

Set Relations in Social Research: Evaluating Their Consistency and Coverage

Charles C. Ragin

Department of Sociology, University of Arizona, Tucson, AZ 85721
e-mail: cragin@email.arizona.edu

Because of its inherently asymmetric nature, set-theoretic analysis offers many interesting contrasts with analysis based on correlations. Until recently, however, social scientists have been slow to embrace set-theoretic approaches. The perception was that this type of analysis is restricted to primitive, binary variables and that it has little or no tolerance for error. With the advent of “fuzzy” sets and the recognition that even rough set-theoretic relations are relevant to theory, these old barriers have crumbled. This paper advances the set-theoretic approach by presenting simple descriptive measures that can be used to evaluate set-theoretic relationships, especially relations between fuzzy sets. The first measure, “consistency,” assesses the degree to which a subset relation has been approximated, whereas the second measure, “coverage,” assesses the empirical relevance of a consistent subset. This paper demonstrates further that set-theoretic coverage can be partitioned in a manner somewhat analogous to the partitioning of explained variation in multiple regression analysis.

1 Introduction

The bulk of social science theory is verbal in nature. Verbal theories typically describe ideal typical cases or situations, which in turn provide grist for social scientists’ formalization of theory as hypotheses about empirical patterns and connections. Verbal theory is largely set theoretic in nature. Because set relations are the building blocks of verbal statements, they are also the building blocks of most social science theories. Unfortunately, set relations described in theories 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), oriented toward the evaluation of the “net effects” of causal variables (Ragin 2006). This paper is premised on the idea that a theory that is formulated in terms of set relations should be evaluated on its own terms, that is, as statements about set relations, not about correlations.

The reformulation of theoretical arguments describing set relations as correlational hypotheses is one of the most common but dubious practices in all of contemporary social science. Consider the “democratic peace” argument: democracies do not go to war against each other. This statement is essentially a claim that country dyads in which both parties

Author’s note: I thank Gary Goertz for his very generous help with this paper and for his contributions to the ideas I develop. He bears no responsibility for the shortcomings that remain. I also thank Jay Verkuilen and the journal’s anonymous referees for their useful comments and suggestions.

are democratic constitute a perfect (or near-perfect) subset of “nonwarring” country dyads. Of course, the rate of warring may be very low 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, which may be relatively small, but that the rate of warring is “zero or very close to zero” in the set of “democratic dyads.” The fact that democratic dyads constitute a perfect or near-perfect subset of nonwarring dyads signals that this arrangement (international relations between democracies) may be sufficient for peaceful coexistence. Of course, there are many other ways to effect peaceful coexistence, and the correlation between democratic dyad and nonwarring is weak because of these many alternate routes (see Peceny and Beer 2002). The set-theoretic claim that democracies do not go to war against each other, however, is not refuted in any way by the existence of these alternate routes to peace.

In this paper I present simple descriptive measures for evaluating the strength of the empirical support for theoretical arguments describing set relations. I focus primarily on arguments stating that a specific cause or combination of causal conditions constitutes one of several possible paths to an outcome. When this is true, cases displaying the causal combination constitute a subset of the cases displaying the outcome, as in the example just presented. I quantify two aspects of set relations. *Set-theoretic consistency* assesses the degree to which the cases sharing a given condition or combination of conditions (e.g., democratic dyad) agree in displaying the outcome in question (e.g., nonwarring). That is, consistency indicates how closely the 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 any given causal combination may be small. Thus, coverage gauges empirical relevance or importance. These two measures also can be used to evaluate set relations suggesting 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. In this context, consistency assesses the degree to which instances of an outcome agree in displaying the causal condition thought to be necessary, whereas coverage assesses the “relevance” of the causal condition—the degree to which instances of the causal condition are paired with instances of the outcome. My discussion of necessary conditions builds on the work of Gary Goertz (Braumoeller and Goertz 2000; Goertz and Starr 2002; Goertz 2003, 2006).

These measures provide vital tools for refining set-theoretic analysis in the social sciences. They are neither complex nor entirely “new,” for there are many similar and parallel measures in the vast literature on quantitative analysis and statistical methods. What is novel is their explicit coupling with set-theoretic reasoning, especially those forms that are at odds with the inherently symmetric, correlational reasoning that undergirds most forms of conventional quantitative analysis. Thus, this paper not only presents measures, it also addresses key differences between correlational reasoning and set-theoretic reasoning. Note that set-theoretic reasoning is not limited to nominal-scale distinctions, as is so often assumed by many quantitative researchers. As this paper demonstrates, with “fuzzy” sets it is a simple matter to apply set-theoretic reasoning to phenomena that vary by level or degree.

2 Set-Theoretic Consistency

Perfectly consistent set relations are relatively rare in social research. Perfect consistency usually requires small *N*s, macrolevel data, or both. Generally, social scientists are able

to identify only rough subsets because there are almost always exceptions (e.g., a war between two democracies). 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. I first address the evaluation of the consistency of “crisp”-set relations, where a very simple measure suffices, and then turn to fuzzy sets. A crisp set is a conventional binary set with two categories (e.g., presence versus absence of a presidential form of government). A fuzzy set allows calibration of degree of set membership, using scores in the interval from 0.0 to 1.0 (e.g., degree of membership in the set of countries with competitive elections; see Ragin 2000).

A simple, straightforward measure of the consistency of a crisp-set relation with sufficiency is the proportion of cases with a given cause or combination of causes that also display the outcome. For example, if 17 out of the 20 cases displaying a cause or causal combination also display the outcome, then the proportion consistent is 0.85. In general, consistency scores should be as close to 1.0 (perfect consistency) as possible. With observed consistency scores below 0.75, it becomes increasingly difficult on substantive grounds to maintain that a subset relation exists, even a very rough one (see also Ragin 2004).

