

البريد الإلكتروني: ibm@ibm.com

IBM

Note

Before using this information and the product it supports, read the information in "Notices" on page 265.

Product Information

This edition applies to version 22, release 0, modification 0 of IBM SPSS Statistics and to all subsequent releases and modifications until otherwise indicated in new editions.

Contents

Chapter 1. Overview 1

What's new in version 22?	1
Windows	2
Designated window versus active window	3
Status Bar	3
Dialog boxes	4
Variable names and variable labels in dialog box lists	4
Resizing dialog boxes	4
Dialog box controls	4
Selecting variables	5
Data type, measurement level, and variable list icons	5
Getting information about variables in dialog boxes	5
Basic steps in data analysis	6
Statistics Coach	6
Finding out more	6

Chapter 2. Getting Help 7

Getting Help on Output Terms	8
--	---

Chapter 3. Data files 9

Opening data files	9
To open data files.	9
Data file types	10
Opening file options	10
Reading Excel Files	10
Reading older Excel files and other spreadsheets	11
Reading dBASE files	11
Reading Stata files	11
Reading Database Files	11
Text Wizard	17
Reading Cognos data	20
Reading IBM SPSS Data Collection Data.	22
File information	24
Saving data files.	24
To save modified data files	24
Saving data files in external formats	24
Saving data files in Excel format	27
Saving data files in SAS format.	27
Saving data files in Stata format	28
Saving Subsets of Variables	29
Encrypting data files	29
Exporting to a Database	30
Exporting to IBM SPSS Data Collection	35
Comparing datasets	36
Compare Datasets: Compare tab	36
Compare Datasets: Attributes tab	37
Comparing datasets: Output tab	37
Protecting original data	38
Virtual Active File	38
Creating a Data Cache.	39

Chapter 4. Distributed Analysis Mode 41

Server Login	41
Adding and Editing Server Login Settings	41
To Select, Switch, or Add Servers	42

Searching for Available Servers	43
Opening Data Files from a Remote Server	43
File Access in Local and Distributed Analysis Mode	43
Availability of Procedures in Distributed Analysis Mode	44
Absolute versus Relative Path Specifications	44

Chapter 5. Data Editor 47

Data View	47
Variable View.	47
To display or define variable attributes	48
Variable names	48
Variable measurement level	49
Variable type	49
Variable labels	51
Value labels	51
Inserting line breaks in labels	51
Missing values	51
Roles	52
Column width	52
Variable alignment	52
Applying variable definition attributes to multiple variables	52
Custom Variable Attributes	53
Customizing Variable View	55
Spell checking	56
Entering data.	56
To enter numeric data	56
To enter non-numeric data	57
To use value labels for data entry	57
Data value restrictions in the data editor	57
Editing data	57
Replacing or modifying data values	57
Cutting, copying, and pasting data values	57
Inserting new cases.	58
Inserting new variables	58
To change data type	59
Finding cases, variables, or imputations	59
Finding and replacing data and attribute values	60
Obtaining Descriptive Statistics for Selected Variables	60
Case selection status in the Data Editor	61
Data Editor display options	61
Data Editor printing	61
To print Data Editor contents	62

Chapter 6. Working with Multiple Data Sources 63

Basic Handling of Multiple Data Sources	63
Working with Multiple Datasets in Command Syntax	63
Copying and Pasting Information between Datasets	63
Renaming Datasets	64
Suppressing Multiple Datasets	64

Chapter 7. Data preparation 65

Variable properties	65
Defining Variable Properties	65
To Define Variable Properties	66
Defining Value Labels and Other Variable Properties	66
Assigning the Measurement Level	67
Custom Variable Attributes	68
Copying Variable Properties	68
Setting measurement level for variables with unknown measurement level	68
Multiple Response Sets	69
Defining Multiple Response Sets	69
Copy Data Properties	71
Copying Data Properties	71
Identifying Duplicate Cases	74
Visual Binning	75
To Bin Variables	75
Binning Variables	76
Automatically Generating Binned Categories	77
Copying Binned Categories	78
User-Missing Values in Visual Binning	79

Chapter 8. Data Transformations 81

Data Transformations	81
Computing Variables	81
Compute Variable: If Cases	81
Compute Variable: Type and Label	82
Functions	82
Missing Values in Functions	82
Random Number Generators	83
Count Occurrences of Values within Cases	83
Count Values within Cases: Values to Count	83
Count Occurrences: If Cases	83
Shift Values	84
Recoding Values	84
Recode into Same Variables	84
Recode into Same Variables: Old and New Values	85
Recode into Different Variables	86
Recode into Different Variables: Old and New Values	86
Automatic Recode	87
Rank Cases	88
Rank Cases: Types	89
Rank Cases: Ties	89
Date and Time Wizard	90
Dates and Times in IBM SPSS Statistics	90
Create a Date/Time Variable from a String	91
Create a Date/Time Variable from a Set of Variables	91
Add or Subtract Values from Date/Time Variables	92
Extract Part of a Date/Time Variable	94
Time Series Data Transformations	94
Define Dates	94
Create Time Series	95
Replace Missing Values	97

Chapter 9. File handling and file transformations 99

File handling and file transformations	99
Sort cases	99
Sort variables	100
Transpose	101
Merging Data Files	101
Add Cases	101
Add Variables	102
Aggregate Data	104
Aggregate Data: Aggregate Function	105
Aggregate Data: Variable Name and Label	105
Split file	105
Select cases	106
Select cases: If	107
Select cases: Random sample	107
Select cases: Range	107
Weight cases	107
Restructuring Data	108
To Restructure Data	108
Restructure Data Wizard: Select Type	108
Restructure Data Wizard (Variables to Cases): Number of Variable Groups	111
Restructure Data Wizard (Variables to Cases): Select Variables	111
Restructure Data Wizard (Variables to Cases): Create Index Variables	112
Restructure Data Wizard (Variables to Cases): Create One Index Variable	113
Restructure Data Wizard (Variables to Cases): Create Multiple Index Variables	114
Restructure Data Wizard (Variables to Cases): Options	114
Restructure Data Wizard (Cases to Variables): Select Variables	115
Restructure Data Wizard (Cases to Variables): Sort Data	115
Restructure Data Wizard (Cases to Variables): Options	115
Restructure Data Wizard: Finish	116

Graphics only options	129
Graphics format options	129
Viewer printing	130

Chart definition options	175
Adding and Editing Titles and Footnotes	175
Setting General Options	176

Chapter 20. Creating and Managing Custom Dialogs 209

Chapter 16. Scoring data with predictive models 177

Scoring Wizard	177
Matching model fields to dataset fields	178
Selecting scoring functions	179
Scoring the active dataset	180
Merging model and transformation XML files	180

Chapter 17. Utilities 183

Utilities	183
Variable information	183
Data file comments	183
Variable sets	184
Defining variable sets	184
Using variable sets to show and hide variables	184
Reordering target variable lists	185
Extension bundles	185
Creating and editing extension bundles	185
Installing local extension bundles	188
Viewing installed extension bundles	191
Downloading extension bundles	191

Chapter 18. Options 193

Options	193
General options	193
Viewer Options	194
Data Options	195
Changing the default variable view	196
Language options	196
Currency options	197
To create custom currency formats	198
Output options	198
Chart options	198
Data Element Colors	199
Data Element Lines	199
Data Element Markers	199
Data Element Fills	200
Pivot table options	200
File locations options	202
Script options	203
Syntax editor options	203
Multiple imputations options	204

Chapter 19. Customizing Menus and Toolbars 207

Customizing Menus and Toolbars	207
Menu Editor	207
Customizing Toolbars	207
Show Toolbars	207
To Customize Toolbars	207
Toolbar Properties	208
Edit Toolbar	208
Create New Tool	208

m o a a b i o o e t a e e t a e r E h o w i e d o o g T D 1 . 0 8 1 3 T c [(.) 5 8 5 . 6 (2) 1 0 8 3 . 3 4 0 6 1 0 8 3 . 3 (7)] J - 9 . 7 5 3 9 - 1 . 2 C r e a t e d 0 7 / 0 8 / 0 7 1 3 : 0 8 3 3 B e [(2 2 2 5 T 3 4 5 8 5 (6 2) 1 0 8 3 3 4 0 6 1 0 8 3 3 (7)] J - 9

Copying OMS identifiers from the viewer

Chapter 1. Overview

What's new in version 22?

Automated output modification

Automated output modification applies formatting and other changes to the contents of the active Viewer window. Changes that can be applied include:

- All or selected viewer objects
- Selected types of output objects (for example, charts, logs, pivot tables)
- Pivot table content that is based on conditional expressions
- Outline (navigation) pane content

The types of changes you can make include:

- Delete objects
- Index objects (add a sequential numbering scheme)
- Change the visible property of objects
- Change the outline label text
- Transpose rows and columns in pivot tables
- Change the selected layer of pivot tables
- Change the formatting of selected areas or specific cells in a pivot table based on conditional expressions. For example, make all significance values less than 0.05 bold.

For more information, see the topic Chapter 13, “Automated Output Modification,” on page 153.

Web reports

A web report is an interactive document that is compatible with most browsers. Many of the interactive features of pivot tables available in the Viewer are also available in web reports. You can distribute results in a single HTML file that can be viewed and explored on most mobile devices.

For more information, see the topic “Web report options” on page 125.

Simulation enhancements

- You can now simulate data without a predictive model.
- Associations between categorical fields can now be captured from historical data and used when simulating data for those fields.
- There is now full support for simulating categorical string fields.

Nonparametric tests (NPTESTS) enhancements

- As an alternative to model viewer output, you can create pivot table and chart output for nonparametric tests.
- Ordinal fields can now be used in nonparametric tests.

Essentials for Python

IBM® SPSS® Statistics - Essentials for Python is now installed by default with IBM SPSS Statistics, and it also now includes Python 2.7 for all supported operating systems. By default, the IBM SPSS Statistics - Integration Plug-in for Python uses the version of Python 2.7 that is installed by IBM SPSS Statistics, but

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 typeional tables,

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. See the topic “General options” on page 193 for more information.

You can also change the variable list display attributes within dialogs. The method for changing the display attributes depends on the dialog:

- If the dialog provides sorting and display controls above the source variable list, use those controls to change the display attributes.
- If the dialog does not contain sorting controls above the source variable list, right-click any variable in the source list and select the display attributes from the pop-up 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. (In dialogs with sorting controls above the source variable list, the default selection of **None** sorts the list in file order.)

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.

Dialog box controls

There are five standard controls in most dialog boxes:

OK or Run. Runs the procedure. After you select your variables and choose any additional specifications, click **OK** to run the procedure and close the dialog box. Some dialogs have a **Run** button instead of the **OK** button.

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.

Table 1. Measurement level icons

	Numeric	String	Date	Time
Scale (Continuous)		n/a		
Ordinal				
Nominal				

- For more information on measurement level, see “Variable measurement level” on page 49.
- For more information on numeric, string, date, and time data types, see “Variable type” on page 49.

Getting information about variables in dialog boxes

Many dialogs provide the ability to find out more about the variables displayed in the variable lists.

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

Chapter 2. 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.
- **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, when you click it, directly opens a Help dialog box that provides information about the dialog box and its controls.

Chapter 3. 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 IBM SPSS Statistics format created on other operating systems
- SYSTAT data files. SYSTAT SYZ files are not supported.
- SAS data files
- Stata data files
- IBM Cognos Business Intelligence data packages and list reports

Opening data files

In addition to files saved in IBM 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. See the topic Chapter 6, “Working with Multiple Data Sources,” on page 63 for more information.
- 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. See the topic Chapter 4, “Distributed Analysis Mode,” on page 41 for more information.

To open data files

1. 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. See the topic “General options” on page 193 for more information.
- 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 page 11. For information on reading data from text data files, see “Text Wizard” on page 17. For information on reading IBM Cognos data, see “Reading Cognos data” on page 20.

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 IBM 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 IBM 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*, ..., and so forth).
- **Variable labels.** Stata variable labels are converted to IBM SPSS Statistics variable labels.
- **Value labels.** Stata value labels are converted to IBM 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 IBM SPSS Statistics format (d-m-y) values. Stata "time-series" date format values (weeks, months, quarters, and so on) are converted to simple numeric (F) format, preserving the original, internal integer value, which is the number of weeks, months, quarters, and so on, 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 IBM SPSS Statistics Server), the drivers must be installed on the remote server. See the topic Chapter 4, "Distributed Analysis Mode," on page 41 for more information.

Note: If you are running the Windows 64-bit version of IBM SPSS Statistics, you cannot read Excel, Access, or dBASE database sources, even though they may appear on the list of available database sources. The 32-bit ODBC drivers for these products are not compatible.

To Read Database Files

1. 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.

Connection Pooling

If you access the same database source multiple times in the same session or job, you can improve performance with connection pooling.

1. In the last step of the wizard, paste the command syntax into a syntax window.

- . At the end of the quoted string, add `OPTIONAL`.

To Edit Saved Database Queries

1. 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.

- In distributed analysis mode (available with IBM SPSS Statistics Server), this button is not available. To add data sources in distributed analysis mode, see your system administrator.

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.

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.

Computing New Fields

If you are in distributed mode, connected to a remote server (available with IBM SPSS Statistics Server), you can compute new fields before you read the data into IBM SPSS Statistics.

You can also compute new fields after you read the data into IBM SPSS Statistics, but computing new fields in the database can save time for large data sources.

New Field Name. The name must comply with IBM SPSS Statistics variable name rules.

Expression. Enter the expression to compute the new field. You can drag existing field names from the **Fields** list and functions from the **Functions** list.

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.

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

1. 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 clauses with expressions for case selection, dates and times in expressions

Note: If you use random sampling, aggregation (available in distributed mode with IBM 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.

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**).

Date and time values must be entered in special manner:

- Date values must use the general form yyyy-mm-dd.
- Time values must use the general form: hh:mm:ss.
- Date/time values (timestamps) must use the general form yyyy-mm-dd hh:mm:ss.

Aggregating Data

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

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

1. 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 IBM 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.

Sorting Cases

Show me

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

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

Text Wizard: Step 3 (Fixed-Width Files)

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?

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.

The default format is determined from the data values in the first 250 rows. If more than one format (e.g., numeric, date, string) is encountered in the first 250 rows, the default format is set to string.

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) in the first 250 rows of the file. 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

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 Cognos data

If you have access to an IBM Cognos Business Intelligence server, you can read IBM Cognos Business Intelligence data packages and list reports into IBM SPSS Statistics.

To read IBM Cognos Business Intelligence data:

1. From the menus choose:

File > Read Cognos Data

- . Specify the URL for the IBM Cognos Business Intelligence server connection.
- . Specify the location of the data package or report.
- . Select the data fields or report that you want to read.

Optionally, you can:

- Select filters for data packages.
- Import aggregated data instead of raw data.
- Specify parameter values.

Mode. Specifies the type of information you want to read: **Data** or **Report**. The only type of report that can be read is a list report.

Connection. The URL of the Cognos Business Intelligence server. Click the **Edit** button to define the details of a new Cognos connection from which to import data or reports. See the topic “Cognos connections” for more information.

Location. The location of the package or report that you want to read. Click the **Edit** button to display a list of available sources from which to import content. See the topic “Cognos location” on page 22 for more information.

Content. For data, displays the available data packages and filters. For reports, display the available reports.

Fields to import. For data packages, select the fields you want to include and move them to this list.

Report to import. For reports, select the list report you want to import. The report must be a list report.

Filters to apply. For data packages, select the filters you want to apply and move them to this list.

Parameters. If this button is enabled, the selected object has parameters defined. You can use parameters to make adjustments (for example, perform a parameterized calculation) before importing the data. If parameters are defined but no default is provided, the button displays a warning triangle.

Aggregate data before performing import. For data packages, if aggregation is defined in the package, you can import the aggregated data instead of the raw data.

Cognos connections

The Cognos Connections dialog specifies the IBM Cognos Business Intelligence server URL and any required additional credentials.

Cognos server URL. The URL of the IBM Cognos Business Intelligence server. This is the value of the "external dispatcher URI" environment property of IBM Cognos Configuration on the server. Contact your system administrator for more information

Mode. Select **Set Credentials** if you need to log in with a specific namespace, username and password (for example, as an administrator). Select **Use Anonymous connection** to log in with no user credentials, in which case you do not fill in the other fields.

