


CHARLES C. RAGIN

Redesigning Social Inquiry
Fuzzy Sets and Beyond

University of Chicago Press: Chicago and London



CHARLES C. RAGIN is professor of sociology and political science at the University of Arizona, Tucson, positions he has held since 2001. Prior to that, Ragin was professor of sociology at Northwestern University (1981–2001) and Indiana University (1975–1981). He has also taught at the University of Oslo. Ragin has written several books, including *Fuzzy-Set Social Science* (University of Chicago Press, 2000), *Constructing Social Research: The Unity and Diversity of Method* (Pine Forge, 1994), and *The Comparative Method: Moving beyond Qualitative and Quantitative Strategies* (University of California Press, 1987), and he has edited or coedited several more. Ragin has also published more than one hundred articles.

The University of Chicago Press, Chicago 60637
The University of Chicago Press, Ltd., London
© 2008 by The University of Chicago
All rights reserved. Published 2008
Printed in the United States of America

16 15 14 13 12 11 10 09 08 07 1 2 3 4 5

ISBN-13: 978-0-226-70273-5 (cloth)
ISBN-13: 978-0-226-70275-9 (paper)
ISBN-10: 0-226-70273-1 (cloth)
ISBN-10: 0-226-70275-8 (paper)

Ragin, Charles C.
Redesigning social inquiry : fuzzy sets and beyond /
Charles C. Ragin.

p. cm.

Includes bibliographical references and index.

ISBN-13: 978-0-226-70273-5 (cloth : alk. paper)

ISBN-13: 978-0-226-70275-9 (pbk. : alk. paper)

ISBN-10: 0-226-70273-1 (cloth : alk. paper)

ISBN-10: 0-226-70275-8 (pbk. : alk. paper) 1. Social sciences—Research—Methodology. 2. Fuzz sets. I. Title.

H62.R235 2008

300.72—dc22

2008009112

∞ The paper used in this publication meets the minimum requirements of the American National Standard for Information Sciences—Permanence of Paper for Printed Library Materials, ANSI Z39.48–1992.

CONTENTS

Acknowledgments, vii
Introduction, 1

PART I

Set-Theoretic versus Correlational Connections

1: Set Relations in Social Research: Basic Concepts, 13

2: Fuzzy Sets and Fuzzy-Set Relations, 29

3: Evaluating Set Relations: Consistency and Coverage, 44

PART II

Calibration versus Measurement

4: Why Calibrate? 71

5: Calibrating Fuzzy Sets, 85

PART III

Configurations of Conditions versus “Independent” Variables

6: Configurational Thinking, 109

7: Configurational Analysis Using Fuzzy Sets
and Truth Tables, 124

PART IV

Analysis of Causal Complexity versus Analysis of Net Effects

8: Limited Diversity and Counterfactual Cases
coauthored with John Sonnett, 147

9: Easy versus Difficult Counterfactuals, 160

- 10: The Limitations of Net Effects Thinking, 176
11: Net Effects versus Configurations:
An Empirical Demonstration
coauthored with Peer Fiss, 190

References, 213
Index, 221

ACKNOWLEDGMENTS

It is impossible to list all the people who have contributed over the past several years to the ideas that I present in this book. I am sure I will miss many, including some who have been very influential. Perhaps my greatest debt is to unidentified questioners at talks and workshops where I have presented this material. I hope that my response, presented in the pages that follow, is finally adequate.

I owe a deep debt to my colleagues, local and long distance, who have commented on various versions of the ideas presented in this book. I thank Edwin Amenta, Howard Becker, Henry Brady, Ron Breiger, David Byrne, David Collier, Barry Cooper, Gary Goertz, Lane Kenworthy, Bruce Kogut, Jim Mahoney, Michael Minkenberg, Lars Mjølset, Benoit Rihoux, Carsten Schneider, Svend-Erik Skaaning, Steve Vaisey, Claudius Wagemann, and the readers for the University of Chicago Press. Special thanks also go to Sean Davey, who implemented my ideas in easy-to-use software and swatted logic bugs that almost got away, and to Sarah Strand, who helped with the bibliography and the index. I also thank Doug Mitchell and his colleagues at the University of Chicago Press for their steadfast support and encouragement.

A variety of institutions and organizations also have supported this work. Direct and indirect support came from the National Science Foundation, the University of Oslo, and various units of the University of Arizona, including the Udall Center for Studies in Public Policy, the Social and Behavioral Sciences Research Institute, and the departments of Sociology and Political Science. This project was initiated while I was a fellow at the Center for Advanced Study in the Behavioral Sciences in Stanford, California.

in conventional quantitative research). Chapter 8 presents a set-theoretic approach to counterfactual analysis, using truth tables to elaborate the idea of limited diversity. Viewed through this lens, it is clear that counterfactual analysis is almost always an issue in non-experimental social research, regardless of the number of cases examined. I emphasize the theory and knowledge dependence of all social research and criticize conventional quantitative researchers for ignoring both the need for explicit counterfactual analysis and its knowledge-dependent nature. Chapter 9 presents the distinction between “easy” and “difficult” counterfactuals and shows how the incorporation of easy counterfactuals (a process that is implicit in much case-oriented research) can be formalized using set-theoretic methods. Bridging my discussion of counterfactual analysis (chapters 8 and 9) and my empirical demonstration (chapter 11) is an examination in chapter 10 of the limitations of what I call *net effects thinking*, the analytic metatheory that dominates the social sciences today. Chapter 11 concludes the book by providing a demonstration, using a large-*N* data set known as the *Bell Curve* data (Herrnstein and Murray 1994). I present a fuzzy-set analysis of the combinations of individual-level characteristics linked to poverty and contrast these results with a conventional, net-effects analysis of the same data.

PART I

Set-Theoretic versus Correlational Connections

1: Set Relations in Social Research

Basics Concepts

When quantitative social scientists think sets, they usually do not get very far. They think, "OK, nominal-scale variables. I can transform them into dummy variables and use them in linear models." Or perhaps they think, "Hmmm, I've got subpopulations." Qualitative researchers are not much different. They think, "OK, typologies of cases. I can construct (and deconstruct) those." Missing in both views is recognition of the importance of the *analysis of set relations* in social research. Consider that almost all social science theory is verbal and, as such, is formulated in terms of sets and set relations. When a theory states, for example, that "small farmers are risk averse," the claim is set theoretic: small farmers constitute a rough subset of risk-averse individuals. Such statements are usually transformed by social scientists into hypotheses about correlations between variables, which are then evaluated using standard correlational techniques (e.g., multiple regression analysis). This chapter argues that theory formulated in terms of set relations should be evaluated on its own terms, that is, as statements about set relations, not about correlations. In the process, it offers a general overview of set relations in social research.

The Nature of Set Relations

The simplest and most basic set relation is the subset, which is easiest to grasp when it involves nested categories. Dogs are a subset of the set of mammals; Protestants are a subset of the set of Christians, who in turn are a subset of the set of monotheists. These subset relations are straightforward and easy to accept as valid because they are

definitional in nature: dogs have all the characteristics of mammals; the set of Christians is partially constituted by the set of Protestants. These examples also involve conventional, *crisp* sets and thus they are easy to grasp and simple to represent using Venn diagrams. The circle representing the set of dogs, for example, is entirely contained within a larger circle representing the set of mammals. (Of course, many observers would argue that the set of Protestants is not *crisp*, but instead is a truly *fuzzy* set. I address the question of fuzzy sets in chapter 2.)

More important than these simple, definitional subsets are subset relations that describe social phenomena that are connected causally or in some other integral manner. When researchers argue, for example, that “religious fundamentalists are politically conservative,” they are stating, in effect, that they believe that religious fundamentalists form a rough subset of the set of political conservatives, and may even go so far as to argue that their fundamentalism is the cause of their conservatism. Likewise, a researcher who argues that having a strong “civil society” is a necessary or essential part of being a “developed country” implies that the developed countries constitute a consistent subset of those with strong civil societies. In this example, the connection is *constitutive*, as opposed to causal.

When set relations reflect integral social or causal connections and are not merely definitional in nature, they require explication—that is, they are theory and knowledge dependent. Assume, for example, that among third-wave democracies, all those that adopted parliamentary governments soon failed. Thus, third-wave democracies with parliamentary governments form a subset of failed third-wave democracies. Were the failures just bad luck, a coincidence? Or did a causal or some other kind of integral connection exist between third-wave democracies adopting a parliamentary form of government and their subsequent failure? The set-theoretic connection in this example is not definitional; it must be explicated in some way. This type of set relation, the kind that is central to almost all social science theorizing, is the main focus of this chapter and this book.

Set-Theoretic Connections Are Asymmetrical

An important aspect of set-theoretic connections, as opposed to correlational connections, is that they are *asymmetrical*. For example, the fact that there are many political conservatives who are not religious fundamentalists does not in any way challenge the claim that religious fundamentalists are politically conservative. In another example, if my theory states that the developed countries are democratic, in essence I am stating that the set of developed countries is a subset of the set of democratic countries. The fact that there are less-developed countries that are also democracies does not undermine my set-theoretic claim. Of course, such cases *do* undermine the correlation between development and democracy—that is, they would count against my argument if it had been formulated symmetrically. The fully symmetric version would be, “the developed countries are democratic, and the less developed countries are not democratic.” However, this reformulation of the argument extends it in ways that may not be warranted or intended. As originally stated, the argument is asymmetric, as are set-theoretic formulations in general.

Set-theoretic arguments are often erroneously reformulated as correlational hypotheses. This mistake is, in fact, one of the most common in all of contemporary social science. For example, a theory may claim that because of the many external vagaries faced by newly formed democracies, third-wave democracies adopting parliamentary governments are unlikely to endure. After reading this argument, the conventional social scientist would try to test it by examining the correlation between “parliamentary government” and “failure” using data on third-wave democracies. Suppose, again, that the set-theoretic evidence supports the theory; that is, third-wave democracies adopting parliamentary governments are a subset of failed third-wave democracies. Despite this clear connection, the correlation between “parliamentary form” and “failure” still might be relatively weak, due to the fact that there are many other paths to failure and thus there are instances of failed democracies with presidential

or other forms of nonparliamentary government. The set-theoretic claim that “third-wave democracies with parliamentary governments fail” is not refuted by these cases. However, these nonparliamentary paths to failure seriously undermine the correlation between “parliamentary form” and “failure.”

Consider the “democratic peace” argument that democracies do not go to war against each other. This statement is essentially a claim that country dyads in which both parties are democratic constitute a perfect (or near-perfect) subset of nonwarring country dyads. Of course, the rate of warring may be very low both in the set of democratic dyads and in the set of dyads in which at least one of the parties is not a democracy. The point of the argument is not the *difference* between these two rates but that the rate of warring is *zero or close to zero* in the set of democratic dyads. The fact that democratic dyads constitute a near-perfect subset of nonwarring dyads signals that this arrangement (international relations between democracies) may be *sufficient* for peaceful coexistence. Of course, many other paths may be taken to peaceful coexistence, and the correlation between “democratic dyad” and “nonwarring” may be weak because of these many alternate paths.

The key difference between correlational and set-theoretic connections is illustrated in tables 1.1 and 1.2. Table 1.1 shows a pattern of results consistent with the existence of a correlational connection between parliamentary government and failure among third-wave democracies. The first column shows the tendency for nonparliamentary governments to survive; the second column shows the tendency for the parliamentary governments to fail. While very satisfying from a correlational viewpoint, this table would be unsatisfying to a researcher interested in set-theoretic connections, for there are no connections in the table that could be described as explicit or consistent. Table 1.2, however, would be of great interest to this researcher because it shows a consistent connection between parliamentary form and failure—all sixteen cases with this governmental form failed, as shown in the second column of this table. While significant to the researcher interested in set-theoretic connections, this table would disappoint the

Table 1.1: A correlational connection

	Presidential form	Parliamentary form
Third-wave democracy failed	7	11
Third-wave democracy survived	17	5

Table 1.2: A set-theoretic connection

	Presidential form	Parliamentary form
Third-wave democracy failed	15	16
Third-wave democracy survived	9	0

researcher interested in correlational connections, for the correlation between form of government and survival versus failure is relatively weak.

To summarize, set relations in social research (1) involve causal or other integral connections linking social phenomena (i.e., are not merely definitional), (2) are theory and knowledge dependent (i.e., require explication), (3) are central to social science theorizing (because theory is primarily verbal in nature, and verbal statements are often set theoretic), (4) are asymmetric (and thus should not be reformulated as correlational arguments), and (5) can be very strong despite relatively modest correlations (as illustrated in tables 1.1 and 1.2).

Two Important Types of Set-Theoretic Relations

Case-oriented researchers—and qualitative researchers more generally—are centrally concerned with the analysis of set relations, which is evident in their efforts to identify *explicit* connections (Ragin and Rihoux 2004). They rarely see their work in set-theoretic or formal terms, however, so I sometimes refer to them as *implicit* Booleans. For example, case-oriented researchers often seek to identify *commonalities* across a set of cases, usually while focusing on a relatively small number of purposefully selected cases (e.g., Vaughan 1986). Why look for commonalities? They are suggestive of important empirical


connections. For example, suppose all (or almost all) of the anorectic teenage girls in my study say that they get a sense of accomplishment from their food and body practices. A light bulb goes on in my head, and I explore the connection further. These kinds of commonalities are, in fact, set-theoretic relations. Consider another example: an examination of social revolutions indicates that some form of prior state breakdown occurred in every case (Skocpol 1979). The evidence indicates that a set-theoretic relation exists between state breakdown and social revolution. In this instance, it might be reasonable to speculate that the set of social revolutions is a subset of the set of state breakdowns and that an important causal link exists between the two.

Two general analytic strategies involve searching for commonalities. The first strategy is to examine cases sharing a given outcome (e.g., consolidated third-wave democracies) and attempt to identify their shared causal conditions (e.g., the possibility that they share the presidential form of government).¹ The second strategy is to examine cases sharing a specific causal condition or, more commonly, a specific combination of causal conditions, and assess whether these cases exhibit the same outcome (e.g., do the countries that combine party fractionalization, a weak executive, and a low level of economic development all suffer democratic breakdown?). Both strategies are set theoretic in nature. The first is an examination of whether instances of a specific outcome constitute a subset of instances of a cause. The second is an examination of whether instances of a specific causal condition or combination of causal conditions constitute a subset of instances of an outcome. These two strategies are illustrated with Venn diagrams in figure 1.1.

Both strategies are methods for establishing *explicit* connections. If it is found, for example, that all (or nearly all) consolidated third-wave democracies have presidential systems, then an explicit connection

1. The term *causal condition* is used generically here and elsewhere in this book to refer to an aspect of a case that is relevant in some way to the researcher's account or explanation of some outcome.

A. Identify causal conditions shared by instances of an outcome



B. Assess whether cases with the same causal conditions share the same outcome




Figure 1.1 Venn diagram representing two different kinds of case-oriented research

has been established between presidentialism and consolidation.² Likewise, if it is found that all (or nearly all) third-wave democracies that share a low level of economic development, party fractionalization, and a weak executive failed as democracies, then an explicit connection has been established between this combination of conditions and democratic breakdown. As previously noted, establishing explicit

2. Neither of these two strategies expects or depends on demonstrations of *perfect* set-theoretic relations. For example, if *almost all* (as opposed to *all*) instances of democratic consolidation involved presidential systems, then the researcher would no doubt accept this as evidence of an integral connection between presidentialism and democratic consolidation. Specific procedures for probabilistic assessment of set-theoretic patterns using benchmarks are presented in Ragin (2000).

connections is not the same as establishing correlations. For example, assume that the survival rate for third-wave democracies with presidential systems is 60 percent, while the survival rate for third-wave democracies with parliamentary systems is 35 percent. Clearly, a correlation exists between these two aspects conceived as variables (presidential versus parliamentary system, survival versus failure). However, the evidence does not come close to approximating a set-theoretic relation. Thus, in this instance evidence would show a correlational connection, but not an explicit connection between presidential systems and democratic survival.

As explained in Ragin (2000), the first analytic strategy—identifying causal conditions shared by cases with the same outcome—is appropriate for the assessment of necessary conditions. The second strategy—examining cases with the same causal conditions to see if they also share the same outcome—is suitable for the assessment of sufficient conditions, especially sufficient combinations of conditions. Establishing conditions that are necessary or sufficient is a long-standing interest of many researchers, especially those working at the macro-social or macropolitical level (see, e.g., Goertz and Starr 2002). However, it is important to point out that the use of set-theoretic methods to establish explicit connections does not necessarily entail the use of the concepts or the language of necessity and sufficiency, or any other language of causation. A researcher might observe, for example, that instances of democratic breakdown are all former colonies without drawing any causal connection from this observation. In a simpler example, colleagues might “act out” only in faculty meetings, but that does not mean that analysts must therefore interpret faculty meetings as a necessary condition for acting out. Demonstrating explicit connections is central to social science, whether or not there is interest in necessary or sufficient causation or any other kind of causation.

How Correlational Methods Sometimes Miss Connections

The mismatch between correlational methods and the study of explicit connections is clearly visible in the simplest form of variable-

Table 1.3: Conventional cross-tabulation of presence/absence of an outcome against presence/absence of a causal condition

		Causal condition absent		Causal condition present	
Outcome present		cell 1: cases here undermine researcher's argument	cell 2: cases here support researcher's argument		
Outcome absent		cell 3: cases here support researcher's argument	cell 4: cases here undermine researcher's argument		

oriented analysis, the 2×2 cross-tabulation of the presence/absence of an outcome against the presence/absence of a hypothesized cause, as illustrated in table 1.3. The correlation focuses simultaneously and equivalently on the degree to which instances of the cause produce instances of the outcome (the number of cases in cell 2 relative to the sum of cases in cells 2 and 4) and on the degree to which instances of the absence of the cause are linked to the absence of the outcome (the number of cases in cell 3 relative to the sum of cases in cells 1 and 3).³ (Alternatively, it could be stated that the correlation focuses simultaneously and equivalently on the degree to which instances of the outcome are linked to instances of the cause and on the degree to which instances of the absence of the outcome are linked to instances of the absence of the cause.) The central point is that the correlation is an omnibus statistic that rewards researchers for producing an abundance of cases in either cells 2 or 3 and penalizes them for depositing cases in either cell 1 or 4. Thus, it is a good tool for studying general cross-case tendencies.

