

The SQL Query Window

Prepared by
destiny
CORPORATION

Introduction

The SQL Query Window allows you to perform queries based on the standardized SQL language (Structured Query Language) without having any prior experience with SQL.


To enter the SQL Query Window from SAS/ASSIST, make the following selections:

1. Select **Data Management** from the Primary Menu.
2. Select **Query** from the Data Management Menu.
3. Choose the **SQL Query** option.

Selecting Data

Selecting a Table

Select **Active data set** to choose from a list of available data sets. Once a data set has been selected, the **RUN** button will continue the query.



Selecting Columns


The next window to be displayed is the **SQL QUERY COLUMNS** window. Here you can select which columns you would like to be included in your query.

To select all columns, simply choose the **<all columns>** option.

By choosing the  (right arrow), all column names will be copied over to the **Selected Columns** list.

To select columns individually:


1. Select the desired column names from the **Available Columns** list and copy them to the **Selected Columns** list by choosing the  (right arrow).
2. To remove columns from the **Selected Columns** list, select the desired column names from the **Selected Columns** list and choose the  (left arrow) button. The column names will then be deleted from the current selections.



The **Actions** pulldown, **Run Query** option can be used to run the query at any time.

Once a query is run, the output is displayed immediately.

The **File** pulldown, **End** option will allow you to exit from the **OUTPUT** window and go back to the **COLUMNS** screen.




CATNUM	SUPPLIER	TYPE	CPU	DISK	WHOLESAL	RETAIL
1	FLOPPY COMPUTERS	DESKTOP	285	20	550	750
2	FLOPPY COMPUTERS	DESKTOP	285	40	600	800
3	FLOPPY COMPUTERS	DESKTOP	285	100	750	950
4	FLOPPY COMPUTERS	DESKTOP	385DX	40	750	950
5	FLOPPY COMPUTERS	DESKTOP	385DX	100	950	1150
6	FLOPPY COMPUTERS	DESKTOP	385DX	40	950	1150
7	FLOPPY COMPUTERS	DESKTOP	385DX	100	1250	1450
8	FLOPPY COMPUTERS	DESKTOP	485DX	40	1350	1550
9	FLOPPY COMPUTERS	DESKTOP	485DX	100	1550	1750
10	FLOPPY COMPUTERS	DESKTOP	485DX	100	1950	2250
11	FLOPPY COMPUTERS	PORTABLE	285	40	1000	1200
12	FLOPPY COMPUTERS	PORTABLE	385DX	40	1200	1400
13	FLOPPY COMPUTERS	PORTABLE	385DX	100	1400	1600
14	FLOPPY COMPUTERS	PORTABLE	385DX	100	1600	1800
15	FLOPPY COMPUTERS	PORTABLE	385DX	200	1850	2150
16	FLOPPY COMPUTERS	PORTABLE	285	60	1100	1300
17	KEYCHIP COMPUTERS	DESKTOP	285	20	1300	1500
18	KEYCHIP COMPUTERS	DESKTOP	285	40	1400	1600
19	KEYCHIP COMPUTERS	DESKTOP	285	100	1600	1800

Building a Column


Suppose we would like to query our **COMPUTER** database to produce a report showing **SUPPLIER**, **CPU**, **DISK**, and **TYPE** columns, but we would also like to know a profit figure for each of these rows, determined by **RETAIL - WHOLESAL**.

We can add such a column to our query by building a column expression.

1. Select columns **SUPPLIER**, **TYPE**, **CPU**, **DISK** from **Available Columns**.



2. Select the **Build a Column** button. This takes us to the **Build a Column Expression** window.



To create the RETAIL-WHOLESALE column in the Build a Column Expression window:

1. Select RETAIL from the Available Columns list.
2. Select the subtraction operator (-) from the operator pop-up menu.
3. Select WHOLESALE from the Available Columns list. Since this is the last entry we need for our column expression, you can ignore the operator pop-up.
4. Select OK. This will take us to the Express Column Attributes window.

Note: A column can be built from any valid SAS expression, including data step functions.

Expression Column Attributes

We can now enter the attributes of our column into the Express Column Attributes window. (The Column Attributes button on the previous screen would also take us into this window.)

1. Enter an Alias Name for the column. Here we name our column PROFIT.
2. Select the arrow next to the Format field to enter the Format Names window. Or, if you know the name of the format you want, you can simply enter it here.
3. Select a format from the list of Format Names. Here we select the dollar format. We also specify a width and decimal range.
4. Select OK to return to the Expression Column Attributes window.
5. Enter a Label in the Expression Column Attributes window to label the PROFIT column, if desired.
6. Select OK to return to the Build a column Expression window.
7. Select OK again to return to the Columns window.

