **Frame.** Controls the display of inner and outer frames on new charts.
- **Grid Lines.** Controls the display of scale and category axis grid lines on new charts.

Style Cycles. Customizes the colors, line styles, marker symbols, and fill patterns for new charts. You can change the order of the colors and patterns that are used when a new chart is created.

Data Element Colors

Specify the order in which colors should be used for the data elements (such as bars and markers) in your new chart. Colors are used whenever you select a choice that includes *color* in the Style Cycle Preference group in the main Chart Options dialog box.

For example, if you create a clustered bar chart with two groups and you select Cycle through colors, then patterns in the main Chart Options dialog box, the first two colors in the Grouped Charts list are used as the bar colors on the new chart.

To Change the Order in Which Colors Are Used

- ▶ Select Simple Charts and then select a color that is used for charts without categories.
- ▶ Select Grouped Charts to change the color cycle for charts with categories. To change a category's color, select a category and then select a color for that category from the palette.

Optionally, you can:

- Insert a new category above the selected category.
- Move a selected category.
- Remove a selected category.
- Reset the sequence to the default sequence.
- Edit a color by selecting its well and then clicking Edit.

Data Element Lines

Specify the order in which styles should be used for the line data elements in your new chart. Line styles are used whenever your chart includes line data elements and you select a choice that includes *patterns* in the Style Cycle Preference group in the main Chart Options dialog box.

For example, if you create a line chart with two groups and you select Cycle through patterns only in the main Chart Options dialog box, the first two styles in the Grouped Charts list are used as the line patterns on the new chart.

To Change the Order in Which Line Styles Are Used

- ▶ Select Simple Charts and then select a line style that is used for line charts without categories.
- ▶ Select Grouped Charts to change the pattern cycle for line charts with categories. To change a category's line style, select a category and then select a line style for that category from the palette.

Optionally, you can:

- Insert a new category above the selected category.
- Move a selected category.

- Remove a selected category.
- Reset the sequence to the default sequence.

Data Element Markers

Specify the order in which symbols should be used for the marker data elements in your new chart. Marker styles are used whenever your chart includes marker data elements and you select a choice that includes *patterns* in the Style Cycle Preference group in the main Chart Options dialog box.

For example, if you create a scatterplot chart with two groups and you select Cycle through patterns only in the main Chart Options dialog box, the first two symbols in the Grouped Charts list are used as the markers on the new chart.

To Change the Order in Which Marker Styles Are Used

- ▶ Select Simple Charts and then select a marker symbol that is used for charts without categories.
- ▶ Select Grouped Charts to change the pattern cycle for charts with categories. To change a category's marker symbol, select a category and then select a symbol for that category from the palette.

Optionally, you can:

- Insert a new category above the selected category.
- Move a selected category.
- Remove a selected category.
- Reset the sequence to the default sequence.

Data Element Fills

Specify the order in which fill styles should be used for the bar and area data elements in your new chart. Fill styles are used whenever your chart includes bar or area data elements, and you select a choice that includes *patterns* in the Style Cycle Preference group in the main Chart Options dialog box.

For example, if you create a clustered bar chart with two groups and you select Cycle through patterns only in the main Chart Options dialog box, the first two styles in the Grouped Charts list are used as the bar fill patterns on the new chart.

To Change the Order in Which Fill Styles Are Used

- ▶ Select Simple Charts and then select a fill pattern that is used for charts without categories.
- ▶ Select Grouped Charts to change the pattern cycle for charts with categories. To change a category's fill pattern, select a category and then select a fill pattern for that category from the palette.

Optionally, you can:

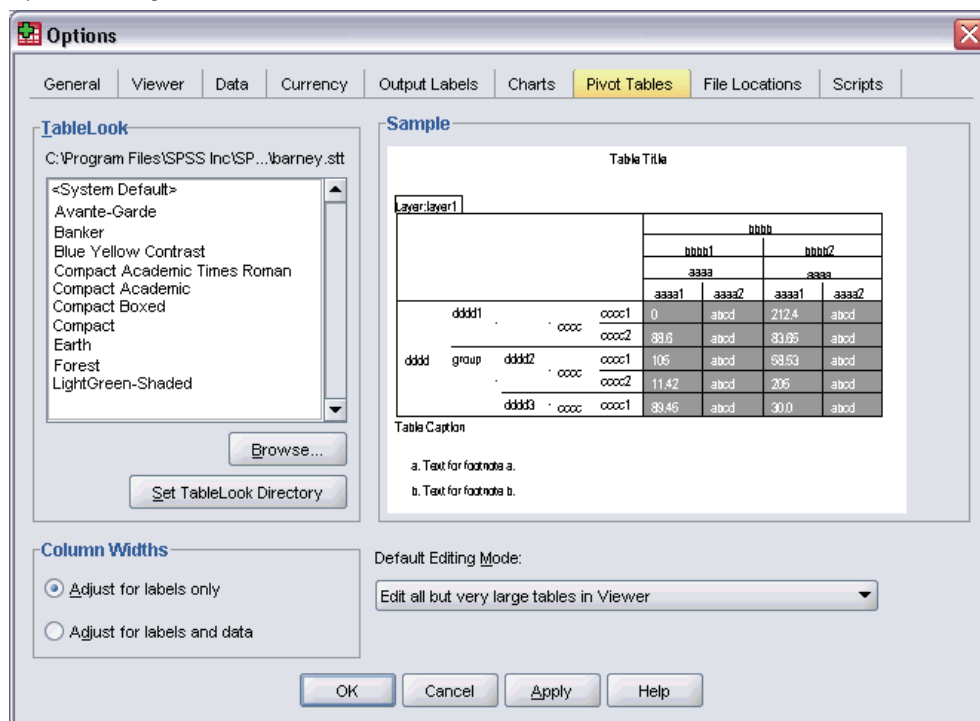
- Insert a new category above the selected category.
- Move a selected category.

- Remove a selected category.
- Reset the sequence to the default sequence.

Pivot Table Options

Pivot Table options set the default TableLook used for new pivot table output. TableLooks can control a variety of pivot table attributes, including the display and width of grid lines; font style, size, and color; and background colors.

Figure 48-8
Options dialog box: Pivot Tables tab



TableLook. Select a TableLook from the list of files and click OK or Apply. You can use one of the TableLooks provided with SPSS Statistics, or you can create your own in the Pivot Table Editor (choose TableLooks from the Format menu).

- **Browse.** Allows you to select a TableLook from another directory.
- **Set TableLook Directory.** Allows you to change the default TableLook directory.

Note: TableLooks created in earlier versions of SPSS Statistics cannot be used in version 16.0 or later.

Column Widths. Controls the automatic adjustment of column widths in pivot tables.

- **Adjust for labels and data except for extremely large tables.** For tables that don't exceed 10,000 cells, adjusts column width to whichever is larger—the column label or the largest data value. For tables with more than 10,000 cells, adjusts column width to the width of the column label.

- **Adjust for labels only.** Adjusts column width to the width of the column label. This produces more compact tables, but data values wider than the label may be truncated.
- **Adjust for labels and data.** Adjusts column width to whichever is larger: the column label or the largest data value. This produces wider tables, but it ensures that all values will be displayed.

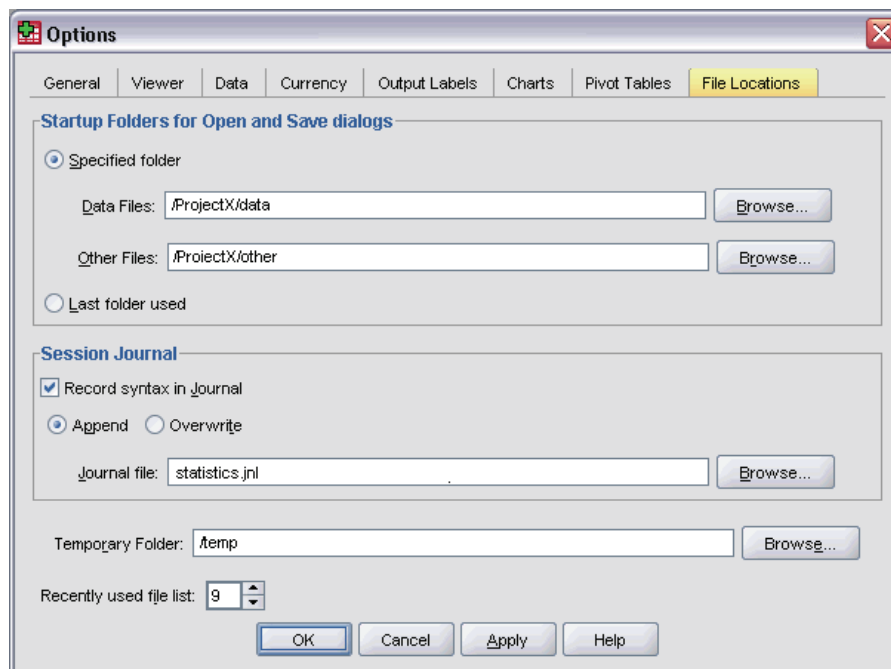
Default Editing Mode. Controls activation of pivot tables in the Viewer window or in a separate window. By default, double-clicking a pivot table activates all but very large tables in the Viewer window. You can choose to activate pivot tables in a separate window or select a size setting that will open smaller pivot tables in the Viewer window and larger pivot tables in a separate window.

Copying wide tables to the clipboard in rich text format. When pivot tables are pasted in Word/RTF format, tables that are too wide for the document width will either be wrapped, scaled down to fit the document width, or left unchanged.

File Locations Options

Options on the File Locations tab control the default location that the application will use for opening and saving files at the start of each session, the location of the journal file, the location of the temporary folder, and the number of files that appear in the recently used file list.

Figure 48-9
Options dialog box: File Locations tab



Startup Folders for Open and Save Dialog Boxes

- **Specified folder.** The specified folder is used as the default location at the start of each session. You can specify different default locations for data and other (non-data) files.
- **Last folder used.** The last folder used to either open or save files in the previous session is used as the default at the start of the next session. This applies to both data and other (non-data) files.

These settings apply only to dialog boxes for opening and saving files—and the “last folder used” is determined by the last dialog box used to open or save a file. Files opened or saved via command syntax have no effect on—and are not affected by—these settings. These settings are only available in local analysis mode. In distributed analysis mode connected to a remote server (requires SPSS Statistics Server), you cannot control the startup folder locations.

Session Journal

You can use the session journal to automatically record commands run in a session. This includes commands entered and run in syntax windows and commands generated by dialog box choices. You can edit the journal file and use the commands again in other sessions. You can turn journaling off and on, append or overwrite the journal file, and select the journal filename and location. You can copy command syntax from the journal file and save it in a syntax file.

Temporary Folder

This specifies the location of temporary files created during a session. In distributed mode (available with the server version), this does not affect the location of temporary data files. In distributed mode, the location of temporary data files is controlled by the environment variable *SPSSTMPDIR*, which can be set only on the computer running the server version of the software. If you need to change the location of the temporary directory, contact your system administrator.

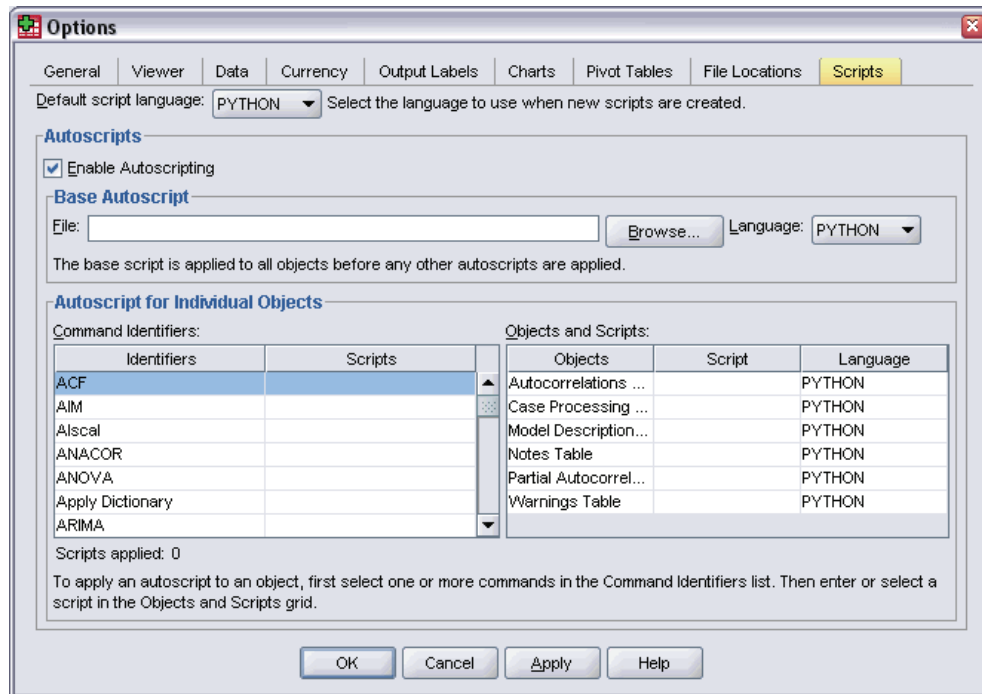
Recently Used File List

This controls the number of recently used files that appear on the File menu.

Script Options

Use the Scripts tab to specify the default script language and any autoscripts you want to use. You can use scripts to automate many functions, including customizing pivot tables.

Figure 48-10
Options dialog box: Scripts tab



Note: Legacy Sax Basic users must manually convert any custom autoscripts. The autoscripts installed with pre-16.0 versions are available as a set of separate script files located in the *Samples* subdirectory of the directory where SPSS Statistics is installed. By default, no output items are associated with autoscripts. You must manually associate all autoscripts with the desired output items, as described below. For information on converting legacy autoscripts, see [Compatibility with Versions Prior to 16.0](#) on p. 589.

Default script language. The default script language determines the script editor that is launched when new scripts are created. It also specifies the default language whose executable will be used to run autoscripts. The available scripting languages depend on your platform. For Windows, the available scripting languages are Basic, which is installed with the Base system, and the Python programming language. For all other platforms, scripting is available with the Python programming language.

Note: To enable scripting with the Python programming language, you must separately install Python and the SPSS Statistics-Python Integration Plug-in. The Python programming language is not installed with the Base system, but is available from the installation CD for Windows. The SPSS Statistics-Python Integration Plug-in is available from the installation CD for all platforms..

Enable Autoscripting. This check box allows you to enable or disable autoscripting. By default, autoscripting is enabled.

Base Autoscript. An optional script that is applied to all new Viewer objects before any other autoscripts are applied. Specify the script file to be used as the base autoscript as well as the language whose executable will be used to run the script.

To Apply Autoscripts to Output Items

- ▶ In the Command Identifiers grid, select a command that generates output items to which autoscripts will be applied.

The *Objects* column, in the Objects and Scripts grid, displays a list of the objects associated with the selected command. The *Script* column displays any existing script for the selected command.

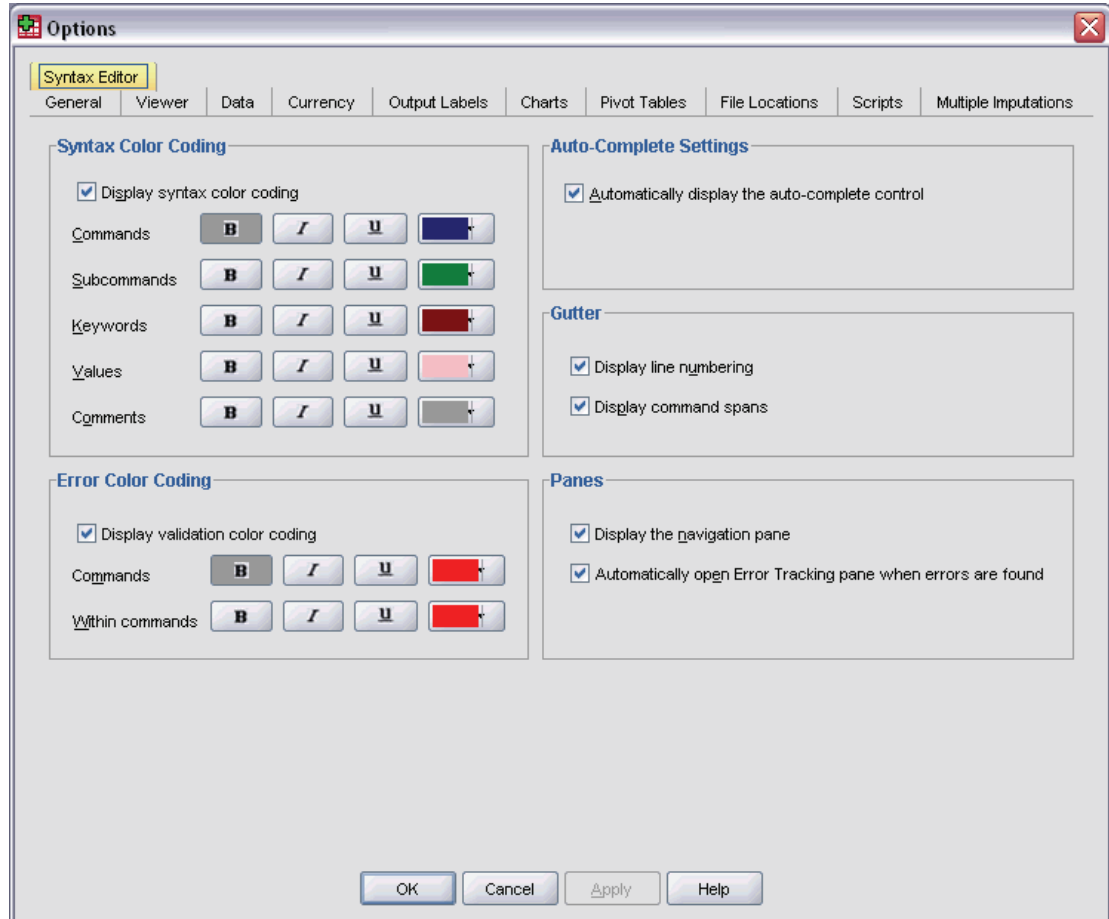
- ▶ Specify a script for any of the items shown in the *Objects* column. Click on the corresponding *Script* cell. Enter the path to the script or click the ellipsis (...) button to browse for the script.
- ▶ Specify the language whose executable will be used to run the script. *Note:* The selected language is not affected by changing the default script language.
- ▶ Click Apply or OK.

To Remove Autoscript Associations

- ▶ In the Objects and Scripts grid, click on the cell in the Script column corresponding to the script to dissociate.
- ▶ Delete the path to the script and then click on any other cell in the Objects and Scripts grid.
- ▶ Click Apply or OK.

Syntax Editor Options

Figure 48-11
Options dialog box: Syntax Editor tab



Syntax Color Coding

You can turn color coding of commands, subcommands, keywords, keyword values, and comments off or on and you can specify the font style and color for each.

Error Color Coding

You can turn color coding of certain syntactical errors off or on and you can specify the font style and color used. Both the command name and the text (within the command) containing the error are color coded, and you can choose different styles for each. [For more information, see Color Coding in Chapter 13 on p. 260.](#)

Auto-Complete Settings

You can turn automatic display of the auto-complete control off or on. The auto-complete control can always be displayed by pressing Ctrl+Spacebar. [For more information, see Auto-Completion in Chapter 13 on p. 259.](#)

Gutter

These options specify the default behavior for showing or hiding line numbers and command spans within the Syntax Editor's gutter—the region to the left of the text pane that is reserved for line numbers, bookmarks, breakpoints, and command spans. Command spans are icons that provide visual indicators of the start and end of a command.

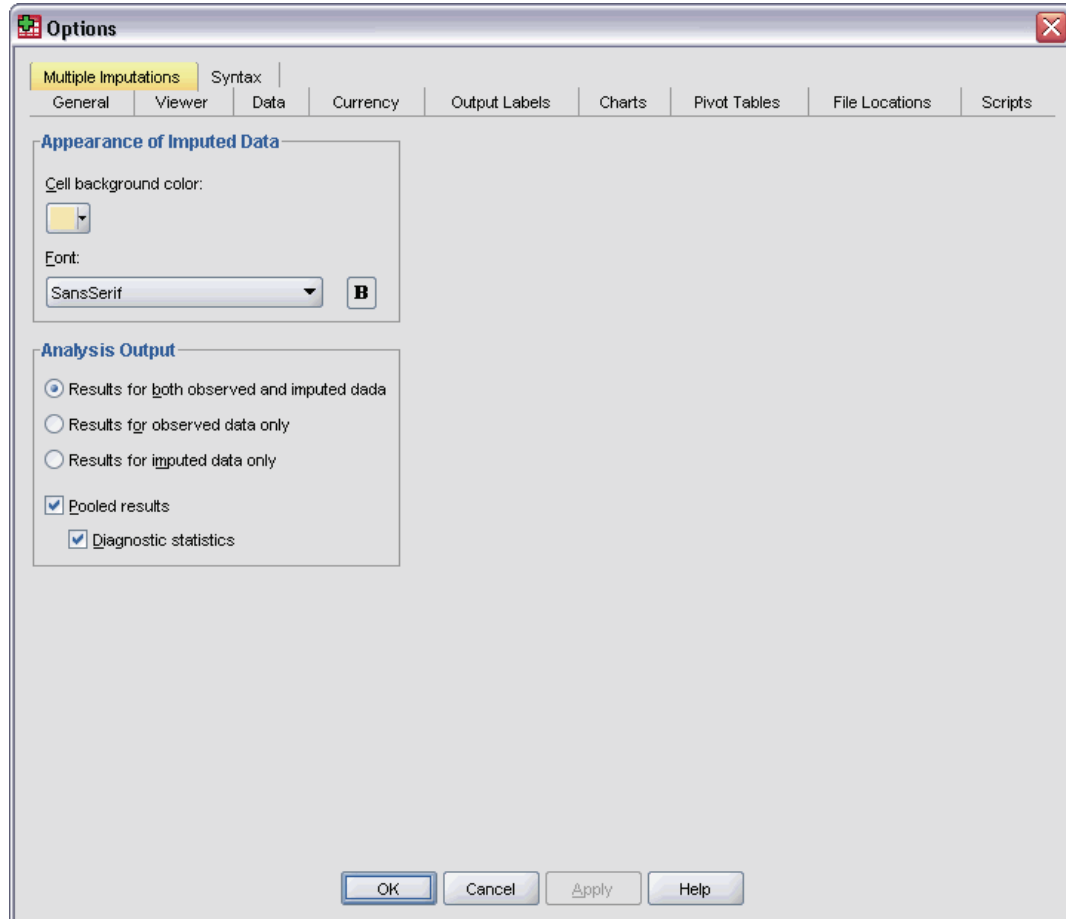
Panes

Display the navigation pane. This specifies the default for showing or hiding the navigation pane. The navigation pane contains a list of all recognized commands in the syntax window, displayed in the order in which they occur. Clicking on a command in the navigation pane positions the cursor at the start of the command.

Automatically open Error Tracking pane when errors are found. This specifies the default for showing or hiding the error tracking pane when run-time errors are found.

Multiple Imputations Options

Figure 48-12
Options dialog box: Multiple Imputations tab



The Multiple Imputations tab controls two kinds of preferences related to Multiple Imputations:

Appearance of Imputed Data. By default, cells containing imputed data will have a different background color than cells containing nonimputed data. The distinctive appearance of the imputed data should make it easy for you to scroll through a dataset and locate those cells. You can change the default cell background color, the font, and make the imputed data display in bold type.

Analysis Output. This group controls the type of Viewer output produced whenever a multiply imputed dataset is analyzed. By default, output will be produced for the original (pre-imputation) dataset and for each of the imputed datasets. In addition, for those procedures that support pooling of imputed data, final pooled results will be generated. When univariate pooling is performed, pooling diagnostics will also display. However, you can suppress any output you do not want to see.

To Set Multiple Imputation Options

From the menus, choose:

Edit
Options

Click the Multiple Imputation tab.

Customizing Menus and Toolbars

Menu Editor

You can use the Menu Editor to customize your menus. With the Menu Editor you can:

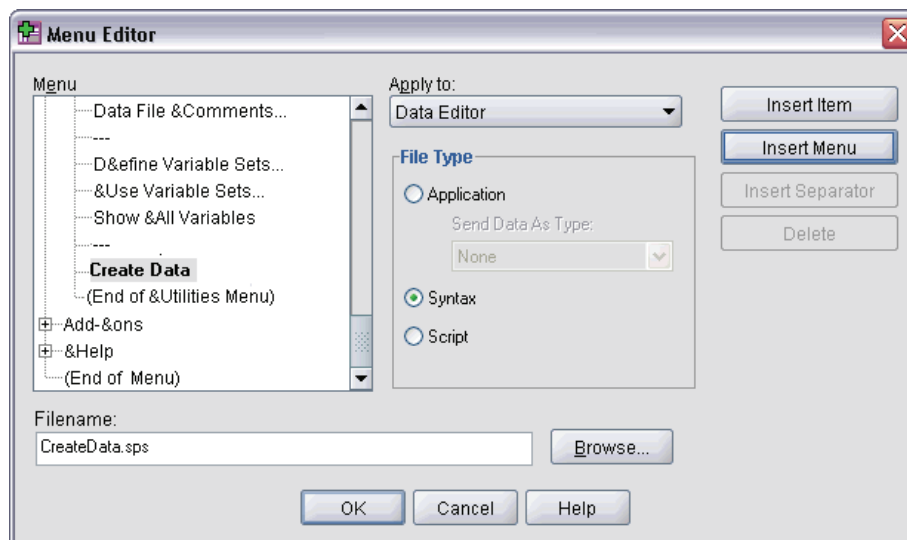
- Add menu items that run customized scripts.
- Add menu items that run command syntax files.
- Add menu items that launch other applications and automatically send data to other applications.

You can send data to other applications in the following formats: SPSS Statistics, Excel, Lotus 1-2-3, tab-delimited, and dBASE IV.

To Add Items to Menus

- ▶ From the menus choose:
View
Menu Editor..
- ▶ In the Menu Editor dialog box, double-click the menu (or click the plus sign icon) to which you want to add a new item.
- ▶ Select the menu item above which you want the new item to appear.
- ▶ Click Insert Item to insert a new menu item.
- ▶ Select the file type for the new item (script file, command syntax file, or external application).
- ▶ Click Browse to select a file to attach to the menu item.

Figure 49-1
Menu Editor dialog box



You can also add entirely new menus and separators between menu items. Optionally, you can automatically send the contents of the Data Editor to another application when you select that application on the menu.

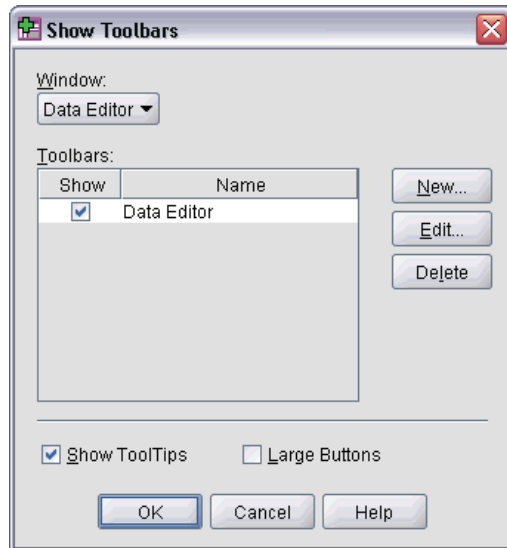
Customizing Toolbars

You can customize toolbars and create new toolbars. Toolbars can contain any of the available tools, including tools for all menu actions. They can also contain custom tools that launch other applications, run command syntax files, or run script files.

Show Toolbars

Use Show Toolbars to show or hide toolbars, customize toolbars, and create new toolbars. Toolbars can contain any of the available tools, including tools for all menu actions. They can also contain custom tools that launch other applications, run command syntax files, or run script files.

Figure 49-2
Show Toolbars dialog box



To Customize Toolbars

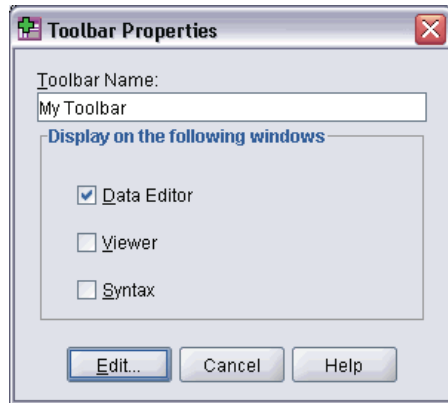
- ▶ From the menus choose:
 - View
 - Toolbars
 - Customize
- ▶ Select the toolbar you want to customize and click Edit, or click New to create a new toolbar.
- ▶ For new toolbars, enter a name for the toolbar, select the windows in which you want the toolbar to appear, and click Edit.
- ▶ Select an item in the Categories list to display available tools in that category.
- ▶ Drag and drop the tools you want onto the toolbar displayed in the dialog box.
- ▶ To remove a tool from the toolbar, drag it anywhere off the toolbar displayed in the dialog box.
To create a custom tool to open a file, to run a command syntax file, or to run a script:
- ▶ Click New Tool in the Edit Toolbar dialog box.
- ▶ Enter a descriptive label for the tool.
- ▶ Select the action you want for the tool (open a file, run a command syntax file, or run a script).
- ▶ Click Browse to select a file or application to associate with the tool.

New tools are displayed in the User-Defined category, which also contains user-defined menu items.

Toolbar Properties

Use Toolbar Properties to select the window types in which you want the selected toolbar to appear. This dialog box is also used for creating names for new toolbars.