When assessing consistency, it is important to consider the number of cases. 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, although interesting and 100% 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 subsequent failure if the tally was, say, 17 out of 20, instead of three out of three.¹ Although not 100%, 17 out of 20 (85%) is substantial enough to indicate, to a social scientist at least, that there may be some sort of integral connection.

The assessment of the consistency of fuzzy set relations is more interesting and more challenging than 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). 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. Following the 2000 presidential election, for instance, a disgruntled observer might score the United States well below 1.0 in its membership in this set (e.g., a score of 0.75). Membership scores greater than 0.5 indicate that a case is more in than out, scores close to 1.0 indicate that a case is mostly in, scores close to 0.0 indicate that a case is mostly out, and so on. Full membership (1.0) and full nonmembership (0.0) are understood as qualitative states, not arbitrary values (e.g., the highest and lowest observed scores). Thus, the calibration of membership in a fuzzy set involves both quantitative and qualitative assessment and must be grounded in theoretical and substantive knowledge (Ragin 2000, 2004; Smithson and Verkuilen 2006).

A fuzzy subset relation exists when the membership scores in one set are consistently less than or equal to their membership scores in another. For example, if degree of

¹As 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 1.00 is not significantly greater than a benchmark proportion of 0.65, using an alpha of 0.05. However, a proportion of 0.85 with an N of 20 passes this test.




Fig. 1 Fuzzy subset relation consistent with sufficiency.

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 in the crisp-set example, not all instances of “failure” were instances of “parliamentary government” because there were (hypothetically) many ways to fail. With fuzzy sets a case might have a score of 0.9 in failure but a score of only 0.2 in parliamentary government. As in the crisp example, this case is not inconsistent with the set-theoretic argument because there is more than one way to fail, including paths to failure for countries with weak membership in the set of countries with parliamentary governments. By contrast, a country with a membership score of 0.8 in parliamentary government and a membership of only 0.3 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 Fig. 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. Note that cases in the upper left-hand corner of the plot do not contradict sufficiency, for these are cases that have substantial membership in the outcome due to the operation of causal conditions other than X , and an argument of causal sufficiency 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, there are many more opportunities to violate the subset relation, as the range of permissible Y_i values narrows. Of course, in a conventional correlational analysis, cases in the upper left-hand corner would be considered errors and these cases, 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 fall on or above the main diagonal of the plot, yielding a consistency score of 1.0 (or 100% 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, defined as a set relation (see Ragin, Drass, and Davey 2003; Ragin and Giesel 2003). For example, if significantly greater than 80% of the cases fall 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 Ragin (2000) 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 procedure closely follows the crisp-set template described above. 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 memberships in the causal combination are weighted equally in the calculation, yet they differ substantially in their relevance to the set-theoretic argument. 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 just described (the proportion of cases on or above the main diagonal of the plot).² 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 commonsense 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 paper. This alternate procedure, like the first, differentiates between consistent and inconsistent cases using the main 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

²This criticism applies, in a slightly modified form, to some of the set-theoretic procedures recommended in Goertz (2003, 2006). He suggests using ratios of membership scores, but small membership scores can easily produce large ratios; and thus cases that are low in relevance to a given argument may be weighted the same as cases that are high in relevance.

Table 1 Illustration of a simple fuzzy subset relation

<i>Religious heterogeneity</i>	<i>Weak class voting</i>
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

calculating the raw proportion of consistent cases, the alternate procedure uses fuzzy membership scores.

Consider, for example, the hypothetical set membership data on religious heterogeneity and weak class voting for 12 countries shown in Table 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. Based on this evidence, a researcher could claim that this causal condition (religious heterogeneity) is a subset of the outcome (weak class voting). Thus, religious heterogeneity could be interpreted (hypothetically) as a sufficient condition for weak class voting. As previously noted, however, social science data are rarely this uniform. When there are cases that are inconsistent with the subset relation, it is important to assess the degree to which the empirical evidence is consistent with the set relation in question. For example, suppose the first value of religious heterogeneity in Table 1 was 1.0 instead of 0.7. It would be inconsistent with the set relation because this value exceeds the corresponding outcome membership score, 0.9. Although the set relation would no longer hold consistently across the cases listed in Table 1, it would still be very close to perfect, with 11 out of the 12 cases consistent.

One straightforward measure of set-theoretic consistency using the fuzzy membership scores is simply 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 2003). In Table 1, as presented, the value of this measure is 1 (4.7/4.7) because all the membership scores in column 1 are consistent. If the first value of religious heterogeneity in Table 1 is changed to 1.0, however, consistency drops to 0.8 (4/5). The numerator is 1.0 fuzzy units 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 0.7 in the second 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 mark, the outcome membership score, by a wide margin.³ This adjustment can be accomplished by adding to the

³The formula described here is the one implemented in the fuzzy-truth table algorithm of **fsQCA** (Ragin, Drass, and Davey 2003).