Namespace ID. The security authentication provider used to log on to the server. The authentication provider is used to define and maintain users, groups, and roles, and to control the authentication process.

User name. Enter the user name with which to log on to the server.

Password. Enter the password associated with the specified user name.

Save as Default. Saves these settings as your default, to avoid having to re-enter them each time.

Cognos location

The Specify Location dialog box enables you to select a package from which to import data, or a package or folder from which to import reports. It displays the public folders that are available to you. If you select **Data** in the main dialog, the list will display folders containing data packages. If you select **Report** in the main dialog, the list will display folders containing list reports. Select the location you want by navigating through the folder structure.

Specifying parameters for data or reports

If parameters have been defined, either for a data object or a report, you can specify values for these parameters before importing the data or report. An example of parameters for a report would be start and end dates for the report contents.

Name. The parameter name as it is specified in the IBM Cognos Business Intelligence database.

Type. A description of the parameter.

Value. The value to assign to the parameter. To enter or edit a value, double-click its cell in the table. Values are not validated here; any invalid values are detected at run time.

Automatically remove invalid parameters from table. This option is selected by default and will remove any invalid parameters found within the data object or report.

Changing variable names

For IBM Cognos Business Intelligence data packages, package field names are automatically converted to valid variable names. You can use the Fields tab of the Read Cognos Data dialog to override the default names. Names must be unique and must conform to variable naming rules. See the topic “Variable names” on page 48 for more information.

Reading IBM SPSS Data Collection Data

On Microsoft Windows operating systems, you can read data from IBM SPSS Data Collection products. (*Note:* This feature is only available with IBM SPSS Statistics installed on Microsoft Windows operating systems.)

To read IBM SPSS Data Collection data sources, 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>.
- IBM SPSS Data Collection Survey Reporter Developer Kit. An installable version of IBM SPSS Data Collection Survey Reporter Developer Kit is available with the installation media.

You can read IBM SPSS Data Collection data sources only in local analysis mode. This feature is not available in distributed analysis mode using IBM SPSS Statistics Server.

To read data from a IBM SPSS Data Collection data source:

1. In any open IBM SPSS Statistics window, from the menus choose:
 - File > Open IBM SPSS Data Collection 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 IBM SPSS Data Collection Data Import dialog box, select the variables that you want to include and select any case selection criteria.

. Click **OK**

- **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.

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

1. 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 IBM 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

1. 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.

Saving data files in external formats

1. 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:

1. Click **Write variable names to spreadsheet** in the Save Data As dialog box.

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

1. 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):

1. 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 page 30.

Saving data: Data file types

You can save data in the following formats:

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

- Data files saved in IBM SPSS Statistics format cannot be read by versions of the software prior to version 7.5. Data files saved in Unicode encoding cannot be read by releases of IBM SPSS Statistics prior to version 16.0 See the topic “General options” on page 193 for more information.
- 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.

IBM SPSS Statistics Compressed (*.zsav). Compressed IBM SPSS Statistics format.

- ZSAV files have the same features as SAV files, but they take up less disk space.
- ZSAV files may take more or less time to open and save, depending on the file size and system configuration. Extra time is needed to de-compress and compress ZSAV files. However, because ZSAV files are smaller on disk, they reduce the time needed to read and write from disk. As the file size gets larger, this time savings surpasses the extra time needed to de-compress and compress the files.
- Only IBM SPSS Statistics version 21 or higher can open ZSAV files.
- The option to save the data file with your local code page encoding is not available for ZSAV files. These files are always saved in UTF-8 encoding.

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. This format is available only on Windows operating systems.

IBM SPSS Statistics Portable (*.por). Portable format that can be read by other versions of IBM 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 IBM SPSS Statistics data files should be platform/operating system independent. You cannot save data files in portable file in Unicode mode. See the topic “General options” on page 193 for more information.

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.) You can save files in Unicode encoding or local code page encoding.

Comma-delimited (*.csv). Text files with values separated by commas or semicolons. If the current IBM 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. You can save files in Unicode encoding or local code page encoding.

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. You can save files in Unicode encoding or local code page encoding.

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 v9+ Windows (*.sas7bdat). SAS versions 9 for Windows. You can save files in Unicode (UTF-8) or local code page encoding.

SAS v9+ UNIX (*.sas7bdat). SAS versions 9 for UNIX. You can save files in Unicode (UTF-8) or local code page encoding.

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

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

SAS v7-8 for UNIX (*.sas7bdat). 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).

Note: SAS data file names can be up to 32 characters in length. Blank spaces and non-alphanumeric characters other than the underscore ("_") are not allowed and names have to start with a letter or an underscore, numbers can follow.

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 IBM SPSS Statistics and the exported data in Excel.

Table 2. How Excel data formats map to IBM SPSS Statistics variable types and formats

IBM 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 IBM 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.
- IBM 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.

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 IBM 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 included with the installation media.

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.

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. See the topic “Replacing Values in Existing Fields” on page 33 for more information.
- **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. See the topic “Adding New Fields” on page 34 for more information. 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. See the topic “Appending New Records (Cases)” on page 34 for more information.
- **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. See the topic “Creating a New Table or Replacing a Table” on page 34 for more information.
- **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. See the topic “Creating a New Table or Replacing a Table” on page 34 for more information.

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.

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.

For information on defining a filter condition for case selection, see “Select cases” on page 106.

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:

1. 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.

Note: The variable names and database field names may not be identical (since database field names may contain characters not allowed in IBM 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:

1. 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.
- 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:

1. 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.

For information on field names and data types, see the section on creating database fields from IBM SPSS Statistics variables in “Exporting to a Database” on page 30.

Show existing fields. Select this option to display a list of existing fields. You cannot use this panel in the

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 IBM SPSS Statistics variables in “Exporting to a Database” on page 30.

- For original variables read from the IBM SPSS Data Collection data source, any metadata attributes not recognized by IBM SPSS Statistics are preserved in their original state. For example, IBM SPSS Statistics converts grid variables to regular IBM SPSS Statistics variables, but the metadata that defines these grid variables is preserved when you save the new metadata file.
- If any IBM SPSS Data Collection variables were automatically renamed to conform to IBM SPSS Statistics variable naming rules, the metadata file maps the converted names back to the original IBM SPSS Data Collection 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 IBM SPSS Data Collection 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 IBM SPSS Data Collection.

This feature is only available with IBM SPSS Statistics installed on Microsoft Windows operating systems, and is only available in local analysis mode. This feature is not available in distributed analysis mode using IBM SPSS Statistics Server.

To write IBM SPSS Data Collection metadata files, 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>.
- IBM SPSS Data Collection Survey Reporter Developer Kit. An installable version of IBM SPSS Data Collection Survey Reporter Developer Kit is available with the installation media.

Comparing datasets

Compare Datasets compares the active dataset to another dataset in the current session or an external file in IBM SPSS Statistics format.

To compare datasets

1. Open a data file and make sure it is the active dataset. (You can make a dataset the active dataset by clicking on the Data Editor window for that dataset.)
 - . From the menus choose:
Data > Compare Datasets
 - . Select the open dataset or IBM SPSS Statistics data file that you want to compare to the active dataset.
 - . Select one or more fields (variables) that you want to compare.

Optionally, you can:

- Match cases (records) based on one or more case ID values.
- Compare data dictionary properties (field and value labels, user-missing values, measurement level, etc).
- Create a flag field in the active dataset that identifies mismatched cases.
- Create new datasets that contain only matched cases or only mismatched cases.

Compare Datasets: Compare tab

The Matched fields list displays a list of fields that have the same name and same basic type (string or numeric) in both datasets.

1. Select one or more fields (variables) to compare. The comparison of the two datasets is based on the selected fields only.
 - . To see a list of fields that either do not have matching names or are not the same basic type in both datasets, click **Unmatched fields**. Unmatched fields are excluded from the comparison of the two datasets.

- . Optionally, select one or more case (record) ID fields that identify each case.
- If you specify multiple case ID fields, each unique combination of values identifies a case.
- Both files must be sorted in ascending order of the case ID fields. If the datasets are not already sorted, select (check) **Sort Cases** to sort both datasets in case ID order.
- If you do not include any case ID fields, cases are compared in file order. That is, the first case (record) in the active dataset is compared to the first case in the other dataset, and so on.

Compare Datasets: Unmatched Fields

The Unmatched Fields dialog displays a list of fields (variables) that are considered unmatched in the two datasets. An **unmatched** field is a field that either is missing from one of the datasets or is not the same basic type (string or numeric) in both files. Unmatched fields are excluded from the comparison of the two datasets.

Compare Datasets: Attributes tab

By default, only data values are compared, and field attributes (data dictionary properties) such as value labels, user-missing values, and measurement level, are not compared. To compare field attributes:

1. In the Compare Datasets dialog, click the **Attributes** tab.
 - . Click Compare the Data Dictionaries.
 - . Select the attributes you want to compare.
- **Width.** For numeric fields, the maximum number of characters displayed (digits plus formatting characters, such as currency symbols, grouping symbols, and decimal indicator). For string fields, the maximum number of bytes allowed.
- **Label.** Descriptive field label.
- **Value Label.** Descriptive value labels.
- **Missing.** Defined user-missing values.
- **Columns.** Column width in Data view of the Data Editor.
- **Align.** Alignment in Data view of the Data Editor.
- **Measure.** Measurement level.
- **Role.** Field role.
- **Attributes.** User-defined custom field attributes.

Comparing datasets: Output tab

By default, Compare Datasets creates a new field in the active dataset that identifies mismatches and produces a table that provides details for the first 100 mismatches. You can use the Output tab to change the output options.

Flag mismatches in a new field. A new field that identifies mismatches is created in the active dataset.

- The value of the new field is 1 if there are differences and 0 if all the values are the same. If there are cases (records) in the active dataset that are not present in the other dataset, the value is -1.
- The default name of the new field is *CasesCompare*. You can specify a different field name. The name must conform to field (variable) naming rules. See the topic “Variable names” on page 48 for more information.

Copy matched cases to a new dataset. Creates a new dataset that contains only cases (records) from the active dataset that have matching values in the other dataset. The dataset name must conform to field (variable) naming rules. If the dataset already exists, it will be overwritten.

Copy mismatched cases to a new dataset. Creates a new dataset that contains only cases from the active dataset that have different values in the other dataset. The dataset name must conform to field (variable) naming rules. If the dataset already exists, it will be overwritten.

Limit the case-by-case table. For cases (records) in the active dataset that also exist in the other dataset and also have the same basic type (numeric or string) in both datasets, the case-by-case table provides details on the mismatched values for each case. By default, the table is limited to the first 100 mismatches. You can specify a different value or deselect (clear) this item to include all mismatches.

Protecting original data

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

1. 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.

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

- Reading IBM SPSS Statistics data files
- Merging two or more IBM SPSS Statistics data files
- Reading database tables with the Database Wizard
- Merging IBM 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` or `DATA LIST` commands
- Using the Cache Data facility or the `SAVE` command
- Launching other applications from IBM SPSS Statistics that read the data file (for example, AnswerTree, DecisionTime)

Note: The `SAVE` command provides functionality comparable to `SAVE AS` without creating an entire copy of the data file in temporary disk space. The

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

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.

Chapter 4. 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

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

1. From the menus choose:

File > Switch Server...

To select a default server:

1. In the server list, select the box next to the server that you want to use.

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.

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

If you're not sure if you're using local analysis mode or distributed analysis mode, look at the title bar in

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*, `./data` is not valid; you must specify the entire path, as in:

```
./bin/data
```

Chapter 5. 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 Vie

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 Vie

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 IBM 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

1. 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, postal 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 (continuous) 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.

For new numeric variables created with transformations, data from external sources, and IBM SPSS Statistics data files created prior to version 8, default measurement level is determined by the conditions in the following table. Conditions are evaluated in the order listed in the table. The measurement level for the first condition that matches the data is applied.

Table 6. Rules for determining measurement level

Condition	Measurement Level
All values of a variable are missing	Nominal
Format is dollar or custom-currency	Continuous
Format is date or time (excluding Month and Wkday)	Continuous
Variable contains at least one non-integer value	Continuous
Variable contains at least one negative value	Continuous
Variable contains no valid values less than 10,000	Continuous
Variable has <i>N</i> or more valid, unique values*	Continuous
Variable has no valid values less than 10	Continuous
Variable has less than <i>N</i> valid, unique values*	Nominal

* *N* is a user-specified cut-off value. The default is 24.

- You can change the cutoff value in the Options dialog box. See the topic “Data Options” on page 195 for more information.
- The Define Variable Properties dialog box, available from the Data menu, can help you assign the correct measurement level. See the topic “Assigning the Measurement Level” on page 67 for more information.

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.

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.

Restricted numeric. A variable whose values are restricted to non-negative integers. Values are displayed with leading zeros padded to the maximum width of the variable. Values can be entered in scientific notation.

To define variable type

1. 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**.

Inputers 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 *No* and is not equivalent to *No*.
- 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

1. 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*).

To specify value labels

1. 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.

1. 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 ¶.

The ¶ 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.

To define missing values

1. 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.

Roles

Some dialogs support predefined roles that can be used to pre-select variables for analysis. When you open one of these dialogs, variables that meet the role requirements will be automatically displayed in the destination list(s). Available roles are:

Input. The variable will be used as an input (e.g., predictor, independent variable).

Target. The variable will be used as an output or target (e.g., dependent variable).

Both. The variable will be used as both input and output.

None. The variable has no role assignment.

Partition. The variable will be used to partition the data into separate samples for training, testing, and validation.

Split. Included for round-trip compatibility with IBM SPSS Modeler. Variables with this role are *not* used as split-file variables in IBM SPSS Statistics.

- By default, all variables are assigned the *Input* role. This includes data from external file formats and data files from versions of IBM SPSS Statistics prior to version 18.
- Role assignment only affects dialogs that support role assignment. It has no effect on command syntax.

To assign roles

1. Select the role from the list in the *Role* cell for the variable.

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

- . 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. See the topic “Variable names” on page 48 for more information.
- . 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.

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.

- 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

1. In Variable View, from the menus choose:
View > Customize Variable View...

- . Use the up and down arrow buttons to change the display order of the attributes.

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:

1. Select the Variable View tab in the Data Editor window.
 - . Right-click the *Labels* or *Values* column and from the pop-up 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:

1. 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.

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. See the topic “Creating Custom Variable Attributes” on page 53 for more information.
- Customized display settings are saved with IBM SPSS Statistics data files.
- You can also control the default display and order of attributes in Variable View. See the topic “Changing the default variable view” on page 196 for more information.

To customize Variable View

1. 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.

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:

1. Select the Variable View tab in the Data Editor window.
 - . Right-click the *Labels* or *Values* column and from the pop-up menu choose:
Spelling
or
 - . In Variable View, from the menus choose:
Utilities >

To enter non-numeric data

1. 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**

- 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

1. 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. The Data Editor inserts the system-missing value for all cases in a blank cell in a blank row. For example, if you enter the value 17.9 in the first cell in a blank row, the Data Editor automatically creates a new variable named \$V17.9. The Data Editor also inserts the system-missing value for all cases in a blank cell in a blank row.

- . From the menus choose:

Edit > Insert Variable

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

To move variables

1. 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

1. 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

1. For variables,]TJ-1.9441-2.4TD[(V)91.7(ariables)]TJ/F81Tf0-1.58j/F41Tf1missing value for all cases.
 - 1.

Finding and replacing data and attribute values

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

1. 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.
- The numeric system-missing value is represented by a single period (.). To find system-missing values, enter a single period as the search value and select **Entire cell**.
- 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.

Obtaining Descriptive Statistics for Selected Variables

To obtain descriptive statistics for selected variables:

1. Right-click the selected variables in either Data View or Variable View.
 - From the pop-up menu select **Descriptive Statistics**.

value4. statistics are letminted by variable(meat)-333levbel and dat
•String

Case selection statistics 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 1. Filtered cases in the Data Editor

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:

1. 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

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

Chapter 6. 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

By default, each data source that you open is displayed in a new Data Editor window. (See “General options” on page 193 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.
- 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, IBM 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 DATA`), you need to use the `FILE` command to name each dataset explicitly in order to have more than one data source open

Renaming Datasets

When you open a data source through the menus and dialog boxes, each data source is automatically

Chapter 7. Data preparation

Once you've opened a data file or entered data in the Data Editor, you can start creating reports, charts,

To Define Variable Properties

1. From the menus choose:

Data > Define Variable Properties...