Displaying the query

1. Choose the Actions pulldown menu, Show option to view the query expression that has been built.

2. The Run Query button will execute the query from here. (Recall that the Actions pulldown, Run option will also execute the query.)

Other options available here include:

- Save Query** to save the query expression in an external file, or to a .QUERY or .SOURCE catalog entry.
- Create Table** to save the data selected by the query.
- Create View** to save the query expression as a view.
- Goback** to return to the Columns window.



The OUTPUT window shows the results of having added a PROFIT column.

Recall that we can return to the Columns window by choosing the File pulldown, End option.


SUPPLIER	TYPE	CPU	DISK	PROFIT
FLOPPY COMPUTERS	DESKTOP	286	20	\$200.00
FLOPPY COMPUTERS	DESKTOP	286	40	\$200.00
FLOPPY COMPUTERS	DESKTOP	286	100	\$200.00
FLOPPY COMPUTERS	DESKTOP	386SX	40	\$200.00
FLOPPY COMPUTERS	DESKTOP	386SX	100	\$200.00
FLOPPY COMPUTERS	DESKTOP	386SX	40	\$200.00
FLOPPY COMPUTERS	DESKTOP	386SX	100	\$200.00
FLOPPY COMPUTERS	DESKTOP	486SX	40	\$200.00
FLOPPY COMPUTERS	DESKTOP	486SX	100	\$200.00
FLOPPY COMPUTERS	DESKTOP	486SX	100	\$300.00
FLOPPY COMPUTERS	PORTABLE	286	40	\$200.00
FLOPPY COMPUTERS	PORTABLE	386SX	40	\$200.00
FLOPPY COMPUTERS	PORTABLE	386SX	100	\$200.00
FLOPPY COMPUTERS	PORTABLE	386SX	100	\$300.00
FLOPPY COMPUTERS	PORTABLE	386SX	200	\$500.00
FLOPPY COMPUTERS	PORTABLE	286	50	\$200.00
KETCHUP COMPUTERS	DESKTOP	286	20	\$200.00
KETCHUP COMPUTERS	DESKTOP	286	40	\$200.00
KETCHUP COMPUTERS	DESKTOP	286	100	\$200.00
KETCHUP COMPUTERS	DESKTOP	386SX	20	\$200.00
KETCHUP COMPUTERS	DESKTOP	386SX	40	\$200.00

Selecting Rows

Thus far we have seen how to select and create columns for our query. We can subset our query by selecting rows.

Where Conditions for Subset - Eliminating Rows


The Locals pull down, Where Conditions for Subset option takes us to the WHERE EXPRESSION window. This window allows us to set up a where expression to subset our data.



We wish to select only rows where SUPPLIER is equal to KETCHUP COMPUTERS.

To set up this expression:

1. Select SUPPLIER from the Available Columns box.
2. Select EQ from the Operator pop-up that appears on the screen.
3. Select <CONSTANT enter value> to get to the pop-up box to enter constants.
4. Enter KETCHUP COMPUTERS and then select the OK button. (You do not need to enter the quotes.)
5. When the WHERE EXPRESSION is completed, select the OK button to return to the Columns window.



Upon running the query, our output shows the subsetting data.

SUPPLIER	TYPE	CPU	DISK	PROFIT
KETCHUP COMPUTERS	DESKTOP	286	20	\$200.00
KETCHUP COMPUTERS	DESKTOP	286	40	\$200.00
KETCHUP COMPUTERS	DESKTOP	286	100	\$200.00
KETCHUP COMPUTERS	DESKTOP	386SX	20	\$200.00
KETCHUP COMPUTERS	DESKTOP	386SX	40	\$200.00
KETCHUP COMPUTERS	DESKTOP	386SX	100	\$200.00
KETCHUP COMPUTERS	DESKTOP	386SX	20	\$200.00
KETCHUP COMPUTERS	DESKTOP	386SX	40	\$200.00
KETCHUP COMPUTERS	DESKTOP	386SX	100	\$200.00
KETCHUP COMPUTERS	DESKTOP	486SX	20	\$200.00
KETCHUP COMPUTERS	DESKTOP	486SX	40	\$200.00
KETCHUP COMPUTERS	DESKTOP	486SX	120	\$200.00
KETCHUP COMPUTERS	DESKTOP	486SX	200	\$200.00
KETCHUP COMPUTERS	PORTABLE	286	20	\$150.00
KETCHUP COMPUTERS	PORTABLE	386SX	20	\$200.00
KETCHUP COMPUTERS	PORTABLE	386SX	40	\$200.00
KETCHUP COMPUTERS	PORTABLE	386SX	120	\$200.00
KETCHUP COMPUTERS	PORTABLE	386SX	60	\$400.00
KETCHUP COMPUTERS	PORTABLE	386SX	120	\$250.00
KETCHUP COMPUTERS	PORTABLE	386SX	200	\$400.00
KETCHUP COMPUTERS	PORTABLE	486SX	200	\$400.00

