



CHAPTER 1

Getting Started Using SAS[®] Software

- 1.1 The SAS Language 2
- 1.2 SAS Data Sets 4
- 1.3 The Two Parts of a SAS Program 6
- 1.4 The DATA Step's Built-in Loop 8
- 1.5 Choosing a Mode for Submitting SAS Programs 10
- 1.6 Windows and Commands in the SAS Windowing Environment 12
- 1.7 Submitting a Program in the SAS Windowing Environment 14
- 1.8 Reading the SAS Log 16
- 1.9 Viewing Your Results in the Output Window 18
- 1.10 Creating HTML Output 20
- 1.11 Using SAS System Options 22

1.1 The SAS Language

Many software applications are either menu driven, or command driven (enter a command—see the result). SAS is neither. With SAS, you use statements to write a series of instructions called a SAS program. The program communicates what you want to do and is written using the SAS language. There are some menu-driven front ends to SAS, Enterprise Guide software for example, which make SAS appear like a point-and-click program. However, these front ends still use the SAS language to write programs for you. You will have much more flexibility using SAS if you learn to write your own programs using the SAS language. Maybe learning a new language is the last thing you want to do, but be assured that although there are parallels between SAS and languages you know (be they English or FORTRAN), SAS is much easier to learn.

SAS programs A SAS program is a sequence of statements executed in order. A statement gives information or instructions to SAS and must be appropriately placed in the program. An everyday analogy to a SAS program is a trip to the bank. You enter your bank, stand in line, and when you finally reach the teller's window, you say what you want to do. The statements you give can be written down in the form of a program:

```
I would like to make a withdrawal.  
My account number is 0937.  
I would like $200.  
Give me five 20s and two 50s.
```

Note that you first say what you want to do, then give all the information the teller needs to carry out your request. The order of the subsequent statements may not be important, but you must start with the general statement of what you want to do. You would not, for example, go up to a bank teller and say, "Give me five 20s and two 50s." This is not only bad form, but would probably make the teller's heart skip a beat or two. You must also make sure that all the subsequent statements belong with the first. You would not say, "I want the largest box you have" when making a withdrawal from your checking account. This statement belongs with "I would like to open a safe deposit box." A SAS program is an ordered set of SAS statements like the ordered set of instructions you use when you go to the bank.

SAS statements As with any language, there are a few rules to follow when writing SAS programs. Fortunately for us, the rules for writing SAS programs are much fewer and simpler than those for English.

The most important rule is

Every SAS statement ends with a semicolon.

This sounds simple enough. But while children generally outgrow the habit of forgetting the period at the end of a sentence, SAS programmers never seem to outgrow forgetting the semicolon at the end of a SAS statement. Even the most experienced SAS programmer will at least occasionally forget the semicolon. You will be two steps ahead if you remember this simple rule.

Layout of SAS programs There really aren't any rules about how to format your SAS program. While it is helpful to have a neat looking program with each statement on a line by itself and indentions to show the various parts of the program, it isn't necessary.

- ◆ SAS statements can be in upper- or lowercase.
- ◆ Statements can continue on the next line (as long as you don't split words in two).
- ◆ Statements can be on the same line as other statements.
- ◆ Statements can start in any column.

So you see, SAS is so flexible that it is possible to write programs so disorganized that no one can read them, not even you. (Of course, we don't recommend this.)

Comments To make your programs more understandable, you can insert comments into your programs. It doesn't matter what you put in your comments—SAS doesn't look at it. You could put your favorite cookie recipe in there if you want. However, comments are usually used to annotate the program, making it easier for someone to read your program and understand what you have done and why.

There are two styles of comments you can use: one starts with an asterisk (*) and ends with a semicolon (;). The other style starts with a slash asterisk (/*) and ends with an asterisk slash (*//). The following SAS program shows the use of both of these style comments:

```
* Read animals' weights from file;
DATA animals;
  INFILE 'c:\MyRawData\Zoo.dat';
  INPUT Lions Tigers;
PROC PRINT DATA = animals; /* Print the results */
RUN;
```

Since some operating environments interpret a slash asterisk (/*) in the first column as the end of a job, be careful when using this style of comment not to place it in the first column. For this reason, we chose the asterisk-semicolon style of comment for this book.

Errors People who are just learning a programming language often get frustrated because their programs do not work correctly the first time they write them. To make matters worse, SAS errors often come up in bright red letters, and for the poor person whose results turn out more red than black, this can be a very humbling experience. You should expect errors. Most programs simply don't work the first time, if for no other reason than that you are human. You forget a semicolon, misspell a word, have your fingers in the wrong place on the keyboard. It happens. Often one small mistake can generate a whole list of errors. Don't panic if you see red.

1.2 SAS Data Sets

Before you run an analysis, before you write a report, before you do anything with your data, SAS must be able to read your data. Before SAS can read your data, the data must be in a special form called a SAS data set.¹ Getting your data into a SAS data set is usually quite simple as SAS is very flexible and can read almost any data. Once your data have been read into a SAS data set, SAS keeps track of what is where and in what form. All you have to do is specify the name and location of the data set you want, and SAS figures out what is in it.

Variables and observations Data, of course, are the primary constituent of any data set. In traditional SAS terminology the data consist of variables and observations. Adopting the terminology of relational databases, SAS data sets are also called tables, observations are also called rows, and variables are also called columns. Below you see a rectangular table containing a small data set. Each line represents one observation, while Id, Name, Height, and Weight are variables. The data point Charlie is one of the values of the variable Name and is also part of the second observation.

		Variables (Also Called Columns)			
		Id	Name	Height	Weight
Observations (Also Called Rows)	1	53	Susie	42	41
	2	54	Charlie	46	55
	3	55	Calvin	40	35
	4	56	Lucy	46	52
	5	57	Dennis	44	.
	6	58		43	50

Data types Raw data come in many different forms, but SAS simplifies this. In SAS there are just two data types: numeric and character. Numeric fields are, well, numbers. They can be added and subtracted, can have any number of decimal places, and can be positive or negative. In addition to numerals, numeric fields can contain plus signs (+), minus signs (-), decimal points (.), or E for scientific notation. Character data are everything else. They may contain numerals, letters, or special characters (such as \$ or !) and can be up to 32,767 characters long.