- . 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.
- .

Note

Custom Variable Attributes

The **Attributes** button in Define Variable Properties opens the Custom Variable Attributes dialog box. In

These conditions apply primarily to reading data or creating new variables via command syntax. Dialogs for reading data and creating new transformed variables automatically perform a data pass that sets the measurement level, based on the default measurement level rules.

To set the measurement level for variables with an unknown measurement level

- 1.

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.

1. 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 pop-up 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.

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).

Cop Data Properties

Cop ing Data Properties

The Copy Data Properties Wizard provides the ability to use an external IBM 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 IBM 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 Cop Data Properties

1. From the menus in the Data Editor window choose:
 - Data > Copy Data Properties...**
 - . 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 IBM 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.

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 Cop

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.

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. See the topic “Custom Variable Attributes” on page 53 for more information.

- **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.

Role. Some dialogs support the ability to pre-select variables for analysis based on defined roles. See the topic “Roles” on page 52 for more information.

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.

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

1. 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:

.

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:

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

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. See the topic “Variable measurement level” on page 49 for more information.

To Bin Variables

1. 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. See the topic “Variable names” on page 48 for more information.
 - . Define the binning criteria for the new variable. See the topic “Binning Variables” on page 76 for more information.
 - . Click **OK**.

Binning Variables

Equal Percentiles Based on Scanned Cases. Generates binned categories with an equal number of cases in each bin (using the **empirical 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.

Cop ing Binned Categories

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



- 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.

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

For more information and a detailed description of each function, type   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:

the result is missing if a case has a missing value for any of the three variables.

In the expression:

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:

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.

Some procedures, such as Linear Models, have internal random number generators.

To select the random number generator and/or set the initialization value:

1. 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.

To Count Occurrences of Values within Cases

1. 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.

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.

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

1. 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.

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.

To Recode Values of a Variable into a New Variable

1. 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.

Save template as.

- . 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.

Table 7. Ranking methods and results (continued)

Value	Mean	Low	High	Sequential
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

1. From the menus choose:

Transform > Date and Time Wizard...

- . Select the task you want to accomplish and follow the steps to define the task.
- **Learn how dates and times are represented.** This choice leads to a screen that provides a brief overview of date/time variables in IBM 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 IBM SPSS Statistics

Variables that represent dates and times in IBM 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

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 IBM SPSS Statistics is 1–13.

Specify Date/Time Variable Created by Merging Variables

1. Enter a name for the Result Variable. This cannot be the name of an existing variable.

1. 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:

1. 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.

Table 8. Date difference for years

Date 1	Date 2	Truncate	Round	Fraction
10/21/2006	10/28/2007	1	1	1.02
10/28/2006	10/21/2007	0	1	.98
2/1/2007	3/1/2007	0	0	.08
2/1/2008	3/1/2008	0	0	.08
3/1/2007	4/1/2007	0	0	.08
4/1/2007	5/1/2007	0	0	.08

Table 9. Date difference for months

Date 1	Date 2	Truncate	Round	Fraction
10/21/2006	10/28/2007	12	12	12.22
10/28/2006	10/21/2007	11	12	11.76
2/1/2007	3/1/2007	1	1	.92
2/1/2008	3/1/2008	1	1	.95
3/1/2007	4/1/2007	1	1	1.02
4/1/2007	5/1/2007	1	1	.99

Specify Result of Subtracting Two Date-Format Variables:

1. 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.

S btract D t D t ,variable.

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. For hours, minutes, and seconds the maximum is the displayed value minus one.

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

1.

Teated:

- . 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.

To Replace Missing Values for Time Series Variables

1. 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

Chapter 9. File handling and file transformations

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 IBM 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 active dataset 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).

To Sort Cases

1. From the menus choose:

Data > Sort Cases...

. Select one or more sorting variables.

Optionally, you can do the following:

Iedyuhc333D[(T)91.7(o)-3335

To save the sorted file with encryption:

- . Select **Save file with sorted data** and click **File**.
- . Select **Encrypt file with password** in the Save Sorted Data As dialog box.
- . Click **Save**.
- . In the Encrypt File dialog box, provide a password and re-enter it in the Confirm password text box. Passwords are limited to 10 characters and are case-sensitive.

Warning: Passwords cannot be recovered if they are lost. If the password is lost the file cannot be opened.

Creating strong passwords

- Use eight or more characters.
- Include numbers, symbols and even punctuation in your password.
- Avoid sequences of numbers or characters, such as "123456" and "1, 2, 3", and avoid repetition, such as "11111111".
- Do not create passwords that use personal information such as birthdays or nicknames.
- Periodically change the password in your or nicknames.vvv

- 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.

1. 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 IBM 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

1. 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.)

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 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 IBM 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.

- If you are not matching cases based on the values of key variables, file order determines how cases are matched.

- If one or more key variables are used to match cases and you indicate that the files are already sorted, the two datasets must be sorted by ascending order of the key variable(s).
- Variable names in the second dataset,

- Use **Aggregate** to merge multiple files that don't contain key variables or multiple files already sorted on key variable values.
- Use **Aggregate** to merge multiple files where there is one case data file and multiple table lookup files. Files do not need to be sorted in order of key variable values, and each table lookup file can use a different key variable.

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.

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

1. 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, fraction or count of values above or below a specified value
- Percentage, fraction, or count of values inside or outside of a specified range

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. See the topic “Variable names” on page 48 for more information.

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**.

Compare groups. Split-file groups are presented together for comparison purposes. For pivot tables, a

Organize output by groups. All results from each procedure are displayed separately for each split-file group.

To Split a Data File for Analysis

1. 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

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

1. 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.

- 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.

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).

Note: If unselected cases are filtered (rather than deleted), subsequently sorting the dataset will turn off filtering applied by this dialog.

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.

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.

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

1. 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

1. 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.

- **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:

Table 10. Data with variables and conditions in separate columns

var	factor
8	1
9	1
3	2
1	2

In this example, the first two rows are a **case group** b.p6[(case)- Data w332.7(w332.7(w332.7(w332.7(w3fp9TSw332.7(w19

- **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*.

Table 12. Test scores recorded in separate columns for each factor

score_a	score_b
1014	864
684	636
810	638

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

in the Restructure Data Wizard, restructure one variable group into a new variable and create an index named *group*. The new data file is shown in the following figure.

Table 13. New, restructured data for variables to cases

group	score
SCORE_A	1014
SCORE_B	864
SCORE_A	684
SCORE_B	636
SCORE_A	810
SCORE_B	638

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.

Table 14. Current data for cases to variables

id	score	time
1	1014	bef
1	864	aft
2	684	bef
2	636	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.

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.

See the topic “Example of One Index for Variables to Cases” on page 112 for more information.

- **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.

Restrict re 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.

See the topic “Example of Two Indices for Variables to Cases” on page 113 for more information.

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.

Restrict re 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.

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

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.

Chapter 10. Working with text

Working with text

To delete output in the Viewer

1. 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:

1. 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

1. 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

1. 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

1. 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

1. From the menus choose:
View > Outline Size
 - . Select the outline size (**Small**, **Medium**, or **Large**).

To change the font in the outline

1. 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.

1. 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

1. 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

1. To find or replace information in the Viewer, from the menus choose:

Edit > Find

or

Edit > Replace

You can use Find and Replace to:

- Search the entire document or just the selected items.
- Search down or up from the current location.
-

Cop ing o tp t into other applications

Output objects can be copied and pasted into other applications, such as a word-processing program or a

To Export Output

1. 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.

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.** 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.** 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. Output can be exported as *Excel 97-2004* or *Excel 2007 and higher*.
- **HTML.** Pivot tables are exported as HTML tables. Text output is exported as preformatted HTML. Charts, tree diagrams, and model views are embedded in the document in the selected graphic format. A browser compatible with HTML 5 is required for viewing output exported in HTML format.
- **Web Report.** A web report is an interactive document that is compatible with most browsers. Many of the interactive features of pivot tables available in the Viewer are also available in web reports. You can also export a web report as an IBM Cognos Active Report.
- **Portable Document Format.** All output is exported as it appears in Print Preview, with all formatting attributes intact.
- **PowerPoint file.** 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.** 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.

Open the containing folder. Opens the folder that contains the files that are created by the export.

Output Management System. You can also automatically export all output or user-specified types of output as Word, Excel, PDF, HTML, text or IBM SPSS Statistics-format data files. See the topic Chapter 22, "Output Management System," on page 239 for more information.

HTML options

HTML export requires a browser that is compatible with HTML 5. 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. See the topic "Table properties: printing" on page 140 for more information.

Export layered tables as interactive. Layered tables are displayed as they appear in the Viewer, and you can interactively change the displayed layer in the browser. If this option is not selected, each table layer is displayed as a separate table.

Tables as HTML. This controls style information included for exported pivot tables.

-

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. See the topic “Table properties: printing” on page 140 for more information.

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.

Preserve break points. If you have defined break points, these settings will be preserved in the Word tables.

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. See the topic “Model properties” on page 148 for more information. (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

1. Select **Word/RTF** as the export format.
 - . Click **Change 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.

When exporting to Excel 97-2004, if you modify an existing worksheet, charts, model views, and tree diagrams are not included in the exported output.

Location diagramssa33(speci(forwar)17]TJ9142.91330TD[(Contr)17.9(ols)-333(the)ID[(Location)-333-333(in)-333(the)T*[(wor

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. See the topic “Model properties” on page 148 for more information. (Note: all model views, including tables, are exported as graphics.)

To set Excel export options

1. Select **Excel** as the export format.
 - . Click **Change Options**.

PowerPoint options

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. See the topic “Table properties: printing” on page 140 for more information.

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. See the topic “Model properties” on page 148 for more information. (Note: all model views, including tables, are exported as graphics.)

To set PDF export options

1. Select **Portable Document Format**

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. See the topic “Model properties” on page 148 for more information. (Note: all model views, including tables, are exported as graphics.)

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:

1. 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 method that reduces the size of the image file.
 - **Compress image to reduce file size.** A lossless compression method that reduces the size of the image file.

.

- **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

1. 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

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
- **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 `Page Title` and `Page 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

1. 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.)

To change printed chart size, page numbering, and space between printed items

1. 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 in several formats:

- **Viewer Files (*.spv).** The format that is used to display files in the Viewer window.
- **SPSS Web Report (*.htm).** A web report is an interactive document that is compatible with most browsers. Many of the interactive features of pivot tables available in the Viewer are also available in web reports. This format requires a browser that is compatible with HTML 5.
- **Cognos Active Report (*.mht).** This format requires a browser that supports MHT format files or the Cognos Active Report application.

To control options for saving web reports or save results in other formats (for example, text, Word, Excel), use **Export** on the **File** menu.

To save a Viewer document

1. From the Viewer window menus choose:

File > Save

- . Enter the name of the document, and then click **Save**.

Optionally, you can do the following:

Lock files to prevent editing in IBM SPSS Smartreader. 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 IBM SPSS Smartreader (a separate product for working with Viewer documents). This setting has no effect on Viewer documents opened in IBM SPSS Statistics.

Encrypt files with a password. You can protect confidential information stored in a Viewer document by encrypting the document with a password. Once encrypted, the document can only be opened by providing the password. IBM SPSS Smartreader users will also be required to provide the password in order to open the file.

To encrypt a Viewer document:

- . Select **Encrypt file with password** in the Save Output As dialog box.
- . Click **Save**.
- . In the Encrypt File dialog box, provide a password and re-enter it in the Confirm password text box. Passwords are limited to 10 characters and are case-sensitive.

Warning: Passwords cannot be recovered if they are lost. If the password is lost the file cannot be opened.

Creating strong passwords

- Use eight or more characters.
- Include numbers, symbols and even punctuation in your password.
- Avoid sequences of numbers or characters, such as

Chapter 11. Pivot tables

Moving rows and columns within a dimension element

1. 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.
For legacy tables, perform the addition step:
 - . From the pop-up menu choose **Insert Before** or **Swap**.

Transposing rows and columns

If you just want to flip the rows and columns, there's a simple alternative to using the pivoting trays:

1. 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

1. 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.

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

1. 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.

1. From the menus choose:
Format > Rotate Inner Column Labels
or
Format > Rotate Outer Row Labels

Only the innermost column labels and the outermost row labels can be rotated.

Sorting rows

To sort the rows of a pivot table:

1. Activate the table.
 - . Select any cell in the column you want to use to sort on. To sort just a selected group of rows, select two or more contiguous cells in the column you want to use to sort on.
 - . From the menus choose:

Edit > Sort rows

. Select **Ascending** or **Descending** from the submenu.

•

Changing the output language

To change the output language in a pivot table:

1. Activate the table
 - . From the menus choose:
View > Language
 - . Select one of the available languages.

Changing the language affects only text generated by the application, such as table titles, row and column labels, and footnote text. Variable names and descriptive variable and value labels are not affected.

Note: This feature is not available for legacy tables. See the topic “Legacy tables” on page 145 for more information.

To control the default language for pivot tables and other output objects, use Edit>Options>General. See the topic “General options” on page 193 for more information.

Navigating large tables

To use the navigation window to navigate large tables:

1. Activate the table.
 - . From the menus choose:
View > Navigation

Undoing changes

You can undo the most recent change or all changes to an activated pivot table. Both actions only apply to changes made since the most recent activation of the table.

To undo the most recent change:

1. From the menus choose:
Edit > Undo
To undo all changes:
 - . From the menus choose:
Edit > Restore

Note: Edit > Restore is not available for legacy tables.

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.

Changing the displayed layer

1. Choose a category from the drop-down list of layers (in the pivot table itself, not the pivoting tray).

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.

Note: TableLooks created in earlier versions of IBM SPSS Statistics cannot be used in version 16.0 or later.

To apply a TableLook

1. Activate a pivot table.
 - . From the menus choose:
Format > TableLooks...
 - . 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

1. 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.
- Only table properties defined in the Table Properties dialog are saved in TableLooks. TableLooks do not include individual cell modifications.

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

1. 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 default number of rows to display in long tables. To display all the rows in a table, regardless of how long it is, deselect (clear) **Display table by rows**. *Note:* This feature only applies to legacy tables.
- 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).

To change general table properties:

1. Click the **General** tab.
 - . Select the options that you want.
 - . Click **OK** or **Apply**.

Set rows to display

Note: This feature only applies to legacy tables.

By default, tables with many rows are displayed in sections of 100 rows. To control the number of rows displayed in a table:

1. Select **Display table by rows**.
 - . Click **Set Rows to Display**.
 - or
 - . From the View menu of an activated pivot table, choose **Display table by rows** and **Set Rows to Display**.

Rows to display. Controls the maximum number of rows to display at one time. Navigation controls allow you move to different sections of the table. The minimum value is 10. The default is 100.

Widow/orphan tolerance. Controls the maximum number of rows of the inner most row dimension of the table to split across displayed views of the table. For example, if there are six categories in each group of the inner most row dimension, specifying a value of six would prevent any group from splitting across displayed views. This setting can cause the total number of rows in a displayed view to exceed the specified maximum number of rows to display.

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.

To change cell formats:

1. 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**.

Alternating row colors

To apply a different background and/or text color to alternate rows in the Data area of the table:

1. Select **Data** from the Area drop-down list.
 - . Select (check) **Alternate row color** in the Background Color group.
 - . Select the colors to use for the alternate row background and text.

Alternate row colors affect only the Data area of the table. They do not affect row or column label areas.

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.

To change table borders:

1. 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.

To control pivot table printing properties:

Marker. By default, standard footnote markers are sequential letters or numbers, depending on the table properties settings. To assign a special marker, simply enter the new marker value in the Marker column. Special markers are not affected when you renumber footnotes or switch between numbers and letters for standard markers. The display of numbers or letters for standard markers and the subscript or superscript position of footnote markers are controlled by the Footnotes tab of the Table Properties dialog. See the topic “Table properties: notes” on page 139 for more information.

To change a special marker back to a standard marker, right-click on the marker in the Edit Footnotes dialog, select **Footnote Marker** from the pop-up menu, and select Standard marker in the Footnote Marker dialog box.