Other Enhancements


ORDER BY - ORDERING THE RESULTS

The Locals pulldown, Order By option takes us to the ORDER BY COLUMNS window. This window allows us to order our query result by selected columns.

Note that all columns are available to order by regardless of whether we have selected them as columns for our query result. Also, a Column Expression may be built with the Build A Column button, and then used to order results.

To select columns to order by:

1. Copy the column names from the Available Columns list to the Order By Columns list. Ascending Order is the default sequence.
2. To specify Descending Order, select the column to be sorted descending from the Order By Columns list, and then choose the Descending Order button.




Formatting and Labelling Columns

Column formats and labels can be added to enhance our query.

To create a column label for SUPPLIER:

1. Select SUPPLIER from the Selected Columns list in the Columns window.
2. Select the Column Alias/Label button.
3. In the Column Alias and Label window, enter a label in the Label area.
4. Select OK.

Here we label our SUPPLIER column with SUPPLIER NAME.




From the Columns window, the Column Formats button can be used to format a column in a similar fashion.

Titles and Footnotes

Report titles and footnotes can be modified through the Globals pulldown, Options option.


To add a title to our report:

1. Select the Globals pulldown.
2. Select Options.
3. Select Titles.




4. Enter a title in the Title window.
5. To exit the Title window, use the File pulldown, End option. (Depending on what system you are running SAS/ASSIST, you may have other options to end this window.)

Here, we enter a title of PROFITABILITY REPORT.



Viewing the query again from the Action pull down, Show query option, we can see the enhancements made to our query.




Running the query shows our output now subsetting for KETCHUP COMPUTERS, sorted by CPU type, and within each CPU type, sorted by descending values of DISK.

PROFITABILITY REPORT				
SUPPLIER NAME	TYPE	CPU	DISK	PROFIT
KETCHUP	COMPUTERS	DESKTOP	286	100
KETCHUP	COMPUTERS	DESKTOP	286	40
KETCHUP	COMPUTERS	PORTABLE	286	20
KETCHUP	COMPUTERS	DESKTOP	286	20
KETCHUP	COMPUTERS	PORTABLE	386DX	200
KETCHUP	COMPUTERS	PORTABLE	386DX	120
KETCHUP	COMPUTERS	DESKTOP	386DX	100
KETCHUP	COMPUTERS	PORTABLE	386DX	60
KETCHUP	COMPUTERS	DESKTOP	386DX	40
KETCHUP	COMPUTERS	DESKTOP	386DX	20
KETCHUP	COMPUTERS	DESKTOP	386DX	20
KETCHUP	COMPUTERS	DESKTOP	386DX	200
KETCHUP	COMPUTERS	DESKTOP	386DX	100
KETCHUP	COMPUTERS	PORTABLE	386DX	40
KETCHUP	COMPUTERS	DESKTOP	386DX	40
KETCHUP	COMPUTERS	PORTABLE	386DX	20
KETCHUP	COMPUTERS	DESKTOP	386DX	20
KETCHUP	COMPUTERS	DESKTOP	486DX	200
KETCHUP	COMPUTERS	DESKTOP	486DX	120
KETCHUP	COMPUTERS	DESKTOP	486DX	60
KETCHUP	COMPUTERS	PORTABLE	486DX	200
KETCHUP	COMPUTERS	PORTABLE	486DX	200

DISTINCT - SELECTING DISTINCT COLUMN VALUES

Another interesting query option is the capability of pulling out all distinct values of a column or combination of columns. The Locals pull down, Distinct option allows us to specify that we want only distinct column values for all columns listed in the Selected Columns list.



Here we get all distinct combinations of SUPPLIER and TYPE in our data.