Figure 49-3
Toolbar Properties dialog box



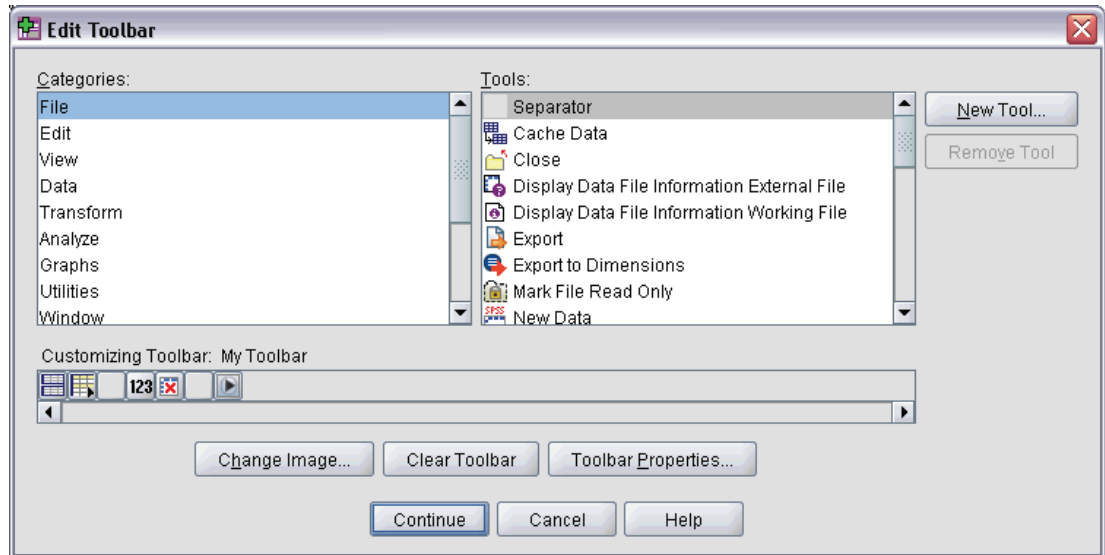
To Set Toolbar Properties

- ▶ From the menus choose:
 - View
 - Toolbars
 - Customize
- ▶ For existing toolbars, click Edit, and then click Properties in the Edit Toolbar dialog box.
- ▶ For new toolbars, click New Tool.
- ▶ Select the window types in which you want the toolbar to appear. For new toolbars, also enter a toolbar name.

Edit Toolbar

Use the Edit Toolbar dialog box to customize existing toolbars and create new toolbars. Toolbars can contain any of the available tools, including tools for all menu actions. They can also contain custom tools that launch other applications, run command syntax files, or run script files.

Figure 49-4
Customize Toolbar dialog box



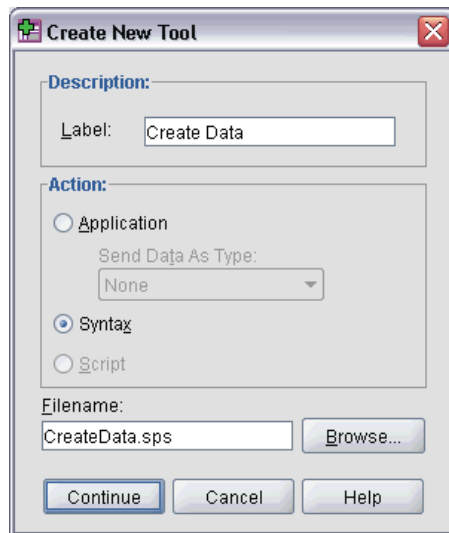
To Change Toolbar Images

- ▶ Select the tool for which you want to change the image on the toolbar display.
- ▶ Click Change Image.
- ▶ Select the image file that you want to use for the tool. The following image formats are supported: BMP, PNG, GIF, JPG.
 - Images should be square. Images that are not square are clipped to a square area.
 - Images are automatically scaled to fit. For optimal display, use 16x16 pixel images for small toolbar images or 32x32 pixel images for large toolbar images.

Create New Tool

Use the Create New Tool dialog box to create custom tools to launch other applications, run command syntax files, and run script files.

Figure 49-5
Create New Tool dialog box



Creating and Managing Custom Dialogs

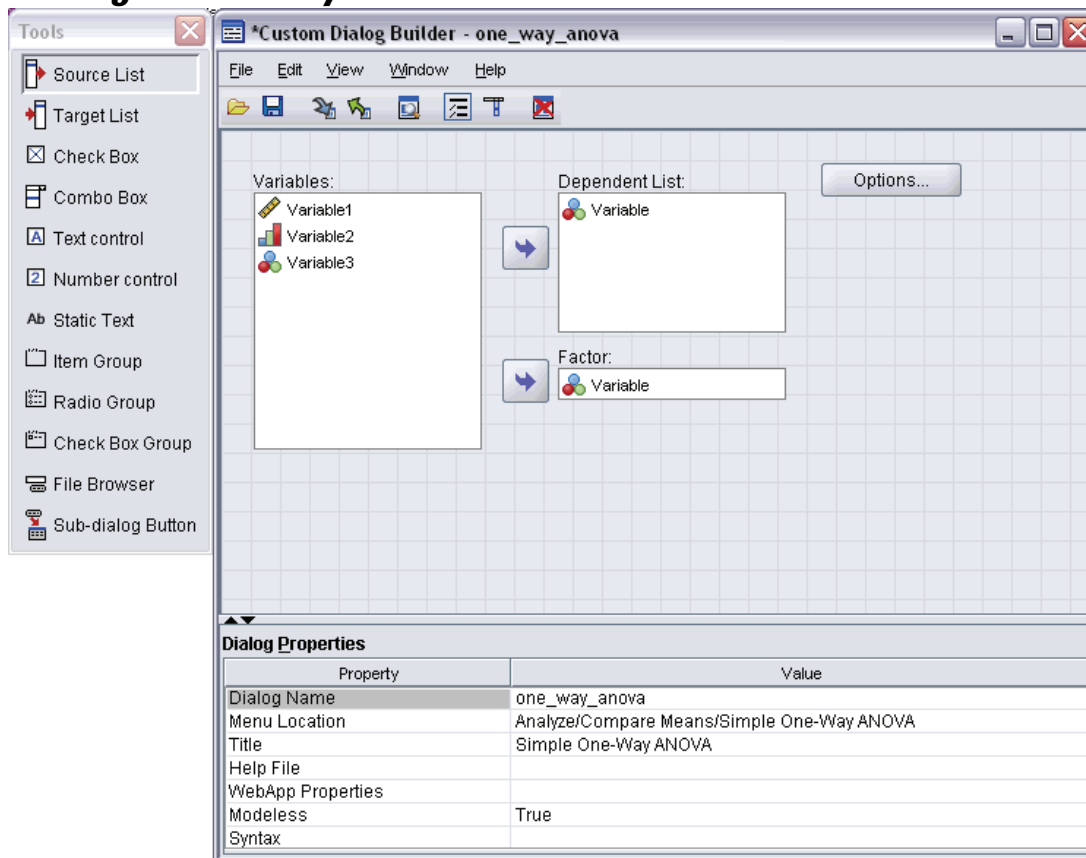
The Custom Dialog Builder allows you to create and manage custom dialogs for generating command syntax. Using the Custom Dialog Builder you can:

- Create your own version of a dialog for a built-in SPSS Statistics procedure. For example, you can create a dialog for the Frequencies procedure that only allows the user to select the set of variables and then generates command syntax with pre-set options that standardize the output.
- Create a user interface that generates command syntax for an extension command. Extension commands are user-defined SPSS Statistics commands that are implemented in either the Python programming language or R. [For more information, see Custom Dialogs for Extension Commands on p. 540.](#)
- Open a file containing the specification for a custom dialog—perhaps created by another user—and add the dialog to your installation of SPSS Statistics, optionally making your own modifications.
- Save the specification for a custom dialog so that other users can add it to their installations of SPSS Statistics.

How to Start the Custom Dialog Builder

- ▶ From the menus choose:
 - Utilities
 - Custom Dialog Builder

Custom Dialog Builder Layout



Canvas

The canvas is the area of the Custom Dialog Builder where you design the layout of your dialog.

Properties Pane

The properties pane is the area of the Custom Dialog Builder where you specify properties of the controls that make up the dialog as well as properties of the dialog itself, such as the menu location.

Tools Palette

The tools palette provides the set of controls that can be included in a custom dialog. You can show or hide the Tools Palette by choosing Tools Palette from the View menu.

Building a Custom Dialog

The basic steps involved in building a custom dialog are:

- ▶ Specify the properties of the dialog itself, such as the title that appears when the dialog is launched and the location of the new menu item for the dialog within the SPSS Statistics menus. [For more information, see Dialog Properties on p. 522.](#)
- ▶ Specify the controls, such as source and target variable lists, that make up the dialog and any sub-dialogs. [For more information, see Control Types on p. 529.](#)
- ▶ Create the syntax template that specifies the command syntax to be generated by the dialog. [For more information, see Building the Syntax Template on p. 525.](#)
- ▶ Install the dialog to SPSS Statistics and/or save the specification for the dialog to a custom dialog package (.spd) file. [For more information, see Managing Custom Dialogs on p. 528.](#)

You can preview your dialog as you're building it. [For more information, see Previewing a Custom Dialog on p. 527.](#)

Dialog Properties

To view and set Dialog Properties:

- ▶ Click on the canvas in an area outside of any controls. With no controls on the canvas, Dialog Properties are always visible.

Dialog Name. The Dialog Name property is required and specifies a unique name to associate with the dialog. This is the name used to identify the dialog when installing or uninstalling it. To minimize the possibility of name conflicts, you may want to prefix the dialog name with an identifier for your organization, such as a URL.

Menu Location. Click the ellipsis (...) button to open the Menu Location dialog box, which allows you to specify the name and location of the menu item for the custom dialog. Adding a menu location for a new custom dialog or changing the menu location of an installed custom dialog requires a restart of SPSS Statistics.

Title. The Title property specifies the text to be displayed in the title bar of the dialog box.

Help File. The Help File property is optional and specifies the path to a help file for the dialog. This is the file that will be launched when the user clicks the Help button on the dialog. Help files must be in HTML format. A copy of the specified help file is included with the specifications for the dialog when the dialog is installed or saved to a custom dialog package file. The Help button on the run-time dialog is hidden if there is no associated help file.

Any supporting files, such as image files and style sheets, must be stored along with the main help file once the custom dialog has been installed. By default, the specifications for an installed custom dialog are stored in the *ext/lib/<Dialog Name>* folder of the installation directory—located at the root of the installation directory for Windows and Linux, and under the *Contents* directory in the application bundle for Mac. Supporting files should be located at the root of the folder and not in sub-folders. They must be manually added to any custom dialog package files you create for the dialog.

If you have specified alternate locations for installed dialogs—using the `SPSS_CDIALOGS_PATH` environment variable—then store any supporting files under the `<Dialog Name>` folder at the appropriate alternate location. For more information, see [Managing Custom Dialogs on p. 528](#).

Note: When working with a dialog opened from a custom dialog package (`.spd`) file, the Help File property points to a temporary folder associated with the `.spd` file. Any modifications to the help file should be made to the copy in the temporary folder.

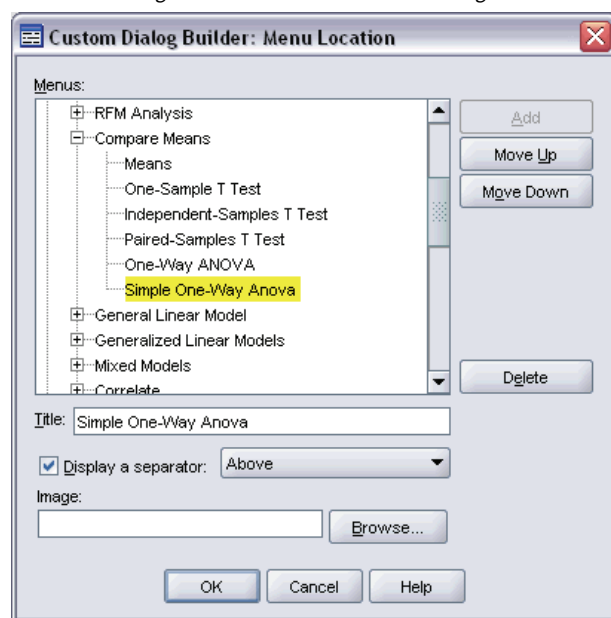
Web Deployment Properties. Allows you to associate a properties file with this dialog for use in building thin client applications that are deployed over the web.

Modeless. Specifies whether the dialog is modal or modeless. When a dialog is modal, it must be closed before the user can interact with the main application windows (Data, Output, and Syntax) or with other open dialogs. Modeless dialogs do not have that constraint. The default is modeless.

Syntax. The Syntax property specifies the syntax template, used to create the command syntax generated by the dialog at run-time. Click the ellipsis (...) button to open the Syntax Template. For more information, see [Building the Syntax Template on p. 525](#).

Specifying the Menu Location for a Custom Dialog

Figure 50-1
Custom Dialog Builder Menu Location dialog box



The Menu Location dialog box allows you to specify the name and location of the menu item for a custom dialog. Adding a menu location for a new custom dialog or changing the menu location of an installed custom dialog requires a restart of SPSS Statistics. Menu items for custom dialogs do not appear in the Menu Editor within SPSS Statistics.

- Double-click the menu (or click the plus sign icon) to which you want to add the item for the new dialog. You can also add items to the top-level menu labelled Custom (located between the Graphs and Utilities items), which is only available for menu items associated with custom dialogs.

If you want to create custom menus or sub-menus, use the Menu Editor. [For more information, see Menu Editor in Chapter 49 on p. 514.](#) Note, however, that other users of your dialog will have to manually create the same menu or sub-menu from their Menu Editor; otherwise, the dialog will be added to their Custom menu.

Note: The Menu Location dialog box displays all menus, including those for all add-on modules. Be sure to add the menu item for your custom dialog to a menu that will be available to you or end users of your dialog.

- ▶ Select the menu item above which you want the item for the new dialog to appear. Once the item is added you can use the Move Up and Move Down buttons to reposition it.
- ▶ Enter a title for the menu item. Titles within a given menu or sub-menu must be unique.
- ▶ Click Add.

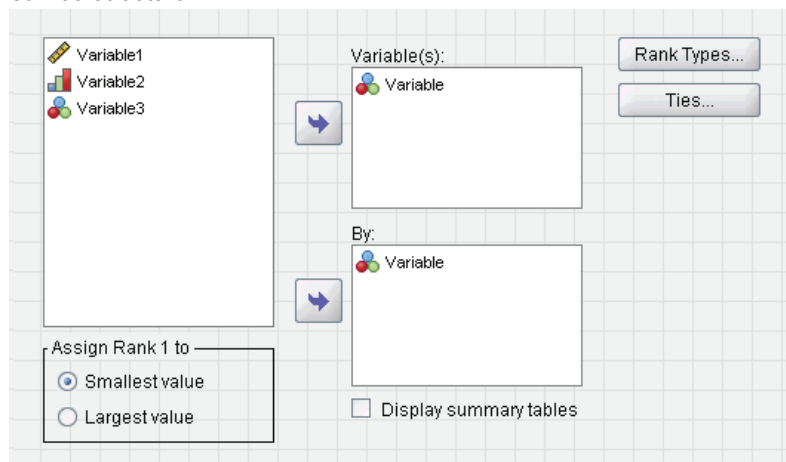
Optionally, you can:

- Add a separator above or below the new menu item.
- Specify the path to an image that will appear next to the menu item for the custom dialog. The supported image types are *gif* and *png*. The image cannot be larger than 16 x 16 pixels.

Laying Out Controls on the Canvas

You add controls to a custom dialog by dragging them from the tools palette onto the canvas. To ensure consistency with built-in dialogs, the canvas is divided into three functional columns in which you can place controls.

Figure 50-2
Canvas structure



Left Column. The left column is primarily intended for a source list control. All controls other than target lists and sub-dialog buttons can be placed in the left column.

Center Column. The center column can contain any control other than source lists and sub-dialog buttons.

Right Column. The right column can only contain sub-dialog buttons.

Although not shown on the canvas, each custom dialog contains OK, Paste, Cancel, and Help buttons positioned across the bottom of the dialog. The presence and locations of these buttons is automatic, however, the Help button is hidden if there is no help file associated with the dialog (as specified by the Help File property in Dialog Properties).

You can change the vertical order of the controls within a column by dragging them up or down, but the exact position of the controls will be determined automatically for you. At run-time, controls will resize in appropriate ways when the dialog itself is resized. Controls such as source and target lists automatically expand to fill the available space below them.

Building the Syntax Template

The syntax template specifies the command syntax that will be generated by the custom dialog. A single custom dialog can generate command syntax for any number of built-in SPSS Statistics commands or extension commands.

The syntax template may consist of static text that is always generated and control identifiers that are replaced at run-time with the values of the associated custom dialog controls. For example, command names and subcommand specifications that don't depend on user input would be static text, whereas the set of variables specified in a target list would be represented with the control identifier for the target list control.

To Build the Syntax Template

- ▶ From the menus in the Custom Dialog Builder choose:
 - Edit
 - Syntax Template
- (Or click the ellipsis (...) button in the Syntax property field in Dialog Properties)
- ▶ For static command syntax that does not depend on user-specified values, enter the syntax as you would in the Syntax Editor.
 - ▶ Add control identifiers of the form `%%Identifier%%` at the locations where you want to insert command syntax generated by controls, where `Identifier` is the value of the Identifier property for the control. You can select from a list of available control identifiers by pressing Ctrl+Spacebar. If you manually enter identifiers, retain any spaces, since all spaces in identifiers are significant.

At run-time, and for all controls other than check boxes and check box groups, each identifier is replaced with the current value of the Syntax property of the associated control. For check boxes and check box groups, the identifier is replaced by the current value of the Checked Syntax or Unchecked Syntax property of the associated control, depending on the current state of the control—checked or unchecked. [For more information, see Control Types on p. 529.](#)

Example: Including Run-time Values in the Syntax Template

Consider a simplified version of the Frequencies dialog that only contains a source list control and a target list control, and generates command syntax of the following form:

```
FREQUENCIES VARIABLES=var1 var2...  
/FORMAT = NOTABLE
```

```
/BARCHART.
```

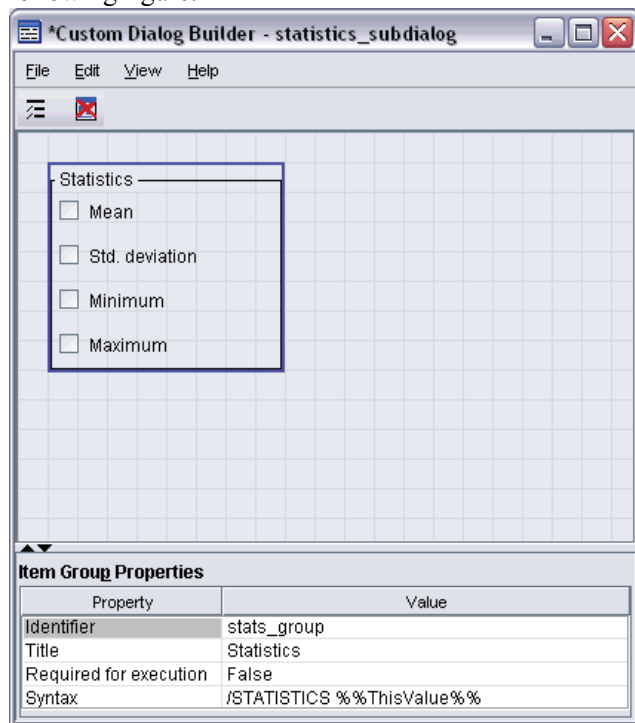
The syntax template to generate this might look like:

```
FREQUENCIES VARIABLES=%%target_list%%
  /FORMAT = NOTABLE
  /BARCHART.
```

- %%target_list%% is the value of the Identifier property for the target list control. At run-time it will be replaced by the current value of the Syntax property of the control.
- Defining the Syntax property of the target list control to be %%ThisValue%% specifies that at run-time, the current value of the property will be the value of the control, which is the set of variables in the target list.

Example: Including Command Syntax from Container Controls

Building on the previous example, consider adding a Statistics sub-dialog that contains a single group of check boxes that allow a user to specify mean, standard deviation, minimum, and maximum. Assume the check boxes are contained in an item group control, as shown in the following figure.



An example of the generated command syntax would be:

```
FREQUENCIES VARIABLES=var1 var2...
  /FORMAT = NOTABLE
  /STATISTICS MEAN STDDEV
  /BARCHART.
```

The syntax template to generate this might look like:

```
FREQUENCIES VARIABLES=%%target_list%%
  /FORMAT = NOTABLE
  %%stats_group%%
  /BARCHART.
```

- %%target_list%% is the value of the Identifier property for the target list control and %%stats_group%% is the value of the Identifier property for the item group control.

The following table shows one way to specify the Syntax properties of the item group and the check boxes contained within it in order to generate the desired result. The Syntax property of the target list would be set to %%ThisValue%%, as described in the previous example.

Syntax property of item group	/STATISTICS %%ThisValue%%
Checked Syntax property of mean check box	MEAN
Checked Syntax property of stddev check box	STDDEV
Checked Syntax property of min check box	MINIMUM
Checked Syntax property of max check box	MAXIMUM

At run-time, %%stats_group%% will be replaced by the current value of the Syntax property for the item group control. Specifically, %%ThisValue%% will be replaced by a blank-separated list of the values of the Checked or Unchecked Syntax property of each check box, depending on its state—checked or unchecked. Since we only specified values for the Checked Syntax property, only checked boxes will contribute to %%ThisValue%%. For example, if the user checks the mean and standard deviation boxes, the run-time value of the Syntax property for the item group will be /STATISTICS MEAN STDDEV.

If no boxes are checked, then the Syntax property for the item group control will be empty, and the generated command syntax will not contain any reference to %%stats_group%%. This may or may not be desirable. For example, even with no boxes checked you may still want to generate the STATISTICS subcommand. This can be accomplished by referencing the identifiers for the check boxes directly in the syntax template, as in:

```
FREQUENCIES VARIABLES=%%target_list%%
  /FORMAT = NOTABLE
  /STATISTICS %%stats_mean%% %%stats_stddev%% %%stats_min%% %%stats_max%%
  /BARCHART.
```

Previewing a Custom Dialog

You can preview the dialog that is currently open in the Custom Dialog Builder. The dialog appears and functions as it would when run from the menus within SPSS Statistics.

- Source variable lists are populated with dummy variables that can be moved to target lists.
- The Paste button pastes command syntax into the designated syntax window.
- The OK button closes the preview.
- If a help file is specified, the Help button is enabled and will open the specified file. If no help file is specified, the help button is disabled when previewing, and hidden when the actual dialog is run.

To preview a custom dialog:

- ▶ From the menus in the Custom Dialog Builder choose:
 - File
 - Preview Dialog

Managing Custom Dialogs

The Custom Dialog Builder allows you to manage custom dialogs created by you or by other users. You can install, uninstall, or modify installed dialogs; and you can save specifications for a custom dialog to an external file or open a file containing the specifications for a custom dialog in order to modify it. Custom dialogs must be installed before they can be used.

Installing a Custom Dialog

You can install the dialog that is currently open in the Custom Dialog Builder or you can install a dialog from a custom dialog package (*.spd*) file. Re-installing an existing dialog will replace the existing version.

To install the currently open dialog:

- ▶ From the menus in the Custom Dialog Builder choose:
 - File
 - Install

To install from a custom dialog package file:

- ▶ From the menus choose:
 - Utilities
 - Install Custom Dialog

To use a new custom dialog, you will need to restart SPSS Statistics. Re-installing an existing dialog does not require a restart unless the menu location changes.

By default, installing a dialog requires write permission to the SPSS Statistics installation directory. If you do not have write permissions to that directory or would like to store installed dialogs elsewhere, you can specify one or more alternate locations by defining the *SPSS_CDIALOGS_PATH* environment variable. When present, the paths specified in *SPSS_CDIALOGS_PATH* take precedence over the SPSS Statistics installation directory. For multiple locations, separate each with a semicolon on Windows and a colon on Linux and Mac.

Opening an Installed Custom Dialog

You can open a custom dialog that you have already installed, allowing you to modify it and/or save it externally so that it can be distributed to other users.

- ▶ From the menus in the Custom Dialog Builder choose:
 - File
 - Open Installed

Note: If you are opening an installed dialog in order to modify it, choosing File > Install will re-install it, replacing the existing version.

Uninstalling a Custom Dialog

- ▶ From the menus in the Custom Dialog Builder choose:

- File
 - Uninstall

The menu item for the custom dialog will be removed the next time you start SPSS Statistics.

Saving to a Custom Dialog Package File

You can save the specifications for a custom dialog to an external file, allowing you to distribute the dialog to other users or save specifications for a dialog that is not yet installed. Specifications are saved to a custom dialog package (*.spd*) file.

- ▶ From the menus in the Custom Dialog Builder choose:

- File
 - Save

Opening a Custom Dialog Package File

You can open a custom dialog package file containing the specifications for a custom dialog, allowing you to modify the dialog and re-save it or install it.

- ▶ From the menus in the Custom Dialog Builder choose:

- File
 - Open

Manually Copying an Installed Custom Dialog

By default, the specifications for an installed custom dialog are stored in the *ext/lib/<Dialog Name>* folder under the installation directory—located at the root of the installation directory for Windows and Linux, and under the *Contents* directory in the application bundle for Mac. You can copy this folder to the same relative location in another instance of SPSS Statistics and it will be recognized as an installed dialog the next time that instance is launched.

- If you have specified alternate locations for installed dialogs—using the *SPSS_CDIALOGS_PATH* environment variable—then copy the *<Dialog Name>* folder from the appropriate alternate location.
- If alternate locations for installed dialogs have been defined for the instance of SPSS Statistics you are copying to, then you can copy to any one of the specified locations and the dialog will be recognized as an installed dialog the next time that instance is launched.

Control Types

The tools palette provides the controls that can be added to a custom dialog.

- **Source List:** A list of source variables from the active dataset. For more information, see [Source List on p. 530](#).
- **Target List:** A target for variables transferred from the source list. For more information, see [Target List on p. 531](#).
- **Check Box:** A single check box. For more information, see [Check Box on p. 532](#).
- **Combo Box:** A combo box for creating drop-down lists. For more information, see [Combo Box on p. 533](#).
- **Text control:** A text box that accepts arbitrary text as input. For more information, see [Text Control on p. 534](#).
- **Number control:** A text box that is restricted to numeric values as input. For more information, see [Number Control on p. 534](#).
- **Static Text control:** A control for displaying static text. For more information, see [Static Text Control on p. 535](#).
- **Item Group:** A container for grouping a set of controls, such as a set of check boxes. For more information, see [Item Group on p. 535](#).
- **Radio Group:** A group of radio buttons. For more information, see [Radio Group on p. 536](#).
- **Check Box Group:** A container for a set of controls that are enabled or disabled as a group, by a single check box. For more information, see [Check Box Group on p. 537](#).
- **File Browser:** A control for browsing the file system to open or save a file. For more information, see [File Browser on p. 538](#).
- **Sub-dialog Button:** A button for launching a sub-dialog. For more information, see [Sub-dialog Button on p. 539](#).

Source List

The Source Variable List control displays the list of variables from the active dataset that are available to the end user of the dialog. You can display all variables from the active dataset (the default) or you can filter the list based on type and measurement level—for instance, numeric variables that have a measurement level of scale. Use of a Source List control implies the use of one or more Target List controls. The Source List control has the following properties:

Identifier. The unique identifier for the control.

Title. An optional title that appears above the control. For multi-line titles, use `\n` to specify line breaks.

ToolTip. Optional ToolTip text that appears when the user hovers over the control. The specified text only appears when hovering over the title area of the control. Hovering over one of the listed variables will display the variable name and label.

Mnemonic Key. An optional character in the title to use as a keyboard shortcut to the control. The character appears underlined in the title. The shortcut is activated by pressing `Alt+[mnemonic key]`. The Mnemonic Key property is not supported on Mac.

Variable Transfers. Specifies whether variables transferred from the source list to a target list remain in the source list (Copy Variables), or are removed from the source list (Move Variables).

Variable Filter. Allows you to filter the set of variables displayed in the control. You can filter on variable type and measurement level, and you can specify that multiple response sets are included in the variable list. Click the ellipsis (...) button to open the Filter dialog. [For more information, see Filtering Variable Lists on p. 532.](#)

Note: The Source List control cannot be added to a sub-dialog.

Target List

The Target List control provides a target for variables that are transferred from the source list. Use of the Target List control assumes the presence of a Source List control. You can specify that only a single variable can be transferred to the control or that multiple variables can be transferred to it, and you can constrain which types of variables can be transferred to the control—for instance, only numeric variables with a measurement level of nominal or ordinal. The Target List control has the following properties:

Identifier. The unique identifier for the control. This is the identifier to use when referencing the control in the syntax template.

Title. An optional title that appears above the control. For multi-line titles, use \n to specify line breaks.

ToolTip. Optional ToolTip text that appears when the user hovers over the control. The specified text only appears when hovering over the title area of the control. Hovering over one of the listed variables will display the variable name and label.

Target list type. Specifies whether multiple variables or only a single variable can be transferred to the control.

Mnemonic Key. An optional character in the title to use as a keyboard shortcut to the control. The character appears underlined in the title. The shortcut is activated by pressing Alt+[mnemonic key]. The Mnemonic Key property is not supported on Mac.

Required for execution. Specifies whether a value is required in this control in order for execution to proceed. If True is specified, the OK and Paste buttons will be disabled until a value is specified for this control. If False is specified, the absence of a value in this control has no effect on the state of the OK and Paste buttons. The default is True.

Variable Filter. Allows you to constrain the types of variables that can be transferred to the control. You can constrain by variable type and measurement level, and you can specify whether multiple response sets can be transferred to the control. Click the ellipsis (...) button to open the Filter dialog. [For more information, see Filtering Variable Lists on p. 532.](#)

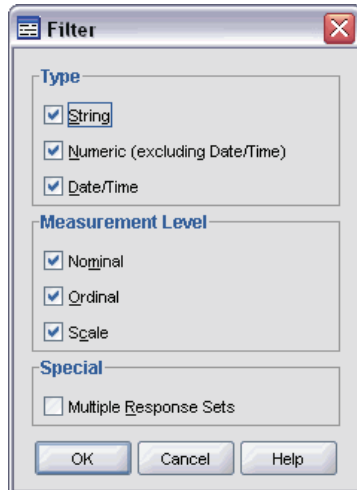
Syntax. Specifies command syntax that is generated by this control at run-time and can be inserted in the syntax template.

- You can specify any valid command syntax and you can use \n for line breaks.
- The value %%ThisValue%% specifies the run-time value of the control, which is the list of variables transferred to the control. This is the default.

Note: The Target List control cannot be added to a sub-dialog.

Filtering Variable Lists

Figure 50-3
Filter dialog box



The Filter dialog box, associated with source and target variable lists, allows you to filter the types of variables from the active dataset that can appear in the lists. You can also specify whether multiple response sets associated with the active dataset are included. Numeric variables include all numeric formats except date and time formats.

Check Box

The Check Box control is a simple check box that can generate different command syntax for the checked versus the unchecked state. The Check Box control has the following properties:

Identifier. The unique identifier for the control. This is the identifier to use when referencing the control in the syntax template.

Title. The label that is displayed with the check box. For multi-line titles, use \n to specify line breaks.

ToolTip. Optional ToolTip text that appears when the user hovers over the control.