Footnote. The content of the footnote. The display reflects the current font and background settings. The font settings can be changed for individual footnotes using the Format subdialog. See the topic “Footnote font and color settings” for more information. A single background color is applied to all footnotes and can be changed in the Font and Background tab of the Cell Properties dialog. See the topic “Font and background” on page 141 for more information.

Visible. All footnotes are visible by default. Deselect (clear) the Visible checkbox to hide a footnote.

Footnote font and color settings

For legacy tables, you can use the Format dialog to change the font family, style, size and color for one or more selected footnotes:

1. In the Edit Footnotes dialog, select (click) one or more footnotes in the Footnotes grid.
 - . Click the **Format** button.

The selected font family, style, size, and colors are applied to all the selected footnotes.

Background color, alignment, and margins can be set in the Cell Properties dialog and apply to all footnotes. You cannot change these settings for individual footnotes. See the topic “Font and background” on page 141 for more information.

Data cell idths

Set Data Cell Width is used to set all data cells to the same width.

To set the width for all data cells:

1. From the menus choose:
 - Format > Set Data Cell Widths...**
 - . Enter a value for the cell width.

Changing col mn idth

1. Click and drag the column border.

Displa ing hidden borders in a pi ,ot table

For tables without many visible borders, you can display the hidden borders. This can simplify tasks like changing column widths.

1. From the View menu choose **Gridlines**.

Selecting rows, columns and cells in a pivot table

You can select an entire row or column or a specified set of data and label cells.

To select multiple cells:

Select > Data and Label Cells

Notes

- The visual highlight that indicates a selected row or column may span noncontiguous areas of the table.
- For legacy tables, you can Ctrl+Alt+click the row or column label to select the row or column.

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. See the topic “Table properties: printing” on page 140 for more information.
- 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. See the topic “Table properties: printing” on page 140 for more information.
- 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

To specify rows or columns to keep together

1. 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 > Breakpoints > Keep Together

Note: For legacy tables, choose Format > Keep Together.

To view breakpoints and keep together groups

You can toggle whether breakpoints and keep together groups are displayed.

1. From the menus choose:

Format > Breakpoints > Display Breakpoints

Breakpoints are shown as vertical or horizontal lines. Keep together groups appear as greyed out rectangular regions enclosed by a darker border.

Note

Chapter 12. 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:

1. Double-click the model.
or
- . Right-click the model and from the pop-up menu choose **Edit Content**.
- . From the submenu choose **In Separate Window**.

Activating the model displays the model in the Model Viewer. See the topic “Working with the Model Viewer” for more information.

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”.)

There are two different styles of Model Viewer:

- **Split into main/auxiliary views.** In this style 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.
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.
- **One view at a time, with thumbnails.** In this style there is only one view visible, and other views are accessed via thumbnails on the left of the Model Viewer. Each view displays some visualization for the model.

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.

Setting model properties

Within the Model Viewer, you can set specific properties for the model. See the topic “Model properties” for more information.

Copying model views

You can also copy individual model views within the Model Viewer. See the topic “Copying model views” for more information.

Model properties

Depending on your Model Viewer, choose:

File > Properties

or

File > Print View

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 (except for auxiliary views based on selection in the main view; these are not printed). Note that you can also print individual model views within the Model Viewer itself. See the topic “Printing a model” for more information.

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

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 page 148. 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 page 123. Note that all model views, including tables, are exported as graphics. Also note that auxiliary views based on selections in the main view are never exported.

Saving fields used in the model to a new dataset

You can save fields used in the model to a new dataset.

1. Activate the model in the Model Viewer. See the topic “Interacting with a model” on page 147 for more information.
 - . From the menus choose:

Generate > Field Selection (Model input and target)

Dataset name. Specify a valid dataset name. 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. See the topic “Variable names” on page 48 for more information.

Saving predictors to a new dataset based on importance

You can save predictors to a new dataset based on the information in the predictor importance chart.

1. Activate the model in the Model Viewer. See the topic “Interacting with a model” on page 147 for more information.
 - . From the menus choose:

Generate > Field Selection (Predictor Importance) but are not saved as files unless explicitly saved prior to the end of the session. Dataset names must conform to variable naming rules. See the topic “Interactable names” on page 48 for more information.

Combining Rule. When scoring an ensemble, this is the rule used to combine the predicted values from the base models to compute the ensemble score value.

- Ensemble predicted values for **categorical** targets can be combined using voting, highest probability, or highest mean probability. **Voting** selects the category that has the highest probability most often across the base models. **Highest probability** selects the category that achieves the single highest probability across all base models. **Highest mean probability** selects the category with the highest value when the category probabilities are averaged across base models.
- Ensemble predicted values for **continuous** targets can be combined using the mean or median of the predicted values from the base models.

The default is taken from the specifications made during model building. Changing the combining rule recomputes the model accuracy and updates all views of model accuracy. The Predictor Importance chart also updates. This control is disabled if the reference model is selected for scoring.

Show All Combining rules. When selected, results for all available combining rules are shown in the model quality chart. The Component Model Accuracy chart is also updated to show reference lines for each voting method.

Model Summary

The Model Summary view is a snapshot, at-a-glance summary of the ensemble quality and diversity.

Quality. The chart displays the accuracy of the final model, compared to a reference model and a naive model. Accuracy is presented in larger is better format; the "best" model will have the highest accuracy. For a categorical target, accuracy is simply the percentage of records for which the predicted value matches the observed value. For a continuous target, accuracy is 1 minus the ratio of the mean absolute error in prediction (the average of the absolute values of the predicted values minus the observed values) to the range of predicted values (the maximum predicted value minus the minimum predicted value).

For bagging ensembles, the reference model is a standard model built on the whole training partition. For boosted ensembles, the reference model is the first component model.

The naive model represents the accuracy if no model were built, and assigns all records to the modal category. The naive model is not computed for continuous targets.

Diversity. The chart displays the "diversity of opinion" among the component models used to build the ensemble, presented in larger is more diverse format. It is a measure of how much predictions vary across the base models. Diversity is not available for boosted ensemble models, nor is it shown for continuous targets.

Predictor Importance

Typically, you will want to focus your modeling efforts on the predictor fields that matter most and consider dropping or ignoring those that matter least. The predictor importance chart helps you do this by indicating the relative importance of each predictor in estimating the model. Since the values are relative, the sum of the values for all predictors on the display is 1.0. Predictor importance does not relate to model accuracy. It just relates to the importance of each predictor in making a prediction, not whether or not the prediction is accurate.

Predictor importance is not available for all ensemble models. The predictor set may vary across component models, but importance can be computed for predictors used in at least one component model.

Predictor Frequency

The predictor set can vary across component models due to the choice of modeling method or predictor selection. The Predictor Frequency plot is a dot plot that shows the distribution of predictors across component models in the ensemble. Each dot represents one or more component models containing the predictor. Predictors are plotted on the y-axis, and are sorted in descending order of frequency; thus the

topmost predictor is the one that is used in the greatest number of component models and the bottommost one is the one that was used in the fewest. The top 10 predictors are shown.

Predictors that appear most frequently are typically the most important. This plot is not useful for methods in which the predictor set cannot vary across component models.

Component Model Accuracy

The chart is a dot plot of predictive accuracy for component models. Each dot represents one or more component models with the level of accuracy plotted on the y-axis. Hover over any dot to obtain information on the corresponding individual component model.

Reference lines. The plot displays color coded lines for the ensemble as well as the reference model and naïve models. A checkmark appears next to the line corresponding to the model that will be used for scoring.

Interactivity. The chart updates if you change the combining rule.

Boosted ensembles. A line chart is displayed for boosted ensembles.

Component Model Details

The table displays information on component models, listed by row. By default, component models are sorted in ascending model number order. You can sort the rows in ascending or descending order by the values of any column.

Model. A number representing the sequential order in which the component model was created.

Accuracy. Overall accuracy formatted as a percentage.

Method. The modeling method.

Predictors. The number of predictors used in the component model.

Model Size. Model size depends on the modeling method: for trees, it is the number of nodes in the tree; for linear models, it is the number of coefficients; for neural networks, it is the number of synapses.

Records. The weighted number of input records in the training sample.

Automatic Data Preparation

This view shows information about which fields were excluded and how transformed fields were derived in the automatic data preparation (ADP) step. For each field that was transformed or excluded, the table lists the field name, its role in the analysis, and the action taken by the ADP step. Fields are sorted by ascending alphabetical order of field names.

The action **Trim outliers**, if shown, indicates that values of continuous predictors that lie beyond a cutoff value (3 standard deviations from the mean) have been set to the cutoff value.

Split Model Viewer

The Split Model Viewer lists the models for each split, and provides summaries about the split models.

Split. The column heading shows the field(s) used to create splits, and the cells are the split values. Double-click any split to open a Model Viewer for the model built for that split.

Accuracy. Overall accuracy formatted as a percentage.

Model Size. Model size depends on the modeling method: for trees, it is the number of nodes in the tree; for linear models, it is the number of coefficients; for neural networks, it is the number of synapses.

Records. The weighted number of input records in the training sample.

Chapter 13. Automated Output Modification

Automated output modification applies formatting and other changes to the contents of the active Viewer window. Changes that can be applied include:

- All or selected viewer objects
- Selected types of output objects (for example, charts, logs, pivot tables)
- Pivot table content based on conditional expressions
- Outline (navigation) pane content

The types of changes you can make include:

- Delete objects
- Index objects (add a sequential numbering scheme)
- Change the visible property of objects
- Change the outline label text
- Transpose rows and columns in pivot tables
- Change the selected layer of pivot tables
- Change the formatting of selected areas or specific cells in a pivot table based on conditional expressions (for example, make all significance values less than 0.05 bold)

To specify automated output modification:

1. From the menus, choose: **Utilities > Style Output**
 - . Select one or more objects in the Viewer.
 - . Select the options you want in the Select dialog. (You can also select objects before you open the dialog.)
 - . Select the output changes you want in the Style Output dialog.

Style Output: Select

The Style Output: Select dialog specifies basic selection criteria for changes that you specify on the Style Output dialog.

You can also select objects in the Viewer after opening the Style Output: Select dialog.

Selected only. Changes are applied only to the selected objects that meet the specified criteria.

- **Select as last command.** If selected, changes are only applied to the output from the last procedure. If not selected, changes are applied to the specific instance of the procedure. For example, if there are three instances of the Frequencies procedure and you select the second instance, changes are applied only to that instance. If you paste the syntax based on your selections, this option selects the second instance of that procedure. If output from multiple procedures is selected, this option applies only to the procedure for which the output is the last block of output in the Viewer.
- **Select as a group.** If selected, all of the objects in the selection are treated as a single group on the main Style Output dialog. If not selected, the selected objects are treated as individual selections, and you can set the properties for each object individually.

All objects of this type. Changes are applied to all objects of the selected type that meet the specified criteria. This option is only available if there is a single object type selected in the Viewer. Object types include tables, warnings, logs, charts, tree diagrams, text, models, and outline headers.

All objects of this sub-type. Changes are applied to all tables of the same subtype as the selected tables that meet the specified criteria. This option is only available if there is a single table subtype selected in the Viewer. For example, the selection can include two separate Frequencies tables, but not a Frequencies table and a Descriptives table.

Objects with a similar name. Changes are applied to all objects with a similar name that meet the specified criteria.

- **Criteria.** The options are Contains, Exactly, Starts with, and Ends with.
- **Value.** The name as it is displayed in the outline pane of the Viewer.
- **Update.** Selects all the objects in the Viewer that meet the specified criteria for the specified value.

St le O tp t

The Style Output dialog specifies the changes you want to make to the selected output objects in the Viewer.

Create a Backup of the Output. Changes made by the automated output modification process cannot be undone. To preserve the original Viewer document, create a backup copy.

Selections and Properties

The list of objects or groups of objects that you can modify is determined by the objects you select in the Viewer and the selections you make in the Style Output: Select dialog.

Selection. The name of the selected procedure or group of object types. When there is an integer in parentheses after the selection text, changes are applied only to that instance of that procedure in the sequence of objects in the Viewer. For example, "Frequencies(2)" applies change only to the second instance of the Frequencies procedure in the Viewer output.

Type. The type of object. For example, log, title, table, chart. For individual table types, the table subtype is also displayed.

Delete. Specifies whether the selection should be deleted.

Visible. Specifies whether the selection should be visible or hidden. The default option is "As Is," which means that current visibility property of the selection is preserved.

Properties. A summary of changes to apply to the selection.

Add. Adds a row to the list and opens the Style Output: Select dialog. You can select other objects in the Viewer and specify the selection conditions.

Duplicate. Duplicates the selected row.

Move Up and **Move Down.** Moves the selected row up or down in the list. The order can be important since changes specified in subsequent rows can overwrite changes that are specified in previous rows.

Create a report of the property changes. Displays a table that summarizes the changes in the Viewer.

Object Properties

You specify the changes that you want to make to each selection from the Selections and Properties section in the Object Properties section. The available properties are determined by the selected row in the Selections and Properties section.

Command. The name of the procedure if the selection refers to a single procedure. The selection can include multiple instances of the same procedure.

Type. The type of object.

Subtype. If the selection refers to a single table type, the table subtype name is displayed.

Outline Label. The label in the outline pane that is associated with the selection. You can replace the label text or add information to the label. For more information, see the topic [Style Output: Labels and Text](#).

Indexing Format. Adds a sequential number, letter, or roman numeral to the objects in the selection. For more information, see the topic [Style Output: Indexing](#).

Table Title. The title of table or tables. You can replace the title or add information to the title. For more information, see the topic [Style Output: Labels and Text](#).

TableLook. The TableLook used for tables. For more information, see the topic [Style Output: TableLooks](#).

Transpose. Transposes rows and columns in tables.

Top Layer. For tables with layers, the category that is displayed for each layer.

Special Variables in Comment Text

You can include special variables to insert date, time, and other values in the **Comment Text** field.

)DATE

Current date in the form dd-mmm-yyyy.

)ADATE

Current date in the form mm/dd/yyyy.

)SDATE

Current date in the form yyyy/mm/dd.

)EDATE

Current date in the form dd.mm.yyyy.

)TIME Current 12-hour clock time in the form hh:mm:ss.

)ETIME

Current 24-hour clock time in the form hh:mm:ss.

)INDEX

The defined index value. For more information, see the topic [Style Output: Indexing](#).

)TITLE

The text of the outline label for the table.

)PROCEDURE

The name of the procedure that created the table.

)DATASET

The name of the dataset used to create the table.

\n Inserts a line break.

Style Output: Labels and Text

The **Style Output: Labels and Text** dialog replaces or adds text to outline labels, text objects, and table titles. It also specifies the inclusion and placement of index values for outline labels, text objects, and table titles.

Add text to the label or text object.. You can add the text before or after the existing text or replace the existing text.

Add indexing. Adds a sequential letter, number, or roman numeral. You can place the index before or after the text. You can also specify one or more characters that are used as a separator between the text and the index. For information about index formatting, see the topic [Style Output: Indexing](#).

Style Output: Indexing

The **Style Output: Indexing** dialog specifies the index format and the starting value.

Type: The sequential index values can be numbers, lower or uppercase letters, lower or uppercase roman numerals.

Starting value. The starting value can be any value that is valid for the type that is selected.

To display index values in the output, you must select **Add indexing** in the **Style Output: Labels and Text** dialog for the selected object type.

- For outline labels, select **Outline label** in the Properties column in the **Style Output** dialog.
- For table titles, select **Table Title** in the Properties column in the **Style Output** dialog.
- For text objects, select **Contents** in the Properties column in the **Style Output** dialog.

Style Options: 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

Dimension. Specifies whether to search rows, columns, or both for a label with the specified value.

Condition. Specifies the condition to find. For more information, see the topic Table Style: Condition.

Formatting. Specifies the formatting to apply to the table cells or areas that meet the condition. For more information, see the topic Table Style: Format.

Add. Adds a row to the list.

Duplicate. Duplicates the selected row.

Move Up and **Move Down.** Moves the selected row up or down in the list. The order can be important since changes specified in subsequent rows can overwrite changes that are specified in previous rows.

Create a report of the conditional styling. Displays a table that summarizes the changes in the Viewer. This option is available when the Table Style dialog is accessed from a statistical procedure dialog. The Style Output dialog has a separate option to create a report.