If a variable contains letters or special characters, it must be character data. However, if it contains only numbers, then it may be numeric or character. You should base your decision on how you will use the variable.² Sometimes data that consist solely of numerals make more sense as character data than as numeric. ZIP codes, for example, are made up of numerals, but it just doesn't make sense to add, subtract, multiply, or divide ZIP codes. Such numbers make more sense as character data. In the previous data set, Name is obviously a character variable, and Height and Weight are numeric. Id, however, could be either numeric or character. It's your choice.

¹There are exceptions. If your data are in a format written by another software product, you may be able to read your data directly without creating a SAS data set. For database management systems and spreadsheets, you may be able to use SAS/ACCESS. See sections 2.19 and 2.20 for more information. For SPSS you can use the SPSS data engine. See appendix D.

²If disk space is a problem, you may also choose to base your decision on storage size. You can use the LENGTH statement, discussed in section 9.15, to control the storage size of variables.

Missing data Sometimes despite your best efforts, your data may be incomplete. The value of a particular variable may be missing for some observations. In those cases, missing character data are represented by blanks, and missing numeric data are represented by a single period (.). In the data set above, the value of Weight for observation 5 is missing, and its place is marked by a period. The value of Name for observation 6 is missing and is just left blank.

Size of SAS data sets SAS can handle up to 300,000 variables in a single data set. The number of observations, on the other hand, is limited only by your computer's capacity to handle and store them.

Rules for SAS names You make up names for the variables in your data and for the data sets themselves. It is helpful to make up names that identify what the data represent, especially for variables. While the variable names A, B, and C might seem like perfectly fine, easy-to-type names when you write your program, the names Sex, Height, and Weight will probably be more helpful when you go back to look at the program six months later. Follow these simple rules when making up names for variables and data set members:

- ◆ Names must be 32 characters or fewer in length.³
- ◆ Names must start with a letter or an underscore (_).
- ◆ Names can contain only letters, numerals, or underscores (_). No %\$!*&#@, please.⁴
- ◆ Names can contain upper- and lowercase letters.

This last point is an important one. SAS is insensitive to case so you can use uppercase, lowercase or mixed case—whichever looks best to you. SAS doesn't care. The data set name heightweight is the same as HEIGHTWEIGHT or HeightWeight. Likewise, the variable name BirthDate is the same as BIRTHDATE and birThDaTe. However, there is one difference for variable names. SAS remembers the case of the first occurrence of each variable name and uses that case when printing results. That is why, in this book, we use mixed case for variable names but lowercase for other SAS names.

Documentation stored in SAS data sets In addition to your actual data, SAS data sets contain information about the data set such as its name, the date that you created it, and the version of SAS you used to create it. SAS also stores information about each variable, including its name, type (numeric or character), length (or storage size), and position within the data set. This information is sometimes called the descriptor portion of the data set, and it makes SAS data sets self-documenting.

³Format names can also be 32 characters long, but informat names must be 31 characters or fewer. Librefs and filerefs must be 8 characters or fewer in length, and member names for versioned data sets must be 28 characters or fewer.

⁴It is possible to use special characters, including spaces, in variable names if you use the system option VALIDVARNAMES=ANY and a name literal of the form 'variable-name'N. See the online documentation for details.

1.3 The Two Parts of a SAS Program



SAS programs are constructed from two basic building blocks: DATA steps and PROC steps. A typical program starts with a DATA step to create a SAS data set and then passes the data to a PROC step for processing. Here is a simple program that converts miles to kilometers in a DATA step and prints the results with a PROC step:

```

DATA step [ DATA distance;
           Miles = 26.22;
           Kilometer = 1.61 * Miles;

PROC step [ PROC PRINT DATA = distance;
           RUN;

```

DATA and PROC steps are made up of statements. A step may have as few as one or as many as hundreds of statements. Most statements work in only one type of step—in DATA steps but not PROC steps, or vice versa. A common mistake made by beginners is to try to use a statement in the wrong kind of step. You're not likely to make this mistake if you remember that DATA steps read and modify data while PROC steps analyze data, perform utility functions, or print reports.

DATA steps start with the DATA statement, which starts, not surprisingly, with the word DATA. This keyword is followed by a name that you make up for a SAS data set. The DATA step above produces a SAS data set named DISTANCE. In addition to reading data from external, raw data files, DATA steps can include DO loops, IF-THEN/ELSE logic, and a large assortment of numeric and character functions. DATA steps can also combine data sets in just about any way you want, including concatenation and match-merge.

Procedures, on the other hand, start with a PROC statement in which the keyword PROC is followed by the name of the procedure (PRINT, SORT, or PLOT, for example). Most SAS procedures have only a handful of possible statements. Like following a recipe, you use basically the same statements or ingredients each time. SAS procedures do everything from simple sorting and printing to analysis of variance and 3D graphics. Other SAS procedures perform utility functions such as importing data files and data entry.

A step ends when SAS encounters a new step (marked by a DATA or PROC statement), a RUN statement, or, if you are running in batch mode, the end of the program.¹ RUN statements tell SAS to run all the preceding lines of the step and are among those rare, global statements that are not part of a DATA or PROC step. In the program above, SAS knows that the DATA step has ended when it reaches the PROC statement. The PROC step ends with a RUN statement, which coincides with the end of the program.

¹ If you use SAS long enough, you may run into an exception. Steps can also terminate with a QUIT, STOP, or ABORT statement.

While a typical program starts with a DATA step to input or modify data and then passes the data to a PROC step, that is certainly not the only pattern for mixing DATA and PROC steps. Just as you can stack building blocks in any order, you can arrange DATA and PROC steps in any order. A program could even contain only DATA steps or only PROC steps.

To review, the table below outlines the basic differences between DATA and PROC steps:

DATA steps	PROC steps
▶ begin with DATA statements	▶ begin with PROC statements
▶ read and modify data	▶ perform specific analysis or function
▶ create a SAS data set	▶ produce results or report

As you read this table, keep in mind that it is a simplification. Because SAS is so flexible, the differences between DATA and PROC steps are, in reality, more blurry. The table above is not meant to imply that PROC steps never create SAS data sets (many do), or that DATA steps never produce reports (they can). Nonetheless, you will find it much easier to write SAS programs if you understand the basic functions of DATA and PROC steps.