numerator in the formula just sketched the part of each inconsistent causal membership score that is consistent with the outcome. For example, if the first value of religious heterogeneity in Table 1 is changed to 1.00, then most of it is consistent, up to the value of the outcome membership score, 0.9. 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 just sketched prescribes substantial penalties for large inconsistencies. Suppose again that the first value of religious heterogeneity in Table 1 is 1.0, but this time assume that the corresponding value of the outcome, weak class voting, 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 (see also Kosko 1993; Smithson and Verkuilen 2006):

$$\text{Consistency}(\mathbf{X}_i \leq \mathbf{Y}_i) = \sum(\min(\mathbf{X}_i, \mathbf{Y}_i)) / \sum(\mathbf{X}_i),$$

where “min” indicates the selection of the lower of the two values. When the \mathbf{X}_i values are all less than or equal to their corresponding \mathbf{Y}_i values, the consistency score is 1.00; when there are only a few near misses, the score is slightly less than 1.00; and when there are many inconsistent scores, with some \mathbf{X}_i values greatly exceeding their corresponding \mathbf{Y}_i values, consistency drops below 0.5. It is important to point out that when the formula for the calculation of fuzzy set-theoretic consistency just presented is applied to crisp-set data, it returns the simple proportion of consistent cases. Thus, the formula can be applied to crisp and fuzzy data 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. The study of necessary conditions has become an important focus in political science, especially in comparative politics and international relations. In fact, upon close inspection it is clear that many long-standing arguments and hypotheses about macropolitical phenomena address necessary conditions (Goertz 2002). 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, $\mathbf{Y}_i \leq \mathbf{X}_i$. This inequality is the reverse of the inequality defining the consistency of the sufficient condition relationship. Thus, a simple measure of the consistency of the subset relationship indicating necessity is:

$$\text{Consistency}(\mathbf{Y}_i \leq \mathbf{X}_i) = \sum(\min(\mathbf{X}_i, \mathbf{Y}_i)) / \sum(\mathbf{Y}_i).$$

When all \mathbf{Y}_i values are less than or equal to their corresponding \mathbf{X}_i values, this formula returns a value of 1.0. When many \mathbf{Y}_i exceed their corresponding \mathbf{X}_i values by wide margins, it returns a value less than 0.5.

For illustration, consider Fig. 2, which depicts a subset pattern compatible with a necessary condition relation. In this figure the outcome (\mathbf{Y}) is a subset of the causal




Fig. 2 Fuzzy subset relation consistent with necessity.

condition (\mathbf{X}); thus, all \mathbf{Y}_i values are less than or equal to their corresponding \mathbf{X}_i values. Note that cases in the lower right-hand 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. Of course, in a conventional correlational analysis cases in the lower right-hand corner would be considered errors, which in turn would undermine the correlation between \mathbf{X} and \mathbf{Y} . Note, however, that when membership in \mathbf{X} is low, membership in \mathbf{Y} also must be low. Thus, in the low range of \mathbf{X} there are many opportunities to violate the subset relation, with only a narrow range of permissible \mathbf{Y}_i values.

As Smithson and Verkuilen (2006) demonstrate, the measure of set-theoretic consistency just described is influenced by the average membership scores in sets \mathbf{X} and \mathbf{Y} . Imagine a causal condition (\mathbf{X}) with very low average membership scores and an outcome (\mathbf{Y}) with very high average membership scores. It may appear that there is support for sufficiency because the set relation $\mathbf{X}_i \leq \mathbf{Y}_i$ holds, when in fact this evidence in favor of sufficiency is primarily the product of skewed membership scores. Researchers, therefore, should be cognizant of the potential impact of skewed membership scores when assessing the consistency of either set-theoretic relation ($\mathbf{X}_i \leq \mathbf{Y}_i$ or $\mathbf{Y}_i \leq \mathbf{X}_i$). The impact of skewed membership scores on the assessment of set-theoretic consistency points to the importance of careful calibration of fuzzy membership scores (Ragin 2000). These scores should reflect researchers' best assessments of the degree of membership of cases in well-defined sets. Researchers should pay close attention to the three calibration anchors (the thresholds for full membership and for full nonmembership and the location of the crossover point, 0.5 membership) when assigning these scores. Fuzzy set calibration is much more demanding than conventional social science measurement and requires thorough utilization of theoretical and substantive knowledge coupled with careful concept formation and elaboration (see Ragin 2000; Goertz 2005). Because calibration is central to fuzzy set

analysis, the procedures researchers use to assign membership scores should be open, transparent, and replicable.

Perhaps even more important is the fact 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 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 integral connection.

3 Set-Theoretic Coverage

When researchers allow for equifinality (Mackie 1965; George 1979; George and Bennett 2005) and causal complexity, a common finding is that a given outcome may result from several different combinations of conditions, with each combination sufficient but not necessary for the outcome. These combinations are generally understood as alternate paths or recipes for the outcome, and they are treated as logically equivalent (i.e., as substitutable). Still, it is common in crisp-set analyses to assess the proportion of instances 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 for crisp sets and is a clear 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.⁴

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% of these individuals are able to avoid poverty. Perfect set-theoretic consistency is unusual with individual-level data, but 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 of the set of individuals who successfully avoid poverty is minuscule.

Although there is often a trade-off between consistency and coverage, it is important to understand that it is reasonable to calculate coverage only after establishing that a set relation is consistent.⁵ 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 calculation has different meanings depending on the context of the calculation. Thus, it is important to adhere to the protocol described here for

⁴Note that coverage gauges only empirical importance, not theoretic 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.

⁵In the truth table algorithm of fsQCA (Ragin 1987; De Meur and Rihoux 2002; Ragin, Drass, and Davey 2003), the assessment of set-theoretic coverage comes at the end of the procedure, after combinations of causal conditions that are consistent subsets of the outcome have been identified, based on criteria specified by the investigator.

Table 2 Cross-tabulation of poverty status and educational achievement: preliminary frequencies

	<i>Low/average educational achievement</i>	<i>High educational achievement</i>
Not in poverty	a. 3046	b. 1474
In poverty	c. 625	d. 55

the results of assessments of consistency and coverage to be meaningful: consistency must be established before coverage can be assessed.


