

USER'S GUIDE TO

Fuzzy-Set / Qualitative Comparative Analysis

Charles C. Ragin
Department of Sociology
University of Arizona
Tucson, AZ 85721
cragin@email.arizona.edu

Assisted by:

Sarah Ilene Strand
Claude Rubinson

September 2008

Based on: fsQCA 2.0
Copyright © 1999-2003, Charles Ragin and Kriss Drass
Copyright © 2004-2008, Charles Ragin and Sean Davey

CONTENTS

1. Data Files

A) Opening a Data File	1
B) Opening Data Files of Various Formats	1
Excel	
SPSS	
Word / Notepad	
Other	
C) Saving File Options	3
D) Opening fsQCA Data in Other Formats	3
SPSS	
Excel	

2. Data Editor

A) Entering Data	5
Variable Names	
Define Variable Type	
Define Column Format	
Define Variable Precision	
Number of Cases	
B) Editing Data	7
Add / Delete Variables	7
Compute Variables	8
1) Arithmetic Operators	
2) Relational Operators	
3) Arithmetic Functions	
4) Other Operators	
Compute Variables If	12
Recode Variables	13
1) Recode Variables Into Same Variables	
2) Recode Variables Into Different Variables	
Calibrating Fuzzy Sets.....	16
Add / Insert Cases	18
Delete Cases	18
Drop / Keep Cases If	19
Select Cases If	20
C) Working with Output	21

Printing Viewer Documents
Saving Results

3. Basic Statistics and Graphs

A) Frequencies	23
Obtaining Frequencies and Statistics	
Sample Output	
B) Descriptives	24
Obtaining Descriptive Statistics	
Sample Output	
C) Crosstabulations	25
Obtaining Crosstabulations	
Sample Output	
D) Graphs	27
Bar chart	
Histogram	
XY Plot	

4. Crisp-Set Analysis


A) Basic Concepts	33
1) Use of Binary Data	
2) Boolean Negation	
3) Use of Truth Tables to Represent Data	
4) Groupings	
5) Boolean Addition	
6) Boolean Multiplication	
7) Combinatorial Logic	
Minimization	37
1) Use of Prime Implicants	39
2) Use of De Morgan's Law	42
3) Necessary and Sufficient Causes	42
B) Data	43
C) Analysis	44
Truth Table Algorithm.....	44
Limited Diversity and Counterfactual Analysis.....	50
Quine Algorithm (QCA 3.0).....	53
Listing Truth Table Rows.....	55

Rules for Including Configurations.....	57
Setting Configurations to Minimize.....	61
The Eight Most Common Ways of Specifying a Crisp Analysis.....	63
Choosing Solution Terms Using the Prime Implicant Chart	66
Frequencies for Solution Terms	66
D) Output (Crisp Sets).....	67
Output from Quine Procedure.....	67
Output from Truth Table Procedure.....	68
5. Fuzzy-Set Analysis	71
A) Operations on Fuzzy Sets.....	71
Logical AND	
Logical OR	
Negation	
B) Fuzzy Sets, Necessity, and Sufficiency (Fuzzy Subset Relation).....	72
C) Using the Fuzzy Truth Table Algorithm.....	74
D) Output for “Specify Analysis” Option	81
E) Output for “Standard Analyses” Option	84
F) Consistency and Coverage.....	85

1. DATA FILES

A) Opening a Data File

- FsQCA opens with the following window:



- From the menu choose:
 - File
 - Open
 - Data...

- In the Open File dialog box, select the file you want to open.

- Click *Open*.

B) Opening Data Files of Various Formats

Data files come in a wide variety of formats, and the software is designed to handle the following:

- QCA (*.qdm and *.qvn): these formats are produced when data was entered in QCA 3.0 (DOS) software, whereby .QDM is the

data file, and .QVN the file that contains the list of variable names. **Only crisp-set data can be input or saved in this format. It is a “legacy” format and soon will be dropped.**

- Comma separated (*.csv) comma delimited file, produced by Excel and other spreadsheet software
- Space separated (*.txt) space delimited file; can be created in WORD or other word processing software **and saved as text only**
- Tab separated (*.dat) tab delimited file, can be created using SPSS and other statistical software packages

The recommended formats are *.csv (Excel) and *.dat (SPSS).

Please note that fsQCA makes the following assumptions about the structure of *.csv, *.dat and *.txt data files. First, and most important, fsQCA assumes that the cells in the first row of the spreadsheet contain variable names for their respective columns. Second, fsQCA assumes that the data begin in the second row of the spreadsheet and that each case is a single row. Finally, fsQCA assumes that each column contains cells of the same type of data. Data types can vary across columns, but they must be consistent within columns. **Please remember to use very simple variables names, using only alphanumeric characters with no embedded punctuation or spaces. For example, “GNP1990” is OK, but “GNP 1990” and “GNP-1990” are not.**

- Saving / Opening data originally created in **Excel**:
Save the Excel file in CSV (comma delimited) format. Make sure that the first row of the Excel data spreadsheet contains the variable names. Open in fsQCA.
- Saving / Opening data originally created in **SPSS**:
Save the SPSS file in DAT (tab delimited) format. SPSS will ask you whether you want to “Write variable names to spreadsheet.” **Do not uncheck this option.**
- Saving / Opening data originally created in **Word / Notepad**:
Enter the data delimited by spaces. Make sure that the first line contains the variable names, also separated by spaces. Save the file in a TXT (Text only) format, TXT (Text with Line Breaks), TXT (MS-DOS), or TXT (MS-DOS with Line Breaks). Open in fsQCA.

- Saving / Opening data originally created in **other programs/ formats**:

By using STAT TRANSFER (Version 6 or later) you can transfer data from a variety of different formats into a CSV (comma separated) file. Open in fsQCA.

<u>Program</u>	<u>File Extension</u>	<u>Program</u>	<u>File Extension</u>
1-2-3	WK*	OSIRIS	DICT
Access	MDB	Paradox	DB
DBASE and compatibles	DBF	Quattro Pro	WQ*, WB*
Epi Info	REC	SAS for Windows and OS/2	SD2
EXCEL	XLS	SAS for Unix	SSD*
FoxPro	DBF	SAS Transport Files	XPT
Gauss	DAT	S-PLUS	[none]
JMP	JMP	SPSS Data Files	SAV
LIMDEP	CPJ	SPSS Portable Files	POR
Matlab	MAT	Stata	DTA
Mineset	SCHEMA	Statistica	STA
Minitab	MTW	SYSTAT	SYS
ODBC Data Source	[none]		

➔ Covert to: ASCII - Delimited TXT, CSV

C) Saving File Options

- From the menu choose:

File
Save...

- The modified data file is saved, overwriting the previous version of the file of the same name and location.

Or: ➤ To save a new data file or save data in a different format, from the menu choose:

File
Save As...

- Select a file type from the drop-down list.
- Enter a filename for the new data file. We recommend using *.csv or *.dat.

D) Opening fsQCA Data in Other Formats

Once you have your data in the fsQCA program and have completed some preliminary analyses, you have the option to either edit the data in fsQCA (see Chapter2), or edit your data with the help of software packages you are more familiar with (i.e. SPSS or Excel). Similarly, you can either display the data graphically with the fsQCA program (see section 3), or turn to SPSS or Excel for more elaborate graphical representations. If you choose SPSS or Excel for these operations, you need to save the fsQCA file and transfer it to the program of your choice.

SPSS

- In order to open fsQCA data in **SPSS**, save the fsQCA data spreadsheet in tab separated format (*.dat). Make sure that the string variables in the fsQCA data file are written without spaces in between them (no embedded spaces are allowed).
- In SPSS choose:
 - File
 - Read Text Data.....
- Open the fsQCA file you have just saved.
- SPSS will ask you several questions regarding your file. Check the following options:

Does your text file match a predefined format?	No
How are your variables arranged?	Delimited
Are variable names included at the top of your file?	Yes
The first case of data begins with line number?	2
How are your cases represented?	Each line represents a case
How many cases do you want to import?	All of the cases
Which delimiters appear between variables?	Tab
What is the text qualifier?	None
Would you like to save this file format for future use?	Y/N
Would you like to paste the syntax?	No

Then click FINISH

- You can now edit the data and display it graphically in SPSS.
- In order to transfer the SPSS file back to fsQCA, see Chapter 1) B) SPSS.

Excel

- In order to open fsQCA data in **Excel**, save the fsQCA data spreadsheet in comma separated format (*.csv). Make sure that the string variables in the fsQCA data file are written without spaces in between them (no embedded spaces).
- In Excel choose:
 - File
 - Open....
- Open the fsQCA file you have just saved.
- You can now edit the data and display it graphically in Excel.
- In order to transfer the Excel file back to fsQCA, see above.

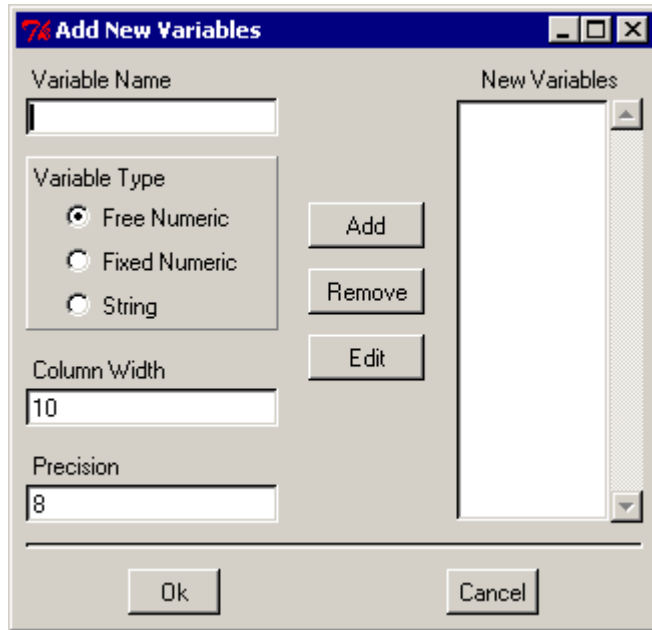
2. DATA EDITOR

A) Entering Data (creating a data file from scratch, in fsQCA)

➤ From the menu choose:

File
New...

➤ The *Add New Variables* window will open.




➤ Enter the **variable name**. The following rules apply to variable names:

- The length of the name cannot exceed fifteen characters.
- Each variable name must be unique; duplication is not allowed.
- Variable names are not case sensitive. The names NEWVAR, NewVAR and newvar are all considered identical.
- Variable names cannot include spaces or hyphens.
- Only alphanumeric characters may be used (0-9, a-Z)

➤ Choose between the three possible **variable types**. By default, all new variables are assumed to be free numeric. You can change the variable type by clicking on one of the three options. The available data types are:

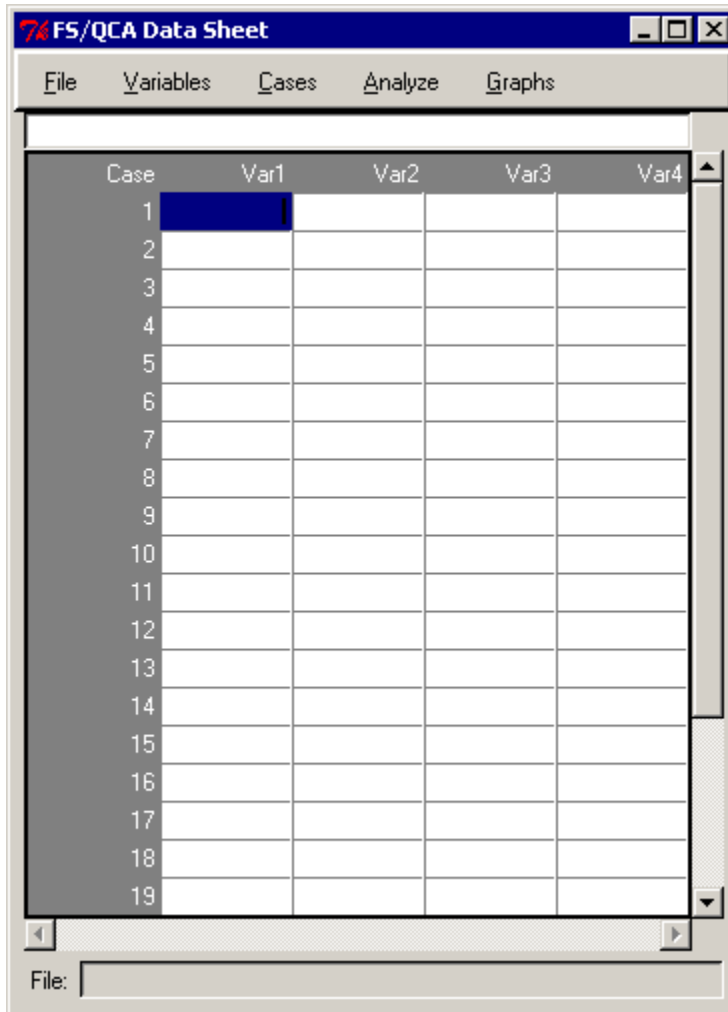
- Free numeric: data that is entered under this type can have any kind of precision (up to the limit of decimals indicated below).
- Fixed numeric: data of this type will always be fixed to an indicated number of decimals
- String

- Determine **column width**. This option controls the width of columns in the Data Sheet.
- Choose the **level of precision**. With this option you can control the number of decimal places.
- Add the variable to the New Variable List by clicking the *Add* button.
- In addition to adding new variables, you can delete variables already added to the New Variable List by highlighting the variable and clicking on the *Remove* button. Likewise, if you need to make changes regarding the type, column width or precision of a variable you've already added to the New Variable List, you can do so by highlighting this variable and clicking on the *Edit* button.
- After entering all variables click *Ok*. A new window will open and ask you for the **number of cases** in your data set:



Note: In general, fsQCA is able to process a large number of cases. However, an important feature of fsQCA is that it deals with combinations of causes; therefore adding more variables will influence computational time much more than adding more cases. The number of possible combinations is 2 to the k power, where k is the number of causal conditions. As a rule of thumb, 10 or fewer causal conditions (i.e., 1024 possible combinations) is not a problem. When dealing with more than 10 conditions, it is basically just a matter of the amount of time you are willing to wait for the program to do the analyses.

- Enter the number of cases of your data set, press the *Ok* button, and the Data Sheet window will open:



- **Enter the data values.** You can enter data in any order. You can enter data by case or by variable, for selected areas or individual cells. The active cell is highlighted with a darker color. When you select a cell and enter a data value, the value is displayed in the cell editor under the menu bar. Values can be numeric or string. Data values are not recorded until after you press *Enter*.
- Before closing the Data Sheet you need to save it in order not to lose the entered information.

B) Editing Data

Add / Delete Variables

- In order to **add variables** to an already existing Data Sheet, choose:
 - Variables
 - Add...
- Enter the variable name and its specifications and press the *Add* button.

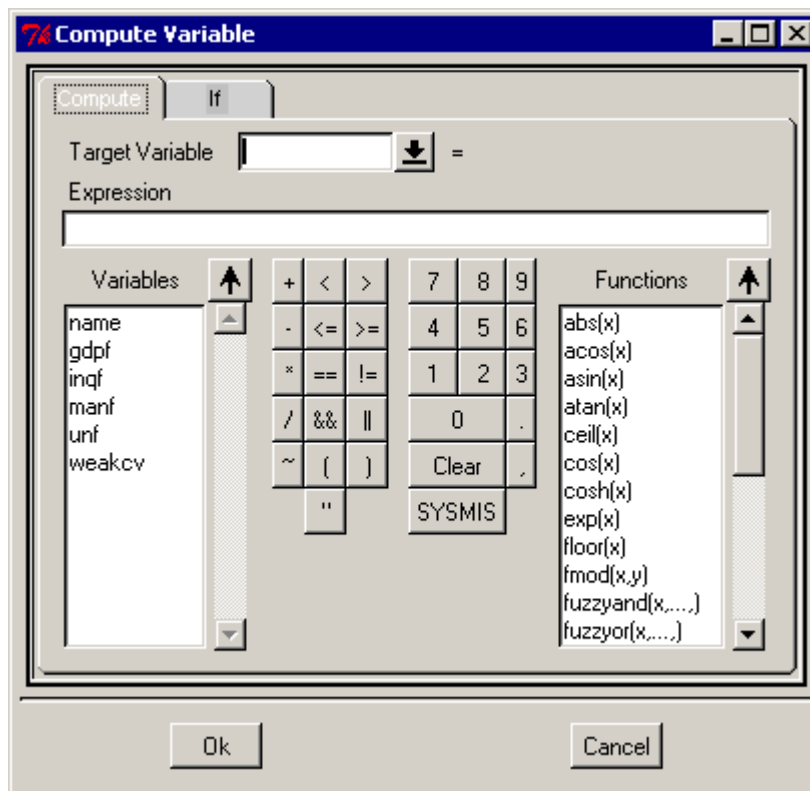
- In order to **delete existing variables** in the Data Sheet, choose:
Variables
Delete...

- Highlight the variables that you want to delete, transfer them into the Delete Column and click *Ok*.

Compute Variables

- In order to **compute new variables** out of existing ones or numeric or logical expressions, choose:
Variables
Compute...

- The following window will open (with the names of the variables in your data file listed in the window):



- Type the name of a single target variable. It can be an existing variable or a new variable to be added to the working data file. ***Do not use a single letter as a variable name (e.g., "X"). This will cause the compute function to crash.***
- To build an expression, either paste components into the Expression field or type directly in the Expression field.

1) Arithmetic Operators

- +** **Addition.** The preceding term is added to the following term. Both terms must be numeric.
- **Subtraction.** The following term is subtracted from the preceding term. Both terms must be numeric.
- *** **Multiplication.** The preceding and the following term are multiplied. Both terms must be numeric.
- /** **Division.** The preceding term is divided by the following term. Both terms must be numeric, and the second must not be 0.

2) Relational Operators

- <** **Logical Less Than.** True (=1) for numeric terms if the preceding term is less than the following term. True for string terms if the preceding term appears earlier than the following term in the collating sequence (in alphabetical order). This operator is normally used only in a logical condition.
- >** **Logical Greater Than.** True (=1) for numeric terms if the preceding term is greater than the following term. True for string terms if the preceding term appears later than the following term in the collating sequence (in alphabetical order). This operator is normally used only in a logical condition.
- <=** **Logical Less Than Or Equal.** True (=1) for numeric terms if the preceding term is less or equal than the following term. True for string terms if the preceding term appears earlier than the following term in the collating sequence (in alphabetical order), or if the two are equal. This operator is normally used only in a logical condition.
- >=** **Logical Greater Than Or Equal.** True (=1) for numeric terms if the preceding term is greater or equal than the following term. True for string terms if the preceding term appears later than the following term in the collating sequence (in alphabetical order), or if the two are equal. This operator is normally used only in a logical condition.
- ==** **Logical Equality.** True (=1) for terms that are exactly equal. If string terms are of unequal length, the shorter term is padded on

the right with spaces before the comparison. This operator is normally used only in a logical condition.

!= **Logical Inequality.** True (=1) for terms that are not exactly equal. If string terms are of unequal length, the shorter term is padded on the right with spaces before the comparison. This operator is normally used only in a logical condition.

&& **Logical And.** True (=1) if both the preceding and the following term are logically true. The terms may be logical or numeric; numeric terms greater than 0 are treated as true. This operator is normally used only in a logical condition.

|| **Logical Or.** True if either the preceding or the following term are logically true. The terms may be logical or numeric; numeric terms greater than 0 are treated as true. This operator is normally used only in a logical condition. This operator only works by pasting the symbol into the Expression Field.

~ **Logical Not.** True if the following term is false. $1 -$ (numeric term). This operator is normally used only in a logical condition.

3) Arithmetic Functions

abs (x)	Returns the absolute value of x, which must be numeric.
acos (x)	Returns the arc cosine (inverse function of cosine) of radians, which must be a numeric value between 0 and 1, measured in radians.
asin (x)	Returns the arc sine (inverse function of sine) of radians, which must be a numeric value between 0 and 1, measured in radians.
atan (x)	Returns the arc tangent (inverse function of tangent) of radians, which must be a numeric value, measured in radians.
ceil (x)	Returns the integer that results from rounding x up (x must be numeric). Example: $\text{ceil}(2.5) = 3.0$
calibrate	Transforms an interval or ratio scale variable into a fuzzy set; see below for details
cos (x)	Return the cosine of radians, which must be a numeric value, measured in radians.

cosh (x)	Returns the hyperbolic cosine $[(e^x + e^{-x})/2]$ of radians, which must be a numeric value, measured in radians. X cannot exceed the value of 230.
exp (x)	Returns e raised to the power x, where e is the base of the natural logarithms and x is numeric. Large values of x ($x > 230$) produce results that exceed the capacity of the machine.
floor (x)	Returns the integer that results from rounding x down (x must be numeric). Example: floor (2.5) = 2.0
fmod (x,y)	Returns the remainder when x is divided by modulus (y). Both arguments must be numeric, and modulus must not be 0.
fuzzyand (x,...)	Returns the minimum of two or more fuzzy sets. Example: fuzzyand (1.0, 0.1) = 0.1
fuzzyor (x,...)	Returns the maximum of two or more fuzzy sets. Example: fuzzyor (1.0, 0.1) = 1.0
fuzzynot (x)	Returns the negation (1-x) of fuzzy sets (same as Logical Not '~'). Example: fuzzynot (0.8) = 0.2
int (x)	Returns the integer part of x. Numbers are rounded down to the nearest integer.
log (x)	Returns the base-e logarithm of x, which must be numeric and greater than 0.
log10 (x)	Returns the base-10 logarithm of x, which must be numeric and greater than 0.
pow (x,y)	Returns the preceding term raised to the power of the following term. If the preceding term is negative, the following term must be an integer. This operator can produce values too large or too small for the computer to process, particularly if the following term (the exponent) is very large or very small.
round (x)	Returns the integer that results from rounding x, which must be numeric. Numbers ending in .5 exactly are rounded away from 0. Example: round (2.5) = 3.0
sin (x)	Returns the sine of radians, which must be a numeric value, measured in radians.

sinh (x)	Returns the hyperbolic sine $[(e^x - e^{-x})/2]$ of radians, which must be a numeric value, measured in radians. X cannot exceed the value of 230.
square (x)	Returns the square of x, which must be numeric.
sqrt (x)	Returns the positive square root of x, which must be numeric and not negative.
tan (x)	Returns the tangent [sine/cosine] of radians, which must be a numeric value, measured in radians.
tanh (x)	Returns the hyperbolic tangent $[(e^x - e^{-x}) / (e^x + e^{-x})]$ of radians, which must be a numeric value, measured in radians.

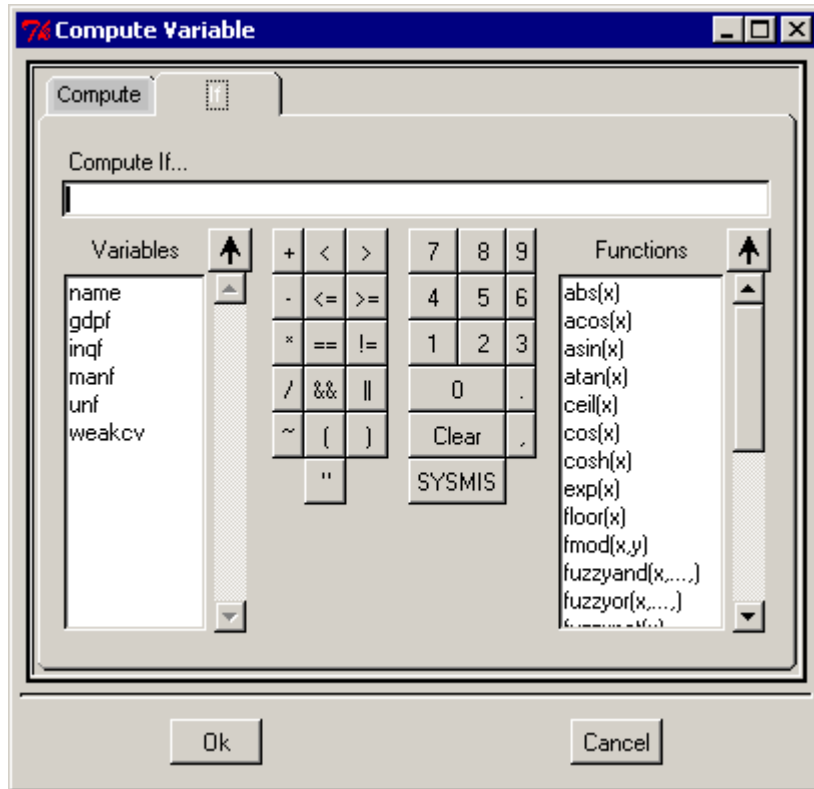
4) Other Operators

()	Grouping. Operators and functions within parentheses are evaluated before operators and functions outside the parentheses.
"	Quotation Mark. Used to indicate the values of string variables. Example: Compute if...: Variable == "NA"
SYSMIS	System Missing. Used when selecting subsets of cases. Example: Select if...: Variable == SYSMISS
Clear	Deletes the text in the Expression Field.