Table Style: Condition

The Table Style: Condition dialog specifies the conditions under which the changes are applied. There are two options.

- **To all values of this type.** The only condition is the value specified in the Value column in the Table Style dialog. This option is the default.
- **Based on the following conditions.** Within the table area specified by the Value and Dimension columns in the Table Style dialog, find values that meet the specified conditions.

Values. The list contains comparison expressions, such as **Exactly**, **Less than**, **Greater Than**, and **Between**.

- **Absolute value** is available for comparison expressions that only require one value. For example, you could find correlations with an absolute value greater than 0.5.
- **Top** and **Bottom** are the highest and lowest *n* values in the specified table area. The value for **Number** must be an integer.
- **System-missing.** Finds system-missing values in the specified table area.

Table Style: Format

The Table Style: Format dialog specifies changes to apply, based on the conditions specified in the Table Style: Conditions dialog.

Use TableLook defaults. If no previous format changes have been made, either manually or by automated output modification, this is equivalent to making no format changes. If previous changes have been made, this removes those changes and restores the affected areas of the table to their default format.

Apply new formatting. Applies the specified format changes. Format changes include font style and color, background color, format for numeric values (including dates and times), and number of decimals displayed.

Apply to. Specifies the area of the table to which to apply the changes.

- **Cells only.** Applies changes only to table cells that meet the condition.
- **Entire column.** Applies changes to the entire column that contains a cell that meet the condition. This option includes the column label.
- **Entire row.** Applies changes to the entire row that contains a cell that meets the condition. This option includes the row label.

Replace value. Replaces values with the specified new value. For **Cells only**, this option replaces individual cell values that meet the condition. For **Entire row** and **Entire column**, this option replaces all values in the row or column.

Chapter 14. 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 `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 `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 `batch` or `run` commands in an interactive session), each line of command syntax should not exceed 256 characters.

- 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,

```
U      -  
_      -
```

and

```
U      -      -  
_      -      -
```

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 characters; any additional characters are truncated.

Unless you have existing command files that already use the `INCLUDE` command, you should probably use the `INCLUDE` 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

1. 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. By default, the syntax is pasted after the last command. You can choose to have syntax pasted at the position of the cursor or to overwrite selected syntax. The setting is specified from the Syntax Editor tab in the Options 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.

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 `include` and `require` 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. See the topic “Running Command Syntax” on page 169 for more information.

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 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. See the topic “Color Coding” on page 165 for more information.
- 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 starting line number of the command containing the error.
- 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 first line of the command that generated the error.
- You can show or hide the error pane by choosing View > Show Error Pane from the menus.

Using Multiple Views

You can split the editor pane into two panes arranged with one above the other.

1. From the menus choose:

Window > Split

Actions in the navigation and error panes--such as clicking on an error--act on the pane where the cursor is positioned.

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

```
76  * * * * *
```

Comments. Text within a comment is colored gray by default.

Quotes. Quotes and text within quotes are colored black by default.

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 `/*`, `*/`, `/*`, `*/`, and `/*`, `*/`. 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, `END`) or require a matching `END` command at some point later in the command stream (for example, `END`). In both cases, the command will be colored red, by default, until the required `END` statement is added.

Note: You can navigate to the next or previous syntactical error by choosing Next Error or Previous error from the Validation Errors submenu of the Tools menu.

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 `/*`, `*/`, `/*`, `*/`, and `/*`, `*/` 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-IBM SPSS Statistics command syntax, such as occur within `/*`, `*/`, and `/*`, `*/` 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

1. Click anywhere in the gutter to the left of the command text.
or
. Position the cursor within the 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:

1. Click the icon representing the breakpoint in the gutter to the left of the command text.

or

- . Position the cursor within the 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 page 169 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

1. 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:

1. Position the cursor on the line containing the bookmark.

- . From the menus choose:

Tools > Toggle Bookmark

To clear all bookmarks:

1. 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.

1. 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:

1. From the menus choose:

Tools > Next Bookmark

or

Tools > Previous Bookmark

To automatically indent text

1. Select the text.
 - . From the menus choose:
Tools > Indent Syntax > Auto Indent

When you automatically indent text, any existing indentation is removed and replaced with the automatically generated indents. Note that automatically indenting code within a block may break the code if it depends on specific indentation to function, such as Python code containing loops and conditional blocks.

Syntax formatted with the auto-indent feature may not run in batch mode. For example, auto-indenting an or block will cause the syntax to fail in batch mode because commands in the block will be indented and will not start in column 1 as required for batch mode. You can, however, use the -i switch in batch mode to force the Batch Facility to use interactive syntax rules. See the topic “Syntax Rules” on page 161 for more information.

Running Command Syntax

1. 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.
 - . and 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 IBM SPSS Statistics prior to 16.0. For more information on Unicode mode, see “General options” on page 193.

To save a syntax file in a format compatible with earlier releases:

1. 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. See SET command, LOCALE subcommand for more information.

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 IBM 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,

```
EXECUTE.
  COMPUTE var1 = VAR1 - LAG1(VAR1).
EXECUTE.
```

and

```
COMPUTE var1 = VAR1 - LAG1(VAR1).
EXECUTE.
```

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.

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 IBM SPSS Statistics prior to 16.0. For more information on Unicode mode, see “General options” on page 193.

To save a syntax file in a format compatible with earlier releases:

- Avoid sequences of numbers or characters, such as "123456" and "12345678", and avoid repetition, such as "11111111".
- Do not create passwords that use personal information such as birthdays or nicknames.
- Periodically change the password.

Note: Storing encrypted files to an IBM SPSS Collaboration and Deployment Services Repository is not supported.

Modifying encrypted files

- If you open an encrypted file, modify the file and choose File > Save, the modified file is saved with the same password.
- You can change the password on an encrypted file by opening the file, repeating the steps for encrypting it, and specifying a different password in the Encrypt File dialog box.
- You can save an unencrypted version of an encrypted file by opening the file, choosing File > Save As and selecting **Syntax** from the Save as type drop-down list.

Note: Encrypted syntax files cannot be opened in versions of IBM SPSS Statistics prior to version 22.

Chapter 15. 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.”

How to Start the Chart Builder

1. From the menus choose:

Graphs > Chart Builder

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.

1. 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 chart you want 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.
 - . If you need to change statistics or modify attributes of the axes or legends (such as the scale range), click **Element Properties**.

- 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.

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 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

1. Create a chart in IBM SPSS Statistics, or open a Viewer file with charts.
 - Double-click a chart in the Viewer.

Chart Editor Fundamentals

The Chart Editor provides various methods for manipulating charts.

Menus

- Dou333(in)-3iasilyDou(Chasily)-33.8ai333b(Y)91.7(ou33(in)-3i(you)-333(can)-3PThe)-333-33tiart yoT8(.)JTJ 2.4TD[(How)-333(
- Dou(Chate)-3330f105(Chart)-33 wit35 charts.

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

1. 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

1. 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.

1. 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.

1. 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**

Chapter 16. Scoring data with predictive models

The process of applying a predictive model to a set of data is referred to as **scoring** the data. IBM SPSS Statistics has 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 in a file that contains all of the information necessary to reconstruct the model. You can then use that model file to generate predictive scores in other datasets. *Note:*

some model files created by other applications, such as IBM SPSS Modeler, but some model files created by other applications cannot be read by IBM SPSS Statistics, including any models that have multiple

- **Timestamp.** Numeric fields with a display format that includes both the date and the time in the active dataset match the timestamp data type in the model. This corresponds to the Datetime format (dd-mm-yyyy hh:mm:ss) in the active dataset.

Note: In addition to field name and type, you should make sure that the actual data values in the dataset being scored are recorded in the same fashion as the data values in the dataset used to build the model. For example, if the model was built with an *Income* field that has income divided into four categories, and *IncomeCategory* in the active dataset has income divided into six categories or four different categories, those fields don't really match each other and the resulting scores will not be reliable.

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 contains no value. For numeric fields (variables), this means the system-missing value. For string fields, 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 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.

- **Linear Regression and Discriminant 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.
- **Decision Tree 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.
- **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.

Selecting scoring functions

The scoring functions are the types of "scores" available for the selected model. For example, predicted value of the target, probability of the predicted value, or probability of a selected target value.

Scoring function. The scoring functions available are dependent on the model. One or more of the following will be available in the list:

- **Predicted value.** The predicted value of the target outcome of interest. This is available for all models, except those that do not have a target.

- . Select the transformation XML file.
- . Enter a path and name for the new merged model XML file, or use **Browse** to select the location and name.

Note: You cannot merge model .zip files for models that contain splits (separate model information for each split group) or ensemble models with transformation XML files.

Chapter 17. Utilities

Utilities

This chapter describes the functions found on the Utilities menu and the ability to reorder target variable lists.

- For information on the Scoring Wizard, see Chapter 16, “Scoring data with predictive models,” on page 177.
- For information on merging model and transformation XML files, see “Merging model and transformation XML files” on page 180.

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

Visible.

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.

Defining 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 IBM SPSS Statistics data files.

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

1. 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**.
-

Using variable sets to show and hide variables

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.

- 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 IBM 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 tws(or)17.9(der)-333(of)-333

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

1. 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.

Extension bundles

Extension bundles package custom components, such as custom dialogs and extension commands, so that they can be easily installed. For example, IBM SPSS Statistics - Essentials for Python, which is installed by default with IBM SPSS Statistics, includes a set of Python extension commands that are packaged in extension bundles and installed with SPSS Statistics. And IBM SPSS Statistics - Essentials for R (available from the SPSS Community website) includes a set of R extension commands that are packaged in extension bundles. Many more extension bundles, hosted on the SPSS Community website, are available from the Download Extension Bundles dialog, which is accessed from **Utilities > Extension Bundles >you Extensstension**

- . Modify values for any fields on the Required tab.
- . Modify values for any fields on the Optional tab.
- . Specify a target file for the extension bundle.
- . Click **Save** to save the extension bundle to the specified location. This closes the Edit Extension Bundle dialog box.

Required fields for extension bundles

Name. A unique name to associate with the extension bundle. It can consist of up to three words and is not case sensitive. Characters are restricted to seven-bit ASCII. To minimize the possibility of name conflicts, you may want to use a multi-word name, where the first word is an identifier for your organization, such as a URL.

Files. Click **Add** to add the files associated with the extension bundle. An extension bundle must at least include a custom dialog specification (*.spd*) file, or an XML specification file for an extension command. If an XML specification file is included then the bundle must include at least one Python, R or Java code file--specifically, a file of type .py, pyc, .pyo, .R, .class, or .jar.

- Translation files for extension commands (implemented in Python or R) included in the extension bundle are added from the Translation Catalogues Folder field on the Optional tab. See the topic “Optional fields for extension bundles” for more information.
- You can add a readme file to the extension bundle. Specify the filename as *ReadMe.txt*FD(0(u)-3f3(a)-3IBM34.08IPSS34.08

Required Plug-ins. Check the boxes for any Plug-ins (Python or R) that are required in order to run the custom components associated with the extension bundle. Users will be alerted at install time if they don't have the required Plug-ins. For reference, there is no option for specifying the Java Plug-in because it is always installed with IBM SPSS Statistics.

Required R Packages. Enter the names of any R packages, from the CRAN package repository, that are required for the extension bundle to run.

The `resources` folder that contains the localized files must have a subfolder named `resources` - that contains the localized `resources` file for a particular language. For example, the French `resources` file must be in the `resources` folder.

Language Identifiers

de. German

Required R packages

The extension bundle installer attempts to download and install any R packages that are required by the extension bundle and not found on your machine. If you do not have internet access, you will need to obtain the necessary packages from someone who does. Packages can be downloaded from <http://www.r-project.org/> and then installed from within R. For details, see the *R Installation and Administration* guide, distributed with R. If installation of the packages fails, you will be alerted with the list of required packages. You can also view the list from the Extension Bundle Details dialog box, once the bundle is installed. See the topic “Viewing installed extension bundles” on page 191 for more information.

Note

Viewing installed extension bundles


To view details for the extension bundles installed on your machine:

1. From the menus choose:

Utilities > Extension Bundles > View Installed Extension Bundles...

. Click the highlighted text in the *Summary* column for the extension bundle.

The Extension Bundle Details dialog box displays the information provided by the author of the extension bundle. In addition to required information, such as Summary, Description, and Version, the author may have included URL's to locations of relevance, such as the author's home page. If the extension bundle was downloaded from the Download Extension Bundles dialog, then it includes "terms of use" that can be viewed by clicking **View terms of use**.

Components. The Components group lists the custom dialog, if any, and the names of any extension commands included in the extension bundle. Extension commands included with the bundle can be run from the syntax editor in the same manner as built-in IBM SPSS Statistics commands. Help for an extension command may be available by running  in the syntax editor.

Note: Installing an extension bundle that contains a custom dialog might require a restart of IBM SPSS Statistics in order to see the entry for the dialog in the Components group.

Dependencies. The Dependencies group lists add-ons that are required to run the components included in the extension bundle.

- **Integration Plug-Ins for Python and R.** The components for an extension bundle may require the Integration Plug-in for Python, or the Integration Plug-in for R, or both. The Integration Plug-in for Java is installed with the Core system and does not require separate installation.
- **R packages.** Lists any R packages required by the extension bundle. During installation of the extension bundle, the installer attempts to download and install the necessary packages on your machine. If this process fails, you will be alerted and will need to manually install the packages. See the topic "Required R packages" on page 190 for more information.
- **Python modules.** Lists any Python modules required by the extension bundle. Any such modules may be available from the SPSS Community at <http://www.ibm.com/developerworks/spssdevcentral>. For

By default, the selected extension bundles are downloaded and installed on your computer. You can also choose to download the selected extension bundles to a specified location without installing them. You can install them later by choosing **Utilities > Extension Bundles > Install Local Extension Bundle**.

Results of the request to install, or only download, each selected extension bundle are displayed in the Extension Bundles Download Log in the Viewer.

Important:

- Extension bundles are always installed, or downloaded, to your local computer. If you work in distributed analysis mode, then see the topic “Installing local extension bundles” on page 188 for further details.
- For users of Windows Vista and later versions of Windows, installing an updated version of an existing extension bundle might require running IBM SPSS Statistics with administrator privileges. You can start IBM SPSS Statistics with administrator privileges by right-clicking the icon for IBM SPSS Statistics and choosing **Run as administrator**. In particular, if you receive an error message that states that one or more extension bundles could not be installed, then try running with administrator privileges.

Filter by search ilti (updated)-56Vstalling a1-.33(search)-333(ilh)-333(particul(a1dhnstalling)]TJe333(your)-333f3(by)-333(s

Chapter 18. Options

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

1. 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

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.

Roles

Some dialogs support the ability to pre-select variables for analysis based on defined roles. See the topic “Roles” on page 52 for more information.

- **Use predefined roles.** By default, pre-select variables based on defined roles.
- **Use custom assignments.** By default, do *not* use roles to pre-select variables.

You can also switch between predefined roles and custom assignment within the dialogs that support this functionality. The setting here controls only the initial default behavior in effect for each dataset.

Windows

Look and feel. Controls the basic appearance of windows and dialog boxes. If you notice any display issues after changing the look and feel, try shutting down and restarting the application.

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 `open` commands to control multiple datasets. See the topic Chapter 6, “Working with Multiple Data Sources,” on page 63 for more information.

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).

Apply locale's digit grouping format to numeric values. Applies the current locale's digit grouping format to numeric values in pivot tables and charts as well as in the Data Editor. For example, in a French locale with this setting enabled the value 123456789 will be displayed as 12 345 678 900 .

The grouping format does not apply to trees, Model Viewer items, numeric values with the `format` or `currency` format, or numeric values with a `format` or custom currency1Tf-44.7721dt7F or 

Data Options

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 **¶** command is pasted after each transformation command. See the topic “Multiple Execute Commands” on page 170 for more information.

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.

Assigning Measurement Level. For data read from external file formats, older IBM SPSS Statistics data files (prior to release 8.0) and new fields created in a session, measurement level for numeric fields is determined by a set of rules, including the number of unique values. You can specify the minimum number of data values for a numeric variable used to classify the variable as continuous (scale) or nominal. Variables with fewer than the specified number of unique values are classified as nominal.