1.4 The DATA Step's Built-in Loop

DATA steps read and modify data, and they do it in a way that is flexible, giving you lots of control over what happens to your data. However, DATA steps also have an underlying structure, an implicit, built-in loop. You don't tell SAS to execute this loop: SAS does it automatically. Memorize this:

DATA steps execute line by line and observation by observation.


This basic concept is rarely stated explicitly. Consequently, new users often grow into old users before they figure this out on their own.

The idea that DATA steps execute line by line is fairly straightforward and easy to understand. It means that, by default, SAS executes line one of your DATA step before it executes line two, and line two before line three, and so on. That seems common sense, and yet new users frequently run into problems because they try to use a variable before they create it. If a variable named Z is the product of X and Y, then you better make sure that the statements creating X and Y come before the statements creating Z.

What is not so obvious is that while DATA steps execute line by line, they also execute observation by observation. That means SAS takes the first observation and runs it all the way through the DATA step (line by line, of course) before looping back to pick up the second observation. In this way, SAS sees only one observation at a time.

Imagine a SAS program running in slow motion: SAS reads observation number one from your input data set. Then SAS executes your DATA step using that observation. If SAS reaches the end of the DATA step without encountering any serious errors, then SAS writes the current observation to a new, output data set and returns to the beginning of the DATA step to process the next observation. After the last observation has been written to the output data set, SAS terminates the DATA step and moves on to the next step, if there is one. End of slow motion; please return to normal megahertz.

This diagram illustrates how an observation flows through a DATA step:



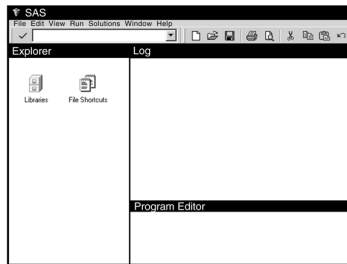
SAS reads observation number one and processes it using line one of the DATA step, then line two, and so on until SAS reaches the end of the DATA step. Then SAS writes the observation in the output data set. This diagram shows the first execution of the line-by-line loop. Once SAS finishes with the first observation, it loops back to the top of the DATA step and picks up observation two. When SAS reaches the last observation, it automatically stops.¹

Here is an analogy. DATA step processing is a bit like voting. When you arrive at your polling place, you stand in line behind other people who have come to vote. When you reach the front of the line you are asked standard questions: “What is your name? Where do you live?” Then you sign your name, and you cast your vote. In this analogy, the people are observations, and the voting process is the DATA step. People vote one at a time (or observation by observation). Each voter’s choices are secret, and peeking at your neighbor’s ballot is definitely frowned upon. In addition, each person completes each step of the process in the same order (line by line). You cannot cast your vote before you give your name and address. Everything must be done in the proper order.

¹ If this seems a bit too structured, don’t worry. You can override the line-by-line and observation-by-observation structure in a number of ways. For example, you can use the RETAIN statement, discussed in section 3.9, to make data from the previous observation available to the current observation. You can also use the OUTPUT statement, discussed in sections 6.11 and 6.12, to control when observations are written to the output data set.

1.5 Choosing a Mode for Submitting SAS Programs

So far we have talked about writing SAS programs, but simply writing a program does not give you any results. Just like writing a letter to your representative in Congress does no good unless you mail it, a SAS program does nothing until you submit or execute it. You can execute a SAS program several ways, but not all methods are available for all operating environments. Check in the online documentation for your operating environment or with your SAS Support Representative to find out which methods are available to you. The method you choose for executing a SAS program will depend on your preferences and on what is most appropriate for your application and your environment. If you are using SAS at a large site with many users, then ask around and find out which is the most accepted method of executing SAS. If you are using SAS on your own personal computer, then choose the method that suits you.



SAS windowing environment If you type SAS at your system prompt, or click on the SAS icon, you will most likely get into the SAS windowing environment. In this interactive environment, you can write and edit SAS programs, submit programs for processing, and view and print your results. In addition, there are many SAS windows for performing different tasks such as managing SAS files, customizing the interface, accessing SAS Help, and importing or exporting data. Exactly what your windowing

environment looks like depends on the type of computer or terminal you are using, the operating environment on the computer, and what options are in effect when you start up SAS. If you are using a personal computer, then the SAS windowing environment will look similar to other programs on your computer, and many of the features will be familiar to you. Submitting programs from the SAS windowing environment is covered in section 1.7.



Enterprise Guide If you have Enterprise Guide software, you may choose to submit your programs from within Enterprise Guide. To do this, you need to use the Insert menu to open a Code window where you can either enter your SAS program or open an existing SAS program. Then you can choose to run (submit) your code on the local machine if SAS is installed, or on a remote server. To run your SAS program on a remote server, you must have SAS Integration Technologies installed. Also, Enterprise Guide

can write SAS code for you through its extensive menu system. Enterprise Guide runs under Windows only.



Noninteractive mode Noninteractive mode is where your SAS program statements are in a file on your system, and you start up SAS specifying that you want to execute that file. SAS immediately starts to process your file and ties up your computer, or window, until it is finished. The results are usually placed in a file or files, and you are returned to your system prompt.

Noninteractive mode is useful in many situations. This mode is good if you want your program to execute immediately, but you do not

want to or cannot use a windowing environment. Noninteractive mode is usually started by typing SAS at your system prompt (shown here as \$), followed by the filename containing your program statements:


```
$ SAS MyFile.sas
```

Batch or background mode With batch or background mode, your SAS program is in a file. You submit the file for processing with SAS. Your SAS program may start executing immediately, or it could be put in a queue behind other jobs. Batch processing is used a lot on mainframe computers, which are capable of executing many processes at one time. You can continue to work on your computer while your job is being processed, or better yet, you can go to the baseball game and let the computer work in your absence. Batch processing is usually less expensive than other methods and is especially good for large jobs which can be set up to execute at off hours when the rates are at their lowest. When your job is complete, the results will be placed in a file or files, which you can display or print at any time.



Batch processing may not be available for your operating environment. Check in the online documentation for your operating environment to see if it is available, then check with your SAS Support Representative to find out how to submit SAS programs for batch processing. Even sites with the same operating environment may have different ways of submitting jobs in batch mode.