Mnemonic Key. An optional character in the title to use as a keyboard shortcut to the control. The character appears underlined in the title. The shortcut is activated by pressing Alt+[mnemonic key]. The Mnemonic Key property is not supported on Mac.

Default Value. The default state of the check box—checked or unchecked.

Checked/Unchecked Syntax. Specifies the command syntax that is generated when the control is checked and when it is unchecked. To include the command syntax in the syntax template, use the value of the Identifier property. The generated syntax, whether from the Checked Syntax or Unchecked Syntax property, will be inserted at the specified position(s) of the identifier. For example, if the identifier is *checkbox1*, then at run-time, instances of %%checkbox1%% in the

syntax template will be replaced by the value of the Checked Syntax property when the box is checked and the Unchecked Syntax property when the box is unchecked.

- You can specify any valid command syntax and you can use \n for line breaks.

Combo Box

The Combo Box control allows you to create a drop-down list that can generate command syntax specific to the selected list item. The Combo Box control has the following properties:

Identifier. The unique identifier for the control. This is the identifier to use when referencing the control in the syntax template.

Title. An optional title that appears above the control. For multi-line titles, use \n to specify line breaks.

ToolTip. Optional ToolTip text that appears when the user hovers over the control.

List Items. Click the ellipsis (...) button to open the List Item Properties dialog box, which allows you to specify the list items of the combo box control.

Mnemonic Key. An optional character in the title to use as a keyboard shortcut to the control. The character appears underlined in the title. The shortcut is activated by pressing Alt+[mnemonic key]. The Mnemonic Key property is not supported on Mac.

Syntax. Specifies command syntax that is generated by this control at run-time and can be inserted in the syntax template.

- The value `%%ThisValue%%` specifies the run-time value of the control, which is the value of the Syntax property for the selected list item. This is the default.
- You can specify any valid command syntax and you can use \n for line breaks.

Specifying List Items for Combo Boxes

The List Item Properties dialog box allows you to specify the list items of a combo box control.

Identifier. A unique identifier for the list item.

Name. The name that appears in the drop-down list for this list item. The name is a required field.

Default. Specifies whether the list item is the default item displayed in the combo box.

Syntax. Specifies the command syntax that is generated when the list item is selected.

- You can specify any valid command syntax and you can use \n for line breaks.

You can add a new list item in the blank line at the bottom of the existing list. Entering any of the properties other than the identifier will generate a unique identifier, which you can keep or modify. You can delete a list item by clicking on the *Identifier* cell for the item and pressing delete.

Text Control

The Text control is a simple text box that can accept arbitrary input, and has the following properties:

Identifier. The unique identifier for the control. This is the identifier to use when referencing the control in the syntax template.

Title. An optional title that appears above the control. For multi-line titles, use \n to specify line breaks.

ToolTip. Optional ToolTip text that appears when the user hovers over the control.

Mnemonic Key. An optional character in the title to use as a keyboard shortcut to the control. The character appears underlined in the title. The shortcut is activated by pressing Alt+[mnemonic key]. The Mnemonic Key property is not supported on Mac.

Text Content. Specifies whether the contents are arbitrary or whether the text box must contain a string that adheres to rules for SPSS Statistics variable names.

Default Value. The default contents of the text box.

Required for execution. Specifies whether a value is required in this control in order for execution to proceed. If True is specified, the OK and Paste buttons will be disabled until a value is specified for this control. If False is specified, the absence of a value in this control has no effect on the state of the OK and Paste buttons. The default is False.

Syntax. Specifies command syntax that is generated by this control at run-time and can be inserted in the syntax template.

- You can specify any valid command syntax and you can use \n for line breaks.
- The value `%%ThisValue%%` specifies the run-time value of the control, which is the content of the text box. This is the default.
- If the Syntax property includes `%%ThisValue%%` and the run-time value of the text box is empty, then the text box control does not generate any command syntax.

Number Control

The Number control is a text box for entering a numeric value, and has the following properties:

Identifier. The unique identifier for the control. This is the identifier to use when referencing the control in the syntax template.

Title. An optional title that appears above the control. For multi-line titles, use \n to specify line breaks.

ToolTip. Optional ToolTip text that appears when the user hovers over the control.

Mnemonic Key. An optional character in the title to use as a keyboard shortcut to the control. The character appears underlined in the title. The shortcut is activated by pressing Alt+[mnemonic key]. The Mnemonic Key property is not supported on Mac.

Numeric Type. Specifies any limitations on what can be entered. A value of Real specifies that there are no restrictions on the entered value, other than it be numeric. A value of Integer specifies that the value must be an integer.

Default Value. The default value, if any.

Minimum Value. The minimum allowed value, if any.

Maximum Value. The maximum allowed value, if any.

Required for execution. Specifies whether a value is required in this control in order for execution to proceed. If True is specified, the OK and Paste buttons will be disabled until a value is specified for this control. If False is specified, the absence of a value in this control has no effect on the state of the OK and Paste buttons. The default is False.

Syntax. Specifies command syntax that is generated by this control at run-time and can be inserted in the syntax template.

- You can specify any valid command syntax and you can use \n for line breaks.
- The value %%ThisValue%% specifies the run-time value of the control, which is the numeric value. This is the default.
- If the Syntax property includes %%ThisValue%% and the run-time value of the number control is empty, then the number control does not generate any command syntax.

Static Text Control

The Static Text control allows you to add a block of text to your dialog, and has the following properties:

Identifier. The unique identifier for the control.

Title. The content of the text block. For multi-line content, use \n to specify line breaks.

Item Group

The Item Group control is a container for other controls, allowing you to group and control the syntax generated from multiple controls. For example, you have a set of check boxes that specify optional settings for a subcommand, but only want to generate the syntax for the subcommand if at least one box is checked. This is accomplished by using an Item Group control as a container for the check box controls. The following types of controls can be contained in an Item Group: check box, combo box, text control, number control, static text, radio group, and file browser. The Item Group control has the following properties:

Identifier. The unique identifier for the control. This is the identifier to use when referencing the control in the syntax template.

Title. An optional title for the group. For multi-line titles, use \n to specify line breaks.

Required for execution. True specifies that the OK and Paste buttons will be disabled until at least one control in the group generates command syntax. If False is specified, the absence of run-time syntax from this control has no effect on the state of the OK and Paste buttons. The default is False.

For example, the group consists of a set of check boxes, each of which only generates syntax in the checked state. If Required for execution is set to True and all of the boxes are unchecked, then OK and Paste will be disabled.

Syntax. Specifies command syntax that is generated by this control at run-time and can be inserted in the syntax template.

- You can specify any valid command syntax and you can use \n for line breaks.
- You can include identifiers for any controls contained in the item group. At run-time the identifiers are replaced with the syntax generated by the controls.
- The value %%ThisValue%% generates a blank-separated list of the syntax generated by each control in the item group, in the order in which they appear in the group (top to bottom). This is the default. If %%ThisValue%% is specified and no syntax is generated by any of the controls in the item group, then the item group as a whole does not generate any command syntax.

Radio Group

The Radio Group control is a container for a set of radio buttons, and has the following properties:

Identifier. The unique identifier for the control. This is the identifier to use when referencing the control in the syntax template.

Title. An optional title for the group. If omitted, the group border is not displayed. For multi-line titles, use \n to specify line breaks.

ToolTip. Optional ToolTip text that appears when the user hovers over the control.

Radio Buttons. Click the ellipsis (...) button to open the Radio Group Properties dialog box, which allows you to specify the properties of the radio buttons as well as to add or remove buttons from the group.

Syntax. Specifies command syntax that is generated by this control at run-time and can be inserted in the syntax template.

- You can specify any valid command syntax and you can use \n for line breaks.
- The value %%ThisValue%% specifies the run-time value of the radio button group, which is the value of the Syntax property for the selected radio button. This is the default.

Defining Radio Buttons

The Radio Button Group Properties dialog box allows you to specify a group of radio buttons.

Identifier. A unique identifier for the radio button.

Name. The name that appears next to the radio button. The name is a required field.

ToolTip. Optional ToolTip text that appears when the user hovers over the control.

Mnemonic Key. An optional character in the name to use as a mnemonic. The specified character must exist in the name.

Default. Specifies whether the radio button is the default selection.

Syntax. Specifies the command syntax that is generated when the radio button is selected.

- You can specify any valid command syntax and you can use \n for line breaks.

You can add a new radio button in the blank line at the bottom of the existing list. Entering any of the properties other than the identifier will generate a unique identifier, which you can keep or modify. You can delete a radio button by clicking on the *Identifier* cell for the button and pressing delete.

Check Box Group

The Check Box Group control is a container for a set of controls that are enabled or disabled as a group, by a single check box. The following types of controls can be contained in a Check Box Group: check box, combo box, text control, number control, static text, radio group, and file browser. The Check Box Group control has the following properties:

Identifier. The unique identifier for the control. This is the identifier to use when referencing the control in the syntax template.

Title. An optional title for the group. If omitted, the group border is not displayed. For multi-line titles, use \n to specify line breaks.

Checkbox Title. An optional label that is displayed with the controlling check box. Supports \n to specify line breaks.

ToolTip. Optional ToolTip text that appears when the user hovers over the control.

Mnemonic Key. An optional character in the title to use as a keyboard shortcut to the control. The character appears underlined in the title. The shortcut is activated by pressing Alt+[mnemonic key]. The Mnemonic Key property is not supported on Mac.

Default Value. The default state of the controlling check box—checked or unchecked.

Checked/Unchecked Syntax. Specifies the command syntax that is generated when the control is checked and when it is unchecked. To include the command syntax in the syntax template, use the value of the Identifier property. The generated syntax, whether from the Checked Syntax or Unchecked Syntax property, will be inserted at the specified position(s) of the identifier. For example, if the identifier is *checkboxgroup1*, then at run-time, instances of `%%checkboxgroup1%%` in the syntax template will be replaced by the value of the Checked Syntax property when the box is checked and the Unchecked Syntax property when the box is unchecked.

- You can specify any valid command syntax and you can use \n for line breaks.
- You can include identifiers for any controls contained in the check box group. At run-time the identifiers are replaced with the syntax generated by the controls.
- The value `%%ThisValue%%` can be used in either the Checked Syntax or Unchecked Syntax property. It generates a blank-separated list of the syntax generated by each control in the check box group, in the order in which they appear in the group (top to bottom).
- By default, the Checked Syntax property has a value of `%%ThisValue%%` and the Unchecked Syntax property is blank.

File Browser

The File Browser control consists of a text box for a file path and a browse button that opens a standard SPSS Statistics dialog to open or save a file. The File Browser control has the following properties:

Identifier. The unique identifier for the control. This is the identifier to use when referencing the control in the syntax template.

Title. An optional title that appears above the control. For multi-line titles, use \n to specify line breaks.

ToolTip. Optional ToolTip text that appears when the user hovers over the control.

Mnemonic Key. An optional character in the title to use as a keyboard shortcut to the control. The character appears underlined in the title. The shortcut is activated by pressing Alt+[mnemonic key]. The Mnemonic Key property is not supported on Mac.

File System Operation. Specifies whether the dialog launched by the browse button is appropriate for opening files or for saving files. A value of Open indicates that the browse dialog validates the existence of the specified file. A value of Save indicates that the browse dialog does not validate the existence of the specified file.

Browser Type. Specifies whether the browse dialog is used to select a file (Locate File) or to select a folder (Locate Folder).

File Filter. Click the ellipsis (...) button to open the File Filter dialog box, which allows you to specify the available file types for the open or save dialog. By default, all file types are allowed.

File System Type. In distributed analysis mode, this specifies whether the open or save dialog browses the file system on which SPSS Statistics Server is running or the file system of your local computer. Select Server to browse the file system of the server or Client to browse the file system of your local computer. The property has no effect in local analysis mode.

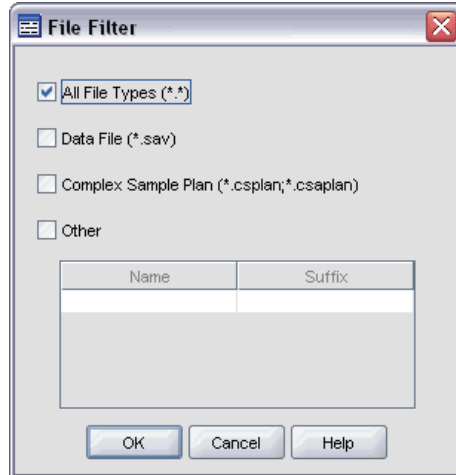
Required for execution. Specifies whether a value is required in this control in order for execution to proceed. If True is specified, the OK and Paste buttons will be disabled until a value is specified for this control. If False is specified, the absence of a value in this control has no effect on the state of the OK and Paste buttons. The default is False.

Syntax. Specifies command syntax that is generated by this control at run-time and can be inserted in the syntax template.

- You can specify any valid command syntax and you can use \n for line breaks.
- The value %%ThisValue%% specifies the run-time value of the text box, which is the file path, specified manually or populated by the browse dialog. This is the default.
- If the Syntax property includes %%ThisValue%% and the run-time value of the text box is empty, then the file browser control does not generate any command syntax.

File Type Filter

Figure 50-4
File Filter dialog box



The File Filter dialog box allows you to specify the file types displayed in the Files of type and Save as type drop-down lists for open and save dialogs accessed from a File System Browser control. By default, all file types are allowed.

To specify file types not explicitly listed in the dialog box:

- ▶ Select Other.
- ▶ Enter a name for the file type.
- ▶ Enter a file type using the form *.suffix—for example, *.xls. You can specify multiple file types, each separated by a semicolon.

Sub-dialog Button

The Sub-dialog Button control specifies a button for launching a sub-dialog and provides access to the Dialog Builder for the sub-dialog. The Sub-dialog Button has the following properties:

Identifier. The unique identifier for the control.

Title. The text that is displayed in the button.

ToolTip. Optional ToolTip text that appears when the user hovers over the control.

Sub-dialog. Click the ellipsis (...) button to open the Custom Dialog Builder for the sub-dialog. You can also open the builder by double-clicking on the Sub-dialog button.

Mnemonic Key. An optional character in the title to use as a keyboard shortcut to the control. The character appears underlined in the title. The shortcut is activated by pressing Alt+[mnemonic key]. The Mnemonic Key property is not supported on Mac.

Note: The Sub-dialog Button control cannot be added to a sub-dialog.

Dialog Properties for a Sub-dialog

To view and set properties for a sub-dialog:

- ▶ Open the sub-dialog by double-clicking on the button for the sub-dialog in the main dialog, or single-click the sub-dialog button and click the ellipsis (...) button for the Sub-dialog property.
- ▶ In the sub-dialog, click on the canvas in an area outside of any controls. With no controls on the canvas, the properties for a sub-dialog are always visible.

Sub-dialog Name. The unique identifier for the sub-dialog. The Sub-dialog Name property is required.

Note: If you specify the Sub-dialog Name as an identifier in the Syntax Template—as in `%%My Sub-dialog Name%%`—it will be replaced at run-time with a blank-separated list of the syntax generated by each control in the sub-dialog, in the order in which they appear (top to bottom and left to right).

Title. Specifies the text to be displayed in the title bar of the sub-dialog box. The Title property is optional but recommended.

Help File. Specifies the path to an optional help file for the sub-dialog. This is the file that will be launched when the user clicks the Help button on the sub-dialog, and may be the same help file specified for the main dialog. Help files must be in HTML format. See the description of the Help File property for [Dialog Properties](#) for more information.

Syntax. Click the ellipsis (...) button to open the Syntax Template. [For more information, see Building the Syntax Template on p. 525.](#)

Custom Dialogs for Extension Commands

Extension commands are user-defined SPSS Statistics commands that are implemented in either the Python programming language or R. Once deployed to an instance of SPSS Statistics, an extension command is run in the same manner as any built-in SPSS Statistics command.

You can use the Custom Dialog Builder to create dialogs for extension commands, and you can install custom dialogs for extension commands created by other users.

Creating Custom Dialogs for Extension Commands

Whether an extension command was written by you or another user, you can create a custom dialog for it. The syntax template for the dialog should generate the command syntax for the extension command. If the custom dialog is only for your use, install the dialog. Assuming the extension command is already deployed on your system, you will be able to run the command from the installed dialog.

If you are creating a custom dialog for an extension command and wish to share it with other users, you should save the specifications for the dialog to a custom dialog package (*.spd*) file and distribute that file. The custom dialog package file will not contain the specifications for the extension command itself, only the specifications for the custom dialog. You will separately need to supply end users with the XML file that specifies the syntax of the extension command and the implementation code written in Python or R.

Installing Extension Commands with Associated Custom Dialogs

An extension command with an associated custom dialog consists of three pieces: an XML file that specifies the syntax of the command; a code file (a Python module or R package) that implements the command; and a custom dialog package file that contains the specifications for the custom dialog.

Custom Dialog Package File. Install the custom dialog package file from the Install Custom Dialog item on the Utilities menu.

XML Syntax Specification File and Implementation Code. The XML file specifying the syntax of the extension command and the implementation code (Python module or R package) should both be placed in the *extensions* directory under the SPSS Statistics installation directory. For Mac, the installation directory refers to the *Contents* directory in the SPSS Statistics application bundle.

- If you do not have write permissions to the SPSS Statistics installation directory or would like to store the XML file and the implementation code elsewhere, you can specify one or more alternate locations by defining the `SPSS_EXTENSIONS_PATH` environment variable. When present, the paths specified in `SPSS_EXTENSIONS_PATH` take precedence over the SPSS Statistics installation directory. For multiple locations, separate each with a semicolon on Windows and a colon on Linux and Mac.
- For an extension command implemented in Python, you can always store the associated Python module to a location on the Python search path, such as the Python *site-packages* directory.
- For an extension command implemented in R, the R package containing the implementation code should be installed to the directory containing the XML syntax specification file. Alternatively, you can always install it to the default location for the associated platform—for example, `R_Home/library` on Windows, where `R_Home` is the installation location of R and `library` is a subdirectory under that location. For help with installing R packages, consult the R Installation and Administration guide, distributed with R.

For information on creating extension commands, see the article “Writing SPSS Statistics Extension Commands”, available from Developer Central at <http://www.spss.com/devcentral>. You may also want to consult the chapter on Extension Commands in *Programming and Data Management for SPSS Statistics*, available in PDF from http://www.spss.com/spss/data_management_book.htm.

Creating Localized Versions of Custom Dialogs

You can create localized versions of custom dialogs for any of the languages supported by SPSS Statistics. You can localize any string that appears in a custom dialog and you can localize the optional help file.

To Localize Dialog Strings

- ▶ Make a copy of the properties file associated with the dialog. The properties file contains all of the localizable strings associated with the dialog. By default, it is located in the `ext/lib/<Dialog Name>` folder of the SPSS Statistics installation directory—located at the root of the installation directory for Windows and Linux, and under the *Contents* directory in the SPSS Statistics application bundle for Mac. The copy should reside in that same folder and not in a sub-folder.

If you have specified alternate locations for installed dialogs—using the `SPSS_CDIALOGS_PATH` environment variable—then the copy should reside in the `<Dialog Name>` folder at the appropriate alternate location. For more information, see [Managing Custom Dialogs on p. 528](#).

- ▶ Rename the copy to `<Dialog Name>_<language identifier>.properties`, using the language identifiers in the table below. For example, if the dialog name is `mydialog` and you want to create a Japanese version of the dialog, then the localized properties file should be named `mydialog_ja.properties`. Localized properties files must be manually added to any custom dialog package files you create for the dialog. A custom dialog package file is simply a ZIP file that can be opened and modified with an application such as WinZip on Windows.
- ▶ Open the new properties file with a text editor that supports UTF-8, such as Notepad on Windows, or the TextEdit application on Mac. Modify the values associated with any properties that need to be localized, but do not modify the names of the properties. Properties associated with a specific control are prefixed with the identifier for the control. For example, the ToolTip property for a control with the identifier `options_button` is `options_button_tooltip_LABEL`. Title properties are simply named `<identifier>_LABEL`, as in `options_button_LABEL`. Colons used in property values should be escaped, as in `\ :`.

When the dialog is launched, SPSS Statistics will search for a properties file whose language identifier matches the current language, as specified by the Language drop-down on the General tab in the Options dialog box. If no such properties file is found, the default file `<Dialog Name>.properties` will be used.

To Localize the Help File

- ▶ Make a copy of the help file associated with the dialog and localize the text for the desired language. The copy must reside in the same folder as the help file and not in a sub-folder. The help file should reside in the `ext/lib/<Dialog Name>` folder of the SPSS Statistics installation directory—located at the root of the installation directory for Windows and Linux, and under the `Contents` directory in the SPSS Statistics application bundle for Mac.

If you have specified alternate locations for installed dialogs—using the `SPSS_CDIALOGS_PATH` environment variable—then the copy should reside in the `<Dialog Name>` folder at the appropriate alternate location. For more information, see [Managing Custom Dialogs on p. 528](#).

- ▶ Rename the copy to `<Help File>_<language identifier>`, using the language identifiers in the table below. For example, if the help file is `myhelp.htm` and you want to create a German version of the file, then the localized help file should be named `myhelp_de.htm`. Localized help files must be manually added to any custom dialog package files you create for the dialog. A custom dialog package file is simply a ZIP file that can be opened and modified with an application such as WinZip on Windows.

If there are supplementary files such as image files that also need to be localized, you will need to manually modify the appropriate paths in the main help file to point to the localized versions, which should be stored along with the original versions.

When the dialog is launched, SPSS Statistics will search for a help file whose language identifier matches the current language, as specified by the Language drop-down on the General tab in the Options dialog box. If no such help file is found, the help file specified for the dialog (the file specified in the Help File property of Dialog Properties) is used.

Language Identifiers

de	German
en	English
es	Spanish
fr	French
it	Italian
ja	Japanese
ko	Korean
pl	Polish
ru	Russian
zh_cn	Simplified Chinese
zh_tw	Traditional Chinese

Note: Text in custom dialogs and associated help files is not limited to the languages supported by SPSS Statistics. You are free to write the dialog and help text in any language without creating language-specific properties and help files. All users of your dialog will then see the text in that language.

Production Jobs

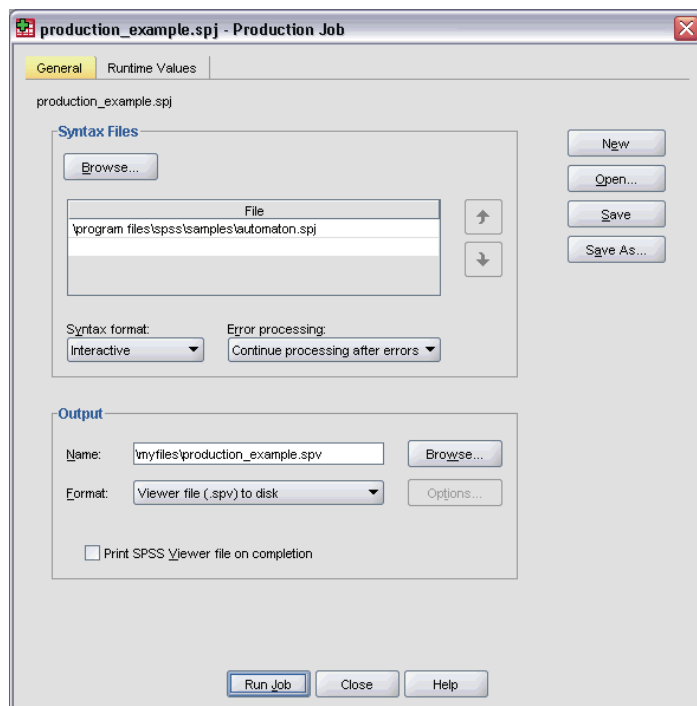
Production jobs provide the ability to run SPSS Statistics in an automated fashion. The program runs unattended and terminates after executing the last command, so you can perform other tasks while it runs or schedule the production job to run automatically at scheduled times. Production jobs are useful if you often run the same set of time-consuming analyses, such as weekly reports.

Production jobs use command syntax files to tell SPSS Statistics what to do. A command syntax file is a simple text file containing command syntax. You can use any text editor to create the file. You can also generate command syntax by pasting dialog box selections into a syntax window or by editing the journal file. [For more information, see Working with Command Syntax in Chapter 13 on p. 253.](#)

To create a production job:

- ▶ From the menus in any window choose:
 - Utilities
 - Production Job

Figure 51-1
Production Job dialog box



- ▶ Click **New** to create a new production job or **Open** to open an existing production job.

Note: Production Facility job files (.spp) created in releases prior to 16.0 will not run in release 16.0 or later. A conversion utility is available to convert Windows and Macintosh Production Facility job files to production jobs (.spj). [For more information, see *Converting Production Facility Files* on p. 551.](#)

- ▶ Select one or more command syntax files.

Note: In distributed analysis mode, syntax file locations must be entered manually and relative paths are relative to the remote server. [For more information, see *Absolute versus Relative Path Specifications* in Chapter 4 on p. 69.](#)

- ▶ Select the output file name, location, and format.
- ▶ Click **Save** or **Save As** to save the production job.

Syntax format. Controls the form of the syntax rules used for the job.

- **Interactive.** Each command must end with a period. Periods can appear anywhere within the command, and commands can continue on multiple lines, but a period as the last nonblank character on a line is interpreted as the end of the command. Continuation lines and new commands can start anywhere on a new line. These are the “interactive” rules in effect when you select and run commands in a syntax window.
- **Batch.** Each command must start at the beginning of a new line (no blank spaces before the start of the command), and continuation lines must be indented at least one space. If you want to indent new commands, you can use a plus sign, dash, or period as the first character at the start of the line and then indent the actual command. The period at the end of the command is optional. This setting is compatible with the syntax rules for command files included with the `INCLUDE` command.

Note: Do *not* use the Batch option if your syntax files contain `GGRAPH` command syntax that includes `GPL` statements. `GPL` statements will only run under interactive rules.

Error Processing. Controls the treatment of error conditions in the job.

- **Continue processing after errors.** Errors in the job do not automatically stop command processing. The commands in the production job files are treated as part of the normal command stream, and command processing continues in the normal fashion.
- **Stop processing immediately.** Command processing stops when the first error in a production job file is encountered. This is compatible with the behavior of command files included with the `INCLUDE` command.

Output. Controls the name, location, and format of the production job results. The following format options are available:

- **Viewer file to disk (.spv).** Results are saved in SPSS Statistics Viewer format in the specified file location.
- **Viewer file (.spv) to PES Repository.** This requires the Adaptor for Enterprise Services option.
- **Web Reports (.spw) to PES Repository.** This requires the Adaptor for Enterprise Services option.
- **Word/RTF (*.doc).** Pivot tables are exported as Word tables with all formatting attributes intact—for example, cell borders, font styles, and background colors. Text output is exported as formatted RTF. Charts, tree diagrams, and model views are included in PNG format.

Note: Microsoft Word may not display extremely wide tables properly.

- **Excel (*.xls).** Pivot table rows, columns, and cells are exported as Excel rows, columns, and cells, with all formatting attributes intact—for example, cell borders, font styles, and background colors. Text output is exported with all font attributes intact. Each line in the text output is a row in the Excel file, with the entire contents of the line contained in a single cell. Charts, tree diagrams, and model views are included in PNG format.
- **HTML (*.htm).** Pivot tables are exported as HTML tables. Text output is exported as preformatted HTML. Charts, tree diagrams, and model views are embedded by reference, and you should export charts in a suitable format for inclusion in HTML documents (for example, PNG and JPEG).
- **Portable Document Format (*.pdf).** All output is exported as it appears in Print Preview, with all formatting attributes intact.
- **PowerPoint file (*.ppt).** Pivot tables are exported as Word tables and are embedded on separate slides in the PowerPoint file, with one slide for each pivot table. All formatting attributes of the pivot table are retained—for example, cell borders, font styles, and background colors. Charts, tree diagrams, and model views are exported in TIFF format. Text output is not included.
Export to PowerPoint is available only on Windows operating systems.
- **Text (*.txt).** Text output formats include plain text, UTF-8, and UTF-16. Pivot tables can be exported in tab-separated or space-separated format. All text output is exported in space-separated format. For charts, tree diagrams, and model views, a line is inserted in the text file for each graphic, indicating the image filename.

Print Viewer file on completion. Sends the final Viewer output file to the printer on completion of the production job.

Run Job. This runs the production job in a separate session. This is particularly useful for testing new production jobs before deploying them.

Production Jobs with OUTPUT Commands

Production jobs honor `OUTPUT` commands, such as `OUTPUT SAVE`, `OUTPUT ACTIVATE`, and `OUTPUT NEW`. `OUTPUT SAVE` commands executed during the course of a production job will write the contents of the specified output documents to the specified locations. This is in addition to the output file created by the production job. When using `OUTPUT NEW` to create a new output document, it is recommended that you explicitly save it with the `OUTPUT SAVE` command.

A production job output file consists of the contents of the active output document as of the end of the job. For jobs containing `OUTPUT` commands, the output file may not contain all output created in the session. For example, suppose the production job consists of a number of procedures followed by an `OUTPUT NEW` command, followed by more procedures but no more `OUTPUT` commands. The `OUTPUT NEW` command defines a new active output document. At the end of the production job, it will contain output from only the procedures executed after the `OUTPUT NEW` command.

HTML Options

Table Options. No table options are available for HTML format. All pivot tables are converted to HTML tables.

Image Options. The available image types are: EPS, JPEG, TIFF, PNG, and BMP. On Windows operating systems, EMF (enhanced metafile) format is also available. You can also scale the image size from 1% to 200%.

PowerPoint Options

Table Options. You can use the Viewer outline entries as slide titles. Each slide contains a single output item. The title is formed from the outline entry for the item in the outline pane of the Viewer.

Image Options. You can scale the image size from 1% to 200%. (All images are exported to PowerPoint in TIFF format.)

Note: PowerPoint format is only available on Windows operating systems and requires PowerPoint 97 or later.

PDF Options

Embed bookmarks. This option includes bookmarks in the PDF document that correspond to the Viewer outline entries. Like the Viewer outline pane, bookmarks can make it much easier to navigate documents with a large number of output objects.

Embed fonts. Embedding fonts ensures that the PDF document will look the same on all computers. Otherwise, if some fonts used in the document are not available on the computer being used to view (or print) the PDF document, font substitution may yield suboptimal results.