There are numerous other conditions that are evaluated at this her condition 80aru3(the)-333(number)

filesar the number (scale) or

Table 22. Rules for determining default measurement level (continued)

Condition	Measurement Level
All values of a variable are missing	Nominal
Variable contains at least one non-integer value	Continuous
Variable contains at least one negative value	Continuous
Variable contains no valid values less than 10,000	Continuous
Variable has <i>N</i> or more valid, unique values*	Continuous
Variable has no valid values less than 10	Continuous
Variable has less than <i>N</i> valid, unique values*	Nominal

* *N* is the user-specified cut-off value. The default is 24.

Rounding and Truncation of Numeric Values. For the `ROUND` 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 `ROUND` 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. See the topic “Changing the default variable view” for more information.

Change Dictionary. Controls the language version of the dictionary used for checking the spelling of items in Variable View. See the topic “Spell checking” on page 55 for more information.

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**.

1. Select (check) the variable attributes you want to display.
 - Use the up and down arrow buttons to change the display order of the attributes.

Language options

Language

Output language. Controls the language that is 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 setting does not affect the user interface language.) Depending on the language, you might also need to use Unicode character encoding for characters to render properly.

Note: Custom scripts that rely on language-specific text strings in the output might not run correctly when you change the output language. For more information, see the topic “Script options” on page 203.

User Interface. This setting controls the language that is used in menus, dialogs, and other user interface features. (Note: This setting does not affect the output language.)

Character Encoding and Locale

This controls the default behavior for determining the encoding for reading and writing data files and syntax files. You can change these settings only when there are no open data sources, and the settings remain in effect for subsequent sessions until explicitly changed.

To create custom currency formats

1. 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 options

Output options control the default setting for a number of output options.

Outline labeling. Controls display of variable names, variable labels, data values, and value labels in the outline pane of the Viewer.

Pivot table labeling. Controls display of variable names, variable labels, data values, and value labels in pivot tables.

Descriptive variable and value labels (V

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

1. 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.

Column Widths

These options control 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. (*Note:* This option is only available if you select the option to render as legacy tables.)
- **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 for all tables.** 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.

Table Comments

You can automatically include comments for each table.

- Comment text is displayed in a tooltip when you hover over a table in the Viewer.
- Screen readers read the comment text when the table has focus.
- The tooltip in the Viewer displays only the first 200 characters of the comment, but screen readers read the entire text.
- When you export output to HTML or a web report, the comment text is used as alt text.

Title. Include the table title in the comment.

Procedure. Include the name of the procedure that created the table.

Date. Includes the date the table was created.

Dataset. Includes the name of the dataset that was used to create the table.

Table Rendering

This option render tables as legacy tables. Legacy tables may render slowly and are only recommended if you require compatibility with IBM SPSS Statistics releases prior to 20. For release 20 and higher, all tables have full support for pivoting and editing.

- Tables, other than legacy tables, created in IBM SPSS Statistics release 20 or later and lightweight tables in output documents that are modified in release 20 or later (but created in release 19) cannot be viewed in releases prior to 19.0.0.2. Such tables are viewable in release 19.0.0.2, where they are rendered as lightweight tables; however, they may not render the same as in release 20 or later.
- Lightweight tables created in IBM SPSS Statistics release 19 automatically have full support for pivoting and editing in release 20 or later.

Default Editing Mode

This option 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.

Copied 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, the number of files that appear in the recently used file list, and the installation of Python 2.7 that is used by the IBM SPSS Statistics - Integration Plug-in for Python.

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.

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 IBM SPSS Statistics is installed. By default, no output items are associated with autoscripts. You must manually associate all autoscripts with the output items, as described below. For information on converting legacy autoscripts, see “Compatibility with Versions Prior to 16.0” on page 255.

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 Core system, and the Python programming language. For all other platforms, scripting is available with the Python programming language.

To use scripting with the Python programming language, you need IBM SPSS Statistics - Essentials for Python, which is installed by default with your IBM SPSS Statistics product.

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

1. 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 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

1. In the Objects and Scripts grid, click the cell in the Script column corresponding to the script to dissociate.
 - . Delete the path to the script and then click any other cell in the Objects and Scripts grid.
 - . Click **Apply** or **OK**.

Script editor options

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. See the topic “Color Coding” on page 165 for more information.

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. See the topic “Auto-Completion” on page 165 for more information.

Indent size

Specifies the number of spaces in an indent. The setting applies to indenting selected lines of syntax as well as to automatic indentation.

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.

Optimize for right to left languages. Check this box for the best user experience when working with right to left languages.

Paste syntax from dialogs. Specifies the position at which syntax is inserted into the designated syntax window when pasting syntax from a dialog. **After last command** inserts pasted syntax after the last command. **At cursor or selection** inserts pasted syntax at the position of the cursor; or, if a block of syntax is selected, then the selection will be replaced by the pasted syntax.

Multiple imputations options

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.

Chapter 19. Customizing Menus and Toolbars

Customizing Menus and Toolbars

Menu Editor

- . 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).
10. 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.

To Set Toolbar Properties

1. 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.

To Change Toolbar Images

1. 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.

Chapter 20. 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 IBM 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 IBM SPSS Statistics commands that are implemented in the Python programming language, R or Java. See the topic “Custom Dialogs for Extension Commands” on page 226 for more information.
- Open a file containing the specification for a custom dialog--perhaps created by another user--and add the dialog to your installation of IBM 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 IBM SPSS Statistics.



Figure 2. Custom Dialog Builder

How to Start the Custom Dialog Builder

1. From the menus choose:

Utilities > Custom Dialogs > Custom Dialog Builder...

Custom Dialog Builders 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:

1. 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 IBM SPSS Statistics menus. See the topic “Dialog Properties” for more information.
 - . Specify the controls, such as source and target variable lists, that make up the dialog and any sub-dialogs. See the topic “Control Types” on page 217 for more information.
 - . Create the syntax template that specifies the command syntax to be generated by the dialog. See the topic “Building the Syntax Template” on page 212 for more information.
 - . Install the dialog to IBM SPSS Statistics and/or save the specification for the dialog to a custom dialog package (.spd) file. See the topic “Managing Custom Dialogs” on page 215 for more information.

You can preview your dialog as you're building it. See the topic “Previewing a Custom Dialog” on page 215 for more information.

Dialog Properties

To view and set Dialog Properties:

1. 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.

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 for Windows and Linux. For Mac, specifications are stored under the */Library/Application Support/IBM/SPSS/Statistics/<version>/CustomDialogs/<Dialog Name>* folder, where *<version>* is the two digit IBM SPSS Statistics version--for example, 22. 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 alternative 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. See the topic “Managing Custom Dialogs” on page 215 for more information.

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. See the topic “Building the Syntax Template” on page 212 for more information.

Required Add-Ons. Specifies one or more add-ons--such as the Integration Plug-in for Python or the Integration Plug-in for R--that are required in order to run the command syntax generated by the dialog. For example, if the dialog generates command syntax for an extension command implemented in R, then check the box for the Integration Plug-in for R. Users will be warned about required add-ons that are missing when they try to install or run the dialog.

Specifying the Menu Location for a Custom Dialog

The Menu Location dialog box allows you to specify the name and location of the menu item for a custom dialog. Menu items for custom dialogs do not appear in the Menu Editor within IBM SPSS Statistics.

1. 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. See the topic “Menu Editor” on page 207 for more information. 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 other 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.

Placing Object 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.

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 IBM 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

1. 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. The Syntax Template dialog box supports the auto-completion and color coding features of the Syntax Editor. See the topic "Using the Syntax Editor" on page 163 for more information.
- Add control identifiers of the form `<control identifier>` at the locations where you want to insert command syntax generated by controls, where `<control 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. The list contains the control identifiers followed by the items available with the syntax auto-completion feature. 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. See the topic "Control Types" on page 217 for more information.

Note: The syntax generated at run-time automatically includes a command terminator (period) as the very last character if one is not present.

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:

```

    .FREQ SOURCE=U TARGET=V
  
```

The syntax template to generate this might look like:


```

    .FREQ SOURCE=U TARGET=U
  
```

- `U` 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 `U` 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 in the following figure.



An example of the generated command syntax would be:

```

  * U
  * U
  * U
  
```

The syntax template to generate this might look like:

```

  * U
  * U
  * U
  
```

- `U` is the value of the Identifier property for the target list control and `U` 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 correct result. The Syntax property of the target list would be set to `U`, as described in the previous example.

Syntax property of item group: `U`

Checked Syntax property of mean check box: `U`

Checked Syntax property of stddev check box: `U`

Checked Syntax property of min check box: `U` `U`

Checked Syntax property of max check box: `U` `U`

At run-time, `!CHECKED` will be replaced by the current value of the Syntax property for the item group control. Specifically, `!CHECKED` 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 `!CHECKED`. 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

`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 `!CHECKED`. This may or may not be desirable. For example, even with no boxes checked you may still want to generate the subcommand. This can be accomplished by referencing the identifiers for the check boxes directly in the syntax template, as in:

```
!CHECKED
  MEAN STDDEV
  !CHECKED
  MEAN STDDEV
```

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 IBM 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:

1. 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.

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:

1. From the menus in the Custom Dialog Builder choose:

File > Install

To install from a custom dialog package file:

- From the menus choose:

Utilities > Custom Dialogs > Install Custom Dialog...

For Windows XP and Linux, and by default, installing a dialog requires write permission to the IBM SPSS Statistics installation directory (for Windows 8, Windows 7, Windows Vista, and Mac, dialogs are installed to a general user-writable location). If you do not have write permissions to the required location 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

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 of the installation directory for Windows and Linux. For Mac, specifications are stored under the */Library/Application Support/IBM/SPSS/Statistics/<version>/CustomDialogs/<Dialog Name>* folder, where *<version>* is the two digit IBM SPSS Statistics version--for example, 22. You can copy this folder to the same relative location in another instance of IBM SPSS Statistics and it will be recognized as an installed dialog the next time that instance is launched.

- If you have specified alternative locations for installed dialogs--using the *SPSS_CDIALOGS_PATH* environment variable--then copy the *<Dialog Name>* folder from the appropriate alternate location.
- If alternative locations for installed dialogs have been defined for the instance of IBM 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. See the topic “Source List” for more information.
- **Target List:** A target for variables transferred from the source list. See the topic “Target List” on page 218 for more information.
- **Check Box:** A single check box. See the topic “Check Box” on page 219 for more information.
- **Combo Box:** A combo box for creating drop-down lists. See the topic “Combo Box and List Box Controls” on page 219 for more information.
- **List Box:** A list box for creating single selection or multiple selection lists. See the topic “Combo Box and List Box Controls” on page 219 for more information.
- **Text control:** A text box that accepts arbitrary text as input. See the topic “Text Control” on page 221 for more information.
- **Number control:** A text box that is restricted to numeric values as input. See the topic “Number Control” on page 221 for more information.
- **Static Text control:** A control for displaying static text. See the topic “Static Text Control” on page 222 for more information.
- **Item Group:** A container for grouping a set of controls, such as a set of check boxes. See the topic “Item Group” on page 222 for more information.
- **Radio Group:** A group of radio buttons. See the topic “Radio Group” on page 223 for more information.
- **Check Box Group:** A container for a set of controls that are enabled or disabled as a group, by a single check box. See the topic “Check Box Group” on page 224 for more information.
- **File Browser:** A control for browsing the file system to open or save a file. See the topic “File Browser” on page 224 for more information.
- **Sub-dialog Button:** A button for launching a sub-dialog. See the topic “Sub-dialog Button” on page 225 for more information.

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 ¶ 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. You can also open the Filter dialog by double-clicking the Source List control on the canvas. See the topic “Filtering Variable Lists” on page 219 for more information.

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 ¶ 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

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 ¶ 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 IBM 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 the 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 ¶ for line breaks.
- The value `***` specifies the run-time value of the control, which is the content of the text box. This is the default.
- If the `gaA2.0260TD(Pastoperty)-333(is)-333(nn4(is)-333(nn4(ox.))-333(anVy)-333(va61Tf0-1r)1746.83320TD[(specifies)-333`

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 the 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 \backslash for line breaks.
- The value \backslash specifies the run-time value of the control, which is the numeric value. This is the default.
- If the Syntax property includes \backslash 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 \backslash 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 \backslash 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 has a value. The default is **False**.

For example, the group consists of a set of check boxes. If Required for execution is set to **True** and all of the boxes are unchecked, then **OK** and **Paste** will be disabled.

Syntax. Specifies the 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 \backslash 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 \backslash 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 the Syntax property includes \backslash 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, each of which can contain a set of nested controls. The Radio 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.

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.

Check Box Group

The Check Box Group control is a container for a set of controls that are enabled or disabled as a group,

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. *Note:* You can also open the File Filter dialog by double-clicking the File Browser control on the canvas.

File System Type. In distributed analysis mode, this specifies whether the open or save dialog browses the file system on which IBM 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

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 S b-dialog

To view and set properties for a sub-dialog:

1. 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

An extension command with an associated custom dialog consists of three pieces: an XML file that specifies the syntax of the command; one or more code files (Python, R or Java) that implement the command; and a custom dialog package file that contains the specifications for the custom dialog. If the extension command and its associated custom dialog are distributed in an extension bundle (.spe) file, then you can simply install the bundle from Utilities>Extension Bundles>Install Extension Bundle. Otherwise, you need to install the custom dialog and the extension command files separately as follows:

Custom Dialog Package File. Install the custom dialog package file from Utilities>Custom Dialogs>Install Custom Dialog.

XML Syntax Specification File and Implementation Code. For Windows and Linux, the XML file specifying the syntax of the extension command and the implementation code (Python, R or Java) should be placed in the *extensions* directory under the IBM SPSS Statistics installation directory. For Mac, the XML and code files should be placed in the */Library/Application Support/IBM/SPSS/Statistics/22/extensions* directory.

- For Windows and Linux, if you do not have write permissions to the IBM 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 IBM SPSS Statistics installation directory. Note that Mac users may also utilize the *SPSS_EXTENSIONS_PATH* environment variable. For multiple locations, separate each with a semicolon on Windows and a colon on Linux and Mac.

To create the *SPSS_EXTENSIONS_PATH* environment variable on Windows, follow the same general steps used to create the *SPSS_CDIALOGS_PATH* variable. See the section on Installing a Custom Dialog in “Managing Custom Dialogs” on page 215.

To view the current locations for custom dialogs, run the following command syntax:

- For an extension command implemented in Python, you can always store the associated Python module(s) to a location on the Python search path, such as the Python *site-packages* directory.

Note: To use a new extension command, restart IBM SPSS Statistics.

Creating Localized Versions of Custom Dialogs

You can create localized versions of custom dialogs for any of the languages supported by IBM 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

1. 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 IBM SPSS Statistics installation directory for Windows and Linux, and under the */Library/Application Support/IBM/SPSS/Statistics/22/CustomDialogs/<Dialog Name>* folder 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. See the topic “Managing Custom Dialogs” on page 215 for more information.

- 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.

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*.

When the dialog is launched, IBM 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

1. Make a copy of the help file associated with the dialog and localize the text for the language you want. 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 IBM SPSS Statistics installation directory for Windows and Linux, and under the */Library/Application Support/IBM/SPSS/Statistics/22/CustomDialogs/<Dialog Name>* folder 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. See the topic "Managing Custom Dialogs" on page 215 for more information.

- . 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, IBM 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

pt_BR. Brazilian Portuguese

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 IBM 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.

Chapter 21. Production jobs

Production jobs provide the ability to run IBM 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.

You can run production jobs in two different ways:

Interactively. The program runs unattended in a separate session on either your local computer or a remote server. Your local computer must remain on (and connected to the remote server, if applicable) until the job is complete.

In the background on a server. The program runs in a separate session on a remote server. Your local computer does not have to remain on and does not have to remain connected to the remote server. You can disconnect and retrieve the results later.

Note: Running a production job on a remote server requires access to a server running IBM SPSS Statistics Server.

Creating and running production jobs

To create and run a production job:

1. From the menus in any window choose:

Utilities > Production Job

- . Click **New** to create a new production job.

or

- . Select a production job to run or modify from the list. Click **Browse** to change the directory location for the files that will appear in the list.

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). See the topic “Converting Production Facility files” on page 237 for more information.

- . Specify one or more command syntax files to include in the job. Click the button with the plus sign (+) icon to select command syntax files.
- . Select the output file name, location, and format.
- . Click **Run** to run the production job interactively or in the background on a server.