Remote submit If you have SAS/CONNECT software, it is possible to write and develop your SAS programs on one system, then submit them for processing on another. Using this method, you write your program on your local machine, then you establish a connection to the remote machine, run the program on the remote machine, and the results are delivered back to your local machine. You might want to do this if your remote machine is much more powerful than your local machine, and you are running very large programs. Also, you might need to access large or shared data files on the remote machine. Check with your SAS Support Representative to find out if this is an option at your site.



Interactive line mode This mode is mentioned only because you might see it in the SAS documentation, and you might get into it by accident. In interactive line mode, you are prompted for SAS statements one line at a time. There is no easy way to correct mistakes once you have entered them, so unless you are an excellent typist, and an excellent programmer, interactive line mode is exceedingly frustrating.



If you do find yourself in this mode (you will know when you get a 1? as a prompt), you can get out by typing ENDSAS; and pressing ENTER. For example

```
1? ENDSAS;
```

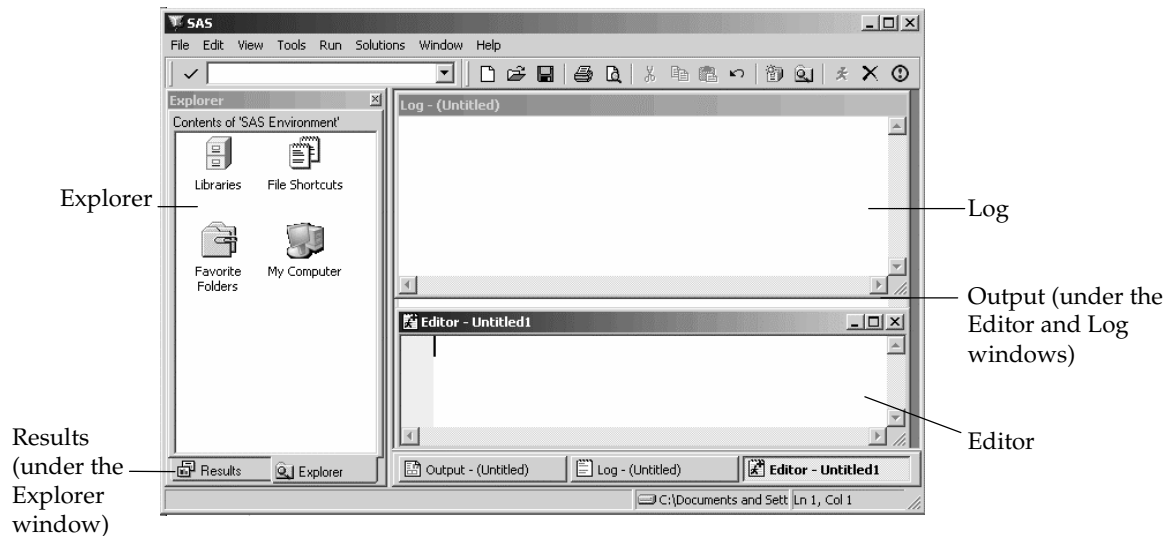
Seek assistance from your SAS Support Representative to find out why you got into line mode and how to avoid it in the future.

1.6 Windows and Commands in the SAS Windowing Environment

It used to be that SAS looked pretty much the same on all platforms, and you couldn't change its appearance. But now SAS adopts the look and feel of your operating environment, and there are many ways in which you can customize your SAS environment. This is good for you because many aspects of the SAS windowing environment will be familiar, and if you don't like the default view, you can change it. It makes writing about it more difficult, because we can't tell you exactly what your SAS session will look like and how it will behave. However, there are many common elements between the various operating environments, and you will probably already be familiar with those elements which are different.

The SAS Windows

There are five basic SAS windows: the Results and Explorer windows, and three programming windows: Editor, Log, and Output. It is possible to bring up SAS without all these windows, and sometimes the windows are not immediately visible (for example, in the Windows operating environment, the Output window comes up behind the Editor and Log windows), but all these windows do exist in your SAS session. There are also many other SAS windows that you may use for tasks such as getting help, changing SAS system options, and customizing your SAS session. The following figure shows the default view for a Microsoft Windows SAS session, with pointers to the five main SAS windows.



Editor This window is a text editor. You can use it to type in, edit, and submit SAS programs as well as edit other text files such as raw data files. In Windows operating environments, the default editor is the Enhanced Editor. The Enhanced Editor is syntax sensitive and color codes your programs making it easier to read and find mistakes. The Enhanced Editor also has the capability of collapsing and expanding the various steps in your program. For other operating environments, the default editor is the Program Editor. The Program Editor will color code your programs, but does not have a collapse and expand feature.

Log The Log window contains notes about your SAS session, and after you submit a SAS program, any notes, errors, or warnings associated with your program as well as the program statements themselves will appear in the Log window.


Output If your program generates any printable results, then they will appear in the Output window.

Results The Results window is like a table of contents for your Output window; the results tree lists each part of your results in an outline form.

Explorer The Explorer window gives you easy access to your SAS files and libraries.

The SAS Commands

There are SAS commands for performing a variety of tasks. Some tasks are probably familiar, such as opening and saving files, cutting and pasting text, and accessing Help. Other commands are specific to the SAS System, such as submitting a SAS program, or starting up a SAS application. You may have up to three ways to issue commands: menus, the tool bar, or the SAS command bar (or command line). The following figure shows the location of these three methods of issuing SAS commands in the Windows operating environment default view.



Menus Most operating environments will have pull-down menus located either at the top of each window, or at the top of your screen. If your menus are at the top of your screen, then the menus will change when you activate the different windows (usually by clicking on them). You may also have, for each window, context-sensitive pop-up menus that appear when you press the right or center button of your mouse.

Tool bar The tool bar, if you have one, gives you quick access to commands that are already accessible through the pull-down menus. Not all operating environments have a tool bar.

SAS command bar The command bar is a place that you can type in SAS commands. In some operating environments the command bar is located with the tool bar (as shown here); in other operating environments you may have a command line with each of the SAS windows (usually indicated by `Command=>`). Most of the commands that you can type in the command bar are also accessible through the pull-down menus or the tool bar.


Controlling your windows The Window pull-down menu gives you choices on how the windows are placed on your screen. You can also activate any of the programming windows by selecting it from the Window pull-down menu, typing the name of the window in the command line area of your SAS session, or simply clicking on the window.