A researcher who is interested in explicit connections, however, is interested in only specific components of the information that is pooled and conflated in a correlation. For example, researchers interested in causal conditions shared by instances of an outcome would focus on cells 1 and 2 of table 1.3. Their goal would be to identify relevant causal conditions that deposit as few cases as possible (ideally none) in cell 1. Likewise, researchers interested in whether cases that

3. I use the term *correlation* generically here to refer to the examination of the strength of the association between two variables, and not as a specific reference to Pearson's r or the calculations used to produce Pearson's r .

are similar with respect to causal conditions experience the same outcome would focus on cells 2 and 4. Their goal would be to identify relevant combinations of causal conditions that deposit as few cases as possible (ideally none) in cell 4. It is clear from these examples that the correlation has two major shortcomings when viewed from the perspective of explicit connections: (1) it attends only to relative differences (e.g., relative survival rates for presidential versus parliamentary systems), and (2) it conflates different kinds of causal assessment. Notice further that a cell that is very important in correlational analysis, cell 3—where neither the cause nor the outcome is present—is not directly relevant to the assessment of either of the two types of explicit connections.

Thus, the study of explicit connections involves a *decomposition* of the most basic unit of variable-oriented analysis—the correlation. This decomposition makes it possible to employ qualitative research strategies that, I have argued, are fundamentally set theoretic in nature: (1) studying cases with the same outcome and in order to identify their causally relevant features and (2) studying cases with the same combination of causally relevant conditions in order to see if they exhibit the same outcome.

It is important to point out that the correlation is not simply a bivariate statistic. It is the cornerstone of most forms of conventional variable-oriented social research, including some of the most sophisticated forms of quantitative analysis available today. A matrix of bivariate correlations, along with the means and standard deviations of the variables included in the correlation matrix, is all that is needed to compute complex regression analyses, factor analyses, and even structural equation models. In essence, these varied techniques offer diverse ways of representing the bivariate correlations in a matrix and the various partial relations (e.g., the net effect of an independent variable in a multiple regression) that can be constructed using formulas based on three or more bivariate correlations. Because they rely on the bivariate correlation as the cornerstone of empirical analysis, these sophisticated quantitative techniques eschew the study of explicit connections, as described here. This underlying, fundamental

shortcoming of the correlation is at the root of the rejection of correlational methods by many scholars who conduct qualitative and case-oriented research.

Qualitative Comparative Analysis and Explicit Connections

In contrast to correlational techniques, qualitative comparative analysis (QCA) is grounded in set theory and thus is ideally suited for studying explicit connections, such as those sketched in figure 1.1. An especially useful feature of QCA is its capacity for analyzing complex causation, defined as a situation in which an outcome may follow from several different combinations of causal conditions, that is, from different causal “recipes.” For example, a researcher may have good reason to suspect that there are several different recipes for the consolidation of third-wave democracies. By examining the fate of cases with different configurations of causally relevant conditions, it is possible, using QCA, to identify the decisive recipes and thereby unravel causal complexity.

The key tool for analyzing causal complexity using QCA is the truth table, a tool that allows structured, focused comparisons (George 1979). Truth tables list the logically possible combinations of causal conditions and the empirical outcome associated with each configuration.⁴ Thus, they directly implement the second type of explicit connection described above. For example, based on theoretical and substantive knowledge, a scholar might argue that a key recipe for democratic consolidation involves the following combination of conditions: a presidential form of government, a strong executive, a low level of party fractionalization, and a noncommunist past. Table 1.4 illustrates the truth table operationalizing this argument. With

4. It is important to point out that the procedures described here are not dependent on the use of dichotomies. Truth tables can be built from fuzzy sets (with set membership scores ranging from 0 to 1) without dichotomizing the fuzzy scores. These procedures take full advantage of the graded membership scores central to the fuzzy-set approach (see chapters 3 and 7 of this book and also Ragin 2000, 2004a; Rihoux and Ragin 2008).

Table 1.4: Hypothetical truth table showing causal conditions relevant to democratic consolidation

Presidential form	Strong executive	Low party fractionalization	Noncommunist	Consolidated
no	no	no	no	—
no	no	no	yes	no
no	no	yes	no	no
no	no	yes	yes	—
no	yes	no	no	no
no	yes	no	yes	no
no	yes	yes	no	—
no	yes	yes	yes	yes
yes	no	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	yes

Note: The “—” entries indicate that there are no empirical cases with the combination of conditions listed in the row.

four causal conditions, sixteen combinations of conditions (causal configurations) are logically possible. In more complex analyses, the rows (representing combinations of causal conditions) may be quite numerous, for the number of causal combinations is an exponential function of the number of causal conditions (number of causal combinations = 2^k , where k is the number of causal conditions).

The use of truth tables to unravel causal complexity is described in detail elsewhere (e.g., Ragin 1987, 2000; De Meur and Rihoux 2002). The essential point is that the truth table elaborates and formalizes one of the key analytic strategies of comparative research—examining cases sharing specific combinations of causal conditions to see if they share the same outcome. Indeed, the main goal of truth table analysis is to identify explicit connections between combinations of

causal conditions and outcomes. By listing the different logically possible combinations of conditions, it is possible to assess not only the sufficiency of a specific recipe (e.g., the recipe presented in the last row of table 1.4, with all four causal conditions present) but also the sufficiency of the other logically possible combinations of conditions that can be constructed from these causal conditions. For example, if the cases with all four conditions present experience democratic consolidation and the cases with three of four conditions present (and one absent) also experience consolidation, then the researcher can conclude that the causal condition that varies across these two combinations is irrelevant to the recipe. The key ingredients for the outcome are the remaining three conditions. Various techniques and procedures for logically simplifying patterns in truth tables, in addition to the simple one just described, are detailed in Ragin (1987, 2000), De Meur and Rihoux (2002), and Rihoux and Ragin (2008).

Often the move from a hypothesized recipe to a truth table stimulates a reformulation or an expansion of a recipe, based on an examination of relevant cases. For example, suppose the truth table revealed substantial inconsistency in the last row—that is, suppose some cases in the last row failed to consolidate, in addition to the several that did. This inconsistency in outcomes signals to the investigator that more in-depth study of cases is needed. By comparing the cases in this row failing to consolidate with those that did consolidate, it would be possible to elaborate the recipe. Suppose this comparison revealed that the cases that failed to consolidate all had severe elite divisions. This ingredient (absence of elite divisions) could then be added to the recipe, and the truth table could then be respecified with five causal conditions (and thus thirty-two rows).

The task of truth table refinement is demanding, for it requires in-depth knowledge of cases and many iterations between theory, cases, and truth table construction. In effect, the truth table disciplines the research process, providing a framework for comparing cases as configurations of similarities and differences while exploring patterns of consistency and inconsistency with respect to case outcomes.

Looking Ahead

The set-theoretic principles described here provide the foundation for the techniques of social inquiry presented in chapters that follow. One limitation of the set-theoretic principles described in this chapter is that they involve only crisp (Boolean) sets and thus may seem crude. Indeed, one reason that social scientists disdain set-theoretic analysis is the perception that it is restricted to nominal-scale variables. Chapter 2 addresses this limitation, showing that key set-theoretic principles can be applied as well to fuzzy sets, which scale degree of membership in sets to values ranging from 0.0 to 1.0.

Practical Appendix: Constructing Truth Tables

Using fuzzy-set qualitative comparative analysis (fsQCA) (Ragin, Drass, and Davey 2007), which can be downloaded from <http://www.fsqca.com>, it is a simple matter to construct a crisp truth table from dichotomous data. The assumption here is that the researcher has a simple data set composed of binary variables, coded 1 for “present” and 0 for “absent” (the construction of truth tables from fuzzy sets is addressed in chapter 7). The goals of crisp truth table construction are to (1) examine the distribution of cases across the logically possible combinations of a given set of dichotomous causal conditions and (2) examine the degree to which cases with each combination of causal conditions agree with respect to a given outcome.

1. Create the data set. This task can be accomplished using fsQCA, which includes procedures for inputting data directly and for importing data sets from other programs (e.g., comma-delimited files from Excel or tab-delimited files from SPSS). The imported files must have simple variables names (with no embedded spaces or punctuation) as the first line of the data set. Missing data should be entered as blanks.

2. Once the data set has been inputted or imported and appears in the data spreadsheet window, click *Analyze*, then *Crisp Sets*, and then *Truth Table Algorithm*. A dialogue box labeled *Select Variables* will open, which allows specification of the outcome and the causal condi-

tions. A box can also be checked for analyzing the negation (reverse) of the chosen outcome.

3. After specifying the outcome and the causal conditions, click *Run*, and fsQCA will generate the full truth table for the specified outcome using the specified causal conditions. A separate window with the truth table spreadsheet opens.

4. The first item of interest is the *number* column, which shows the distribution of cases across causal combinations. The truth table is first presented with the causal combinations sorted according to frequency, along with the cumulative percentage of cases (shown in the *number* column). The information in this column should be used to select any frequency threshold that might be used as a cut-off value. When the total number of cases in a study is small, the threshold should be at least one case, and truth table rows with no cases (*number* = 0) should be deleted. When the total number of cases is large, however, a higher threshold may be used, to allow for measurement and assignment error or to generate a “coarse grained” analysis. To delete rows, simply click on the first (top-most) row to be deleted, click *Edit*, and then click *Delete current row to last row*.

5. The second main item of interest is the set-theoretic consistency scores, shown in the *consistency* column. With crisp sets, this calculation is simply the proportion of cases in a given row that display the outcome in question. A score of 1.0 (or close to 1.0) indicates high consistency—that the cases in the row agree in displaying the outcome. A score of 0.0 (or close to 0.0) indicates that the cases in the row agree in *not* displaying the outcome. With crisp sets, consistency scores in the middle (0.30 to 0.70) indicate that the cases in a given row are strongly divided with respect to presence/absence of the outcome.

Consistency scores have two main uses. They can be used to code the outcome column in the truth table, which is done manually by entering 1s and 0s in the column labeled with the outcome name, or they can be used to guide further research. Suppose, for example, that several rows have consistency scores indicating that the cases are contradictory—that is, many display the outcome and many do not. By

identifying these cases and studying them closely, it is often possible to specify a causal condition that can be added to the truth table in order to resolve contradictions. The researcher can then re-specify the truth table, including this additional condition. Further details on the refinement and use of truth tables constructed from crisp sets can be found in Ragin (1987).

2: Fuzzy Sets and Fuzzy-Set Relations

Many of the phenomena that interest social scientists vary by level or degree. For example, while it is clear that some countries are democracies and some are not, many in-between cases can be found as well. These latter countries are not fully in the set of democracies, nor are they fully excluded from this set. The fact that so many of the things that interest social scientists do not fit neatly into crisp sets may seem to nullify all the good reasons sketched in chapter 1 for analyzing social phenomena in terms of set relations. Does it make sense, for example, to think about the developed countries as a subset of the democratic countries if both development and democracy are measured on fine-grained, interval scales? One reason social scientists are reluctant to study social phenomena in terms of set relations is that they think that the study of set relations is restricted to nominal-scale measures. Not only are such scales considered "primitive," but interval and ratio scales that have been recoded to nominal scales (and thus "downgraded") are almost always suspect. Has a researcher selected cut-points in a biased way, to favor a particular conclusion?

Fortunately, a well-developed mathematical system is available for addressing degree of membership in sets: fuzzy-set theory (Zadeh 1965). Fuzzy sets are especially powerful because they allow researchers to calibrate partial membership in sets using values in the interval between 0.0 (nonmembership) and 1.0 (full membership) without abandoning core set theoretic principles and operations (e.g., the subset relation). As explained in chapter 1, set relations are central to social science theory, yet the assessment of set relations is outside the scope of conventional correlational methods.

The Nature of Fuzzy Sets

Fuzzy sets are simultaneously qualitative and quantitative, for they incorporate both kinds of distinctions in the calibration of degree of set membership. Thus, fuzzy sets have many of the virtues of conventional interval- and ratio-scale variables, but at the same time they permit qualitative assessment. Consider an example: The United States might receive a membership score of 1.0 (full membership) in the set of developed countries but a score of only 0.9 (slightly less than full membership) in the set of democratic countries, especially in the wake of its performance in the 2000 presidential election. A membership score of 1.0 indicates full membership in a set; scores close to 1.0 (e.g., 0.8 or 0.9) indicate strong but not quite full membership in a set; scores less than 0.5 but greater than 0.0 (e.g., 0.2 or 0.3) indicate that objects are more "out" than "in" a set, but still weak members of the set; a score of 0.0 indicates full nonmembership in a set. The 0.5 score is also qualitatively anchored, for it indicates the point of maximum ambiguity (i.e., fuzziness) in the assessment of whether a case is more in or out of a set.

A fuzzy set can be seen as a continuous variable that has been purposefully calibrated to indicate degree of membership in a well-defined and specified set (see chapters 4 and 5). Such calibration is possible only through the use of theoretical and substantive knowledge, which is essential to the specification of the three qualitative breakpoints (full membership, full nonmembership, and maximum ambiguity). For example, cases in the lower ranges of a conventional continuous variable may all be fully out of the set in question, with fuzzy membership scores truncated to 0.0, while cases in the upper ranges of this same continuous variable may be all fully in the set, with fuzzy membership scores truncated to 1.0.

For illustration of the general idea of fuzzy sets, consider a simple three-value set that allows cases to be in the gray zone between in and out of a set. As shown in table 2.1, instead of using only two scores, 0.0 and 1.0, three-value logic adds a third value, 0.5, to identify objects that are neither fully in nor fully out of the set in question (compare

Table 2.1: Crisp versus fuzzy sets

Crisp set	Three-value fuzzy set	Four-value fuzzy set	Six-value fuzzy set	"Continuous" fuzzy set
1 = fully in	1 = fully in	1 = fully in	1 = fully in	1 = fully in
		0.67 = more in than out	0.8 = mostly but not fully in	Degree of membership is more "in" than "out"; $0.5 < X_i < 1$
	0.5 = neither fully in nor fully out	0.33 = more out than in	0.6 = more or less in	0.5 = cross-over: neither in nor out (maximum ambiguity)
0 = fully out	0 = fully out	0 = fully out	0 = fully out	Degree of membership is more "out" than "in"; $0 < X_i < 0.5$

columns 1 and 2 of table 2.1). This three-value set is a rudimentary fuzzy set. A more elegant but still simple fuzzy set uses four numerical values, as shown in column 3 of table 2.1. The four-value scheme uses the numerical values 0.1, 0.67, 0.33, and 0.0 to indicate "fully in," "more in than out," "more out than in," and "fully out," respectively. The four-value scheme is especially useful in situations where researchers have a substantial amount of information about cases, but the evidence is not systematic or strictly comparable from case to case. A more fine-grained fuzzy set uses six values, as shown in column 4 of table 2.1. Like the four-value fuzzy set, the six-value fuzzy set utilizes two qualitative states (fully out and fully in). The six-value fuzzy set inserts two intermediate levels between fully out and the crossover

point ("mostly out" = 0.2 and "more or less out" = 0.4) and two intermediate levels between the crossover point and fully in ("more or less in" = 0.6 and "mostly in" = 0.8).

At first glance, the four-value and six-value fuzzy sets might seem equivalent to ordinal scales. In fact, they are fundamentally different from such scales. An ordinal scale is a mere ranking of categories, usually without reference to such criteria as set membership. When constructing ordinal scales, researchers do not peg categories to degree of membership in sets; rather, the categories are simply arrayed relative to each other, yielding a rank order. For example, a researcher might develop a six-level ordinal scheme of country wealth, using categories that range from destitute to super rich. It is unlikely that this scheme would translate automatically to a six-value fuzzy set, with the lowest rank set to 0.0, the next rank to 0.2, and so on (as in column 4 of table 2.1). Assume the relevant fuzzy set is the set of rich countries. The lowest two ranks of the ordinal variable might both translate to fully out of the set of rich countries (fuzzy score = 0.0). The next rank up might translate to 0.3 rather than 0.2. The top two ranks might translate to fully in (fuzzy score = 1.0), and so on. In short, the specific translation of ordinal ranks to fuzzy membership scores depends on the fit between the specific content of the ordinal categories and the researcher's conceptualization and labeling of the fuzzy set. This point underscores the fact that researchers must use substantive and theoretical knowledge to calibrate membership in fuzzy sets. Calibration of degree of membership in sets should be purposeful and thoughtful, never mechanical.¹

Finally, a continuous fuzzy set permits cases to take values anywhere in the interval from 0.0 to 1.0, as shown in the last column of table 2.1. The continuous fuzzy set, like all fuzzy sets, utilizes the two qualitative states (fully out and fully in) and also uses the crossover point to distinguish cases that are more out from those that are more in. As an example of a continuous fuzzy set, consider membership in

1. Specific techniques of fuzzy-set calibration are discussed in detail in chapter 5.

the set of rich countries, based on gross national product (GNP) per capita. The translation of this variable to fuzzy membership scores is neither automatic nor mechanical. It would be a serious mistake, for instance, to score the poorest country 0, score the richest country 1, and then array all the other countries between 0 and 1, depending on their positions in the range of GNP per capita values. Likewise, it would be a serious mistake to base fuzzy membership scores on the rank order of GNP per capita values. Instead, the first task in this translation would be to specify three important qualitative anchors: the point on the GNP per capita distribution at which full membership is reached (i.e., definitely a rich country), the point at which full nonmembership is reached (i.e., definitely not a rich country), and the point of maximum ambiguity in considering whether a country is more in or more out of the set of rich countries (a membership score of 0.5, the crossover point). When specifying these qualitative anchors, the investigator should have an explicit rationale for each breakpoint.

Qualitative anchors make it possible to distinguish between relevant and irrelevant variation. Variation in GNP per capita among the unambiguously rich countries is *not* relevant to degree of membership in the set of rich countries, at least from the perspective of fuzzy sets. If a country is unambiguously rich, then it is accorded full membership. Similarly, variation in GNP per capita among the unambiguously not-rich countries is also *not* relevant to membership in the set of rich countries. Thus, in research using fuzzy sets, it is not enough simply to develop scales that show the positions of cases relative to each other (e.g., a conventional index of country wealth such as GNP per capita). It is also necessary to use substantive and theoretical knowledge to map the links between specific scores on continuous variables (e.g., an index of wealth) and specific fuzzy-set membership scores (e.g., full membership in the set of rich countries). It follows that when a researcher reconceptualizes and relabels a set (e.g., shifting the focus from the set of rich countries to the set of "middle income countries"), the membership scores change accordingly, even though the underlying index variable (e.g., GNP per capita) may be the same.