Default encoding

By default, IBM SPSS Statistics runs in Unicode mode. You can run production jobs in Unicode mode or the current local encoding. The encoding affects how data and syntax files are read. See the topic “General options” on page 193 for more information.

Syntax files

Production jobs use command syntax files to tell IBM SPSS Statistics what to do. A command syntax file is a simple text file containing command syntax. You can use the syntax editor or any text editor to create the file. You can also generate command syntax by pasting dialog box selections into a syntax window. See the topic Chapter 14, “Working with Command Syntax,” on page 161 for more information.

If you include multiple command syntax files, the files are concatenated together in the order they appear in the list and run as a single 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

- **PowerPoint file.**

Production jobs with OUTPUT commands

Production jobs honor `OUTPUT` commands, such as `OUTPUT`, `OUTPUT`, and `OUTPUT`. `OUTPUT` 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` to create a new output document, it is recommended that you explicitly save it with the `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` command, followed by more procedures but no more `OUTPUT` commands. The `OUTPUT` 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` command.

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.

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. See the topic "Variable names" on page 48 for more information.

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 `DEFAULT` keyword when running the production job with command line switches, unless you also use the `NODEFAULT` switch to specify runtime values. See the topic "Running production jobs from a command line" on page 236 for more information.

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.

Example command syntax file with user prompt symbols

```
SYNTAX FILE WITH USER PROMPT SYMBOLS
  FILE='@datafile'
  VARIABLE='@varlist'
```

Run options

You can run production jobs in two different ways:

Interactively. The program runs unattended in a separate session on either your local computer or a remote server. Your local computer must remain on (and connected to the remote server, if applicable) until the job is complete.

In the background on a server. The program runs in a separate session on a remote server. Your local computer does not have to remain on and does not have to remain connected to the remote server. You can disconnect and retrieve the results later.

Note: Running a production job on a remote server requires access to a server running IBM SPSS Statistics Server.

Statistics server. If you select to run the production job in the background on a remote server, you must specify the server on which it will run. Click **Select Server** to specify the server. This applies only to jobs run in the background on a remote server, not jobs run interactively on a remote server.

Server login

Use the Server login dialog to add and modify remote servers and to select the server to use to run the current production job. Remote servers usually require a user ID and password, and a domain name may also be required. To add a remote server, click **Add** in the **Server Login** dialog. To modify a remote server, click **Modify** in the **Server Login** dialog. To select a remote server to use to run the current production job, click **Select** in the **Server Login** dialog.

Refresh. Updates the job status information.

Get job output. Retrieves the output from the selected production job. The output for each job resides on the server the job was run on; so you must switch to the status for that server to select the job and retrieve the output. This button is disabled if the job status is *Running*.

Cancel job. Cancels the selected production job. This button is only enabled if the job status is *Running*.

Remove job. Removes the selected production job. This removes the job from the list and removes the associated files from the remote server. This button is disabled if the job status is *Running*.

Note: Background job status does not reflect the status of any jobs run interactively on a remote server.

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:

```
stats -s <server> -j <job> -e <executable>
```

Depending on how you invoke the production job, you may need to include directory paths for the *stats* executable file (located in the directory in which the application is installed) and/or the production job 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 IBM 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 IBM SPSS Statistics data file format output. Dataset names must conform to variable-naming rules. See the topic “Variable names” on page 48 for more information.

. 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. See the topic “Excluding output display from the viewer” on page 245 for more information.
- 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 (\$).

Tips for selecting multiple items in a list

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:

1. In the Requests list, click any cell in the row for the request.
 - . Click **End**.

To end all active OMS requests:

1. Click **End All**.

To delete a new request (a request that has been added but is not yet active):

1. 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 objects

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.

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

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

1. In the Output Management System Control Panel, select one or more output types and then select one or more commands.
 - . Click **Labels**.
 - . 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 specificn abard include information about variaifn abar iharacter-333(thebelth)-333b333(Generl)-333, yo

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.

IBM 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. See the topic "Routing output to IBM SPSS Statistics data files" on page 245 for more information.

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 IBM 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 `` element with an `src="imagefile.png"` attribute of the general form `src="imagefile.png"` 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, and BMP.

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 IBM SPSS Statistics data file format, table columns become variables, and rows become cases.

- 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 (clear) **Log OMS activity**.

How to specify OMS logging

To specify OMS logging:

1. Click **Logging** in the Output Management System Control Panel.

Excluding output displayed 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:

1. 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.

Note: This setting has no effect on OMS output saved to external formats or files, including the Viewer SPV and SPW formats. It also has no effect on output saved to SPV format in a batch job executed with the Batch Facility (available with IBM SPSS Statistics Server).

Routing output to IBM SPSS Statistics data files

A data file in IBM 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. See the topic "Command identifiers and table subtypes" on page 241 for more information. *Label_* contains the table title text.
- Rows in the table become cases in the data file.

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).

Var1

- 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.

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.

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*.

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:


```
<table command="Frequencies" subtype="Statistics" data-bbox="115 815 505 845">
```
- 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 `command` attribute value of "frequencies" is *not* the same as a `command` 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:




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 example shows a simple frequency table and the complete output XML representation of that table.

Table 23. 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	



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:

```
• <table> <tr> <td> </td> </tr> </table>
```


- 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 `lang` 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 `label` attribute and one or more additional attribute values. An example is as follows:

```
<table>
  <thead>
    <tr>
      <th>
        <code>label="Variable Name" data-bbox="131 178 560 203" />
      </th>
    </tr>
  </thead>
  <tbody>
    <tr>
      <td>
        <code>label="Variable Name" data-bbox="131 178 560 203" />
      </td>
    </tr>
  </tbody>
</table>
```

- For a numeric variable, there would be a `numdec` attribute instead of a `label` attribute. The `numdec` attribute is present only if the variable or values have defined labels.
- The `numdec` elements that contain cell values for numbers will contain the `numdec` attribute and one or more additional attribute values. An example is as follows:

```
<table>
  <tbody>
    <tr>
      <td>
        <code>numdec="2" data-bbox="111 269 453 286" />
      </td>
    </tr>
  </tbody>
</table>
```

The `numdec` attribute is the actual, unrounded numeric value, and the `numdec` 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 `label="Variable Name" data-bbox="111 328 896 391" />` 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.

To use the oms identifiers dialog box

- From the menus choose:
Utilities > OMS Identifiers...

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 `label` 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

1. 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:

```
[ "label" ]
```

Chapter 23. Scripting Facilit

A toscripts

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. See the topic “Script options” on page 203 for more information.

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 A toscripts

You can create an autoscript by starting with the output object that you want to serve as the trigger--for instance, a frequency table.

1. In the Viewer, select the object that will trigger the autoscript.

. From the menus choose:

Utilities > Create/Edit 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. See the topic “Script options” on page 203 for more information.

. Type the code.

For help with converting custom Sax Basic autoscripts used in pre-16.0 versions, see “Compatibility with Versions Prior to 16.0” on page 255.

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.

1. 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...
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. See the topic “Script options” on page 203 for more information.

Scripting with the Python Programming Language

IBM SPSS Statistics provides two separate interfaces for programming with the Python language on Windows, Linux, Mac OS, and for IBM SPSS Statistics Server. Use of these interfaces requires the IBM SPSS Statistics - Integration Plug-in for Python, which is installed by default with your IBM SPSS Statistics product. 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 `run` 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 IBM 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 IBM SPSS Statistics client is running.

Complete documentation of the IBM SPSS Statistics classes and methods available for Python scripts can be found in the Scripting Guide for IBM SPSS Statistics, available under Integration Plug-in for Python in the Help system.

Python Programs

Python programs make use of the interface exposed by the Python `run` module. They operate on the IBM 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 `run` 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 IBM SPSS Statistics Server), Python programs execute on the computer where IBM SPSS Statistics Server is running.

More information about Python programs, including complete documentation of the IBM SPSS Statistics functions and classes available for them, can be found in the documentation for the Python Integration Package for IBM SPSS Statistics, available under Integration Plug-in for Python in the Help system.

Running Python Scripts and Python programs

Both Python scripts and Python programs can be run from within IBM SPSS Statistics or from an external Python process, such as a Python IDE or the Python interpreter.

Python Scripts

Python Script Run from IBM 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 IBM SPSS Statistics operate on the IBM 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 IBM SPSS Statistics, or the Python interpreter. The script will attempt to connect to an existing IBM 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 IBM 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 `runpython` block in command syntax. The command syntax can be run from the IBM SPSS Statistics client or from the IBM SPSS Statistics Batch Facility for IBM Moe0kpa-333(IBM)-executa(with)-333 and connVnew client. YStatistics client clients Batch

To migrate a pre-16.0 version of a script that called functions in the global procedures file, add the statement `# *** "C:\Program Files\IBM\SPSS\Statistics\GlobalProcedures.sps"` to the declarations section of the script, where `C:\Program Files\IBM\SPSS\Statistics` is the directory where IBM SPSS Statistics is installed. `# ***` 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 `# ***` statement. You can also use `!INCLUDE` instead of `# ***`.

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 IBM 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. See the topic "Script options" on page 203 for more information.

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:

1. 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 `!INCLUDE` and remove the parameter specification, keeping track of which parameters are required by the script, such as a pivot table applied file the way a procedure which script, -333(items,)-333



- 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.
- `GetObject` gets the output item (an `OutputObject` object) that triggered the autoscript.
- The object returned by `GetObject` is not activated. If your script requires an activated object, you'll need to activate it, as done in this example with the `Activate` 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. See the topic "The scriptContext Object" for more information.

Note: To trigger a script from the application creation event, see "Startup Scripts" on page 258.

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 IBM SPSS Statistics Basic Script Editor is a standalone application that is launched from within IBM 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 IBM SPSS Statistics from which it was launched. Once opened, the editor will remain open after exiting IBM SPSS Statistics, but scripts that use IBM 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 IBM 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 IBM SPSS Statistics from an external COM client is `SPSSWINLIB.APPLICATION16`. Application objects should be declared as `spsswinLib.Application16`. For example:

```
Dim spsswinLib As New spsswinLib.Application16
```

To connect to a running instance of the IBM SPSS Statistics client from an external COM client, use:

```
spsswinLib.Connect
```

If more than one client is running, `Connect` will connect to the most recently launched one.

Note: For post-16.0 versions, the identifier is still `SPSSWINLIB.APPLICATION16`.

The scriptContext Object

Detecting When a Script is Run as an Autoscript

Using the `isRunningAsAutoscript` object, 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.

```

function myFunction() {
    if (isRunningAsAutoscript) {
        // Code to run when the script is an autoscript
    } else {
        // Code to run when the script is not an autoscript
    }
}

```

- When a script is not run as an autoscript, the `isRunningAsAutoscript` object will have a value of `false`.
- Given the logic in this example, you would include your autoscript-specific code in the `if` clause. Any code that is not to be run in the context of an autoscript would be included in the `else` clause. Of course you can also include code that is to be run in either context.

Getting Values Required by Autoscripts

The `isRunningAsAutoscript` object provides access to values required by an autoscript, such as the

```

# 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31 32 33 34 35 36 37 38 39 40 41 42 43 44 45 46 47 48 49 50 51 52 53 54 55 56 57 58 59 60 61 62 63 64 65 66 67 68 69 70 71 72 73 74 75 76 77 78 79 80 81 82 83 84 85 86 87 88 89 90 91 92 93 94 95 96 97 98 99 100 101 102 103 104 105 106 107 108 109 110 111 112 113 114 115 116 117 118 119 120 121 122 123 124 125 126 127 128 129 130 131 132 133 134 135 136 137 138 139 140 141 142 143 144 145 146 147 148 149 150 151 152 153 154 155 156 157 158 159 160 161 162 163 164 165 166 167 168 169 170 171 172 173 174 175 176 177 178 179 180 181 182 183 184 185 186 187 188 189 190 191 192 193 194 195 196 197 198 199 200 201 202 203 204 205 206 207 208 209 210 211 212 213 214 215 216 217 218 219 220 221 222 223 224 225 226 227 228 229 230 231 232 233 234 235 236 237 238 239 240 241 242 243 244 245 246 247 248 249 250 251 252 253 254 255 256 257 258 259 260 261 262 263 264 265 266 267 268 269 270 271 272 273 274 275 276 277 278 279 280 281 282 283 284 285 286 287 288 289 290 291 292 293 294 295 296 297 298 299 300 301 302 303 304 305 306 307 308 309 310 311 312 313 314 315 316 317 318 319 320 321 322 323 324 325 326 327 328 329 330 331 332 333 334 335 336 337 338 339 340 341 342 343 344 345 346 347 348 349 350 351 352 353 354 355 356 357 358 359 360 361 362 363 364 365 366 367 368 369 370 371 372 373 374 375 376 377 378 379 380 381 382 383 384 385 386 387 388 389 390 391 392 393 394 395 396 397 398 399 400 401 402 403 404 405 406 407 408 409 410 411 412 413 414 415 416 417 418 419 420 421 422 423 424 425 426 427 428 429 430 431 432 433 434 435 436 437 438 439 440 441 442 443 444 445 446 447 448 449 450 451 452 453 454 455 456 457 458 459 460 461 462 463 464 465 466 467 468 469 470 471 472 473 474 475 476 477 478 479 480 481 482 483 484 485 486 487 488 489 490 491 492 493 494 495 496 497 498 499 500 501 502 503 504 505 506 507 508 509 510 511 512 513 514 515 516 517 518 519 520 521 522 523 524 525 526 527 528 529 530 531 532 533 534 535 536 537 538 539 540 541 542 543 544 545 546 547 548 549 550 551 552 553 554 555 556 557 558 559 560 561 562 563 564 565 566 567 568 569 570 571 572 573 574 575 576 577 578 579 580 581 582 583 584 585 586 587 588 589 590 591 592 593 594 595 596 597 598 599 600 601 602 603 604 605 606 607 608 609 610 611 612 613 614 615 616 617 618 619 620 621 622 623 624 625 626 627 628 629 630 631 632 633 634 635 636 637 638 639 640 641 642 643 644 645 646 647 648 649 650 651 652 653 654 655 656 657 658 659 660 661 662 663 664 665 666 667 668 669 670 671 672 673 674 675 676 677 678 679 680 681 682 683 684 685 686 687 688 689 690 691 692 693 694 695 696 697 698 699 700 701 702 703 704 705 706 707 708 709 710 711 712 713 714 715 716 717 718 719 720 721 722 723 724 725 726 727 728 729 730 731 732 733 734 735 736 737 738 739 740 741 742 743 744 745 746 747 748 749 750 751 752 753 754 755 756 757 758 759 760 761 762 763 764 765 766 767 768 769 770 771 772 773 774 775 776 777 778 779 780 781 782 783 784 785 786 787 788 789 790 791 792 793 794 795 796 797 798 799 800 801 802 803 804 805 806 807 808 809 810 811 812 813 814 815 816 817 818 819 820 821 822 823 824 825 826 827 828 829 830 831 832 833 834 835 836 837 838 839 840 841 842 843 844 845 846 847 848 849 850 851 852 853 854 855 856 857 858 859 860 861 862 863 864 865 866 867 868 869 870 871 872 873 874 875 876 877 878 879 880 881 882 883 884 885 886 887 888 889 890 891 892 893 894 895 896 897 898 899 900 901 902 903 904 905 906 907 908 909 910 911 912 913 914 915 916 917 918 919 920 921 922 923 924 925 926 927 928 929 930 931 932 933 934 935 936 937 938 939 940 941 942 943 944 945 946 947 948 949 950 951 952 953 954 955 956 957 958 959 960 961 962 963 964 965 966 967 968 969 970 971 972 973 974 975 976 977 978 979 980 981 982 983 984 985 986 987 988 989 990 991 992 993 994 995 996 997 998 999 1000

```

The `StartServer_` method is used to run a `Windows` command that calls the `StartServer_` command to perform the mapping. When the `StartServer_` script runs, IBM SPSS Statistics is in distributed mode so the `StartServer_` command runs on the IBM SPSS Statistics Server machine.