1.7 Submitting a Program in the SAS Windowing Environment

Naturally after going to the trouble of writing SAS programs, you want to see some results. As we have already discussed, there are several ways of submitting your SAS programs. If you use the SAS windowing environment, then you can edit and submit your program, and see your results all within the windowing environment.

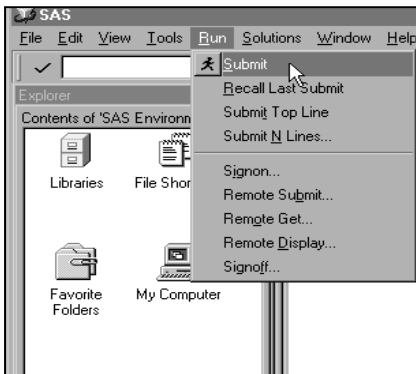
Getting your program into the editor The first thing you need to do is get your program into the Editor window. You can either type your program into the editor, or you can bring the program into the Editor window from a file. The commands for editing in the editor and for opening files should be familiar. SAS tries to follow conventions that are common for your operating environment. For example, to open a file in the editor, you can select Open from the File pull-down menu. For some operating environments you may have an Open icon on your tool bar, and you may also have the option of pasting your file into the editor from the clipboard.

Submitting your program Once your program appears in the editor, you execute it using the SUBMIT command. Depending on your operating environment, you have a few choices on how to execute the SUBMIT command.

 Use the Submit icon on the tool bar.

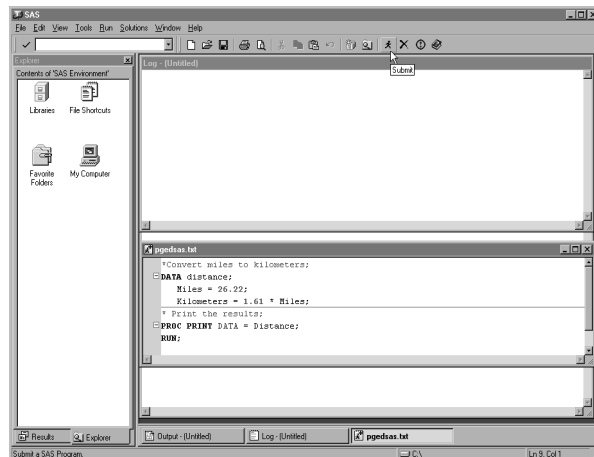


Make the Editor window active and enter SUBMIT in the command line area of your SAS session.



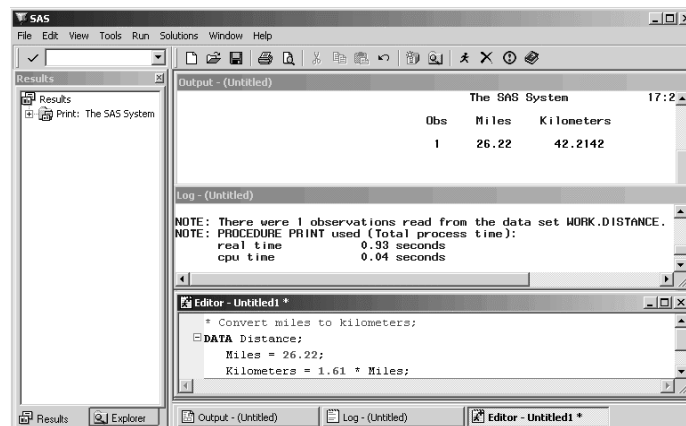
Make the Editor window active and select Submit from the Run pull-down menu.

This figure shows a program in the Enhanced Editor in the Windows operating environment ready to be submitted using the Submit icon on the tool bar.



Viewing the SAS Log and Output If you are using the Enhanced Editor (Windows operating environment), after you submit your program, the program remains in the Enhanced Editor window and the results of your program go into the Log and Output windows. If you are using the Program Editor (all other operating environments) then your results also go into the Log and Output windows, but your program disappears from the Program Editor window. At first it may be a shock for you to see your program disappear in front of your eyes. Don't worry; the program you spent so long writing is not gone forever. If your program produced any output, then you will also get new entries in your Results window (your Results window may not be visible if you do not have the SAS Explorer window active). The Results window is like a table of contents for your SAS output and is discussed in more detail in section 1.9. This figure is an example of what your screen might look like after you submit a program from the Enhanced Editor.

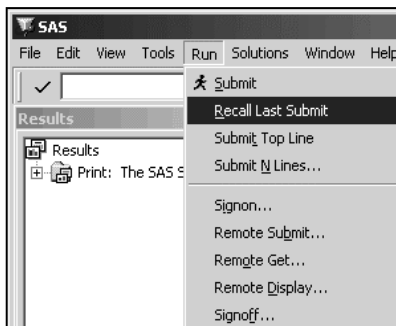
You may not see all three of the programming windows (Editor, Log, and Output) at the same time. In some operating environments, the windows are placed one on top of the other. You can bring a window to the top by clicking on it, typing its name in the command line area, or selecting it from the Window menu.



Getting your program back Unfortunately for most of us, our programs do not run perfectly every time. If you have an error in your program, you will most likely want to edit the program and run it again. If you are using the Program Editor window, you will need to get your program back in the Program Editor window using the RECALL command. You have two choices for executing the RECALL command.



Make the Program Editor the active window, then enter RECALL in the command line area of your SAS session.



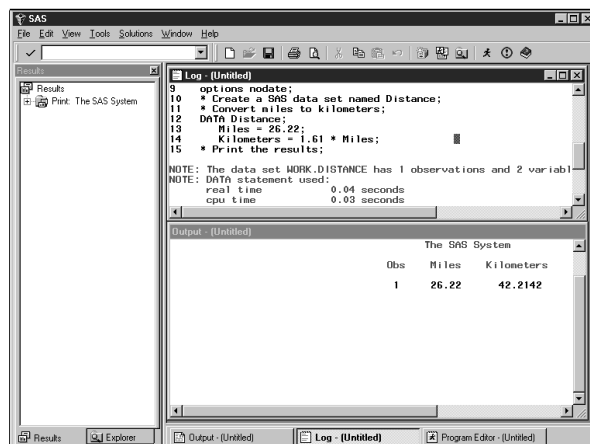
Make the Program Editor the active window, then select Recall Last Submit from the Run pull-down menu.