SUPPLIER	TYPE
FLOPPY	COMPUTERS
FLOPPY	COMPUTERS
KETCHUP	COMPUTERS
KETCHUP	COMPUTERS

THE TABLES WINDOW

Choosing a Table

From the Locals pulldown, Tables option we can select tables for our query. This step must be done before any other steps in formulating the query are allowed.


Note: If the SQL Query window was entered from SAS/ASSIST, then you already selected a table from the SAS/ASSIST: SQL Query window and therefore can bypass the TABLES window. (You will actually have been taken directly to the COLUMNS window anyway.) You can, however, come back to this window at any time to change tables.

To select a table to query:

1. Select the table from the Available Tables list.

2. Move the table to the Selected Tables list with the  (right arrow) button.

NOTE: More than one table can be selected.




ALIASING A TABLE


It is sometimes useful to give a table a name other than the full two-level *libref.dataset* name shown. This is especially true when multiple tables with some matching variable names are being queried. By simplifying the table name, the query is easier to structure. The new name we give the table is called an *alias*.

To give a table an alias:

1. Select the table from the **Selected Tables** list.
2. Choose the **Alias** button. We are then prompted to enter an alias for the selected table.



3. Enter an alias. Here we simply enter the letter D as the alias for `SAVED.DEMOGRAF`.
4. Select OK.



Note that once we are returned to the TABLES window, our alias is listed in the Selected Tables list alongside the original table name.

5. Select OK from the TABLES window to go directly to the COLUMNS window. You could also get to the COLUMNS window from the Locals pulldown, Columns option.

THE COLUMNS WINDOW

Applying Summary Functions


in section QW1 we saw how to use the COLUMNS window to select columns and build column expressions. Here we look at how to apply **Summary Functions** to a column.

The COUNT(*) Column


The COUNT (*) column will use the COUNT function to return a count of rows that our query has produced. In it's simplest form, we get an output of one number showing a frequency count of rows in our query. The COUNT function is analogous to the FREQ or N statistic.

To select the COUNT(*) column:

Copy the COUNT(*) column from Available Columns to Selected Columns.




Upon running the query, our output shows the result of all rows counted.




What happens when we add other columns to our query, still including the COUNT(*) function?

Here we add the GENDER, SALARY AND STATUS columns to our query.



This time when we run the query, a new window pops up asking what we want to do about the summary function we have requested. Since our COUNT (*) function returns only one row as a result, but the remainder of our query will

Here we will select to **REMERGE** the statistic onto all rows of the query output.



Our output shows that the COUNT(*) result has been "remerged" onto every row of our query.

COUNT(*)	GENDER	SALARY	STATUS
35	F	0	S
35	F	0	S
35	F	0	S
35	F	13000	S
35	F	55000	M
35	F	8000	O
35	F	15000	M
35	F	18000	SEP
35	F	30000	S
35	F	0	M
35	F	12000	M
35	F	0	M
35	F	0	M
35	F	10000	M
35	F	30000	M
35	F	0	M
35	F	15000	M
35	F	13000	M
35	F	15000	M
35	F	20000	M
35	F	10000	M
35	M	0	S
35	M	0	S

Other Summary Functions

To get an average salary figure:

1. Select the **SALARY** column from the **Selected Columns** list.
2. Select the **Summary Functions** button. A pop-up list appears allowing us to select a statistic.

Summary Functions

3. Select the **AVG** statistic. (This is the same as the **MEAN** statistic.)

The Selected Columns now show that the AVG summary function will be applied to the SALARY column. We no longer are requesting the individual SALARY values, but instead an average of all salaries.



Upon running the query, we choose to “remerge” the statistic onto our result. The result shows the AVG(SALARY) column with the average salary for the table remerged onto every row.

```

OUTPUT - [Untitled]

```


GENDER	AUG(SALARY)	STATUS
F	11391.43	S
F	11391.43	S
F	11391.43	S
F	11391.43	S
F	11391.43	N
F	11391.43	D
F	11391.43	N
F	11391.43	DEP
F	11391.43	S
C	11391.43	N

THE GROUP BY WINDOW

What if we would like to produce an average salary for all values of GENDER, i.e. one for males and one for females? To do this we can select to group by GENDER. Any variables selected in the GROUP BY window are used to group summary function results.

To select to group by GENDER in the GROUP BY COLUMNS window:

1. Select GENDER from the Available Columns list.
2. Move GENDER to the Selected Columns list with the  (right arrow) button.
3. When done, select OK to return to the COLUMNS window.