Using Fuzzy Sets: The Basics

When using fuzzy sets to assess set-theoretic relations, both the outcome and the causal conditions can be represented in terms of membership scores.² Consider, for example, the first five columns of table 2.2, which show a simple data matrix containing fuzzy membership scores. This data set, which is used in the examples that follow, addresses class voting in the advanced industrial democracies. In this example, the outcome of interest is the degree of membership in the set of countries with weak class voting (W). This fuzzy set was constructed from survey evidence, compiled by Paul Nieuwbeerta (1995), covering the post-World War II era. While levels of class voting have generally declined across the advanced industrial countries, the rank order of these countries with regard to levels of class voting has remained relatively stable over time. This analysis focuses on the conditions linked to persistently low levels of class voting. The causal conditions used in this example are (1) degree of membership in the set of countries with strong unions (U), (2) degree of membership in the set of countries with a high percentage of workers employed in manufacturing (M), (3) degree of membership in the set of highly affluent countries (A), and (4) degree of membership in the set of countries with substantial income inequality (I). Strong unions and manufacturing employment tend to strengthen class voting, while affluence and inequality tend to undermine it. All fuzzy sets used in this analysis are six-value sets and are based on general characterizations of these countries over the post-World War II period. While finer gradations are possible with these data (as in column 5 of table 2.1), the intent here is to demonstrate operations on fuzzy sets with a simple data set.³

2. Crisp-set causal conditions can be included along with fuzzy-set causal conditions in a fuzzy-set analysis.

3. The primary goal here is to illustrate fuzzy-set principles. Accordingly, this presentation does not focus on how these fuzzy sets were calibrated or even on the issue of which causal conditions might provide the best possible specification of the social structural circumstances linked to persistently low levels of class voting. Instead, the focus is on practical procedures.

Table 2.2: Fuzzy-set data on class voting in the advanced industrial societies

Country	Weak class voting (W)	Affluent (A)	Income inequality (I)	Manufacturing (M)	Strong unions (U)	-M	A-M	A+M
Australia	0.6	0.6	0.2	0.4	0.6	0.6	0.6	0.8
Belgium	0.6	0.6	0.2	0.2	0.8	0.8	0.6	0.8
Denmark	0.2	0.6	0.4	0.2	0.2	0.8	0.6	0.8
France	0.8	0.6	0.8	0.2	0.4	0.6	0.6	0.8
Germany	0.6	0.6	0.6	0.4	0.8	0.6	0.6	0.6
Ireland	0.8	0.2	0.8	0.8	0.2	0.8	0.2	0.2
Italy	0.6	0.4	0.8	0.6	0.6	0.8	0.4	0.8
Netherlands	0.8	0.6	0.4	0.2	0.4	0.8	0.6	0.8
Norway	0.2	0.6	0.4	0.6	0.8	0.4	0.4	0.6
Sweden	0.0	0.8	0.4	0.8	1.0	0.2	0.2	0.8
United Kingdom	0.4	0.6	0.6	0.8	0.6	0.2	0.2	0.6
United States	1.0	0.8	0.8	0.4	0.2	0.6	0.6	1.0

Three common operations on fuzzy sets are set negation, set intersection (logical *and*), and set union (logical *or*). The discussion of these three operations provides important background knowledge for understanding how to work with fuzzy sets.

Negation. Like conventional crisp sets, fuzzy sets can be negated. With crisp sets, negation switches membership scores from 1.0 to 0.0 and from 0.0 to 1.0. The negation of the crisp set of democracies, for example, is the crisp set of not-democracies. This simple mathematical principle holds in fuzzy algebra as well, but the relevant numerical values are not restricted to the two Boolean values 0.0 and 1.0; rather, they extend to values between 0.0 and 1.0. To calculate the membership of a case in the negation of fuzzy set *M*, simply subtract its membership in set *M* from 1.0, as follows:

$$\text{(membership in set } \sim M) = 1.0 - \text{(membership in set } M)$$

or

$$\sim M = 1.0 - M$$

where \sim signals negation. Thus, for example, the United States has a membership score of 0.4 in the set of countries with high manufacturing employment; therefore, it has a score of 0.6 in the set of cases with not-high manufacturing employment. For further illustration, examine the sixth data column of table 2.2, which shows the negated membership scores of set *M* (high manufacturing employment) for all twelve countries. The negated set is labeled $\sim M$, for not-high manufacturing employment.

Logical and. Compound sets are formed when two or more sets are combined, an operation commonly known as *set intersection*. A researcher interested in the fate of class voting in relatively inhospitable settings might want to draw up a list of countries that combine not-high manufacturing employment ($\sim M$) with "highly affluent" (*A*). Conventionally, these countries would be identified using crisp sets by cross-tabulating the two dichotomies (not-high versus high manufacturing employment and highly affluent (not-high not-highly affluent) and seeing which countries are in the not-high manufacturing/highly affluent cell of this 2×2 table. This cell, in

effect, would show the cases that exist in the intersection of the two crisp sets.

With fuzzy sets, logical *and* is accomplished by taking the minimum membership score of each case in the sets that are combined. For example, if a country's membership in the set of countries with not-high manufacturing employment is 0.6 and its membership in the set of highly affluent countries is 1.0, its membership in the set of countries that combine these two traits is the lesser of these two scores, 0.6. A score of 0.6 indicates that this case is still more in than out of the intersection. For further illustration of this principle, consider the seventh data column of table 2.2, which demonstrates the operation of logical *and*. This column shows the intersection of the $\sim M$ (not-high manufacturing) and *A* (highly affluent) sets, yielding membership in the set of countries that combine these two traits. The algebraic expression for this intersection is $A \cdot \sim M$; the midlevel dot is used to indicate set intersection (combinations of aspects).

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 have *either* not-high manufacturing employment ($\sim M$) or high affluence (*A*), based on the conjecture that these two conditions might offer equivalent, substitutable bases for some outcome (e.g., weak class voting, *W*). When using fuzzy sets, logical *or* directs the researcher's attention to the *maximum* of each case's memberships in the component sets. That is, a case's membership in the set formed from the *union* of two or more fuzzy sets is the maximum value of its memberships in the component sets. Thus, if a country has a score of 0.2 in the set of affluent countries and a score of 0.8 in the set of not-high manufacturing countries, it has a score of 0.8 in the set of countries that have *either* of these two traits. For illustration of the use of logical *or*, consider the eighth data column of table 2.2. This column shows countries that have either a not-high percentage of workers in manufacturing ($\sim M$) or high affluence (*A*); the algebraic expression is $\sim M + A$, where the addition sign is used to indicate logical *or*.

Fuzzy-Set Relations

Chapter 1 presents my rationale for studying set relations in social research. Set relations reflect the explicit connections that are central to social science theorizing. Theory is largely verbal in nature; thus, set relations are central to social theory, just as they are to verbal statements in general. One of the great strengths of fuzzy sets is that they make set-theoretic analysis possible while retaining fine-gained empirical gradations. In short, it is possible to determine if one set is a subset of another (e.g., do the developed countries constitute a subset of the democratic countries?) *without* reverting to nominal-scale measurement (i.e., crisp sets).

With crisp sets, determining whether the cases sharing a specific combination of conditions share the same outcome, and thus constitute a subset of the cases with the outcome, is a simple matter. (Recall that this is one of the two important types of explicit connections described in chapter 1.) The researcher simply examines cases sharing the relevant combination of conditions and assesses whether they agree in displaying the outcome. This assessment can be seen as an evaluation of the second column of the cross-tabulation of the presence/absence of the outcome in question against the presence/absence of a given combination of causal conditions (see table 2.3). The subset relation is indicated when the cell corresponding to the presence of the causal combination and the absence of the outcome is empty and the cell corresponding to the presence of the outcome is empty and the presence of the outcome is populated with cases.⁴

Obviously, these procedures cannot be duplicated with fuzzy sets. There is no simple way to isolate the cases sharing a specific combination of causal conditions because each case's array of membership scores may be unique. Cases also have different degrees of membership in the outcome, complicating the assessment of whether they "agree" on the outcome. While these properties of fuzzy sets make

4. Of course, cell 4 may not be completely empty. In case-oriented research, however, the researcher should be able to explain the errant cases that may find their way into cell 4.

Table 2.3: Crisp-set assessment of the connection between a combination of causal conditions and an outcome (causal combination is a subset of the outcome)

	Causal combination absent	Causal combination present
Outcome present	cell 1: not directly relevant to the assessment	cell 2: cases here support researcher's argument that this is a connection
Outcome absent	cell 3: not directly relevant to the assessment	cell 4: should be empty or nearly empty; cases here undermine argument

it difficult to duplicate crisp-set procedures for assessing subset relationships, the fuzzy subset relation can be assessed using fuzzy algebra. With fuzzy sets, a subset relation is indicated when membership scores in one set (e.g., a causal condition or combination of causal conditions) are consistently less than or equal to their corresponding membership scores in another set (e.g., the outcome).

For illustration, consider the data listed in table 2.4 and plotted in figure 2.1. Table 2.4 shows membership scores in two fuzzy sets, the set of countries with weak class voting (W) and the set of countries lacking strong unions ($\sim U$), using data from table 2.2. Observe that the weak class voting membership scores are consistently greater than or equal to not-strong unions scores. This pattern is consistent with the fuzzy subset relation. If membership in the causal condition is high, then membership in the outcome also must be high. Note, however, that the reverse does not have to be true. That is, the fact that there are cases with relatively low membership in the causal condition but substantial membership in the outcome (e.g., Ireland and Belgium) is not problematic from the viewpoint of set theory because the expectation is that there may be several different ways to generate high membership in the outcome (i.e., there are causal pathways to weak class voting in addition to the one illustrated). Cases with low scores in the causal condition (or combination of conditions) but high scores in the outcome indicate the operation of alternate causal conditions or alternate combinations of causal conditions.

Table 2.4: Illustration of fuzzy subset relation ($\sim U \leq W$)

Country	Weak class voting (W)	Not-strong unions ($\sim U$)
Australia	0.6	0.4
Belgium	0.6	0.2
Denmark	0.2	0.2
France	0.8	0.8
Germany	0.6	0.6
Ireland	0.8	0.2
Italy	0.6	0.4
Netherlands	0.8	0.6
Norway	0.2	0.2
Sweden	0	0
United Kingdom	0.4	0.4
United States	1	0.8

Figure 2.1 illustrates the fuzzy subset relation using the membership scores from table 2.4. The characteristic upper-left triangular plot indicates that the set plotted on the horizontal axis (not-strong unions, $\sim U$) is a subset of the set plotted on the vertical axis (weak class voting, W). The points in the upper-left region of the plot are not "errors," as they would be regarded in a linear regression analysis. Rather, these points have strong membership in the outcome due to the operation of other causal conditions or other combinations of causal conditions. The vacant lower triangle in this plot of fuzzy sets corresponds to empty cell 4 of table 2.3, which uses crisp sets. Just as cases in cell 4 of table 2.3 would violate the crisp subset relation, cases in the lower-right triangle of figure 2.1 would violate the fuzzy subset relation.

Table 2.4 and figure 2.1 illustrate the fuzzy subset relation using a single causal condition. Note, however, that this same assessment could have been conducted using degree of membership in a combination of causal conditions. As noted previously, in order to compute a case's degree of membership in a combination of conditions, it is necessarily simply to use the lowest (minimum) membership score among the causal conditions, which follows from the application of fuzzy algebra's logical *and* operation. Degree of membership in a causal combination can then be used to assess the fuzzy subset relation

by comparing scores in the causal combination (horizontal axis) with membership scores in the outcome (vertical axis). This examination establishes whether degree of membership in a combination of causal conditions is a fuzzy subset of degree of membership in the outcome, a pattern of results consistent with an argument of sufficient causation (Ragin 2000). An upper-left triangular plot, with degree of membership in the causal combination on the horizontal axis and degree of membership in the outcome on the vertical axis, signals the fuzzy subset relation. (See especially chapter 6 for an in-depth examination of the study of configurations of set memberships.)

Recall several of the main points about set relations in social research presented in chapter 1, namely (1) set relations often involve causal or other integral connections between social phenomena, (2) they are fundamentally asymmetric, and (3) they can be strong despite relatively weak correlations. Consider the statement, "Among the advanced industrial democracies, those lacking strong unions have weak




Figure 2.1 Plot of fuzzy sets showing subset relationship

class voting." The statement hypothesizes an explicit link from weak unions to weak class voting. Like many such statements, it lists the subset first (weak class voting) and claims, in essence, that the set of countries with weak unions constitutes a subset of the countries with weak class voting. The statement is fundamentally asymmetric. It does not claim that countries with *strong* unions are somehow barred or prevented from having weak class voting, and thus it leaves open the possibility that other obstacles to class voting may exist. Such evidence does not directly challenge the claim that there is a path through weak unions. Finally, from a set-theoretic viewpoint, the evidence presented is *perfectly* consistent with the set-theoretic claim: all cases are in the upper-left triangle of the plot. From a correlational viewpoint, however, the evidence is imperfect, as indicated in the scatter of the points (Pearson's $r = .766$).

Looking Ahead

Chapter 3 presents two simple descriptive measures for evaluating set theoretic connections, such as the one shown in figure 2.1. Specifically, chapter 3 demonstrates that it is possible to assess both the degree to which empirical evidence is consistent with the claim that a set-theoretic connection exists and the empirical importance or relevance of that connection.

Practical Appendix: Fuzzy-Set Relations

Fuzzy-set relations are easy to spot using simple XY plots of fuzzy membership scores. In general, triangular plots, with points consistently above or consistently below either diagonal of the plot, signal a fuzzy subset relation of some sort. Fuzzy-set qualitative comparative analysis (fsQCA) includes facilities for plotting fuzzy-set relations.

1. Create a data set with fuzzy membership scores. Fuzzy membership scores can be assigned, or they can be computed using the procedures detailed in chapters 4 and 5. Data can be entered directly into fsQCA or imported into fsQCA as comma-delimited files (e.g.,

from Excel) or tab-delimited files (e.g., from SPSS). Simple variable names should appear on the first line of the data set (see chapter 1, Practical Appendix).

2. Once the data set is visible in the data spreadsheet window of fsQCA, click *Graphs*, then *Fuzzy*, then *XY Plot*. Specify the fuzzy sets to be plotted on the X and Y axes by clicking the adjacent down arrows and then clicking the relevant variable names. It is also possible to negate fuzzy sets before plotting them; click the *Negate* box next to the variable name. Specify an optional *case Id variable* so that the case or cases that reside on specific points in the plot can be readily identified.

3. Click *Plot*. Examine the pattern. Click on any point in the plot, and its information will appear at the bottom of the plot.

4. The numbers shown in the boxes above the upper-left corner and below the lower-right corner of the plot are consistency and coverage scores, which are explained in chapter 3.

3: Evaluating Set Relations *Consistency and Coverage*

This chapter presents simple descriptive measures for evaluating the strength of the empirical support for arguments specifying set-theoretic connections. To structure the presentation, I focus primarily on arguments stating that a specific cause or combination of causal conditions constitutes one of several possible paths to an outcome (see discussion of explicit connections in chapter 1). When this is true, cases displaying a specified causal combination should constitute a subset of the cases displaying the outcome, as illustrated in chapter 2. I present measures for assessing two distinct aspects of set-theoretic connections. Set-theoretic *consistency* gauges the degree to which the cases sharing a given combination of conditions (e.g., democratic dyad) agree in displaying the outcome in question (e.g., peaceful coexistence). That is, consistency indicates how closely a perfect subset relation is approximated. Set-theoretic *coverage*, by contrast, assesses the degree to which a cause or causal combination “accounts for” instances of an outcome. When there are several paths to the same outcome, the coverage of a given causal combination may be small. Thus, coverage gauges empirical relevance or importance.

These same measures can be used to evaluate situations where the researcher suspects that a causal condition is necessary (but not sufficient) for an outcome, that is, where instances of an outcome constitute a subset of instances of a cause. (This set relation is the other type of explicit connection discussed in chapter 1.) In this context, *consistency* assesses the degree to which instances of the outcome agree in displaying the causal condition thought to be necessary, while *cover-*

age assesses the relevance of the necessary condition—the degree to which instances of the condition are paired with instances of the outcome. This discussion of necessary conditions builds on the work of Goertz (2002, 2003), Goertz and Starr (2002), and Braumoeller and Goertz (2000).

These assessments of set relations are important in the analysis of explicit connections in the same way that assessments of significance and strength are important in the analysis of correlational connections. Consistency, like significance, signals whether an empirical connection merits the close attention of the investigator. If a hypothesized subset relation is not consistent, then the researcher's theory or conjecture is not supported. Coverage, like strength, indicates the empirical relevance or importance of a set-theoretic connection. As shown in this chapter, just as it is possible in correlational analysis to have a significant but weak correlation, it is possible in set-theoretic analysis to have a set relation that is highly consistent but low in coverage. I argue here and show in subsequent chapters that these set-theoretic measures provide vital tools for refining both crisp-set and fuzzy-set analysis in social research.

Set-Theoretic Consistency

Perfectly consistent set relations are relatively rare in social research. They usually require either small *Ns* or macrolevel data or both. Generally, social scientists are able to identify only rough subset relations because exceptions are almost always present (e.g., a war between two democratic countries). It is important, therefore, to develop useful descriptive measures of the degree to which a set relation has been approximated, that is, the degree to which the evidence is consistent with the argument that a set relation exists. First, the chapter addresses the evaluation of the consistency of crisp-set relations, where a very simple measure suffices, and then turns to fuzzy sets.

When conducting consistency evaluations, it is prudent to take the number of cases into account. Perfect consistency does not guarantee that a meaningful set-theoretic connection exists. Suppose, for

example, that all three third-wave democracies that adopted parliamentary governments subsequently failed. The prudent conclusion would be that this connection, while interesting and perfectly consistent from a set-theoretic viewpoint, might well be happenstance (see also Dion 1998; Ragin 2000). Most social scientists would be more convinced of an explicit connection between parliamentary government and failure if the tally was, say, seventeen out of twenty, instead of three out of three. This example also underscores the fact that "close counts" in social science. While not 100 percent, a rate of 17 out of 20 (85 percent) is substantial enough to indicate, to a social scientist at least, that some sort of integral connection may exist that is worthy of further investigation.