For illustration of the general idea of coverage consider Table 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 third 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 pathway 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 achievements—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, which has the same total number of cases as Table 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 is 0.967, about the same as in Table 2 (0.964). Thus, from a set-theoretic viewpoint, the evidence is again highly consistent. But how important is this path, using the hypothetical frequencies presented in Table 3? This can be ascertained by computing the proportion of cases avoiding poverty that are covered by the set-theoretic argument, which is only 0.0325 (147/4520). Thus, in Table 3 the set-theoretic pattern is highly consistent, but coverage is very low, indicating (hypothetically) that having high educational achievement is not an important path to the outcome, avoiding poverty.

The procedures for calculating coverage using fuzzy sets parallel the computations for crisp sets just presented. Another way to understand the calculation of coverage using conventional binary sets (cell b divided by the sum of cells a and b) is to visualize Table 2 as a Venn diagram, showing a subset relationship, as in Fig. 3. The basic idea behind the calculation of coverage is to assess the degree to which the smaller set (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

Table 3 Cross-tabulation of poverty status and educational achievement: altered frequencies

	<i>Low/average educational achievement</i>	<i>High educational achievement</i>
Not in poverty	a. 4373	b. 147
In poverty	c. 675	d. 5



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

Fig. 3 Venn diagram illustrating concept of coverage using hypothetical data (from Table 2).

(representing the outcome). The calculation of the size of the overlap of two fuzzy sets is given by:

$$\text{Overlap} = \sum (\min(\mathbf{X}_i, \mathbf{Y}_i)),$$

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 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, $\sum(\mathbf{Y}_i)$. This calculation parallels the simple counting of the number of cases in a set (e.g., the set 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 (\mathbf{Y}):

$$\text{Coverage}(\mathbf{X}_i \leq \mathbf{Y}_i) = \frac{\sum (\min(\mathbf{X}_i, \mathbf{Y}_i))}{\sum (\mathbf{Y}_i)}.$$

In short, the formula for coverage substitutes the $\sum(\mathbf{Y}_i)$ for $\sum(\mathbf{X}_i)$ in the denominator of the formula for the consistency of $\mathbf{X}_i \leq \mathbf{Y}_i$.

Observe that this formula is identical to the formula for the consistency of \mathbf{Y}_i as a subset of \mathbf{X}_i ($\mathbf{Y}_i \leq \mathbf{X}_i$) presented in the discussion of the assessment of the consistency of a necessary conditions relationship. Recall, however, that in the context of sufficiency, the coverage of \mathbf{Y}_i by \mathbf{X}_i (i.e., the formula just presented) is calculated only after it has been established that \mathbf{X} is a consistent subset of \mathbf{Y} . Thus, the purpose of the calculation in the context of sufficiency is to assess the magnitude of \mathbf{X} relative to \mathbf{Y} , given that most if not all




Fig. 4 Scatterplot illustrating the concept of coverage.

X_i values are less than or equal to their corresponding Y_i values. When the goal is to assess the consistency of Y as a subset of X , however, the expectation is that most Y_i values will be less than or equal to their corresponding X_i values—indicating a possible necessary conditions relationship. Indeed, if this is not the case, then the result will be a consistency score (for $Y_i \leq X_i$) that is substantially less than 1.0, indicating that Y is not a subset of X .⁶

Figure 4 depicts the concept of coverage relevant to the fuzzy subset relation, with $X_i \leq Y_i$. As in Fig. 1, condition X is a subset of outcome Y . Points below the main diagonal constitute violations of 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 an X_i 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 Fig. 4 are above the main diagonal, and thus consistent with $X_i \leq Y_i$. When these X_i values are small relative to their corresponding Y_i values, they are closer to the Y axis than to the main diagonal. Although these points are consistent with the subset relation depicted in the figure, they contribute relatively little to coverage, especially when the Y_i values are above 0.5. The dotted horizontal lines in the figure show the portions of the X_i values counted as consistent; these values are 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 portions of set Y that are 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 paper. A key focus in his work is the distinction between trivial and nontrivial necessary conditions. A trivial

⁶In this context, it would be possible to find a $\sum(\min(X_i, Y_i))$ that is close to $\sum(Y_i)$ —thus yielding a very high coverage score—only if the values of X_i are roughly equal to their corresponding Y_i values. This situation would correspond to a close *coincidence* of the two sets. Set coincidence is not the same as the 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 plot. A simple measure of the degree to which the membership scores in two sets coincide is $\sum(\min(X_i, Y_i)) / \sum(\max(X_i, Y_i))$, where “max” indicates using the larger of the two scores. See also Smithson and Verkuilen (2006), who contrast comorbidity, covariation, and co-occurrence.

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 organization, but grievances are almost always present, and the absence of grievances rarely gets the chance to act in a constraining manner on social movement 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 concerted government repression) could be seen as a nontrivial necessary condition, for social movements routinely encounter government repression. Although the specific computational formula recommended in this paper for assessing the relevance of necessary conditions differs in its details from the one suggested by Goertz (2003, 2006), the underlying goals are similar.⁷

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

$$\text{Coverage}(\mathbf{Y}_i \leq \mathbf{X}_i) = \frac{\sum (\min(\mathbf{X}_i, \mathbf{Y}_i))}{\sum (\mathbf{X}_i)}.$$

When the coverage of \mathbf{X} by \mathbf{Y} is small, then the constraining effect of \mathbf{X} on \mathbf{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 \mathbf{X} by \mathbf{Y} is substantial, then the constraining effect of \mathbf{X} as a necessary condition may be great. For example, if a substantial proportion of people associating with heroine addicts later become addicts and there is only a very small number of people who 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 Fig. 5a and 5b. Figure 5a depicts a necessary condition that exerts some constraint on the outcome (coverage is nontrivial). Figure 5b depicts an empirically trivial necessary condition (very low set-theoretic coverage). Using fuzzy sets, the situation depicted in Fig. 5b would appear as a plot in which almost all cases have very strong membership in \mathbf{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, as just described) only after establishing that the subset relation is consistent. That is, it must first be established that \mathbf{Y} is a rough subset of \mathbf{X} before assessing the size of \mathbf{Y} relative to the size of \mathbf{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, whereas the calculation of the coverage of a sufficiency relationship is identical to the calculation of the consistency of a necessity relationship.