Text Options

Table Options. Pivot tables can be exported in tab-separated or space-separated format. For space-separated format, you can also control:

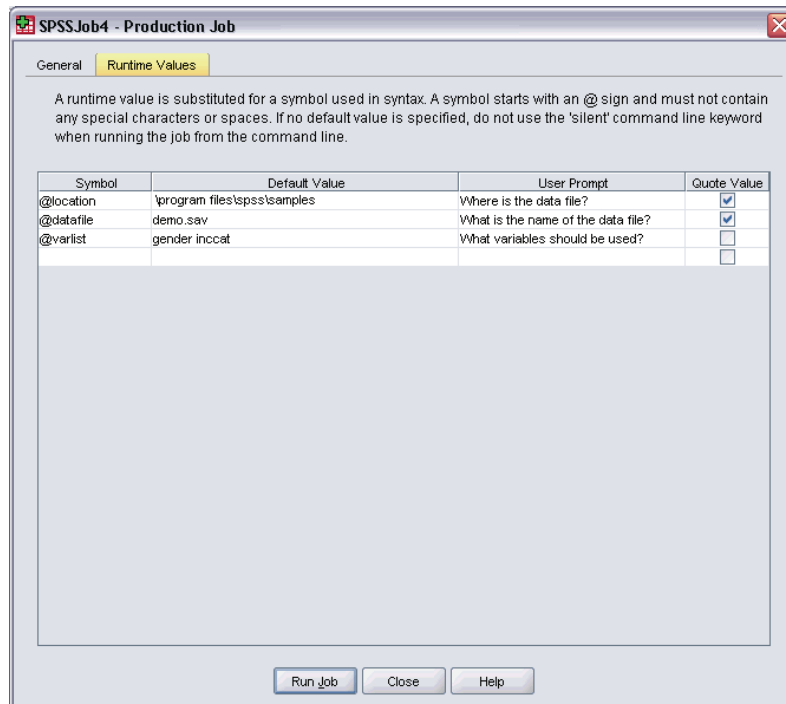
- **Column Width.** Autofit does not wrap any column contents, and each column is as wide as the widest label or value in that column. Custom sets a maximum column width that is applied to all columns in the table, and values that exceed that width wrap onto the next line in that column.
- **Row/Column Border Character.** Controls the characters used to create row and column borders. To suppress display of row and column borders, enter blank spaces for the values.

Image Options. The available image types are: EPS, JPEG, TIFF, PNG, and BMP. On Windows operating systems, EMF (enhanced metafile) format is also available. You can also scale the image size from 1% to 200%.

Runtime Values

Runtime values defined in a production job file and used in a command syntax file simplify tasks such as running the same analysis for different data files or running the same set of commands for different sets of variables. For example, you could define the runtime value `@datafile` to prompt you for a data filename each time you run a production job that uses the string `@datafile` in place of a filename in the command syntax file.

Figure 51-2
Production Job dialog box, Runtime Values tab



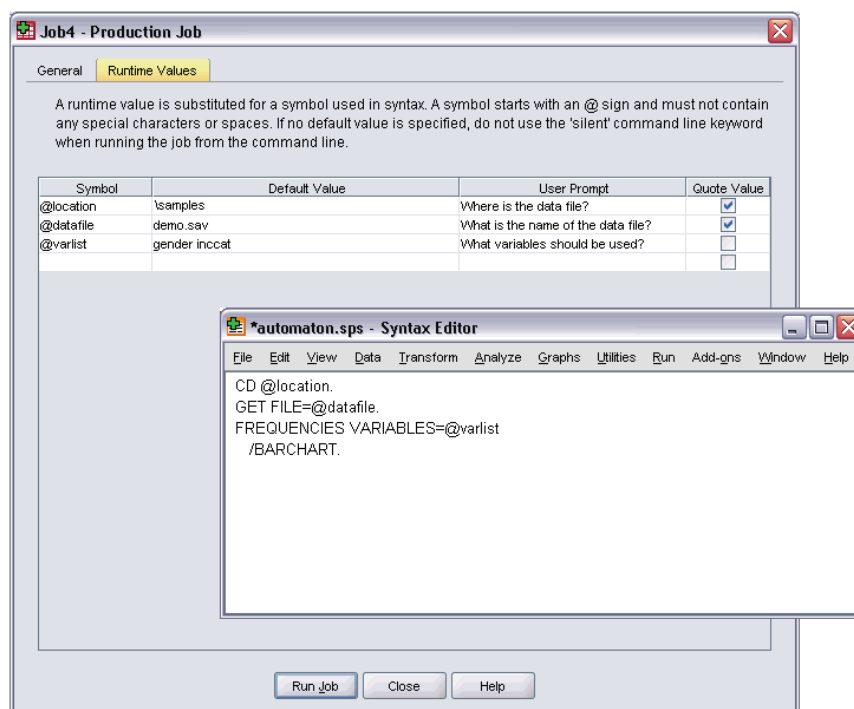
Symbol. The string in the command syntax file that triggers the production job to prompt the user for a value. The symbol name must begin with an `@` sign and must conform to variable naming rules. For more information, see [Variable Names in Chapter 5 on p. 72](#).

Default Value. The value that the production job supplies by default if you don't enter a different value. This value is displayed when the production job prompts you for information. You can replace or modify the value at runtime. If you don't provide a default value, don't use the `silent` keyword when running the production job with command line switches, unless you also use the `-symbol` switch to specify runtime values. For more information, see [Running Production Jobs from a Command Line on p. 549](#).

User Prompt. The descriptive label that is displayed when the production job prompts you to enter information. For example, you could use the phrase "What data file do you want to use?" to identify a field that requires a data filename.

Quote Value. Encloses the default value or the value entered by the user in quotes. For example, file specifications should be enclosed in quotes.


Figure 51-3
Runtime symbols in a command syntax file



User Prompts

A production job prompts you for values whenever you run a production job that contains defined runtime symbols. You can replace or modify the default values that are displayed. Those values are then substituted for the runtime symbols in all command syntax files associated with the production job.

Figure 51-4
Production User Prompts dialog box



Running Production Jobs from a Command Line

Command line switches enable you to schedule production jobs to run automatically at certain times, using scheduling utilities available on your operating system. The basic form of the command line argument is:

```
statistics filename.spj -production
```

Depending on how you invoke the production job, you may need to include directory paths for the *statistics* executable file (located in the directory in which the application is installed) and/or the production job file.

You can run production jobs from a command line with the following switches:

-production [prompt|silent]. Start the application in production mode. The `prompt` and `silent` keywords specify whether to display the dialog box that prompts for runtime values if they are specified in the job. The `prompt` keyword is the default and shows the dialog box. The `silent` keyword suppresses the dialog box. If you use the `silent` keyword, you can define the runtime symbols with the `-symbol` switch. Otherwise, the default value is used. The `-switchserver` and `-singleseat` switches are ignored when using the `-production` switch.

-symbol <values>. List of symbol-value pairs used in the production job. Each symbol name starts with `@`. Values that contain spaces should be enclosed in quotes. Rules for including quotes or apostrophes in string literals may vary across operating systems, but enclosing a string that includes single quotes or apostrophes in double quotes usually works (for example, `'a quoted value'`).

To run production jobs on a remote server in distributed analysis mode, you also need to specify the server login information:

-server <inet:hostname:port>. The name or IP address and port number of the server. Windows only.

-user <name>. A valid user name. If a domain name is required, precede the user name with the domain name and a backslash (`\`). Windows only.

-password <password>. The user's password. Windows only.

Example

```
statistics \production_jobs\prodjob1.spj -production silent -symbol @datafile /data/July_data.sav
```

- This example assumes that you are running the command line from the installation directory, so no path is required for the *statistics* executable file.
- This example also assumes that the production job specifies that the value for `@datafile` should be quoted (Quote Value checkbox on the Runtime Values tab), so no quotes are necessary when specifying the data file on the command line. Otherwise, you would need to specify something like `"/data/July_data.sav"` to include quotes with the data file specification, since file specifications should be quoted in command syntax.
- The directory path for the location of the production job uses the Windows back slash convention. On Macintosh and Linux, use forward slashes. The forward slashes in the quoted data file specification will work on all operating systems since this quoted string is inserted into the command syntax file and forward slashes are acceptable in commands that include file specifications (for example, `GET FILE`, `GET DATA`, `SAVE`) on all operating systems.
- The `silent` keyword suppresses any user prompts in the production job, and the `-symbol` switch inserts the quoted data file name and location wherever the runtime symbol `@datafile` appears in the command syntax files included in the production job.

Converting Production Facility Files

Production Facility job files (.spp) created in releases prior to 16.0 will not run in release 16.0 or later. For Windows and Macintosh Production Facility job files created in earlier releases, you can use *prodconvert*, located in the installation directory, to convert those files to new production job files (.spj). Run *prodconvert* from a command window using the following specifications:

```
[installpath]\prodconvert [filepath]\filename.spp
```

where *[installpath]* is the location of the folder in which SPSS Statistics is installed and *[filepath]* is the location of the folder where the original production job file. A new file with the same name but with the extension *.spj* is created in the same folder as the original file. (Note: If the path contains spaces, enclose each path and file specification in double quotes. On Macintosh operating systems, use forward slashes instead on back slashes.)

Limitations

- WMF and EMF chart formats are not supported. PNG format is used in place of these formats.
- The export options Output Document (No Charts), Charts Only, and Nothing are not supported. All output objects supported by the selected format are included.
- Remote server settings are ignored. To specify remote server settings for distributed analysis, you need to run the production job from a command line, using command line switches to specify the server settings. [For more information, see Running Production Jobs from a Command Line on p. 549.](#)
- Publish to Web settings are ignored.

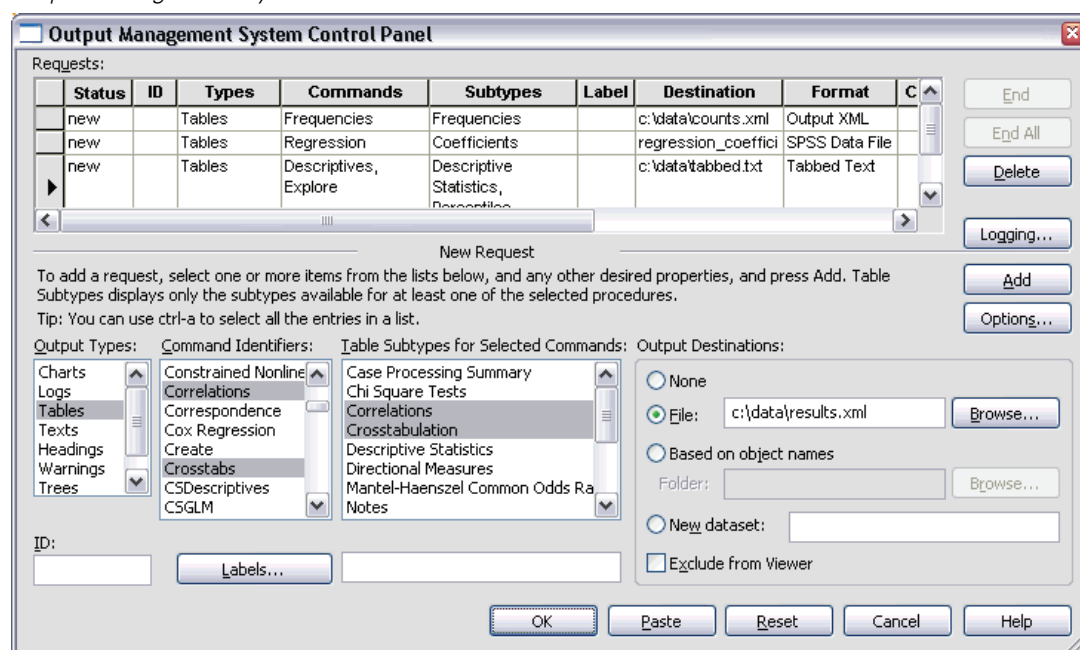
Output Management System

The Output Management System (OMS) provides the ability to automatically write selected categories of output to different output files in different formats. Formats include: Word, Excel, PDF, SPSS Statistics data file format (.sav), Viewer file format (.spv), web report format (.spw), XML, HTML, and text. [For more information, see OMS Options on p. 558.](#)

To Use the Output Management System Control Panel

- From the menu choose:
 - Utilities
 - OMS Control Panel...

Figure 52-1
Output Management System Control Panel



You can use the control panel to start and stop the routing of output to various destinations.

- Each OMS request remains active until explicitly ended or until the end of the session.
- A destination file that is specified on an OMS request is unavailable to other procedures and other applications until the OMS request is ended.

- While an OMS request is active, the specified destination files are stored in memory (RAM), so active OMS requests that write a large amount of output to external files may consume a large amount of memory.
- Multiple OMS requests are independent of each other. The same output can be routed to different locations in different formats, based on the specifications in different OMS requests.
- The order of the output objects in any particular destination is the order in which they were created, which is determined by the order and operation of the procedures that generate the output.

To Add New OMS Requests

- ▶ Select the output types (tables, charts, etc.) that you want to include. [For more information, see Output Object Types on p. 554.](#)
- ▶ Select the commands to include. If you want to include all output, select all items in the list. [For more information, see Command Identifiers and Table Subtypes on p. 556.](#)
- ▶ For commands that produce pivot table output, select the specific table types to include.
The list displays only the tables that are available in the selected commands; any table type that is available in one or more of the selected commands is displayed in the list. If no commands are selected, all table types are displayed. [For more information, see Command Identifiers and Table Subtypes on p. 556.](#)
- ▶ To select tables based on text labels instead of subtypes, click Labels. [For more information, see Labels on p. 557.](#)
- ▶ Click Options to specify the output format (for example, SPSS Statistics data file, XML, or HTML). By default, Output XML format is used. [For more information, see OMS Options on p. 558.](#)
- ▶ Specify an output destination:
 - **File.** All selected output is routed to a single file.
 - **Based on object names.** Output is routed to multiple destination files based on object names. A separate file is created for each output object, with a filename based on either table subtype names or table labels. Enter the destination folder name.
 - **New dataset.** For SPSS Statistics data file format output, you can route the output to a dataset. The dataset is available for subsequent use in the same session but is not saved unless you explicitly save it as a file prior to the end of the session. This option is available only for SPSS Statistics data file format output. Dataset names must conform to variable-naming rules. [For more information, see Variable Names in Chapter 5 on p. 72.](#)
- ▶ Optionally:
 - Exclude the selected output from the Viewer. If you select Exclude from Viewer, the output types in the OMS request will not be displayed in the Viewer window. If multiple active OMS requests include the same output types, the display of those output types in the Viewer

is determined by the most recent OMS request that contains those output types. [For more information, see Excluding Output Display from the Viewer on p. 563.](#)

- Assign an ID string to the request. All requests are automatically assigned an ID value, and you can override the system default ID string with a descriptive ID, which can be useful if you have multiple active requests that you want to identify easily. ID values that you assign cannot start with a dollar sign (\$).

The following tips are for selecting multiple items in a list:

- Press Ctrl+A to select all items in a list.
- Use Shift+click to select multiple contiguous items.
- Use Ctrl+click to select multiple noncontiguous items.

To End and Delete OMS Requests

Active and new OMS requests are displayed in the Requests list, with the most recent request at the top. You can change the widths of the information columns by clicking and dragging the borders, and you can scroll the list horizontally to see more information about a particular request.

An asterisk (*) after the word Active in the *Status* column indicates an OMS request that was created with command syntax that includes features that are not available in the Control Panel.

To end a specific, active OMS request:

- ▶ In the Requests list, click any cell in the row for the request.
- ▶ Click End.

To end all active OMS requests:

- ▶ Click End All.

To delete a new request (a request that has been added but is not yet active):

- ▶ In the Requests list, click any cell in the row for the request.
- ▶ Click Delete.

Note: Active OMS requests are not ended until you click OK.

Output Object Types

There are different types of output objects:

Charts. This includes charts created with the Chart Builder, charting procedures, and charts created by statistical procedures (for example, a bar chart created by the Frequencies procedure).

Headings. Text objects that are labeled *Title* in the outline pane of the Viewer. Heading text objects are not included with Output XML format.

Logs. Log text objects. Log objects contain certain types of error and warning messages. Depending on your Options settings (Edit menu, Options, Viewer tab), log objects may also contain the command syntax that is executed during the session. Log objects are labeled *Log* in the outline pane of the Viewer.

Models. Output objects displayed in the Model Viewer. A single model object can contain multiple views of the model, including both tables and charts.

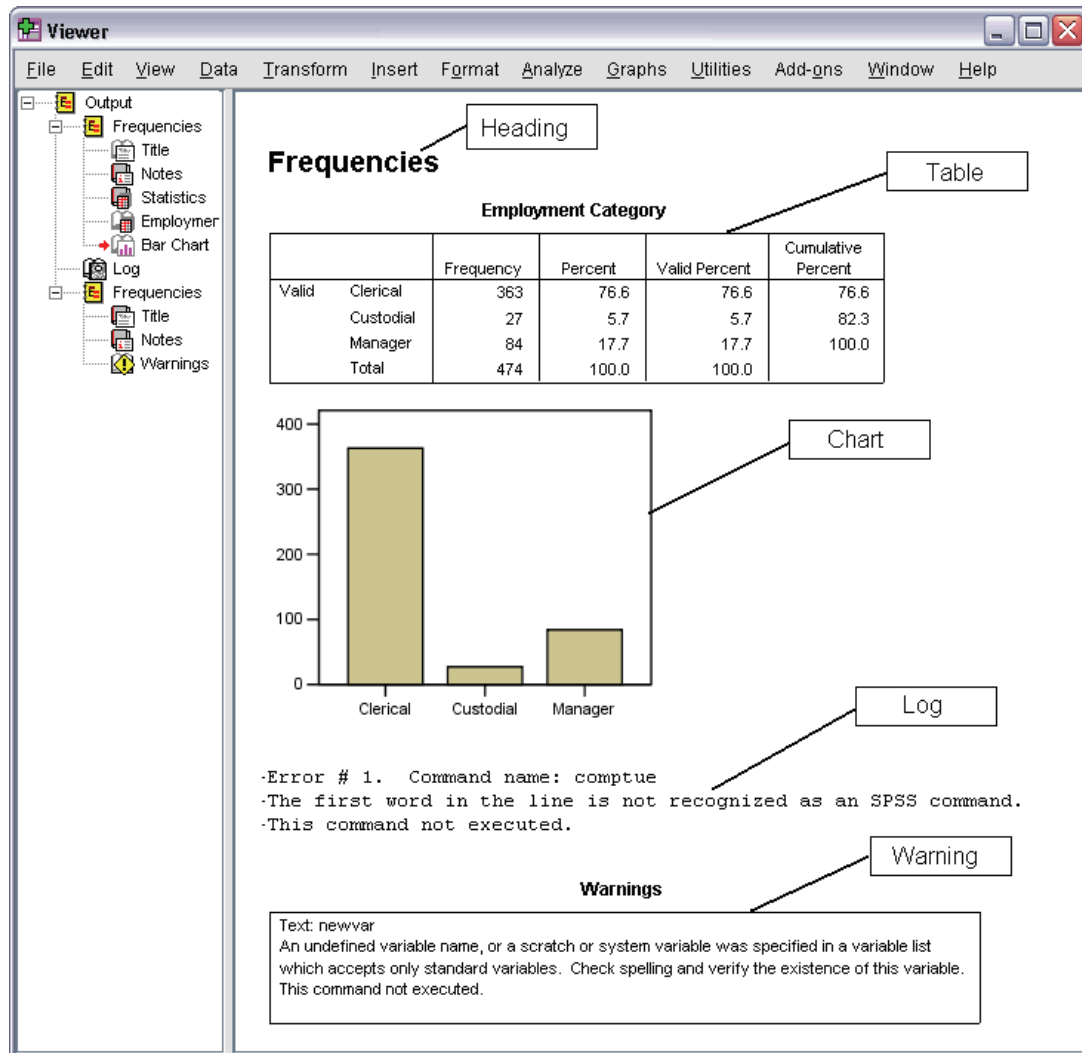
Tables. Output objects that are pivot tables in the Viewer (includes Notes tables). Tables are the only output objects that can be routed to SPSS Statistics data file (.sav) format.

Texts. Text objects that aren't logs or headings (includes objects labeled *Text Output* in the outline pane of the Viewer).

Trees. Tree model diagrams that are produced by the Decision Tree option.

Warnings. Warning objects contain certain types of error and warning messages.

Figure 52-2
Output object types



Command Identifiers and Table Subtypes

Command Identifiers

Command identifiers are available for all statistical and charting procedures and any other commands that produce blocks of output with their own identifiable heading in the outline pane of the Viewer. These identifiers are usually (but not always) the same or similar to the procedure names on the menus and dialog box titles, which are usually (but not always) similar to the underlying command names. For example, the command identifier for the Frequencies procedure is "Frequencies," and the underlying command name is also the same.

There are, however, some cases where the procedure name and the command identifier and/or the command name are not all that similar. For example, all of the procedures on the Nonparametric Tests submenu (from the Analyze menu) use the same underlying command, and the command identifier is the same as the underlying command name: Npar Tests.

Table Subtypes

Table subtypes are the different types of pivot tables that can be produced. Some subtypes are produced by only one command; other subtypes can be produced by multiple commands (although the tables may not look similar). Although table subtype names are generally descriptive, there can be many names to choose from (particularly if you have selected a large number of commands); also, two subtypes may have very similar names.

To Find Command Identifiers and Table Subtypes

When in doubt, you can find command identifiers and table subtype names in the Viewer window:

- ▶ Run the procedure to generate some output in the Viewer.
- ▶ Right-click the item in the outline pane of the Viewer.
- ▶ Choose Copy OMS Command Identifier or Copy OMS Table Subtype.
- ▶ Paste the copied command identifier or table subtype name into any text editor (such as a Syntax Editor window).

Labels

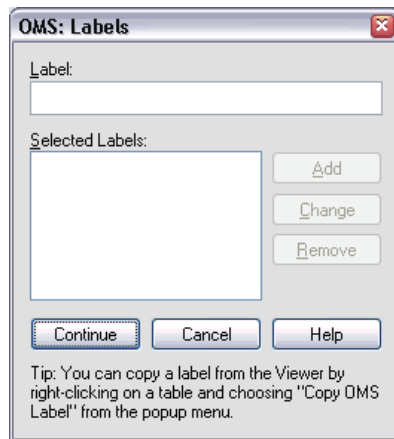
As an alternative to table subtype names, you can select tables based on the text that is displayed in the outline pane of the Viewer. You can also select other object types based on their labels. Labels are useful for differentiating between multiple tables of the same type in which the outline text reflects some attribute of the particular output object, such as the variable names or labels. There are, however, a number of factors that can affect the label text:

- If split-file processing is on, split-file group identification may be appended to the label.
- Labels that include information about variables or values are affected by your current output label options settings (Edit menu, Options, Output Labels tab).
- Labels are affected by the current output language setting (Edit menu, Options, General tab).

To Specify Labels to Use to Identify Output Objects

- ▶ In the Output Management System Control Panel, select one or more output types and then select one or more commands.
- ▶ Click Labels.

Figure 52-3
OMS Labels dialog box



- ▶ Enter the label *exactly* as it appears in the outline pane of the Viewer window. (You can also right-click the item in the outline, choose Copy OMS Label, and paste the copied label into the Label text field.)
- ▶ Click Add.
- ▶ Repeat the process for each label that you want to include.
- ▶ Click Continue.

Wildcards

You can use an asterisk (*) as the last character of the label string as a wildcard character. All labels that begin with the specified string (except for the asterisk) will be selected. This process works only when the asterisk is the last character, because asterisks can appear as valid characters inside a label.

OMS Options

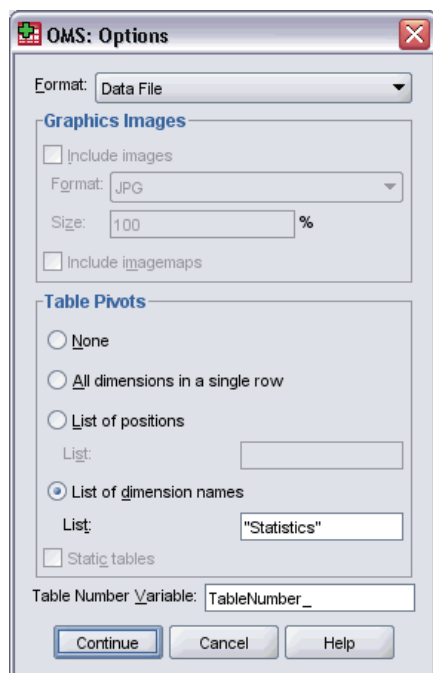
You can use the OMS Options dialog box to:

- Specify the output format.
- Specify the image format (for HTML and Output XML output formats).
- Specify what table dimension elements should go into the row dimension.
- Include a variable that identifies the sequential table number that is the source for each case (for SPSS Statistics data file format).

To Specify OMS Options

- ▶ Click Options in the Output Management System Control Panel.

Figure 52-4
OMS Options dialog box



Format

Excel. Excel 97-2003 format. Pivot table rows, columns, and cells are exported as Excel rows, columns, and cells, with all formatting attributes intact — for example, cell borders, font styles, and background colors. Text output is exported with all font attributes intact. Each line in the text output is a row in the Excel file, with the entire contents of the line contained in a single cell. Charts, tree diagrams, and model views are included in PNG format.

HTML. Output objects that would be pivot tables in the Viewer are converted to simple HTML tables. No TableLook attributes (font characteristics, border styles, colors, etc.) are supported. Text output objects are tagged <PRE> in the HTML. Charts, tree diagrams, and model views are exported as separate files in the selected graphics format and are embedded by reference. Image file names use the HTML file name as the root name, followed by a sequential integer, starting with 0.

Output XML. XML that conforms to the *spss-output* schema.

PDF. Output is exported as it would appear in Print Preview, with all formatting attributes intact. The PDF file includes bookmarks that correspond to the entries in the Viewer outline pane.

SPSS Statistics Data File. This format is a binary file format. All output object types other than tables are excluded. Each column of a table becomes a variable in the data file. To use a data file that is created with OMS in the same session, you must end the active OMS request before you can open the data file. [For more information, see Routing Output to SPSS Statistics Data Files on p. 564.](#)

Text. Space-separated text. Output is written as text, with tabular output aligned with spaces for fixed-pitch fonts. Charts, tree diagrams, and model views are excluded.

Tabbed Text. Tab-delimited text. For output that is displayed as pivot tables in the Viewer, tabs delimit table column elements. Text block lines are written as is; no attempt is made to divide them with tabs at useful places. Charts, tree diagrams, and model views are excluded.

Viewer File. This is the same format used when you save the contents of a Viewer window.

Web Report File. This output file format is designed for use with Predictive Enterprise Services. It is essentially the same as the SPSS Statistics Viewer format except that tree diagrams are saved as static images.

Word/RTF. Pivot tables are exported as Word tables with all formatting attributes intact—for example, cell borders, font styles, and background colors. Text output is exported as formatted RTF. Charts, tree diagrams, and model views are included in PNG format.

Graphics Images

For HTML and Output XML formats, you can include charts, tree diagrams, and model views as image files. A separate image file is created for each chart and/or tree.

- For HTML document format, standard `` tags are included in the HTML document for each image file.
- For Output XML document format, the XML file contains a `chart` element with an `imageFile` attribute of the general form `<chart imageFile="filepath/filename"/>` for each image file.
- Image files are saved in a separate subdirectory (folder). The subdirectory name is the name of the destination file, without any extension and with `_files` appended to the end. For example, if the destination file is `julydata.htm`, the images subdirectory will be named `julydata_files`.

Format. The available image formats are PNG, JPG, EMF, BMP, and VML.

- EMF (enhanced metafile) format is available only on Windows operating systems.
- VML image format is available only for HTML document format.
- VML image format does not create separate image files. The VML code that renders the image is embedded in the HTML.
- VML image format does not include tree diagrams.

Size. You can scale the image size from 10% to 200%.

Include Imagemaps. For HTML document format, this option creates image map ToolTips that display information for some chart elements, such as the value of the selected point on a line chart or bar on a bar chart.

Table Pivots

For pivot table output, you can specify the dimension element(s) that should appear in the columns. All other dimension elements appear in the rows. For SPSS Statistics data file format, table columns become variables, and rows become cases.

- If you specify multiple dimension elements for the columns, they are nested in the columns in the order in which they are listed. For SPSS Statistics data file format, variable names are constructed by nested column elements. [For more information, see Variable Names in OMS-Generated Data Files on p. 571.](#)
- If a table doesn't contain any of the listed dimension elements, all dimension elements for that table will appear in the rows.
- Table pivots that are specified here have no effect on tables that are displayed in the Viewer.

Each dimension of a table—row, column, layer—may contain zero or more elements. For example, a simple two-dimensional crosstabulation contains a single row dimension element and a single column dimension element, each of which contains one of the variables that are used in the table. You can use either positional arguments or dimension element “names” to specify the dimension elements that you want to put into the column dimension.

All dimensions in rows. Creates a single row for each table. For SPSS Statistics format data files, this means each table is a single case, and all the table elements are variables.

List of positions. The general form of a positional argument is a letter indicating the default position of the element—C for column, R for row, or L for layer—followed by a positive integer indicating the default position within that dimension. For example, R1 would indicate the outermost row dimension element.

- To specify multiple elements from multiple dimensions, separate each dimension with a space—for example, R1 C2.
- The dimension letter followed by ALL indicates all elements in that dimension in their default order. For example, CALL is the same as the default behavior (using all column elements in their default order to create columns).
- CALL RALL LALL (or RALL CALL LALL, and so on) will put all dimension elements into the columns. For SPSS Statistics data file format, this creates one row/case per table in the data file.

Figure 52-5
Row and column positional arguments

Model	Correlations	Previous Experience (months)	Months since Hire	Beginning Salary	Date of Birth
1		1.000	.067	-.087	.805
		.067	1.000	.012	.085
		-.087	.012	1.000	-.075
		.805	.085	-.075	1.000
	Covariances	Previous Experience (months)	12.940	-0.22	7.096E-06
		Months since Hire	1205.248	.019	4.635E-06
		Beginning Salary	.019	.002	-5.236E-09
		Date of Birth	4.635E-06	-5.236E-09	2.485E-12

List of dimension names. As an alternative to positional arguments, you can use dimension element “names,” which are the text labels that appear in the table. For example, a simple two-dimensional crosstabulation contains a single row dimension element and a single column dimension element,

each with labels based on the variables in those dimensions, plus a single layer dimension element labeled *Statistics* (if English is the output language).

- Dimension element names may vary, based on the output language and/or settings that affect the display of variable names and/or labels in tables.
- Each dimension element name must be enclosed in single or double quotation marks. To specify multiple dimension element names, include a space between each quoted name.

The labels that are associated with the dimension elements may not always be obvious.