This example suggests a straightforward measure of the consistency of set relations using crisp sets: the proportion of cases with a given cause or combination of causes that also display the outcome. With three out of three cases consistent, the proportion is 1.0; with seventeen out of twenty cases consistent, the proportion is 0.85. As explained in *Fuzzy-Set Social Science* (Ragin 2000), the N of cases can be taken into account by using benchmarks and an exact probability test. For example, with three cases, a proportion consistent of 1.0 is not significantly greater than a benchmark proportion of 0.65, using a significance level (alpha) of 0.05. However, a proportion of 0.85 with an N of 20 passes this test. In general, consistency scores should be as close to 1.0 (perfect consistency) as possible. When observed consistency scores are below 0.75, maintaining on substantive grounds that a set relation exists, even a very rough one, becomes increasingly difficult (see also Ragin 2004a).

The assessment of the consistency of *fuzzy-set* relations is more interesting and more challenging than that of the crisp-set case. An overview of the use of fuzzy sets in social research is presented in *Fuzzy-Set Social Science* (Ragin 2000; see also Smithson and Verkuilen 2006, and chapter 2 of this book). The key point for present purposes is that with fuzzy sets, cases can have varying degrees of membership in sets, with membership scores ranging from 0.0 to 1.0. For example, a country might have only partial membership in the set of democracies. The

calibration of degree of membership in a fuzzy set involves both quantitative and qualitative assessment and must be grounded in theoretical and substantive knowledge (Ragin 2000; Smithson and Verkuilen 2006; see also chapters 4 and 5 of this book).

As explained in chapter 2, a fuzzy subset relation exists when the membership scores in one set are consistently less than or equal to their corresponding membership scores in another. For example, if degree of membership in "parliamentary form of government" is consistently less than or equal to degree of membership in "failure of democracy" across relevant third-wave democracies, then the former is a subset of the latter. Recall that with crisp sets, it does not matter that instances of failure exist that are not also instances of parliamentary government because there are (hypothetically) many ways to fail (see table 1.2). With fuzzy sets, the parallel situation occurs when cases display outcome membership scores that greatly exceed their membership scores in the causal condition. For example, a case might have a score of 0.90 in failure but a score of only 0.20 in parliamentary government. As in the crisp analysis, this case is not inconsistent with the set-theoretic argument because there may be several ways to fail, including paths for countries with weak membership in the set of countries with parliamentary governments. By contrast, a country with a membership score of 0.80 in parliamentary government but a membership of only 0.30 in failure clearly contradicts the set-theoretic claim.

The fuzzy subset relation has a triangular form when depicted as a plot of two fuzzy sets, as shown in figure 3.1. In this figure, the causal condition (X) is a subset of the outcome (Y); thus, all X_i values are less than or equal to their corresponding Y_i values, where i indicates reference to individual X or Y values or specific observations of X or Y . Note that cases in the upper-left corner of the plot do not contradict the idea that this cause may be sufficient but not necessary for the outcome, because these cases have high membership in the outcome due to the operation of causal conditions other than X (an argument of sufficiency without necessity permits multiple paths). Thus, when membership in X is low, a wide range of Y_i values is permissible. When membership in X is high, however, many more opportunities may be




Figure 3.1 Fuzzy subset relation consistent with sufficiency

present to violate the subset relation, as the range of permissible Y_i values narrows. Of course, in a conventional correlational analysis, points in the upper-left corner would be considered errors, which in turn would undermine the correlation between X and Y .

In *Fuzzy-Set Social Science* (Ragin 2000), the definition of the consistency of a fuzzy-set relation is straightforward but simplistic. In the plot of membership in the outcome (Y) against membership in a causal condition or combination of causal conditions (X), consistency is defined as the proportion of cases on or above the main diagonal of the plot. If membership in X is consistently less than or equal to membership in Y , then all the cases will plot on or above the main diagonal, yielding a consistency score of 1.0 (or 100 percent consistent). In the “fuzzy inclusion” algorithm described in Ragin (2000), consistency scores are computed for different combinations of causal conditions, and these scores provide the basis for evaluating sufficiency (see Ragin, Drass, and Davey 2007; Ragin 2007). For example, if significantly

greater than 80 percent of the cases plot on or above the main diagonal in the plot just described, then the investigator might claim that the cause or causal combination X is “almost always” sufficient for the outcome Y .

The procedures presented in *Fuzzy-Set Social Science* for the evaluation of the sufficiency of causal combinations are based on the simple categorization of cases as either consistent or inconsistent and the computation of the simple proportion of consistent cases. In short, the procedures closely follow the crisp-set template. One issue in the use of this procedure concerns the contrast between cases with strong versus weak membership in the causal condition or combination of causal conditions (X). Specifically, cases with strong and weak membership in the causal combination are weighted equally in the calculation, yet they differ substantially in their relevance to the set-theoretic claim and thus to the argument that X is sufficient for Y . For example, a case with a membership of only 0.25 in the set of cases with the causal combination (X) and a score of 0.0 in the outcome set (Y) is just as inconsistent as a case with a score of 1.0 in the causal combination and a score of 0.75 in the outcome. (A membership score of 0.25 indicates that a case is more out than in a set; 0.5 is the crossover point.) In fact, however, the second inconsistent case, with full membership in X , clearly has more bearing on the set-theoretic argument because it is a much better instance of the causal combination. It thus constitutes a more glaring inconsistency than the first case despite the equal gaps—the X_i values exceed the Y_i values by the same amount.

The same reasoning holds for consistent cases. A consistent case with two high membership scores (e.g., 0.9 in the causal combination and 1.0 in the outcome) is clearly more relevant to the set-theoretic argument than a consistent case with two low scores (e.g., 0.1 in the causal combination and 0.2 in the outcome) or a consistent case with a low score in the causal combination (say, 0.15) and a high score in the outcome (say, 0.8). Yet all are counted equally in the formula for consistency used in Ragin (2000) (the proportion of cases on or above the main diagonal in the fuzzy plot). Imagine trying to support an argument in an oral presentation to colleagues using in-depth evidence on

a case with only weak membership in the relevant sets. The common-sense thinking that indicates that this presentation would be a waste of time is precisely formalized in fuzzy membership scores. Cases with strong membership in the causal condition provide the most relevant consistent cases and the most relevant inconsistent cases.

This commonsense idea is operationalized in the alternate measure of the consistency of fuzzy-set data with set-theoretic arguments recommended in this chapter. This alternate procedure, like the one proposed in Ragin (2000), differentiates between consistent and inconsistent cases using the diagonal of the plot. A case on or above the main diagonal is consistent because its membership in the causal condition is less than or equal to its membership in the outcome. A case below the main diagonal is inconsistent because its membership in the causal condition is greater than its membership in the outcome. However, rather than simply calculating the raw proportion of consistent cases, the alternate procedure uses fuzzy membership scores.

Consider, for example, the hypothetical fuzzy-set data on degree of membership in "strong left parties" and "generous welfare states" for twelve advanced industrial countries in table 3.1. Notice that the data in this table are perfectly consistent from a set-theoretic viewpoint; that is, all the membership scores in the causal condition are less than or equal to their corresponding membership scores in the outcome (see chapter 2). Based on this evidence, a researcher could claim that this causal condition (having strong left parties) is a subset of the outcome (having generous welfare states). Thus, having strong left parties could be interpreted (hypothetically, with these data) as a sufficient condition for having a generous welfare states. As previously noted, however, social science data are rarely this uniform. When cases are inconsistent with the subset relation, the researcher must assess the *dégré* to which the empirical evidence is consistent with the set relation in question. For example, suppose the score for strong left parties in the first row of table 3.1 was 1.0 instead of 0.70. It would be inconsistent with the set relation because this value exceeds the corresponding outcome membership score, 0.90. While the set relation would no longer hold consistently across the cases listed in table 3.1, it would

Table 3.1: Illustration of a simple fuzzy-subset relation (hypothetical data for twelve countries)

Strong left parties	Generous welfare states
0.7	0.9
0.1	0.9
0.1	0.1
0.3	0.3
0.9	0.9
0.7	0.7
0.3	0.9
0.3	0.7
0.3	0.7
0.1	0.1
0.0	0.0
0.9	1.0

still be very close to perfectly consistent, with eleven out of the twelve cases consistent and only one near miss.

One straightforward measure of set-theoretic consistency using fuzzy membership scores is the sum of the *consistent* membership scores in a causal condition or combination of causal conditions divided by the sum of *all* the membership scores in a cause or causal combination (Ragin 2003b). In table 3.1, the value of this measure is 1.0 (4.7/4.7) because all the membership scores in column 1 are consistent. If the score for strong left parties in the first row of table 3.1 is changed to 1.0, however, consistency drops to 0.8 (4/5). The numerator is 1.0 fuzzy unit lower than the denominator because of the one inconsistent score of 1.0. The reduction of consistency to 0.8 (from perfect consistency, 1.0) is substantial because 1.0 (the value substituted for 0.70 in the first row) is a large membership score.

This consistency measure can be refined further so that it gives credit for near misses and penalties for causal membership scores that exceed their marks, the corresponding outcome membership scores, by wide margins.¹ This adjustment can be accomplished by adding to

1. The formula described here is the one implemented in the fuzzy-truth table algorithm of fQCA (Ragin, Drass, and Davey 2007).

the numerator in the formula just sketched (the sum of the consistent scores divided by the sum of all membership scores in the cause or causal combination) the part of each *inconsistent* causal membership score that is consistent with the outcome. For example, if the score for strong left parties in the first row of table 3.1 is changed to 1.0, then most of its score is consistent, up to the value of the outcome membership score, 0.90. This portion is added to the numerator of the consistency measure. Using this more refined measure of consistency yields an overall consistency score of 0.98 (4.9/5). This adjusted consistency score is more compatible with the evidence. After all, only one of the scores is inconsistent, and it is a very near miss. Thus, a consistency score close to 1.0 should be expected.

Notice that the revised measure of consistency prescribes substantial penalties for *large* inconsistencies. Suppose again that the score for strong left parties in the first row of table 3.1 is 1.0, but this time assume that the corresponding value of the outcome, generous welfare states, is only 0.3. The consistent portion of the 1.0 membership score is 0.3, yielding an overall addition of only 0.3 to the numerator. The resulting consistency score in this instance would be 0.86 (4.3/5). This lower score reflects the fact that the one inconsistent score exceeds its target by a wide margin.

It is possible to formalize the calculation of fuzzy set-theoretic consistency as follows:

$$\text{Consistency } (X_i \leq Y_i) = \Sigma[\min(X_i, Y_i)]/\Sigma(X_i)$$

where min indicates the selection of the lower of the two values (see also Kosko 1993; Smithson and Verkuilen 2006). When the X_i values are all less than or equal to their corresponding Y_i values, the consistency score is 1.0; when there are only a few near misses, the score is slightly less than 1.0; when there are many inconsistent scores, with some X_i values greatly exceeding their corresponding Y_i values, consistency may drop below 0.5. Note that the same procedures for incorporating probabilistic criteria, mentioned above and discussed in detail in Ragin (2000), can be applied here. These probabilistic tests require a benchmark value (e.g., 0.75 consistency) and an alpha level (e.g., 0.05 significance). Finally, when the formula for the calculation

of fuzzy set-theoretic consistency is applied to crisp-set data, it returns the simple proportion of consistent cases. Thus, the formula can be applied to crisp and fuzzy membership scores alike.

This same general formula also can be applied to the assessment of the consistency of a set relation indicating that a causal condition is a necessary condition for an outcome. An argument of causal necessity is supported when it can be demonstrated that instances of an outcome constitute a subset of instances of a causal condition. With fuzzy sets, the consistency of the necessary condition relationship depends on the degree to which it can be shown that membership in the outcome is consistently less than or equal to membership in the cause, $Y_i \leq X_i$. Figure 3.2 illustrates this fuzzy-set relation. In this figure, the outcome (Y) is a subset of the causal condition (X); thus, all Y_i values are less than or equal to their corresponding X_i values. Note that cases in the lower-right corner of the plot do not contradict necessity, for these are cases that have low membership in the outcome because they lack some other, unspecified causal condition. After all, the causal condition in this example is only necessary, not sufficient. Of course, in a conventional correlational analysis, cases in the lower-right corner would be considered errors, which in turn would undermine the correlation between X and Y. Note, however, that when membership in X is low, membership in Y also must be low. Thus, in the low range of X, many opportunities exist to violate the subset relation, with only a narrow range of permissible Y_i values.

Because the inequality signaling necessity ($Y_i \leq X_i$) is the reverse of the inequality that defines sufficiency ($X_i \leq Y_i$), a simple measure of the consistency of the subset relationship for a necessary condition is:

$$\text{Consistency } (Y_i \leq X_i) = \Sigma[\min(X_i, Y_i)]/\Sigma(Y_i)$$

When all Y_i values are less than or equal to their corresponding X_i values, this formula returns a value of 1.0. When many Y_i exceed their corresponding X_i values by wide margins, the computation returns a value less than 0.5.

Of course, it is important to remember that the interpretation of any set-theoretic relation as either necessary or sufficient must be built on a solid foundation of theoretical and substantive knowledge. Causal




Figure 3.2 Fuzzy subset relation consistent with necessity

connections are not inherent in data. Set-theoretic consistency is only one piece of evidence in the array of support that must be brought to bear when a researcher makes a claim of either sufficiency or necessity or any other kind of causal or constitutive connection.

Set-Theoretic Coverage

When researchers allow for equifinality (Mackie 1965; George 1979; George and Bennett 2005) and causal complexity (Ragin 1987), a common finding is that a given outcome may result from several different combinations of conditions. These combinations are generally understood as alternate causal paths or "recipes" for the outcome. Usually, these alternate paths are treated as logically equivalent (i.e., as substitutable). However, it is common in crisp-set analyses to assess the proportion of cases following each path—that is, the number of cases following a specific path to the outcome divided by the total number

of instances of the outcome. This simple proportion is a direct measure of set-theoretic coverage and is a straightforward indicator of the empirical importance of a causal combination. Clearly, a causal combination that covers or accounts for only a small proportion of the instances of an outcome is not as empirically important as one that covers a large proportion.²

Coverage is distinct from consistency, and the two sometimes work against each other because high consistency may yield low coverage. Complex set-theoretic arguments involving the intersection of many sets can achieve remarkable consistency but low coverage. For example, consider the adults in the United States who combine excellent school records, high achievement test scores, college-educated parents, high parental income, graduation from Ivy League universities, and so on. It would not be surprising to learn that 100 percent of these individuals are able to avoid poverty. Perfect set-theoretic consistency is unusual with individual-level data, but it is certainly not impossible. There are, however, relatively few individuals with this specific combination of highly favorable circumstances among the many who successfully avoid poverty. From a practical viewpoint, therefore, this high level of set-theoretic consistency is not compelling because the causal combination is so narrowly formulated that its coverage is trivial.

While there is often a trade-off between consistency and coverage, it is reasonable to calculate coverage *only after establishing that a set relation is consistent*. It is pointless to compute the coverage of a cause or combination of causes that is not a consistent subset of the outcome. Also, as will become clear in the discussion that follows, the same set-theoretic calculation has different meanings, depending on the context of the calculation. Thus, it is important to adhere to the protocol described here for the results of assessments of consistency and coverage

2. Note that coverage gauges only empirical importance, not theoretical importance. A sufficient relation may be quite "rare" from an empirical point of view (and thus exhibit low coverage), but it still could be centrally relevant to theory. For example, the sufficient relation might be proof that a path that was thought to be empirically impossible, at least from the perspective of theory, in fact is not.

to be meaningful: set-theoretic consistency must be established before coverage can be assessed.

For illustration of the general idea of coverage, consider table 3.2, which shows a hypothetical cross-tabulation of poverty status (in poverty versus not-in poverty) against educational achievement (high versus not-high), using crisp sets and individual-level data. This crude analysis using binary data supports the argument that individuals with high educational achievement are able to avoid poverty. This set-theoretic argument is supported by the high proportion of cases in the second column that are not in poverty (cell b divided by the sum of cells b and d yields a consistency score of 0.964). But how important is this path when it comes to avoiding poverty? The simplest way to answer this question is to calculate the proportion of the individuals not in poverty who have high educational achievement—that is, cell b divided by the sum of cells a and b, which is 0.326. This calculation shows that the path in question covers almost a third of the cases not in poverty, which is substantial.

For comparison purposes, consider table 3.3, which has the same total number of cases as table 3.2, but some of the cases have been shifted from cell b to cell a and from cell d to cell c. The proportion of cases consistent with the set-theoretic argument in table 3.3 is 0.967, about the same as in table 3.2 (0.964). Thus, from a set-theoretic point of view, the evidence is again highly consistent. But how important is this path, using the hypothetical frequencies presented in table 3.3? Its importance can be ascertained by computing the proportion of cases avoiding poverty that are covered by the set-theoretic argument, which is only 0.0325 (147/4,520). Thus, in table 3.3 the set-theoretic pattern is again highly consistent, but coverage is very low, indicating

Table 3.2: Cross-tabulation of poverty status and educational achievement: preliminary frequencies

	Low/average educational achievement	High educational achievement
Not-in poverty	a. 3,046	b. 1,474
In poverty	c. 625	d. 55

Table 3.3: Cross-tabulation of poverty status and educational achievement: altered frequencies

	Low/average educational achievement	High educational achievement
Not-in poverty	a. 4,373	b. 147
In poverty	c. 675	d. 5

(hypothetically) that having high educational achievement is not an important path to the outcome, avoiding poverty.