The RECALL command will bring back the last block of statements you submitted. If you use the RECALL command again, it will insert the block of statements submitted before the last one, and so on and so on, until it retrieves all the statements you submitted.

1.8 Reading the SAS Log

Every time you run a SAS job, SAS writes messages in your log. Many SAS programmers ignore the SAS log and go straight to the output. That's understandable, but dangerous. It is possible—and sooner or later it happens to all of us—to get bogus results that look fine in the output. The only way to know they are bad is to check the SAS log. Just because it runs doesn't mean it's right.

Where to find the SAS log The location of the SAS log varies depending on the operating environment you use, the mode you use (SAS windowing environment, noninteractive, or batch), and local settings. If you submit a program in the windowing environment, you will, by default, see the SAS log in your Log window as in the following figure.



If you submit your program in batch or noninteractive mode, the log will be written to a file that you can view or print using your operating environment's commands for viewing and printing. The name given to the log file is generally some permutation of the name you gave the original program. For example, if you named your SAS program `Marathon.sas`, then it is a good bet that your log file will be `Marathon.log`. At some installations the log and output files are written to a single file, so don't be surprised if you find them together.

What the log contains People tend to think of the SAS log as either a rehash of their program or as just a lot of gibberish. OK, we admit, there is some technical trivia in the SAS log, but there is also plenty of important information. Here is a simple program that converts miles to kilometers and prints the result:

```

* Create a SAS data set named distance;
* Convert miles to kilometers;
DATA distance;
  Miles = 26.22;
  Kilometers = 1.61 * Miles;
* Print the results;
PROC PRINT DATA = distance;
RUN;

```


If you run this program, SAS will produce a log similar to this:

```

① NOTE: Copyright (c) 2002 by SAS Institute Inc., Cary, NC, USA.
NOTE: SAS (r) Proprietary Software Pre-Production Version 9 (TS B0)
      Licensed to XYZ Inc., Site 0098541001.
NOTE: This session is executing on the WIN_PRO platform.

NOTE: SAS initialization used:
      real time          2.51 seconds
      cpu time           0.81 seconds

② 1   * Create a SAS data set named distance;
   2   * Convert miles to kilometers;
   3   DATA distance;
   4     Miles = 26.22;
   5     Kilometers = 1.61 * Miles;
   6   * Print the results;

③ NOTE: The data set WORK.DISTANCE has 1 observations and 2 variables.
NOTE: DATA statement used (Total process time):
      real time          0.11 seconds
      cpu time           0.02 seconds

② 7   PROC PRINT DATA = distance;
   8   RUN;

④ NOTE: There were 1 observations read from the data set WORK.DISTANCE
NOTE: PROCEDURE PRINT used:
      real time          0.06 seconds
      cpu time           0.04 seconds

```

The SAS log above is a blow-by-blow account of how SAS executes the program.

- ① It starts with notes about the version of SAS and your SAS site number.
- ② It contains the original program statements with line numbers added on the left.
- ③ The DATA step is followed by a note containing the name of the SAS data set created (WORK.DISTANCE), and the number of observations (1) and variables (2). A quick glance is enough to assure you that you did not lose any observations or accidentally create a lot of unwanted variables.
- ④ Both DATA and PROC steps produce a note about the computer resources used. At first you probably won't care in the least. But if you run on a multi-user system or have long jobs with large data sets, these statistics may start to pique your interest. If you ever find yourself wondering why your job takes so long to run, a glance at the SAS log will tell you which steps are the culprits.

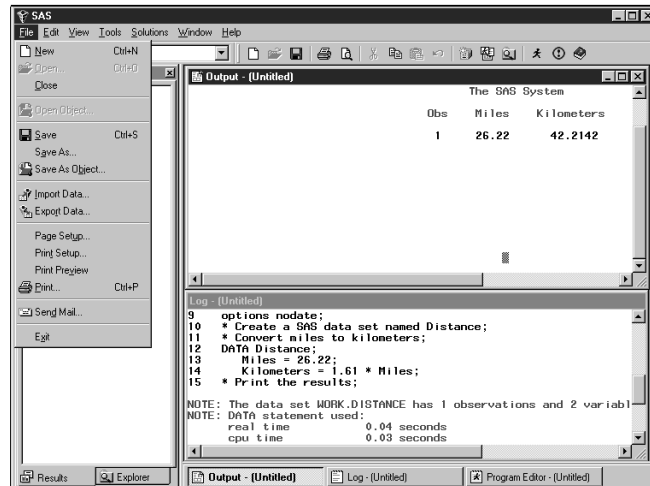
If there were error messages, they would appear in the log, indicating where SAS got confused and what action it took. You may also find warnings and other types of notes which sometimes indicate errors and other times just provide useful information.

1.9 Viewing Your Results in the Output Window

How you view or print your output depends on how you submit your programs. If you submit your program in the SAS windowing environment, then your output will, by default, go to the Output window. If you choose another way to submit your programs, either batch or non-interactive, then your output will probably be in a file on your computer. Use your operating environment's commands to view and print the output file (also called the listing). For example, if you execute your SAS program in non-interactive mode on a UNIX system, then your output will be in a file with an extension `.lst`. To view the file, you can use either the `cat` or `more` commands, and to print the file you would use your system's command for printing files (usually you would type either `lp` or `lpr`).

The Output window After submitting your program in the SAS windowing environment, your results will go to your Output window. If you have the SAS Explorer option turned on (some operating environments have this turned on by default, while others do not), then you will also see a listing of the different parts of your output in your Results window. The following figure shows what your screen might look like after submitting a simple program under Windows.

Printing or saving the contents of the Output window If you want to print or save the entire contents of the Output window, first make the Output window active by clicking in it, then select either `Print` or `Save As` from the `File` pull-down menu. If you are not using a personal computer, then your environment may not be set up for printing from within SAS. If you cannot print from within SAS, then save the output to a file and use your system's command for printing files.



The Results window When you have a lot of output, the Results window can be very helpful. The Results window is like a table of contents for your output. It lists each procedure that produces output, and if you open, or expand, the procedure in the Results tree, you can see each part of the procedure output. The following figure shows what your screen might look like if you ran the Analysis of Variance (ANOVA) procedure.