⁷In 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 in this paper, 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 the patterns observed in the best instances.

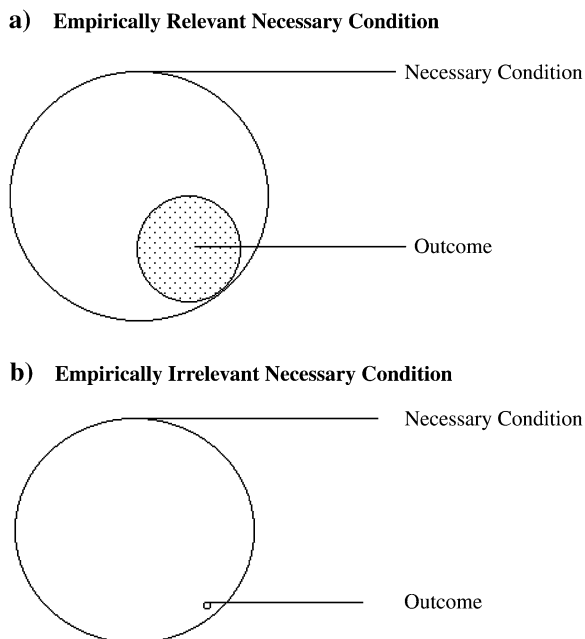


Fig. 5 Venn diagrams illustrating necessary conditions. (a) Empirically relevant necessary condition. (b) Empirically irrelevant necessary condition.

4 Partitioning Coverage

When there is more than one condition or combination of conditions sufficient for an outcome (equifinality), the assessment of the coverage of alternate 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 somewhat analogous to the partitioning of explained variation in multiple regression. The discussion of the partitioning of coverage that follows assumes that the researcher has demonstrated that the relevant conditions or combinations of conditions are highly consistent subsets of the outcome.⁸

For purposes of illustration, consider evidence from a fuzzy set analysis of individual-level data. The data set is the National Longitudinal 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 ($\sim\mathbf{P}$, where “ \mathbf{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 achievement test scores (\mathbf{T}), the fuzzy set of cases with high parental income (\mathbf{I}), and the fuzzy set of cases with college education (\mathbf{C}). (The calibration of these fuzzy sets is described in Ragin [2006].) Applying fsQCA (Ragin, Drass, and Davey 2003) to these data yields two recipes for avoiding poverty, namely, the combination of high test scores and high parental income ($\mathbf{T}\mathbf{I}$) and college education (\mathbf{C}) by itself.⁹

⁸For reasons to be explained subsequently, this section focuses only on the raw versus unique coverage of sufficient conditions and not on the relative importance of necessary conditions. The latter topic is addressed in Goertz (2003, 2006).

⁹Membership in $\mathbf{T}\mathbf{I}$ is the minimum (lower) of degree of membership in \mathbf{T} and degree of membership in \mathbf{I} .

Table 4 Calculation of coverage

<i>Causal conditions</i>	<i>Sum of consistent scores</i>	<i>Sum of outcome scores</i>	<i>Coverage</i>
T·I	192.90	947	0.2037
C	402.85	947	0.4254
T·I + C	438.56	947	0.4631

The calculation of the raw coverage of these two recipes for the outcome, avoiding poverty ($\sim\mathbf{P}$), is shown in Table 4. The first row reports the coverage calculation for the combination of high test scores and high parental income (**T·I**). The sum of the overlap between **T·I** and the outcome is 192.9. The sum of the memberships in the outcome is 947. Thus, this combination covers about 20.4% of the total membership in the outcome ($192.90/947 = 0.2037$). Using these procedures, condition **C** covers about 42.5% of the total membership of the outcome (see row 2 of Table 4). Thus, both combinations cover a substantial proportion of the outcome. However, the raw coverage of condition **C** (college education) is substantially greater.

For comparison purposes, Table 4 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** and 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 4, which reports a coverage of 46.3%, greater than the coverage of either of the two components (compare row 3 of Table 4 with rows 1 and 2). However, the coverage of the two-path model (46.3%) is only modestly superior to the raw coverage of the best single path (path **C**, with 42.5%).

Table 4 provides all the information that is needed to partition coverage, following the template provided by 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 \mathbf{X}_1 to explained variation in \mathbf{Y} , it is necessary to compute the multiple regression equation with all relevant independent variables included, and then to recompute the equation excluding \mathbf{X}_1 . The difference in explained variation between the first and second equations shows the unique contribution of \mathbf{X}_1 . These procedures ensure that the explained variation that \mathbf{X}_1 shares with correlated independent variables is not credited to \mathbf{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.