The procedures for calculating set-theoretic coverage using fuzzy sets parallel the computations for crisp sets presented above. Another way to understand the calculation of coverage using conventional crisp sets (cell b divided by the sum of cells a and b) is to visualize table 3.2 as a Venn diagram showing a subset relationship, as in figure 3.3. The basic idea behind the calculation of coverage is to assess the degree to which the subset (the set of cases with high educational achievement in this example) physically covers the larger set (the set of cases avoiding poverty). Thus, *coverage*, a gauge of empirical weight or importance, can be seen as the size of the overlap of the two sets relative to the size of the larger set (representing the outcome). The calculation of the size of the overlap of two fuzzy sets is given by their intersection:

$$\text{Overlap} = \Sigma[\min(X_i, Y_i)]$$

which is the same as the numerator in the calculation of fuzzy set-theoretic consistency described previously. With fuzzy sets, the size of the larger set (the relevant denominator) is given directly by the sum of the membership scores in that set, that is, the sum of the membership scores in the outcome, $\Sigma(Y_i)$. This calculation parallels the simple counting of the number of cases in a set (e.g., the number of cases not in poverty) using crisp sets. Thus, the measure of fuzzy-set coverage is simply the overlap expressed as a proportion of the sum of the membership scores in the outcome:

$$\text{Coverage}(X_i \leq Y_i) = \Sigma[\min(X_i, Y_i)] / \Sigma(Y_i)$$

In short, the formula for coverage of Y by X substitutes $\Sigma(Y_i)$ for $\Sigma(X_i)$ in the denominator of the formula for consistency.



- Area a = Cases with low/average educational achievement, not in poverty
- Area b = Cases with high educational achievement, not in poverty
- Area c = Cases with low/average educational achievement, in poverty
- Area d = Cases with high educational achievement, in poverty

Figure 3.3 Venn diagram illustrating concept of coverage using hypothetical data (from table 3.2)

Observe that this formula is identical to the formula for the consistency of Y as a subset of X (i.e., $Y_1 \leq X_1$), presented in the discussion of the assessment of the consistency of a necessary conditions relationship. Recall, however, that in the present context (assessing sufficiency), the coverage of Y by X (i.e., as in the equation above) is calculated *only after* it has been established that X is a consistent subset of Y . Thus, the purpose of the calculation in the context of suffi-

ciency is to assess the magnitude of X relative to Y , given that most, if not all, X_1 values are less than or equal to their corresponding Y_1 values. In the necessary-conditions context, by contrast, the goal is to assess the consistency of Y as a subset of X . Thus, in that context the expectation is that most, if not all, Y_1 values will be less than or equal to their corresponding X_1 values—indicating a necessary-conditions relationship. Indeed, if this is not the case, then the result of the calculation will be a consistency score (for $Y_1 \leq X_1$) that falls far short of perfect consistency (i.e., substantially less than 1.0), indicating that Y is *not* a consistent subset of X .³ In short, context must be taken into account when conducting these assessments.

Figure 3.4 depicts the concept of coverage relevant to the fuzzy subset relation, with $X_1 \leq Y_1$. As in figure 3.1, condition X is a subset of outcome Y . Points below the main diagonal constitute violations of consistency and thus undermine the argument that X is a subset of Y . However, there are only two such points, and the subset relationship is largely consistent. When calculating coverage, only the portion of the X_1 score that is above the main diagonal is counted as consistent (and thus included as part of the overlap between X and Y). Most of the points in figure 3.4 are above the main diagonal and thus consistent with $X_1 \leq Y_1$. When X_1 values are small relative to their corresponding Y_1 values, they are closer to the Y axis than they are to the main diagonal. While these points are consistent with the subset relationship $X_1 \leq Y_1$, they contribute relatively little to coverage, especially when the Y_1 values are above 0.5. The dotted horizontal lines in the figure show the portions of the X_1 values counted as consistent; these values are

3. It follows that it is possible to find a $\Sigma[\min(X_1, Y_1)]$ that is close to $\Sigma(Y_1)$ —thus yielding a very high coverage score—only if the values of X_1 are roughly equal to their corresponding Y_1 values. This situation would correspond to a close *coincidence* of the two sets. Set coincidence is not the same as correlation, but rather is a special case of correlation. In a plot of two fuzzy sets, any straight line that is neither vertical nor horizontal yields a perfect correlation coefficient. However, perfect set coincidence occurs only when all the cases plot exactly on the main diagonal of the fuzzy plot. A simple measure of the degree to which the membership scores in two sets coincide is $\Sigma[\min(X_1, Y_1)] / \Sigma[\max(X_1, Y_1)]$, where max indicates using the larger of the two scores. See also Smithson and Verkuilen (2006), who contrast comorbidity, covariation, and co-occurrence.

ment organization. Thus, the existence of grievances could be seen as an empirically trivial necessary condition. By contrast, an open and permissive political climate (i.e., the absence of government repression) could be seen as a nontrivial necessary condition, for social movements routinely encounter government repression. While the specific computational formula recommended in this chapter for assessing the relevance of necessary conditions differs in its details from the one suggested by Goertz (2003, 2006), the underlying goals are similar.⁴

A simple measure of the importance or relevance of X as a necessary condition for Y is given by the degree of coverage of X by Y :

$$\text{Coverage } (Y_i \leq X_i) = \sum[\min(X_i, Y_i)] / \sum(X_i)$$

When the coverage of X by Y is small, the constraining effect of X on Y is negligible. Conceptually, very low coverage corresponds to an empirically irrelevant or even meaningless necessary condition. For example, almost all heroine addicts in the United States are former milk drinkers, but it would be difficult to portray milk drinking as a relevant necessary condition (i.e., as a gateway substance) for heroine addiction because the set of former milk drinkers completely dwarfs the set of heroine addicts. By contrast, when the coverage of X by Y is substantial, then the constraining effect of X as a necessary condition may be great. For example, if a substantial proportion of people who associate with heroine addicts later become addicts and only a very small number of people become addicted to heroine without first associating with heroine addicts, then coverage is high and “associating with heroine addicts” may be considered a relevant necessary condition for heroine addiction.

The contrast between these two situations, high versus low relevance in the analysis of necessary conditions, is depicted in figure 3.5.

4. In Goertz's (2003, 2006) approach, membership scores are divided at the case level, and then these ratios are averaged. In effect, this procedure assigns cases equal importance in the computation of a given measure. In the approach advocated here, however, cases with low fuzzy membership scores are given less weight because they are weak instances of the phenomenon in question. This computational strategy makes the resulting measures more reflective of patterns observed in the best instances.




Figure 3.4 Illustration of the concept of coverage


added to the numerator of the formula for coverage. The denominator is the sum of the Y_i values. The gaps from the consistent X_i values to the main diagonal show the part of set Y that is *not* covered by set X .

The calculation of coverage also can be applied to the assessment of necessary conditions, where the outcome is a subset of the cause. Goertz (2003, 2006), building on Braumoeller and Goertz (2000), presents an approach to the assessment of necessary conditions that addresses some of the same issues discussed in this chapter. A key focus in his work is the distinction between trivial and nontrivial necessary conditions. A trivial necessary condition is one that is strongly present in most cases, whether or not these cases display the outcome. For example, “grievances” may be a necessary condition for the organization and activation of a social movement, but grievances are almost always present, and the absence of grievances rarely gets the chance to act in a constraining manner on social move-

Figure 3.5(a) depicts a necessary condition that exerts some constraint on the outcome (coverage is nontrivial). Figure 3.5(b) depicts an empirically trivial necessary condition (very low set-theoretic coverage). Using fuzzy sets, the situation depicted in figure 3.5(b) would appear as a plot in which almost all cases have very strong membership in X (the causal condition) and thus would plot to the far right (see also Goertz 2003, 2006).

As with the assessment of the coverage of a sufficient condition, it is important to assess the relevance of a necessary condition (i.e., its constraining impact) *only after* establishing that the subset relation is consistent. That is, it must first be established that Y is a consistent subset

a. Empirically Relevant Necessary Condition



b. Empirically Irrelevant Necessary Condition




Figure 3.5 Venn diagram illustrating necessary conditions

Table 3.4: Protocol for assessing consistency and coverage

Procedure	Type of set-theoretic relation	
	Cause (X) is a subset of outcome (Y) (sufficiency)	Outcome (Y) is a subset of cause (X) (necessity)
Step 1	Assess consistency using $\Sigma[\min(X_i, Y_i)]/\Sigma(X)$	Assess consistency using $\Sigma[\min(X_i, Y_i)]/\Sigma(Y)$
Step 2	If consistent, assess coverage using $\Sigma[\min(X_i, Y_i)]/\Sigma(Y)$	If consistent, assess coverage using $\Sigma[\min(X_i, Y_i)]/\Sigma(X)$

of X before assessing the size of Y relative to size of X. Adherence to this protocol prevents confusion regarding the interpretation of what are essentially identical calculations: the calculation of the consistency of a sufficiency relationship is identical to the calculation of the coverage (relevance) of a necessity relationship, while the calculation of the coverage of a sufficiency relationship is identical to the calculation of the consistency of a necessity relationship. The protocol for assessing consistency and coverage for these two types of set-theoretic relations is summarized in table 3.4.

Partitioning Coverage

When more than one condition or combination of conditions is sufficient for an outcome (i.e., when there is equifinality), the assessment of the coverage of alternate causal combinations provides direct evidence of their relative empirical importance. Further, the assessments of "raw" coverage can be complemented with assessments of each combination's "unique" coverage, for it is possible to partition coverage in set-theoretic analysis in a manner that is analogous to the partitioning of explained variation in multiple regression. The discussion of the partitioning of coverage that is presented here assumes that the researcher has demonstrated that the relevant conditions or combinations of conditions are consistent subsets of the outcome.

For purposes of illustration, consider evidence from a fuzzy-set analysis of individual-level data. The data set is the National Longitudinal

Table 3.5: Calculation of coverage

Causal conditions	Sum of consistent		Coverage
	scores	Sum of outcome scores	
T-I	307.387	1385.25	0.2219
C	548.559	1385.25	0.3960
T-I + C	607.709	1385.25	0.4387

Survey of Youth (better known as the *Bell Curve* data; see Herrnstein and Murray 1994). The sample is white males, interviewed as young adults. The outcome is the fuzzy set of cases not in poverty ($\sim P$, where P indicates degree of membership in the set of cases in poverty and \sim indicates negation). The three causal conditions are the fuzzy set of cases with high test scores (T), the fuzzy set of cases with high parental income (I), and the fuzzy set of cases with college education (C). (The calibration of these fuzzy sets is described in chapter 11.) Applying fuzzy-set qualitative comparative analysis (Ragin, Drass, and Davey 2007) to these data yields two recipes for avoiding poverty, namely, the combination of high test scores and high parental income ($T-I$) and college education (C) by itself.

The calculation of the raw coverage of these two recipes for the outcome, avoiding poverty ($\sim P$), is shown in table 3.5. The first row reports the coverage calculation for the combination of high test scores and high parental income ($T-I$). The sum of the overlaps between $T-I$ and the outcome is 307.387. The sum of the memberships in the outcome is 1385.25. Thus, this combination covers about 22.19 percent of the total membership in the outcome ($307.387/1,385.25 = 0.2219$). Using these same procedures, condition C covers about 39.6 percent of the total membership of the outcome (see row two of table 3.5). Thus, both combinations cover a substantial proportion of the outcome. However, the raw coverage of condition C (college education) is much greater.

For comparison purposes, table 3.5 also shows the coverage of the two combinations ($T-I$ and C) conceived as alternate paths to the same outcome, using logical *or*. When causal combinations are joined by logical *or*, each case's score in the union is the maximum value of the

two paths (i.e., the larger of the two scores, membership in $T-I$ versus membership in C). In other words, when there is more than one path to an outcome, it is possible to calculate how close a case is to the outcome by finding its highest membership score among the possible paths. The degree of coverage of the outcome by this maximum score, in turn, can be calculated using the same procedures applied separately to the two components. This calculation is shown in the third row of table 3.5, which reports a coverage of 43.87 percent, greater than the coverage of either of the two components (compare row 3 of table 3.5 with the first two rows). However, the coverage of the two-path model (43.87 percent) is only modestly superior to the raw coverage of the best single path (path C , with 39.6 percent).

Table 3.5 provides all the information that is needed to partition coverage, following the procedure that is used to partition explained variation in multiple regression analysis. To assess an independent variable's separate or unique contribution to explained variation in a multiple regression involving several correlated predictor variables, researchers calculate the *decrease* in explained variation that occurs once the variable in question is removed from the fully specified multivariate equation. For example, to find the unique contribution of X_1 to explained variation in Y , it is necessary to compute the multiple regression equation with all relevant independent variables included and then to recompute the equation *excluding* X_1 . The difference in explained variation between these two equations shows the unique contribution of X_1 . These procedures ensure that the explained variation that X_1 shares with correlated independent variables is not credited to X_1 . The goal of partitioning in fuzzy-set analysis, by contrast, is to assess the relative importance of different combinations of causally relevant conditions. Thus, the issue in set-theoretic analyses is not correlated independent variables because causal conditions are not viewed in isolation from one another, as they are in multiple regression analysis. Rather, partitioning coverage is important because some cases conform to more than one path. Using our example, there are surely many individuals who combine high test scores, high-income parents, and college education.

Table 3.6: Partitioning coverage

	Total coverage	Without term	Unique
Unique to T-I	0.4387	0.3960	0.0427
Unique to C	0.4387	0.2219	0.2168

Consider the crisp-set case. Suppose a researcher finds that two combinations of conditions generate outcome Y: A·B and C·D. The researcher calculates the coverage of these two paths and finds that the first embraces 25 percent of the instances of Y (coverage = 0.25), while the second embraces 30 percent (coverage = 0.3). However, when calculating their coverage as alternate paths (i.e., their union: A·B + C·D, where addition indicates logical *or*), the researcher finds that together they embrace only 35 percent of the instances of the outcome (coverage = 0.35). The reason that this quantity is substantially less than the sum of the two separate coverage scores (i.e., $0.35 < 0.25 + 0.3 = 0.55$) is because the two paths partially overlap. That is, there are cases that combine all four causal conditions (i.e., instances of A·B·C·D) and the coverage of these instances is counted twice when raw coverage is calculated separately for the two causal combinations.

Fortunately, it is a simple matter to partition total coverage (0.35 in this example) into its three components: uniquely due to A·B, uniquely due to C·D, and overlapping (i.e., due to the existence of cases of A·B·C·D). The unique coverage of each term can be calculated by subtraction, following the template provided by regression analysis. The unique coverage of path A·B is $0.35 - 0.3 = 0.05$; the unique coverage of path C·D is $0.35 - 0.25 = 0.10$; and the remainder of total coverage is due to the overlap between these two terms. In short, these simple calculations indicate that 20 percent of the instances of the outcome are A·B·C·D; 5 percent of the instances of the outcome are A·B without C·D; and 10 percent are C·D without A·B.

The calculation of the unique coverage of a combination of conditions in fuzzy-set analysis is exactly parallel, as shown in table 3.6, which uses the same individual-level data used in table 3.5. The cover-

age of the outcome (avoiding poverty) that is uniquely due to path T-I is the difference between the coverage of the two-path model (0.4387) and the coverage that is obtained once this path (T-I) is removed from the two-path model, which in this example is equivalent to the coverage of the other path (C) by itself. Thus, the unique coverage of path T-I is 0.0427, that is, 0.4387 (the combined coverage of the two paths) less 0.3960 (the single coverage of path C). Likewise, the coverage of the outcome that is uniquely due to path C is the difference between the coverage of the two-path model (0.4387) and the coverage of path T-I by itself (0.2219), or 0.2168. These calculations reveal that the unique coverage of path C is much greater than the unique coverage of path T-I. In fact, the coverage of T-I is almost entirely a subset of the coverage of C. (In other words, most of T-I is T-I·C.) Much of the coverage of the two-path model is overlapping. This proportion can be calculated by computing the difference between the coverage of the two-path model (0.4387) and the sum of the two unique portions ($0.0427 + 0.2168 = 0.2595$), which is 0.1792. Figure 3.6 illustrates these results using a Venn diagram.

When many different paths can lead to the same outcome, it is important to calculate both the raw and unique coverage of each causal

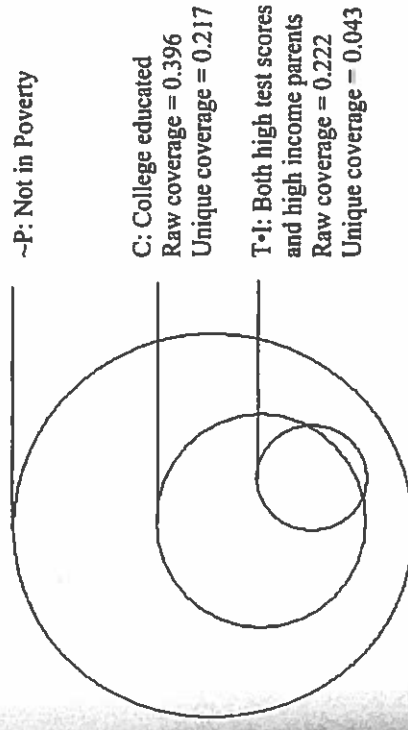


Figure 3.6 Venn diagram representing the partitioning of set-theoretic coverage using fuzzy sets

combination. These calculations often reveal that only a few high-coverage causal combinations exist, even in analyses that have many sufficient combinations. While it is useful to know all the different causal combinations linked to an outcome, it is also important to assess their relative empirical weight. Calculations of raw and unique coverage provide these assessments directly.

Looking Ahead

As will become clear in the chapters that follow, these measures of set-theoretic consistency and coverage have many uses. They can, for example, aid the construction of “crisp” truth tables from fuzzy-set data, which is the foundation of the “fuzzy truth table” algorithm described in chapter 7.

PART II

Calibration versus Measurement

4: Why Calibrate?

Fuzzy sets are relatively new to social science. The first comprehensive introduction of fuzzy sets to the social sciences was offered by Michael Smithson in 1987. However, applications were few and far between until the basic principles of fuzzy-set analysis were elaborated through qualitative comparative analysis (QCA) (see Ragin 1987, 2000), an analytic system that is fundamentally set theoretic, as opposed to correlational, in both inspiration and design. The marriage of these two yields fuzzy-set QCA (fsQCA), a family of methods that offers social scientists an alternative to conventional quantitative methods, which are based almost exclusively on correlational reasoning.

The key to useful fuzzy-set analysis is well-constructed fuzzy sets, which in turn raise the issue of calibration. How does a researcher calibrate degree of membership in a set, for example, the set of Democrats? How should this set be defined? What constitutes full membership? What constitutes full nonmembership? What would a person with 0.75 membership in this set (more in than out, but not fully in) be like? How would this person differ from someone with 0.9 membership? The main message of this chapter is that fuzzy sets, unlike conventional variables, must be calibrated. Because they must be calibrated, they are superior in many respects to conventional measures, as they are used today in both quantitative and qualitative social science. In essence, I argue that fuzzy sets offer a middle path between quantitative and qualitative measurement. However, this middle path is not a compromise between these two; rather, it transcends many of the limitations of both.