To See All Dimension Elements and Their Labels for a Pivot Table

- ▶ Activate (double-click) the table in the Viewer.
 - ▶ From the menus choose:
 - View
 - Show All
- and/or*
- ▶ If the pivoting trays aren't displayed, from the menus choose:
 - Pivot
 - Pivoting Trays

The element labels are displayed in the pivoting trays.

Figure 52-6
Dimension element names displayed in table and pivoting trays

Table without Dimension Element Labels Displayed

		Current Salary	Beginning Salary
Female	Mean	\$26,031.92	\$13,091.97
	Median	\$24,300.00	\$12,375.00
Male	Mean	\$41,441.78	\$20,301.40
	Median	\$32,850.00	\$15,750.00
Total	Mean	\$34,419.57	\$17,016.09
	Median	\$28,875.00	\$15,000.00

Table with Dimension Element Labels Displayed

		Variables	
Gender	Statistics	Current Salary	Beginning Salary
Female	Mean	\$26,031.92	\$13,091.97
	Median	\$24,300.00	\$12,375.00
Male	Mean	\$41,441.78	\$20,301.40
	Median	\$32,850.00	\$15,750.00
Total	Mean	\$34,419.57	\$17,016.09
	Median	\$28,875.00	\$15,000.00

Logging

You can record OMS activity in a log in XML or text format.

- The log tracks all new OMS requests for the session but does not include OMS requests that were already active before you requested a log.
- The current log file ends if you specify a new log file or if you deselect (uncheck) Log OMS activity.

To specify OMS logging:

- ▶ Click Logging in the Output Management System Control Panel.

Excluding Output Display from the Viewer

The Exclude from Viewer check box affects all output that is selected in the OMS request by suppressing the display of that output in the Viewer window. This process is often useful for production jobs that generate a lot of output and when you don't need the results in the form of a Viewer document (.spv file). You can also use this functionality to suppress the display of

particular output objects that you simply never want to see, without routing any other output to some external file and format.

To suppress the display of certain output objects without routing other output to an external file:

- ▶ Create an OMS request that identifies the unwanted output.
- ▶ Select Exclude from Viewer.
- ▶ For the output destination, select File—but leave the File field blank.
- ▶ Click Add.

The selected output will be excluded from the Viewer while all other output will be displayed in the Viewer in the normal fashion.

Routing Output to SPSS Statistics Data Files

A data file in SPSS Statistics format consists of variables in the columns and cases in the rows, which is essentially the format in which pivot tables are converted to data files:

- Columns in the table are variables in the data file. Valid variable names are constructed from the column labels.
- Row labels in the table become variables with generic variable names (*Var1*, *Var2*, *Var3*, and so on) in the data file. The values of these variables are the row labels in the table.
- Three table-identifier variables are automatically included in the data file: *Command_*, *Subtype_*, and *Label_*. All three are string variables. The first two variables correspond to the command and subtype identifiers. [For more information, see Command Identifiers and Table Subtypes on p. 556.](#) *Label_* contains the table title text.
- Rows in the table become cases in the data file.

Example: Single Two-Dimensional Table

In the simplest case—a single two-dimensional table—the table columns become variables, and the rows become cases in the data file.

Figure 52-7
Single two-dimensional table

Report

Gender		Current Salary	Beginning Salary
Female	Mean	\$26,031.92	\$13,091.97
	Median	\$24,300.00	\$12,375.00
	Minimum	\$15,750	\$9,000
	Maximum	\$58,125	\$30,000
Male	Mean	\$41,441.78	\$20,301.40
Total			

	Command_	Subtype_	Label_	Var1	Var2	CurrentSalary	BeginningSal
1	Means	Report	Report	Female	Mean	\$26,031.92	\$13,091.97
2	Means	Report	Report	Female	Median	\$24,300.00	\$12,375.00
3	Means	Report	Report	Female	Minimum	\$15,750	\$9,000
4	Means	Report	Report	Female	Maximum	\$58,125	\$30,000
5	Means	Report	Report	Male	Mean	\$41,441.78	\$20,301.40
6	Means	Report	Report	Male	Median	\$32,850.00	\$15,750.00
7	Means	Report	Report	Male	Minimum	\$19,650	\$9,000
8	Means	Report	Report	Male	Maximum	\$135,000	\$79,980
9	Means	Report	Report	Total	Mean	\$34,419.57	\$17,016.09
10	Means	Report	Report	Total	Median	\$28,875.00	\$15,000.00

- The first three variables identify the source table by command, subtype, and label.
- The two elements that defined the rows in the table—values of the variable *Gender* and statistical measures—are assigned the generic variable names *Var1* and *Var2*. These variables are both string variables.
- The column labels from the table are used to create valid variable names. In this case, those variable names are based on the variable labels of the three scale variables that are summarized in the table. If the variables didn't have defined variable labels, or you chose to display variable names instead of variable labels as the column labels in the table, the variable names in the new data file would be the same as in the source data file.

Example: Tables with Layers

In addition to rows and columns, a table can also contain a third dimension: the layer dimension.

Figure 52-8
Table with layers

Layer	Minority Classification Yes	Gender	
		Female	Male
Employment Category	Clerical	40	47
	Custodial		
	Manager		

Layer	Minority Classification No	Gender	
		Female	Male
Employment Category	Clerical	166	110
	Custodial	0	14
	Manager	10	70

	Var1	Var2	Var3	Var4	Female	Male
1	Employment Category	Clerical	Minority Classification	No	166	110
2	Employment Category	Clerical	Minority Classification	Yes	40	47
3	Employment Category	Custodial	Minority Classification	No	0	14
4	Employment Category	Custodial	Minority Classification	Yes	0	13
5	Employment Category	Manager	Minority Classification	No	10	70
6	Employment Category	Manager	Minority Classification	Yes	0	4

- In the table, the variable labeled *Minority Classification* defines the layers. In the data file, two additional variables are created: one variable that identifies the layer element and one variable that identifies the categories of the layer element.
- As with the variables that are created from the row elements, the variables that are created from the layer elements are string variables with generic variable names (the prefix *Var* followed by a sequential number).

Data Files Created from Multiple Tables

When multiple tables are routed to the same data file, each table is added to the data file in a fashion that is similar to merging data files by adding cases from one data file to another data file (Data menu, Merge Files, Add Cases).

- Each subsequent table will always add cases to the data file.
- If column labels in the tables differ, each table may also add variables to the data file, with missing values for cases from other tables that don't have an identically labeled column.

Example: Multiple Tables with the Same Column Labels

Multiple tables that contain the same column labels will typically produce the most immediately useful data files (data files that don't require additional manipulation). For example, two or more frequency tables from the Frequencies procedure will all have identical column labels.

Figure 52-9
Two tables with identical column labels

Gender					
		Frequency	Percent	Valid Percent	Cumulative Percent
Valid	Female	216	45.6	45.6	45.6
	Male	258	54.4	54.4	100.0

Employment Category					
		Frequency	Percent	Valid Percent	Cumulative Percent
Valid	Clerical	363	76.6	76.6	76.6
	Custodial	27	5.7	5.7	82.3
	Manager	84	17.7	17.7	100.0

Command_	Subtype_	Label_	Var1	Var2	Frequency	Percent	ValidPercent	CumulativePercent
1	Frequencies	Gender	Valid	Female	216	45.6	45.6	45.6
2	Frequencies	Gender	Valid	Male	258	54.4	54.4	100.0
3	Frequencies	Gender	Valid	Total	474	100.0	100.0	.
4	Frequencies	Employment Categor	Valid	Clerical	363	76.6	76.6	76.6
5	Frequencies	Employment Categor	Valid	Custodial	27	5.7	5.7	82.3
6	Frequencies	Employment Categor	Valid	Manager	84	17.7	17.7	100.0
7	Frequencies	Employment Categor	Valid	Total	474	100.0	100.0	.
8								
9								

- The second table contributes additional cases (rows) to the data file but contributes no new variables because the column labels are exactly the same, so there are no large patches of missing data.
- Although the values for *Command_* and *Subtype_* are the same, the *Label_* value identifies the source table for each group of cases because the two frequency tables have different titles.

Example: Multiple Tables with Different Column Labels

A new variable is created in the data file for each unique column label in the tables that are routed to the data file. This process results in blocks of missing values if the tables contain different column labels.

Figure 52-10
Two tables with different column labels

Table 1			
Gender		Beginning Salary	Current Salary
Female	Mean	\$13,091.97	\$26,031.92
	Median	\$12,375.00	\$24,300.00
Male	Mean	\$20,301.40	\$41,441.78

Table 2			
Gender		Education level	Months since Hire
Female	Mean	12.37	80.38
	Median	12.00	81.00
Male	Mean	14.43	81.72

The screenshot shows the SPSS Data Editor window for a file named *temp.sav [DataSet5]. The 'Means' table is displayed, showing the results of a merge operation. The table has columns for Subtype, Label, Var1, Var2, Current Salary, Beginning Salary, Educational Level (years), and Monthssince Hire. The rows represent different combinations of gender and salary/education metrics. Missing values (indicated by dots) are present for variables that were not in the source table for that specific case.

	Subtype_	Label_	Var1	Var2	Current Salary	Beginning Salary	Educational Level (years)	Monthssince Hire
1	Report	Report	Female	Mean	\$26032	\$13092.0	.	.
2	Report	Report	Female	Median	\$24300	\$12375.0	.	.
3	Report	Report	Male	Mean	\$41442	\$20301.4	.	.
4	Report	Report	Male	Median	\$32850	\$15750.0	.	.
5	Report	Report	Total	Mean	\$34420	\$17016.1	.	.
6	Report	Report	Total	Median	\$28875	\$15000.0	.	.
7	Report	Report	Female	Mean	.	.	12.37	80.38
8	Report	Report	Female	Median	.	.	12.00	81.00
9	Report	Report	Male	Mean	.	.	14.43	81.72
10	Report	Report	Male	Median	.	.	15.00	82.00
11	Report	Report	Total	Mean	.	.	13.49	81.11
12	Report	Report	Total	Median	.	.	12.00	81.00

- The first table has columns labeled *Beginning Salary* and *Current Salary*, which are not present in the second table, resulting in missing values for those variables for cases from the second table.
- Conversely, the second table has columns labeled *Education level* and *Months since Hire*, which are not present in the first table, resulting in missing values for those variables for cases from the first table.
- Mismatched variables like the variables in this example can occur even with tables of the same subtype. In this example, both tables are the same subtype.

Example: Data Files Not Created from Multiple Tables

If any tables do not have the same number of row elements as the other tables, no data file will be created. The number of rows doesn't have to be the same; the number of row *elements* that become variables in the data file must be the same. For example, a two-variable crosstabulation and a three-variable crosstabulation contain different numbers of row elements, because the "layer" variable is actually nested within the row variable in the default three-variable crosstabulation display.

Figure 52-11
Tables with different numbers of row elements

Employment Category * Gender Crosstabulation

		Gender		Total
		Female	Male	
Employment Category	Clerical	206	157	363
	Custodial	0	27	27
	Manager	10	74	84
Total		216	258	474

Employment Category * Gender * Minority Classification Crosstabulation

Minority Classification			Gender		Total
			Female	Male	
No	Employment Category	Clerical	166	110	276
		Custodial	0	14	14
		Manager	10	70	80
	Total	176	194	370	
Yes	Employment Category	Clerical	40	47	87
		Custodial	0	13	13
		Manager	0	4	4
	Total	40	64	104	

Controlling Column Elements to Control Variables in the Data File

In the Options dialog box of the Output Management System Control Panel, you can specify which dimension elements should be in the columns and therefore will be used to create variables in the generated data file. This process is equivalent to pivoting the table in the Viewer.

For example, the Frequencies procedure produces a descriptive statistics table with statistics in the rows, while the Descriptives procedure produces a descriptive statistics table with statistics in the columns. To include both table types in the same data file in a meaningful fashion, you need to change the column dimension for one of the table types.

Because both table types use the element name “Statistics” for the statistics dimension, we can put the statistics from the Frequencies statistics table in the columns simply by specifying “Statistics” (in quotation marks) in the list of dimension names in the OMS Options dialog box.

Figure 52-12
OMS Options dialog box

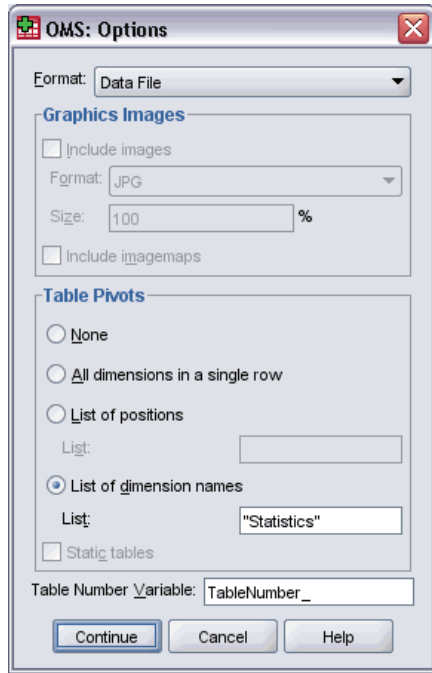


Figure 52-13
Combining different table types in a data file by pivoting dimension elements

		Beginning Salary	Current Salary
N	Valid	474	474
	Missing	0	0
Mean		\$17,016.09	\$34,419.57

	N		Mean	Minimum	Maximum
	Valid	Missing			
Beginning Salary	474	0	\$17,016.09	\$9,000	\$79,980
Current Salary					\$135,000

	N	Minimum	Maximum	Mean
Months since Hire	474	63	98	81.11
Previous Experience (months)	474	0	476	95.86

Some of the variables will have missing values, because the table structures still aren't exactly the same with statistics in the columns.

Variable Names in OMS-Generated Data Files

OMS constructs valid, unique variable names from column labels:

- Row and layer elements are assigned generic variable names—the prefix *Var* followed by a sequential number.
- Characters that aren't allowed in variable names (spaces, parentheses, etc.) are removed. For example, “This (Column) Label” would become a variable named *ThisColumnLabel*.
- If the label begins with a character that is allowed in variable names but not allowed as the first character (for example, a number), “@” is inserted as a prefix. For example, “2nd” would become a variable named *@2nd*.
- Underscores or periods at the end of labels are removed from the resulting variable names. The underscores at the end of the automatically generated variables *Command_*, *Subtype_*, and *Label_* are not removed.
- If more than one element is in the column dimension, variable names are constructed by combining category labels with underscores between category labels. Group labels are not included. For example, if *VarB* is nested under *VarA* in the columns, you would get variables like *CatA1_CatB1*, not *VarA_CatA1_VarB_CatB1*.

Figure 52-14

Variable names constructed from table elements

Layer Variable	Column Variable			Total	
	In	(Out)			
High	Row Variable	Yes	1	1	2
		No	0	1	1
	Total		1	2	3
Low	Row Variable	Yes	1	0	1
		No	1	1	2
	Total		2	1	3

Layer Variable		High			Low		
Row Variable	Yes	Column Variable		Total	Column Variable		Total
		In	(Out)		In	(Out)	
	No	1	1	2	1	0	1
		0	1	1	1	1	2
Total		1	2	3	2	1	3

OXML Table Structure

Output XML (OXML) is XML that conforms to the *spss-output* schema. For a detailed description of the schema, see the Output Schema section of the Help system.

- OMS command and subtype identifiers are used as values of the command and subType attributes in OXML. An example is as follows:

```
<command text="Frequencies" command="Frequencies"...>
<pivotTable text="Gender" label="Gender" subType="Frequencies"...>
```

- OMS command and subType attribute values are not affected by output language or display settings for variable names/labels or values/value labels.
- XML is case sensitive. A subType attribute value of “frequencies” is *not* the same as a subType attribute value of “Frequencies.”
- All information that is displayed in a table is contained in attribute values in OXML. At the individual cell level, OXML consists of “empty” elements that contain attributes but no “content” other than the content that is contained in attribute values.
- Table structure in OXML is represented row by row; elements that represent columns are nested within the rows, and individual cells are nested within the column elements:

```
<pivotTable...>
<dimension axis='row'...>
<dimension axis='column'...>
<category...>
<cell text='...' number='...' decimals='...'/>
</category>
<category...>
<cell text='...' number='...' decimals='...'/>
</category>
</dimension>
</dimension>
...
</pivotTable>
```

The preceding example is a simplified representation of the structure that shows the descendant/ancestor relationships of these elements. However, the example does not necessarily show the parent/child relationships, because there are typically intervening nested element levels.

The following figure shows a simple frequency table and the complete output XML representation of that table.

Figure 52-15
Simple frequency table

		Gender			
		Frequency	Percent	Valid Percent	Cumulative Percent
Valid	Female	216	45.6	45.6	45.6
	Male	258	54.4	54.4	100.0
Total		474	100.0	100.0	

Figure 52-16
Output XML for the simple frequency table

```

<?xml version="1.0" encoding="UTF-8" ?>
<outputTreeoutputTree xmlns="http://xml.spss.com/spss/oms"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://xml.spss.com/spss/oms
  http://xml.spss.com/spss/oms/spss-output-1.0.xsd">
<command text="Frequencies" command="Frequencies"
  displayTableValues="label" displayOutlineValues="label"
  displayTableVariables="label" displayOutlineVariables="label">
<pivotTable text="Gender" label="Gender" subType="Frequencies"
  varName="gender" variable="true">
<dimension axis="row" text="Gender" label="Gender"
  varName="gender" variable="true">
<group text="Valid">
<group hide="true" text="Dummy">
<category text="Female" label="Female" string="f"
  varName="gender">
<dimension axis="column" text="Statistics">
<category text="Frequency">
<cell text="216" number="216"/>
</category>
<category text="Percent">
<cell text="45.6" number="45.569620253165" decimals="1"/>
</category>
<category text="Valid Percent">
<cell text="45.6" number="45.569620253165" decimals="1"/>
</category>
<category text="Cumulative Percent">
<cell text="45.6" number="45.569620253165" decimals="1"/>
</category>
</dimension>
</category>
<category text="Male" label="Male" string="m" varName="gender">
<dimension axis="column" text="Statistics">
<category text="Frequency">
<cell text="258" number="258"/>
</category>
<category text="Percent">
<cell text="54.4" number="54.430379746835" decimals="1"/>
</category>
<category text="Valid Percent">
<cell text="54.4" number="54.430379746835" decimals="1"/>
</category>
<category text="Cumulative Percent">
<cell text="100.0" number="100" decimals="1"/>
</category>
</dimension>
</category>
</group>
<category text="Total">

```

```

<dimension axis="column" text="Statistics">
  <category text="Frequency">
    <cell text="474" number="474"/>
  </category>
  <category text="Percent">
    <cell text="100.0" number="100" decimals="1"/>
  </category>
  <category text="Valid Percent">
    <cell text="100.0" number="100" decimals="1"/>
  </category>
</dimension>
</category>
</group>
</dimension>
</pivotTable>
</command>
</outputTree>

```

As you may notice, a simple, small table produces a substantial amount of XML. That's partly because the XML contains some information that is not readily apparent in the original table, some information that might not even be available in the original table, and a certain amount of redundancy.

- The table contents as they are (or would be) displayed in a pivot table in the Viewer are contained in text attributes. An example is as follows:

```
<command text="Frequencies" command="Frequencies"...>
```

- Text attributes can be affected by both output language and settings that affect the display of variable names/labels and values/value labels. In this example, the text attribute value will differ, depending on the output language, whereas the command attribute value remains the same, regardless of output language.
- Wherever variables or values of variables are used in row or column labels, the XML will contain a text attribute and one or more additional attribute values. An example is as follows:

```

<dimension axis="row" text="Gender" label="Gender" varName="gender">
  ...<category text="Female" label="Female" string="f" varName="gender">

```

- For a numeric variable, there would be a number attribute instead of a string attribute. The label attribute is present only if the variable or values have defined labels.
- The <cell> elements that contain cell values for numbers will contain the text attribute and one or more additional attribute values. An example is as follows:

```
<cell text="45.6" number="45.569620253165" decimals="1"/>
```

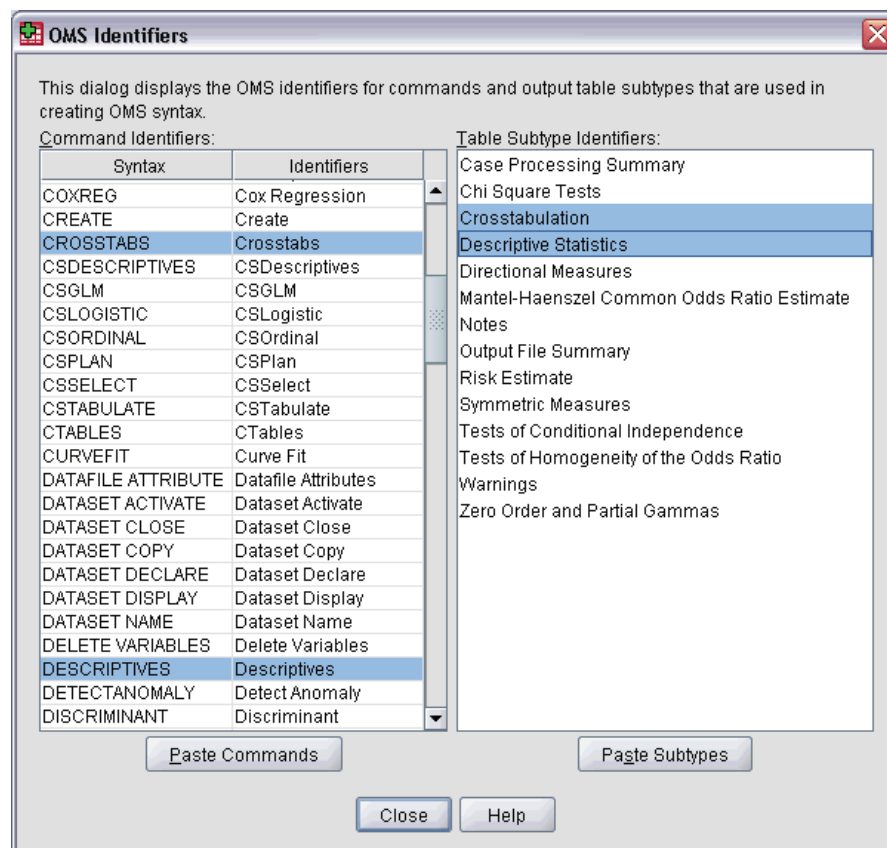
The number attribute is the actual, unrounded numeric value, and the decimals attribute indicates the number of decimal positions that are displayed in the table.

- Because columns are nested within rows, the category element that identifies each column is repeated for each row. For example, because the statistics are displayed in the columns, the element `<category text="Frequency">` appears three times in the XML: once for the male row, once for the female row, and once for the total row.

OMS Identifiers

The OMS Identifiers dialog box is designed to assist you in writing OMS command syntax. You can use this dialog box to paste selected command and subtype identifiers into a command syntax window.

Figure 52-17
OMS Identifiers dialog box



To Use the OMS Identifiers Dialog Box

- From the menu choose:
Utilities
OMS Identifiers...

- ▶ Select one or more command or subtype identifiers. (Use Ctrl+click to select multiple identifiers in each list.)
- ▶ Click Paste Commands and/or Paste Subtypes.
 - The list of available subtypes is based on the currently selected command(s). If multiple commands are selected, the list of available subtypes is the union of all subtypes that are available for any of the selected commands. If no commands are selected, all subtypes are listed.
 - The identifiers are pasted into the designated command syntax window at the current cursor location. If there are no open command syntax windows, a new syntax window is automatically opened.
 - Each command and/or subtype identifier is enclosed in quotation marks when pasted, because OMS command syntax requires these quotation marks.
 - Identifier lists for the `COMMANDS` and `SUBTYPES` keywords must be enclosed in brackets, as in:

```
/IF COMMANDS=['Crosstabs' 'Descriptives']
    SUBTYPES=['Crosstabulation' 'Descriptive Statistics']
```

Copying OMS Identifiers from the Viewer Outline

You can copy and paste OMS command and subtype identifiers from the Viewer outline pane.

- ▶ In the outline pane, right-click the outline entry for the item.
- ▶ Choose Copy OMS Command Identifier or Copy OMS Table Subtype.

This method differs from the OMS Identifiers dialog box method in one respect: The copied identifier is not automatically pasted into a command syntax window. The identifier is simply copied to the clipboard, and you can then paste it anywhere you want. Because command and subtype identifier values are identical to the corresponding command and subtype attribute values in Output XML format (OXML), you might find this copy/paste method useful if you write XSLT transformations.

Copying OMS Labels

Instead of identifiers, you can copy labels for use with the `LABELS` keyword. Labels can be used to differentiate between multiple graphs or multiple tables of the same type in which the outline text reflects some attribute of the particular output object, such as the variable names or labels. There are, however, a number of factors that can affect the label text:

- If split-file processing is on, split-file group identification may be appended to the label.
- Labels that include information about variables or values are affected by the settings for the display of variable names/labels and values/value labels in the outline pane (Edit menu, Options, Output Labels tab).
- Labels are affected by the current output language setting (Edit menu, Options, General tab).

To copy OMS labels

- ▶ In the outline pane, right-click the outline entry for the item.
- ▶ Choose Copy OMS Label.

As with command and subtype identifiers, the labels must be in quotation marks, and the entire list must be enclosed in square brackets, as in:

```
/IF LABELS=['Employment Category' 'Education Level']
```

Scripting Facility

The scripting facility allows you to automate tasks, including:

- Opening and saving data files.
- Exporting charts as graphic files in a variety of formats.
- Customizing output in the Viewer.

Scripting Languages

The available scripting languages depend on your platform. For Windows, the available scripting languages are Basic, which is installed with the Base system, and the Python programming language. For all other platforms, scripting is available with the Python programming language.

To enable scripting with the Python programming language, you must separately install Python and the SPSS Statistics-Python Integration Plug-in. The Python programming language is not installed with the Base system, but is available from the installation CD for Windows. The SPSS Statistics-Python Integration Plug-in is available from the installation CD for all platforms.

Note: SPSS Inc. is not the owner or licensor of the Python software. Any user of Python must agree to the terms of the Python license agreement located on the Python Web site. SPSS Inc. does not make any statement about the quality of the Python program. SPSS Inc. fully disclaims all liability associated with your use of the Python program.

Default Script Language

The default script language determines the script editor that is launched when new scripts are created. It also specifies the default language whose executable will be used to run autoscripts. On Windows, the default script language is Basic. You can change the default language from the Scripts tab in the Options dialog. [For more information, see Script Options in Chapter 48 on p. 507.](#)

Sample Scripts

A number of scripts are included with the software, in the *Samples* subdirectory of the directory where SPSS Statistics is installed. You can use these scripts as they are or you can customize them to your needs.

To Create a New Script

- ▶ From the menus choose:

- File
 - New Script

The editor associated with the default script language opens.

To Run a Script

- ▶ From the menus choose:

- Utilities
 - Run Script...

- ▶ Select the script you want.
- ▶ Click Run.

Python scripts can be run in a number of ways, other than from Utilities>Run Script. [For more information, see Scripting with the Python Programming Language on p. 582.](#)

To Edit a Script

- ▶ From the menus choose:

- File
 - Open Script...

- ▶ Select the script you want.
- ▶ Click Open.

The script is opened in the editor associated with the language in which the script is written.

Autoscripts

Autoscripts are scripts that run automatically when triggered by the creation of specific pieces of output from selected procedures. For example, you can use an autoscript to automatically remove the upper diagonal and highlight correlation coefficients below a certain significance whenever a Correlations table is produced by the Bivariate Correlations procedure.

Autoscripts can be specific to a given procedure and output type or apply to specific output types from different procedures. For example, you might have an autoscript that formats the ANOVA tables produced by One-Way ANOVA as well as ANOVA tables produced by other statistical procedures. On the other hand, Frequencies produces both a frequency table and a table of statistics, and you might choose to have a different autoscript for each.

Each output type for a given procedure can only be associated with a single autoscript. You can, however, create a base autoscript that is applied to all new Viewer items prior to the application of any autoscripts for specific output types. [For more information, see Script Options in Chapter 48 on p. 507.](#)

The Scripts tab in the Options dialog box (accessed from the Edit menu) displays the autoscripts that have been configured on your system and allows you to set up new autoscripts or modify the settings for existing ones. Optionally, you can create and configure autoscripts for output items directly from the Viewer.

Events that Trigger Autoscripts

The following events can trigger autoscripts:

- Creation of a pivot table
- Creation of a Notes object
- Creation of a Warnings object

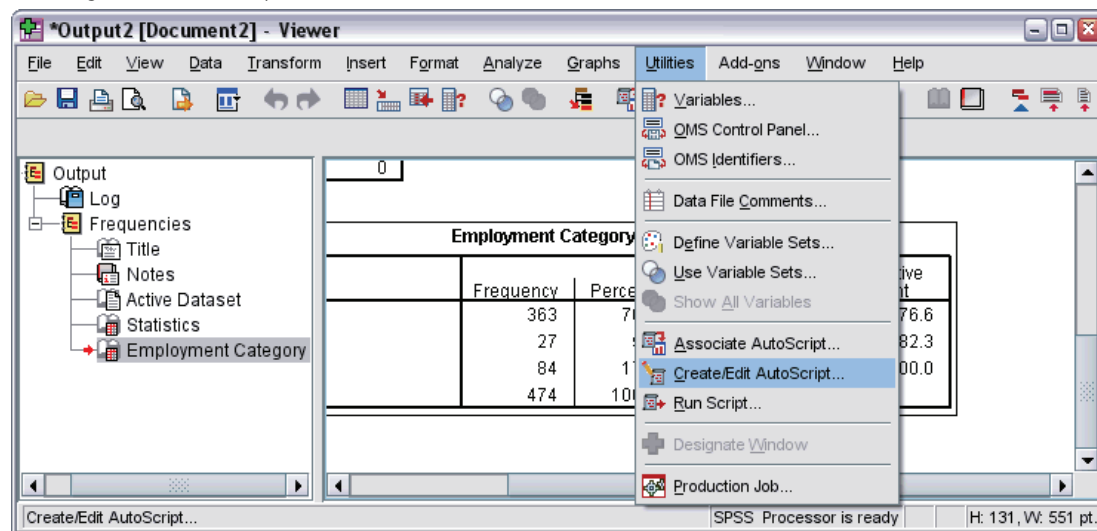
You can also use a script to trigger an autoscript indirectly. For example, you could write a script that invokes the Correlations procedure, which in turn triggers an autoscript registered to the resulting Correlations table.

Creating Autoscripts

You can create an autoscript by starting with the output object that you want to serve as the trigger—for instance, a frequency table.

- ▶ In the Viewer, select the object that will trigger the autoscript.
- ▶ From the menus choose:
 - Utilities
 - Create/Edit AutoScript...

