STRATEGIES FOR SOCIAL INQUIRY

Set-Theoretic Methods for the Social Sciences

A Guide to Qualitative Comparative Analysis

CARSTEN Q. SCHNEIDER
CLAUDIUS WAGEMANN

CAMBRIDGE

5 Parameters of fit

Easy reading guide

In Chapter 4, we introduced the use of truth tables in the analysis of sufficient conditions. A central point of our previous chapter was to assess for every single truth table row whether it represents a sufficient condition for the outcome. If yes, then such a row has been included in the logical minimization. If not, then it has not been included.

So far, we have assumed an ideal world that presents itself in clear and neat patterns. In reality, social science research based on observational data is characterized by noisy data. The coming chapters deal with issues that derive from this fact and describe strategies for how set-theoretic methods need to react to this. One fruitful way of looking at the discrepancy between neat set theory and the underlying empirical evidence is to frame this in terms of *incomplete truth tables*. A truth table is incomplete if it shows one or both of the following features. First, it might consist of rows that contain cases whose membership scores in that row and the outcome contradict the statement of sufficiency. These are *contradictory* or *inconsistent rows*. Second, a truth table might contain rows for which no (or, at least, not enough) empirical evidence is available. These rows are called *logical remainders*, and the presence of such rows is referred to as the phenomenon of *limited diversity*. The analytic problem caused by both forms of incomplete truth tables is that it becomes impossible to decide whether certain truth table rows represent sufficient conditions for a given outcome. This means that for some truth table rows it is not a straightforward business to establish whether they are sufficient for the outcome. Put differently, it is difficult to decide whether to include a given row in the Boolean minimization process. This represents an analytic problem, since the solution formula greatly depends on the decision of which rows are included in the minimization.

This chapter discusses the phenomenon of less-than-perfect subset relations, while Chapter 6 will deal with limited diversity. We start by introducing the notion of logically contradictory truth table rows and outline strategies of dealing with them (5.1). We introduce the consistency measure as one important strategy, and one which exists both for sufficient and for necessary conditions. After this, we also introduce the parameter of coverage, which expresses the empirical importance (sufficiency) and relevance (necessity) of a condition. We first introduce consistency and coverage formulas for sufficient conditions (5.2 and 5.3, respectively) and then for necessary conditions (5.4 and 5.5, respectively). As in previous chapters, we introduce each argument by starting from crisp sets and by then extending it to fuzzy sets.

The notion of consistency is indispensable for understanding the logic of the Truth Table Algorithm (Chapter 7), which is at the core of QCA. A solid knowledge of the meaning and measure of the parameters of fit is therefore indispensable. More advanced readers might want to consult the At-a-glance boxes in order to assess whether they are familiar enough with these issues and, if so, simply skim through this chapter.

5.1 Defining and dealing with contradictory truth table rows

The notion of a contradictory truth table row is easier to understand with crisp sets. It describes a situation in which those cases that are members in a truth table row do not share the same membership in the outcome. Put differently, the same row leads to both the occurrence and the non-occurrence of the outcome. Since truth table rows are, in essence, statements of sufficiency, such an empirical situation suggests a logical contradiction, for it would mean that the very same combination of conditions (aka truth table row) produces both Y and ~Y. The analytic problem is that, based on the empirical evidence, it is not straightforward to decide whether this row is sufficient for Y, ~Y, or neither and, consequently, whether it should be included in the logical minimization for outcome Y, outcome ~Y, or neither. It cannot, however, be included in both minimization procedures.

There are several, mutually non-exclusive strategies for dissolving logically contradictory truth table rows in either csQCA or fsQCA *prior* to the logical minimization, and there is another set of strategies for handling such contradictory rows *during* the minimization procedure (Ragin 1987: 113–18; Rihoux and De Meur 2009). Let us first turn to the strategies for dissolving the contradiction.

The first strategy consists in *adding a condition* to the truth table. If those cases in the contradictory row that display qualitatively different outcome membership scores also show qualitatively different membership scores in the new condition, then the contradiction is resolved. This is because by adding a new condition, the contradictory row is split in two rows, separating the cases with different outcome membership scores in these two new rows. Of course, the downside of this strategy is that not only the contradictory row, but also all the other rows are split in two, thus doubling the number of truth table rows. Remember, the number of truth table rows is a direct function of the number of conditions (k), as expressed in the formula 2^k (Chapter 4). This, in turn, increases the problem of limited diversity (Chapter 6).

A second strategy is to *respecify the definition of the population of interest*. By virtue of this, some cases might be excluded and/or new ones included. Such a change of the set of cases via a redefinition of the scope conditions (Walker and Cohen 1985) must be based in theoretical arguments. Cases cannot be excluded in an ad hoc manner simply because they contradict a statement of sufficiency. Instead, theoretical and substantive arguments must be explicitly brought forward as to why such cases are of a qualitatively different kind and therefore fall outside the scope of the analysis (Ragin and Becker 1992). The difficulty of this strategy might consist in the lack of plausible theoretical arguments. Even if these do exist, such a redefinition of the scope conditions might have to be accompanied by a change in the relevant theories. This, in turn, would have to have an influence on the choice of conditions and the outcome and their respective calibration functions, which could create new contradictory rows.