What Is Calibration?

Calibration is a necessary and routine research practice in such fields as chemistry, astronomy, and physics (Pawson 1989, 135–37). In these and other natural sciences, researchers *calibrate* their measuring devices and the readings these instruments produce by adjusting them so that they match or conform to dependably known standards. These standards make measurements directly interpretable (Byrne 2002). A temperature of 20 °Celsius is interpretable because it is situated in between 0 degrees (water freezes) and 100 degrees (water boils). By contrast, the calibration of measures according to agreed upon standards is relatively rare in the social sciences.¹ 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. With an uncalibrated measure of temperature, for example, it is possible to know that one object has a higher temperature than another or even that it has a higher temperature than average for a given set of objects but still not know whether it is hot or cold. Likewise, 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.

Calibration is especially important in situations where one condition sets or shapes the context for other conditions. For example, the relationship between the temperature and volume of H₂O changes qualitatively at 0 °C and then again at 100 °C. Volume decreases as temperature crosses 0 °C and then increases as temperature crosses

1. Perhaps the greatest calibration efforts have been exerted in the field of poverty research, where the task of establishing external standards (i.e., defining who is poor) has deep policy relevance. Another example of a calibrated measure is the Human Development Index developed by the United Nations and published in its *Human Development Report*. In economics, by contrast, *calibration* has a different meaning altogether. Researchers “calibrate” parameters in models by fixing them to particular values, so that the properties and behavior of other parameters in the model can be observed. This type of calibration is very different from the explicit calibration of measures, the central concern of this chapter.

100 °C. The Celsius scale is purposefully calibrated to indicate these “phase shifts,” and researchers studying the properties of H₂O know not to examine the relationships between its properties without taking these two qualitative breakpoints into account. Knowledge of these phase shifts, which is external to the measurement of temperature per se, provides the basis for its calibration.²

Context-setting conditions that operate like phase shifts abound in the study of social phenomena. The most basic context-setting condition is the scope condition (Walker and Cohen 1985). When researchers state that a certain property or relationship holds or exists only for cases of a certain type (e.g., only for countries that are democracies), they have used a scope condition to define an enabling context. Another example of a context-setting condition in social science is the use of empirical populations as enabling conditions. For instance, when researchers argue that a property or relationship holds only for Latin American countries, they have used an empirically delineated population as a context-setting condition. While the distinction between scope conditions and populations is sometimes blurred, their use as context-setting conditions is parallel. In both usages, they act as conditions that enable or disable specific properties or relationships.

Tests for statistical interaction are usually motivated by this same concern for conditions that alter the relationships between other variables, that is, by this same concern for context-setting conditions. If the effect of X on Y increases from no effect to a substantial effect as the level of a third variable Z increases, then Z operates as a context-setting condition, enabling a relationship between X and Y. Unlike scope conditions and population boundaries, the interaction variable Z in this example varies by level and is not a simple presence/absence dichotomy. While having context-setting conditions vary by level or degree complicates their study, the logic is the same in all three situations. In fact, it could be argued that dichotomous context-setting conditions such as scope conditions are special cases of statistical interaction.

2. I thank Henry Brady for pointing out the importance of the idea of phase shifts as a way to elaborate my argument.

The fact that the interaction variable Z varies by level as a context-setting condition automatically raises the issue of calibration. At what level of Z does a relationship between X and Y become possible? At what level of Z does a strong connection exist between X and Y ? To answer these questions, it is necessary to specify the relevant values of Z , which is a *de facto* calibration of Z . Over a specific range of values of Z , there is no relation between X and Y , while over another range there is a strong relation between X and Y . Perhaps over intermediate values of Z , there is a weak to moderate relation between X and Y . To specify these values or levels, it is necessary to bring in external, substantive knowledge in some way—to interpret these different levels as context-setting conditions. Researchers who test for statistical interaction have largely ignored this issue and have been content to conduct broad tests of statistical interaction, focusing on increments to explained variation in a dependent variable without attending to issues of calibration and context.

To set the stage for the discussion of fuzzy sets and their calibration, I first examine common measurement practices in quantitative and qualitative social research. After sketching these practices, I show that fuzzy sets resonate with both the measurement concerns of qualitative researchers, where the goal often is to distinguish between relevant and irrelevant variation (that is, to *interpret* variation), and with the measurement concerns of quantitative researchers, where the goal is the precise placement of cases relative to each other.

Common Measurement Practices in Quantitative Research

Measurement, as practiced in the social sciences today, remains relatively haphazard and unsystematic, despite the efforts and exhortations of many distinguished scholars (e.g., Duncan 1984; Pawson 1989). The dominant approach is the indicator approach, in which social scientists seek to identify the best possible empirical indicators of their theoretical concepts. For example, national income per capita (in constant U.S. dollars, adjusted for differences in purchasing power) is often used as an empirical indicator of the theoretical concept of devel-

opment applied to countries. In the indicator approach, the main requirement is that the indicator must vary across cases, ordering them in a way that is consistent with the underlying concept. The values of national income per capita, for example, must distinguish less-developed from more-developed countries in a systematic manner.

In this approach, fine gradations and equal measurement intervals are preferred to coarse distinctions and mere ordinal rankings. Indicators such as income per capita are especially prized not only because they offer fine gradations (e.g., an income per capita value of \$5,500 is exactly \$100 less than a value of \$5,600) but also because the distance between two cases is considered the “same” regardless of whether it is the difference between \$1,000 and \$2,000 or between \$21,000 and \$22,000 (i.e., a \$1,000 difference).³ Such interval- and ratio-scale indicators are well suited for the most widely used analytic techniques for assessing relationships between variables, such as multiple regression and related linear techniques.⁴

More sophisticated versions of the indicator model use multiple indicators and rely on psychometric theory (Nunnally and Bernstein 1994). The core idea in psychometric theory is that an index that is composed of multiple, correlated indicators of the same underlying concept is likely to be more accurate and more reliable than any single indicator. Consider this simple example: National income per capita could easily overstate the level of development of oil-exporting countries,

3. Actually, there is a world of difference between living in a country with a gross national product (GNP) per capita of \$2,000 and living in one with a GNP per capita of \$1,000; however, there is virtually no difference between living in one with a GNP per capita of \$22,000 and living in one with a GNP per capita of \$21,000. Such fine points are rarely addressed by researchers who use the conventional indicator approach, but they must be confronted directly in research that uses calibrated measures (e.g., fuzzy sets).

4. While most textbooks assert that ratio scales are the highest form of measurement because they are anchored by a meaningful zero point, it is important to note that fuzzy sets have three numerical anchors: 1.0 (full membership), 0.0 (full non-membership), and 0.5 (the crossover point separating cases that are more in versus more out of the set in question); see Ragin (2000). If it is accepted that such “anchoring” signals a higher level of measurement, then it follows that a fuzzy set is a higher level of measurement than a ratio-scale variable.

making them appear to be more developed than they "really are." Such anomalies challenge the face validity of income per capita as an indicator of the underlying concept. However, using an index of development composed of multiple indicators (e.g., including such factors as literacy, life expectancy, energy consumption, and labor force composition) would address these anomalies, because many oil-exporting countries have relatively lower scores on some of these alternate indicators of development. Ideally, the various indicators of an underlying concept should correlate very strongly with each other. If they do not, then they may be indicators of different underlying concepts (Nunnally and Bernstein 1994). Only cases with consistently high scores across all indicators obtain the highest scores on an index built from multiple indicators. Correspondingly, only those cases with consistently low scores across all indicators obtain the lowest scores on an index. Cases in the middle, of course, are a mixed bag.

Perhaps the most sophisticated implementation of the indicator approach is through an analytic technique known as structural equation modeling (SEM) (see Bollen 1989). SEM extends the use of multiple indicators of a single concept (the basic psychometric model) to multiple concepts and their interrelationships. In essence, the construction of indexes from multiple indicators takes place within the context of an analysis of the interrelationships among concepts. Thus, index construction is adjusted in ways that optimize hypothesized relationships. Using SEM, researchers can evaluate the coherence of their constructed indexes within the context of the model in which they are embedded. Simultaneously, they can evaluate the coherence of the model as a whole.

All techniques in the indicator family share a deep reliance upon observed variation, which in turn is almost always sample specific in its definition and construction. As mentioned previously, in the conventional approach, the key requirement that an indicator must meet is that it must order cases in a way that reflects the underlying concept. It is important to point out that these orderings are entirely relative in nature. That is, cases are defined *relative to each other* in the distribution of scores on the indicator (i.e., as having "higher" versus "lower"

scores). For example, if the United States' national income per capita is \$1,000 higher than Italy's, then the United States correspondingly is considered *relatively* more developed. The greater the gap between countries, the more different their relative positions in the development hierarchy. Furthermore, the definition of "high" versus "low" scores is defined *relative to the observed distribution of scores*, usually conceived as a sample of scores drawn from a well-defined population. Thus, a case with a score that is above the sample's central tendency (usually the mean) has a high score; the greater this positive gap, the higher the score. Likewise, a case with a score that is below the mean has a low score; the greater this negative gap, the lower the score. Notice that the use of deviations from sample-specific measures of central tendency offers a very crude but passive form of calibration. Its crudeness lies in the fact that the calibration standards (e.g., the mean and standard deviation) vary from one sample to the next and are inductively derived. By contrast, the routine practice in the physical sciences is to base calibration on external, dependably known standards (e.g., the boiling point of water).

At first glance, these conventional practices with respect to the use of indicators in the social sciences appear to be entirely straightforward and uncontroversial. It seems completely reasonable, for example, that countries should be ranked relative to each other and that some measure of central tendency, based on the sample or population in question, should be used to define high versus low scores. Again, the fundamental requirement of the indicator model is simply variation, which in turn requires only (1) a sample (or population) displaying a variety of scores and (2) a measure of central tendency based on the sample (or population). Note, however, that in this view all variation is considered equally *relevant*.⁵ That is, variation in the entire range of the indicator is considered pertinent with respect to what it

5. Of course, researchers sometimes transform their variables (e.g., using logarithmic transformations of raw data) in order to reduce skew and shift the weight of the variation. However, such adjustments are relatively uncommon and, in any event, are usually understood mechanically, as a way to improve the robustness of a model.

reveals about the underlying concept. For example, the two countries at the very top of the income distribution are both "highly developed countries." Yet, the difference that separates them indicates that one is still more highly developed than the other. In the indicator approach, this difference is usually taken at face value, meaning that usually no attempt is made to look at the cases and ask whether this difference—or any other difference, regardless of magnitude—is a relevant or meaningful difference with respect to the underlying concept.⁶ By contrast, the interpretation of scores relative to agreed upon, external standards is central to measurement calibration. These external standards provide a context for the interpretation of scores.

Common Measurement Practices in Qualitative Research

In conventional quantitative research, measures are indicators of concepts, which in turn are components of models, which in turn are derived from theories. Thus, the quantitative approach to measurement is strongly theory centered. Much qualitative research, by contrast, is more knowledge and case centered and thus tends to be more grounded in empirical evidence and also more "iterative" in nature. That is, there is an interplay between concept formation and measurement on the one hand and research strategy on the other hand (see, e.g., Glaser and Strauss 1967). The researcher begins with orienting ideas and broad concepts and uses empirical cases to help refine and elaborate concepts (Becker 1958). This process of progressive refinement involves an iterative "back and forth" movement between ideas and evidence (Katz 1982; Ragin 1994). In this back-and-forth process, researchers specify and refine their empirical indicators and measures.

6. Notice also that the idea that variation at either end of a distribution should be deemphasized or truncated in some way is usually viewed with great suspicion by quantitative researchers because truncating such variation tends to attenuate correlations.

Consider this simple example: Macrolevel researchers often distinguish between countries that experienced "early" versus "late" state formation (see, e.g., Rokkan 1975). Those that developed early had certain advantages over those that developed late and vice versa. David Laitin (1992, xi), for example, notes that coercive nation-building practices available earlier to monarchs (e.g., the draconian imposition of a national language) are not available to leaders of new states today, in part because of the international censure these policies might generate. But what is early state formation? The occurrence of state formation, of course, can be dated. Thus, it is possible to develop a relatively precise ratio-scale measure of the "age" of a state. But most of the variation captured by this simple and direct measure is not relevant to the concept of early versus late state formation. Suppose, for example, that one state has been around for 500 years and another for 250 years. The first is twice as old as the second, but both are fully early when viewed through the lens of accumulated substantive and theoretical knowledge about state formation. Thus, much of the variation captured by the ratio-scale indicator age is simply irrelevant to the distinction between early and late state formation. Age in years must be adjusted on the basis of accumulated substantive knowledge in order to be able to interpret early versus late in a way that resonates appropriately with existing theory.

Such calibrations are routine in qualitative work, even though they are rarely modeled or even stated explicitly. Indeed, from the perspective of conventional quantitative research, it appears that qualitative researchers skew their measurements to fit their preconceptions. In fact, however, the qualitative researcher's goal is simply to interpret "mere indicators" such as age in years in the light of knowledge about cases and the interests of the investigator (e.g., whether a state is early or late from the standpoint of state formation theory).

A second essential feature of measurement in qualitative research is that it is more case oriented than measurement in quantitative research. This observation goes well beyond the previous observation

that qualitative researchers pay more attention to the details of cases. In case-oriented research, the conceptual focus is on specific kinds of cases, for example, the developed countries. In variable-oriented research, by contrast, the focus is on dimensions of variation in a defined sample or population of cases, for example, variation in level of development across currently constituted nation-states. The distinction is subtle but important because cases can vary not only along a given dimension but also in how well they satisfy the requirements for membership in a category or set. For example, countries vary in how well they satisfy requirements for membership in the set of developed countries—some cases satisfy them fully, some partially, and some not at all. In order to assess how well cases satisfy membership requirements, it is necessary to invoke external standards, for example, regarding what it takes for a country to be considered developed. Thus, in the case-oriented view, the main focus is on sets of cases, the members of which can be identified and studied individually (e.g., the developed countries). In the variable-oriented view, by contrast, cases are usually understood simply as sites for taking measurements (that is, they are often seen as mere “observations”), which in turn provide the necessary raw material for studying relationships between variables, viewed as cross-case patterns.

It follows that the case-oriented view is more compatible with the idea that measures should be calibrated, for the focus is on the degree to which cases satisfy membership criteria, which in turn are usually externally determined, not inductively derived (e.g., using the sample mean). These membership criteria must reflect agreed upon standards; otherwise, the constitution of a category or set will be contested. In the variable-oriented view, the members of a population simply vary in the degree to which they express a given trait or phenomenon, and there is usually no special motivation for specifying the criteria for membership in a set or for identifying specific cases as instances. Thus, a fundamental difference between the qualitative approach to measurement and the quantitative approach is that, in the qualitative approach, meaning is attached to or imposed upon specific measure-

ments, for example, what constitutes early state formation or what it takes to warrant designation as a developed country. In short, measurement in qualitative research is interpreted.

The qualitative sociologist Aaron Cicourel was an early proponent of the understanding of measurement described here. In his classic text, *Method and Measurement in Sociology*, Cicourel (1964, 24) argues that it is necessary to consider the three “media” through which social scientists develop categories and link them to observable properties of objects and events: language, cultural meaning, and the properties of measurement systems. In his view, the problem of establishing equivalence classes (like “democracies” or “developed countries”) cannot be seen as independent from or separate from problems of language and cultural meaning. Cicourel (1964, 33) argues, “Viewing variables as quantitative because available data are expressed in numerical form or because it is considered more ‘scientific’ does not provide a solution to the problems of measurement but avoids them in favor of measurement by fiat. Measurement by fiat is not a substitute for examining and re-examining the structure of our theories so that our observations, descriptions, and measures of the properties of social objects and events have a literal correspondence with what we believe to be the structure of social reality.” In simple terms, Cicourel argues that measures and their properties must be evaluated in the context of both theoretical and substantive knowledge. The fact that social scientists may possess a ratio-scale indicator of a theoretical concept does not mean that this aspect of “social reality” has the mathematical properties of this type of scale.

Thus, in qualitative research, the idea that social scientists should use external standards to evaluate and interpret their measures has much greater currency than it does in conventional quantitative research. An important difference with quantitative research, however, is that measurement in qualitative research is typically lacking in precision, and the context-sensitive and case-oriented way of measuring that is typical of qualitative research often appears haphazard and unscientific.

Fuzzy Sets: A Bridge between the Two Approaches

With fuzzy sets, it is possible to have the best of both worlds, namely, the precision that is prized by quantitative researchers and the use of substantive knowledge to calibrate measures that is central to qualitative research. With fuzzy sets, precision comes in the form of quantitative assessments of degree of set membership, which can range from a score of 0.0 (full exclusion from a set) to 1.0 (full inclusion). Substantive knowledge provides the external criteria that make it possible to calibrate measures. This knowledge indicates what constitutes full membership, full nonmembership, and the point at which cases are more in a given set than out (Ragin 2000; Smithson and Verkuilen 2006; see also chapter 2).

The external criteria that are used to calibrate measures and translate them into set membership scores may reflect standards based on social knowledge (e.g., the fact that twelve years of education constitutes an important educational threshold), collective social scientific knowledge (e.g., about variation in economic development and what it takes to be considered fully in the set of developed countries), or the researcher's own accumulated knowledge, derived from the study of specific cases. These external criteria should be stated explicitly, and they also must be applied systematically and transparently. This requirement separates the use of fuzzy sets from conventional qualitative work, where the standards that are applied usually remain implicit.

Fuzzy sets bridge quantitative and qualitative approaches to measurement because they are simultaneously qualitative and quantitative. Full membership and full nonmembership are qualitative states. In between these two qualitative states are varying degrees of membership ranging from more out (closer to 0.0) to more in (closer to 1.0). Fuzzy sets are also simultaneously qualitative and quantitative because they are both case oriented and variable oriented. They are case oriented in their focus on sets and set membership. In case-oriented work, the identity of cases matters, as do the sets to which a case may belong (e.g., the set of democracies). Fuzzy sets are also variable oriented in

their allowance for degrees of membership and thus for fine-grained variation across cases. This aspect of fuzzy sets also provides a basis for precise measurement, which is greatly prized in quantitative research.