Consider the crisp-set case. Suppose a researcher finds that there are two combinations of conditions that generate outcome \mathbf{Y} , $\mathbf{A}\cdot\mathbf{B}$ and $\mathbf{C}\cdot\mathbf{D}$. The researcher calculates the coverage of these two paths and finds that the first embraces 25% of the instances of \mathbf{Y} (coverage = 0.25), whereas the second embraces 30% (coverage = 0.3). However, when

Table 5 Partitioning coverage

	<i>Total coverage</i>	<i>Without term</i>	<i>Unique</i>
Unique to T·I	0.4631	0.4254	0.0377
Unique to C	0.4631	0.2037	0.2594

calculating their coverage as alternate paths (i.e., their union: $\mathbf{A}\cdot\mathbf{B} + \mathbf{C}\cdot\mathbf{D}$, where addition indicates logical *or*), the researcher finds that together they embrace only 35% 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., $0.35 < 0.25 + 0.3$) is because the two paths partially overlap. That is, there are cases that combine all four causal conditions (i.e., instances of $\mathbf{A}\cdot\mathbf{B}\cdot\mathbf{C}\cdot\mathbf{D}$) and the coverage of these instances is counted twice when 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 $\mathbf{A}\cdot\mathbf{B}$, uniquely due to $\mathbf{C}\cdot\mathbf{D}$, and overlapping (i.e., due to the existence of cases of $\mathbf{A}\cdot\mathbf{B}\cdot\mathbf{C}\cdot\mathbf{D}$). The unique coverage of each term can be calculated by subtraction, following the template provided by regression analysis. The unique coverage of path $\mathbf{A}\cdot\mathbf{B}$ is $0.35 - 0.3 = 0.05$, the unique coverage of path $\mathbf{C}\cdot\mathbf{D}$ 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% of the instances of the outcome is $\mathbf{A}\cdot\mathbf{B}\cdot\mathbf{C}\cdot\mathbf{D}$, 5% of the instances of the outcome is $\mathbf{A}\cdot\mathbf{B}$ without $\mathbf{C}\cdot\mathbf{D}$, and 10% is $\mathbf{C}\cdot\mathbf{D}$ without $\mathbf{A}\cdot\mathbf{B}$.

The calculation of the unique coverage of a combination of conditions in fuzzy set analysis is exactly parallel, as shown in Table 5, which uses the same individual-level data used in Table 4. The coverage of the outcome (avoiding poverty) that is uniquely due to path $\mathbf{T}\cdot\mathbf{I}$ is the difference between the coverage of the two-path model (0.4631) and the coverage that is obtained once this path ($\mathbf{T}\cdot\mathbf{I}$) is removed from the two-path model, which in this example is equivalent to the coverage of the other path (\mathbf{C}) by itself. Thus, the unique coverage of path $\mathbf{T}\cdot\mathbf{I}$ is 0.0377, that is, 0.4631 (the combined coverage of the two paths) less 0.4254 (the single coverage of path \mathbf{C}). Likewise, the coverage of the outcome that is uniquely due to path \mathbf{C} is the difference between the coverage of the two-path model (0.4631) and the coverage of path $\mathbf{T}\cdot\mathbf{I}$ by itself (0.2037), or 0.2594. These calculations reveal that the unique coverage of path \mathbf{C} is much greater than the unique coverage of path $\mathbf{T}\cdot\mathbf{I}$. In fact, the coverage of $\mathbf{T}\cdot\mathbf{I}$ is almost entirely a subset of the coverage of \mathbf{C} . (In other words, most of $\mathbf{T}\cdot\mathbf{I}$ is $\mathbf{T}\cdot\mathbf{I}\cdot\mathbf{C}$.) The remaining 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 (0.4631) and the sum of the two unique portions ($0.0377 + 0.2594 = 0.2971$), which is 0.166. Figure 6 illustrates these results using a Venn diagram.

When there are many different paths to the same outcome, it is very important to calculate both the raw and unique coverage of each causal combination. These calculations often reveal that there are only a few high-coverage causal combinations, even in analyses that have many sufficient combinations. Although it is useful to know all the different causal combinations linked to an outcome, it is also important to have an assessment of their relative empirical weight. Calculations of raw and unique coverage provide these assessments directly.

Note that it is possible to partition the coverage of overlapping sufficient conditions only because the denominator is the same in all the different coverage calculations: $\sum(\mathbf{Y}_i)$.

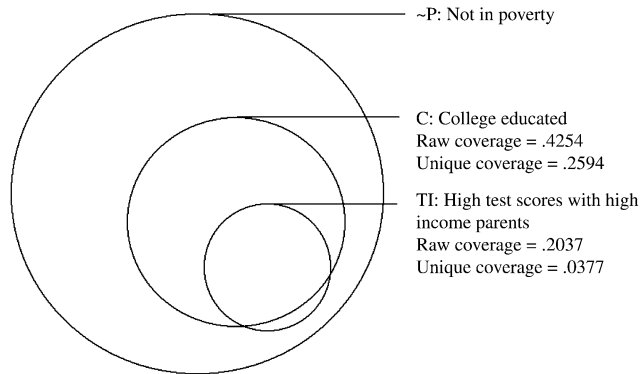


Fig. 6 Venn diagram representation of the partitioning of set-theoretic coverage using fuzzy sets.