Compute Variables If ...

The IF dialog box allows you to apply data transformations to select subsets of cases, using conditional expressions. A conditional expression returns a value of *true*, *false*, or *missing* for each case.

- In order to select cases relevant for the data transformation, choose:
 - Variables
 - Compute...
- Click on *If* and the following dialog box will appear:



- If the result of a conditional expression is true, the transformation is applied to the case.
- If the result of a conditional expression is false or missing, the transformation is not applied to the case.
- 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.

Recode Variables

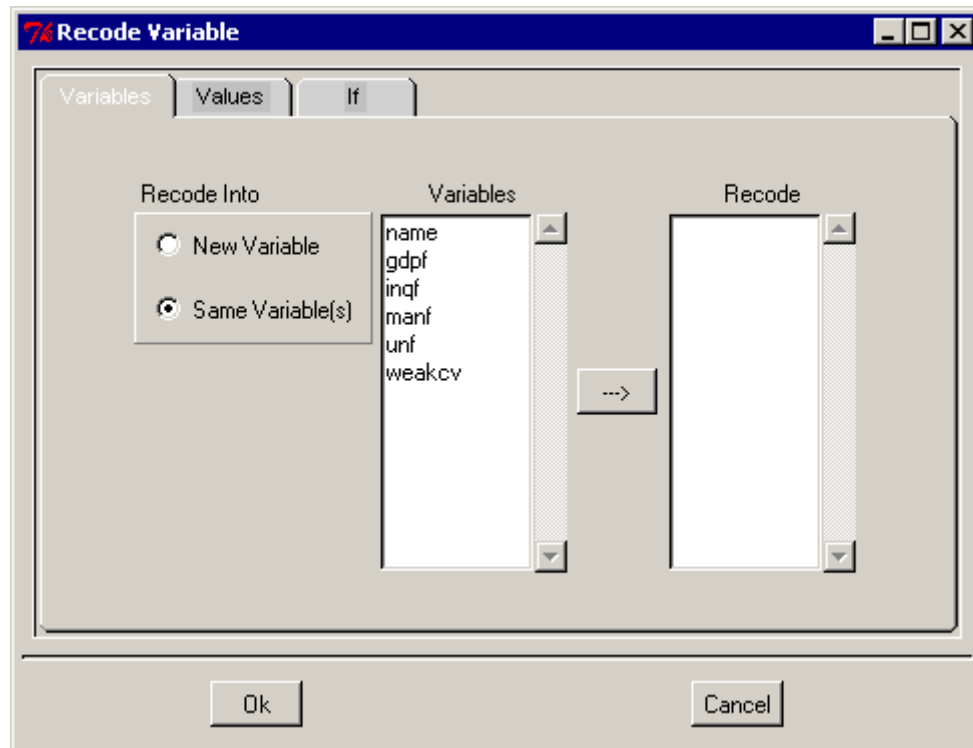
You can modify data values by recoding them. This is particularly useful for collapsing or combining categories. You can recode the values within existing variables, or you can create new variables based on the recorded values of existing variables.

1) Recode Into Same Variables reassigns the values of existing variables or collapses ranges of existing values into new values. You can recode numeric and string variables. You can recode single or multiple variables—they do not have to be all the same type. You can recode numeric and string variables together.

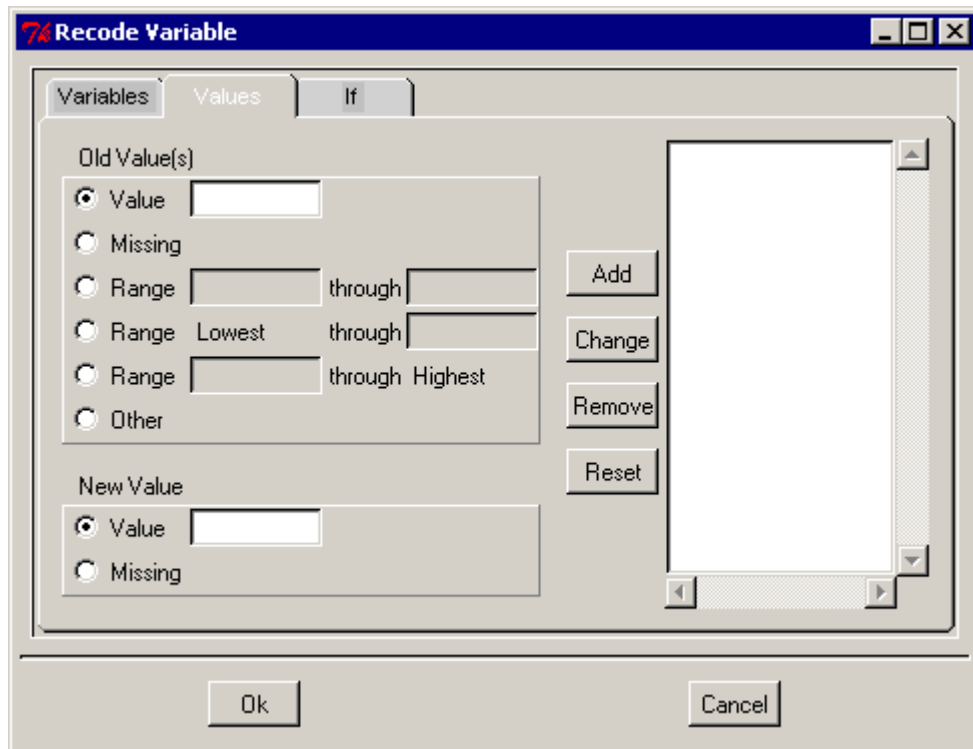
- In order to recode the values of a variable choose:

Variables
Recode...

- The following window will open:



- Select the *Same Variable(s)* option.
- Select the variables you want to recode (numeric or string).
- Optionally, you can define a subset of cases to recode. The *If* dialog box for defining subsets of cases is the same as the one described for *Compute Variables*.
- Click *Values* to specify how to recode the values. The following dialog box will appear:



➤ You can define values to recode in this dialog box.

Old Value(s). The value(s) to be recoded. You can recode single values, ranges of values, and missing values. Ranges cannot be selected for string variables, since the concept does not apply to string variables. Ranges include their endpoints and any user-missing values that fall within the range.

New Value. The single value into which each old value or range of values is recoded. You can enter a value or assign the missing value.

➤ Add your specifications to the list on the right. Optionally, you can remove or change already entered specifications. Use the *Reset* button to clear the list of specifications.

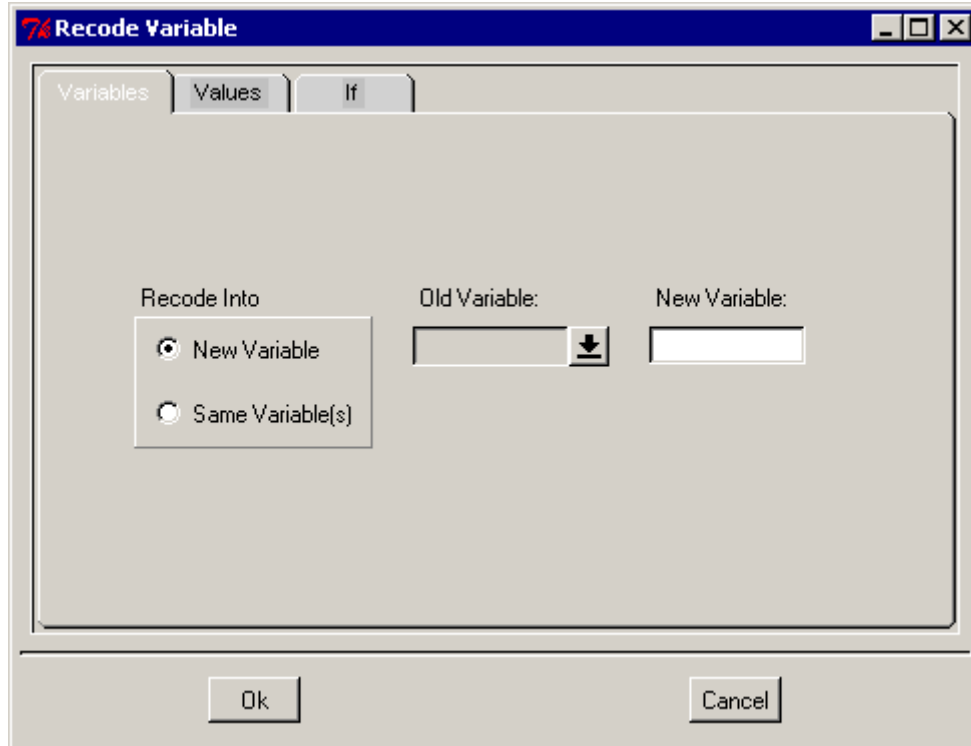
2) **Recode Into Different Variables** reassigns the values of existing variables or collapses ranges of existing values into new values for a new variable.

- You can recode numeric and string variables.

- You can recode numeric variables into string variables and vice versa.

➤ In order to recode the values of an old variable into a new variable, choose:
 Variables
 Recode
 New Variable...

- The following window will appear:



- Select the (old) variable you want to recode either by selecting it from the drop-down menu or by typing it into the field.
- Enter an output (new) variable name.
- Click *Values* and specify how to recode values.
- Optionally, you can define a subset of cases to recode. The *IF* dialog box for defining subsets of cases is the same as the one described for *Compute Variables*.

Calibrating Fuzzy Sets

In order to transform conventional variables into fuzzy variables, it is necessary to calibrate them, so that the variables match or conform to external standards. Most social scientists are content to use uncalibrated measures, which simply show the positions of cases relative to each other. Uncalibrated measures, however, are clearly inferior to calibrated measures. For example, with an uncalibrated measure of democracy it is possible to know that one country is more democratic than another or more democratic than average, but still not know if it is more a democracy or an autocracy.

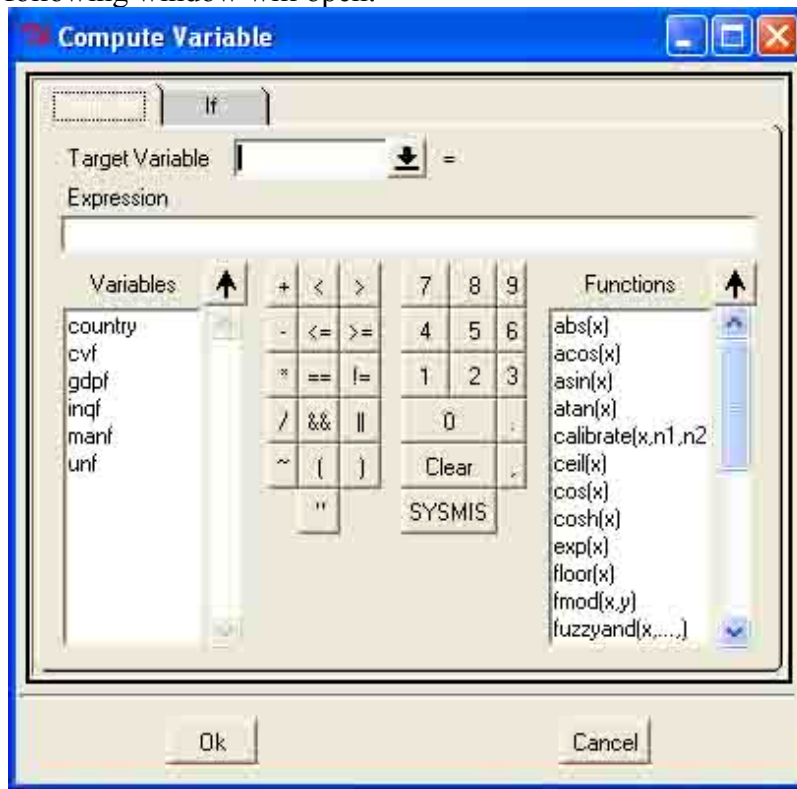
Fuzzy sets are calibrated using theoretical and substantive criteria external to the data, and take into account the researcher's conceptualization, definition, and labeling of the

set in question. The end product is the fine-grained calibration of the degree of membership of cases in sets, with scores ranging from 0.0 to 1.0.

The researcher must specify the values of an interval-scale variable that correspond to three qualitative breakpoints that structure a fuzzy set: the threshold for full membership (fuzzy score = 0.95), the threshold for full nonmembership (fuzzy score = 0.05), and the cross-over point (fuzzy score = 0.5). These three benchmarks are used to transform the original ratio or interval-scale values into fuzzy membership scores, using transformations based on the log odds of full membership.

- From the menu choose:
Variables
Compute....

- The following window will open:



- Name the target variable (using 2-8 standard alphanumeric characters and no spaces, dashes, or punctuation) for the fuzzy set.
- Click *calibrate(x,n1,n2,n3)* in the Functions menu. Click the arrow up next to the word “Functions”.
- Edit the expression *calibrate(,,)*, for example, “*calibrate(oldvar,25,10,2)*”. Here, *oldvar* is the name of the existing interval or ratio scale variable already in the file, the first number is the value of *oldvar* that corresponds to the threshold for full


membership in the target set (0.95), the second number is value of oldvar that corresponds to the cross-over point (0.5) in the target set, and the third number is the value of oldvar that corresponds to the threshold for nonmembership in the target set (0.05).

- Click “OK”.
- Check the data spreadsheet to make sure the fuzzy scores correspond to the original values in the manner intended. It may be useful to sort the variable in descending or ascending order, using the pull-down menus. The result is a fine-grained calibration of the degree of membership of cases in sets, with scores ranging from 0 to 1.

Add / Insert Cases

- In order to **add cases** into an already existing data sheet, choose:
Cases
Add....

- The following window will appear:




- Enter the number of cases you want to add to the existing number of cases. The additional case(s) will appear at the end (bottom) of the data set.
- In order to **insert cases** into an already existing data sheet, choose:
Cases
Insert....
- The *Add* window will appear.
- Enter the number of cases you want to insert above the case you highlighted in the data sheet. The additional case(s) will appear above the case (row) in which you have highlighted a cell.

Delete Cases

- In order to **delete single cases** from an already existing data sheet, choose:
Cases
Delete...

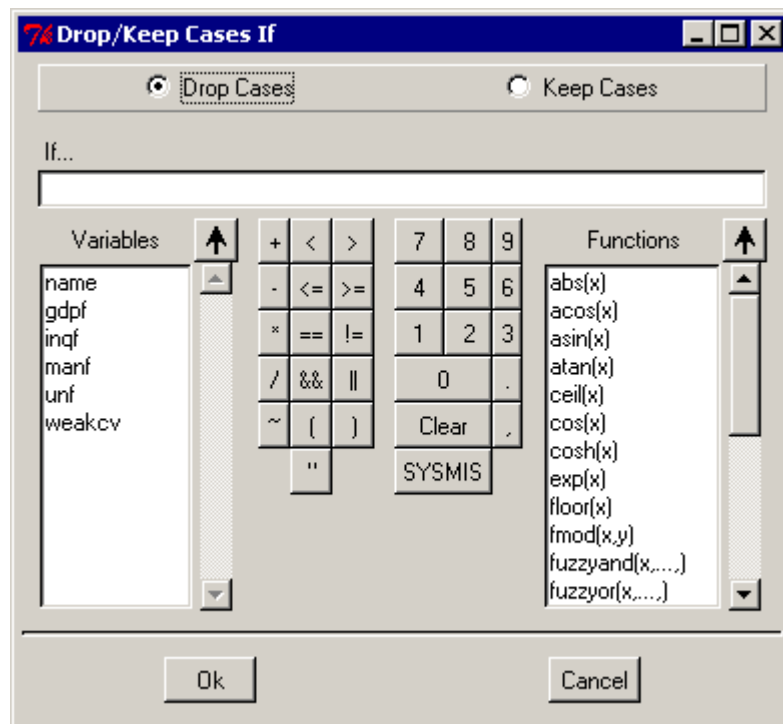
- The following window will appear:



- With this function you can only delete one case at a time.
- The program will ask you whether you want to delete the case in which you highlighted a cell in the data sheet. In the example above, a cell in case #4 was highlighted.

Drop / Keep Cases If

- In order to **drop or keep groups of cases** that have certain specifications, choose:
Cases
Drop / Keep If...
- The following window will appear:



- Specify the cases you want to keep or drop and click *OK*.

- If the result of a conditional expression is true, the case is dropped or kept. If the result of a conditional expression is false or missing, the case is not dropped or kept.

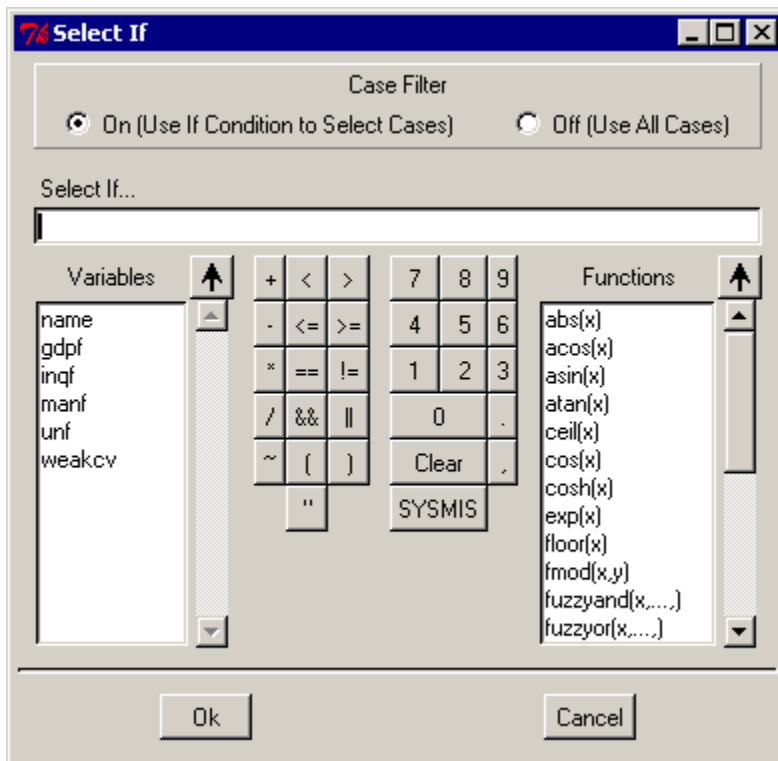
Select Cases If

Select Cases If provides several methods for selecting a subgroup of cases based on criteria that include variables and complex expressions, like:

- Variable values and ranges
- Arithmetic expressions
- Logical expressions
- Functions

Unselected cases remain in the data file but are excluded from analysis. Unselected cases are indicated with parentheses around the row number in the Data Sheet.

- In order to **select a subset of cases for analysis**, choose:
Cases
Select If...
- The following window will open:



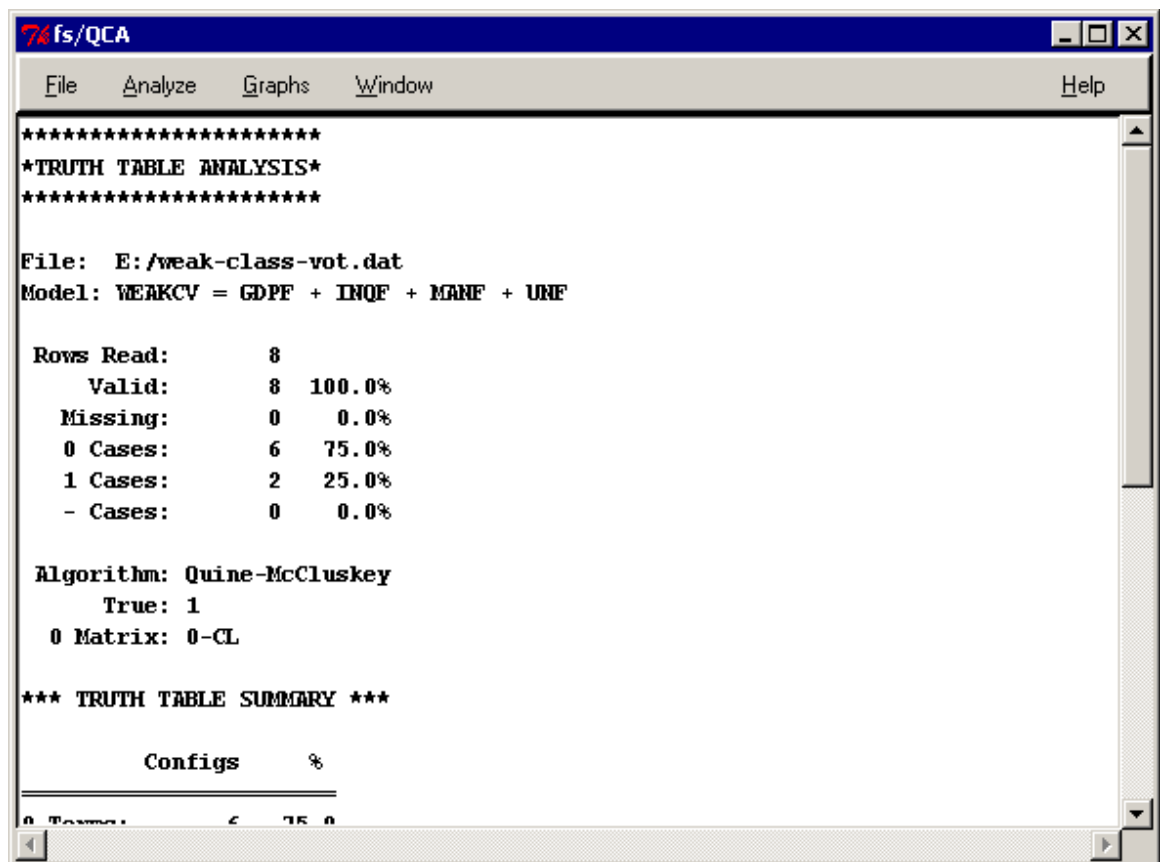
- Specify the criteria for selecting cases.

- If the result of a conditional expression is true, the case is selected. If the result of a conditional expression is false or missing, the case is not selected.
- Most conditional expressions use one or more of the six relational operators (<, >, <=, >=, ==, !=) on the calculator pad.
- Conditional expressions can include variable names, constants, arithmetic operators, numeric and other functions, logical variables, and relational operators.

Note: “Select If” works best when it is univariate. For example, if you want to use the “Select If” function combining two logical statements, e.g., both a logical AND and a logical NOT, try creating a new variable (with compute or recode) that reflects your selection criteria and then use the new variable with “Select If.”

C) Working with Output

When you run a procedure, the results are displayed in the fsQCA window. You can use the scroll bars to browse the results.



- In order to **print output**, choose:
File
Print...

- Your computer specific printer options window will appear, in which you can specify your printing options.
- All output will be printed with the day, date, time, and page number on the top of each printed page.
- The output is written in monospace New Courier (10) in order allow simple transport between programs. Therefore, if you open the *.out file in SPSS or some other program, the numbers in the tables will be slightly dislocated, unless you specify the appropriate font.
- Output may also be copied and pasted into Word, Wordpad, Text, or other files. In addition, it is possible to type directly in the output window in order to insert text or simple tables. You can also delete parts of the output by backspacing in the output window, or highlighting parts and pressing the backspace or delete key on your keyboard.
- In order to **save output**, choose:
 - File
 - Save As
 - Output....
- fsQCA will give you the option to save your output as either a fsQCA specific Output File (*.out), or any other text format you may choose (e.g., txt).

3. BASIC STATISTICS AND GRAPHS

Note: Statistics, histograms, and XY plots cannot be obtained for a variable if, for any case, the variable value is assigned as “don’t care” (a dash) in the data spreadsheet. Barcharts, however, may still be created. (The use of “don’t care” codings in the data spreadsheet is rare and is usually only associated with the entry of a truth table directly into the data spreadsheet.)