Figure 53-1
Creating a new autoscript



If the selected object does not have an associated autoscript, an Open dialog prompts you for the location and name of a new script.

- ▶ Browse to the location where the new script will be stored, enter a file name and click Open. The editor for the default script language opens. You can change the default script language from the Scripts tab on the Options dialog. [For more information, see Script Options in Chapter 48 on p. 507.](#)
- ▶ Type the code.

For help with coding autoscripts in Python, see [Getting Started with Autoscripts in Python](#) on p. 585. For help with converting custom Sax Basic autoscripts used in pre-16.0 versions, see [Compatibility with Versions Prior to 16.0](#) on p. 589.

Note: By default, the executable associated with the default script language will be used to run the autoscript. You can change the executable from the Scripts tab in the Options dialog.

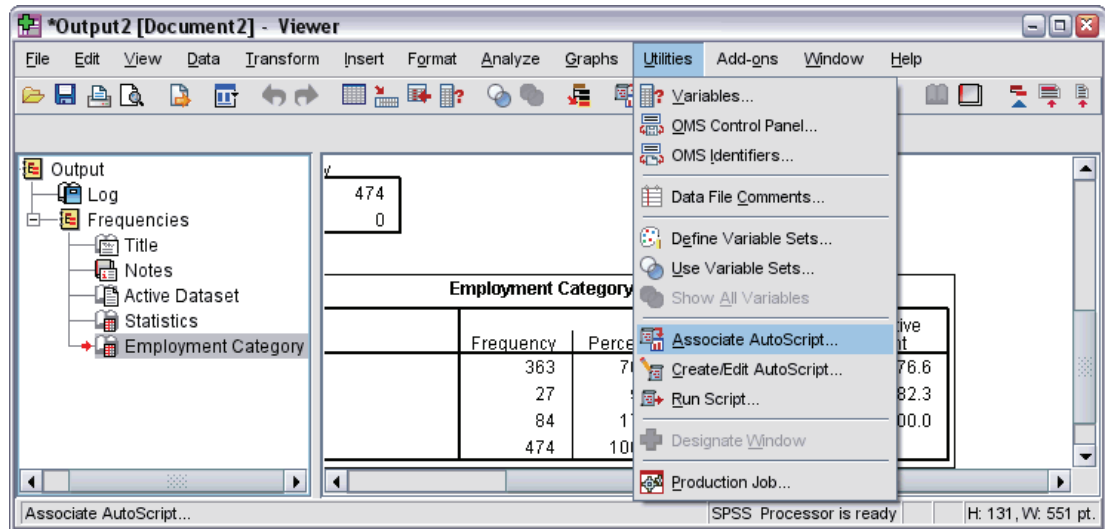
If the selected object is already associated with an autoscript, the script is opened in the script editor associated with the language in which the script is written.

Associating Existing Scripts with Viewer Objects

You can use existing scripts as autoscripts by associating them with a selected object in the Viewer—for instance, a frequency table.

- ▶ In the Viewer, select an object to associate with an autoscript (multiple Viewer objects can trigger the same autoscript, but each object can only be associated with a single autoscript).
- ▶ From the menus choose:
 - Utilities
 - Associate AutoScript...

Figure 53-2
Associating an autoscript



If the selected object does not have an associated autoscript, the Select Autoscript dialog opens.

- ▶ Browse for the script you want and select it.
- ▶ Click Apply.

If the selected object is already associated with an autoscript, you are prompted to verify that you want to change the association. Clicking OK opens the Select Autoscript dialog.

Optionally, you can configure an existing script as an autoscript from the Scripts tab in the Options dialog box. The autoscript can be applied to a selected set of output types or specified as the base autoscript that is applied to all new Viewer items. For more information, see [Script Options in Chapter 48 on p. 507](#).

Scripting with the Python Programming Language

SPSS Statistics provides two separate interfaces for programming with the Python language on Windows, Linux, and Mac OS. Use of these interfaces requires the SPSS Statistics-Python Integration Plug-in, which is available from the installation CD for all platforms. For help getting started with the Python programming language, see the Python tutorial, available at <http://docs.python.org/tut/tut.html>.

Python Scripts

Python scripts make use of the interface exposed by the Python `SpssClient` module. They operate on user interface and output objects and can also run command syntax. For instance, you would use a Python script to customize a pivot table.

- Python scripts are run from Utilities>Run Script, from the Python editor launched from SPSS Statistics (accessed from File>Open>Script), or from an external Python process, such as a Python IDE or the Python interpreter.

- Python scripts can be run as autoscripts.
- Python scripts run on the machine where the SPSS Statistics client is running.

A brief overview of working with Python scripts is provided in this help system. For complete documentation of the SPSS Statistics classes and methods available for Python scripts, see *SPSS Statistics Scripting Guide.pdf*, accessed from Help>Programmability>Scripting, and available once the SPSS Statistics-Python Integration Plug-in is installed.

Python Programs

Python programs make use of the interface exposed by the Python `spss` module. They operate on the SPSS Statistics processor and are used to control the flow of a command syntax job, read from and write to the active dataset, create new datasets, and create custom procedures that generate their own pivot table output.

- Python programs are run from command syntax within `BEGIN PROGRAM-END PROGRAM` blocks, or from an external Python process, such as a Python IDE or the Python interpreter.
- Python programs cannot be run as autoscripts.
- In distributed analysis mode (available with SPSS Statistics Server), Python programs execute on the computer where SPSS Statistics Server is running.

For information on working with Python programs, see *SPSS Statistics-Python Integration package.pdf*, accessed from Help>Programmability>Python Plug-in, and available once the SPSS Statistics-Python Integration Plug-in is installed.

Running Python Scripts and Python programs

Both Python scripts and Python programs can be run from within SPSS Statistics or from an external Python process, such as a Python IDE or the Python interpreter.

Python Scripts

Python Script Run from SPSS Statistics. You can run a Python script from Utilities>Run Script or from the Python script editor which is launched when opening a Python file (.py) from File>Open>Script. Scripts run from the Python editor that is launched from SPSS Statistics operate on the SPSS Statistics client that launched the editor. This allows you to debug your Python code from a Python editor.

Python Script Run from an External Python Process. You can run a Python script from any external Python process, such as a Python IDE that is not launched from SPSS Statistics, or the Python interpreter. The script will attempt to connect to an existing SPSS Statistics client. If more than one client is found, a connection is made to the most recently launched one. If an existing client is not found, the Python script starts up a new instance of the SPSS Statistics client. By default, the Data Editor and Viewer are invisible for the new client. You can choose to make them visible or work in invisible mode with datasets and output documents.

Python Programs

Python Program Run from Command Syntax. You can run a Python program by embedding Python code within a `BEGIN PROGRAM-END PROGRAM` block in command syntax. The command syntax can be run from the SPSS Statistics client or from the SPSS Statistics Batch Facility—a separate executable provided with SPSS Statistics Server.

Python Program Run from an External Python Process. You can run a Python program from any external Python process, such as a Python IDE or the Python interpreter. In this mode, the Python program starts up a new instance of the SPSS Statistics processor without an associated instance of the SPSS Statistics client. You can use this mode to debug your Python programs using the Python IDE of your choice.

Invoking Python Scripts from Python Programs and Vice Versa

Python Script Run from Python Program. You can run a Python script from a Python program by importing the Python module containing the script and calling the function in the module that implements the script. You can also call Python script methods directly from within a Python program. [For more information, see *Running Python Scripts from Python Programs* on p. 586.](#) These features are not available in distributed mode or when running a Python program from an external Python process.

Python Autoscript Triggered from Python Program. A Python script specified as an autoscript will be triggered when a Python program executes the procedure containing the output item associated with the autoscript. For example, you associate an autoscript with the Descriptive Statistics table generated by the Descriptives procedure. You then run a Python program that executes the Descriptives procedure. The Python autoscript will be executed.

Python Program Run from Python Script. Python scripts can run command syntax, which means they can run command syntax containing Python programs.

Limitations and Warnings

- Running a Python program from the Python editor launched by SPSS Statistics will start up a new instance of the SPSS Statistics processor and will not interact with the instance of SPSS Statistics that launched the editor.
- Python programs are not intended to be run from Utilities>Run Script.
- Python programs cannot be run as autoscripts.
- The interfaces exposed by the `spss` module cannot be used in a Python script.

Getting Started with Python Scripts

The basic structure of a Python script is:

```
import SpssClient
SpssClient.StartClient()
<Python language statements>
SpssClient.StopClient()
```

- The `import SpssClient` statement imports the Python module containing the SPSS Statistics classes and methods available in the Python scripting interface.

- `SpssClient.StartClient()` provides a connection to the associated SPSS Statistics client, enabling the script to retrieve information from the client and to perform operations on objects managed by the client, such as pivot tables. Whether the script connects to an existing client or starts up a new client depends on how the script was invoked. [For more information, see Running Python Scripts and Python programs on p. 583.](#)
- `SpssClient.StopClient()` terminates the connection to the SPSS Statistics client and should be called at the completion of each Python script.

Note: If you're running a Python script from an external Python process that starts up a new client, call `SpssClient.Exit()` before `SpssClient.StopClient()`.

Example

This script accesses the designated output document and sets each of the pivot tables as selected.

```
import SpssClient
SpssClient.StartClient()

OutputDoc = SpssClient.GetDesignatedOutputDoc()
OutputItems = OutputDoc.GetOutputItems()

for index in range(OutputItems.Size()):
    OutputItem = OutputItems.GetItemAt(index)
    if OutputItem.GetType() == SpssClient.OutputItemType.PIVOT:
        OutputItem.SetSelected(True)
SpssClient.StopClient()
```

Target for Standard output

The Python `print` statement writes output to Python's standard output. When you run a Python script from Utilities>Run Script, Python's standard output is directed to a log item in the SPSS Statistics Viewer.

Getting Started with Autoscripts in Python

Autoscripts typically require a reference to the object that triggered the script. They may also require a reference to the associated output document and possibly the index of the output item in the output document. These values are obtained from the `SpssScriptContext` object, as shown in this example of an autoscript that transposes the rows and columns of a pivot table.

```
import SpssClient
SpssClient.StartClient()

SpssScriptContext = SpssClient.GetScriptContext()
SpssOutputItem = SpssScriptContext.GetOutputItem()
SpssPivotTable = SpssOutputItem.GetSpecificType()
SpssPivotMgr = SpssPivotTable.PivotManager()
SpssPivotMgr.TransposeRowsWithColumns()

SpssClient.StopClient()
```

- `SpssClient.GetScriptContext` returns an `SpssScriptContext` object that provides values for use by the autoscript.
- The `GetOutputItem` method of the `SpssScriptContext` object returns the output item that triggered the current autoscript—in this example, the pivot table whose rows and columns are to be transposed.

Although not used in this example, the `GetOutputDoc` method of the `SpssScriptContext` object returns the associated output document, and the `GetOutputItemIndex` method returns the index (in the associated output document) of the output item that triggered the autoscript.

Detecting When a Script is Run as an Autoscript

Using the `GetScriptContext` method, you can detect when a script is being run as an autoscript. This allows you to code a script so that it functions in either context (autoscript or not). This trivial script illustrates the approach.

```
import SpssClient
SpssClient.StartClient()

SpssScriptContext = SpssClient.GetScriptContext()
if SpssScriptContext == None:
    print "I'm not an autoscript"
else:
    print "I'm an autoscript"

SpssClient.StopClient()
```

- When a script is not run as an autoscript, the `GetScriptContext` method will return a value of `None`.
- Given the `if-else` logic in this example, you would include your autoscript-specific code in the `else` clause. Any code that is not to be run in the context of an autoscript would be included in the `if` clause. Of course you can also include code that is to be run in either context.

Running Python Scripts from Python Programs

You can run Python scripts from Python programs and you can call Python script methods from within a Python program. This allows you to write Python programs that operate on user interface and output objects.

Note: This feature is only available when running a Python program from the SPSS Statistics client—within a `BEGIN PROGRAM-END PROGRAM` block in command syntax. It is not available in distributed mode or when running a Python program from an external Python process.

Example: Calling a Python Script from a Python Program

This example shows a Python program that creates a custom pivot table and calls a Python script to make the column labels of the table bold.

```

BEGIN PROGRAM.
import spss, MakeColsBold
spss.StartProcedure("Demo")
table = spss.BasePivotTable("Sample Table", "OMS subtype")
table.SimplePivotTable(rowlabels = ["1", "2"],
                       collabels = ["A", "B"],
                       cells = ["1A", "1B", "2A", "2B"])
spss.EndProcedure()
MakeColsBold.Run("Sample Table")
END PROGRAM.

```

- Python programs use the interface exposed by the Python `spss` module, so the first line of the program contains an `import` statement for that module. The Python script is assumed to be contained in a Python module named `MakeColsBold`, so the `import` statement also includes that module.
- The code from `spss.StartProcedure` to `spss.EndProcedure` creates a pivot table titled “Sample Table”.
- `MakeColsBold.Run("Sample Table")` calls the `Run` function in the `MakeColsBold` module and passes the value “Sample Table” as the argument. The `Run` function implements the Python script to make the column labels of the specified table bold.

The content of the `MakeColsBold` module is as follows:

```

import SpssClient

def Run(tableName):
    SpssClient.StartClient()
    OutputDoc = SpssClient.GetDesignatedOutputDoc()
    OutputItems = OutputDoc.GetOutputItems()
    for index in range(OutputItems.Size()):
        OutputItem = OutputItems.GetItemAt(index)
        if OutputItem.GetType() == SpssClient.OutputItemType.PIVOT \
            and OutputItem.GetDescription() == tableName:
            PivotTable = OutputItem.GetSpecificType()
            ColumnLabels = PivotTable.ColumnLabelArray()
            for i in range(ColumnLabels.GetNumColumns()):
                ColumnLabels.SelectLabelAt(1, i)
            PivotTable.SetTextStyle(SpssClient.SpssTextStyleTypes.SpssTSBold)
    SpssClient.StopClient()

```

- The `import SpssClient` statement is needed to access the classes and methods available in the Python scripting interface.
- The module contains a single function named `Run`, which implements the script. It takes a single argument that specifies the name of the table to modify. There is nothing special about the name `Run` and the module is not limited to a single function. You can create a module that contains many functions, each of which implements a different script.
- The `Run` function calls `SpssClient.StartClient()` to provide a connection to the associated SPSS Statistics client and `SpssClient.StopClient()` to terminate the connection at the completion of the script.

Example: Calling Python Scripting Methods Directly from a Python Program

This example shows a Python program that creates a custom pivot table and makes direct calls to Python scripting methods to make the title of the table italic.

```

BEGIN PROGRAM.
import spss, SpssClient
spss.StartProcedure("Demo")
table = spss.BasePivotTable("Sample Table", "OMS subtype")
table.SimplePivotTable(cells = ["A", "B", "C", "D"])
spss.EndProcedure()

SpssClient.StartClient()
OutputDoc = SpssClient.GetDesignatedOutputDoc()
OutputItems = OutputDoc.GetOutputItems()
OutputItem = OutputItems.GetItemAt(OutputItems.Size()-1)
PivotTable = OutputItem.GetSpecificType()
PivotTable.SelectTitle()
PivotTable.SetTextStyle(SpssClient.SpssTextStyleTypes.SpssTSItalic)
SpssClient.StopClient()
END PROGRAM.

```

- The `import spss, SpssClient` statement provides access to the classes and methods available for Python programs (`spss`) as well as those for Python scripts (`SpssClient`).
- The code from `spss.StartProcedure` to `spss.EndProcedure` is the Python program code that creates the pivot table.
- The code from `SpssClient.StartClient()` to `SpssClient.StopClient()` is the Python script code that makes the title italic.

Script Editor for the Python Programming Language

For the Python programming language, the default editor is IDLE, which is provided with Python. IDLE provides an integrated development environment (IDE) with a limited set of features. Many IDE's are available for the Python programming language. For instance, on Windows you may choose to use the freely available PythonWin IDE.

To change the script editor for the Python programming language:

- ▶ Open the file *clientscriptingcfg.ini*, located in the directory where SPSS Statistics is installed.
- ▶ Under the section labeled [Python], change the value of `EDITOR_PATH` to point to the executable for the desired editor.
- ▶ In that same section, change the value of `EDITOR_ARGS` to handle any arguments that need to be passed to the editor. If no arguments are required, remove any existing values.

Scripting in Basic

Scripting in Basic is available on Windows only and is installed with the Base system. Extensive online help for scripting in Basic is available from the SPSS Statistics Basic Script Editor. The editor can be accessed from File>New>Script when the default script language (set from the Scripts tab on the Options dialog) is set to Basic (the system default on Windows). It is also accessed from File>Open>Script, by choosing Basic (`wwd;sbs`) in the Files of type list.

Compatibility with Versions Prior to 16.0

Obsolete Methods and Properties

A number of automation methods and properties are obsolete for version 16.0 and above. In terms of general features, this includes all objects associated with interactive graphs, the Draft Document object, and methods and properties associated with maps. For additional information, see “Release Notes for Version 16.0” in the help system provided with the SPSS Statistics Basic Script Editor. The SPSS Statistics-specific help is accessed from Help>SPSS Statistics Objects Help, in the script editor.

Global Procedures

Prior to version 16.0, the scripting facility included a global procedures file. For version 16.0 and above the scripting facility does not use a global procedures file, although the pre-16.0 version of *Global.sbs* (renamed *Global.wwd*) is installed for backwards compatibility.

To migrate a pre-16.0 version of a script that called functions in the global procedures file, add the statement `'#Uses "<install dir>\Samples\Global.wwd"` to the declarations section of the script, where `<install dir>` is the directory where SPSS Statistics is installed. `'#Uses` is a special comment recognized by the Basic script processor. If you're not sure if a script uses the global procedures file, you should add the `'#Uses` statement. You can also use `'$Include:` instead of `'#Uses`.

Legacy Autoscripts

Prior to version 16.0, the scripting facility included a single autoscript file containing all autoscripts. For version 16.0 and above there is no single autoscript file. Each autoscript is now stored in a separate file and can be applied to one or more output items, in contrast to pre-16.0 versions where each autoscript was specific to a particular output item.

Some of the autoscripts installed with pre-16.0 versions are available as a set of separate script files located in the *Samples* subdirectory of the directory where SPSS Statistics is installed. They are identified by a filename ending in *Autoscript*, with a file type of *wwd*. By default, they are not associated with any output items. The association is done from the Scripts tab of the Options dialog. [For more information, see Script Options in Chapter 48 on p. 507.](#)

Any custom autoscripts used in pre-16.0 versions must be manually converted and associated with one or more output items, from the Scripts tab of the Options dialog. The conversion process involves the following steps:

- ▶ Extract the subroutine specifying the autoscript from the legacy *Autoscript.sbs* file and save it as a new file with an extension of *wwd* or *sbs*. The name of the file is arbitrary.
- ▶ Change the name of the subroutine to `Main` and remove the parameter specification, keeping track of which parameters are required by the script, such as a pivot table that triggers the autoscript.
- ▶ Use the `scriptContext` object (always available) to get the values required by the autoscript, such as the output item that triggered the autoscript.

- From the Scripts tab of the Options dialog, associate the script file with the desired output object.

To illustrate the converted code, consider the autoscript *Descriptives_Table_DescriptiveStatistics_Create* from the legacy *Autoscript.sbs* file.

```
Sub Descriptives_Table_DescriptiveStatistics_Create(objPivotTable As Object,objOutputDoc As Object)
  'Autoscript
  'Trigger Event: DescriptiveStatistics Table Creation after running Descriptives procedure.

  'Purpose: Swaps the Rows and Columns in the currently active pivot table.
  'Assumptions: Selected Pivot Table is already activated.
  'Effects: Swaps the Rows and Columns in the output
  'Inputs: Pivot Table, OutputDoc, Item Index

  Dim objPivotManager As ISpssPivotMgr
  Set objPivotManager=objPivotTable.PivotManager
  objPivotManager.TransposeRowsWithColumns

End Sub
```

Following is the converted script:

```
Sub Main
  'Purpose: Swaps the Rows and Columns in the currently active pivot table.
  'Effects: Swaps the Rows and Columns in the output

  Dim objOutputItem As ISpssItem
  Dim objPivotTable as PivotTable
  Set objOutputItem = scriptContext.GetOutputItem()
  Set objPivotTable = objOutputItem.ActivateTable

  Dim objPivotManager As ISpssPivotMgr
  Set objPivotManager = objPivotTable.PivotManager
  objPivotManager.TransposeRowsWithColumns
  objOutputItem.Deactivate
End Sub
```

- Notice that nothing in the converted script indicates which object the script is to be applied to. The association between an output item and an autoscript is set from the Scripts tab of the Options dialog and maintained across sessions.
- `scriptContext.GetOutputItem` gets the output item (an `ISpssItem` object) that triggered the autoscript.
- The object returned by `scriptContext.GetOutputItem` is not activated. If your script requires an activated object, you'll need to activate it, as done in this example with the `ActivateTable` method. When you're finished with any table manipulations, call the `Deactivate` method.

For version 16.0, there is no distinction between scripts that are run as autoscripts and scripts that aren't run as autoscripts. Any script, appropriately coded, can be used in either context. [For more information, see The scriptContext Object on p. 591.](#)

Note: To trigger a script from the application creation event, see [Startup Scripts](#).

Script Editor

For version 16.0 and above the script editor for Basic no longer supports the following pre-16.0 features:

- The Script, Analyze, Graph, Utilities, and Add-Ons menus.
- The ability to paste command syntax into a script window.

The SPSS Statistics Basic Script Editor is a standalone application that is launched from within SPSS Statistics via File>New>Script, File>Open>Script, or Utilities>Create/Edit AutoScript (from a Viewer window). It allows you to run scripts against the instance of SPSS Statistics from which it was launched. Once opened, the editor will remain open after exiting SPSS Statistics, but scripts that use SPSS Statistics objects will no longer run.

File Types

For version 16.0 and above, the scripting facility will continue to support running and editing scripts with a file type of *sbs*. By default, new Basic scripts created with the SPSS Statistics Basic Script Editor have a file type of *wwd*.

Using External COM Clients

For version 16.0 and above, the program identifier for instantiating SPSS Statistics from an external COM client is `SPSS.Application16`. Application objects should be declared as `spsswinLib.Application16`. For example:

```
Dim objSpssApp As spsswinLib.Application16
Set objSpssApp=CreateObject("SPSS.Application16")
```

To connect to a running instance of the SPSS Statistics client from an external COM client, use:

```
Dim objSpssApp As spsswinLib.Application16
Set objSpssApp=GetObject("", "SPSS.Application16")
```

If more than one client is running, `GetObject` will connect to the most recently launched one.

Note: For post-16.0 versions, the identifier is still `Application16`.

The scriptContext Object

Detecting When a Script is Run as an Autocscript

Using the `scriptContext` object, you can detect when a script is being run as an autocscript. This allows you to code a script so that it functions in either context (autocscript or not). This trivial script illustrates the approach.

```
Sub Main
  If scriptContext Is Nothing Then
    MsgBox "I'm not an autocscript"
  Else
    MsgBox "I'm an autocscript"
  End If
End Sub
```

- When a script is not run as an autoscript, the `scriptContext` object will have a value of `Nothing`.
- Given the `If-Else` logic in this example, you would include your autoscript-specific code in the `Else` clause. Any code that is not to be run in the context of an autoscript would be included in the `If` clause. Of course you can also include code that is to be run in either context.

Getting Values Required by Autoscripts

The `scriptContext` object provides access to values required by an autoscript, such as the output item that triggered the current autoscript.

- The `scriptContext.GetOutputItem` method returns the output item (an `ISpssItem` object) that triggered the current autoscript.
- The `scriptContext.GetOutputDoc` method returns the output document (an `ISpssOutputDoc` object) associated with the current autoscript.
- The `scriptContext.GetOutputItemIndex` method returns the index, in the associated output document, of the output item that triggered the current autoscript.

Note: The object returned by `scriptContext.GetOutputItem` is not activated. If your script requires an activated object, you'll need to activate it—for instance, with the `ActivateTable` method. When you're finished with any manipulations, call the `Deactivate` method.

Startup Scripts

You can create a script that runs at the start of each session and a separate script that runs each time you switch servers. For Windows you can have versions of these scripts in both Python and Basic. For all other platforms the scripts can only be in Python.

- The startup script must be named `StartClient_.py` for Python or `StartClient_.wwd` for Basic.
- The script that runs when switching servers must be named `StartServer_.py` for Python or `StartServer_.wwd` for Basic.
- The scripts must be located in the `scripts` directory of the installation directory—located at the root of the installation directory for Windows and Linux, and under the `Contents` directory in the application bundle for Mac.
- On Windows, if the `scripts` directory contains both a Python and a Basic version of `StartClient_` or `StartServer_`, then both versions are executed. The order of execution is the Python version followed by the Basic version.
- If your system is configured to start up in distributed mode, then at the start of each session any `StartClient_` scripts are run followed by any `StartServer_` scripts. *Note:* The `StartServer_` scripts also run each time you switch servers, but the `StartClient_` scripts only run at the start of a session.

TABLES and IGRAPH Command Syntax Converter

If you have command syntax files that contain `TABLES` syntax that you want to convert to `CTABLES` syntax and/or `IGRAPH` syntax that you want to convert to `GGRAPH` syntax, a simple utility program is provided to help you get started with the conversion process. There are, however, significant functionality differences between `TABLES` and `CTABLES` and between `IGRAPH` and `GGRAPH`. It is likely that you will find that the utility program cannot convert some of your `TABLES` and `IGRAPH` syntax jobs or may generate `CTABLES` and `GGRAPH` syntax that produces tables and graphs that do not closely resemble the originals produced by the `TABLES` and `IGRAPH` commands. For most tables, you can edit the converted syntax to produce a table closely resembling the original.

The utility program is designed to:

- Create a new syntax file from an existing syntax file. The original syntax file is not altered.
- Convert only `TABLES` and `IGRAPH` commands in the syntax file. Other commands in the file are not altered.
- Retain the original `TABLES` and `IGRAPH` syntax in commented form.
- Identify the beginning and end of each conversion block with comments.
- Identify `TABLES` and `IGRAPH` syntax commands that could not be converted.
- Convert command syntax files that follow either interactive or production mode syntax rules.

This utility cannot convert commands that contain errors. The following other limitations also apply.

TABLES Limitations

The utility program may convert `TABLES` commands incorrectly under some circumstances, including `TABLES` commands that contain:

- Parenthesized variable names with the initial letters “sta” or “lab” in the `TABLES` subcommand if the variable is parenthesized by itself—for example, `var1 by (statvar) by (labvar)`. These will be interpreted as the `(STATISTICS)` and `(LABELS)` keywords.
- `SORT` subcommands that use the abbreviations `A` or `D` to indicate ascending or descending sort order. These will be interpreted as variable names.

The utility program cannot convert `TABLES` commands that contain:

- Syntax errors.

- OBSERVATION subcommands that refer to a range of variables using the TO keyword (for example, `var01 TO var05`).
- String literals broken into segments separated by plus signs (for example, `TITLE "My" + "Title"`).
- Macro calls that, in the absence of macro expansion, would be invalid TABLES syntax. Since the converter does not expand the macro calls, it treats them as if they were simply part of the standard TABLES syntax.

The utility program will not convert TABLES commands contained in macros. All macros are unaffected by the conversion process.

IGRAPH Limitations

IGRAPH changed significantly in release 16. Because of these changes, some subcommands and keywords in IGRAPH syntax created before that release may not be honored. See the IGRAPH section in the *Command Syntax Reference* for the complete revision history.

The conversion utility program may generate additional syntax that it stores in the INLINETEMPLATE keyword within the GGRAPH syntax. This keyword is created only by the conversion program. Its syntax is not intended to be user-editable.

Using the Conversion Utility Program

The conversion utility program, *SyntaxConverter.exe*, can be found in the installation directory. It is designed to run from a command prompt. The general form of the command is:

```
syntaxconverter.exe [path]/inputfilename.sps [path]/outputfilename.sps
```

You must run this command from the installation directory.

If any directory names contain spaces, enclose the entire path and filename in quotation marks, as in:

```
syntaxconverter.exe /myfiles/oldfile.sps "/new files/newfile.sps"
```

Interactive versus Production Mode Command Syntax Rules

The conversion utility program can convert command files that use interactive or production mode syntax rules.

Interactive. The interactive syntax rules are:

- Each command begins on a new line.
- Each command ends with a period (.).

Production mode. The Production Facility and commands in files accessed via the `INCLUDE` command in a different command file use production mode syntax rules:

- Each command must begin in the first column of a new line.
- Continuation lines must be indented at least one space.
- The period at the end of the command is optional.

If your command files use production mode syntax rules and don't contain periods at the end of each command, you need to include the command line switch `-b` (or `/b`) when you run *SyntaxConverter.exe*, as in:

```
syntaxconverter.exe -b /myfiles/oldfile.sps /myfiles/newfile.sps
```

SyntaxConverter Script (Windows Only)

On Windows, you can also run the syntax converter with the script *SyntaxConverter.wwd*, located in the *Samples* directory of the installation directory.

- ▶ From the menus choose:
 - Utilities
 - Run Script...
- ▶ Navigate to the *Samples* directory and select *SyntaxConverter.wwd*.

This will open a simple dialog box where you can specify the names and locations of the old and new command syntax files.