Third, one can *respecify the definition, conceptualization, and/or measurement of the outcome or condition(s)*. A closer look at the similarities and differences in the contradictory cases in a given row might reveal that the specification of the outcome or a condition was too vague, imprecise, or just plain wrong. If so, a respecification may contribute to solving inconsistencies. Just as in the case of redefining the scope conditions, a change of the meaning and thus calibration of concepts must also be based on theoretical arguments, without which such a recalibration strategy would degenerate into a blunt data-fitting exercise.

Any of these approaches can help solving contradictions. These strategies belong to the standards of good QCA practice and they represent part of what is meant by the phrase "going back and forth between ideas and evidence" (Ragin 2000), i.e., the process of updating theoretical, conceptual, and research design decisions based on preliminary empirical insights. At the same time, all strategies come at a cost and none can promise to always solve every logical contradiction. Thus, in applied QCA, it usually happens that researchers enter the process of logical minimization with truth tables that contain some logically contradictory truth table rows. There are several, mutually exclusive treatments of logically contradictory rows during the process of logical minimization.

First, one can *exclude all contradictory rows* from the logical minimization process. By doing this, one allows only perfect subset relations to qualify as sufficient conditions. As a consequence, any case that is a member of the outcome but which falls into a contradictory truth table row will not be explained, or covered, by the solution term obtained with this strategy. Second, one can *include all contradictory rows* in the logical minimization process. This strategy is based on the argument that a contradictory row at least makes the occurrence of the outcome possible. The solution formula obtained thus represents the conjunctions of conditions that make the outcome possible. All cases that are members of the outcome will be explained, or covered, by that solution term. The downside, however, is that the solution term will also cover some cases that are not members of the outcome. Third, one can *make all inconsistent rows available for computer-generated assumptions* about their outcome value. It is then up to the computer to decide which of the contradictory rows to include in the process of logical minimization and which ones not to include. The only rationale for selecting some contradictory rows is whether their inclusion makes the resulting solution term more parsimonious. Although all three strategies for handling contradictory rows come at a price, the third strategy is usually least justifiable and is hardly ever encountered in applied QCA.

In the remainder of this chapter and, in fact, throughout the book, we advocate yet another way of dealing with contradictory rows and inconsistent truth table rows and set relations. This strategy takes into account how much, or to what degree, a given row deviates from a perfect set relation. Consider, for example, the following two scenarios. In one truth table row, nine out of ten cases share the same qualitative membership score in the outcome. Hence, one case deviates from the general pattern, or 90 percent of the evidence is in line with a subset relation. In another truth table row, six out of ten cases agree on their membership score in the outcome. Hence, only 60 percent of the empirical evidence is in line with the subset relation of sufficiency. This type of percentage can be seen as an important measure of how *consistent* a particular configuration is with the assertion that it is a sufficient condition for the outcome. We introduce this parameter as the *consistency value* in the remainder of the book.

In sum, strategies that aim at dissolving contradictions directly stem from the anchoring of set-theoretic approaches in qualitative methods. They remind us of the important fact that set-theoretic methods, in general, and QCA, in particular, are not only data analysis techniques but also research approaches with specific requirements for the research process before and after the actual data analysis. They reflect, in other words, the double nature of QCA as both a research approach and a data analysis technique (see the Introduction, section on QCA as a set-theoretic approach) (Berg-Schlosser, De Meur, Rihoux, and Ragin 2008; Wagemann and Schneider 2010). Only if inconsistent rows still exist after these time- and energy-consuming countermeasures should one resort to those strategies that handle such rows during the process of logical minimization. Here, of greatest importance for applied QCA is the use of the consistency measure as a yardstick for guiding the decision on whether or not to include a truth table row into the logical minimization procedure.

At-a-glance: defining and dealing with contradictory truth table rows

In dealing with **contradictory truth table rows**, a decision must be made before any **logical minimization** of the **truth table** is undertaken on how to approach these contradictory rows. Some of the strategies for resolving contradictions include approaches that better specify the conditions in the explanatory model or the case selection with regard to the reference population.

Consistency measures will be of additional help in making decisions about contradictory rows. Consistency scores should not replace, but rather complement, the qualitative strategies for dissolving contradictions.

5.2 Consistency of sufficient conditions

Starting off with csQCA, perhaps the most intuitive way of graphically displaying the notion of consistency of a sufficient condition is by means of a Venn diagram (Ragin 2006). Figure 5.1 displays Venn diagrams for three different conditions (X_1, X_2, X_3) and an outcome Y. In all three scenarios, the size of sets X and Y remains identical; only their relative location changes. Condition X_1 (Venn diagram on the left) is a perfect subset of outcome set Y, whereas both conditions X_2 and X_3 are not. Conditions X_2 and X_3 differ in the degree to which they violate the subset relation with Y. The share of set X_3 that is outside Y (area d) in relation to the overall size of X_3 (areas b and d) is larger than for condition X_2. Therefore, X_2 is more consistent than X_3 as a sufficient condition for Y.