A) Frequencies

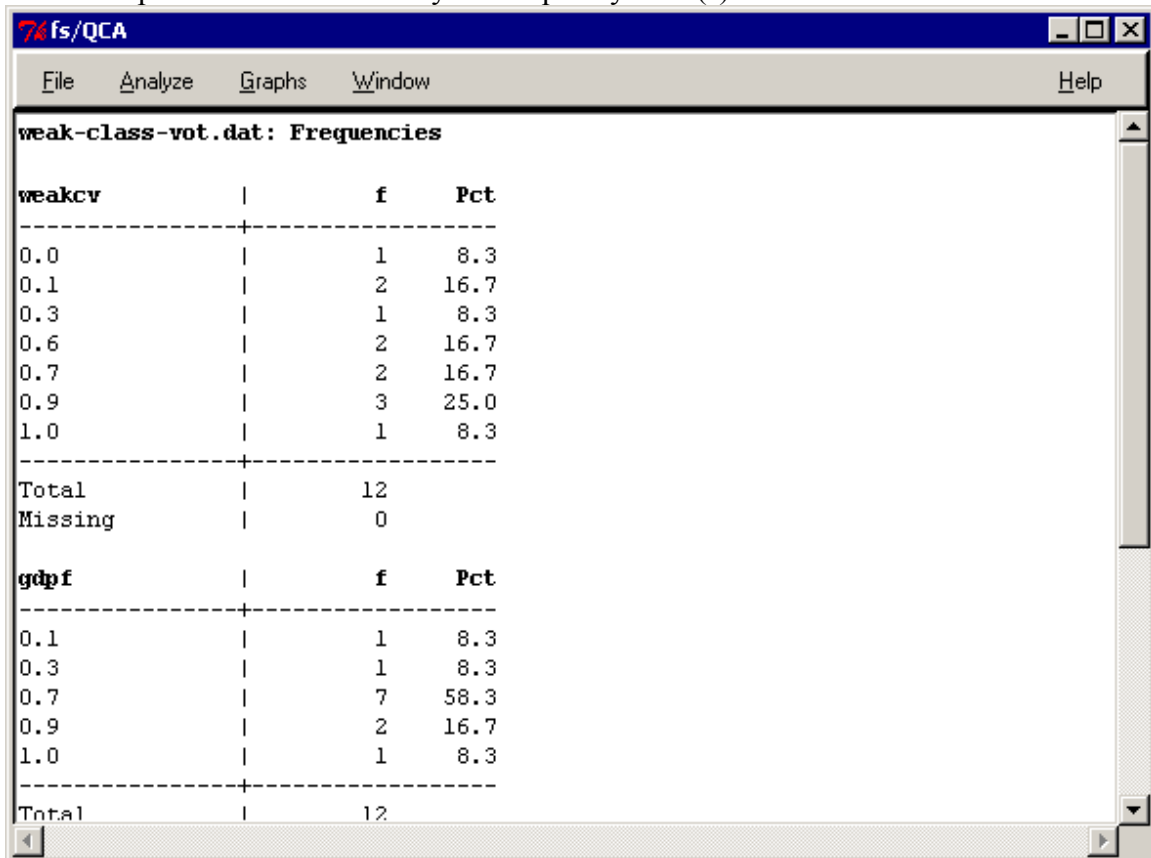
The Frequencies procedure provides frequency charts that are useful for describing many types of variables.

- In order to obtain frequencies, choose:

- Analyze
- Statistics
- Frequencies...

- Select one or more variables from the *Variables* column and transfer them into the *Frequencies* column. Click *Ok*.

- The output window will show your frequency table(s):



The screenshot shows the output window of the fs/QCA software. The window title is "fs/QCA" and it has a menu bar with "File", "Analyze", "Graphs", "Window", and "Help". The main content area displays the output for the "Frequencies" procedure on the file "weak-class-vot.dat".

The output consists of two frequency tables. The first table is for the variable "weakcv" and the second is for "gdpf". Both tables show the frequency (f) and percentage (Pct) for each value.

weakcv	f	Pct
0.0	1	8.3
0.1	2	16.7
0.3	1	8.3
0.6	2	16.7
0.7	2	16.7
0.9	3	25.0
1.0	1	8.3
Total	12	
Missing	0	

gdpf	f	Pct
0.1	1	8.3
0.3	1	8.3
0.7	7	58.3
0.9	2	16.7
1.0	1	8.3
Total	12	

- The first line of your output will read the file name and the procedure you have chosen (Frequencies). The columns in the frequency table indicate the following:
 1. The values of the selected variable(s)
 2. The frequency (f) of this value in the data set
 3. The valid percent (Pct) (the percentage disregarding the missing cases)
- The last two rows indicate the total number of cases (Total) analyzed and the number of missing cases (Missing).

B) Descriptives

The Descriptives procedure displays univariate summary statistics for several variables in a single table.

- In order to obtain descriptive statistics, choose:
 - Analyze
 - Statistics
 - Descriptives...
- Select one or more variables from the *Variables* column and transfer them into the *Descriptives* column. Click *Ok*.
- The output window will show your descriptive statistics:

weak-class-vot.dat: Descriptive Statistics

Variable	Mean	Std. Dev.	Minimum	Maximum	N Cases	Missing
gdpr	0.675	0.2384848	0.1	1	12	0
inqf	0.5833333	0.2881936	0.1	0.9	12	0
manf	0.4	0.3316625	0.1	0.9	12	0
unf	0.6083333	0.3094574	0.1	1	12	0
weakcv	0.5666667	0.3399346	0	1	12	0

➤ The first line of your output will read the file name and the procedure you have chosen (Descriptive Statistics). The columns in the descriptives table indicate the following:

1. The variable chosen (Variable)
2. The mean value (Mean)
3. The standard deviation (Std. Dev.)
4. The lowest value of the variable (Minimum)
5. The highest value of the variable (Maximum)
6. The number of cases (N Cases)
7. The number of missing cases (Missing)

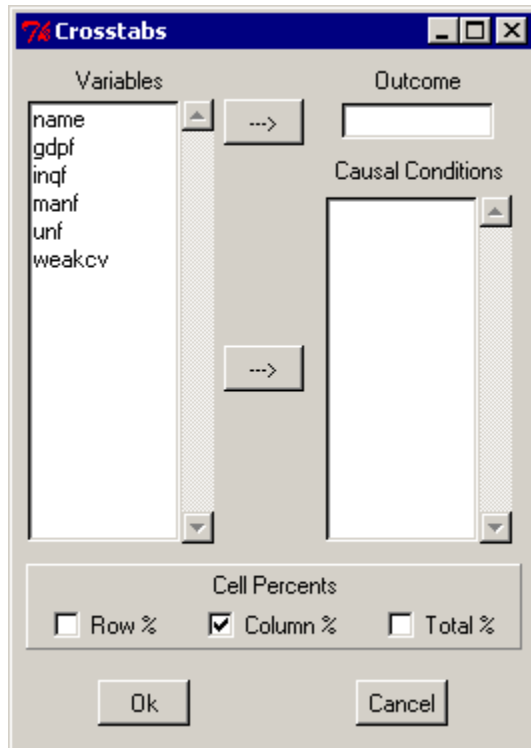
C) Crosstabulations

The Crosstabs procedure forms two-way tables. For multiway tables, statistics, and measures of association you will have to use standard statistical packages, like SPSS, for example.

➤ In order to obtain crosstabulations, choose:

- Analyze
- Statistics
- Crosstabs...

- The following window will open:



- Select one variable from the *Variables* column and transfer it into the *Dependent* column. This variable will appear in the rows of the table.
- Select one or more variables from the *Variables* column and transfer them into the *Independent* column. These variables will appear in the columns of the table.
- You can specify how you want the table percentaged:
 - percentage across rows
 - percentage across columns
 - or percentage the total number of cases represented in the table
- Click *Ok*.
- The output window will show your crosstabulation(s):

weak-class-vot.dat: Crosstabs

N	gdprf					
	0.1	0.3	0.7	0.9	1.0	
0.0	0	0	0	1	0	1
	0.0	0.0	0.0	50.0	0.0	
0.1	0	0	2	0	0	2
	0.0	0.0	28.6	0.0	0.0	
0.3	0	0	1	0	0	1
	0.0	0.0	14.3	0.0	0.0	
0.6	0	1	1	0	0	2
	0.0	100.0	14.3	0.0	0.0	
0.7	0	0	1	1	0	2
	0.0	0.0	14.3	50.0	0.0	
0.9	1	0	2	0	0	3

➤ The first line of your output will read the file name and the procedure you have chosen (Crosstabs). The first number in the cell indicates the number of cases in that cell (N). If requested, additional numbers represent (a) the number of cases across the row (Row %), (b) the number of cases across the column (Column %), and/or (c) The percentage of total number of cases in the table (Total %). See the first cell of the column header for the key.

➤ The last column of the table reads the number of cases across the rows.

➤ The last row of the table reads the number of cases across the columns.

➤ The last two rows in the output indicate the total number of cases (Total) analyzed and the number of missing cases (Missing).

D) Graphs

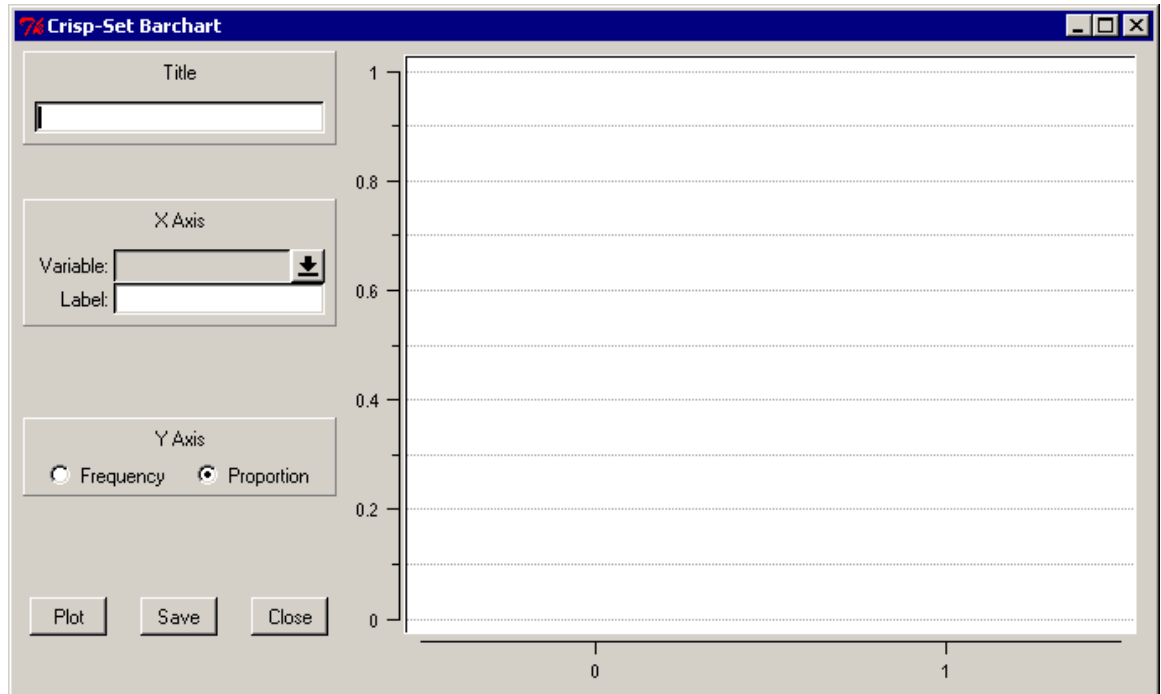
Barchart

The Barchart procedure displays simple bar charts with summaries for groups of cases.

➤ In order to produce a **Barchart**, choose:
Graphs

Crisp Barchart....

- The following window will open:



- Select a variable to define the categories on the X Axis shown in the chart. You can only choose **numeric data**. There is one bar for each value of the variable.
- The Barchart function only displays **values between 0 and 2**. If the values are not integers, the bars will overlap. Any instances where the variable is coded “don’t care” will be represented as having a value of 2.
- Optionally, you can enter a **label** for this variable and / or a **title** for the graph, which will be displayed on the graph.
- Decide whether you want the Y Axis to represent the **frequency of cases** in a category or the **proportion of cases** in a category.
- Once you have entered the specifications, click the **Plot** button. The graph will be displayed in the white field to the right.
- In order to **print** this graph, you have to **save** it as a PostScript file (*.ps). Click the **Save** button and choose ‘Postscript file’ as the saving option. Open it in a program which reads Postscript files, such as Ghostview (GSview). [You can download the Ghostviewer free from the Internet at: <http://www.cs.wisc.edu/~ghost/>].

Histogram

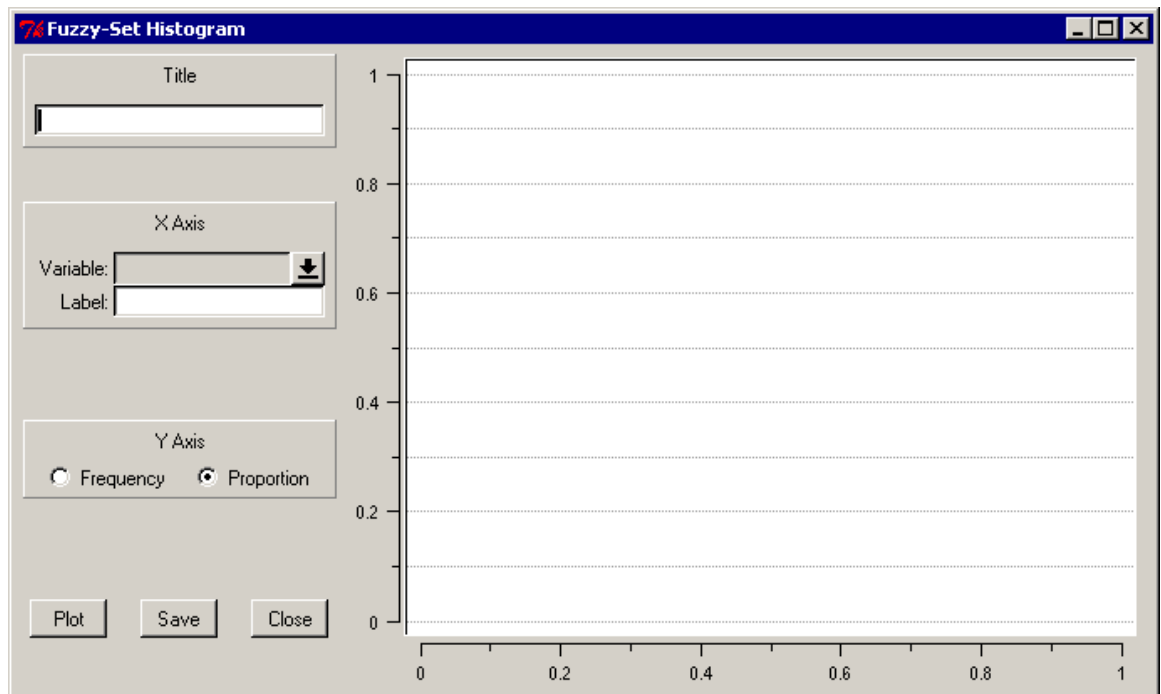
- In order to produce a **Histogram**, choose:

Graphs

Fuzzy

Histogram....

- The following window will open:



- Select a variable to define the categories on the X Axis shown in the chart. You can only choose **numeric data**. There is one bar for each value of the variable.
- Optionally, you can enter a **label** for this variable and / or a **title** for the graph, which will be displayed on the graph.
- Decide whether you want the Y Axis to represent the **frequency of cases** in a category **or the proportion of cases** in a category.
- Once you have entered the specifications, click the **Plot** button. The graph will be displayed in the white field to the right.
- In order to **print** this graph, you have to **save** it as a PostScript file (*.ps). Click the **Save** button and choose 'Postscript file' as the saving option. Open it in a program which reads Postscript files (e.g. Ghostview (GSview) or some versions of MS Word. [You can download the Ghostviewer free from the Internet at: <http://www.cs.wisc.edu/~ghost/>].

XY Plot

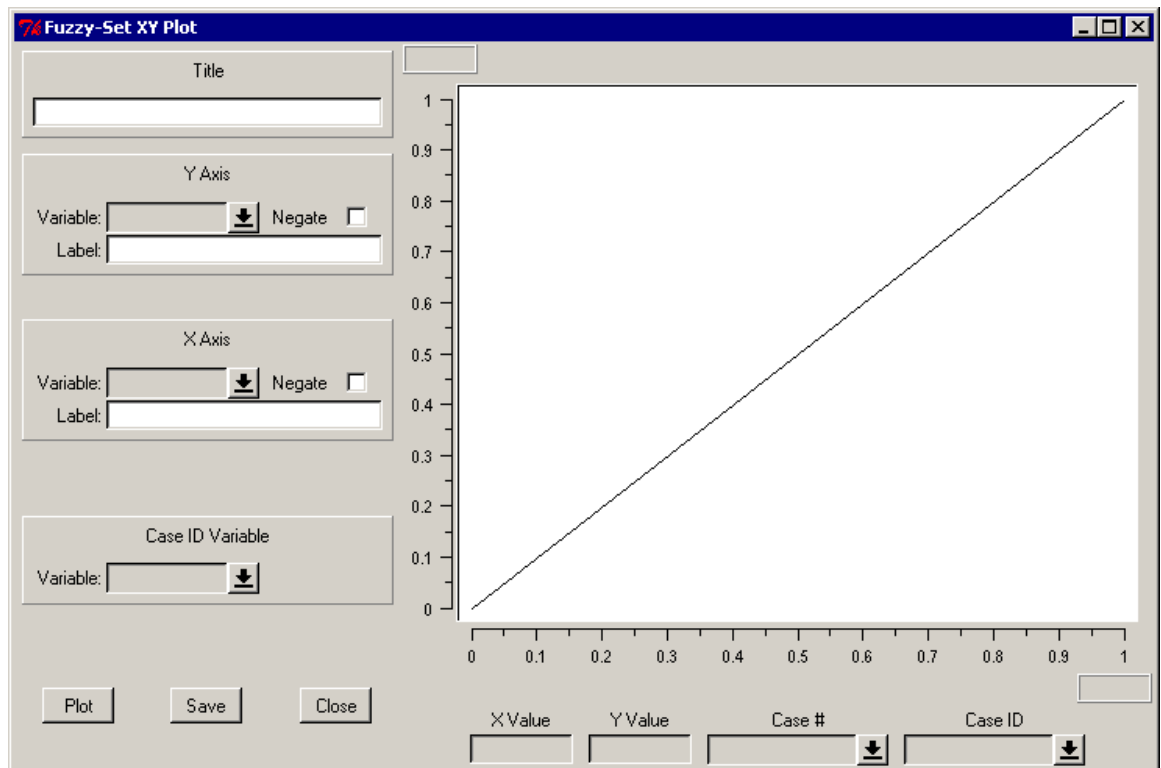
- In order to produce a **XY Plot**, choose:

Graphs

Fuzzy

XY Plot....

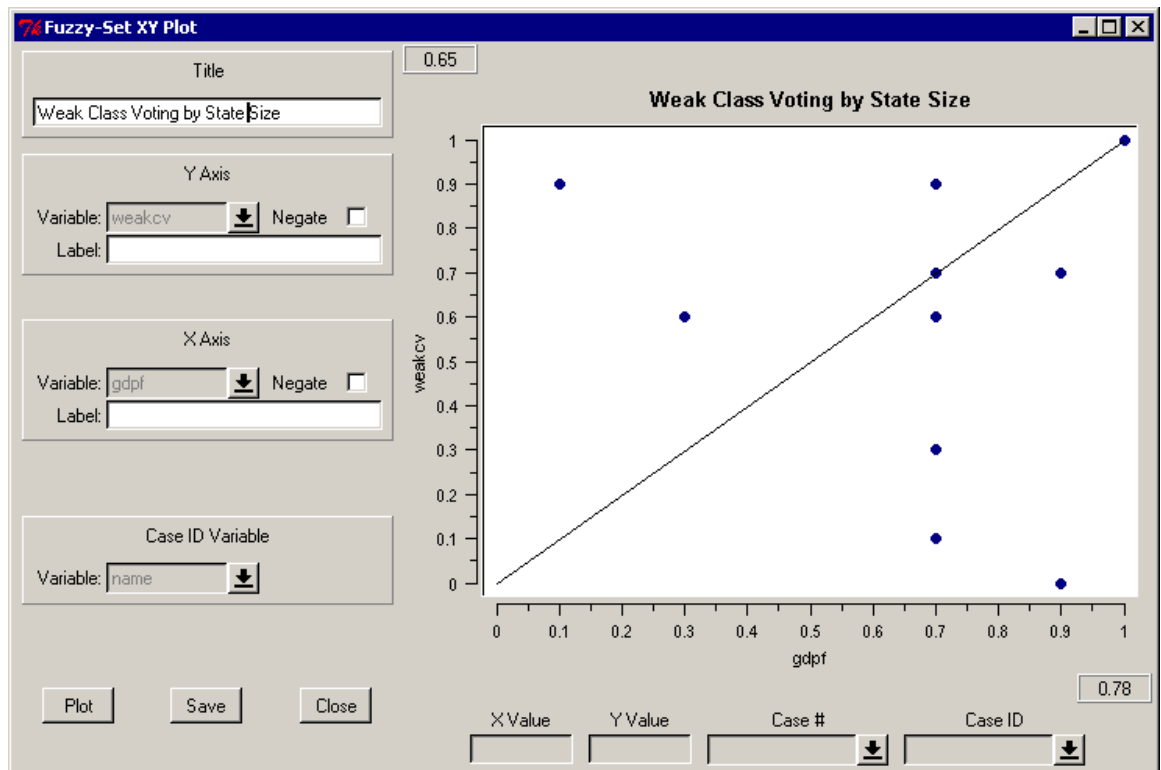
- The following window will open:



- Select a variable to define the values on the **X Axis** shown in the chart. You can only use numeric variables that ranges from 0 to 1 (i.e., fuzzy sets).
- Select a variable to define the values on the **Y Axis** shown in the chart. Again, you can only use numeric variables that range from 0 to 1.
- Optionally, you can enter a **title** for the graph, which will be displayed on the graph when it is saved.
- You can also add more information by choosing a **Case ID Variable**. This variable will not be represented in the graph, but you can determine its value by clicking on a particular point in the graph—the Case ID value will be displayed in the field in the lower right corner labeled Case ID. For example, the Case ID variable could be a string variable with the names of the countries in the data set. Once plotted, it is possible to click on any point in the plot. The Case ID field beneath the graph will

show the country that is represented by this point. Clicking on the arrow next to the field will show other countries that share the same data values in the plot.

- The row beneath the graph will also indicate the x and y values of particular cases and the case numbers that correspond to each point.
- Once you have entered the specifications, click the **Plot** button and the plot will be displayed:



- The numbers in the boxes at two corners of the graph show set-theoretic consistency scores. The lower right box shows the degree to which the data plotted are consistent with $Y \leq X$ (Y is a subset of X). The upper left box shows the degree to which the data plotted are consistent with $X \leq Y$ (X is a subset of Y). If one of these two numbers indicates high consistency, the other can be interpreted as a coverage score. For example, if the number in the upper left corner is .91 and the number in the lower right corner is .63, these calculations indicate that the data are largely consistent with the argument that X is a subset of Y and its coverage of Y is 63%, That is, X accounts for 63% of the sum of the memberships in Y .
- You can negate variables in the graph by clicking on the negate option next to the variable name. This feature will subtract the fuzzy-set value of this variable from 1. Example: Pressure = .33; negation of Pressure = .67. [Same as ‘~’ and ‘fuzzynot(x)’]
- In order to **print** this graph, you have to **save** it as a PostScript file (*.ps). Click the *Save* button and choose ‘PostScript file’ as the saving option. Open it in a program

which reads PostScript files, such as Ghostview (GSview). [You can download the Ghostviewer free from the Internet at: <http://www.cs.wisc.edu/~ghost/>].

4. CRISP-SET ANALYSIS

This part of the manual refers to the analysis of dichotomous social data reflecting the memberships of cases in conventional, crisp sets. In-depth discussions of this method can be found in *The Comparative Method* (Ragin 1987) and in chapter 5 of *Fuzzy-Set Social Science* (Ragin 2000). The data analytic strategy used here is known as *qualitative comparative analysis*, or QCA. QCA is based on Boolean algebra, where a case is either in or out of a set, and QCA uses binary-coded data, with 1 indicating membership and 0 indicating nonmembership. QCA using conventional, crisp sets is also known as csQCA.

A) Basic Concepts

An explicit algebraic basis for qualitative comparison exists in Boolean algebra. Also known as the algebra of logic and as the algebra of sets, Boolean algebra was developed in the mid-nineteenth century by George Boole. The Boolean principles used in qualitative comparative analysis are quite simple. Seven aspects of Boolean algebra are essential for the algorithms and are presented here in rough sequence here, with more difficult concepts following simpler concepts.

1) Use of binary data

There are two conditions or states in Boolean algebra: true (or present) and false (or absent). These two states are represented in base 2: 1 indicates presence; 0 indicates absence. The typical Boolean-based comparative analysis addresses the presence/absence conditions under which a certain outcome is obtained (that is, is true). Thus, in a Boolean analysis of social data all variables, independent and dependent, must be nominal-scale measures. Interval-scale measures are transformed into multi-category nominal-scale measures. Nominal-scale measures with more than two categories are represented with several binary variables.