- Access (Microsoft), 16
- active file, 60, 62
 - caching, 62
 - creating a temporary active file, 62
 - virtual active file, 60
- active window, 4
- adding group labels, 229
- adjusted R^2
 - in Linear Regression, 349
- aggregating data, 175
 - aggregate functions, 178
 - variable names and labels, 178
- algorithms, 10
- alignment, 78, 205, 496
 - in Data Editor, 78
 - output, 205, 496
- alpha coefficient
 - in Reliability Analysis, 460, 462
- alpha factoring, 398
- analysis of variance
 - in Curve Estimation, 359
 - in Linear Regression, 349
 - in Means, 300
 - in One-Way ANOVA, 314
- Anderson-Rubin factor scores, 400
- Andrews' wave estimator
 - in Explore, 283
- ANOVA
 - in GLM Univariate, 320
 - in Means, 300
 - in One-Way ANOVA, 314
 - model, 322
- aspect ratio, 502
- attributes
 - custom variable attributes, 80
- automated production, 544
- autoscripts, 507, 579
 - associating with viewer objects, 581
 - Basic, 591
 - creating, 580
 - Python, 585
 - trigger events, 580
- average absolute deviation (AAD)
 - in Ratio Statistics, 473

- background color, 242
- backward elimination
 - in Linear Regression, 344
- banding, 116
- bar charts
 - in Frequencies, 276

- Bartlett factor scores, 400
- Bartlett's test of sphericity
 - in Factor Analysis, 397
- basic steps, 8
- beta coefficients
 - in Linear Regression, 349
- binning, 116
- Binomial Test, 425
 - command additional features, 427
 - dichotomies, 425
 - missing values, 426
 - options, 426
 - statistics, 426
- Bivariate Correlations
 - command additional features, 334
 - correlation coefficients, 332
 - missing values, 334
 - options, 334
 - significance level, 332
 - statistics, 334
- block distance
 - in Distances, 340
- Blom estimates, 141
- BMP files, 210, 221
 - exporting charts, 210, 221
- Bonferroni
 - in GLM, 326
 - in One-Way ANOVA, 316
- borders, 239, 246
 - displaying hidden borders, 246
- boxplots
 - comparing factor levels, 284
 - comparing variables, 284
 - in Explore, 284
- Box's M test
 - in Discriminant Analysis, 390
- break variables
 - in Aggregate Data, 175
- Brown-Forsythe statistic
 - in One-Way ANOVA, 318
- build terms, 322, 358

- caching, 62
 - active file, 62
- Cancel button, 7
- captions, 244–245
- case-control study
 - Paired-Samples T Test, 310
- cases, 87, 185
 - finding duplicates, 113
 - finding in Data Editor, 88

- inserting new cases, 87
- restructuring into variables, 185
- selecting subsets, 180–181, 183
- sorting, 167
- weighting, 183
- casewise diagnostics
 - in Linear Regression, 349
- categorical data, 101
 - converting interval data to discrete categories, 116
- cell properties, 242–243
- cells in pivot tables, 237, 245–246
 - formats, 237
 - hiding, 232
 - showing, 232
 - widths, 246
- centered moving average function, 160
- centering output, 205, 496
- Chart Builder, 478
 - gallery, 479
- chart creation, 478
- Chart Editor, 482
 - properties, 483
- chart options, 502
- charts, 204, 210, 248, 478, 502
 - aspect ratio, 502
 - case labels, 359
 - Chart Builder, 478
 - copying into other applications, 209
 - creating, 478
 - creating from pivot tables, 248
 - exporting, 210
 - hiding, 204
 - in ROC Curve, 475
 - missing values, 485
 - overview, 478
 - size, 485
 - templates, 485, 502
 - wrapping panels, 485
- chartsoutput, 209
- Chebychev distance
 - in Distances, 340
- chi-square, 422
 - expected range, 424
 - expected values, 424
 - Fisher’s exact test, 289
 - for independence, 289
 - in Crosstabs, 289
 - likelihood-ratio, 289
 - linear-by-linear association, 289
 - missing values, 424
 - one-sample test, 422
 - options, 424
 - Pearson, 289
 - statistics, 424
 - Yates’ correction for continuity, 289
- chi-square distance
 - in Distances, 340
- city-block distance
 - in Nearest Neighbor Analysis, 371
- classification
 - in ROC Curve, 475
- classification table
 - in Nearest Neighbor Analysis, 385
- cluster analysis
 - efficiency, 418
 - Hierarchical Cluster Analysis, 412
 - K-Means Cluster Analysis, 417
- cluster frequencies
 - in TwoStep Cluster Analysis, 410
- clustering
 - choosing a procedure, 403
- Cochran’s Q
 - in Tests for Several Related Samples, 439
- Cochran’s statistic
 - in Crosstabs, 289
- Codebook, 266
 - output, 268
 - statistics, 270
- coefficient of dispersion (COD)
 - in Ratio Statistics, 473
- coefficient of variation (COV)
 - in Ratio Statistics, 473
- Cohen’s kappa
 - in Crosstabs, 289
- collapsing categories, 116
- collinearity diagnostics
 - in Linear Regression, 349
- colors in pivot tables, 239
 - borders, 239
- column percentages
 - in Crosstabs, 291
- column summary reports, 454
- column width, 78, 235, 246, 505
 - controlling default width, 505
 - controlling maximum width, 235
 - controlling width for wrapped text, 235
 - in Data Editor, 78
 - pivot tables, 246
- columns, 246–247
 - changing width in pivot tables, 246
 - selecting in pivot tables, 247
- COMMA format, 74–75
- comma-delimited files, 28
- command identifiers, 556
- command language, 253
- command line switches, 549
 - Production jobs, 549
- command syntax, 253, 263, 514, 518, 544
 - accessing *Command Syntax Reference*, 10
 - adding to menus, 514
 - journal file, 265
 - output log, 255
 - pasting, 255
 - Production jobs rules, 544

- running, 263
- running with toolbar buttons, 518
- syntax rules, 253
- command syntax editor, 257
 - auto-completion, 259
 - bookmarks, 258, 262
 - breakpoints, 258, 261, 263
 - color coding, 260
 - command spans, 258
 - commenting out text, 263
 - line numbers, 258
 - options, 510
- command syntax files, 264
- comparing groups
 - in OLAP Cubes, 305
- comparing variables
 - in OLAP Cubes, 305
- compound model
 - in Curve Estimation, 360
- computing variables, 124
 - computing new string variables, 126
- concentration index
 - in Ratio Statistics, 473
- conditional transformations, 126
- confidence intervals
 - in Explore, 283
 - in GLM, 324, 329
 - in Independent-Samples T Test, 309
 - in Linear Regression, 349
 - in One-Sample T Test, 313
 - in One-Way ANOVA, 318
 - in Paired-Samples T Test, 311
 - in ROC Curve, 477
 - saving in Linear Regression, 347
- contingency coefficient
 - in Crosstabs, 289
- contingency tables, 287
- continuation text, 240
 - for pivot tables, 240
- contrasts
 - in GLM, 324
 - in One-Way ANOVA, 315
- control variables
 - in Crosstabs, 288
- convergence
 - in Factor Analysis, 398–399
 - in K-Means Cluster Analysis, 419
- Cook's distance
 - in GLM, 328
 - in Linear Regression, 347
- correlation matrix
 - in Discriminant Analysis, 390
 - in Factor Analysis, 395, 397
 - in Ordinal Regression, 355
- correlations
 - in Bivariate Correlations, 332
 - in Crosstabs, 289
 - in Partial Correlations, 335
 - zero-order, 336
- counting occurrences, 129
- covariance matrix
 - in Discriminant Analysis, 390, 392
 - in GLM, 328
 - in Linear Regression, 349
 - in Ordinal Regression, 355
- covariance ratio
 - in Linear Regression, 347
- Cox and Snell R^2
 - in Ordinal Regression, 355
- Cramér's V
 - in Crosstabs, 289
- Cronbach's alpha
 - in Reliability Analysis, 460, 462
- Crosstabs, 287
 - cell display, 291
 - clustered bar charts, 289
 - control variables, 288
 - formats, 292
 - fractional weights, 183
 - layers, 288
 - statistics, 289
 - suppressing tables, 287
- crostabulation
 - in Crosstabs, 287
 - multiple response, 444
- CSV format
 - reading data, 28
 - saving data, 42
- CTABLES
 - converting TABLES command syntax to CTABLES, 593
- cubic model
 - in Curve Estimation, 360
- cumulative frequencies
 - in Ordinal Regression, 355
- cumulative sum function, 160
- currency formats, 499
- Curve Estimation, 359
 - analysis of variance, 359
 - forecast, 361
 - including constant, 359
 - models, 360
 - saving predicted values, 361
 - saving prediction intervals, 361
 - saving residuals, 361
- custom attributes, 80
- custom currency formats, 74, 499
- Custom Dialog Builder, 520
 - check box, 532
 - check box group, 537
 - combo box, 533
 - combo box list items, 533
 - custom dialog package (spd) files, 528
 - custom dialogs for extension commands, 540

- dialog properties, 522
- file browser, 538
- file type filter, 539
- filtering variable lists, 532
- help file, 522
- installing dialogs, 528
- item group control, 535
- layout rules, 524
- localizing dialogs and help files, 541
- menu location, 523
- modifying installed dialogs, 528
- number control, 534
- opening dialog specification files, 528
- preview, 527
- radio group, 536
- radio group buttons, 536
- saving dialog specifications, 528
- source list, 530
- static text control, 535
- sub-dialog button, 539
- sub-dialog properties, 540
- syntax template, 525
- target list, 531
- text control, 534
- custom models
 - in GLM, 322
- custom tables
 - converting TABLES command syntax to CTABLES, 593
- custom variable attributes, 80

- d*
 - in Crosstabs, 289
- data analysis, 8
 - basic steps, 8
- data dictionary
 - applying from another file, 106
- Data Editor, 70, 72, 78, 84–88, 90–92, 514
 - alignment, 78
 - changing data type, 88
 - column width, 78
 - data value restrictions, 85
 - Data View, 70
 - defining variables, 72
 - display options, 91
 - editing data, 85–86
 - entering data, 84
 - entering non-numeric data, 85
 - entering numeric data, 84
 - filtered cases, 90
 - inserting new cases, 87
 - inserting new variables, 87
 - moving variables, 87
 - multiple open data files, 93, 494
 - multiple views/panes, 91
 - printing, 92
 - sending data to other applications, 514
 - Variable View, 71
- data entry, 84
- data files, 12–13, 28, 40–42, 47, 62, 67, 185
 - adding comments, 489
 - dictionary information, 40
 - Dimensions, 37
 - file information, 40
 - flipping, 169
 - improving performance for large files, 62
 - mrInterview, 37
 - multiple open data files, 93, 494
 - opening, 12–13
 - protecting, 60
 - Quancept, 37
 - Quanvert, 37
 - remote servers, 67
 - restructuring, 185
 - saving, 41–42
 - saving output as SPSS Statistics data files, 552
 - saving subsets of variables, 47
 - text, 28
 - transposing, 169
- DATA LIST, 60
 - versus GET DATA command, 60
- data transformations, 497
 - computing variables, 124
 - conditional transformations, 126
 - delaying execution, 497
 - functions, 127
 - ranking cases, 140
 - recoding values, 132–133, 135, 137
 - string variables, 126
 - time series, 157, 159
- data types, 74–75, 88, 499
 - changing, 88
 - custom currency, 74, 499
 - defining, 74
 - display formats, 75
 - input formats, 75
- Data View, 70
- databases, 15–16, 19–20, 22, 24–25, 27
 - adding new fields to a table, 55
 - appending records (cases) to a table, 56
 - conditional expressions, 22
 - converting strings to numeric variables, 25
 - creating a new table, 57
 - creating relationships, 20
 - defining variables, 25
 - Microsoft Access, 16
 - parameter queries, 22, 24
 - Prompt for Value, 24
 - random sampling, 22
 - reading, 15–16, 19
 - replacing a table, 57
 - replacing values in existing fields, 54
 - saving, 47
 - saving queries, 27

- selecting a data source, 16
- selecting data fields, 19
- specifying criteria, 22
- SQL syntax, 27
- table joins, 20
- updating, 47
- verifying results, 27
- Where clause, 22
- datasets
 - renaming, 95
- date format variables, 74–75, 497
 - add or subtract from date/time variables, 142
 - create date/time variable from set of variables, 142
 - create date/time variable from string, 142
 - extract part of date/time variable, 142
- date formats
 - two-digit years, 497
- date variables
 - defining for time series data, 158
- dBASE files, 12, 14, 41–42
 - reading, 12, 14
 - saving, 41–42
- default file locations, 506
- Define Multiple Response Sets, 441
 - categories, 441
 - dichotomies, 441
 - set labels, 441
 - set names, 441
- defining variables, 72, 74, 76–79, 97
 - applying a data dictionary, 106
 - copying and pasting attributes, 78–79
 - data types, 74
 - missing values, 77
 - templates, 78–79
 - value labels, 76, 97
 - variable labels, 76
- deleted residuals
 - in GLM, 328
 - in Linear Regression, 347
- deleting multiple EXECUTES in syntax files, 265
- deleting output, 204
- dendrograms
 - in Hierarchical Cluster Analysis, 415
- dependent *t* test
 - in Paired-Samples T Test, 310
- descriptive statistics
 - in Descriptives, 278
 - in Explore, 283
 - in Frequencies, 274
 - in GLM Univariate, 329
 - in Ratio Statistics, 473
 - in Summarize, 296
 - in TwoStep Cluster Analysis, 410
- Descriptives, 278
 - command additional features, 280
 - display order, 279
 - saving *z* scores, 278
 - statistics, 279
- designated window, 4
- detrended normal plots
 - in Explore, 284
- deviation contrasts
 - in GLM, 324
- DfBeta
 - in Linear Regression, 347
- DfFit
 - in Linear Regression, 347
- dialog boxes, 7, 489–490, 494
 - controls, 7
 - defining variable sets, 489
 - displaying variable labels, 5, 494
 - displaying variable names, 5, 494
 - reordering target lists, 492
 - selecting variables, 7
 - using variable sets, 490
 - variable display order, 494
 - variable icons, 7
 - variable information, 8
 - variables, 5
- dictionary, 40
 - Codebook, 266
- difference contrasts
 - in GLM, 324
- difference function, 160
- differences between groups
 - in OLAP Cubes, 305
- differences between variables
 - in OLAP Cubes, 305
- Dimensions data, 37
 - saving, 59
- direct oblimin rotation
 - in Factor Analysis, 399
- Discriminant Analysis, 387
 - command additional features, 393
 - covariance matrix, 392
 - criteria, 391
 - defining ranges, 389
 - descriptive statistics, 390
 - discriminant methods, 391
 - display options, 391–392
 - example, 387
 - exporting model information, 393
 - function coefficients, 390
 - grouping variables, 387
 - independent variables, 387
 - Mahalanobis distance, 391
 - matrices, 390
 - missing values, 392
 - plots, 392
 - prior probabilities, 392
 - Rao's *V*, 391
 - saving classification variables, 393
 - selecting cases, 389
 - statistics, 387, 390

- stepwise methods, 387
 - Wilks' lambda, 391
- disk space, 60, 62
 - temporary, 60, 62
- display formats, 75
- display order, 228
- distance measures
 - in Distances, 340
 - in Hierarchical Cluster Analysis, 413
 - in Nearest Neighbor Analysis, 371
- Distances, 338
 - command additional features, 341
 - computing distances between cases, 338
 - computing distances between variables, 338
 - dissimilarity measures, 340
 - example, 338
 - similarity measures, 341
 - statistics, 338
 - transforming measures, 340–341
 - transforming values, 340–341
- distributed mode, 64–65, 67–69
 - available procedures, 68
 - data file access, 67
 - relative paths, 69
- division
 - dividing across report columns, 456
- DOLLAR format, 74–75
- DOT format, 74–75
- Duncan's multiple range test
 - in GLM, 326
 - in One-Way ANOVA, 316
- Dunnett's *C*
 - in GLM, 326
 - in One-Way ANOVA, 316
- Dunnett's *t* test
 - in GLM, 326
 - in One-Way ANOVA, 316
- Dunnett's *T3*
 - in GLM, 326
 - in One-Way ANOVA, 316
- duplicate cases (records)
 - finding and filtering, 113
- Durbin-Watson statistic
 - in Linear Regression, 349

- editing data, 85–86
- effect-size estimates
 - in GLM Univariate, 329
- eigenvalues
 - in Factor Analysis, 397–398
 - in Linear Regression, 349
- entering data, 84–85
 - non-numeric, 85
 - numeric, 84
 - using value labels, 85
- environment variables, 506
 - SPSSTMPDIR, 506

- EPS files, 210, 221
 - exporting charts, 210, 221
- equamax rotation
 - in Factor Analysis, 399
- error summary
 - in Nearest Neighbor Analysis, 386
- estimated marginal means
 - in GLM Univariate, 329
- eta
 - in Crosstabs, 289
 - in Means, 300
- eta-squared
 - in GLM Univariate, 329
 - in Means, 300
- Euclidean distance
 - in Distances, 340
 - in Nearest Neighbor Analysis, 371
- Excel files, 12, 14, 42, 514
 - adding menu item to send data to Excel, 514
 - opening, 12, 14
 - saving, 41–42
 - saving value labels instead of values, 41
- Excel format
 - exporting output, 210, 214, 558
- excluding output from Viewer with OMS, 563
- EXECUTE (command)
 - pasted from dialog boxes, 265
- expected count
 - in Crosstabs, 291
- expected frequencies
 - in Ordinal Regression, 355
- Explore, 282
 - command additional features, 286
 - missing values, 285
 - options, 285
 - plots, 284
 - power transformations, 285
 - statistics, 283
- exponential model
 - in Curve Estimation, 360
- export data, 41
- exporting
 - models, 252
- exporting charts, 210, 221, 544
 - automated production, 544
- exporting data, 514
 - adding menu items to export data, 514
- exporting output, 210, 216, 219
 - Excel format, 210, 214, 558
 - HTML, 212
 - HTML format, 210
 - OMS, 552
 - PDF format, 210, 217, 558
 - PowerPoint format, 210
 - text format, 558
 - Word format, 210, 213, 558

- extension commands
 - custom dialogs, 540
- extreme values
 - in Explore, 283
- Factor Analysis, 395
 - coefficient display format, 401
 - command additional features, 401
 - convergence, 398–399
 - descriptives, 397
 - example, 395
 - extraction methods, 398
 - factor scores, 400
 - loading plots, 399
 - missing values, 401
 - overview, 395
 - rotation methods, 399
 - selecting cases, 396
 - statistics, 395, 397
- factor scores, 400
- feature selection
 - in Nearest Neighbor Analysis, 383
- feature space chart
 - in Nearest Neighbor Analysis, 379
- file information, 40
- file locations
 - controlling default file locations, 506
- file transformations, 185
 - aggregating data, 175
 - merging data files, 170, 173
 - restructuring data, 185
 - sorting cases, 167
 - split-file processing, 179
 - transposing variables and cases, 169
 - weighting cases, 183
- files, 207
 - adding a text file to the Viewer, 207
 - opening, 12
- filtered cases, 90
 - in Data Editor, 90
- find and replace
 - viewer documents, 208
- first
 - in Means, 300
 - in OLAP Cubes, 303
 - in Summarize, 296
- Fisher's exact test
 - in Crosstabs, 289
- Fisher's LSD
 - in GLM, 326
- fixed format, 28
- fonts, 91, 207, 242
 - in Data Editor, 91
 - in the outline pane, 207
- footers, 223
- footnotes, 236, 244–245
 - charts, 485
 - markers, 236, 245
 - renumbering, 245
- forecast
 - in Curve Estimation, 361
- formatting
 - columns in reports, 449
- forward selection
 - in Linear Regression, 344
 - in Nearest Neighbor Analysis, 372
- freefield format, 28
- Frequencies, 273
 - charts, 276
 - display order, 276
 - formats, 276
 - statistics, 274
 - suppressing tables, 276
- frequency tables
 - in Explore, 283
 - in Frequencies, 273
- Friedman test
 - in Tests for Several Related Samples, 439
- full factorial models
 - in GLM, 322
- functions, 127
 - missing value treatment, 127
- Gabriel's pairwise comparisons test
 - in GLM, 326
 - in One-Way ANOVA, 316
- Games and Howell's pairwise comparisons test
 - in GLM, 326
 - in One-Way ANOVA, 316
- gamma
 - in Crosstabs, 289
- generalized least squares
 - in Factor Analysis, 398
- geometric mean
 - in Means, 300
 - in OLAP Cubes, 303
 - in Summarize, 296
- GET DATA, 60
 - versus DATA LIST command, 60
 - versus GET CAPTURE command, 60
- GGRAPH
 - converting IGRAPH to GGRAPH, 593
- GLM
 - model, 322
 - post hoc tests, 326
 - profile plots, 325
 - saving matrices, 328
 - saving variables, 328
 - sum of squares, 322
- GLM Univariate, 320, 330
 - contrasts, 324
 - diagnostics, 329
 - display, 329
 - estimated marginal means, 329

- options, 329
- Goodman and Kruskal's gamma
 - in Crosstabs, 289
- Goodman and Kruskal's lambda
 - in Crosstabs, 289
- Goodman and Kruskal's tau
 - in Crosstabs, 289
- goodness of fit
 - in Ordinal Regression, 355
- grand totals
 - in column summary reports, 458
- grid lines, 246
 - pivot tables, 246
- group labels, 229
- group means, 298, 302
- grouped median
 - in Means, 300
 - in OLAP Cubes, 303
 - in Summarize, 296
- grouping rows or columns, 229
- grouping variables, 185
 - creating, 185
- growth model
 - in Curve Estimation, 360
- Guttman model
 - in Reliability Analysis, 460, 462

- Hampel's redescending M-estimator
 - in Explore, 283
- harmonic mean
 - in Means, 300
 - in OLAP Cubes, 303
 - in Summarize, 296
- headers, 223
- Helmert contrasts
 - in GLM, 324
- Help button, 7
- Help windows, 10
- hiding, 204, 233, 245, 515
 - captions, 245
 - dimension labels, 233
 - footnotes, 245
 - procedure results, 204
 - rows and columns, 233
 - titles, 233
 - toolbars, 515
- hiding (excluding) output from the Viewer with OMS, 563
- Hierarchical Cluster Analysis, 412
 - agglomeration schedules, 414
 - cluster membership, 414, 416
 - clustering cases, 412
 - clustering methods, 413
 - clustering variables, 412
 - command additional features, 416
 - dendrograms, 415
 - distance matrices, 414
 - distance measures, 413
 - example, 412
 - icicle plots, 415
 - plot orientation, 415
 - saving new variables, 416
 - similarity measures, 413
 - statistics, 412, 414
 - transforming measures, 413
 - transforming values, 413
- hierarchical decomposition, 323
- histograms
 - in Explore, 284
 - in Frequencies, 276
 - in Linear Regression, 346
- Hochberg's GT2
 - in GLM, 326
 - in One-Way ANOVA, 316
- holdout sample
 - in Nearest Neighbor Analysis, 373
- homogeneity-of-variance tests
 - in GLM Univariate, 329
 - in One-Way ANOVA, 318
- Hotelling's T^2
 - in Reliability Analysis, 460, 462
- HTML, 210, 212
 - exporting output, 210, 212
- Huber's M-estimator
 - in Explore, 283

- ICC. *See* intraclass correlation coefficient, 462
- icicle plots
 - in Hierarchical Cluster Analysis, 415
- icons
 - in dialog boxes, 7
- IGRAPH
 - converting IGRAPH to GGRAPH, 593
- image factoring, 398
- import data, 12, 15
- importance chart
 - in TwoStep Cluster Analysis, 409
- imputations
 - finding in Data Editor, 88
- Independent-Samples T Test, 307
 - confidence intervals, 309
 - defining groups, 309
 - grouping variables, 309
 - missing values, 309
 - options, 309
 - string variables, 309
- initial threshold
 - in TwoStep Cluster Analysis, 407
- inner join, 20
- input formats, 75
- inserting group labels, 229
- interaction terms, 322, 358
- interactive charts, 209
 - copying into other applications, 209

- intraclass correlation coefficient (ICC)
 - in Reliability Analysis, 462
- inverse model
 - in Curve Estimation, 360
- iteration history
 - in Ordinal Regression, 355
- iterations
 - in Factor Analysis, 398–399
 - in K-Means Cluster Analysis, 419
- journal file, 506
- JPEG files, 210, 221
 - exporting charts, 210, 221
- justification, 205, 496
 - output, 205, 496
- k and feature selection
 - in Nearest Neighbor Analysis, 385
- k selection
 - in Nearest Neighbor Analysis, 384
- K-Means Cluster Analysis
 - cluster distances, 419
 - cluster membership, 419
 - command additional features, 421
 - convergence criteria, 419
 - efficiency, 418
 - examples, 417
 - iterations, 419
 - methods, 417
 - missing values, 420
 - overview, 417
 - saving cluster information, 419
 - statistics, 417, 420
- kappa
 - in Crosstabs, 289
- Kendall's tau-*b*
 - in Bivariate Correlations, 332
 - in Crosstabs, 289
- Kendall's tau-*c*, 289
 - in Crosstabs, 289
- Kendall's *W*
 - in Tests for Several Related Samples, 439
- keyed table, 173
- Kolmogorov-Smirnov *Z*
 - in One-Sample Kolmogorov-Smirnov Test, 429
 - in Two-Independent-Samples Tests, 432
- KR20
 - in Reliability Analysis, 462
- Kruskal-Wallis *H*
 - in Two-Independent-Samples Tests, 436
- Kruskal's tau
 - in Crosstabs, 289
- Kuder-Richardson 20 (KR20)
 - in Reliability Analysis, 462
- kurtosis
 - in Descriptives, 279
 - in Explore, 283
- in Frequencies, 274
- in Means, 300
- in OLAP Cubes, 303
- in Report Summaries in Columns, 455
- in Report Summaries in Rows, 450
- in Summarize, 296
- labels, 229
 - deleting, 229
 - inserting group labels, 229
 - vs. subtype names in OMS, 557
- LAG (function), 160
- lag function, 131
- lambda
 - in Crosstabs, 289
- Lance and Williams dissimilarity measure, 340
 - in Distances, 340
- language
 - changing output language, 494
 - changing user interface language, 494
- last
 - in Means, 300
 - in OLAP Cubes, 303
 - in Summarize, 296
- layers, 222, 230, 232, 235, 240
 - creating, 230
 - displaying, 230, 232
 - in Crosstabs, 288
 - in pivot tables, 230
 - printing, 222, 235, 240
- lead function, 131, 160
- least significant difference
 - in GLM, 326
 - in One-Way ANOVA, 316
- level of measurement, 73, 101
 - defining, 73
- Levene test
 - in Explore, 284
 - in GLM Univariate, 329
 - in One-Way ANOVA, 318
- leverage values
 - in GLM, 328
 - in Linear Regression, 347
- likelihood-ratio chi-square
 - in Crosstabs, 289
 - in Ordinal Regression, 355
- Lilliefors test
 - in Explore, 284
- line breaks
 - variable and value labels, 77
- linear model
 - in Curve Estimation, 360
- Linear Regression, 343
 - blocks, 343
 - command additional features, 352
 - exporting model information, 347
 - missing values, 351

- plots, 346
- residuals, 347
- saving new variables, 347
- selection variable, 345
- statistics, 349
- variable selection methods, 344, 351
- weights, 343
- linear-by-linear association
 - in Crosstabs, 289
- link
 - in Ordinal Regression, 354
- listing cases, 294
- loading plots
 - in Factor Analysis, 399
- local encoding, 264
- location model
 - in Ordinal Regression, 356
- logarithmic model
 - in Curve Estimation, 360
- logging in to a server, 64
- logistic model
 - in Curve Estimation, 360
- Lotus 1-2-3 files, 12, 41–42, 514
 - adding menu item to send data to Lotus, 514
 - opening, 12
 - saving, 41–42

- M-estimators
 - in Explore, 283
- Mahalanobis distance
 - in Discriminant Analysis, 391
 - in Linear Regression, 347
- Manhattan distance
 - in Nearest Neighbor Analysis, 371
- Mann-Whitney *U*
 - in Two-Independent-Samples Tests, 432
- Mantel-Haenszel statistic
 - in Crosstabs, 289
- marginal homogeneity test
 - in Two-Related-Samples Tests, 434
- matched-pairs study
 - in Paired-Samples T Test, 310
- maximum
 - comparing report columns, 456
 - in Descriptives, 279
 - in Explore, 283
 - in Frequencies, 274
 - in Means, 300
 - in OLAP Cubes, 303
 - in Ratio Statistics, 473
 - in Summarize, 296
- maximum branches
 - in TwoStep Cluster Analysis, 407
- maximum likelihood
 - in Factor Analysis, 398
- McFadden R^2
 - in Ordinal Regression, 355
- McNemar test
 - in Crosstabs, 289
 - in Two-Related-Samples Tests, 434
- mean
 - in Descriptives, 279
 - in Explore, 283
 - in Frequencies, 274
 - in Means, 300
 - in OLAP Cubes, 303
 - in One-Way ANOVA, 318
 - in Ratio Statistics, 473
 - in Report Summaries in Columns, 455
 - in Report Summaries in Rows, 450
 - in Summarize, 296
 - of multiple report columns, 456
 - subgroup, 298, 302
- Means, 298
 - options, 300
 - statistics, 300
- measurement level, 73, 101
 - defining, 73
 - icons in dialog boxes, 7
- measurement system, 494
- measures of central tendency
 - in Explore, 283
 - in Frequencies, 274
 - in Ratio Statistics, 473
- measures of dispersion
 - in Descriptives, 279
 - in Explore, 283
 - in Frequencies, 274
 - in Ratio Statistics, 473
- measures of distribution
 - in Descriptives, 279
 - in Frequencies, 274
- median
 - in Explore, 283
 - in Frequencies, 274
 - in Means, 300
 - in OLAP Cubes, 303
 - in Ratio Statistics, 473
 - in Summarize, 296
- median test
 - in Two-Independent-Samples Tests, 436
- memory, 494
- memory allocation
 - in TwoStep Cluster Analysis, 407
- menus, 514
 - customizing, 514
- merging data files
 - dictionary information, 173
 - files with different cases, 170
 - files with different variables, 173
 - renaming variables, 173
- metafiles, 210
 - exporting charts, 210
- Microsoft Access, 16