Differences between Fuzzy Sets and Conventional Variables

A key difference between a fuzzy set and a conventional variable is how they are conceptualized and labeled. For example, while it is possible to construct a generic variable such as "years of education," it is impossible to transform this variable directly into a fuzzy set without first designating and defining a target set of cases. In this instance, the researcher might be interested in the set of individuals with at least a high school education or perhaps the set of individuals who are college educated. This example makes it clear that the designation of different target sets dictates different calibration schemes. A person who has one year of college education, for example, has full membership (1.0) in the set of people who are at least high school educated, but this same person clearly has less than full membership in the set of people who are college educated. In a parallel fashion, it is clear that "level of economic development" makes sense as a generic variable, but in order to calibrate it as a fuzzy set, a target set must be specified, for example, the set of developed countries. Notice that this requirement—that the researcher designate a target set—not only structures the calibration of the set but it also provides a direct connection between theoretical discourse and empirical analysis. After all, it is more common for theoretical discourse to be organized around designated sets of cases (e.g., developed countries) than it is for it to be organized around generic variables (e.g., level of economic development).

These examples clarify a crucial feature of fuzzy sets central to their calibration—the fact that in order to calibrate a fuzzy set it is necessary for researchers to distinguish between relevant and irrelevant variation. For example, the difference between an individual who has completed one year of college and an individual who has completed two years of college is irrelevant to the set of individuals with at least

a high school education, for both of these individuals are fully in this set (membership = 1.0). Their one-year difference is simply not relevant to the target set as conceptualized and labeled. When calibrating a fuzzy set, variation that is irrelevant to the set must be truncated so that the resulting membership scores faithfully reflect the target set's label. This requirement also establishes a close connection between theoretical discourse and empirical analysis.

In line with the general theme of this book, a great benefit of using carefully calibrated fuzzy sets is that they permit the utilization of set-theoretic principles in social research. These principles include subset relations (which are central to the analysis of necessity and sufficiency), set intersection (which is central to the study of cases as configurations), set union (which is central to the examination of alternate paths to the same outcome), truth tables (which are used to unravel causal complexity), and so on. These set-theoretic operations are off-limits to researchers who use uncalibrated measures, such as conventional interval- and ratio-scale variables.

Looking Ahead

Chapter 5 explores the calibration of fuzzy sets in more detail, with a practical emphasis. It focuses on the calibration of interval- and ratio-scale variables as fuzzy sets and describes two general methods. The first, the direct method, is based on researcher-specified benchmarks for full membership, full nonmembership, and the crossover point. The second, labeled the indirect method, is based on the researcher's sorting of cases into six categories and the use of a regression estimation procedure to translate raw scores into fuzzy membership scores

5: Calibrating Fuzzy Sets

This chapter sketches two techniques for calibrating conventional interval-scale variables as fuzzy sets, using external standards to structure the calibration. As noted in chapter 4, conventional variables are either uncalibrated or only implicitly calibrated using inductively derived, sample-specific standards—the mean and standard deviation. Fuzzy sets, by contrast, are calibrated using external criteria, which in turn must follow from and conform to the researcher's conceptualization, definition, and labeling of the set in question. External standards can be implemented in two different ways. Using the first, *direct*, method, the researcher specifies the values of an interval scale that correspond to the three qualitative breakpoints that structure a fuzzy set: full membership, full nonmembership, and the crossover point. These three benchmarks are then used to transform the original interval-scale values to fuzzy membership scores. Using the second, *indirect*, method, the external standard used is the researcher's qualitative assessment of the degree to which cases with given scores on an interval scale are members of the target set. The researcher assigns each case into one of six categories and then uses a simple estimation technique to rescale the original measure so that it conforms to these qualitative assessments. The end product of both methods is the fine-grained calibration of the degree of membership of cases in sets, with scores ranging from 0.0 to 1.0. The examples provided in this chapter illustrate the responsiveness of these calibration methods to the researcher's conceptualization of the target set.

Transforming Interval-Scale Variables into Fuzzy Sets

Ideally, the calibration of degree of membership in a set should be based entirely on the researcher's substantive and theoretical knowledge. That is, the collective knowledge base of social scientists should provide the basis for the specification of precise calibrations. For example, armed with an adequate knowledge of development, social scientists should be able to specify the per capita income level that signals full membership in the set of developed countries. However, the social sciences are still in their infancy, and this knowledge base does not exist. Furthermore, the dominance of variable-oriented research, with its paramount focus on mean-centered variation and on covariation as the key to assessing relationships between case aspects, undermines scholarly interest in substantively based thresholds and benchmarks. While the problem of specifying thresholds and benchmarks has not attracted the attention it deserves, it is not a daunting task. The primary requirement for useful calibration is simply sustained attention to the substantive issues at hand (e.g., establishing what constitutes full membership in the set of developed countries).

Despite the imperfections of the existing knowledge base, it is still possible to demonstrate techniques of calibration. All that is lacking are precise "agreed upon standards" for calibrating measures. To the extent possible, the calibrations presented in this chapter are based on the existing theoretical and substantive literature. Still, the focus is on techniques of calibration, and not on the specific empirical benchmarks used to structure calibration.

The techniques presented assume that researchers already have at their disposal conventional interval-scale indicators of their concepts, for example, per capita national income as an indicator of development. The techniques also assume that the underlying concept can be structured and labeled in set-theoretic terms, for example, degree of membership in the set of developed countries. Notice that this labeling requirement moves the investigation in a decidedly case-oriented direction. The set of developed countries identifies specific countries,

while level of development does not. The latter simply identifies a dimension of cross-national variation.

The direct method uses estimates of the log of the odds of full membership in a set as an intermediate step. While this translation route—using estimates of the log odds of full membership—may seem roundabout, the value of the approach will become clear as the demonstration proceeds. For now, consider table 5.1, which shows the different metrics that are used in the demonstration of the direct method. The first column shows various verbal labels that can be attached to differing degrees of set membership, ranging from full nonmembership to full membership. The second column shows the degree of set membership linked to each verbal label. For convenience, degree of membership is rounded to three decimal places. The third column shows the odds of full membership that result from the transformation of the set membership scores (column 2) into the odds of full membership, using the following formula:

$$\text{odds of membership} = (\text{degree of membership}) / [1 - (\text{degree of membership})]$$

The last column shows the natural log of the odds reported in column 3. In effect, columns 2 through 4 are different representations of the same numerical values, using different metrics. For example, the membership score attached to "threshold of full membership" is 0.953. Converting it to an odds value yields 20.09. Calculating the natural log of 20.09 yields a score of 3.0.¹

Working in the metric of log odds is useful because this metric is completely symmetric around 0.0 (an odds of 50/50) and suffers neither floor nor ceiling effects. Thus, for example, if a calibration technique returns a value in the log of odds that is either a very large positive number or a very large negative number, its translation to degree of membership stays within the 0.0 to 1.0 bounds, which is a core requirement of fuzzy membership scores. The essential task of calibration

1. The values shown for degree of membership in column 2 have been adjusted (e.g., using 0.993 instead of 0.99 for full membership) so that they correspond to simple, single-digit entries in column 4.

Table 5.1: Mathematical translations of verbal labels

Verbal label	Degree of membership	Associated odds	Log odds of full membership
Full membership	0.993	148.41	5.0
Threshold of full membership	0.953	20.09	3.0
Mostly in	0.881	7.39	2.0
More in than out	0.622	1.65	0.5
Crossover point	0.500	1.00	0.0
More out than in	0.378	0.61	-0.5
Mostly out	0.119	0.14	-2.0
Threshold of full nonmembership	0.047	0.05	-3.0
Full nonmembership	0.007	0.01	-5.0

using the direct method is to transform interval-scale variables into the log odds metric in a way that respects the verbal labels shown in column 1 of table 5.1.²

It is important to note that the set membership scores that result from these transformations (ranging from 0.0 to 1.0) are *not* probabilities, but instead should be seen simply as transformations of interval scales into degree of membership in the target set. In essence, a fuzzy membership score attaches a *truth value*, not a probability, to a statement (e.g., the statement that a country is in the set of developed countries). The difference between a truth value and a probability is easy to grasp, and it is surprising that so many scholars confuse the two. For example, the *truth value* of the statement "beer is a deadly poison" is perhaps about 0.05—that is, this statement is almost but not completely out of the set of true statements, and beer is consumed freely, without concern, by millions and millions of people every day. However, these same millions would be quite unlikely to consume a liquid that has a 0.05 probability of being a deadly poison, with death the outcome, on average, in one in twenty beers.

2. The procedures for calibrating fuzzy membership scores presented in this chapter are mathematically incapable of producing set membership scores of exactly 1.0 or 0.0. These two membership scores would correspond to positive and negative infinity, respectively, for the log of the odds. Instead, scores that are greater than 0.95 may be interpreted as (virtually) full membership in the target set, and scores that are less than 0.05 may be interpreted as (virtually) full nonmembership.

The Direct Method of Calibration

The starting point of any set calibration is clear specification of the target set. The focus of this demonstration is the set of developed countries, and the goal is to use per capita national income data to calibrate degree of membership in this set. Altogether, 136 countries are included in the demonstration; table 5.2 presents data on 24 of these 136 countries, which were chosen to represent a wide range of national income values.

Table 5.2: Calibrating degree of membership in the set of developed countries: Direct method

Country	National income (US\$)	Deviations from crossover	Scalars	Product	Degree of membership
Switzerland	40,110	35,110.00	.0002	7.02	1.00
United States	34,400	29,400.00	.0002	5.88	1.00
Netherlands	25,200	20,200.00	.0002	4.04	0.98
Finland	24,920	19,920.00	.0002	3.98	0.98
Australia	20,060	15,060.00	.0002	3.01	0.95
Israel	17,090	12,090.00	.0002	2.42	0.92
Spain	15,320	10,320.00	.0002	2.06	0.89
New Zealand	13,680	8,680.00	.0002	1.74	0.85
Cyprus	11,720	6,720.00	.0002	1.34	0.79
Greece	11,290	6,290.00	.0002	1.26	0.78
Portugal	10,940	5,940.00	.0002	1.19	0.77
Korea, Rep.	9,800	4,800.00	.0002	.96	0.72
Argentina	7,470	2,470.00	.0002	.49	0.62
Hungary	4,670	-330.00	.0012	-0.40	0.40
Venezuela	4,100	-900.00	.0012	-1.08	0.25
Estonia	4,070	-930.00	.0012	-1.12	0.25
Panama	3,740	-1,260.00	.0012	-1.51	0.18
Mauritius	3,690	-1,310.00	.0012	-1.57	0.17
Brazil	3,590	-1,410.00	.0012	-1.69	0.16
Turkey	2,980	-2,020.00	.0012	-2.42	0.08
Bolivia	1,000	-4,000.00	.0012	-4.80	0.01
Cote d'Ivoire	650	-4,350.00	.0012	-5.22	0.01
Senegal	450	-4,550.00	.0012	-5.46	0.00
Burundi	110	-4,890.00	.0012	-5.87	0.00

The direct method uses three important qualitative anchors to structure calibration: the threshold for full membership, the threshold for full nonmembership, and the crossover point (see Ragin 2000 and chapter 2 of this book). The crossover point is the value of the interval-scale variable where there is maximum ambiguity as to whether a case is more in or more out of the target set. For the purpose of this demonstration, I use a per capita national income value of \$5,000 as the crossover point. An important step in the direct method of calibration is to calculate the deviations of raw scores (shown in column 1) from the crossover point designated by the investigator (\$5,000 in this example). These values are shown in column 2 of table 5.2. Negative scores indicate that a case is more out than in the target set, while positive scores signal that a case is more in than out.

For the threshold of full membership in the target set, I use a per capita national income value of \$20,000, which is a deviation score of \$15,000 (compare columns 1 and 2 of table 5.2). This value corresponds to a set membership score of .95 and a log odds of 3.0. Thus, cases with national income per capita of \$20,000 or greater (i.e., deviation scores of \$15,000 or greater) are considered fully in the target set, with set membership scores ≥ 0.95 and log odds of membership ≥ 3.0 . In the reverse direction, the threshold for full nonmembership in the target set is \$2,500, which is a deviation score of $-$2,500$. This national income value corresponds to a set membership score of .05 and a log odds of -3.0 . Thus, cases with national income per capita of \$2,500 or lower (i.e., deviation scores of $-$2,500$ or lower) are considered fully out of the target set, with set membership scores $\leq .05$ and log odds of membership ≤ -3.0 .

Once these three values (the two thresholds and the crossover point) have been selected, it is possible to calibrate degree of membership in the target set. The main task at this point is to translate the crossover centered national income data (column 2) into the metric of log odds, utilizing the external criteria that have been operationalized in the three qualitative anchors. For deviation scores above the crossover point, this translation can be accomplished by multiplying the relevant deviation scores (in column 2 of table 5.2) by the ratio of the

log odds associated with the verbal label for the threshold of full membership (3.0) to the deviation score designated as the threshold of full membership (i.e., $\$20,000 - \$5,000 = \$15,000$). This ratio is $3/15,000$, or 0.0002. For deviation scores below the crossover point, this translation can be accomplished by multiplying the relevant deviation scores (in column 2 of table 5.2) by the ratio of the log odds associated with the verbal label for the threshold of full nonmembership (-3.0) to the deviation score designated as the threshold of full nonmembership ($\$2,500 - \$5,000 = -\$2,500$). This ratio is $-3/-2,500$, or 0.0012. These two scalars are shown in column 3, and the products of columns 2 and 3 are shown in column 4.³ Thus, column 4 shows the translation of income deviation scores into the log odds metric, using the three qualitative anchors to structure the transformation via the two scalars.

The values in column 4, in effect, are per capita national income values that have been rescaled into values reflecting the log odds of membership in the set of developed countries, in a manner that strictly conforms to the values attached to the three qualitative anchors—the threshold of full membership, the threshold of full nonmembership, and the crossover point. Thus, the values in column 4 are not mere mechanistic rescalings of national income, for they reflect the imposition of external criteria via the three qualitative anchors. The use of such external criteria is the hallmark of measurement calibration (see chapter 4).

It is a small step from the log odds reported in column 4 to the degree of membership values reported in column 5. It is necessary simply to apply the standard formula for converting log odds to scores that range from 0.0 to 1.0, namely:

$$\text{degree of membership} = \exp(\log \text{ odds}) / [1 + \exp(\log \text{ odds})]$$

where exp represents the exponentiation of log odds to simple odds.⁴ Note that the membership values reported in the last column of table

3. These two scalars constitute the slopes of the two lines extending from the origin (0,0) to the two threshold points (15,000,3) and (-2,500,-3) in the plot of the deviations of national income from the crossover point (X axis) against the log odds of full membership in the set of developed countries (Y axis).

4. These procedures may seem forbidding. For the mathematically disinclined, I note that the complex set of computational steps depicted in table 5.2 can be accomplished




Figure 5.1 Plot of degree of membership in the set of developed countries against national income per capita: Direct method

5.2 strictly conform to the distribution dictated by the three qualitative anchors. That is, the threshold for full membership (0.95) is pegged to an income per capita value of \$20,000; the crossover point (0.50) is pegged to an income of \$5,000; and so on. For further illustration of the results of the direct method, consider figure 5.1, which shows a plot of degree of membership in the set of developed countries against per capita national income, using data on all 136 countries included in this demonstration. As the plot shows, the line flattens as it approaches 0.0 (full nonmembership) and 1.0 (full membership), consistent with the conceptualization of degree of set membership. What the plot does not reveal is that most of the world's countries are in the lower-left corner of the plot, with low national incomes and full exclusion from the set of developed countries (i.e., set membership scores < 0.05).

with a simple *compute* command using the software package *fsQCA* (Ragin, Drass, and Davey 2007).

Table 5.3: Calibrating degree of membership in the set of moderately developed countries: Direct method

Country	National income (US\$)	Deviations from crossover	Scalars	Product	Degree of membership
Switzerland	40,110	37,610	0.0006	22.57	1.00
United States	34,400	31,900	0.0006	19.14	1.00
Netherlands	25,200	22,700	0.0006	13.62	1.00
Finland	24,920	22,420	0.0006	13.45	1.00
Australia	20,060	17,560	0.0006	10.54	1.00
Israel	17,090	14,590	0.0006	8.75	1.00
Spain	15,320	12,820	0.0006	7.69	1.00
New Zealand	13,680	11,180	0.0006	6.71	1.00
Cyprus	11,720	9,220	0.0006	5.53	1.00
Greece	11,290	8,790	0.0006	5.27	0.99
Portugal	10,940	8,440	0.0006	5.06	0.99
Korea, Rep	9,800	7,300	0.0006	4.38	0.99
Argentina	7,470	4,970	0.0006	2.98	0.95
Hungary	4,670	2,170	0.0006	1.30	0.79
Venezuela	4,100	1,600	0.0006	0.96	0.72
Estonia	4,070	1,570	0.0006	0.94	0.72
Panama	3,740	1,240	0.0006	0.74	0.68
Mauritius	3,690	1,190	0.0006	0.71	0.67
Brazil	3,590	1,090	0.0006	0.65	0.66
Turkey	2,980	480	0.0006	0.29	0.57
Bolivia	1,000	-1,500	0.0020	-3.00	0.05
Cote d'Ivoire	650	-1,850	0.0020	-3.70	0.02
Senegal	450	-2,050	0.0020	-4.10	0.02
Burundi	110	-2,390	0.0020	-4.78	0.01

To illustrate the importance of external criteria to calibration, consider using the same national income data (column 1 of table 5.2) to calibrate degree of membership in the set of countries that are "at least moderately developed." Because the definition of the target set has changed, so, too, must the three qualitative anchors. Appropriate anchors for the set of at least moderately developed countries are a crossover value of \$2,500; a threshold of full membership value of \$7,500; and a threshold of full nonmembership value of \$1,000. The appropriate scalars in this example are 3/5,000 for cases above the crossover

value and -3/-1,500 for cases below the crossover value. The complete procedure is shown in table 5.3, using the same cases as in table 5.2.

The key point of contrast between tables 5.2 and 5.3 is shown in the last column, the calibrated membership scores. For example, with a national income per capita of \$2,980, Turkey has a membership of 0.08 in the set of developed countries. Its membership in the set of at least moderately developed countries, however, is 0.57, which places it above the crossover point. Notice, more generally, that in table 5.3 many more cases register set membership scores close to 1.0, consistent with the simple fact that more countries have high membership in the set of countries that are at least moderately developed than in tables 5.2 and 5.3 underscores both the knowledge-dependent nature of calibration and the impact of applying different external standards to the same measure (per capita national income). Again, the key to understanding calibration is to grasp the importance of external criteria, which are based, in turn, on the substantive and theoretical knowledge that researchers bring to their research.