2) Boolean negation

In Boolean logic, negation switches membership scores from 1 to 0 and from 0 to 1. The negation of the crisp set of males, for example, is the crisp set of not males. If a case has a Boolean score of 1 in the set of males, then it has a Boolean score of 0 in the set of not males.

3) Use of truth table to represent data

In order to use Boolean algebra as a technique of qualitative comparison, it is necessary to reconstruct a raw data matrix as a truth table. The idea behind a truth table is simple. Once the data have been recoded into nominal-scale variables and represented in binary form (as 1's and 0's), it is necessary only to sort the data into their different combinations of values on the independent variables. Each logical combination of values on the independent variables is represented as one row of the truth table. Once this part of the truth table is constructed, each row is assigned an output value (a score of 1 or 0 on the dependent variable) based on the scores of the cases which share that combination of

input values (that combination of scores on the independent variables). Thus, both the different combinations of input values (independent variables) and their associated output values (the dependent variable) are summarized in a truth table.

Truth tables have as many rows as there are logically possible combinations of values on the causal variables. If there are three binary independent variables, for example, the truth table will contain $2^3 = 8$ rows, one for each logically possible combination of three presence/absence independent variables. The truth table for a moderate-sized data set with three binary independent variables and one binary dependent variable (with 1 = present and 0 = absent) is shown in Table 1. Technically, there is no reason to include the frequency of each combination as part of the truth table. These values are included in the examples to remind the reader that each row is not a single case but a summary of all the cases with a certain combination of input values. In this respect, a row of a truth table is like a cell from a multiway cross-classification of several categorical independent variables.

Table 1: Hypothetical Truth Table Showing Three Causes of Regime Failure

Condition			Regime Failure	Number of Instances
<i>conflict</i>	<i>death</i>	<i>cia</i>	<i>failure</i>	
0	0	0	0	9
1	0	0	1	2
0	1	0	1	3
0	0	1	1	1
1	1	0	1	2
1	0	1	1	1
0	1	1	1	1
1	1	1	1	3

conflict = Conflict between older and younger military officers
death = Death of a powerful dictator
cia = CIA dissatisfaction with the regime

4) Groupings

Just as it is possible to calculate the logically possible number of combinations (2^k), it is also possible to calculate the number of logically possible groupings. The formula is $3^k - 1$, where k again is the number of attributes ($3^3 - 1 = 26$). Table 2 shows the 26 logically possible groupings of the three dichotomies presented in Table 1. Using the formula just described, the 26 possible groupings are formed as follows: 8 involve combinations of three attributes, 12 involve combinations of two attributes, and six involve single attributes.

Table 2: Groupings Using Three Dichotomies (from Table 1)

Initial Configuration	Groupings involving	Groupings evolving a single
-----------------------	---------------------	-----------------------------

(8 combinations of three aspects)	combinations of two aspects (12)	aspect (6)
<i>conflict • death • cia</i> <i>conflict • death • ~cia</i> <i>conflict • ~death • cia</i> <i>conflict • ~death • ~cia</i> <i>~conflict • death • cia</i> <i>~conflict • ~death • cia</i> <i>conflict • death • ~cia</i> <i>~conflict • ~death • ~cia</i>	<i>conflict • death</i> <i>conflict • ~death</i> <i>~conflict • ~death</i> <i>~conflict • death</i> <i>conflict • cia</i> <i>conflict • ~cia</i> <i>~conflict • cia</i> <i>~conflict • ~cia</i> <i>death • cia</i> <i>death • ~cia</i> <i>~death • cia</i> <i>~death • ~cia</i>	<i>conflict</i> <i>~conflict</i> <i>death</i> <i>~death</i> <i>cia</i> <i>~cia</i>

5) Boolean Addition

In Boolean algebra, if $A + B = Z$, and $A = 1$ and $B = 1$, then $Z = 1$. In other words, $1 + 1 = 1$. The basic idea in Boolean addition is that if any of the additive terms is satisfied (present), then the outcome is true (occurs). Addition in Boolean algebra is equivalent to the logical operator OR. (In this discussion uppercase OR is used to indicate logical OR.) Thus, the above statement $A + B = Z$ becomes: if A equals 1 OR B equals 1, then Z equals 1.

The best way to think of this principle is in logical terms, not arithmetically. For example, there might be several things a person could do to lose his or her job. It does not matter how many of these things the person does. If the employee does any one (or all) of them, he or she will be fired. Doing two of them will not cause one employee to be more fired than another employee who does only one of them. Fired is fired, a truly qualitative state. This example succinctly illustrates the nature of Boolean addition: satisfy any one of the additive conditions and the expected outcome follows. This aspect of Boolean addition is very useful in social scientific analysis, especially qualitative comparison, although its value is not generally recognized.

Consider the collapse of military regimes. Assume that there are three general conditions that cause military regimes to fall: sharp conflict between older and younger military officers (*conflict*), death of a powerful dictator (*death*), or CIA dissatisfaction with the regime (*cia*). Any one of these three conditions may be sufficient to prompt a collapse. The truth table for a number of such regimes in different countries is shown in Table 1 (with 1 = present and 0 = absent). Each combination of causes produces either regime failure or an absence of regime failure—there are no contradictory rows.

The "simplified" Boolean equation

$$failure = conflict + death + cia$$

expresses the relation between the three conditions and regime failure simply and elegantly for both negative and positive instances. Simply stated: if any one (or any two or all three) of these conditions obtains, then the regime will fall.

6) Boolean Multiplication

Boolean multiplication differs substantially from normal multiplication. Boolean multiplication is relevant because the typical social science application of Boolean algebra concerns the process of simplifying expressions known as "sums of products." A product is a particular combination of causal conditions. The data on collapsed military regimes from Table 1 can be represented in "primitive" (that is, unreduced) sums-of-products form as follows:

$$\begin{aligned}
 failure = & \textit{conflict} \bullet \sim\textit{death} \bullet \sim\textit{cia} + \\
 & \sim\textit{conflict} \bullet \textit{death} \bullet \sim\textit{cia} + \\
 & \sim\textit{conflict} \bullet \sim\textit{death} \bullet \textit{cia} + \\
 & \textit{conflict} \bullet \textit{death} \bullet \sim\textit{cia} + \\
 & \textit{conflict} \bullet \sim\textit{death} \bullet \textit{cia} + \\
 & \sim\textit{conflict} \bullet \textit{death} \bullet \textit{cia} + \\
 & \textit{conflict} \bullet \textit{death} \bullet \textit{cia}
 \end{aligned}$$

Each of the seven terms represents a combination of causal conditions found in at least one instance of regime failure. The different terms are products because they represent intersections of conditions (conjunctures of causes and absences of causes). The equation shows the different primitive combinations of conditions that are linked to the collapse of military regimes.

Boolean multiplication, like Boolean addition, is not arithmetic. The expression *conflict* • *~death* • *~cia* does not mean that the value of *conflict* (1) is multiplied by the value of *death* (0) and by the value of *cia* (0) to produce a result value of 0. It means simply that a presence of *conflict* is combined with an absence of *death* and an absence of *cia*. The total situation, *failure* = *conflict* • *~death* • *~cia*, occurs in the data twice. This conjunctural character of Boolean multiplication shapes the interpretation of the primitive sums-of-products equation presented above: *failure* (regime failure) occurs if any of seven combinations of three causes is obtained. In Boolean algebra addition indicates logical OR and multiplication indicates logical AND. The three causes are ANDed together in different ways to indicate different empirical configurations. These intersections are ORed together to form an unreduced, sums-of-products equation describing the different combinations of the three causes linked to regime failure.

7) Combinatorial Logic

Boolean analysis is combinatorial by design. In the analysis of regime failures presented above, it appears from casual inspection of only the first four rows of the truth table (Table 1) that if any one of the three causes is present, then the regime will collapse.

While it is tempting to take this shortcut, the route taken by Boolean analysis is much more exacting of the data. This is because the absence of a cause has the same logical status as the presence of a cause in Boolean analysis. As noted above, Boolean multiplication indicates that presence and absence conditions are combined, that they intersect.

Consider the second row of the truth table (Table 1), which describes the two instances of military regime failure linked to causal configuration $conflict \bullet \sim death \bullet \sim cia$. Simple inspection suggests that in this case *failure* (regime failure) resulted from the first cause, *conflict*. But notice that if the investigator had information on only this row of the truth table, and not on any of the other instances of regime failure, he or she might conclude that *conflict* causes *failure* only if causes *death* and *cia* are absent. This is what the $conflict \bullet \sim death \bullet \sim cia$ combination indicates. This row by itself does not indicate whether *conflict* would cause *failure* in the presence of *death* or *cia* or both. All the researcher knows from these two instances of $conflict \bullet \sim death \bullet \sim cia$ is that for *conflict* to cause *failure*, it may be necessary for the other conditions (*death* and *cia*) to be absent. From a Boolean perspective, it is entirely plausible that in the presence of one or both of these other conditions (say, configuration $conflict \bullet \sim death \bullet cia$), *failure* may not result. To return to the original designations, it may be that in the presence of CIA meddling (*cia*), conflict between junior and senior officers (*conflict*) will dissipate as the two factions unite to oppose the attempt by outsiders to dictate events.

To push this argument further, assume the investigator had knowledge of only the first four rows of the truth table. The data would support the idea that the presence of any one of the three conditions causes *failure*, but again the data might indicate that *conflict* causes *failure* only when *death* and *cia* are absent ($conflict \bullet \sim death \bullet \sim cia$); *death* causes *failure* only when *conflict* and *cia* are absent ($\sim conflict \bullet death \bullet \sim cia$), and so on. A strict application of combinatorial logic requires that these limitations be placed on conclusions drawn from a limited variety of cases.

This feature of combinatorial logic is consistent with the idea that cases, especially their causally relevant features, should be viewed holistically. The holistic character of the Boolean approach is consistent with the orientation of qualitative scholars in comparative social science who examine different causes in context. When the second row of the truth table (Table 1) is examined, it is not interpreted as instances of *failure* caused by *conflict*, but as instances of *failure* caused by $conflict \bullet \sim death \bullet \sim cia$. Thus, in Boolean-based qualitative comparison, causes are not viewed in isolation but always within the context of the presence and absence of other causally relevant conditions.

Minimization

The restrictive character of combinatorial logic seems to indicate that the Boolean approach simply compounds complexity on top of complexity. This is not the case. There are simple and straightforward rules for simplifying complexity—for reducing primitive expressions and formulating more succinct Boolean statements. The most fundamental of these rules is:

If two Boolean expressions differ in only one causal condition yet produce the same outcome, then the causal condition that distinguishes the two expressions can be considered irrelevant and can be removed to create a simpler, combined expression.

Essentially this minimization rule allows the investigator to take two Boolean expressions that differ in only one term and produce a combined expression. For example, *conflict • ~death • ~cia* and *conflict • death • ~cia*, which both produce outcome *failure*, differ only in *death*; all other elements are identical. The minimization rule stated above allows the replacement of these two terms with a single, simpler expression: *conflict • ~cia*. In other words, the comparison of these two rows, *conflict • ~death • ~cia* and *conflict • death • ~cia*, as wholes indicates that in instances of *conflict • ~cia*, the value of *death* is irrelevant. Cause *death* may be either present or absent; *failure* will still occur.

The logic of this simple data reduction parallels the logic of experimental design. Only one causal condition, *death*, varies and no difference in outcome is detected (because both *conflict • ~death • ~cia* and *conflict • death • ~cia* are instances of *failure*). According to the logic of experimental design, *death* is irrelevant to *failure* in the presence of *conflict • ~cia* (that is, holding these two conditions constant). Thus, the process of Boolean minimization mimics the logic of experimental design. It is a straightforward operationalization of the logic of the ideal social scientific comparison.

This process of logical minimization is conducted in a bottom-up fashion until no further stepwise reduction of Boolean expressions is possible. Consider again the data on military regime failures presented above. Each of the rows with one cause present and two absent can be combined with rows with two causes present and one absent because all these rows have the same outcome (*failure*) and each pair differs in only one causal condition:

conflict • ~death • ~cia combines with *conflict • death • ~cia* to produce *conflict • ~cia*.
conflict • ~death • ~cia combines with *conflict • ~death • cia* to produce *conflict • ~death*.
~conflict • death • ~cia combines with *conflict • death • ~cia* to produce *death • ~cia*.
~conflict • death • ~cia combines with *~conflict • death • cia* to produce *~conflict • death*.
~conflict • ~death • cia combines with *conflict • ~death • cia* to produce *~death • cia*.
~conflict • ~death • cia combines with *~conflict • death • cia* to produce *~conflict • cia*.

Similarly, each of the rows with two causes present and one absent can be combined with the row with all three present:

conflict • death • ~cia combines with *conflict • death • cia* to produce *conflict • death*.
conflict • ~death • cia combines with *conflict • death • cia* to produce *conflict • cia*.
~conflict • death • cia combines with *conflict • death • cia* to produce *death • cia*.

Further reduction is possible. Note that the reduced terms produced in the first round can be combined with the reduced terms produced in the second round to produce even simpler expressions:

conflict • ~death combines with *conflict • death* to produce *conflict*.

conflict • ~cia combines with *conflict • cia* to produce *conflict*.

~conflict • death combines with *conflict • death* to produce *death*.

death • ~cia combines with *death • cia* to produce *death*.

~conflict • cia combines with *conflict • cia* to produce *cia*.

~death • cia combines with *death • cia* to produce *cia*.

Although tedious, this simple process of minimization produces the final, reduced Boolean equation:

$$failure = conflict + death + cia$$

True enough, this was obvious from simple inspection of the entire truth table, but the problem presented was chosen for its simplicity. The example directly illustrates key features of Boolean minimization. It is bottom-up (that is, inductively oriented). It seeks to identify ever wider sets of conditions (that is, simpler combinations of causal conditions) for which an outcome is true. And it is experiment-like in its focus on pairs of configurations differing in only one cause.

1) Use of “prime implicants”

A further Boolean concept that needs to be introduced is the concept of implication. A Boolean expression is said to imply another if the membership of the second term is a subset of the membership of the first. For example, *a* implies *a • ~b • ~c* because *a* embraces all the members of *a • ~b • ~c* (that is, *a • ~b • ~c* is a subset of *a*). This concept is best understood by example. If *a* indicates economically dependent countries, *b* indicates the presence of heavy industry, and *c* indicates centrally coordinated economies, *a* embraces all dependent countries while *a • ~b • ~c* embraces all dependent countries that lack both centrally coordinated economies and heavy industry. Clearly the membership of *a • ~b • ~c* is included in the membership of *a*. Thus, *a* implies *a • ~b • ~c*.

The concept of implication, while obvious, provides an important tool for minimizing primitive sums-of-products expressions. Consider the hypothetical truth table shown in Table 3, which summarizes data on three causal conditions thought to affect the success of strikes already in progress (*success*): a booming market for the product produced by the strikers (*market*), the threat of sympathy strikes by workers in associated industries (*threat*), and the existence of a large strike fund (*fund*).

The Boolean equation for *success* (successful strikes) showing unreduced (primitive) Boolean expressions is

$$\text{success} = \text{market} \bullet \sim\text{threat} \bullet \text{fund} + \sim\text{market} \bullet \text{threat} \bullet \sim\text{fund} + \text{market} \bullet \text{threat} \bullet \sim\text{fund} + \text{market} \bullet \text{threat} \bullet \text{fund}$$

Table 3: Hypothetical Truth Table Showing Three Causes of Successful Strikes

<i>Market</i>	Condition		Success	Frequency
	<i>Threat</i>	<i>fund</i>	<i>success</i>	
1	0	1	1	6
0	1	0	1	5
1	1	0	1	2
1	1	1	1	3
1	0	0	0	9
0	0	1	0	6
0	1	1	0	3
0	0	0	0	4

The first step in the Boolean analysis of these data is to attempt to combine as many compatible rows of the truth table as possible. (Note that this part of the minimization process uses rows with an output value of 1—strike succeeded.) This first phase of the minimization of the truth table produces the following partially minimized Boolean equation, which in effect turns a primitive Boolean equation with four three-variable terms into an equation with three two-variable terms:

market • *threat* • *fund* combines with *market* • *~threat* • *fund* to produce *market* • *fund*.
market • *threat* • *fund* combines with *market* • *threat* • *~fund* to produce *market* • *threat*.
market • *threat* • *~fund* combines with *~market* • *threat* • *~fund* to produce *threat* • *~fund*.

$$\text{success} = \text{market} \bullet \text{fund} + \text{market} \bullet \text{threat} + \text{threat} \bullet \sim\text{fund}$$

Product terms such as those in the preceding equation which are produced using this simple minimization rule—combine rows that differ on only one cause if they have the same output values—are called prime implicants. Usually, each prime implicant covers (that is, implies) several primitive expressions in the truth table. In the partially minimized equation given above, for example, prime implicant *market* • *fund* covers two primitive Boolean expressions listed in the truth table: *market* • *threat* • *fund* and *market* • *~threat* • *fund*.

This partially reduced Boolean expression illustrates a common finding in Boolean analysis: often there are more reduced expressions (prime implicants) than are needed to cover all the original primitive expressions. Prime implicant *market* • *threat* implies primitive terms *market* • *threat* • *fund* and *market* • *threat* • *~fund*, for example, yet these two primitive terms are also covered by *market* • *fund* and *threat* • *~fund*, respectively. Thus, *market* • *threat* may be redundant from a purely logical point of view; it may not

be an essential prime implicant. In order to determine which prime implicants are logically essential, a minimization device known as a prime implicant chart is used. Minimization of the prime implicant chart is an optional, second phase of Boolean minimization.

Briefly stated, the goal of this second phase of the minimization process is to "cover" as many of the primitive Boolean expressions as possible with a logically minimal number of prime implicants. This objective derives from a straightforward desire for non-redundancy. The prime implicant chart maps the links between prime implicants and primitive expressions. The prime implicant chart describing these links in the data on strike outcomes is presented in Table 4. Simple inspection indicates that the smallest number of prime implicants needed to cover all of the original primitive expressions is two. (For very complex prime implicant charts, sophisticated computer algorithms are needed; see Mendelson 1970, Roth 1975, and McDermott 1985.) Prime implicants $market \bullet fund$ and $threat \bullet \sim fund$ cover all four primitive Boolean expressions. Analysis of the prime implicant chart, therefore, leads to the final reduced Boolean expression containing only the logically essential prime implicants:

$$success = market \bullet fund + threat \bullet \sim fund$$

This equation states simply that successful strikes occur when there is a booming market for the product produced by the workers AND a large strike fund ($market \bullet fund$) or when there is the threat of sympathy strikes by workers in associated industries combined with a low strike fund ($threat \bullet \sim fund$). (Perhaps the threat of sympathy strikes is taken seriously only when the striking workers badly need the support of other workers.)

Table 4: Prime Implicant Chart Showing Coverage of Original Terms by Prime Implicants (Hypothetical Strike Data)

		Primitive Expressions			
		$market \bullet threat \bullet fund$	$market \bullet \sim threat \bullet fund$	$market \bullet threat \bullet \sim fund$	$\sim market \bullet threat \bullet \sim fund$
Prime Implicants	$market \bullet fund$	X	X		
	$market \bullet threat$	X		X	
	$threat \bullet \sim fund$			X	X

These simple procedures allow the investigator to derive a logically minimal equation describing the different combinations of conditions associated with an outcome. The final, reduced equation shows the two (logically minimal) combinations of conditions that cause successful strikes and thus provides an explicit statement of multiple conjunctural causation.

2) Use of De Morgan's Law

The application of De Morgan's Law is straightforward. Consider the solution to the hypothetical analysis of successful strikes presented above: $success = market \bullet fund + threat \bullet \sim fund$. Elements that are coded present in the reduced equation (say, $market$ in the term $market \bullet fund$) are recoded to absent, and elements that are coded absent (say, $\sim fund$ in the term $threat \bullet \sim fund$) are recoded to present. Next, logical AND is recoded to logical OR, and logical OR is recoded to logical AND. Applying these two rules,

$$success = market \bullet fund + threat \bullet \sim fund$$

becomes:

$$\begin{aligned} \sim success &= (\sim market + \sim fund) \bullet (\sim threat + fund) \\ &= \sim market \bullet \sim threat + \sim market \bullet fund + \sim threat \bullet \sim fund \end{aligned}$$

According to this equation, strikes fail when (1) the market for the relevant product is not booming AND there is no serious threat of sympathy strikes, (2) the market for a product is not booming AND there is a large strike fund, OR (3) there is no threat of sympathy strikes AND only a small strike fund. (The combination $\sim market \bullet fund$ —nonbooming market and large strike fund, which seems contradictory—may suggest an economic downturn after a period of stability. In this situation a shutdown might be welcomed by management.)

De Morgan's Law produces the exact negation of a given logical equation. If there are "remainder" combinations in the truth table and they are used as "don't cares," then the results of the application of De Morgan Law will yield a logical statement that is **not** the same as the analysis of the absence of the outcome. Likewise, if the remainders are defined as "false" in the initial analysis, then the application of De Morgan's Law to the this solution (of positive cases) will yield a logical statement that embraces not only the negative cases, but also the remainders.

3) Necessary and Sufficient Causes

A cause is defined as necessary if it must be present for an outcome to occur. A cause is defined as sufficient if by itself it can produce a certain outcome. This distinction is meaningful only in the context of theoretical perspectives. No cause is necessary, for example, independent of a theory that specifies it as a relevant cause. Neither necessity nor sufficiency exists independently of theories that propose causes.

Necessity and sufficiency are usually considered jointly because all combinations of the two are meaningful. A cause is both necessary and sufficient if it is the only cause that produces an outcome and it is singular (that is, not a combination of causes). A cause is sufficient but not necessary if it is capable of producing the outcome but is not the only cause with this capability. A cause is necessary but not sufficient if it is capable of producing an outcome in combination with other causes and appears in all such combinations. Finally, a cause is neither necessary nor sufficient if it appears only in a

subset of the combinations of conditions that produce an outcome. In all, there are four categories of causes (formed from the cross-tabulation of the presence/absence of sufficiency against the presence/absence of necessity).

The typical application of QCA (crisp or fuzzy) results in a logical statement describing combinations of conditions that are sufficient for the outcome. The listed combinations may or may not be exhaustive, that is, they may not explain *all* instances of the outcome. Usually, it is best to examine necessity separately, before the analysis of sufficient combinations of conditions. If a necessary condition is identified and it makes sense as a necessary condition, it can be omitted from the truth table analysis (the analysis of sufficient combinations of conditions). It still plays an important part in the presentation of the results, however, and may be considered a component of any combination of conditions identified as sufficient in the truth table analysis.

B) Data

The following window shows a sample crisp-set data sheet:

The screenshot shows a software window titled "FS/QCA Data Sheet" with a menu bar containing "File", "Variables", "Cases", "Analyze", and "Graphs". The main area displays a data table with the following structure:

Case	caseid	wealthy	urban	literate	industrial	unstable	survived
1	AUS	1	0	1	1	1	0
2	BEL	1	1	1	1	0	1
3	CZE	0	1	1	1	0	1
4	EST	0	0	1	0	0	0
5	FIN	0	0	1	0	0	1
6	FRA	1	0	1	1	0	1
7	GER	1	1	1	1	1	0
8	GRE	0	0	0	0	1	0
9	HUN	0	0	1	0	1	0
10	IRE	1	0	1	0	0	1
11	ITA	0	0	0	0	0	0
12	NET	1	1	1	1	0	1
13	POL	0	0	1	0	1	0
14	POR	0	0	0	0	1	0
15	ROM	0	0	0	0	0	0
16	SPA	0	0	0	0	1	0
17	SWE	1	0	1	1	0	1
18	UK	1	1	1	1	0	1

The window also shows a file path at the bottom: "File: benoit.csv".

caseid	abbreviated country name
wealthy	high GDP/cap versus not
urban	highly urban versus not
literate	high level of literacy versus not
industrial	high percentage of industrial workers versus not
unstable	government instability versus not
survived	democracy survived during interwar period versus not

[The example in this section is from Rihoux and Ragin (2008), *Configurational Comparative Analysis*.]

C) Analysis

The current version of the fsQCA software (as of this writing, version 2.0, September 2008) contains two methods of conducting crisp-set analysis: the “Truth Table Algorithm” and “Quine (QCA 3.0).” Both methods make use of the Quine-McCluskey algorithm and can produce the same results; the difference between the two methods regards how the truth table is constructed and managed. **We recommend using the Truth Table Algorithm method as it is more straightforward.** Both are described here; the preferred procedure, the Truth Table Algorithm is described first.