This constant is computationally analogous to the total sums of squares in the dependent variable in conventional multiple regression analysis. It is not possible to apply parallel procedures to the assessment of the relative coverage of necessary conditions because each measure of the coverage of a necessary condition has a different denominator. For example, the coverage of necessary condition X_1 would be:

$$\sum(\min(\mathbf{X}_{1i}, \mathbf{Y}_i)) / \sum(\mathbf{X}_{1i}),$$

whereas the coverage of necessary condition X_2 would be:

$$\sum(\min(\mathbf{X}_{2i}, \mathbf{Y}_i)) / \sum(\mathbf{X}_{2i}).$$

5 Suggested Uses of Consistency and Coverage

5.1 Evaluating Monocausal Arguments

The two measures presented in this paper, set-theoretic consistency and coverage, have many different uses.¹⁰ Any social scientific argument or statement that is formulated in terms of sets (and most are formulated in these terms) can be evaluated using these measures. As noted previously, the argument that democracies do not go to war against each other can be evaluated by assessing whether or not the set of democratic dyads is a consistent subset of the set of nonwarring dyads. More than likely, this subset relation is highly consistent. However, its coverage is likely to be small because there are many ways to avoid war, especially when all possible country dyads are considered (see, e.g., Peceny and Beer 2002). Geographic distance alone suffices for most dyads, for example, regardless of whether the countries in question are democratic. For many scholars, the low coverage of the democratic peace argument is not considered a liability, for the most

¹⁰A very important use, constructing crisp truth tables from fuzzy sets, is not addressed in this paper. Interested readers should consult Ragin (2004).

common paths to peaceful coexistence (e.g., geographic distance) are trivial from the perspective of theory.

5.2 *Evaluating Combinatorial Arguments*

The relevance of consistency and coverage to the democratic peace argument is easy to grasp (and to display graphically) because the argument is monocausal and can be formulated in terms of simple, crisp sets. Most arguments in the social sciences, however, cite combinations of causes, not solitary conditions, and often these conditions are difficult to operationalize as crisp sets. Consider, for example, the conditions thought to spawn mass protest against austerity measures mandated by the International Monetary Fund (IMF) in debtor countries. Most discussions of these conditions present, in effect, a description of the ideal typical protest country. The conditions manifested by this ideal typical case include economic factors (e.g., deprivation and a declining standard of living), social structural factors (e.g., the massing of the poor in urban slums), and political factors (e.g., political corruption, authoritarian rule, and a recent history of political contention). From the viewpoint of theory, it is this potent mix of conditions, their combination, that explains austerity protest.

With fuzzy sets, the assessment of combinatorial arguments, such as the one just sketched, is straightforward. The researcher first assesses each case's degree of membership in each of the components of the ideal typical formulation (e.g., degree of membership of each country in the set of cases with economic deprivation and degree of membership in the set of countries with concentrated urban slums). Next, the researcher assesses each case's degree of membership in the combination of causal conditions specified in the ideal typical formulation. As explained in Ragin (2000), a case's degree of membership in a combination of causal conditions is determined by its lowest component membership score (i.e., its weakest link). For example, a case with only weak membership in "authoritarian rule" (one of the components of the ideal typical formulation just sketched) could have, at best, only weak membership in any combination of conditions that includes this component. Finally, the researcher assesses the consistency and coverage of the causal combination: (1) Is fuzzy membership in the combination of causal conditions (as specified in the formulation) a consistent subset of the outcome (austerity protest)? (2) If so, how much of the outcome is covered by the combination?

5.3 *Comparing Nested Combinatorial Arguments*

Measures of consistency and coverage also can be used to compare alternate ideal typical formulations. A researcher might suspect, for example, that the authoritarian rule component of the ideal typical formulation just sketched is superfluous. Membership scores can be recalculated omitting this component, and new measures of consistency and coverage can then be computed. If the researcher's hunch is correct, then the measure of consistency should remain high, indicating that membership in the combination of causal conditions is a consistent subset of the outcome, and coverage should increase. Increased coverage would indicate that the second formulation does a better job of accounting for membership in the outcome than the first. Note that membership scores in the second ideal typical formulation must be greater than or equal to membership scores in the first (i.e., the one that includes authoritarian rule). This mathematical property follows from (1) the use of the minimum to determine degree of membership in a combination of conditions and (2) the fact that the causal conditions included in the second formulation constitute a subset of those included in the first.

6 Conclusion

Lieberson (2004) comments that it is an “oxymoron” to describe set-theoretic relations in partial or probabilistic terms (e.g., “**X** is *almost always* necessary for **Y**”). For example, he might consider it oxymoronic to observe that “at public universities the most reliable way for faculty members to receive large salary increases is to secure outside offers from peer (or better) universities.” (Thus, faculty receiving such offers constitute a substantial but perhaps slightly imperfect subset of faculty receiving large salary increases.) Everyday experience indicates, of course, that such observations are not oxymoronic. The example describes precisely the situation faced by most faculty at most public universities, especially during periods when public revenues are limited (i.e., the usual situation). The measures of set-theoretic consistency and coverage presented in this paper provide useful tools for evaluating such statements. Do faculty receiving such offers constitute a consistent subset of faculty receiving large salary increases? How consistent? And how important is this path to the desired outcome, large salary increases? What is its coverage? It not only “makes sense” to assess set-theoretic statements in this manner, it is essential to do so. Statements about set relations constitute the bulk of social science theorizing. It follows that set-theoretic analysis is central to the assessment of social science theory.

Set-theoretic analysis is still in its infancy in the social sciences today. The purpose of the measures of set-theoretic consistency and coverage introduced in this essay is to enhance the utility of the approach. These simple measures provide powerful tools for improving research on set relations among social phenomena. When using these measures, it is important to keep in mind three important aspects.

The first is that they are oriented toward the evaluation of set relations reflecting explicit connections (Ragin and Rihoux 2004). Explicit connections are best understood as uniformities or near uniformities in social phenomena. Sometimes these uniformities reflect causal connections (e.g., the operation of a sufficient or necessary condition), but they also may reflect other types of integral connections (e.g., constitutive relationships). This feature explains why (a) it is always important to establish first that a set relation is consistent before evaluating its coverage or relevance (i.e., before gauging its empirical importance) and (b) it is often the case that high consistency yields low coverage. This emphasis on explicit connections in set-theoretic analysis contrasts fundamentally with the emphasis of conventional quantitative methods, where correlational connections are the central focus.