- minimum
 - comparing report columns, 456
 - in Descriptives, 279
 - in Explore, 283
 - in Frequencies, 274
 - in Means, 300
 - in OLAP Cubes, 303
 - in Ratio Statistics, 473
 - in Summarize, 296
- Minkowski distance
 - in Distances, 340
- missing values, 77, 485
 - charts, 485
 - defining, 77
 - in Binomial Test, 426
 - in Bivariate Correlations, 334
 - in Chi-Square Test, 424
 - in column summary reports, 458
 - in Explore, 285
 - in Factor Analysis, 401
 - in functions, 127
 - in Independent-Samples T Test, 309
 - in Linear Regression, 351
 - in Multiple Response Crosstabs, 446
 - in Multiple Response Frequencies, 443
 - in Nearest Neighbor Analysis, 377
 - in One-Sample Kolmogorov-Smirnov Test, 430
 - in One-Sample T Test, 313
 - in One-Way ANOVA, 318
 - in Paired-Samples T Test, 311
 - in Partial Correlations, 336
 - in Report Summaries in Rows, 451
 - in ROC Curve, 477
 - in Runs Test, 428
 - in Tests for Several Independent Samples, 438
 - in Two-Independent-Samples Tests, 433
 - in Two-Related-Samples Tests, 435
 - replacing in time series data, 161
 - string variables, 77
- mode
 - in Frequencies, 274
- model file
 - loading saved models to score data, 164
- model view
 - in Nearest Neighbor Analysis, 378
- Model Viewer, 250
- models, 250
 - activating, 250
 - copying, 251
 - exporting, 252
 - interacting, 250
 - Model Viewer, 250
 - printing, 252
 - properties, 251
- Moses extreme reaction test
 - in Two-Independent-Samples Tests, 432
- moving rows and columns, 228
- mrInterview, 37
- Multidimensional Scaling, 465
 - command additional features, 470
 - conditionality, 468
 - creating distance matrices, 467
 - criteria, 469
 - defining data shape, 467
 - dimensions, 468
 - display options, 469
 - distance measures, 467
 - example, 465
 - levels of measurement, 468
 - scaling models, 468
 - statistics, 465
 - transforming values, 467
- multiple comparisons
 - in One-Way ANOVA, 316
- Multiple Imputation
 - options, 512
- multiple open data files, 93, 494
 - suppressing, 96
- multiple *R*
 - in Linear Regression, 349
- multiple regression
 - in Linear Regression, 343
- Multiple Response
 - command additional features, 447
- multiple response analysis
 - crosstabulation, 444
 - frequency tables, 443
 - Multiple Response Crosstabs, 444
 - Multiple Response Frequencies, 443
- Multiple Response Crosstabs, 444
 - cell percentages, 446
 - defining value ranges, 446
 - matching variables across response sets, 446
 - missing values, 446
 - percentages based on cases, 446
 - percentages based on responses, 446
- Multiple Response Frequencies, 443
 - missing values, 443
- multiple response sets
 - Codebook, 266
 - defining, 104
 - multiple categories, 104
 - multiple dichotomies, 104
- multiple views/panes
 - Data Editor, 91
- multiplication
 - multiplying across report columns, 456
- Nagelkerke R^2
 - in Ordinal Regression, 355
- Nearest Neighbor Analysis, 367
 - feature selection, 372
 - model view, 378
 - neighbors, 371

- options, 377
- output, 376
- partitions, 373
- saving variables, 375
- nearest neighbor distances
 - in Nearest Neighbor Analysis, 382
- new features
 - version 17.0, 2
- Newman-Keuls
 - in GLM, 326
- noise handling
 - in TwoStep Cluster Analysis, 407
- nominal, 73
 - measurement level, 73, 101
- nonparametric tests
 - chi-square, 422
 - One-Sample Kolmogorov-Smirnov Test, 429
 - Runs Test, 427
 - Tests for Several Independent Samples, 436
 - Tests for Several Related Samples, 438
 - Two-Independent-Samples Tests, 431
 - Two-Related-Samples Tests, 434
- normal probability plots
 - in Explore, 284
 - in Linear Regression, 346
- normal scores
 - in Rank Cases, 141
- normality tests
 - in Explore, 284
- number of cases
 - in Means, 300
 - in OLAP Cubes, 303
 - in Summarize, 296
- numeric format, 74–75

- observed count
 - in Crosstabs, 291
- observed frequencies
 - in Ordinal Regression, 355
- observed means
 - in GLM Univariate, 329
- OK button, 7
- OLAP Cubes, 302
 - statistics, 303
 - titles, 306
- OMS, 552, 575
 - command identifiers, 556
 - controlling table pivots, 558, 569
 - Excel format, 558
 - excluding output from the Viewer, 563
 - output object types, 554
 - PDF format, 558
 - SAV file format, 558, 564
 - SPSS Statistics data file format, 558, 564
 - table subtypes, 556
 - text format, 558
 - using XSLT with OXML, 576
 - variable names in SAV files, 571
 - Word format, 558
 - XML, 558, 571
- One-Sample Kolmogorov-Smirnov Test, 429
 - command additional features, 431
 - missing values, 430
 - options, 430
 - statistics, 430
 - test distribution, 429
- One-Sample T Test, 311
 - command additional features, 313
 - confidence intervals, 313
 - missing values, 313
 - options, 313
- One-Way ANOVA, 314
 - command additional features, 319
 - contrasts, 315
 - factor variables, 314
 - missing values, 318
 - multiple comparisons, 316
 - options, 318
 - polynomial contrasts, 315
 - post hoc tests, 316
 - statistics, 318
- online Help, 10
 - Statistics Coach, 8
- opening files, 12–16, 28
 - controlling default file locations, 506
 - data files, 12–13
 - dBASE files, 12, 14
 - Excel files, 12, 14
 - Lotus 1-2-3 files, 12
 - spreadsheet files, 12, 14
 - Stata files, 15
 - SYSTAT files, 12
 - tab-delimited files, 12
 - text data files, 28
- options, 494, 496–497, 499–500, 502, 505–507, 510
 - charts, 502
 - currency, 499
 - data, 497
 - general, 494
 - multiple imputation, 512
 - output labels, 500
 - pivot table look, 505
 - scripts, 507
 - syntax editor, 510
 - temporary directory, 506
 - two-digit years, 497
 - Variable View, 499
 - Viewer, 496
- ordinal, 73
 - measurement level, 73, 101
- Ordinal Regression , 353
 - command additional features, 358
 - link, 354
 - location model, 356

- options, 354
 - scale model, 357
 - statistics, 353
- outer join, 20
- outliers
 - in Explore, 283
 - in Linear Regression, 346
 - in TwoStep Cluster Analysis, 407
- outline, 205–206
 - changing levels, 206
 - collapsing, 206
 - expanding, 206
 - in Viewer, 205
- output, 203–205, 209–210, 226, 496
 - alignment, 205, 496
 - centering, 205, 496
 - changing output language, 494
 - copying, 204
 - copying into other applications, 209
 - deleting, 204
 - exporting, 210
 - hiding, 204
 - moving, 204
 - pasting into other applications, 209
 - saving, 226
 - showing, 204
 - Viewer, 203
- Output Management System (OMS), 552, 575
- output object types
 - in OMS, 554
- OXML, 576

- page control
 - in column summary reports, 457
 - in row summary reports, 451
- page numbering, 225
 - in column summary reports, 458
 - in row summary reports, 451
- page setup, 223, 225
 - chart size, 225
 - headers and footers, 223
- Paired-Samples T Test, 310
 - missing values, 311
 - options, 311
 - selecting paired variables, 310
- pane splitter
 - Data Editor, 91
- parallel model
 - in Reliability Analysis, 460, 462
- parameter estimates
 - in GLM Univariate, 329
 - in Ordinal Regression, 355
- Partial Correlations, 335
 - command additional features, 337
 - in Linear Regression, 349
 - missing values, 336
 - options, 336
 - statistics, 336
 - zero-order correlations, 336
- Partial Least Squares Regression, 363
 - export variables, 366
 - model, 365
- partial plots
 - in Linear Regression, 346
- Paste button, 7
- pattern difference measure
 - in Distances, 340
- pattern matrix
 - in Factor Analysis, 395
- PDF
 - exporting output, 210, 217
- PDF format
 - exporting output, 558
- Pearson chi-square
 - in Crosstabs, 289
 - in Ordinal Regression, 355
- Pearson correlation
 - in Bivariate Correlations, 332
 - in Crosstabs, 289
- Pearson residuals
 - in Ordinal Regression, 355
- peers
 - in Nearest Neighbor Analysis, 381
- percentages
 - in Crosstabs, 291
- percentiles
 - in Explore, 283
 - in Frequencies, 274
- performance, 62
 - caching data, 62
- phi
 - in Crosstabs, 289
- phi-square distance measure
 - in Distances, 340
- pie charts
 - in Frequencies, 276
- pivot tables, 204, 209–210, 222, 227–230, 232, 234–237, 239–240, 246–248, 505
 - alignment, 243
 - background color, 242
 - borders, 239
 - captions, 244–245
 - cell formats, 237
 - cell properties, 242–243
 - cell widths, 246
 - changing display order, 228
 - changing the look, 234
 - continuation text, 240
 - controlling table breaks, 248
 - copying into other applications, 209
 - creating charts from tables, 248
 - default column width adjustment, 505
 - default look for new tables, 505
 - deleting group labels, 229

- displaying hidden borders, 246
- editing, 227
- exporting as HTML, 210
- fonts, 242
- footnote properties, 236
- footnotes, 244–245
- general properties, 235
- grid lines, 246
- grouping rows or columns, 229
- hiding, 204
- inserting group labels, 229
- layers, 230
- manipulating, 227
- margins, 243
- moving rows and columns, 228
- pasting as tables, 209
- pasting into other applications, 209
- pivoting, 227
- printing large tables, 248
- printing layers, 222
- properties, 235
- rotating labels, 229
- scaling to fit page, 235, 240
- selecting rows and columns, 247
- showing and hiding cells, 232
- transposing rows and columns, 229
- ungrouping rows or columns, 229
- using icons, 227
- pivoting
 - controlling with OMS for exported output, 569
- PLUM
 - in Ordinal Regression, 353
- PNG files, 210, 221
 - exporting charts, 210, 221
- polynomial contrasts
 - in GLM, 324
 - in One-Way ANOVA, 315
- port numbers, 65
- portable files
 - variable names, 42
- post hoc multiple comparisons, 316
- PostScript files (encapsulated), 210, 221
 - exporting charts, 210, 221
- power estimates
 - in GLM Univariate, 329
- power model
 - in Curve Estimation, 360
- PowerPoint, 216
 - exporting output as PowerPoint, 216
- PowerPoint format
 - exporting output, 210
- predicted values
 - saving in Curve Estimation, 361
 - saving in Linear Regression, 347
- prediction intervals
 - saving in Curve Estimation, 361
 - saving in Linear Regression, 347
- price-related differential (PRD)
 - in Ratio Statistics, 473
- principal axis factoring, 398
- principal components analysis, 395, 398
- printing, 92, 222–223, 225, 235, 240, 248
 - chart size, 225
 - charts, 222
 - controlling table breaks, 248
 - data, 92
 - headers and footers, 223
 - layers, 222, 235, 240
 - models, 252
 - page numbers, 225
 - pivot tables, 222
 - print preview, 222
 - scaling tables, 235, 240
 - space between output items, 225
 - text output, 222
- prior moving average function, 160
- Production Facility, 494
 - converting files to production jobs, 551
 - using command syntax from journal file, 494
- Production jobs, 544, 548–549
 - command line switches, 549
 - converting Production Facility files, 551
 - exporting charts, 544
 - output files, 544
 - running multiple production jobs, 549
 - scheduling production jobs, 549
 - substituting values in syntax files, 548
 - syntax rules, 544
- profile plots
 - in GLM, 325
- programming with command language, 253
- properties, 235
 - pivot tables, 235
 - tables, 235
- proportion estimates
 - in Rank Cases, 141
- Proximities
 - in Hierarchical Cluster Analysis, 412
- Python
 - autoscripts, 585
 - scripts, 582, 584
- quadrant map
 - in Nearest Neighbor Analysis, 382
- quadratic model
 - in Curve Estimation, 360
- Quancept, 37
- Quanvert, 37
- quartiles
 - in Frequencies, 274
- quartimax rotation
 - in Factor Analysis, 399

- r* correlation coefficient
 - in Bivariate Correlations, 332
 - in Crosstabs, 289
- R* statistic
 - in Linear Regression, 349
 - in Means, 300
- R-E-G-W *F*
 - in GLM, 326
 - in One-Way ANOVA, 316
- R-E-G-W *Q*
 - in GLM, 326
 - in One-Way ANOVA, 316
- R^2
 - in Linear Regression, 349
 - in Means, 300
 - R^2 change, 349
- random number seed, 128
- random sample, 22
 - databases, 22
 - random number seed, 128
 - selecting, 182
- range
 - in Descriptives, 279
 - in Frequencies, 274
 - in Means, 300
 - in OLAP Cubes, 303
 - in Ratio Statistics, 473
 - in Summarize, 296
- rank correlation coefficient
 - in Bivariate Correlations, 332
- ranking cases, 140
 - fractional ranks, 141
 - percentiles, 141
 - Savage scores, 141
 - tied values, 142
- Rankit estimates, 141
- Rao's *V*
 - in Discriminant Analysis, 391
- Ratio Statistics, 471
 - statistics, 473
- recoding values, 116, 132–133, 135, 137
- reference category
 - in GLM, 324
- regression
 - Linear Regression, 343
 - multiple regression, 343
 - plots, 346
- regression coefficients
 - in Linear Regression, 349
- related samples, 434, 438
- relative risk
 - in Crosstabs, 289
- Reliability Analysis, 460
 - ANOVA table, 462
 - command additional features, 463
 - descriptives, 462
 - example, 460
 - Hotelling's T^2 , 462
 - inter-item correlations and covariances, 462
 - intraclass correlation coefficient, 462
 - Kuder-Richardson 20, 462
 - statistics, 460, 462
 - Tukey's test of additivity, 462
- remote servers, 64–65, 67–69
 - adding, 65
 - available procedures, 68
 - data file access, 67
 - editing, 65
 - logging in, 64
 - relative paths, 69
- removing group labels, 229
- renaming datasets, 95
- reordering rows and columns, 228
- repeated contrasts
 - in GLM, 324
- replacing missing values
 - linear interpolation, 162
 - linear trend, 162
 - mean of nearby points, 162
 - median of nearby points, 162
 - series mean, 162
- Report Summaries in Columns, 454
 - column format, 449
 - command additional features, 458
 - grand total, 458
 - missing values, 458
 - page control, 457
 - page layout, 452
 - page numbering, 458
 - subtotals, 457
 - total columns, 456
- Report Summaries in Rows, 448
 - break columns, 448
 - break spacing, 451
 - column format, 449
 - command additional features, 458
 - data columns, 448
 - footers, 453
 - missing values, 451
 - page control, 451
 - page layout, 452
 - page numbering, 451
 - sorting sequences, 448
 - titles, 453
 - variables in titles, 453
- reports
 - column summary reports, 454
 - comparing columns, 456
 - composite totals, 456
 - dividing column values, 456
 - multiplying column values, 456
 - row summary reports, 448
 - total columns, 456
- Reset button, 7

- residual plots
 - in GLM Univariate, 329
- residuals
 - in Crosstabs, 291
 - saving in Curve Estimation, 361
 - saving in Linear Regression, 347
- restructuring data, 184–185, 188–190, 192–197, 199, 201
 - and weighted data, 201
 - creating a single index variable for variables to cases, 194
 - creating index variables for variables to cases, 192
 - creating multiple index variables for variables to cases, 195
 - example of cases to variables, 188
 - example of one index for variables to cases, 193
 - example of two indices for variables to cases, 194
 - example of variables to cases, 188
 - options for cases to variables, 199
 - options for variables to cases, 196
 - overview, 184
 - selecting data for cases to variables, 197
 - selecting data for variables to cases, 190
 - sorting data for cases to variables, 199
 - types of restructuring, 185
 - variable groups for variables to cases, 189
- rho
 - in Bivariate Correlations, 332
 - in Crosstabs, 289
- risk
 - in Crosstabs, 289
- ROC Curve, 475
 - statistics and plots, 477
- rotating labels, 229
- row percentages
 - in Crosstabs, 291
- rows, 247
 - selecting in pivot tables, 247
- running median function, 160
- Runs Test
 - command additional features, 429
 - cut points, 427–428
 - missing values, 428
 - options, 428
 - statistics, 428
- Ryan-Einot-Gabriel-Welsch multiple F
 - in GLM, 326
 - in One-Way ANOVA, 316
- Ryan-Einot-Gabriel-Welsch multiple range
 - in GLM, 326
 - in One-Way ANOVA, 316
- S model
 - in Curve Estimation, 360
- S-stress
 - in Multidimensional Scaling, 465
- sampling
 - random sample, 182
- SAS files
 - opening, 12
 - saving, 41
- SAV file format
 - routing output to an SPSS Statistics data file, 564
 - routing output to SPSS Statistics data file, 558
- Savage scores, 141
- saving charts, 210, 221
 - BMP files, 210, 221
 - EMF files, 210
 - EPS files, 210, 221
 - JPEG files, 210, 221
 - metafiles, 210
 - PICT files, 210
 - PNG files, 221
 - PostScript files, 221
 - TIFF files, 221
- saving files, 41–42
 - controlling default file locations, 506
 - data files, 41–42
 - database file queries, 27
 - SPSS Statistics data files, 41
- saving output, 210, 216, 219
 - Excel format, 210, 214
 - HTML, 210, 212
 - HTML format, 210
 - PDF format, 210, 217
 - PowerPoint format, 210, 216
 - text format, 210, 219
 - Word format, 210, 213
- scale, 73
 - in Multidimensional Scaling, 465
 - in Reliability Analysis, 460
 - measurement level, 73, 101
- scale model
 - in Ordinal Regression, 357
- scale variables
 - binning to create categorical variables, 116
- scaling
 - pivot tables, 235, 240
- scatterplots
 - in Linear Regression, 346
- Scheffé test
 - in GLM, 326
 - in One-Way ANOVA, 316
- scientific notation, 74, 494
 - suppressing in output, 494
- scoring
 - displaying loaded models, 166
 - loading saved models, 164
 - models supported for export and scoring, 163
- scripts, 514, 518, 578
 - adding to menus, 514
 - autoscripts, 579
 - Basic, 588
 - creating, 578
 - default language, 507, 578

- editing, 578
- languages, 578
- Python, 582, 584
- running, 578
- running with toolbar buttons, 518
- startup scripts, 592
- search and replace
 - viewer documents, 208
- seasonal difference function, 160
- select cases, 180
- selecting cases, 180
 - based on selection criteria, 181
 - date range, 183
 - random sample, 182
 - range of cases, 183
 - time range, 183
- selection methods, 247
 - selecting rows and columns in pivot tables, 247
- selection variable
 - in Linear Regression, 345
- servers, 64–65
 - adding, 65
 - editing, 65
 - logging in, 64
 - names, 65
 - port numbers, 65
- session journal, 506
- Shapiro-Wilk's test
 - in Explore, 284
- Shift Values, 131
- showing, 204, 233, 245, 515
 - captions, 245
 - dimension labels, 233
 - footnotes, 245
 - results, 204
 - rows or columns, 233
 - titles, 233
 - toolbars, 515
- Sidak's *t* test
 - in GLM, 326
 - in One-Way ANOVA, 316
- sign test
 - in Two-Related-Samples Tests, 434
- similarity measures
 - in Distances, 341
 - in Hierarchical Cluster Analysis, 413
- simple contrasts
 - in GLM, 324
- size difference measure
 - in Distances, 340
- sizes, 206
 - in outline, 206
- skewness
 - in Descriptives, 279
 - in Explore, 283
 - in Frequencies, 274
 - in Means, 300
- in OLAP Cubes, 303
- in Report Summaries in Columns, 455
- in Report Summaries in Rows, 450
- in Summarize, 296
- smoothing function, 160
- Somers' *d*
 - in Crosstabs, 289
- sorting
 - variables, 168
- sorting cases, 167
- sorting variables, 168
- space-delimited data, 28
- Spearman correlation coefficient
 - in Bivariate Correlations, 332
 - in Crosstabs, 289
- Spearman-Brown reliability
 - in Reliability Analysis, 462
- speed, 62
 - caching data, 62
- spelling, 83
 - dictionary, 497
- split-file processing, 179
- split-half reliability
 - in Reliability Analysis, 460, 462
- splitting tables, 248
 - controlling table breaks, 248
- spp files
 - converting to spj files, 551
- spread-versus-level plots
 - in Explore, 284
 - in GLM Univariate, 329
- spreadsheet files, 12–14, 43
 - opening, 14
 - reading ranges, 13
 - reading variable names, 13
 - writing variable names, 43
- SPSS Statistics data file format
 - routing output to a data file, 558, 564
- SPSSTMPDIR environment variable, 506
- squared Euclidean distance
 - in Distances, 340
- standard deviation
 - in Descriptives, 279
 - in Explore, 283
 - in Frequencies, 274
 - in GLM Univariate, 329
 - in Means, 300
 - in OLAP Cubes, 303
 - in Ratio Statistics, 473
 - in Report Summaries in Columns, 455
 - in Report Summaries in Rows, 450
 - in Summarize, 296
- standard error
 - in Descriptives, 279
 - in Explore, 283
 - in Frequencies, 274
 - in GLM, 328–329

- in ROC Curve, 477
- standard error of kurtosis
 - in Means, 300
 - in OLAP Cubes, 303
 - in Summarize, 296
- standard error of skewness
 - in Means, 300
 - in OLAP Cubes, 303
 - in Summarize, 296
- standard error of the mean
 - in Means, 300
 - in OLAP Cubes, 303
 - in Summarize, 296
- standardization
 - in TwoStep Cluster Analysis, 407
- standardized residuals
 - in GLM, 328
 - in Linear Regression, 347
- standardized values
 - in Descriptives, 278
- Stata files, 15
 - opening, 15
 - reading, 12
 - saving, 41
- Statistics Coach, 8
- status bar, 5
- stem-and-leaf plots
 - in Explore, 284
- stepwise selection
 - in Linear Regression, 344
- stress
 - in Multidimensional Scaling, 465
- strictly parallel model
 - in Reliability Analysis, 460, 462
- string format, 74
- string variables, 77, 85
 - breaking up long strings in earlier releases, 42
 - computing new string variables, 126
 - entering data, 85
 - in dialog boxes, 5
 - missing values, 77
 - recoding into consecutive integers, 137
- Student-Newman-Keuls
 - in GLM, 326
 - in One-Way ANOVA, 316
- Studentized residuals
 - in Linear Regression, 347
- Student's *t* test, 307
- subgroup means, 298, 302
- subsets of cases
 - random sample, 182
 - selecting, 180–181, 183
- subtitles
 - charts, 485
- subtotals
 - in column summary reports, 457
- subtypes, 556
 - vs. labels, 557
- sum
 - in Descriptives, 279
 - in Frequencies, 274
 - in Means, 300
 - in OLAP Cubes, 303
 - in Summarize, 296
- sum of squares, 323
 - in GLM, 322
- Summarize, 294
 - options, 295
 - statistics, 296
- syntax, 253, 263, 517, 544
 - accessing *Command Syntax Reference*, 10
 - journal file, 265
 - output log, 255
 - pasting, 255
 - Production jobs rules, 544
 - running, 263
 - running command syntax with toolbar buttons, 517
 - syntax rules, 253
 - Unicode command syntax files, 264
- syntax converter, 593
- syntax editor, 257
 - auto-completion, 259
 - bookmarks, 258, 262
 - breakpoints, 258, 261, 263
 - color coding, 260
 - command spans, 258
 - commenting out text, 263
 - line numbers, 258
 - options, 510
- SYSTAT files, 12
 - opening, 12
- t* test
 - in GLM Univariate, 329
 - in Independent-Samples T Test, 307
 - in One-Sample T Test, 311
 - in Paired-Samples T Test, 310
- T4253H smoothing, 160
- tab-delimited files, 12–13, 28, 41–43
 - opening, 12
 - reading variable names, 13
 - saving, 41–42
 - writing variable names, 43
- table breaks, 248
- table chart, 248
- table subtypes, 556
 - vs. labels, 557
- TableLooks, 234–235
 - applying, 234
 - creating, 235
- tables, 248
 - alignment, 243
 - background color, 242

- cell properties, 242–243
- controlling table breaks, 248
- converting TABLES command syntax to CTABLES, 593
- fonts, 242
- margins, 243
- TABLES
 - converting TABLES command syntax to CTABLES, 593
- Tamhane's T2
 - in GLM, 326
 - in One-Way ANOVA, 316
- target lists, 492
- tau-*b*
 - in Crosstabs, 289
- tau-*c*
 - in Crosstabs, 289
- templates, 78–79, 485, 502
 - charts, 485
 - in charts, 502
 - using an external data file as a template, 106
 - variable definition, 78–79
- temporary active file, 62
- temporary directory, 506
 - setting location in local mode, 506
 - SPSSTMPDIR environment variable, 506
- temporary disk space, 60, 62
- test of parallel lines
 - in Ordinal Regression, 355
- tests for independence
 - chi-square, 289
- tests for linearity
 - in Means, 300
- Tests for Several Independent Samples, 436
 - command additional features, 438
 - defining range, 437
 - grouping variables, 437
 - missing values, 438
 - options, 438
 - statistics, 438
 - test types, 437
- Tests for Several Related Samples, 438
 - command additional features, 440
 - statistics, 440
 - test types, 439
- text, 28, 207, 210, 219
 - adding a text file to the Viewer, 207
 - adding to Viewer, 207
 - data files, 28
 - exporting output as text, 210, 219, 558
- TIFF files, 221
 - exporting charts, 210, 221
- time series analysis
 - forecast, 361
 - predicting cases, 361
- time series data
 - creating new time series variables, 159
 - data transformations, 157
 - defining date variables, 158
 - replacing missing values, 161
 - transformation functions, 160
- titles, 207
 - adding to Viewer, 207
 - charts, 485
 - in OLAP Cubes, 306
- tolerance
 - in Linear Regression, 349
- toolbars, 515, 517–518
 - creating, 515, 517
 - creating new tools, 518
 - customizing, 515, 517
 - displaying in different windows, 517
 - showing and hiding, 515
- total column
 - in reports, 456
- total percentages
 - in Crosstabs, 291
- training sample
 - in Nearest Neighbor Analysis, 373
- transformation matrix
 - in Factor Analysis, 395
- transposing rows and columns, 229
- transposing variables and cases, 169
- tree depth
 - in TwoStep Cluster Analysis, 407
- trigger events, 580
 - autoscripts, 580
- trimmed mean
 - in Explore, 283
- Tukey estimates, 141
- Tukey's *b* test
 - in GLM, 326
 - in One-Way ANOVA, 316
- Tukey's biweight estimator
 - in Explore, 283
- Tukey's honestly significant difference
 - in GLM, 326
 - in One-Way ANOVA, 316
- Tukey's test of additivity
 - in Reliability Analysis, 460, 462
- Two-Independent-Samples Tests, 431
 - command additional features, 434
 - defining groups, 433
 - grouping variables, 433
 - missing values, 433
 - options, 433
 - statistics, 433
 - test types, 432
- Two-Related-Samples Tests, 434
 - command additional features, 436
 - missing values, 435
 - options, 435
 - statistics, 435
 - test types, 435

- two-sample *t* test
 - in Independent-Samples T Test, 307
- TwoStep Cluster Analysis, 404
 - options, 407
 - plots, 409
 - save to external file, 410
 - save to working file, 410
 - statistics, 410
- uncertainty coefficient
 - in Crosstabs, 289
- Unicode, 12, 41
- Unicode command syntax files, 264
- unstandardized residuals
 - in GLM, 328
- unweighted least squares
 - in Factor Analysis, 398
- user-missing values, 77
- V*
 - in Crosstabs, 289
- value labels, 76, 85, 91, 97, 500
 - applying to multiple variables, 102
 - copying, 102
 - in Data Editor, 91
 - in merged data files, 173
 - in outline pane, 500
 - in pivot tables, 500
 - inserting line breaks, 77
 - saving in Excel files, 41
 - using for data entry, 85
- Van der Waerden estimates, 141
- variable attributes, 78–79
 - copying and pasting, 78–79
 - custom, 80
- variable importance
 - in Nearest Neighbor Analysis, 380
- variable importance plots
 - in TwoStep Cluster Analysis, 409
- variable information, 488
- variable labels, 76, 494, 500
 - in dialog boxes, 5, 494
 - in merged data files, 173
 - in outline pane, 500
 - in pivot tables, 500
 - inserting line breaks, 77
- variable lists, 492
 - reordering target lists, 492
- variable names, 72, 494
 - generated by OMS, 571
 - in dialog boxes, 5, 494
 - mixed case variable names, 72
 - portable files, 42
 - rules, 72
 - truncating long variable names in earlier releases, 42
 - wrapping long variable names in output, 72
- variable pairs, 185
 - creating, 185
- variable sets, 489–490
 - defining, 489
 - using, 490
- Variable View, 71
 - customizing, 82, 499
- variables, 7, 72, 87, 185, 488–489, 494
 - defining, 72
 - defining variable sets, 489
 - definition information, 488
 - display order in dialog boxes, 494
 - finding in Data Editor, 88
 - in dialog boxes, 5
 - inserting new variables, 87
 - moving, 87
 - recoding, 132–133, 135, 137
 - renaming for merged data files, 173
 - restructuring into cases, 185
 - selecting in dialog boxes, 7
 - sorting, 168
 - variable information in dialog boxes, 8
- variance
 - in Descriptives, 279
 - in Explore, 283
 - in Frequencies, 274
 - in Means, 300
 - in OLAP Cubes, 303
 - in Report Summaries in Columns, 455
 - in Report Summaries in Rows, 450
 - in Summarize, 296
- variance inflation factor
 - in Linear Regression, 349
- varimax rotation
 - in Factor Analysis, 399
- vertical label text, 229
- viewer
 - find and replace information, 208
 - search and replace information, 208
- Viewer, 203–207, 225–226, 496, 500
 - changing outline font, 207
 - changing outline levels, 206
 - changing outline sizes, 206
 - collapsing outline, 206
 - deleting output, 204
 - display options, 496
 - displaying data values, 500
 - displaying value labels, 500
 - displaying variable labels, 500
 - displaying variable names, 500
 - excluding output types with OMS, 563
 - expanding outline, 206
 - hiding results, 204
 - moving output, 204
 - outline, 205
 - outline pane, 203
 - results pane, 203

Index

- saving document, 226
- space between output items, 225
- virtual active file, 60
- Visual Bander, 116

- Wald-Wolfowitz runs
 - in Two-Independent-Samples Tests, 432
- Waller-Duncan *t* test
 - in GLM, 326
 - in One-Way ANOVA, 316
- weighted data, 201
 - and restructured data files, 201
- weighted least squares
 - in Linear Regression, 343
- weighted mean
 - in Ratio Statistics, 473
- weighted predicted values
 - in GLM, 328
- weighting cases, 183
 - fractional weights in Crosstabs, 183
- Welch statistic
 - in One-Way ANOVA, 318
- wide tables
 - pasting into Microsoft Word, 209
- Wilcoxon signed-rank test
 - in Two-Related-Samples Tests, 434
- Wilks' lambda
 - in Discriminant Analysis, 391
- window splitter
 - Data Editor, 91
- windows, 3
 - active window, 4
 - designated window, 4
- Word format
 - exporting output, 210, 213, 558
 - wide tables, 210
- wrapping, 235
 - controlling column width for wrapped text, 235
 - variable and value labels, 77

- XML
 - OXML output from OMS, 576
 - routing output to XML, 558
 - saving output as XML, 552
 - table structure in OXML, 571
- XSLT
 - using with OXML, 576

- Yates' correction for continuity
 - in Crosstabs, 289
- years, 497
 - two-digit values, 497

- z* scores
 - in Descriptives, 278

- in Rank Cases, 141
- saving as variables, 278
- zero-order correlations
 - in Partial Correlations, 336