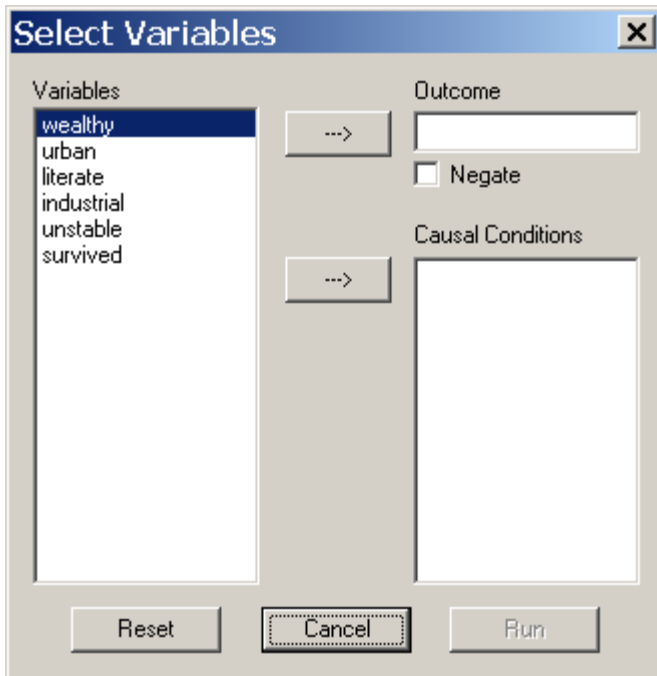
Truth Table Algorithm

Two important tasks structure the application of the crisp-set truth table algorithm: (1) The assessment of the *distribution of cases* across different logically possible combinations of causal conditions. And (2) the assessment of the *consistency of the evidence* for each causal combination with the argument that the cases with this combination of conditions constitute a subset of the cases with the outcome.

The truth table algorithm involves a two-step analytic procedure. The first step consists of creating a truth table spreadsheet from the raw data, which primarily involves specifying the outcome and causal conditions to include in the analysis. The second step consists of preparing the truth table spreadsheet for analysis, by selecting both a frequency threshold and a consistency threshold. These steps must be performed in conjunction and both must be performed for each separate analysis.

- In order to create the truth table spreadsheet, choose:
 - Analyze
 - Fuzzy Sets
 - Truth Table Algorithm.....

The following window will open, listing the variables in your file:



- Identify the case aspect you want to explain and transfer it into the *Outcome* field.
- Select a preliminary list of causal conditions by moving them over to the *Causal Conditions* field.
- Click on the Run button and the following window containing the full truth table will appear:

wealthy	literate	industrial	unstable	number	survived	consist	pre	product
1	1	1	0	5 (27%)		1.000000	1.000000	1.000000
0	0	0	1	3 (44%)		0.000000	0.000000	0.000000
0	0	0	0	2 (55%)		0.000000	0.000000	0.000000
0	1	0	0	2 (66%)		0.500000	0.500000	0.250000
0	1	0	1	2 (77%)		0.000000	0.000000	0.000000
1	1	1	1	2 (88%)		0.000000	0.000000	0.000000
0	1	1	0	1 (94%)		1.000000	1.000000	1.000000
1	1	0	0	1 (100%)		1.000000	1.000000	1.000000
0	0	1	0	0 (100%)				
0	0	1	1	0 (100%)				
0	1	1	1	0 (100%)				
1	0	0	0	0 (100%)				
1	0	0	1	0 (100%)				
1	0	1	0	0 (100%)				
1	0	1	1	0 (100%)				
1	1	0	1	0 (100%)				
1	1	0	1	0 (100%)				

➤ The truth table will have 2^k rows (where k represents the number of causal conditions), reflecting all possible combinations of causal conditions. The 1s and 0s represent full membership and zero membership for each condition, respectively. For each row, a value for each of the following variables is created:

<i>number</i>	the number of cases displaying the combination of conditions
<i>consist</i>	the proportion of cases in each truth table row that display the outcome.
<i>pre</i>	an alternative measure of consistency (developed for fuzzy sets) based on a quasi proportional reduction in error calculation. In crisp set analyses this will be equal to <i>consist</i> .
<i>product</i>	the product of <i>consist</i> and <i>pre</i> . For crisp sets this is simply degree of consistency squared.

Note that the column labeled as the outcome (*survived* in this example) is blank. It is up to the investigator to determine the outcome for each configuration using the following procedure.

➤ The researcher must begin by developing a rule for classifying some combinations (rows) as relevant and others as irrelevant, based on their frequency. This is accomplished by selecting a frequency threshold based on the number of cases in each row, shown in *number* column. When the total number of cases in an analysis is relatively small, the frequency threshold should be 1 or 2. When the total N is large, however, a more substantial threshold should be used. It is very important to examine the distribution of cases across causal combinations.

➤ Configurations (rows) can be sorted by their frequency by clicking any cell in the *number* column and choosing

Sort
Descending

➤ After sorting rows and selecting a frequency threshold, delete all rows that do not meet the threshold. If the cases have been sorted in a descending order according to *number*, click on the first case that falls below the threshold then select

Edit
Delete current row to last row

If cases have not been sorted then those cases that do not meet the threshold can be deleted individually by selecting the row the choosing

Edit
Delete current row

➤ The next step is to distinguish configurations that are subsets of the outcome from those that are not. For crisp sets, this determination is made using the measure of set-theoretic consistency reported in the *consist* column. Values below 0.75 indicate substantial inconsistency. It is useful to sort the consistency scores in descending order to evaluate their distribution (this should be done ***after removing rows that fail to meet***

the frequency threshold). Sorting is accomplished by clicking any cell in the *consist* column and choosing

Sort
Descending

Identify any gaps in the upper range of consistency that might be useful for establishing a consistency threshold. Keep in mind that it is always possible to examine several different thresholds and assess the consequences of lowering and raising the consistency cut-off.

➤ It is now necessary to indicate which configurations can be considered subsets of the outcome and which cannot. Input a 1 in the outcome column (*survived* in this example) for each configuration whose consistency level meets and/or exceeds the threshold. Input a 0 in the outcome column for each configuration whose consistency level does not meet the consistency threshold.

➤ Alternatively, one can use the “Delete and code” function to automate this process. Select

Edit
Delete and code

In the first field, the frequency threshold is selected. The default number of cases is 1, but may be changed by typing the selected frequency threshold into the field. In the second field, the consistency threshold is selected. The default consistency is 0.8, but this may be changed by typing the selected consistency threshold into the field.

Click “OK”. The program will delete rows where the frequency threshold is not met, and will code the outcome as 0 or 1 depending on the selected consistency threshold.

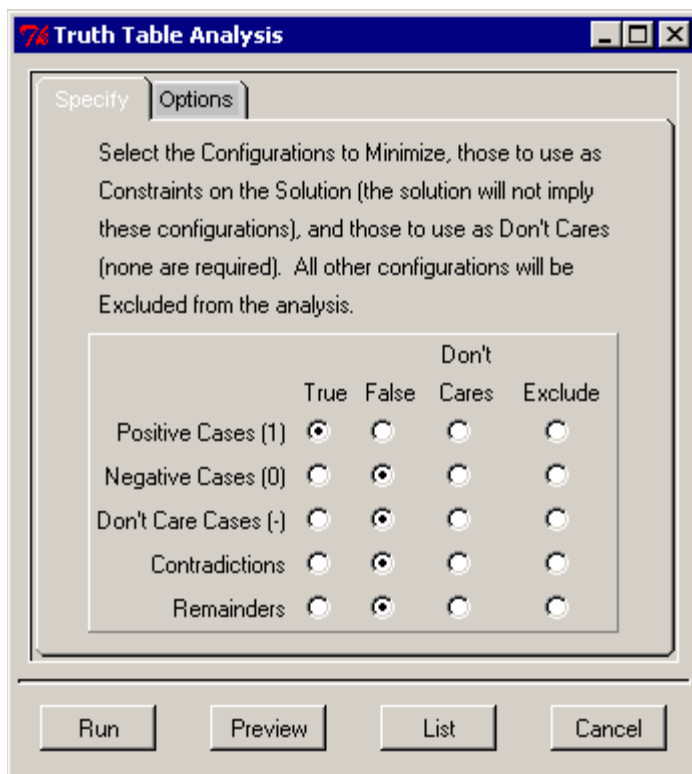
- The following window displays the truth table that would appear after
1. applying a frequency threshold of 1 to the data and eliminating configurations that do not have any observations (8 configurations)
 2. selecting a consistency threshold of 0.9 and placing a 1 in the *survived* column for configurations with 0.90 consistency or greater (3 configurations) and a 0 for cases with lower consistency (5 configurations)

wealthy	literate	industrial	unstable	number	survived	consist	pre	product
1	1	1	0	5	1	1.000000	1.000000	1.000000
1	1	0	0	1	1	1.000000	1.000000	1.000000
0	1	1	0	1	1	1.000000	1.000000	1.000000
0	1	0	0	2	0	0.500000	0.500000	0.250000
0	1	0	1	2	0	0.000000	0.000000	0.000000
1	1	1	1	2	0	0.000000	0.000000	0.000000
0	0	0	0	2	0	0.000000	0.000000	0.000000
0	0	0	1	3	0	0.000000	0.000000	0.000000

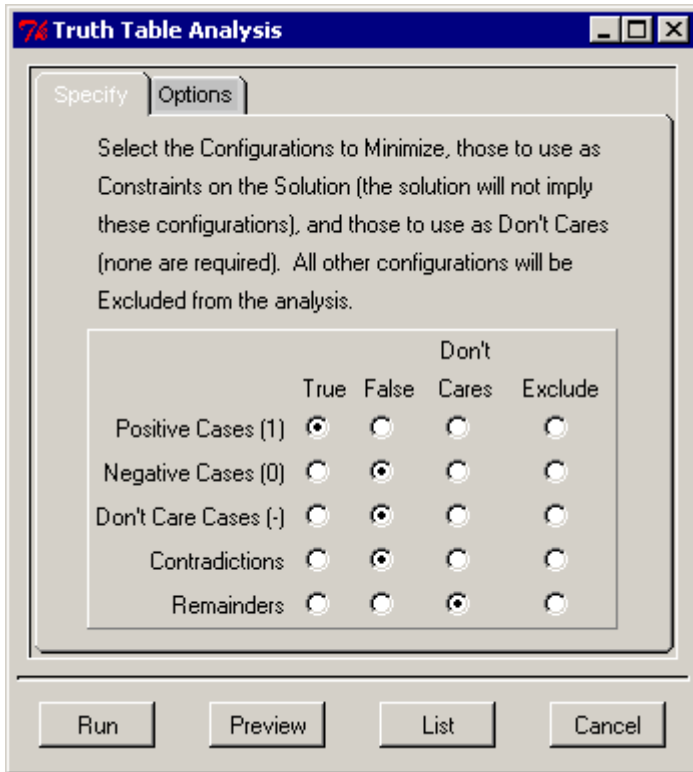
From here, there are two possibilities for the analysis: specifying a single analysis versus deriving the three “standard” analyses (complex, parsimonious and intermediate). Clicking the “Standard Analyses” button (which gives the three solutions) is the recommended procedure.

i) Specify Analysis Option

- Once the truth table is constructed select Specify Analysis to bring up the Truth Table Analysis Window.
- In the Specify panel setting Positive cases to “True” and all the others to “False” yields the “most complex” solution. This window appears as:



- To derive the most parsimonious solution, set Positive cases to “True,” Negative Cases, Don’t Care, and Contradictions to “False,” and Remainders to “Don’t Cares.” The window will appear:



- All other options (for prime implicants and output frequencies) can be set following the procedures outlined below for the Quine algorithm.
- Clicking Preview will present a simple truth table summary in the output window. Generally, this is not a useful button to click for the truth table algorithm because the user has just created and coded the truth table and thus does not need to “preview” it. In any event, the “preview” procedure is a dead end (the truth table *analysis* is not performed) and should not be clicked when using the truth table algorithm.
- Clicking List will give you a list of the configurations of causal conditions found in the data and the consistency of each configuration. Again, this is not a useful procedure for the truth table algorithm, because the user has just constructed and viewed the truth table. Like “Preview” it is a dead end (the truth table analysis is not performed) and should not be clicked when using the truth table algorithm.
- To perform the analysis click the Run button and the output will appear in the output window.

ii) Standard Analyses Option

- Once the truth table is fully constructed, select Standard Analyses. Standard Analyses automatically provides the user with the complex, parsimonious, and intermediate solutions. **“Standard Analyses” is the recommended procedure, as this is the only way to derive the intermediate solution.** To derive the intermediate solution, the software

conducts counterfactual analyses based on information about causal conditions supplied by the user.

Limited Diversity and Counterfactual Analysis

One of the most challenging aspects of comparative research is the simple fact that researchers work with relatively small *Ns*. Investigators often confront "more variables than cases", a situation that is greatly complicated the fact that comparativists typically focus on *combinations* of case aspects—how aspects of cases fit together configurationally. For example, a researcher interested in a causal argument specifying an intersection of four causal conditions ideally should consider all sixteen logically possible combinations of these four conditions in order to provide a thorough assessment of this argument. Naturally occurring social phenomena, however, are profoundly limited in their diversity. The empirical world almost never presents social scientists all the logically possible combinations of causal conditions relevant to their arguments (as shown with hypothetical data in Table 1 below). While limited diversity is central to the constitution of social and political phenomena, it also severely complicates their analysis.

Table 1: Truth table with four causal conditions (A, B, C, and D) and one outcome (Y)

A	B	C	D	Y*
no	no	no	no	no
no	no	no	yes	?
no	no	yes	no	?
no	no	yes	yes	?
no	yes	no	no	no
no	yes	no	yes	no
no	yes	yes	no	?
no	yes	yes	yes	no
yes	no	no	no	?
yes	no	no	yes	?
yes	no	yes	no	?
yes	no	yes	yes	?
yes	yes	no	no	yes
yes	yes	no	yes	yes
yes	yes	yes	no	?
yes	yes	yes	yes	?

* Rows with "?" in the Y column lack cases--the outcome cannot be determined.

As a substitute for absent combinations of causal conditions, comparative researchers often engage in "thought experiments" (Weber [1905] 1949). That is, they imagine counterfactual cases and hypothesize their outcomes, using their theoretical and substantive knowledge to guide their assessments. Because QCA uses truth tables to assess cross-case patterns, this process of considering counterfactual cases (i.e., absent combinations of causal conditions) is explicit and systematic. In fact, this feature of QCA is one of its key strengths. However, the explicit consideration of counterfactual cases and the

systematic incorporation of the results of such assessments into statements about cross-case patterns are relatively new to social science. The specification of best practices with respect to QCA and counterfactual analysis, therefore, is essential.

Imagine a researcher who postulates, based on existing theory, that causal conditions A, B, C, and D are all linked in some way to outcome Y. That is, it is the presence of these conditions, not their absence, that should be linked to the presence of the outcome. The empirical evidence indicates that many instances of Y are coupled with the presence of causal conditions A, B, and C, along with the *absence* of condition D (i.e., $A \cdot B \cdot C \cdot d \rightarrow Y$). The researcher suspects, however, that all that really matters is having the first three causes, A, B and C. In order for $A \cdot B \cdot C$ to generate Y, it is not necessary for D to be absent. However, there are no observed instances of A, B, and C combined with the presence of D (i.e., no observed instances of $A \cdot B \cdot C \cdot D$). Thus, the decisive case for determining whether the *absence* of D is an essential part of the causal mix (with $A \cdot B \cdot C$) simply does not exist.

Through counterfactual analysis (i.e., a thought experiment), the researcher could declare this hypothetical combination ($A \cdot B \cdot C \cdot D$) to be a likely instance of the outcome (Y). That is, the researcher might assert that $A \cdot B \cdot C \cdot D$, if it existed, would lead to Y. This counterfactual analysis would allow the following logical simplification:

$$\begin{aligned} A \cdot B \cdot C \cdot d + A \cdot B \cdot C \cdot D &\rightarrow Y \\ A \cdot B \cdot C \cdot (d + D) &\rightarrow Y \\ A \cdot B \cdot C &\rightarrow Y \end{aligned}$$

How plausible is this simplification? The answer to this question depends on the state of the relevant theoretical and substantive knowledge concerning the connection between D and Y in the presence of the other three causal conditions ($A \cdot B \cdot C$). If the researcher can establish, on the basis of existing knowledge, that there is every reason to expect that the presence of D should contribute to outcome Y under these conditions (or conversely, that the absence of D should not be a contributing factor), then the counterfactual analysis just presented is plausible. In other words, existing knowledge makes the assertion that $A \cdot B \cdot C \cdot D \rightarrow Y$ an "easy" counterfactual, because it involves the addition of a redundant cause (D) to a configuration which is believed to be linked to the outcome ($A \cdot B \cdot C$).

One strength of QCA is that it not only provides tools for deriving the two endpoints of the complexity/parsimony continuum, it also provides tools for specifying intermediate solutions. Consider again the truth table presented in Table 1, which uses A, B, C, and D as causal conditions and Y as the outcome (page 50). Assume, as before, that existing theoretical and substantive knowledge maintains that it is the presence of these causal conditions, not their absence, that is linked to the outcome. The results of the analysis barring counterfactuals (i.e., the complex solution) reveals that combination $A \cdot B \cdot c$ explains Y. The analysis of this same evidence permitting any counterfactual that will yield a more parsimonious result (i.e., the parsimonious solution) is that A by itself accounts for the presence of Y. Conceive of these two results as the two endpoints of the complexity/parsimony continuum, as follows:

A·B·c A

Observe that the solution privileging complexity (A·B·c) is a subset of the solution privileging parsimony (A). This follows logically from the fact that both solutions must cover the rows of the truth table with Y present; the parsimonious solution also incorporates some of the remainders as counterfactual cases and thus embraces additional rows. Along the complexity/parsimony continuum are other possible solutions to this same truth table, for example, the combination A·B. These intermediate solutions are produced when different subsets of the remainders used to produce the parsimonious solution are incorporated into the results. These intermediate solutions constitute *subsets* of the most parsimonious solution (A in this example) and *supersets* of the solution allowing maximum complexity (A·B·c). The subset relation between solutions is maintained along the complexity/parsimony continuum. The implication is that any causal combination that uses at least some of the causal conditions specified in the complex solution (A·B·c) is a valid solution of the truth table as long as it contains all the causal conditions specified in the parsimonious solution (A). It follows that there are two valid intermediate solutions to the truth table:

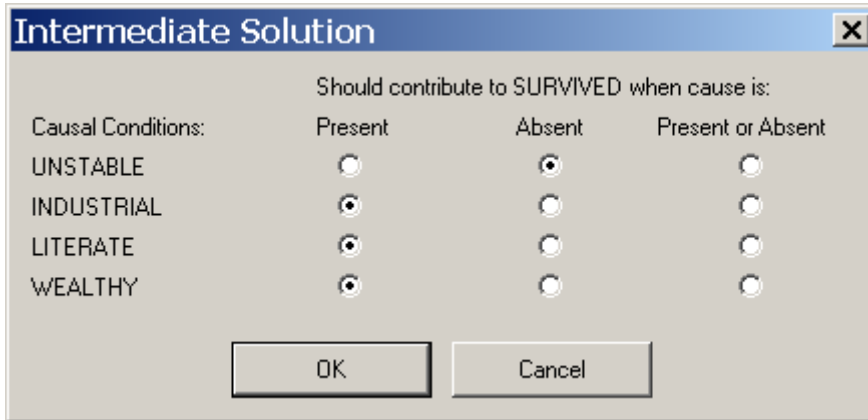
$$\begin{array}{ccc} & A \cdot B & \\ A \cdot B \cdot c & \underline{A \cdot c} & A \end{array}$$

Both intermediate solutions (A·B) and (A·c) are subsets of the solution privileging parsimony and supersets of the solution privileging complexity. The first (A·B) permits counterfactuals A·B·C·D and A·B·C·d as combinations linked to outcome Y. The second permits counterfactuals A·b·c·D and A·b·c·d.

The relative viability of these two intermediate solutions depends on the plausibility of the counterfactuals that have been incorporated into them. The counterfactuals incorporated into the first intermediate solution are "easy" because they are used to eliminate c from the combination A·B·c, and in this example, existing knowledge supports the idea that it is the *presence* of C, not its absence, that is linked to outcome Y. The counterfactuals incorporated into the second intermediate solution, however, are "difficult" because they are used to eliminate B from A·B·c. According to existing knowledge the presence of B should be linked to the presence of outcome Y. The principle that only easy counterfactuals should be incorporated supports the selection of A·B as the optimal intermediate solution. This solution is the same as the one that a conventional case-oriented researcher would derive from this evidence, based on a straightforward interest in combinations of causal conditions that are (1) shared by the positive cases (or at least a subset of the positive cases), (2) believed to be linked to the outcome, and (3) not displayed by negative cases.

➤ After Standard Analysis is selected, a window for guiding the derivation of the intermediate solution will appear. Here, the researcher must select how each causal condition should theoretically contribute to the outcome, as described above. If the condition should contribute to the outcome when present, select "Present." If the condition should contribute to the outcome when absent, select "Absent." If the

condition could contribute to the outcome when it is present OR absent, select “Present or Absent.”



➤ Please note: When the algorithm for selecting prime implicants cannot fully reduce the truth table, the Prime Implicant Window will appear and the user must select the prime implicants to be used, based on theoretical and substantive knowledge. This window is most likely to pop open when the program is deriving the parsimonious solution, but could happen for all three solutions. (See below for a description of how this window operates.

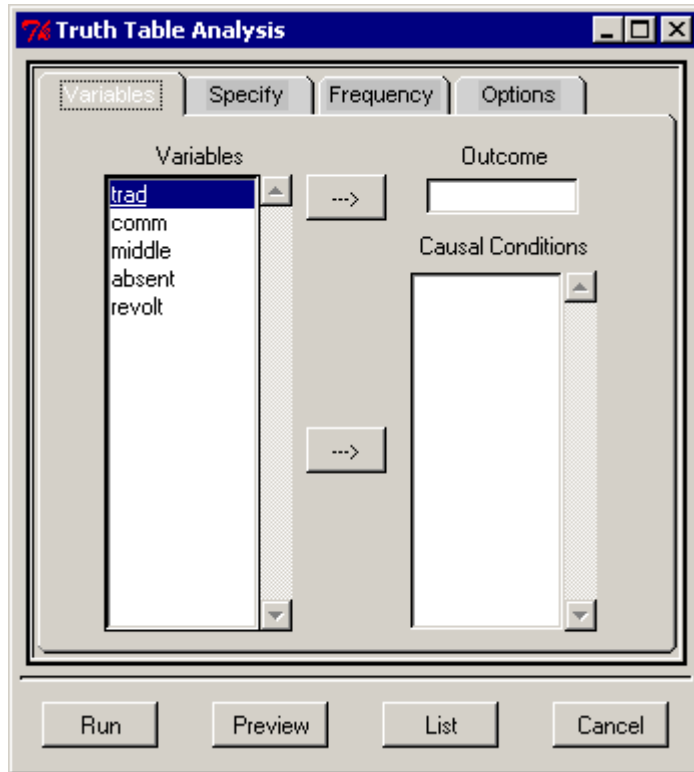
➤ To perform the analysis, click OK and the intermediate solution will appear in the output window. The output window is positioned on the parsimonious solution. The complex solution is directly above the parsimonious, and the intermediate solution is directly below the parsimonious.

The Quine Algorithm (QCA 3.0)

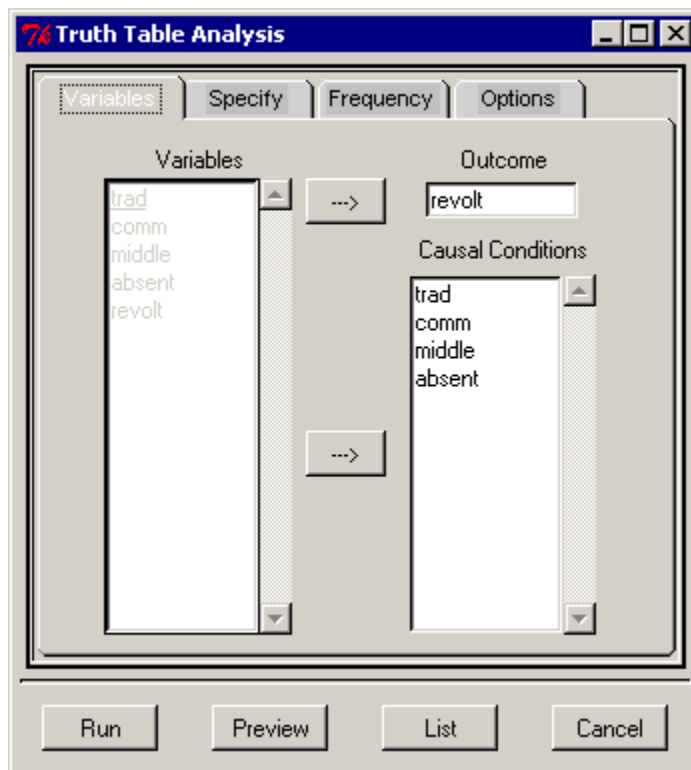
The Quine algorithm for crisp-set analysis is not as easy to use as the truth table algorithm because it does not present the user with a code-able truth table spreadsheet as an intermediate step. It is retained in the software, however, for users who are more familiar with it. (The Quine procedure was implemented first.)