Upon running the query, we now have two values of AVG(SALARY) remerged onto our output. The average salary for GENDER 'F' is remerged onto all rows for females. The average salary for GENDER 'M' is remerged onto all rows for males.

GENDER	AVG(SALARY)	STATUS
F	10580.95	S
F	10580.95	S
F	10580.95	S
F	10580.95	S
F	10580.95	M
F	10580.95	D
F	10580.95	M
F	10580.95	SEP
F	10580.95	S

If we eliminate the STATUS column from our query and simply specify that we want GENDER and AVG(SALARY) columns, we can produce a different style of report.

We now get only two rows of output, one for females and one for males. The summarization was allowed to take place without the need to remerge the results back onto other rows.

GENDER	AVG(SALARY)
F	10580.95
M	12007.14

THE WHERE EXPRESSION WINDOW

Further Subsetting with the WHERE Clause

In order to contrast the WHERE clause and the HAVING clause, we will add a WHERE EXPRESSION to our previous query.

The WHERE expression subsets rows coming *in* to the query and cannot compare against a summary function value. The HAVING expression subset rows going *out* to the output, and can calculate and compare against summary function values. Therefore, any rows that do not meet the WHERE condition are not considered in the calculation of the summary statistic on the SELECT or HAVING clauses.

We saw the <CONSTANT enter value> option earlier. Here we will consider the <LOOKUP distinct values> and <PROMPT at run-time> options to create the WHERE EXPRESSION.

Lookup Distinct Values

The WHERE EXPRESSION window (as well as the HAVING EXPRESSION window) has the capability of looking up and listing for us all distinct values of a variable that we might want to subset by. We can then select from the list which value(s) we want included in our WHERE expression.

1. Begin the WHERE expression by selecting STATUS and the EQ operator. <LOOKUP distinct values> now becomes an option on the Available Columns list.

Note: You can choose from other valid WHERE operators by selecting

These choices are available

2. Select <LOOKUP distinct values>.

3. The Lookup Distinct Values window shows all values for column STATUS. Select M to include 'M'arried people only in our query.

"M" has now been added to our WHERE expression.

Prompt at Run Time


We can also choose be prompted to enter a status value each time we run the query.

1. Again, begin the WHERE expression with the STATUS column, EQ operator.
2. Select <PROMPT at run-time>. You will be asked to Enter a String for this Prompt.
3. Here we enter a string that will appear when we run the query to prompt the user for a STATUS value. The


string can be anything you like. The Prompt String here has been entered as 'Enter a status for this report'.

4. Select OK.


The WHERE EXPRESSION window now shows &PROMPT1 indicating that a value for this variable will be supplied at run time.




Upon running the query, we are prompted with the Prompt String we entered previously. We have the option of entering a STATUS value, or using the Lookup button to view all values of STATUS.



1. Select the Lookup button. The Lookup Distinct Values window appears.
2. Select M for 'M'arried.
3. Select OK.



4. The result of our selection is shown in the Prompt window. Select OK to continue running the query.



Our output shows that only rows for a STATUS of 'M'arried have been included in our query. Also note that the value of our AVG(SALARY) statistic has changed since it is now based on incoming rows where STATUS is 'M'arried only. The rows on our output are those rows that :

1. Met the WHERE expression coming in, AND
2. Met the HAVING expression going out.

GENDER	STATUS	SALARY	AVG(SALARY)
F	M	15000	10123.08
F	M	18000	10123.08
F	M	30000	10123.08
F	M	15000	10123.08
F	M	13000	10123.08
F	M	15000	10123.08
M	M	23000	13716.67
M	M	23000	13716.67
M	M	40000	13716.67

If we run the query again and choose 'S'ingle for status, our result changes again.


GENDER	STATUS	SALARY	AVG(SALARY)
F	S	13000	8500
F	S	30000	8500
M	S	10000	2550
M	S	7800	2550

JOINING TABLES

Resetting the Query

Before beginning a new query, we can select to reset the current query.

1. From the Actions pulldown, select Reset.
2. We are warned to make sure we know what we are doing! Select OK.



3. By default, tables from the SASUSER library are displayed in the Available Tables list. We can remove them from the list, if desired, by selecting SASUSER from the Table Sources list, and selecting Remove SASUSER from Available Tables.

Joining Tables

SQL allows 4 ways to join tables:


- Inner Joins
- Left Joins
- Right Joins
- Full Joins

Simply selecting more than one table implies that the tables are somehow to be joined.