The second important aspect of these measures is that they are descriptive, not inferential. The set-theoretic techniques described here and in Ragin (1987, 2000) were developed as methods of exploring cross-case evidence in a *configurational* manner. Viewing cases as configurations in cross-case analysis (a) maintains a strong link to the study of specific empirical cases and (b) counteracts the veiling of cases that occurs when cross-case evidence is subjected to conventional forms of correlational analysis. Such analyses obscure cases in their emphasis on the net effects of “independent” variables and the competition to explain variation in the dependent variable (Ragin 2006).

The third aspect is that these measures are not ends in themselves. In conventional cross-case analysis, researchers typically focus on measures of explained variation to compare “models.” Too often, maximizing explained variation becomes the central focus of the investigation (Lieberson 1985). By contrast, the goal of configurational analysis using set-theoretic methods is to help researchers make sense of their cases (Ragin 1987; Rihoux 2003; Ragin and Rihoux 2004). Calculations of consistency and coverage do and should provide guidance, but the ultimate “test” of the results of a configurational analysis

is not their consistency or coverage but how well they help researchers make sense of their cases. Do the results resonate with what is known about processes and mechanisms operating at the case level? Do they highlight different aspects of cases or suggest new typologies? Do they group cases in a theoretically interesting or progressive manner? A key goal of social research is to make sense of the diversity of empirical cases in ways that resonate with the researcher's theoretical ideas about social phenomena. Configurational methods are especially well suited for this task.

References

- Braumoeller, Bear, and Gary Goertz. 2000. The methodology of necessary conditions. *American Journal of Political Science* 44(4):844–58.
- De Meur, Gisèle, and Benoît Rihoux. 2002. *L'Analyse Quali-Quantitative Comparée: Approche, Techniques et applications en sciences humaines*. Louvain-la-Neuve: Bruylant-Academia.
- Dion, Douglas. 1998. Evidence and inference in the comparative case study. *Comparative Politics* 30(2):127–45.
- George, Alexander. 1979. Case studies and theory development: The method of structured, focussed comparison. In *Diplomacy: New approaches in history, theory and policy*, ed. Paul G. Lauren, 43–68. New York: Free Press.
- George, Alexander, and Andrew Bennett. 2005. *Case studies and theory development*. Cambridge: MIT Press.
- Goertz, Gary. 2002. The substantive importance of necessary conditions hypotheses. In *Necessary conditions: Theory, methodology, and applications*, ed. Gary Goertz and Harvey Starr, 65–94. New York: Rowman & Littlefield.
- . 2003. Assessing the importance of necessary or sufficient conditions in fuzzy-set social science. Compass Working Paper WP2003-7. www.compass.org/wp.htm.
- . 2005. *Social Science Concepts: A User's Guide*. Princeton: Princeton University Press.
- . 2006. Assessing the trivialness, relevance, and relative importance of necessary or sufficient conditions in social science. *Studies in Comparative International Development*. Forthcoming.
- Goertz, Gary, and Harvey Starr, eds. 2002. *Necessary conditions: Theory, methodology, and applications*. New York: Rowman & Littlefield.
- Herrnstein, Richard, and Charles Murray. 1994. *The bell curve: Intelligence and class structure in American life*. New York: Free Press.
- Kosko, Bart. 1993. *Fuzzy thinking: The new science of fuzzy logic*. New York: Hyperion.
- Lieberson, Stanley. 1985. *Making it count: The improvement of social research and theory*. Berkeley: University of California Press.
- . 2004. Comments on the use and utility of QCA. *Qualitative Methods. Newsletter of the American Political Science Association Organized Section on Qualitative Methods* 2(2):13–4.
- Mackie, John L. 1965. Causes and conditionals. *American Philosophical Quarterly* 2:245–65.
- Peceny, Mark, and Caroline C. Beer. 2002. Dictatorial peace? *American Political Science Review* 96:15–26.
- Ragin, Charles C. 1987. *The comparative method: Moving beyond qualitative and quantitative strategies*. Berkeley: University of California Press.
- . 2000. *Fuzzy-set social science*. Chicago: University of Chicago Press.
- . 2003. Recent advanced in fuzzy-set methods and their application to policy questions. Compass Working Paper WP2003-9. www.compass.org/wp.htm.
- . 2004. From fuzzy sets to crisp truth tables. Compass Working Paper WP2004-28. www.compass.org/wp.htm.
- . 2006. The Limitations of net effects thinking. In *Innovative comparative methods for policy analysis*, ed. Heike Grimm and Benoît Rihoux, 13–41. Springer.
- Ragin, Charles C., Kriss A. Drass, and Sean Davey. 2003. *Fuzzy-set/qualitative comparative analysis*, version 2.0. www.u.arizona.edu/~cragin/fsQCA/download/Setup_fsQCA.exe.
- Ragin, Charles C., and H. Giesel. 2003. *User's guide to fuzzy-set/qualitative comparative analysis*. www.u.arizona.edu/~cragin/fsQCA/download/fsQCAManual.pdf.
- Ragin, Charles C., and Benoît Rihoux. 2004. Qualitative comparative analysis (QCA): State of the art and prospects. *Qualitative Methods. Newsletter of the American Political Science Association Organized Section on Qualitative Methods* 2(2):3–13.
- Rihoux, Benoît. 2003. Bridging the gap between the qualitative and quantitative worlds? A retrospective and prospective view on qualitative comparative analysis. *Field Methods* 15(4):351–65.
- Smithson, Michael, and Jay Verkuilen. 2006. *Fuzzy set theory*. Thousand Oaks, CA: Sage.