➤ In order to analyze your crisp data set using the original Quine algorithm, choose
 Analyze
 Crisp Sets
 Quine...

➤ The following window will open (note that the highlighted box is the *Variables* box):



- Identify the outcome you want to explain and transfer it into the *Outcome* field.
- Highlight among the remaining conditions the ones that you think explain the outcome and transfer them into the *Causal Conditions* field.



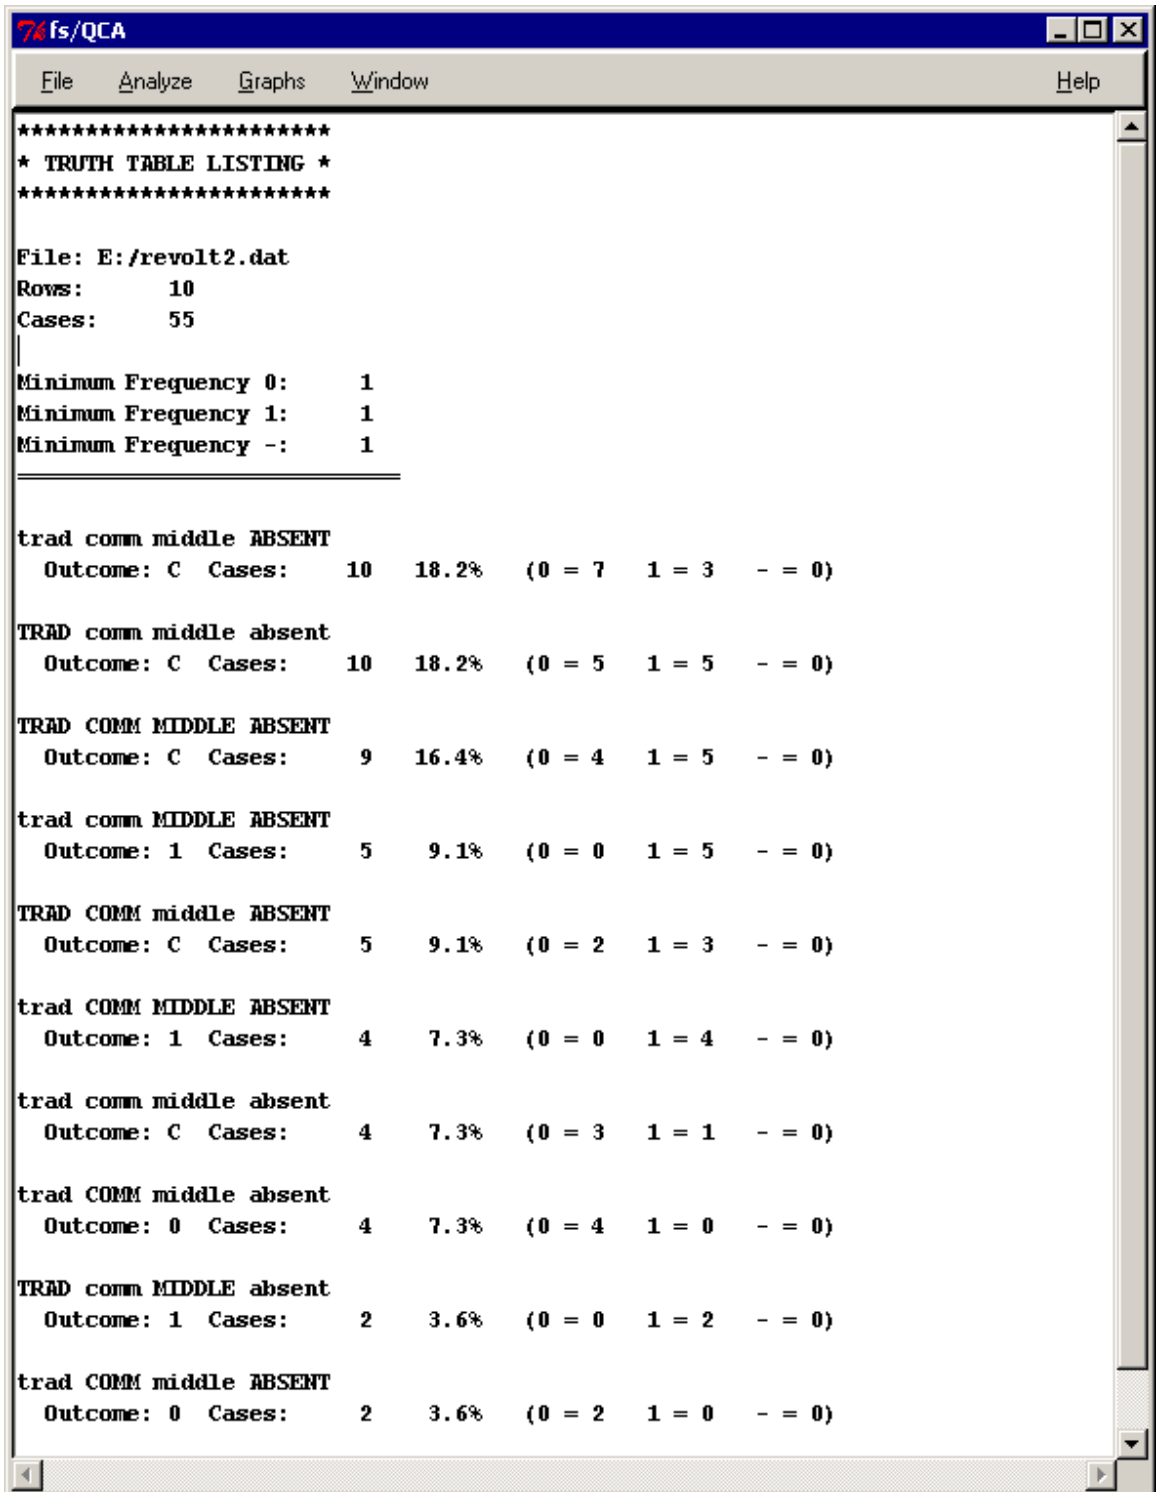
Listing Truth Table Rows

➤ In order to view a listing of the truth table you just specified, choose:

Analyze
Crisp-Sets
Quine.....

➤ Specify your outcome and causal conditions and click the *List* button.

➤ The following output will be shown in the main window:



[Compare to Table 9 in *The Comparative Method*, page 114]

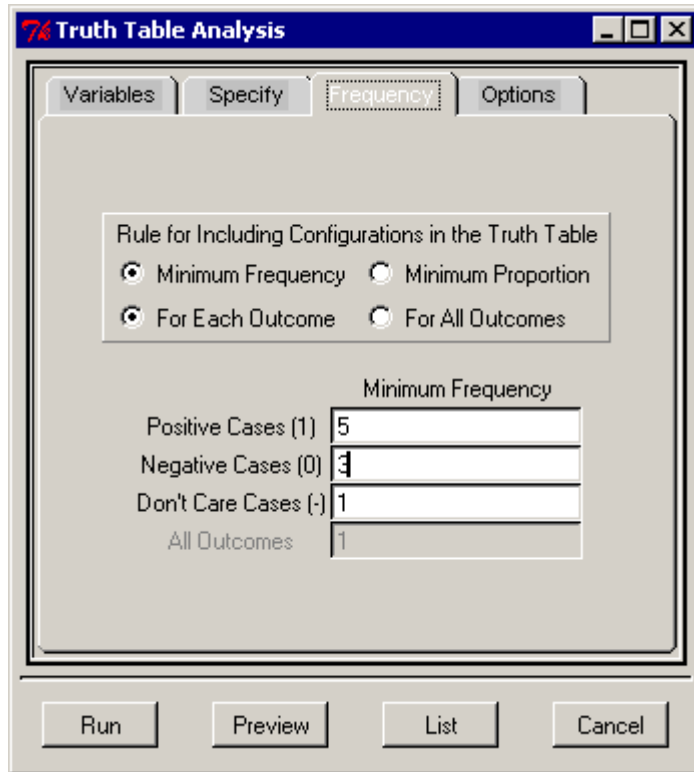
- Cases with missing values for one or more variables are not included in the configurations in the truth table.
- The configurations are sorted according to number of instances (cases).

- Each row displays one configuration of causal conditions. Absence of a condition is indicated by lower-case letters, the presence of a condition by upper-case letters.
- What follows is the accumulated outcome of all cases with this specific configuration of causes. In the example above, 5 rows have clear-cut outcomes (4,6,8,9,10). In row #4, for instance, all five cases have the outcome 1. The four cases of row #8, on the other hand, all have the outcome 0. In the event that not all cases of a certain configuration indicate the same outcome value, the table symbolizes this contradiction by assigning a “C” as the outcome value. In the truth table above, there are five configurations with mixed, or contradictory, outcomes (1,2,3,5,7). The ten cases in row #1, for example, have contradictory outcomes—three exhibit the outcome 1, while seven do not.
- You can now (optionally) specify the minimum frequency or proportion of cases as a requirement for the inclusion in the truth table.

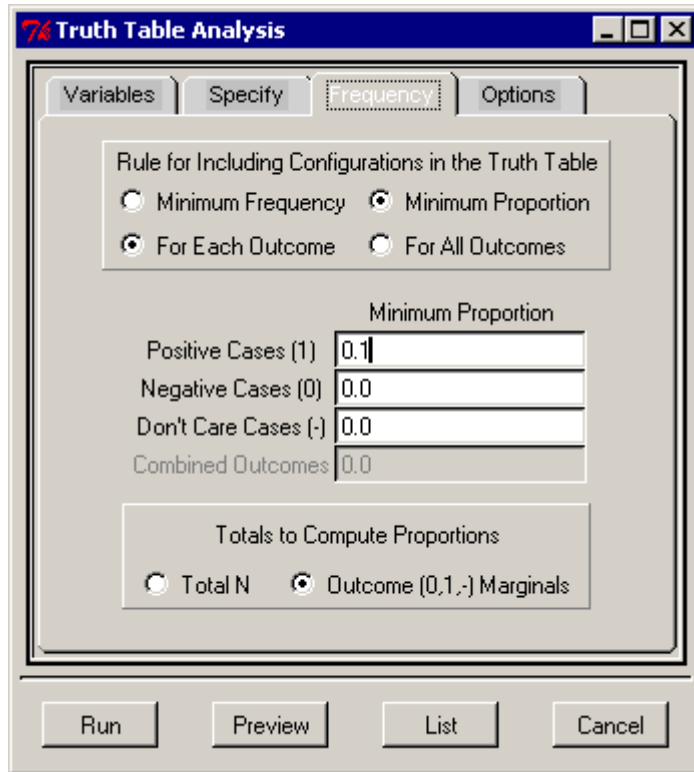
In general, this assessment of the evidence is easier to accomplish with the crisp truth table algorithm

E) Rules for Including Configurations

- In order to specify the rules for including configurations in the truth table, choose:
 - Analyze
 - Crisp-Sets
 - Quine.....
- Click on the *Frequency* option and the following window will open (note that the highlighted box is the *Frequency* box):



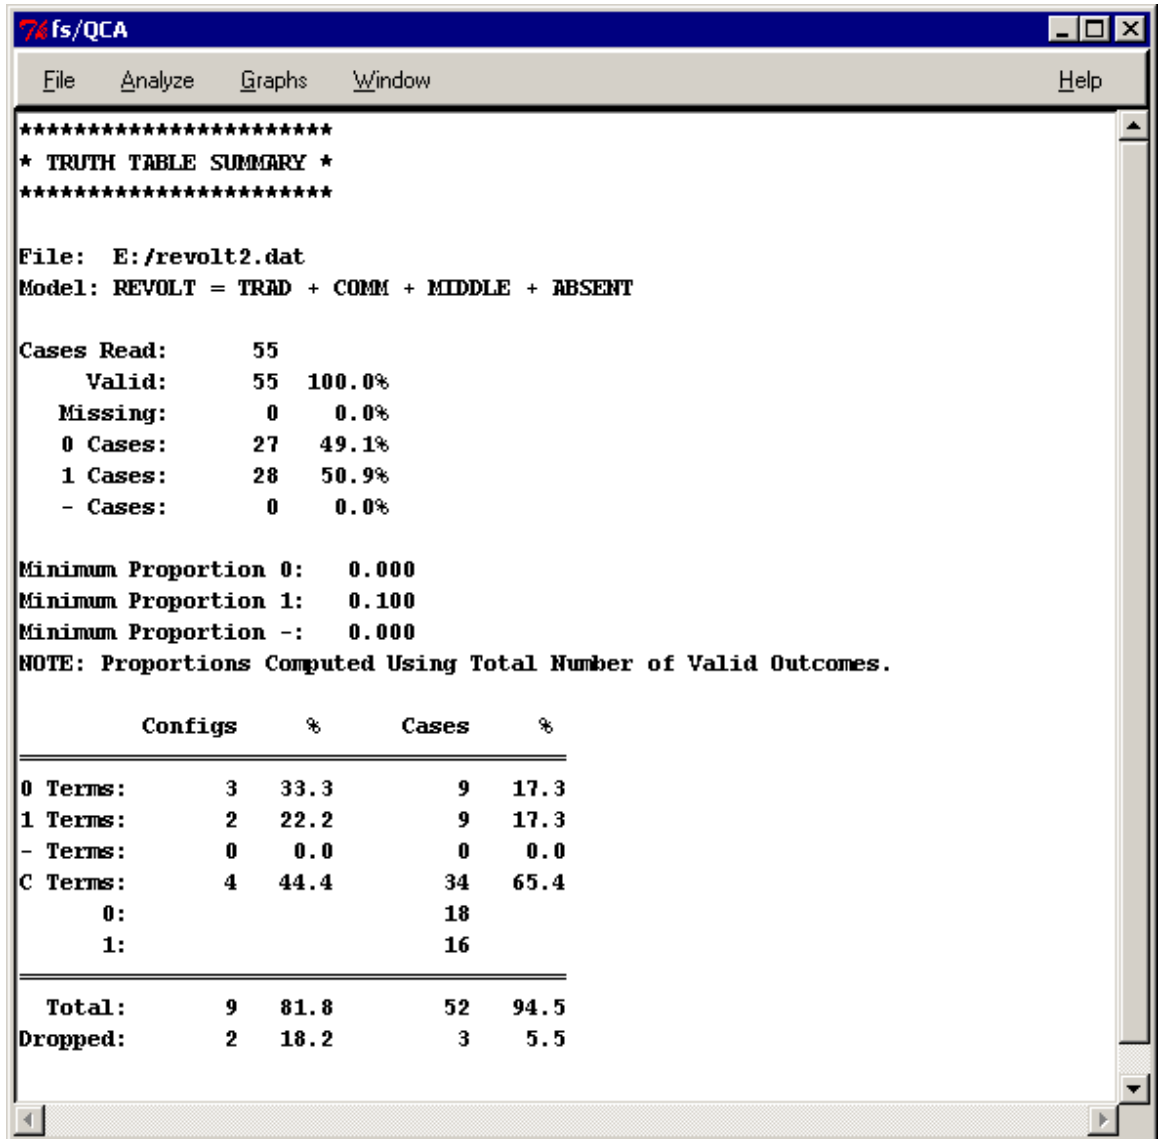
- You can either determine the minimum frequency of cases for each configuration or the minimum proportion. (The above displayed window shows the option of minimum frequency.)
- Additionally, you can determine whether you want different minimum frequencies for the three different kinds of outcomes. In the example above, the researcher chose to only include configurations if the 0 outcome is represented by at least 3 cases, and the 1 outcome by at least 5. Configurations with “Don’t Care” outcomes were excluded via the “Specify” tab.
- If you decide to set the minimum frequency to the same level for all three possibilities of outcomes, choose *All Outcomes*, and set the desired limit.
- If you choose to set the limit of cases for each configuration as a proportion, click on the *Minimum Proportion* option and the following window will appear:



- By choosing to determine the limits via proportions for each outcome, you will also have to determine the denominator the proportions. The options are:
 - Total N
 - Outcomes

In the example above, the researcher decided to include configurations with the outcome value 1 only if they are represented by a minimum proportion of cases of 10%. Applied to the example of this chapter, this limit of 10% excludes all existing configurations with the outcome 1—the configurations of row 4 and row 6 are represented only by 9.1% and 7.3% of all cases, respectively (see truth table on page 41).

- In order to inspect the limitations you have set for the inclusion of configurations, click on the *Preview* button. The following will be displayed in the main window:



The first two rows of this output indicate the data file location and the specified outcome as well as the causal conditions.

What follows is summary of the cases of the data set, including the number and percentage of the valid and missing cases, and the number and percentage of the 0, 1, and “Don’t Care” configurations.

The next three lines, then, indicate the limits set for inclusion in the truth table (compare to the minimum frequency example on page 43).

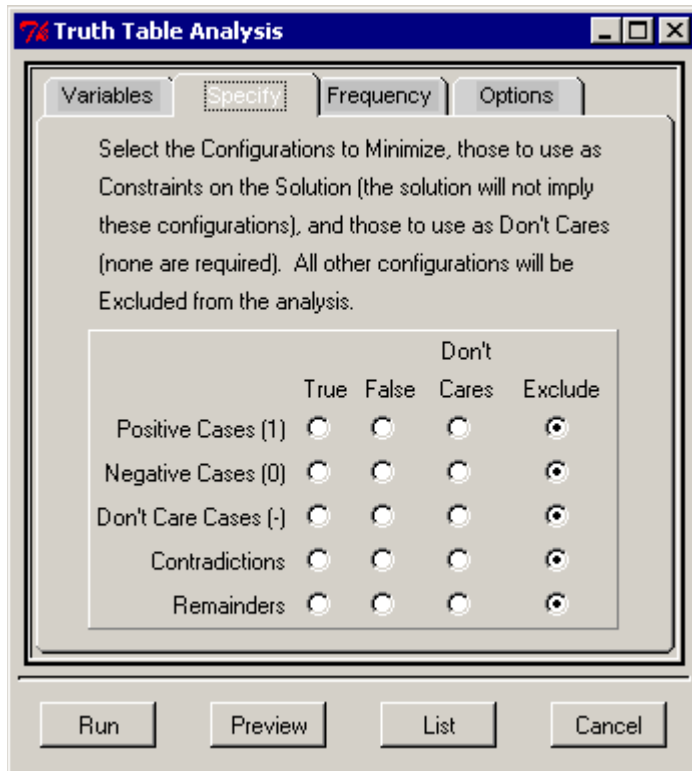
Finally, the table summarizes the modified truth table. The last row of this table indicates the number and proportions of configurations and the cases that have been excluded according to the researcher’s specifications (2 configurations with 3 cases).

F) Setting Configurations to Minimize

- In order to select the configurations to minimize, those to use as constraints on the solution, and those to use as “Don’t Cares,” choose:

Analyze
 Crisp-Sets
 Quine.....

- Click on the *Specify* option and the following window will open (note that the highlighted box is the *Specify* box):



In order to explain the options given in the window above, it is necessary to clarify the different possible output codes for each of the configurations. Imagine the following truth table as part of a data set:

	X_1	X_2	X_3	Cases	Instances	(Output Code)
Row 1	0	0	0	5	0	1
Row 2	0	0	1	8	8	0
Row 3	1	1	0	1	-	-
Row 4	0	1	1	5	2	C
Row 5	1	1	1	0	?	?

The output code of each configuration can be one of the following possibilities:

Positive Cases (1): For all cases in these kinds of configurations the outcome is absent (see row 1).

Negative Cases (0): For all cases in these kinds of configurations the outcome is present (see row 2).

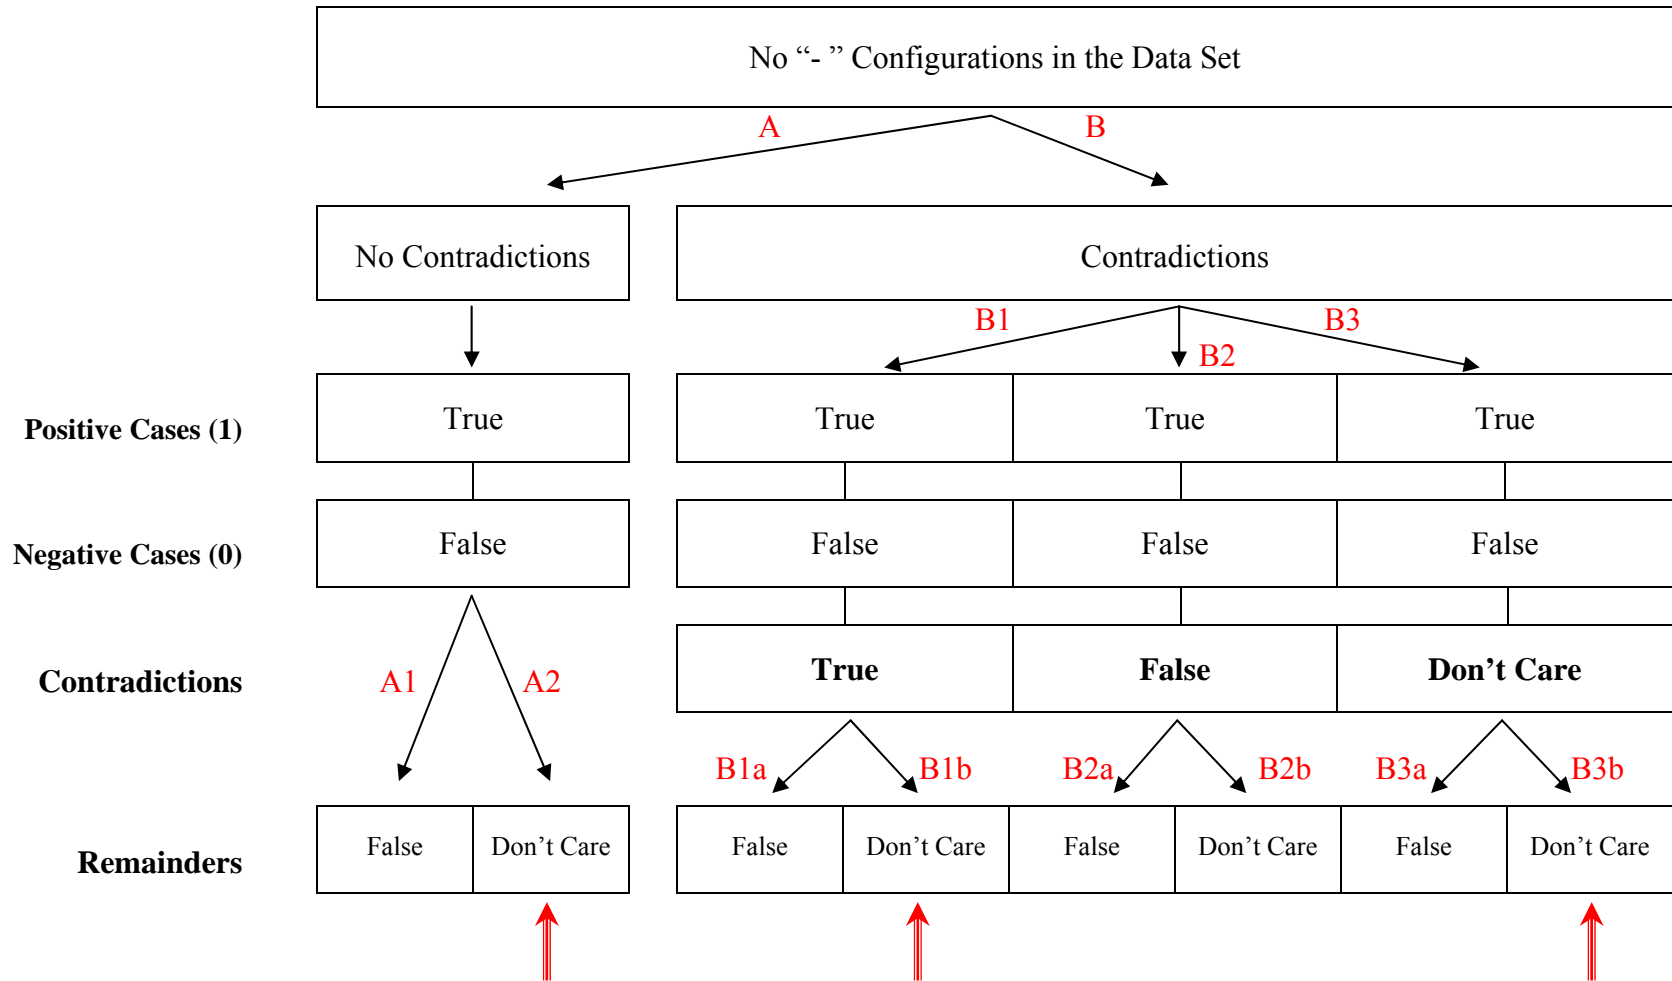
Don't Care Cases (-): For all cases of these configurations the researcher has assigned the "Don't Care" value in the data spreadsheet. This allows the computer to treat these configurations having either outcome present or having the outcome absent, depending on the situation (see row 3). This assignment is rarely used.