The Indirect Method of Calibration

In contrast to the direct method, which relies on the specification of the numerical values linked to three qualitative anchors, the indirect method relies on the researcher's broad groupings of cases according to their degree of membership in the target set. In essence, the researcher performs an initial sorting of cases into different levels of membership, assigns these different levels preliminary membership scores, and then refines these membership scores using the interval-scale data.

Consider again the data on per capita national income, this time presented in table 5.4. The first and most important step in the indirect method is to categorize cases in a qualitative manner, according to their presumed degree of membership in the target set. These qualitative groupings can be preliminary and open to revision. However, they should be based as much as possible on existing theoretical and

Table 5.4: Calibrating degree of membership in the set of developed countries: indirect method

Country	National income (US\$)	Qualitative coding	Predicted value
Switzerland	40,110	1.00	1.000
United States	34,400	1.00	1.000
Netherlands	25,200	1.00	1.000
Finland	24,920	1.00	1.000
Australia	20,060	1.00	1.000
Israel	17,090	0.80	0.999
Spain	15,320	0.80	0.991
New Zealand	13,680	0.80	0.977
Cyprus	11,720	0.80	0.991
Greece	11,290	0.80	0.887
Portugal	10,940	0.80	0.868
Korea, Rep	9,800	0.80	0.852
Argentina	7,470	0.60	0.793
Hungary	4,670	0.60	0.653
Venezuela	4,100	0.40	0.495
Estonia	4,070	0.40	0.465
Panama	3,740	0.20	0.463
Mauritius	3,690	0.20	0.445
Brazil	3,590	0.20	0.442
Turkey	2,980	0.20	0.436
Bolivia	1,000	0.00	0.397
Cote d'Ivoire	650	0.00	0.053
Senegal	450	0.00	0.002
Burundi	110	0.00	0.000

substantive knowledge. The six key qualitative categories used in this demonstration are the following:⁵

1. In the target set (membership = 1.0)
2. Mostly but not fully in the target set (membership = 0.8)
3. More in than out of the target set (membership = 0.6)

5. Of course, other coding schemes are possible, using as few as three qualitative categories. The important point is that the scoring of these categories should reflect the researcher's initial estimate of each case's degree of set membership. These qualitative assessments provide the foundation for finer-grained calibration.

4. More out than in the target set (membership = 0.4)
5. Mostly but not fully out of the target set (membership = 0.2)
6. Out of the target set (membership = 0.0)

These categorizations are shown in column 2 of table 5.4, using explicit numerical values to reflect preliminary estimates of degree of set membership. These six numerical values are not arbitrary, of course, but are chosen as rough estimates of degree of membership specific to each qualitative grouping. The goal of the indirect method is to rescale the interval-scale indicator to reflect knowledge-based, qualitative groupings of cases, categorized according to degree of set membership. These qualitative interpretations of cases must be grounded in substantive knowledge. The stronger the empirical basis for making qualitative assessments of set membership, the more precise the calibration of the values of the interval-scale indicator as set membership scores.

Note that the qualitative groupings implemented in table 5.4 have been structured so that they utilize roughly the same criteria used to structure the calibrations shown in table 5.2. That is, countries with national income per capita greater than \$20,000 have been coded as fully in the set of developed countries; countries with income per capita greater than \$5,000 have been coded as more in than out; and so on. By maintaining fidelity to the qualitative anchors used in table 5.2, it is possible to compare the results of the two methods. The direct method utilizes precise specifications of the key benchmarks, while the indirect method requires only a broad classification of cases.

The next step is to use the two series reported in columns 1 and 2 of table 5.4 to estimate the predicted qualitative coding of each case, using per capita national income as the independent variable and the qualitative codings as the dependent variable. The best technique for this task is a fractional logit model, which is implemented in Stata software in the `fracpoly` (fractional polynomial) regression procedure.⁶

6. In Stata, this estimation procedure can be implemented using the commands `fracpoly glm qualcode interv`, `family(binomial) link(logit)` and then `predict fpred`, where `qualcode` is the variable that implements the researcher's six-value coding of set membership, as shown in table 5.4; `interv` is the name of the interval-scale variable

The predicted values resulting from this analysis are reported in column 3 of table 5.4. The reported values are based on an analysis using all 136 cases, not the subset of 24 presented in the table. The predicted values, in essence, constitute estimates of fuzzy membership in the set of developed countries based on per capita national income (column 1) and the qualitative analysis that produced the codings shown in column 2.

Comparison of the set membership scores in column 5 of table 5.2 (direct method) and column 3 of table 5.4 (indirect method) reveals great similarities, but also some important differences. First, notice that table 5.2 faithfully implements \$20,000 as the threshold for full membership in the set of developed countries (0.95). In table 5.4, however, this threshold value drops well below New Zealand's score (\$13,680). Second, observe that the indirect method reveals a large gap separating Turkey (0.397) and the next case, Bolivia (0.053). Using the direct method, however, this gap is much narrower, with Turkey at 0.08 and Bolivia at 0.01. These differences, which arise despite the use of the same general criteria, follow from the indirectness of the second method and its necessary reliance on regression estimation. Still, if researchers lack the external criteria used in the direct method, the comparison of tables 5.2 and 5.4 confirms that the indirect method produces useful set membership scores.

Using Calibrated Measures

Calibrated measures have many uses. They are especially useful in evaluating theory that is formulated in terms of set relations. As noted in chapter 1, while some social science theory is strictly mathematical, the vast majority of it is verbal. Verbal theory, in turn, is formulated almost entirely in terms of set relations (Ragin 2000, 2006b). Unfortunately, social scientists have been slow to recognize this fact. Consider,

that is used to generate fuzzy membership scores; and `fzpred` is the predicted value showing the resulting fuzzy membership scores. I thank Steve Vaisey for pointing out the robustness of this estimation technique.

for example, the statement, "the developed countries are democratic." As in many statements of this type, the assertion is essentially that instances of the set mentioned first (developed countries) constitute a subset of instances of the set mentioned second (democracies). (It is common in English to state the subset first, as in the statement "ravens are black.") Close examination of most social science theories reveals that they are composed largely of statements describing set relations, such as the subset relation. These set relations, in turn, may involve a variety of different types of empirical connections—descriptive, constitutive, or causal, among others.

The set relation with developed countries as a subset of democratic countries, described above, is also compatible with a specific type of causal argument, namely, that development is sufficient but not necessary for democracy. In arguments of this type, if the cause (development) is present, then the outcome (democracy) should also be present. However, instances of the outcome (democracy) without the cause (development) do not count against or undermine the argument that development is sufficient for democracy (even though such cases dramatically undermine the correlation). Rather, these instances of the outcome without the cause are due to the existence of alternate routes or recipes for that outcome (e.g., the imposition of a democratic form of government by a departing colonial power). Thus, in situations where instances of a causal condition constitute a subset of instances of the outcome, a researcher may claim that the cause is sufficient but not necessary for the outcome.⁷

Before the advent of fuzzy sets (Zadeh 1965, 1972, 2002; Lakoff 1973), many social scientists disdained the analysis of set-theoretic relations because such analyses required the use of categorical-scale variables (i.e., conventional binary or crisp sets), which, in turn, often necessitated the dichotomization of interval and ratio scales. For example, using crisp sets to assess a set-theoretic statement about de-

7. As always, claims of this type cannot be based simply on the demonstration of the subset relation. Researchers should marshal as much corroborating evidence as possible when making any type of causal claim.

veloped countries, a researcher might be required to categorize countries into two groups, developed and not developed, using per capita national income. Such practices are often criticized because researchers may manipulate breakpoints when dichotomizing interval- and ratio-scale variables in ways that enhance the consistency of the evidence with a set-theoretic claim. However, as demonstrated here, it is possible to calibrate degree of membership in sets and thereby avoid arbitrary dichotomizations.

As shown in chapter 2, the fuzzy subset relation is established by demonstrating that membership scores in one set are consistently less than or equal to membership scores in another. In other words, if, for every case, degree of membership in set X is less than or equal to degree of membership in set Y , then set X is a subset of set Y . Of course, social science data are rarely perfect, and some allowance must be made for these imperfections. It is possible to assess the degree of consistency of empirical evidence with the subset relation using the simple formula described in chapter 3:

$$\text{Consistency } (X_1 \leq Y_1) = \Sigma[\min(X_i, Y_i)] / \Sigma(X_i)$$

where X_1 is degree of membership in set X ; Y_1 is degree of membership in set Y ; $(X_i \leq Y_i)$ is the subset relation in question; and min dictates selection of the lower of the two scores.

For illustration, consider the consistency of the empirical evidence with the claim that the set of developed countries (as calibrated in table 5.2) constitutes a subset of the set of democracies, using data on all 136 countries. For this demonstration, the Polity IV democracy/autocracy measure is used, which ranges from -10 to $+10$. (This measure is used because of its popularity, despite its many shortcomings. See, e.g., Goertz 2006, ch. 4.) The calibration of membership in the set of democracies, using the direct method, is shown in table 5.5. Polity scores for 24 of the 136 countries included in the calibration are presented in column 1 of table 5.5. These specific cases were selected in order to provide a range of polity scores. Column 2 shows deviations from the crossover point (a polity score of 2), and the column 3 shows the scalars used to transform the polity deviation scores into the metric of log odds of membership in the set of democracies. The threshold

Table 5.5: Calibrating degree of membership in the set of democratic countries: Direct method

Country	Polity score	Deviations from crossover	Scalars	Product	Degree of membership
Norway	10	8.00	0.43	3.43	0.97
United States	10	8.00	0.43	3.43	0.97
France	9	7.00	0.43	3.00	0.95
Korea, Rep.	8	6.00	0.43	2.57	0.93
Colombia	7	5.00	0.43	2.14	0.89
Croatia	7	5.00	0.43	2.14	0.89
Bangladesh	6	4.00	0.43	1.71	0.85
Ecuador	6	4.00	0.43	1.71	0.85
Albania	5	3.00	0.43	1.29	0.78
Armenia	5	3.00	0.43	1.29	0.78
Nigeria	4	2.00	0.43	0.86	0.70
Malaysia	3	1.00	0.43	0.43	0.61
Cambodia	2	0.00	0.60	0.00	0.50
Tanzania	2	0.00	0.60	0.00	0.50
Zambia	1	-1.00	0.60	-0.60	0.35
Liberia	0	-2.00	0.60	-1.20	0.23
Tajikistan	-1	-3.00	0.60	-1.80	0.14
Jordan	-2	-4.00	0.60	-2.40	0.08
Algeria	-3	-5.00	0.60	-3.00	0.05
Rwanda	-4	-6.00	0.60	-3.60	0.03
Gambia	-5	-7.00	0.60	-4.20	0.01
Egypt	-6	-8.00	0.60	-4.80	0.01
Azerbaijan	-7	-9.00	0.60	-5.40	0.00
Bhutan	-8	-10.00	0.60	-6.00	0.00

of full membership in the set of democracies is a polity score of 9, yielding a scalar of 3/7 for cases above the crossover point; the threshold of full nonmembership in the set of democracies is a polity score of -3, yielding a scalar of -3/-5 for cases below the crossover point. Column 4 shows the product of the deviation scores and the scalars, while column 5 reports the calibrated membership scores, using the procedures previously described (see the discussion surrounding table 5.2).

Applying the formula for set-theoretic consistency described above to all 136 countries, the consistency of the evidence with the argument

that the set of developed countries constitutes a subset of the set of democracies is 0.99. (1.0 indicates perfect consistency). Likewise, the consistency of the evidence with the argument that the set of at least moderately developed countries (as calibrated in table 5.3) constitutes a subset of the set of democratic countries is 0.95. In short, both subset relations are highly consistent, providing ample support for both statements ("developed countries are democratic" and "countries that are at least moderately developed are democratic"). Likewise, both analyses support the argument that development is sufficient but not necessary for democracy. Note, however, that the set of at least moderately developed countries is a much more inclusive set, with higher average membership scores than the set of developed countries. It thus offers a more demanding test of the underlying argument. The greater the average membership in a causal condition, the more difficult it is to satisfy the inequality indicating the subset relation ($X_i \leq Y_i$). The two formulations also differ substantially in their set theoretic "coverage." Coverage is a gauge of empirical importance or weight (see Ragin 2006b and chapter 3 of this book). It shows the proportion of the sum of the outcome membership scores (in this example, the set of democratic countries) that is "covered" by a causal condition. The coverage of democratic countries by developed countries is 0.35, while the coverage of democratic countries by at least moderately developed countries is substantially more, 0.52. These results indicate that the latter gives a much better account of degree of membership in the set of democratic countries. Thus, using set-theoretic methods, it is possible to demonstrate that membership in the set of countries with a moderate level of development is sufficient for democracy; membership in the set of fully developed countries is not required.

As explained in chapter 1, it is very difficult to evaluate set-theoretic arguments using correlational methods. The three main sources of this difficulty are as follows:

1. Set-theoretic statements are about kinds of cases; correlations concern relationships between variables. The statement that developed countries are democratic (i.e., that they constitute a subset of democratic countries) invokes cases, not dimensions of cross-national

variation. This focus on cases as instances of concepts follows directly from the set-theoretic nature of social science theory. The computation of a correlation, by contrast, is premised on an interest in assessing how well dimensions of variation parallel each other across a sample or population, not on an interest in a set of cases per se. To push the argument even further, a data set might not include a single developed country or a single democratic country, yet a correlational researcher could still compute a correlation between degree of development and degree of democracy. Note, however, that this data set would be completely inappropriate for a test of the argument that the developed countries are democratic, for it contains neither developed countries nor democratic countries.

2. Correlational arguments are fully symmetric, while set-theoretic arguments are almost always asymmetric. The correlation between development and democracy (treating both as conventional variables) is weakened by the fact that there are many less-developed countries that are democratic. However, such cases do not challenge the set-theoretic claim or weaken its consistency. The theoretical argument in question addresses the qualities of developed countries—that they are democratic—and does not make *specific* claims about relative differences between less-developed and more-developed countries in their degree of democracy. Again, set-theoretic analysis is faithful to verbal formulations, which are typically asymmetric; correlation is not.

3. Correlations are insensitive to the calibrations implemented by researchers. The contrast between tables 5.2 and 5.3 is meaningful from a set-theoretic point of view. The set represented in table 5.3 is more inclusive and thus provides a more demanding set-theoretic test of the connection between development and democracy. From a correlational perspective, however, there is little difference between the two ways of representing development. Indeed, the Pearson correlation between fuzzy membership in the set of developed countries and fuzzy membership in the set of at least moderately developed countries is 0.911. Thus, from a strictly correlational viewpoint, the difference between these two fuzzy sets is slight. From a set-theoretic viewpoint, however, they are quite different, for the set-theoretic con-

erage of democracy by developed is only 0.35, while the coverage of democracy by at least moderately developed is 0.52. The insensitivity of correlation to calibration follows directly from the fact that correlation is computationally reliant on deviations from the fact that correlation is sample-specific measure of central tendency—the mean. For this reason, correlation is incapable of analyzing set-theoretic relations and, correspondingly, cannot be used to assess causal sufficiency or necessity.

Conclusion

This chapter demonstrates both the power of fuzzy sets and the centrality of calibration to their fruitful use. It is important to be able to assess not only “more versus less” (uncalibrated measurement) but also “a lot versus a little” (calibrated measurement). The use of calibrated measures grounds social science in substantive knowledge and enhances the relevance of the results of social research to practical and policy issues. Fuzzy sets are especially powerful as carriers of calibration. They offer measurement tools that transcend the quantitative/qualitative divide in the social sciences.

Current practices in quantitative social science undercut serious attention to calibration. These difficulties stem from reliance on the indicator approach to measurement, which requires only variation across sample points and treats all variation as equally meaningful. The limitations of the indicator approach are compounded and reinforced by correlational methods, which are insensitive to calibrations implemented by researchers. Reliance on deviations from the mean tends to neutralize the impact of any direct calibration implemented by the researcher. A further difficulty arises when it is acknowledged that almost all social science theory is set-theoretic in nature and that correlational methods are incapable of assessing set-theoretic relations.

The set-theoretic nature of most social science theory is not generally recognized by social scientists today, nor is the fact that the assessment of set-theoretic arguments and set calibration go hand in hand.

Set theoretic analysis without careful calibration of set membership is an exercise in futility. It follows that researchers need to be faithful to their theories by clearly identifying the target sets that correspond to the concepts central to their theories and by specifying useful external criteria that can be used to guide the calibration of set membership.

Practical Appendix: Using fsQCA to Calibrate Fuzzy Sets (Direct Method)

1. In fuzzy-set qualitative comparative analysis (fsQCA), create or retrieve your data set. For example, you might have an SPSS or Excel file with the relevant interval- or ratio-scale data. Save these files as comma-delimited or tab-delimited files with simple variable names on the first row of the file. Make sure missing data are blank and not assigned a special code (e.g., -999).
2. With your data in the data spreadsheet window of fsQCA, click the Variables menu; then click Compute.
3. In the compute dialogue box, name the target fuzzy set. Select a simple name (two to eight characters), using standard alphanumeric characters and no spaces, dashes, or punctuation.
4. Click calibrate(x.n1,n2,n3) in the Functions menu and then click the up arrow that is next to the word Functions. Next, calibrate(,,) will appear in the Expression field of the dialogue box.
5. Edit the expression so that calibrate(,,) becomes something like calibrate(intvar,25,10,2), where intvar is the name of the existing interval- or ratio-scale variable already in the file, the first number is the value of intvar you have chosen as the threshold for full membership in the target set (fuzzy score = 0.95); the second number is the value of intvar that you have selected for the crossover point (fuzzy score = 0.5), and the third number is the value of intvar that you have selected for the threshold for full nonmembership in the target set (fuzzy score = 0.05).
6. Click OK. Check the data spreadsheet to make sure it came out as you expected. It is possible to sort the original interval-scale vari-

able in descending or ascending order using the pull-down menus. Click any case in the column you want to sort, then click Cases, and then Sort Ascending or Sort Descending. You can then check the corresponding fuzzy scores to see if they conform with your interval- or ratio-scale variable in the manner you intended.