Chapter 24. TABLES and IGRAPH Command Syntax Converter

If you have command syntax files that contain `TABLES` syntax that you want to convert to `IGRAPH` syntax and/or `IGRAPH` syntax that you want to convert to `TABLES` 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 `IGRAPH` and between `IGRAPH` and `TABLES`. It is likely that you will find that the utility program cannot convert some of your `TABLES` and `IGRAPH` syntax jobs or may generate `TABLES` and `IGRAPH` 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 " " or " " in the `TABLES` subcommand if the variable is parenthesized by itself—for example, `TABLES (variable) (command)`. These will be interpreted as the `TABLES` and `TABLES` keywords.
- `TABLES(dmti680TD(dmti680TD(dmti680TD(dmt(includelf—for)ized)-384.1(by)-384.1(itself—for)-384)mibTfA09i461ngwdES)680`

The conversion utility program may generate additional syntax that it stores in the keyword within the 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 /I:input /O:output /S:syntax
```

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 /I:"C:\Program Files\IBM\SPSS\Statistics 22\Utility" /O:"C:\Program Files\IBM\SPSS\Statistics 22\Utility" /S:"C:\Program Files\IBM\SPSS\Statistics 22\Utility"
```

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 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 `/P` (or `/S`) when you run *SyntaxConverter.exe*, as in:

```
SYNTAXCONVERTER /I:input /O:output /S:syntax /P
```

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.

1. 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.

Chapter 25. Encrypting data files, output documents, and syntax files

You can protect confidential information stored in a data file, an output document, or a syntax file by encrypting the file with a password. Once encrypted, the file can only be opened by providing the password. The option to encrypt a file is provided on the Save As dialogs for data files, output documents, and syntax files. You can also encrypt a data file when sorting it and saving the sorted file.

- Passwords cannot be recovered if they are lost. If the password is lost the file cannot be opened.
- Passwords are limited to 10 characters and are case-sensitive.

Creating strong passwords

- Use eight or more characters.
- Include numbers, symbols and even punctuation in your password.
- Avoid sequences of numbers or characters, such as "123456" and "1,2,3", and avoid repetition, such as "11111111".
- Do not create passwords that use personal information such as birthdays or nicknames.
- Periodically change the password.

Modifying encrypted files

- If you open an encrypted file, make modifications to it and choose File > Save, the modified file will be saved with the same password.
- You can change the password on an encrypted file by opening the file, repeating the steps for encrypting it, and specifying a different password in the Encrypt File dialog box.
- You can save an unencrypted version of an encrypted data file or output document by opening the file, choosing File > Save As and deselecting **Encrypt file with password** in the associated Save As dialog box. For an encrypted syntax file, select **Syntax** from the Save as type drop-down list to save an unencrypted version of the file.

Note: Encrypted data files and output documents cannot be opened in versions of IBM SPSS Statistics prior to version 21. Encrypted syntax files cannot be opened in versions prior to version 22.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program,

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Software Group
ATTN: Licensing
200 W. Madison St.
Chicago, IL; 60606
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at “Copyright and trademark information” at www.ibm.com/legal/copytrade.shtml.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

- Custom Dialog Builder *(continued)*
 - item group control 222
 - layout rules 212
 - list box 219
 - list box list items 220
 - localizing dialogs and help files 227
 - menu location 211
 - modifying installed dialogs 215
 - number control 221
 - opening dialog specification files 215
 - preview 215
 - radio group 223
 - radio group buttons 223
 - saving dialog specifications 215
 - source list 217
 - static text control 222
 - sub-dialog button 225
 - sub-dialog properties 226
 - syntax template 212
 - target list 218
 - text control 221
- custom tables
 - converting TABLES command syntax to CTABLES 261
- custom variable attributes 53

D

- data analysis 6
 - basic steps 6
- data dictionary
 - applying from another file 71
- Data Editor 47, 48, 52, 56, 57, 58, 59, 61, 207
 - alignment 52
 - changing data type 59
 - column width 52
 - data value restrictions 57
 - Data View 47
 - defining variables 48
 - descriptive statistics 60
 - descriptive statistics options 198
 - display options 61
 - editing data 57, 58
 - entering data 56
 - entering non-numeric data 57
 - entering numeric data 56
 - filtered cases 61
 - inserting new cases 58
 - inserting new variables 58
 - moving variables 59
 - multiple open data files 63, 193
 - multiple views/panes 61
 - printing 61
 - roles 52
 - sending data to other
 - applications 207
 - Variable View 47
- data entry 56
- data files 9, 10, 17, 24, 25, 29, 39, 43, 108
 - adding comments 183
 - dictionary information 24
 - encrypting 29
 - file information 24
 - flipping 101
 - IBM SPSS Data Collection 22

- data files *(continued)*
 - improving performance for large files 39
 - multiple open data files 63, 193
 - opening 9, 10
 - protecting 38
 - Quancept 22
 - Quanvert 22
 - remote servers 43
 - restructuring 108
 - saving 24, 25
 - saving output as IBM SPSS Statistics data files 239
 - saving subsets of variables 29
 - text 17
 - transposing 101
- DATA LIST 38
 - versus GET DATA command 38
- data transformations 195
 - computing variables 81
 - conditional transformations 81
 - delaying execution 195
 - functions 82
 - ranking cases 88
 - recoding values 84, 85, 86, 87
 - string variables 82
 - time series 94, 95
- data types 49, 50, 59, 197
 - changing 59
 - custom currency 49, 197
 - defining 49
 - display formats 50
 - input formats 50
- Data View 47
- databases 12, 13, 14, 15, 16, 17
 - adding new fields to a table 34
 - appending records (cases) to a table 34
 - conditional expressions 15
 - converting strings to numeric variables 16
 - creating a new table 34
 - creating relationships 14
 - defining variables 16

MicrosofcreativEcc100, (pr)17.7(otarameed)-riablesqu00(94.)-332.7(16.)-332.7((pr)17.7(oPecting)-mpformats
Micrcurrdatabasesie00(1ormats)-acces4, data2TD[(data)-332DOLLARor l0(5arTJT*7(50.)-332.7(-1.25ning)D

- entering data (*continued*)
 - numeric 56
 - using value labels 57
- environment variables 202
 - SPSSTMPDIR 202
- EPS files 123, 129
 - exporting charts 123, 129
- Excel files 9, 10, 25, 207
 - adding menu item to send data to
 - Excel 207
 - opening 9, 10
 - saving 24, 25
 - saving value labels instead of values 24
- Excel format
 - exporting output 123, 126, 242
- excluding output from Viewer with OMS 245
- EXECUTE (command)
 - pasted from dialog boxes 170, 171
- export data 24
- exporting
 - models 149
- exporting charts 123, 129, 231
 - automated production 231
- exporting data 207
 - adding menu item07

- missing values (*continued*)
 - charts 176
 - defining 51
 - in functions 82
 - replacing in time series data 97
 - scoring models 178
 - string variables 51
- model viewer
 - split models 151
- Model Viewer 147
- models 147
 - activating 147
 - copying 148
 - exporting 149
 - interacting 147
 - merging model and transformation files 180
 - Model Viewer 147
 - models supported for export and scoring 177
 - printing 148
 - properties 148
 - scoring 177
- moving rows and columns 134
- multiple open data files 63, 193
 - suppressing 64
- multiple response sets
 - defining 69
 - multiple categories 69
 - multiple dichotomies 69
- multiple views/panes
 - Data Editor 61
 - syntax editor 163

N

- new features
 - version 22 1
- nominal 49
 - measurement level 49, 67
- normal scores
 - in Rank Cases 89
- numeric format 49, 50

O

- OK button 4
- OMS 239, 249
 - command identifiers 241
 - controlling table pivots 242, 246
 - Excel format 242
 - excluding output from the Viewer 245
 - IBM SPSS Statistics data file
 - format 242, 245
 - output object types 240
 - PDF format 242
 - SAV file format 242, 245
 - table subtypes 241
 - text format 242
 - using XSLT with OXML 249
 - variable names in SAV files 246
 - Word format 242
 - XML 242, 246
- online Help 7
 - Statistics Coach 6

- opening files 9, 10, 11, 12, 17
 - controlling default file locations 202
 - data files 9, 10
 - dBASE files 9, 11
 - Excel files 9, 10
 - Lotus 1-2-3 files 9
 - spreadsheet files 9, 11
 - Stata files 11
 - SYSTAT files 9
 - tab-delimited files 9
 - text data files 17
- options 193, 194, 195, 197, 198, 200, 202, 203
 - charts 198
 - currency 197
 - data 195
 - descriptive statistics in Data Editor 198
 - general 193
 - language 196
 - output labels 198
 - pivot table look 200
 - scripts 203
 - syntax editor 203
 - temporary directory 202
 - two-digit years 195
 - Variable View 196
 - Viewer 194
- ordinal 49
 - measurement level 49, 67
- outer join 14
- outline 120, 121
 - changing levels 121
 - collapsing 120
 - expanding 120
 - inD[(or)17.8(dinal)7u

- pivoting
 - controlling with OMS for exported output 246
- PNG files 123, 129
 - exporting charts 123, 129
- port numbers 41, 235
- portable files
 - variable names 25
- post-processing output 153, 154, 156, 157, 158
- PostScript files (encapsulated) 123, 129
 - exporting charts 123, 129
- PowerPoint 127
 - exporting output as PowerPoint 127
- PowerPoint format
 - exporting output 123
- printing 61, 130, 131, 138, 140, 144
 - chart size 131
 - charts 130
 - controlling table breaks 144
 - data 61
 - headers and footers 130
 - layers 130, 138, 140
 - models 148
 - page numbers 131
 - pivot tables 130
 - print preview 130
 - scaling tables 138, 140
 - space between output items 131
 - text output 130
- prior moving average function 96
- Production Facility 193
 - converting files to production jobs 237
 - using command syntax from journal file 193
- Production jobs 231, 235, 236
 - command line switches 236
 - converting Production Facility files 237
 - exporting charts 231
 - output files 231
 - running multiple production jobs 236
 - scheduling production jobs 236
 - syntax rules 231
- programming with command language 161
- properties 138
 - pivot tables 138
 - tables 138
- proportion estimates
 - in Rank Cases 89
- Python
 - scripts 253

Q

- Quancept 22
- Quanvert 22

R

- random number seed 83
- random sample 15
 - databases 15

- random sample (*continued*)
 - random number seed 83
 - selecting 107
- ranking cases 88
 - fractional ranks 89
 - percentiles 89
 - Savage scores 89
 - tied values 89
- Rankit estimates 89
- recoding values 75, 84, 85, 86, 87
- remote servers 41, 43, 44, 235
 - adding 41, 235
 - available procedures 44
 - data file access 43
 - editing 41, 235
 - logging in 41
 - relative paths 44
- removing group labels 134
- renaming datasets 64
- reordering rows and columns 134
- replacing missing values
 - linear interpolation 97
 - linear trend 97
 - mean of nearby points 97
 - median of nearby points 97
 - series mean 97
- Reset button 4
- restricted numeric format 49
- restructuring data 108, 110, 111, 112, 113, 114, 115, 116
 - and weighted data 116
 - creating a single index variable for variables to cases 113
 - creating index variables for variables to cases 112
 - creating multiple index variables for variables to cases 114
 - example of cases to variables 110
 - example of one index for variables to cases 112
 - example of two indices for variables to cases 113
 - example of variables to cases 110
 - options for cases to variables 115
 - options for variables to cases 114
 - overview 108
 - selecting data for cases to variables 115
 - selecting data for variables to cases 111
 - sorting data for cases to variables 115
 - types of restructuring 108
 - variable groups for variables to cases 111
- roles
 - Data Editor 52
- rotating labels 134
- rows 144
 - selecting in pivot tables 144
- running median function 96

S

- sampling
 - random sample 107

- SAS files
 - opening 9
 - reading 9
 - saving 24
- SAV file format
 - routing output to IBM SPSS Statistics data file 242
 - routing output to IBM SPSS Statistics data files 245
- Savage scores 89
- saving charts 123, 129
 - BMP files 123, 129
 - EMF files 123
 - EPS files 123, 129
 - JPEG files 123, 129
 - metafiles 123
 - PICT files 123
 - PNG files 129
 - PostScript files 129
 - TIFF files 129
- saving files 24, 25
 - controlling default file locations 202
 - data files 24, 25
 - database file queries 17
 - IBM SPSS Statistics data files 24
- saving output 123, 127, 128
 - Excel format 123, 126
 - HTML 123, 124
 - HTML format 123
 - PDF format 123, 127
 - PowerPoint format 123, 127
 - text format 123, 128
 - web report 125
 - Word format 123, 125
- scale 49
 - measurement level 49, 67
- scale variables
 - binning to create categorical variables 75
- scaling
 - pivot tables 138, 140
- scientific notation 49, 193
 - suppressing in output 193
- scoring 177
 - matching dataset fields to model fields 178
 - merging model and transformation XML files 180
 - missing values 178
 - models supported for export and scoring 177
 - scoring functions 179
- scripts 207, 208, 251
 - adding to menus 207
 - autoscripts 252
 - Basic 255
 - creating 251
 - default language 203, 251
 - editing 251
 - languages 251
 - Python 253
 - running 251
 - running with toolbar buttons 208
 - startup scripts 258
- search and replace
 - Viewer documents 122
- seasonal difference function 96

- select cases 106
- selecting cases 106
 - based on selection criteria 107
 - date range 107
 - random sample 107
 - range of cases 107
 - time range 107
- selection methods 144
 - selecting rows and columns in pivot tables 144
- servers 41, 235
 - adding 41, 235
 - editing 41, 235
 - logging in 41
 - names 41, 235
 - port numbers 41, 235
- session journal 202
- Shift Values 84
- showing 119, 137, 207
 - captions 141
 - dimension labels 137
 - footnotes 142
 - results 119
 - rows or columns 137
 - titles 137
 - toolbars 207
- sizes 121
 - in outline 121
- smoothing function 96
- sorting
 - pivot table rows 134
 - variables 100
- sorting cases 99
- sorting variables 100
- space-delimited data 17
- speed 39
 - caching data 39
- spelling 55, 56
 - dictionary 195
- split-file processing 105
- split-model viewer 151
- splitting tables 144
 - controlling table breaks 144
- spp files
 - converting to spj files 237
- spreadsheet files 9, 10, 11, 27
 - opening 11
 - reading ranges 10
 - reading variable names 10
 - writing variable names 27
- SPSSTMPDIR environment variable 202
- Stata files 11
 - opening 9, 11
 - reading 9
 - saving 24
- Statistics Coach 6
- status bar 3
- string format 49
- string variables 51, 57
 - breaking up long strings in earlier releases 25
 - computing new string variables 82
 - entering data 57
 - in dialog boxes 4
 - missing values 51
 - recoding into consecutive integers 87
- style output 154
- subsets of cases
 - random sample 107
 - selecting 106, 107

- variable labels *(continued)*
 - in dialog boxes 4, 193
 - in merged data files 102
 - in outline pane 198
 - in pivot tables 198
 - inserting line breaks 51
- variable lists 185
 - reordering target lists 185
- variable names 48, 193
 - generated by OMS 246
 - in dialog boxes 4, 193
 - mixed case variable names 48
 - portable files 25
 - rules 48
 - truncating long variable names in earlier releases 25
 - wrapping long variable names in output 48
- variable pairs 108
 - creating 108
- variable sets 184
 - defining 184
 - using 184
- Variable View 47
 - customizing 54, 55, 196
- variables 5, 48, 58, 59, 108, 183, 184, 193
 - defining 48
 - defining variable sets 184
 - definition information 183
 - display order in dialog boxes 193
 - finding in Data Editor 59
 - in dialog boxes 4
 - inserting new variables 58
 - moving 59
 - recoding 84, 85, 86, 87
 - renaming for merged data files 102
 - restructuring into cases 108
 - selecting in dialog boxes 5
 - sorting 100
 - variable information in dialog boxes 5
- vertical label text 134
- Viewer 119, 120, 121, 131, 194, 198
 - changing outline font 121
 - changing outline levels 121
 - changing outline sizes 121
 - collapsing outline 120
 - deleting output 120
 - display options 194
 - displaying data values 198
 - displaying value labels 198
 - displaying variable labels 198
 - displaying variable names 198
 - excluding output types with OMS 245
 - expanding outline 120
 - find and replace information 122
 - hiding results 119
 - moving output 119
 - outline 120
 - outline pane 119
 - results pane 119
 - saving document 131
 - search and replace information 122
 - space between output items 131
- virtual active file 38

IBM

ALCIP