Contradictions: Not all cases have the same outcome value. Row 4 of the truth table demonstrates this by having only 2 of the 5 cases in this configuration having the outcome present (3 cases have the outcome absent).

Remainders: These are configurations with no cases in the data set (see row 5).

Please note that in Quine minimization "Exclude" and "False" are equivalent. Only two assignments really matter for the Quine routine: "True" and "Don't Care."

Eight “Most Common” Ways:



G) Options

1) Prime Implicants

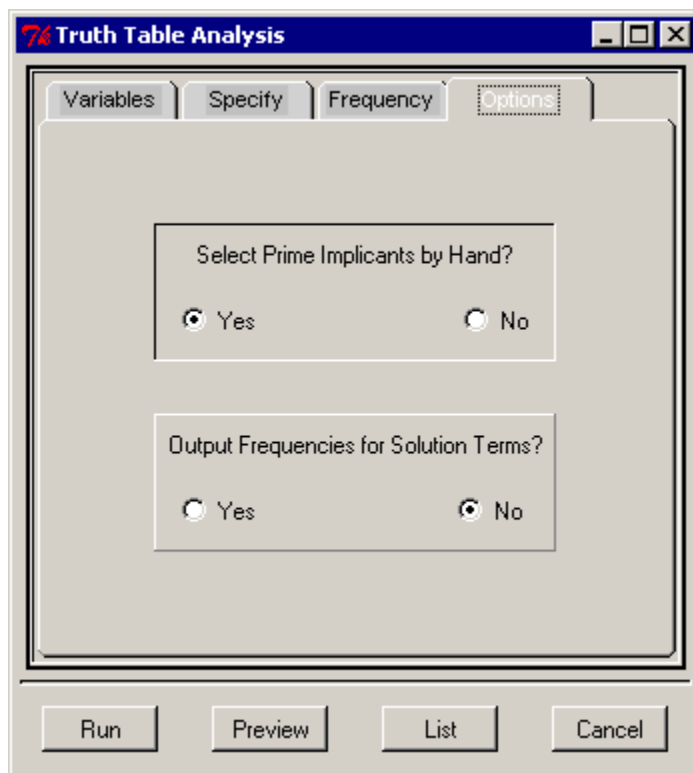
Prime implicants (PIs) are product terms that are produced using minimization rules (e.g. rules that combine rows that differ on only one cause if they have the same output value). For example: ABC combines with AbC to produce AC . AC , therefore, is the prime implicant that covers the two primitive Boolean expressions ABC and AbC . In other words, ABC and AbC are subsets of AC , or AC implies ABC and AbC .

Often, though, there are more reduced prime implicants than are needed to cover all the original primitive expressions and the user has the option to choose from among those that are “logically tied” using the prime implicant chart. (For a more comprehensive discussion of the use of Prime Implicants, refer to *The Comparative Method* page 95.)

➤ In order to specify whether you want to exercise this option, choose:

Analyze
Crisp-Sets
Quine.....

➤ Click on *Options* and the following window will open (note that the highlighted box is the *Options* box):

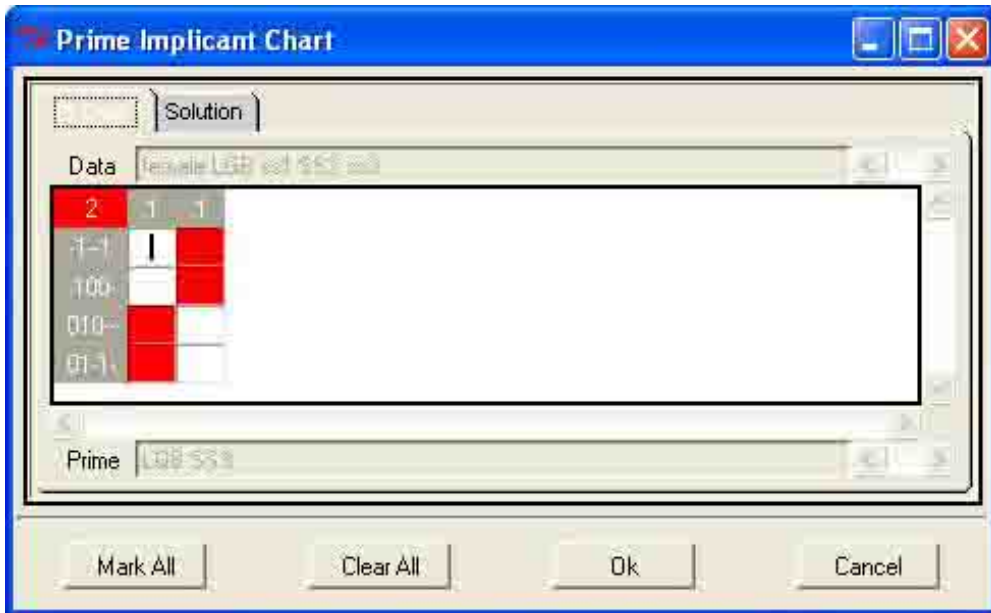


➤ In order to be able to select among prime implicants that are logically equivalent (i.e., that cover the same truth table rows), choose YES when asked to select prime

implicants by hand. If NO is selected, the program will report all of the possible choices in the solution. A partially reduced PI chart will essentially be reproduced in the output, with options under the heading(s) “One of the Following.”

Choosing Solution Terms Using the Prime Implicant Chart

- If the user opts to pick prime implicants by hand, the Prime Implicant Chart window will appear only when there are logically equivalent prime implicants to choose from.
 - To choose prime implicants (PIs), the program employs an algorithm that attempts to reduce the table until no further simplification is possible, beginning with essential PIs (which uniquely cover specific rows in the truth table) that must appear in the solution. If the algorithm is run and the table cannot be fully reduced, the user may select the PIs to be used, based on theoretical and substantive knowledge.
 - The Prime Implicant Chart has two tabs. Selecting “Solution” will display essential PIs that must be included in the solution. If this field is blank, there are no essential PIs.
 - The Prime Implicant Chart tab displays the possible prime implicants for the user to choose. Each column in the chart represents a different truth table row that is covered by more than one prime implicant. The “Data” field across the top displays the truth table row in question (the one that needs to be covered).
 - The “Prime” field describes the PI that the user may select. Each row in the chart represents one PI that may be selected. Clicking on the interior of the table will cause the column information and the row information to be displayed for that cell.
 - A PI is selected by clicking on the cell in the first column with the combination of 0’s, 1’s, and -’s that represents the PI. If you change your mind about your selection, you can click “Clear All” to unselect the PI. “Mark all” may be used to select all PIs.
 - The number in the top left corner of the chart describes how many PIs need to be selected. The numbers at the top of the columns describe how many PIs need to be selected for each truth table row in question. (Remember that the columns in this chart represent truth table rows that still need to be covered.) As an appropriate PI is selected, both the number in the corner and the number at the top of the column will decrease by 1. If there is more than one truth table row with multiple PIs represented in the column of this chart, the user is required to choose at least one PI for each. The cells representing the PIs that correspond to each row are colored red; when the PI is selected the cell colors of the possible PIs for the truth table row will turn grey. When one or more PIs are selected for each truth table row, the cell in the upper-left corner will say “0” and will turn green.



For this chart, the user could select the first and third PI, the first and fourth PI, the second and third PI, or the second and fourth PI. Of course, additional PIs could be selected, but these four combinations of choices would provide minimal coverage of the remaining truth table rows, represented in the columns of the chart.

Frequencies for Solution Terms

- In order to specify whether you want output frequencies for the solution term
 - Analyze
 - Crisp-Sets
 - Quine.....
- Click on *Options* and select YES or NO for “Select Output Frequencies Solution Terms.”

This option will include the number of cases for each of the crisp set solution terms. In the solution term below, for example, 12 cases of the data set have the outcome with a causal combination of the absence of “traditional” combined with the absence of “middle.”

traditional middle (12) +
 traditional ABSENTEE (5) +
 COMMERC ABSENTEE (24) +
 TRADITIONAL commerc absentee (14)

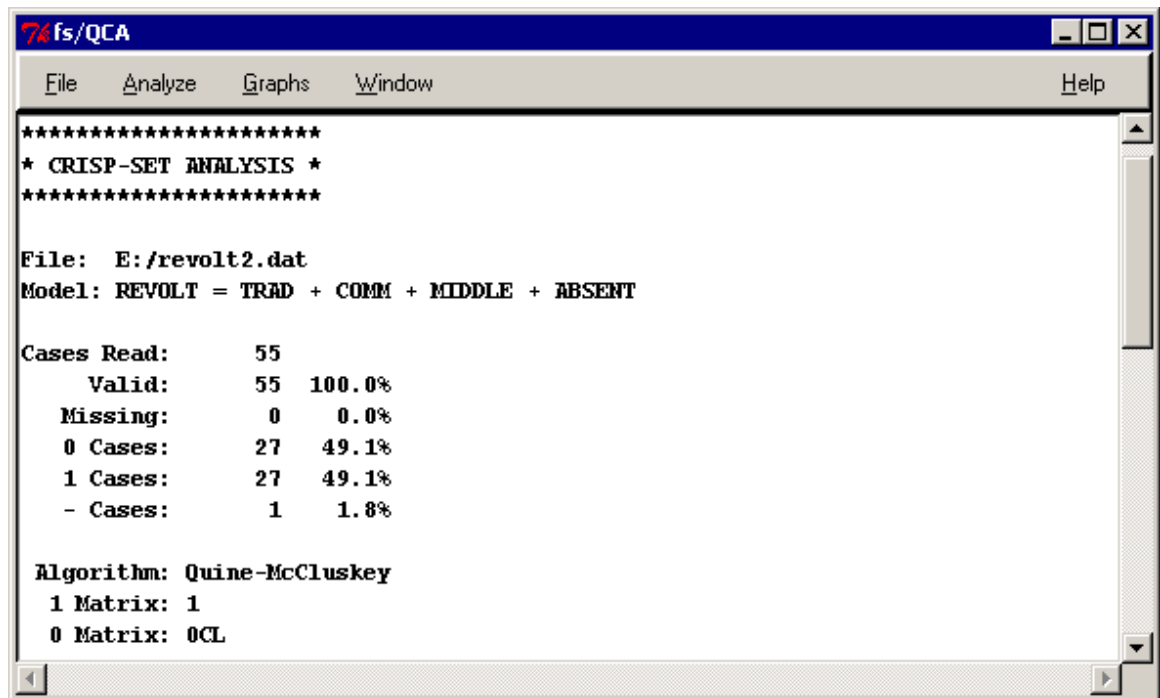
This information is also supplied by the truth-table algorithm via the calculation of the “coverage” of solution terms.

D) Output (Crisp Sets)

Once you have minimized the truth table according to your specifications, the main window will show you the following output.

Output from Quine Procedure

The first part of the output describes your data. The first two rows indicate the file directory and the model you specified. What follows is a summary of the truth table with the number and percentages of cases in each kind of configuration. Then, the output displays what kind of algorithm you used, and which configurations were minimized and which were set to constraint. In the example below, the configurations with the outcome were minimized, whereas the configurations without the outcome, the contradictions, and the remainders were set to constraint.



```
*****
* CRISP-SET ANALYSIS *
*****

File: E:/revolt2.dat
Model: REVOLT = TRAD + COMM + MIDDLE + ABSENT

Cases Read:      55
  Valid:         55 100.0%
  Missing:        0  0.0%
  0 Cases:       27 49.1%
  1 Cases:       27 49.1%
  - Cases:        1  1.8%

Algorithm: Quine-McCluskey
  1 Matrix: 1
  0 Matrix: 0CL
```

The truth table summary tells you whether you have excluded any of your configurations, how many cases are in each of the configurations, and how many are left in your analysis. In the example below, none of the configurations have been dropped.

*** TRUTH TABLE SUMMARY ***

Minimum Frequency 0: 1
 Minimum Frequency 1: 1
 Minimum Frequency -: 1

	Configs	%	Cases	%
0 Terms:	3	30.0	9	16.7
1 Terms:	3	30.0	11	20.4
- Terms:	0	0.0	0	0.0
C Terms:	4	40.0	34	63.0
0:			18	
1:			16	
Total:	10	100.0	54	100.0
Dropped:	0	0.0	0	0.0

The last part of the output yields the solution of your crisp-set analysis.

*** CRISP-SET SOLUTION ***

trad*MIDDLE*ABSENT+
 TRAD*comm*MIDDLE*absent

Revolts are caused by a combination of either (1) the absence of peasant traditionalism, and the presence of a middle class of peasants and absentee landed elites or (2) the presence of peasant traditionalism and a middle class of peasants and the absence of commercial agriculture and absentee landed elites.

Output from Truth Table Procedure

Once the truth table has been specified and minimized, the main window will show you the output. The first part of the output describes the data. The first two rows indicate the file directory and the model you specified. The next row lists the number of truth table rows read for the analysis. Then, the output displays what kind of algorithm you used, and which configurations were minimized.

The complex solution is always presented first, then the parsimonious, then the intermediate. All three are labeled. Also reported are the frequency and consistency cut-offs that were used. Rather than reporting the actual cut-offs used (e.g., in the delete and

code function of the truth table spreadsheet), the output reports the lowest frequency configuration actually used and the lowest consistency score, again, among those actually used in the procedure. For example, if you selected a consistency cut-off of 0.80, but there was a gap in the observed consistency scores between 0.90 and 0.70, then 0.90 would be reported as the consistency cut-off value because it would be the lowest consistency among the rows coded with an output equal to true (or 1.0). The complex solution is as follows:

File: C:/Documents and Settings/Desktop/benoit.csv
 Model: SURVIVED = f(WEALTHY, LITERATE, INDUSTRIAL, UNSTABLE)

Rows: 8

Algorithm: Qui ne-McCluskey
 True: 1

--- COMPLEX SOLUTION ---
 frequency cutoff: 1.000000
 consistency cutoff: 1.000000

	raw coverage	uni que coverage	consi stency
WEALTHY*LITERATE*unstable+	0.750000	0.125000	1.000000
LITERATE*INDUSTRIAL*unstable	0.750000	0.125000	1.000000

solution coverage: 0.875000
 solution consistency: 1.000000

The report of the frequency and consistency cutoff values is followed by the solution. When the user selects “Specify Analysis,” there will be only one Truth Table Solution section. When “Standard Analyses” is selected, solutions will be provided for the complex, parsimonious, and intermediate configurations and will be labeled as such. In addition, the intermediate solution will contain the list of assumptions about causal conditions specified by the user via the dialogue box.

The solution provides a line for each separate path to the outcome. In this example, there are two combinations of conditions linked to the outcome, democratic survival, not being unstable combined with being literate and wealthy, and not being unstable combined with being literate and industrial. Consistency and coverage scores (including raw and unique coverage) are sketched below and explained in detail in *Redesigning Social Inquiry: Fuzzy Sets and Beyond*.

The most parsimonious solution for this truth table is:

File: C:/Documents and Settings/Desktop/benoit.csv
 Model: SURVIVED = f(WEALTHY, LITERATE, INDUSTRIAL, UNSTABLE)

Rows: 8

Algorithm: Qui ne-McCluskey
 True: 1-L

--- PARSIMONIOUS SOLUTION ---
 frequency cutoff: 1.000000
 consistency cutoff: 1.000000

	raw coverage	uni que coverage	consi stency
	-----	-----	-----
WEALTHY*unstab l e+	0. 750000	0. 125000	1. 000000
I NDUSTRI AL*unstab l e	0. 750000	0. 125000	1. 000000
sol uti on coverage:	0. 875000		
sol uti on consi stency:	1. 000000		

Again, the output displays what kind of algorithm was used, and which configurations were minimized. The solution again indicates two paths to survival, not being unstable combined with being wealthy, and not being unstable combined with being industrial.

Intermediate solutions are usually the most interpretable, but the parsimonious solutions show which conditions are essential to distinguishing between positive and negative cases.

5. FUZZY-SET ANALYSIS

This part of the manual addresses the analysis of **fuzzy sets**, discussed in depth in *Fuzzy-Set Social Science* (Ragin, 2000) and *Redesigning Social Inquiry* (Ragin, 2008). Instead of allowing only two mutually exclusive states, membership and nonmembership, fuzzy sets extend crisp sets by permitting membership scores in the interval between 0 and 1. There are many ways to construct fuzzy sets. Three common ways are:

four-value fuzzy sets (0, .33, .67, 1)
six-value fuzzy sets (0, .2, .4, .6, .8, 1)
and continuous fuzzy sets (any value ≥ 0 and ≤ 1)

There are two fuzzy-set algorithms, the “inclusion” algorithm and the “truth table” algorithm. The inclusion algorithm is described in *Fuzzy-Set Social Science*; the truth table algorithm is described in *Redesigning Social Inquiry* (Ragin 2008) and also in *Configurational Comparative Methods* (Rihoux and Ragin 2008). The inclusion algorithm is currently being overhauled to make it more robust and also to make it more consistent with the truth table approach. Thus, access to the inclusion algorithm is currently blocked. The truth table algorithm has proven to be more robust and is the preferred approach for now.

A) Operations on Fuzzy Sets

The logical operations AND and OR are used in the fuzzy-set algorithms but are different from the use in crisp-sets. What follows is an introduction of the common operations: logical AND, and logical OR.

Logical AND. With fuzzy sets, logical AND is accomplished by taking the minimum membership score of each case in the sets that are intersected. For example, if a country’s membership in the set of poor countries is .34 and its membership in the set of democratic countries is .91, its membership in the set of countries that are poor and democratic is the smaller of these two scores, .34.

Logical OR. Two or more sets also can be joined through logical OR—the union of sets. For example, a researcher might be interested in countries that are “developed” OR “democratic” based on the conjecture that these two conditions might offer equivalent bases for some outcome (e.g., bureaucracy-laden government). Conventionally, crisp categories would be used to compile a complete list of countries that are “developed or democratic” (i.e., countries that have one or both characteristics). With fuzzy sets, the researcher focuses on the maximum of each case’s memberships in the component sets. That is, membership in the set formed from the union of two or more component sets is the maximum value of the case’s memberships in the component sets. Thus, if a country has a score of .15 in the set of democratic countries and a score of .93 in the set of developed countries, it has a score of .93 in the set of countries that are “democratic or developed.”

Negation. Like in crisp sets, fuzzy sets can be negated. In crisp set logic, negation switches membership scores from 1 to 0 and from 0 to 1. This simple mathematical principle holds in fuzzy algebra as well. The relevant numerical values are not restricted to the Boolean values 0 and 1 but extend to values between 0 and 1 as well. To calculate the membership of a case in the negation of fuzzy set A, simply subtract its membership in set A from 1, as follows:

Fuzzy membership in set not A = 1 – fuzzy membership in set A.

This can be displayed as $\sim A_i = 1 - A_i$, where the subscript “i” indicates the “ith” case, the set “not A” is represented as $\sim A$, and the symbol “ \sim ” denotes negation. Thus, for example, if the United States has a membership score of .79 in the set of “democratic countries,” it has a score of .21 in the set of “not democratic countries.”

B) Fuzzy Sets, Necessity, and Sufficiency (Fuzzy Subset Relation)

Subset principle and arithmetic relationship between membership scores in CRISP sets. Consider the example of state breakdown being a necessary but not sufficient condition of social revolution (p. 211 in *Fuzzy Set Social Science*). It follows logically that if a condition is necessary but not sufficient for an outcome, then instances of the outcome will constitute a subset of instances of the cause. Another way to understand the subset relationship is in terms of the arithmetic relationship between crisp-set membership scores (1s and 0s). If instances of the outcome are a subset of instances of the cause, then the Boolean value of the outcome (1 versus 0) will be less than or equal to the Boolean value of the cause.

Subset principle and arithmetic relationship between membership scores in FUZZY sets. With fuzzy sets it would be difficult to “select” countries with the outcome (the usual first step in the crisp-set analysis of necessary conditions) because countries vary in their degree of membership in the set displaying social revolution. Likewise, it would be very difficult to evaluate cases’ agreement with respect to the relevant causal condition (state breakdown) because they vary in their membership in this set as well.

Fortunately, the subset principle and the arithmetic relationship between membership scores holds for fuzzy sets as well. With fuzzy sets, set A is a subset of set B if the membership scores of cases in set A are less than or equal to their respective membership scores in set B. Furthermore, when fuzzy membership scores in the outcome are less than or equal to fuzzy membership in the cause, then it is possible to argue that instances of the outcome are a subset of instances of the cause. Figure 1 displays this arithmetic relationship in two dimensions. When researchers find this pattern, they might cite this evidence as support for an argument of causal necessity.




Figure 1: Plot of “social revolution” against “state breakdown”

The evaluation of sufficiency can be seen as a test of whether the cases displaying the causal conditions form a subset of the cases displaying the outcome. As shown above, another way to understand the subset relationship is in terms of the arithmetic relation between membership scores. In order to argue that a cause or causal combination is a sufficient for the outcome, the fuzzy membership scores in cause have to be less than or equal to the fuzzy membership in the outcome.

Consider the following example taken from *Fuzzy-Set Social Science*, p. 236ff. Figure 2 displays the arithmetic relationship between the sufficient causal combinations (\sim cross-class • \sim multiracial) against the outcome (ideological conflict). The upper-triangular plot shown in Figure 2 is a direct reflection of the fact that membership scores in the fuzzy set “race and class homogeneity” are less than or equal to membership scores in the fuzzy set “ideological conflict.”

Note the important difference between the application of the subset principle to the assessment of sufficiency and its applications to the assessment of necessity. To demonstrate necessity the researcher must show that the outcome is a subset of the cause. To support an argument of sufficiency, the researcher must demonstrate that the cause is a subset of the outcome.




Figure 2: Plot of “ideological conflict” against “race and class homogeneity”

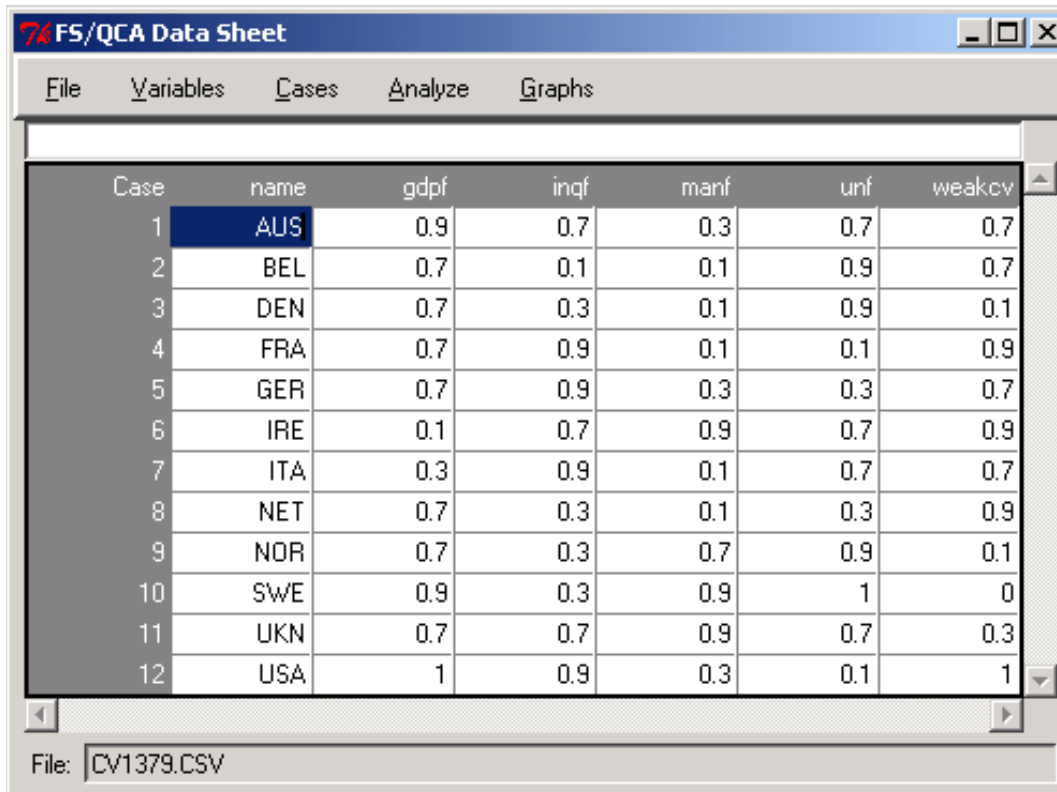
C) Using the Fuzzy Truth Table Algorithm

This method for analyzing fuzzy sets using Truth Tables was introduced in version 2.0 of fsQCA. It is described in detail in Ragin’s (2008) *Redesigning Social Inquiry: Fuzzy Sets and Beyond* and in Rihoux and Ragin’s (2008) *Configurational Comparative Methods: Qualitative Comparative Analysis (QCA) and Related Techniques*.