We will discuss joining the `SAVED.MANUF` and `SAVED.PRODUCTS` tables. The `SAVED.PRODUCTS` table contains a row for each product stocked. The `SAVED.MANUF` table is a subset of products for those that are currently being manufactured.

To begin the join:

1. Select **SAVED.MANUF** and **SAVED.PRODUCTS** from Available Tables.
2. Move the table names to the Selected Tables list.
3. Alias **SAVED.MANUF** as **M** and **SAVED.PRODUCTS** as **P**



4. Columns from both tables are displayed. Select PRODNO, PRODESC, and STOCK from P (SAVED.PRODUCTS).
5. Select BATCHNO and NO_EXP from M (SAVED.MANUF).



If we attempt to run the query at this point, without specifying on what column the tables are to be joined, we are warned that we are about to create a Cartesian Product, which is probably not something we really want to do!

At this point, we can select the **WHERE** button to take us to the WHERE EXPRESSION window and create an Inner Join, or **CANCEL** and select some other join options elsewhere.

Inner Joins

Inner joins require a WHERE EXPRESSION to bring rows together that have matching values for a common column or columns. An Inner Join produces an intersection of the

tables being joined. The SAVED.MANUF and SAVED.PRODUCTS tables have the common column PRODNO which is a product number. Therefore, we can join the tables by this column. The Inner Join will produce only rows where a particular PRODNO existed in both tables.

1. Get to the **WHERE EXPRESSION** window, either by having select the **WHERE** button from the above pop-up warning, or from the **Locals** pulldown, **Where Conditions for Subset** option.
2. Put the **WHERE EXPRESSION** together so that is says:
M.PRODNO EQ P.PRODNO.
3. Select **OK**.

Upon running this query, we get a listing of all products that exist in both tables.

PRODNO	PRODESC	STOCK	BATCHNO	NO_EXP
F40019	TABLE	0	F045	30
F40020	CHAIR	5	F058	4
F40021	LOUNGER	15	F122	10
F40022	UMBRELLA	3	F055	8
H10058	15 NETER DOMESTIC	0	H245	10
H10059	20 NETER DOMESTIC	2	H890	10
H20048	15 NETER COMMERCIAL	0	C340	20
H20057	60 NETER HEAVY DUTY	2	H023	10
T20023	SHLL MDEL SHPRM	0	T450	20
T30024	HOE	5	T340	10
T30025	TRIMEL	6	T557	10
T30026	SPADE	10	T750	15

Outer Joins

Left Joins, Right Joins, and Full Joins are all types of joins that SQL calls Outer Joins.

Left Joins join all rows from the left listed table with any matching rows from the right listed table based on some common column(s) values. Therefore, rows in the left table that have no match in the right table are still produced on the output.

Right Joins join all rows from the right listed table with any matching rows from the left listed table based on some common column(s).

Full Joins produce all rows from both tables, with matching rows matched up by the common column(s) values.

Selecting a Join Type


To select to join tables a method to join tables:

- 1. From the Locals pulldown, select Join Type.**

We are prompted to choose an Inner or Outer Join.

Note: We could have created an Inner Join here instead of using the WHERE EXPRESSION window.

2. Click the Matched Join and Unmatched Rows (Outer Join) selection.
3. Select OK.



Columns for Setting Join Criteria

In the Columns for Setting Join Criteria window we can select which columns we want to join by. If we have chosen to produce an Outer Join, we can also select which type of Outer Join we want.

Note that **SAVED.MANUF** is the (Left) table, and **SAVED.PRODUCTS** is the (Right) table. This is because **SAVED.MANUF** was listed above **SAVED.PRODUCTS** when we selected the tables in the **TABLES** window, and is therefore the first, or left, table.

1. Select **PRODNO** from both the left and right tables. Simply selecting it moves the column name to the bottom set of boxes. If the common columns had different names this would not be a problem.

By selecting a Right Join, we are saying that we want all rows from **SAVED.PRODUCTS** regardless of whether a product is currently being manufactured. If a product is being manufactured, information from **SAVED.MANUF** will be matched up.

Upon running the query, we see that we have listed all of our products, along with any related manufacturing information if it exists.

2. Select the  (down arrow) button next to Join Type.
3. From the pop-up selection, choose **Right (Include Unmatched Rows from the Right Table)**.

4. The window now shows that we have selected a Right Join. Select **OK**.