Scheffe's Test for height

NOTE: This test controls the type I experimentwise error rate.

Alpha	0.05
Error Degrees of Freedom	55
Error Mean Square	13.78182
Critical Value of F	2.53953
Minimum Significant Difference	4.8398

Means with the same letter are not significantly different.

Scheffe Grouping	Mean	N	team
A	54.833	12	pink
A	50.500	12	gold
B	50.333	12	gray
B	49.833	12	blue
B	49.500	12	red

There is one entry in the Results window for the ANOVA procedure. Notice that in the Output window, you see the end of the procedure's output. If you expand the ANOVA procedure in the results tree, by clicking on the plus (+) signs, then you will see all the different parts of the ANOVA output. Double click on the output you want to see, and it will appear at the top of the Output window. The following figure shows what your screen would look like after you double click on the Overall ANOVA item in the Results window.

Printing or saving parts of the output Using the Results window, it is possible to print or save just the parts of the output you want. First highlight the item you want in the Results window, then bring up the context-sensitive menu. In the Windows operating environment you do this with the right mouse button; in other operating environments, it may be the middle or right mouse button. Then select either Print or Save As from the pop-up menu. You may also be able to print or save from the File pull-down menu once you highlight the output part you want. If your SAS environment is not set up for printing from within SAS, then save your results to a file and use your operating environment's command for printing files.

Girls' Heights on Softball Teams

The ANOVA Procedure

Dependent Variable: Height

Source	DF	Sum of Squares	Mean Square
Model	4	228.0000000	57.0000000
Error	55	758.0000000	13.781818
Total	59	986.0000000	

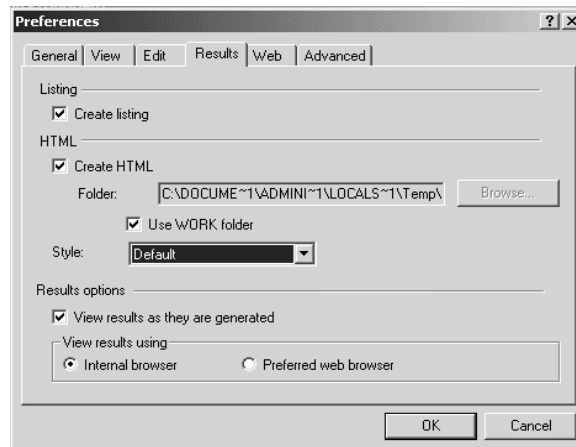
R-Square	Coeff Var	Root MSE	Height
0.231237	7.279190	3.712387	51.1000000

DF	Anova SS	Mean Square
4	228.0000000	57.0000000

1.10 Creating HTML Output

If you are using the SAS windowing environment, then you can create output in Hypertext Markup Language (HTML) format with just a few clicks of your mouse.¹

The Preferences window To turn on HTML output (in Windows, UNIX, or OpenVMS³), select `Options-Preferences` from the Tools menu. This opens the Preferences window. Click on the Results tab to bring it to the front. Here is what the Preferences window looks like in Windows:



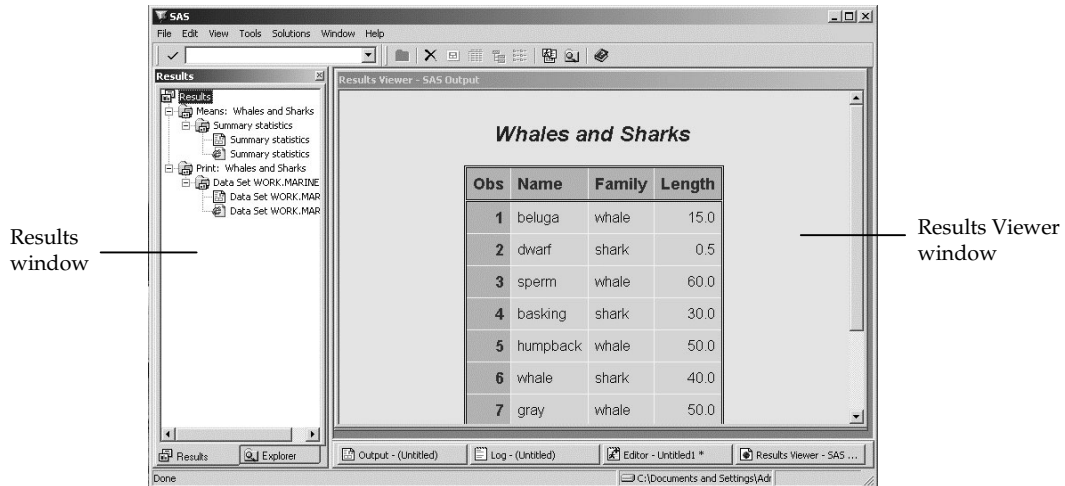
When you first open this window, you will see a check next to `Create Listing`. Listing is the default type of output, and it is what you see in the Output window if you are using the SAS windowing environment, or in the output or listing file if you are running in batch mode. You can turn on HTML output by clicking in the box next to `Create HTML`. To turn off the listing or HTML output, just click to un-check it.

In the Preferences window, you can also select a style for HTML output by clicking on the arrow next to the Style box and scrolling through the list of styles provided with the SAS System. When you are done with the Preferences window, click on the OK button.

The Results Viewer and Results windows Once you have turned on HTML output, then every time you run a program, your output will automatically appear in the Results Viewer window. The following figure shows what you see after running two simple procedures: MEANS and PRINT. Two windows are showing: the Results Viewer window displaying the HTML output, and the Results window listing all the pieces of output in tree form.

¹ If you are not using the SAS windowing environment, you can still produce HTML output by using ODS statements (see chapter 5). In addition, Enterprise Guide allows you to create HTML output in a way that is similar to the one shown in this section with the added bonus that you can also produce RTF and PDF output.