The fuzzy truth table algorithm can be conceptualized as a bridge with three pillars. The first pillar is the direct *correspondence* that exists between the rows of a crisp truth table and the corners of the vector space defined by fuzzy set causal conditions (see Ragin 2000). The second pillar is the assessment of the *distribution of cases* across different logically possible combinations of causal conditions (or corners of the vector space). Some corners of the vector space may have many cases with strong membership while other corners may have cases with only weak membership. The third pillar is the assessment of the *consistency of the evidence* for each causal combination with the argument that it is a fuzzy subset of the outcome. The truth table algorithm involves establishing these three pillars to construct a crisp truth table, at which point the analysis proceeds similar to the crisp algorithm. This section will explain the steps involved in recording the results of multiple fuzzy set analyses in a crisp truth table and then analyzing that table.

Data

Fuzzy set data can be imported from other programs or created in fsQCA as described in Sections 1. and 2. This chapter will use the example of countries with weak class voting from Ragin (2005). The table below depicts the data sheet:



The screenshot shows a software window titled "FS/QCA Data Sheet" with a menu bar containing "File", "Variables", "Cases", "Analyze", and "Graphs". The main area displays a data table with the following content:

Case	name	gdpf	inqf	manf	unf	weakcv
1	AUS	0.9	0.7	0.3	0.7	0.7
2	BEL	0.7	0.1	0.1	0.9	0.7
3	DEN	0.7	0.3	0.1	0.9	0.1
4	FRA	0.7	0.9	0.1	0.1	0.9
5	GER	0.7	0.9	0.3	0.3	0.7
6	IRE	0.1	0.7	0.9	0.7	0.9
7	ITA	0.3	0.9	0.1	0.7	0.7
8	NET	0.7	0.3	0.1	0.3	0.9
9	NOR	0.7	0.3	0.7	0.9	0.1
10	SWE	0.9	0.3	0.9	1	0
11	UKN	0.7	0.7	0.9	0.7	0.3
12	USA	1	0.9	0.3	0.1	1

At the bottom of the window, the "File:" field shows "CV1379.CSV".

name	Country Identifier
gdpf	Affluent
inqf	Substantial Income Inequality
manf	Strong Manufacturing Sector
unf	Strong Unions
weakcv	Weak Class Voting

Interval and ratio scale data can be converted to fuzzy set membership scores using the “calibrate” procedure described in section 2 of the manual (Data Editor) and in Ragin (2008).

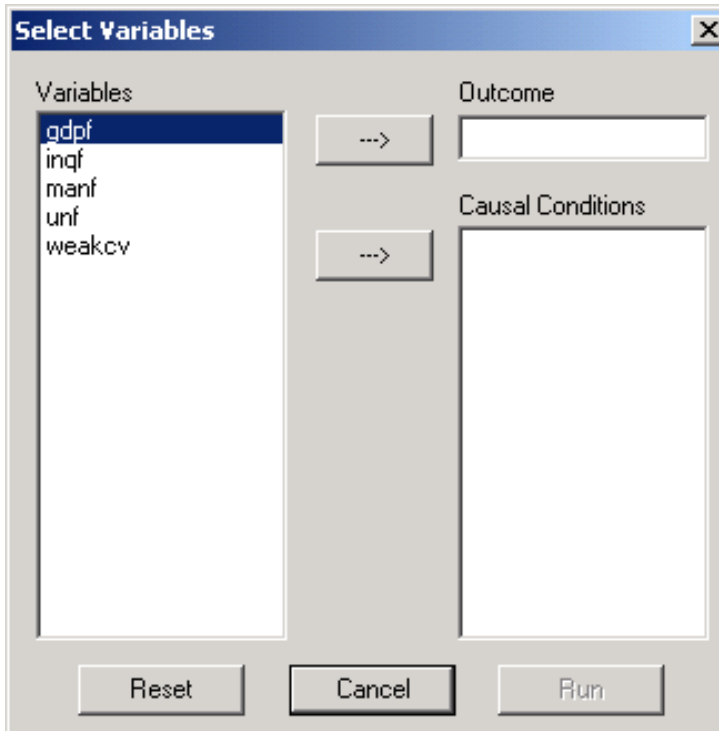
Analysis

The truth table algorithm incorporates a two-stage analytic procedure. The first step consists of creating a truth table from the fuzzy data, which includes specifying the outcome for each configuration and determining which configurations to include in the analysis. The second step involves specifying the causal conditions and outcomes to

minimize. These steps must be performed in conjunction and both must be performed for each separate analysis.

- In order to list the data in the output window, choose:
Analyze
 Fuzzy Sets
 Truth Table Algorithm.....

The following window will open:



- Identify the variable you want to explain and transfer it into the *Outcome* field.
- Choose causal conditions one at a time and click them over to the *Causal Conditions* field.
- Click on the Run button and the following window containing the truth table will appear:

gdpf	inqf	manf	unf	number	weakcv	consist	pre	product
1	1	0	0	3 (25%)		0.950000	0.909091	0.863636
1	0	0	1	2 (41%)		0.736842	0.523810	0.385965
1	0	1	1	2 (58%)		0.666667	0.333333	0.222222
0	1	0	1	1 (66%)		0.764706	0.529412	0.404844
0	1	1	1	1 (75%)		0.812500	0.571429	0.464286
1	0	0	0	1 (83%)		0.928571	0.833333	0.773809
1	1	0	1	1 (91%)		0.789474	0.600000	0.473684
1	1	1	1	1 (100%)		0.722222	0.411765	0.297386
0	0	0	0	0 (100%)		0.909091	0.714286	0.649351
0	0	0	1	0 (100%)		0.733333	0.428571	0.314286
0	0	1	0	0 (100%)		0.916667	0.714286	0.654762
0	0	1	1	0 (100%)		0.785714	0.454546	0.357143
0	1	0	0	0 (100%)		0.928571	0.818182	0.759740

➤ The truth table will have 2^k rows (where k represents the number of causal conditions), reflecting all possible combinations of causal conditions. The 1s and 0s indicate the different corners of the vector space defined by the fuzzy set causal conditions. For each row, a value for each of the following variables is created:

- number* the number of cases with greater than 0.5 membership in that corner of the vector space. Shown in parentheses is the cumulative percentage of cases, beginning with the most populated sector of the vector space
- consist* the degree to which membership in that corner of the vector space is a consistent subset of membership in the outcome. (For crisp sets, this is the proportion of cases in a given crisp truth table row that display the outcome.)
- pre* an alternative measure of consistency for fuzzy sets based on a quasi proportional reduction in error calculation. (In crisp sets this will be equal to *consist*).
- product* the multiplicative product of *consist* and *pre*. This helps identify gaps in the upper ranges of set-theoretic consistency, to help establish a consistency threshold for the outcome.

Note that the column labeled as the outcome (*weakcv* in this example) is blank. It is up to the investigator to determine the outcome for each configuration and to enter it into the spreadsheet using the following procedure.

➤ The researcher must begin by developing a rule for classifying some configurations (vector spaces corners) as relevant and others as irrelevant, based on the number of cases residing in each sector of the vector space defined by the causal conditions. This is accomplished by selecting a frequency threshold based on the number of cases with greater than 0.5 membership in each configuration, as shown in the *number* column. When the total N (number of cases) is relatively small, the frequency threshold should be

1 or 2. When the total N is large, a more substantial threshold should be used. It is very important to examine the distribution of cases across conditions, to identify the most populated sectors of the vector space. In general, the configurations selected should capture *at least* 75-80% of the cases.

➤ Cases can be sorted by their frequency by clicking in the *number* column and choosing
Sort
Descending

➤ After sorting and selecting a threshold, delete all rows that do not meet the threshold. If the cases have been sorted in a descending order according to *number*, click on the first case that falls below the threshold and then choose
Edit
Delete current row to last row

If cases have not been sorted then those cases that do not meet the threshold can be deleted individually by selecting the row and then choosing
Edit
Delete current row

➤ The next step is to distinguish configurations that are consistent subsets of the outcome from those that are not. This determination is made using the measures of set-theoretic consistency reported in the *consist*, *pre*, and/or *product* columns. Values below 0.75 in the *consist* column indicate substantial inconsistency. It is useful to sort the consistency scores in descending order to evaluate their distribution (this should be done *after* removing rows that fail to meet the frequency threshold). Sorting is accomplished by clicking in the *consist*, *pre*, or *product* column and choosing
Sort
Descending

Identify any gaps in the upper range of consistency that might be useful for establishing a consistency threshold. ***Keep in mind that it is always possible to examine several different thresholds and assess the consequences of lowering and raising the consistency cut-off.***

➤ It is now necessary to indicate which configurations exhibit the outcome and which do not. Place a 1 in the outcome column (*weakv* in this example) for each configuration whose consistency level meets and/or exceeds the threshold. Place a 0 in the outcome column for each configuration whose consistency level does not meet the consistency threshold.

➤ Alternatively, one can use the “Delete and code” function to automate this process.
Select
Edit
Delete and code

In the first field, the frequency threshold is selected. The default number of cases is 1, but may be changed by typing the selected frequency threshold into the field. In the second field, the consistency threshold is selected. The default consistency is 0.8, but this may be changed by typing the selected consistency threshold into the field.

Click “OK”. The program will delete rows where the frequency threshold is not met, and will code the outcome as 0 or 1 depending on the selected consistency threshold.

➤ The following window displays the truth table that would appear after:

1. applying a frequency threshold of 1 to the data and eliminating configurations that do not have any observations (8 configurations)
2. selecting a consistency threshold of 0.9 and placing a 1 in the *weakcv* column for configurations with 0.95 consistency or greater (2 configurations) and a 0 for cases with lower consistency (6 configurations)

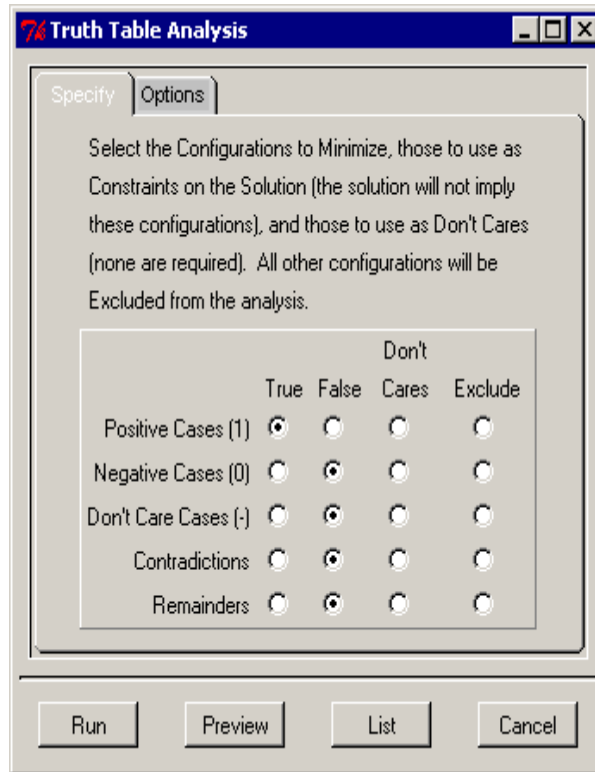
gdpf	inqf	manf	unf	number	weakcv	consist	pre	product
1	1	0	0	3	1	0.950000	0.909091	0.863636
1	0	0	0	1	1	0.928571	0.833333	0.773809
0	1	1	1	1	0	0.812500	0.571429	0.464286
1	1	0	1	1	0	0.789474	0.600000	0.473684
0	1	0	1	1	0	0.764706	0.529412	0.404844
1	0	0	1	2	0	0.736842	0.523810	0.385965
1	1	1	1	1	0	0.722222	0.411765	0.297386
1	0	1	1	2	0	0.666667	0.333333	0.222222

From this point in the procedure, there are two possibilities for analysis: specifying the analysis versus selecting “Standard Analyses.” ***“Standard Analyses” is the recommended choice because this is the only way to generate the “intermediate” solution.***

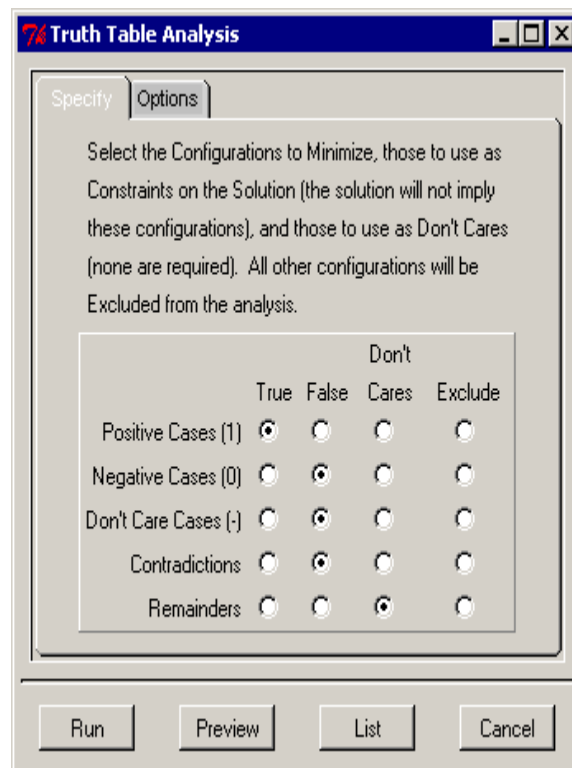
“Specify Analysis” Option

➤ Once the truth table is constructed select Specify Analysis to bring up the Truth Table Analysis Window.

➤ In the Specify panel set Positive cases to True and all the others to False to yield the most complex solution. This window appears as follows:



➤ To derive the most parsimonious solution, set Positive cases to True, Negative Cases, Don't Care, and Contradictions to False, and Remainders to Don't Cares. The window will appear as follows:



➤ All other options (for prime implicants and output frequencies) can be set following the same procedures outlined for crisp sets. *In general, however, these options are not useful when using the truth table algorithm.*

“Standard Analyses” Option

➤ Once the truth table is constructed select “Standard Analyses.” This procedure automatically provides the user with the complex, intermediate, and parsimonious solutions. “Standard Analyses” is recommended over “Specify Analysis.”

Please refer back to the discussion of the “Standard Analyses” procedure for crisp sets. The fuzzy set procedure is parallel. Three solutions are derived, the complex, the parsimonious and the intermediate. Each solution is based on a different treatment of the remainder combinations:

Complex: remainders are all set to false; no counterfactuals

Parsimonious: any remainder that will help generate a logically simpler solution is used, regardless of whether it constitute an “easy” or a “difficult” counterfactual case;

Intermediate: only remainders that are “easy” counterfactual cases are allowed to be incorporated into the solution. The designation of “easy” versus “difficult” is based on user-supplied information regarding the connection between each causal condition and the outcome.

D) Output for “Specify Analysis” Option

Once the truth table has been minimized, the main window will show you the following output. The output shown is for the most complex solution, obtained by using the “Specify Analysis” option as described above.

The first part of the output describes the data. The first two rows indicate the file directory and the model you specified. The next row lists the number of truth table rows (from the coded truth table spreadsheet) read for the analysis. Then, the output displays what kind of algorithm you used, and which configurations were minimized and which were set to constraint. In the example below, the configurations with the outcome (1) were minimized, whereas the configurations without the outcome (0), don’t cares (-), contradictions (C), and remainders (L) were set to constraint (false).

```

fs/QCA
File  Analyze  Graphs  Window  Help

*****
*TRUTH TABLE ANALYSIS*
*****

File:  E:/weak-class-vot.dat
Model:  WEAQCV = f(GDPF, INQF, MANF, UNF)

Rows:      8

Algorithm:  Quine-McCluskey
  True:  1
  0 Matrix:  0-C
Don't Care:  Remainder

```

The last part of the output shows the solution of your analysis. First, frequency and consistency cutoffs are listed. This is followed by the solution. When the “Specify Analysis” option is selected, there will be one Truth Table Solution section. When “Standard Analyses” is selected, three solutions will be provided (complex, parsimonious and intermediate).

In this example, the following solution for the most complex solution was reported:

```

fs/QCA
File  Analyze  Graphs  Window  Help

--- TRUTH TABLE SOLUTION ---
frequency cutoff: 1.000000
consistency cutoff: 0.928571

      raw      unique
      coverage  coverage  consistency
-----
GDPF*manf*unf  0.563380  0.563380  0.952381
solution coverage: 0.563380
solution consistency: 0.952381

```

Weak class voting is a product of high membership in the set of affluent countries, weak membership in the set of countries with a high percentage of workers in manufacturing, and weak membership in the set of countries with strong unions.

The first two lines list the frequency and consistency cutoffs. The consistency cutoff will list the lowest consistency value above the cut-off value specified by the user. Here, 0.9 was given as the consistency cutoff, and the lowest actual value above 0.9 was 0.928571.

The solution provides a line for each separate path to the outcome (in this example only one path exists: GDPF*manf*unf). The output also computes the consistency and coverage for each solution term and the solution as a whole (these computations are discussed below).

The output for the most parsimonious solution of this same truth table is:

```

*****
*TRUTH TABLE ANALYSIS*
*****

File: E:/weak-class-vot.dat
Model: WEAKCV = f(GDPF, INQF, MANF, UNF)

Rows:      8

Algorithm: Quine-McCluskey
  True: 1
  0 Matrix: 0-C
Don't Care: Remainder

--- TRUTH TABLE SOLUTION ---
frequency cutoff: 1.000000
consistency cutoff: 0.928571

      raw      unique
      coverage  coverage  consistency
-----
unf    0.676056  0.676056  0.960000
solution coverage: 0.676056
solution consistency: 0.960000

```

Again, the output displays the algorithm used, and which configurations were minimized and which were set to constraint. In the example above, the configurations with the outcome (1) were minimized, configurations without the outcome (0), don't cares (-), and contradictions (C), were set to constraint (false), and remainders (R) were set to don't care.

The solution indicates only one path to weak class voting. Countries with weak membership in the set of countries with strong unions exhibit weak class voting.

E) Output for "Standard Analyses" Option

Output for the Standard Analysis will look slightly different:

```

fs/QCA
File Analyze Graphs Window Help
*****
*TRUTH TABLE ANALYSIS*
*****
File: C:/Documents and Settings/Sarah/Desktop/Weak_CV.csv
Model: WEAKCV = f(GDPF, INQF, MANF, UNF)

Rows:      10

Algorithm: Quine-McCluskey
  True: 1 |
  0 Matrix: 0
Don't Care: -

--- INTERMEDIATE SOLUTION ---
frequency cutoff: 1.000000
consistency cutoff: 0.928571
Assumptions:
GDPF (present)
INQF (present)
MANF (present)
UNF (present)

           raw      unique
           coverage  coverage  consistency
-----
GDPF*unf   0.619718   0.619718   0.956522
solution coverage: 0.619718
solution consistency: 0.956522

```


Most notably, the complex, intermediate, *and* parsimonious solutions will be displayed. Depicted here is the intermediate solution. The “Assumptions” portion of the output display the options previously selected in the “Intermediate Solution” window; here, each was selected such that when present, the condition should contribute to the outcome.

In this solution, weak class voting is a product of high membership in the set of affluent countries and weak membership in the set of countries with strong unions.

F) Consistency and Coverage

The output includes measures of coverage and consistency for each solution term and for the solution as a whole. Consistency measures the degree to which solution terms and the solution as a whole are subsets of the outcome. Coverage measures how much of the outcome is covered (or explained) by each solution term and by the solution as a whole. These measures are computed by examining the original fuzzy data set in light of the solution (composed of one or more solution terms). The degree to which cases in the original dataset have membership in each solution term and in the outcome form the basis of consistency and coverage measures. More specifically, consider the following data table with three causal conditions (A, B, and C) and an outcome (Y) all measured as fuzzy sets.

Causal Condition Membership			Outcome Membership	Solution Membership			Solution Consistency		
A	B	C	Y	A*B	A*C	A*B + A*C	C _{A*B}	C _{A*C}	C _{A*B + A*C}
.8	.9	.8	.9	.8	.8	.8	.8	.8	.8
.6	.7	.4	.8	.6	.4	.6	.6	.4	.6
.6	.7	.2	.7	.6	.2	.6	.6	.2	.6
.6	.6	.3	.7	.6	.3	.6	.6	.3	.6
.8	.3	.7	.8	.3	.7	.7	.3	.7	.7
.6	.1	.7	.9	.1	.6	.6	.1	.6	.6
.7	.4	.2	.3	.4	.2	.4	.3	.2	.3
.2	.9	.9	.1	.2	.2	.2	.1	.1	.1
.1	.6	.2	.2	.1	.1	.1	.1	.1	.1
.2	.1	.7	.3	.1	.2	.2	.1	.2	.2
.3	.1	.3	.3	.1	.3	.3	.1	.3	.3
.1	.2	.3	.2	.1	.1	.1	.1	.1	.1
Sum:			6.2	4.0	4.1	5.2	3.8	4.0	5.0

The relevant output for this analysis is shown below. The solution is comprised of two terms: $A*B + A*C$. To calculate consistency and coverage, several intermediate values must be calculated first. Membership in the outcome (\sum_Y) is the sum of outcome membership scores across all cases in the data. A case’s membership in each solution term is computed as the minimum of the cases membership in each causal condition of the term. Membership in the first solution term (\sum_{A*B}) is the sum of membership in that solution term across all cases. Similarly, membership in the second solution term (\sum_{A*C})

is the sum of membership in that solution term across all cases. Membership in the solution ($\sum_{A*B + A*C}$) is defined as the maximum of a case's membership across the solution terms.

	raw coverage	uni que coverage	consi stency
A*B+	0. 612903	0. 161290	0. 950000
A*C	0. 645161	0. 193548	0. 975610
sol uti on coverage:	0. 806452		
sol uti on consi stency:	0. 961538		

Consistency measures the degree to which membership in each solution term is a subset of the outcome. Consistency is computed by first computing the consistency of each case. For any solution term, a case is consistent if membership in the solution term is less than or equal to membership in the outcome. If a case's membership in the solution term is greater than its membership in the outcome (i.e., it is inconsistent), then the case is given a score that equals its membership in the outcome. These scores are then summed (yielding $\sum C_{A*B}$) and divided by the sum of memberships in the solution term (\sum_{A*B}). Thus, consistency for the first solution term is $\sum C_{A*B} / \sum_{A*B} = 3.8/4 = .95$ and for the second solution term is $4.0/4.1 = .976$.

Solution Consistency measures the degree to which membership in the solution (the set of solution terms) is a subset of membership in the outcome. The maximum of each case's membership across solution terms ($A*B + A*C$) is compared to membership in the outcome. If membership in the solution is less than or equal to membership in the outcome, then the case is given a score that equals its membership in the solution term. If membership in the solution term is greater than membership in the outcome (i.e., if it is inconsistent), then the case is given the outcome scores (the lower of the two scores). These scores are summed and then divided by the sum of memberships in the solution term ($\sum C_{(A*B + A*C)} / \sum_{(A*B + A*C)}$). The consistency for the solution in this example is $5.0/5.2 = .962$.

Solution coverage measures the proportion of memberships in the outcome that is explained by the complete solution. The consistent membership scores are summed across cases and then divided by the sum of the membership in the outcome: $(\sum C_{(A*B + A*C)} / \sum_Y) = 5.2/6.2 = .806$.

Raw coverage measures the proportion of memberships in the outcome explained by each term of the solution. Raw coverage is computed for each solution term from the original data by dividing the sum of consistent membership in the solution term by the sum of membership in the outcome. Raw coverage for the first solution term is $\sum C_{A*B} / \sum_Y = 3.8/6.2 = .613$ and for the second term is $4.0/6.2 = .645$.

Unique coverage measures the proportion of memberships in the outcome explained solely by each individual solution term (memberships that are not covered by other solution terms). This is computed by first removing the term from the solution and

computing solution coverage. In this example, solution coverage after removing the first solution term (ΣC_{A*B}) is simply ΣC_{A*C} (with n solution terms the reduced solution will contain $n-1$ solution terms). The reduced coverage term is then divided by the full solution coverage and subtracted from the raw coverage to give the unique coverage for the omitted solution term. For the first solution term (ΣC_{A*B}) unique coverage equals: $(\Sigma C_{(A*B + A*C)} / \Sigma Y) - (\Sigma C_{A*C} / \Sigma Y) = (5.0/6.2) - (4.0/6.2) = .161$. Unique coverage for the second term equals $(5.0/6.2) - (3.8/6.2) = .194$.