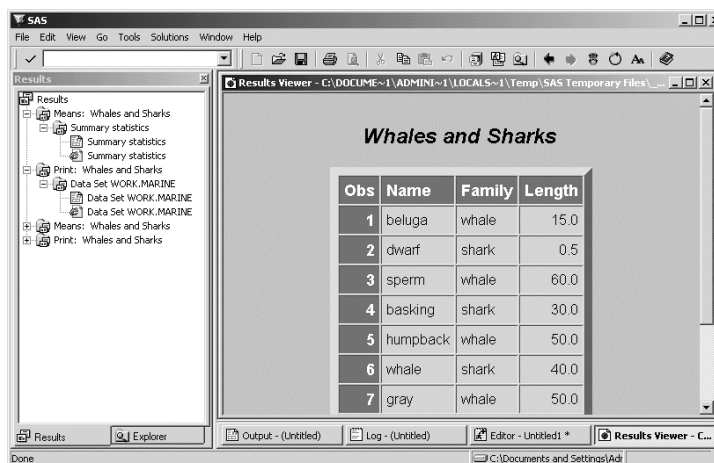
² If you are using OS/390, you will need to modify your registry settings in order to generate HTML interactively. Contact your site's SAS Support Representative for more information.



The Results Viewer window only shows you one piece of output at a time, but you can tell that SAS ran both procedures by looking at the list in the Results window. You can expand the list by clicking on plus (+) signs, or contract it by clicking on minus (-) signs. Since both the listing and HTML output were turned on, each procedure produced two pieces of output: one for listing, and one for HTML. You can display any piece of output by double clicking its name in the Results window.

To save a piece of output in a file, make the Results Viewer window active by clicking on it, then click on the File menu, and choose *Save As...* To print a piece of output, select *Print* from the File menu.

The preceding screen used the DEFAULT style which is the default for HTML output. To see the same output with a different style, just choose a different style in the Preferences window, and re-run your program. Here is the output from the same program using the D3D style.



1.11 Using SAS System Options

System options are parameters you can change that affect the SAS System—how it works, what the output looks like, how much memory is used, error handling, and a host of other things. The SAS System makes many assumptions about how you want it to work. This is good. You do not want to specify every little detail each time you use SAS. However, you may not always like the assumptions SAS makes. System options give you a way to change some of these assumptions.

A long list of system options can be found in your online documentation for the SAS language. The options are grouped by the following general areas:

Communications	Log and procedure output control
Environment control	Macro
Files	Sort
Input control	System administration
Graphics	

Not all options are available for all operating environments. A list of options specific to your operating environment appears in the online documentation for your operating environment. You can see a list of system options and their current values by opening the SAS System Options window or by using the OPTIONS procedure. To use the OPTIONS procedure, submit the following SAS program and view the results in the SAS log:

```
PROC OPTIONS;
RUN;
```

There are four ways to specify system options. Some options can be specified only by using some of these methods. The online documentation for your operating environment tells you which methods are valid for each system option:

1. Your system administrator (this could be you if you are using a PC) can create a SAS configuration file which contains settings for the system options. This file is accessed by the SAS System every time SAS is started.
2. Specify system options at the time you start up SAS from your system's prompt (called the invocation).
3. Change selected options in the SAS System Options window if you are using the SAS windowing environment.
4. Use the OPTIONS statement as a part of your SAS program.

The methods are listed here in order of increasing precedence; method 2 will override method 1, method 3 will override method 2, and so forth. If you are using the SAS windowing environment, methods 3 and 4, the SAS System Options window and OPTIONS statement, will override each other—so whichever was used last will be in effect. Only the last two methods are covered here. The first two are very system dependent; to find out more about these methods see the online documentation for your operating environment.

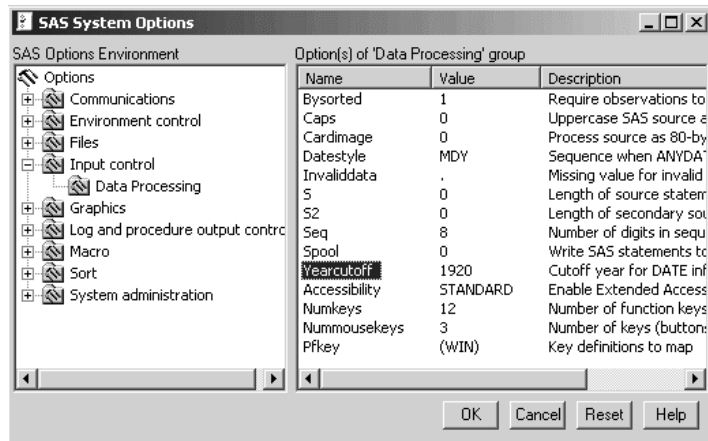
OPTIONS statement The OPTIONS statement is part of a SAS program and affects all steps that follow it. It starts with the keyword OPTIONS and follows with a list of options and their values. For example

```
OPTIONS LINESIZE = 80 NODATE;
```

The OPTIONS statement is one of the special SAS statements which do not belong to either a PROC or a DATA step. This global statement can appear anywhere in your SAS program, but it usually makes the most sense to let it be the first line in your program. This way you can easily see which options are in effect. If the OPTIONS statement is in a DATA or PROC step, then it affects that step and the following steps. Any subsequent OPTIONS statements in a program override previous ones.

The SAS System Options

window You can view and change SAS system options through the SAS System Options window. Open it by either typing OPTIONS in the command line area on your screen, or by selecting it from the Tools pull-down menu. To change the value of an option, first locate the option by clicking on the appropriate category on the left side of the screen. A list of options and their current values will appear on the right side of the screen. Right click on the option itself to modify the value or set it to the default.



Common options The following are some common system options you might want to use:

CENTER NOCENTER	This option works as a switch. CENTER centers your output on the page. NOCENTER left justifies your output.
DATE NODATE	This is also a switch. With DATE, today's date will appear at the top of each page of output; with NODATE it will not.
NUMBER NONUMBER	This switch controls whether or not page numbers appear on each page of SAS output.
LINESIZE = <i>n</i>	With LINESIZE you can control the maximum length of output lines. Possible values for <i>n</i> are 64 to 256.
PAGESIZE = <i>n</i>	PAGESIZE controls the maximum number of lines per page of output. Possible values for <i>n</i> are 15 to 32767.
PAGENO = <i>n</i>	Starts numbering output pages with <i>n</i> .
YEARCUTOFF = <i>yyyy</i>	Specifies the first year in a hundred-year span for interpreting two-digit dates.