While Venn diagrams are good for grasping the basic notion of set-theoretic consistency, two-by-two tables are more powerful when trying to explain how to calculate this parameter of fit. Table 5.1 displays the same three conditions and outcome as Figure 5.1, and the cells (a–d) correspond to the areas (a–d) in the Venn diagrams. The numbers in the cells indicate the number of cases that show the respective membership scores in the condition and outcome.

Table 5.1 Two-by-two tables – consistent and inconsistent sufficient conditions

Outcome Y	X_1 = 0	X_1 = 1	X_2 = 0	X_2 = 1	X_3 = 0	X_3 = 1
1	80 a	100 b	80 a	90 b	80 a	8 b
0	15 c	0 d	15 c	10 d	15 c	92 d

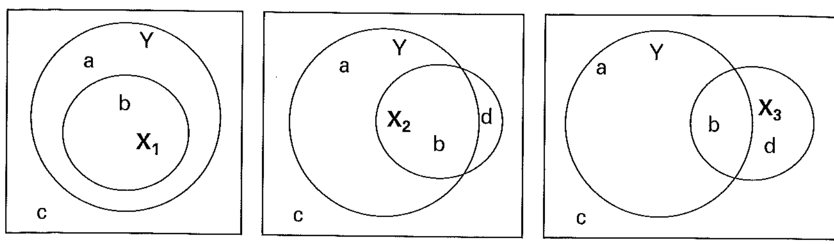

Figure 5.1 Venn diagrams – consistent and inconsistent sufficient conditions

As we can see, the difference between condition X_1, on the one hand, and X_2 and X_3, on the other, is that with the fully consistent condition X_1 all members of X_1 are located in cell b and none in cell d. This is why there is no area d in the first Venn diagram in Figure 5.1. Some cases in X_2 and many cases in X_3 fall into cell d rather than cell b. Recall from Chapter 3 that for statements of sufficiency, only those cases matter that are members of the alleged sufficient conditions (X = 1). Perfectly consistent sufficiency requires that all cases with X = 1 are also members of outcome (Y = 1). Therefore, no case should be in cell d. The more cases fall into cell d, the more consistency decreases.

Ragin (2006) suggests that the consistency of a sufficient condition X for outcome Y be mathematically expressed by dividing the number of cases in cell b by all the cases that matter to measure sufficiency, i.e., the number of cases in cells b and d. In csQCA, the consistency of X as a sufficient condition for Y can therefore be calculated as follows:

$$\begin{array}{l}\textit{Consistency of X as a}\\ \textit{sufficient condition for Y}\end{array} = \frac{\textit{Number of cases where } X=1 \textit{ and } Y=1}{\textit{Number of cases where } X=1}$$

The same can be expressed by making reference to the cells in the two-by-two table:

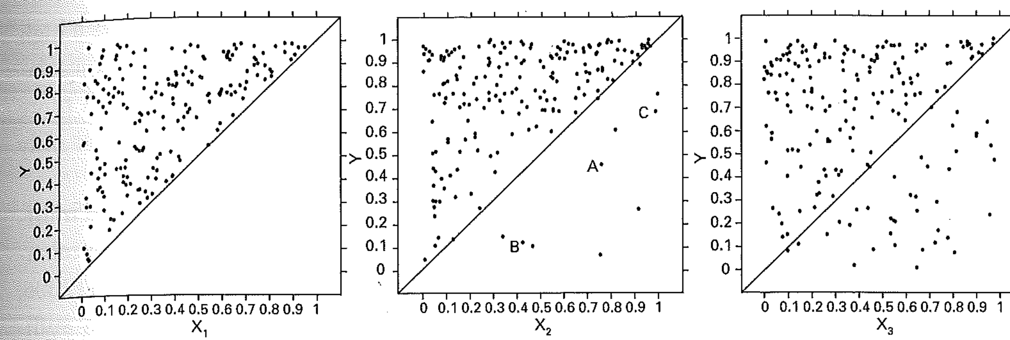

Figure 5.2 XY plot – consistent and inconsistent sufficient conditions

$$\begin{array}{l}\textit{Consistency of X as a}\\ \textit{sufficient condition for Y}\end{array} = \frac{\textit{Number of cases cell b}}{\textit{Number of cases cells } b+d}$$

The consistency value is 1 if a condition is fully consistent and decreases as inconsistency becomes stronger. Applied to our example in Table 5.1, the consistency values are as follows:

$X_1 = 100 / 100 = 1$
$X_2 = 90 / 100 = 0.9$
$X_3 = 8 / 100 = 0.08.$

When shifting to fuzzy sets, the notion of subset relations is graphically best represented in the form of XY plots (section 3.1.2.1). Figure 5.2 shows three such XY plots, which, along the lines of Figure 5.1 and Table 5.1, display three different sufficient conditions with increasing inconsistency from left (X_1) to right (X_3).

When trying to calculate consistency, one approach could be to proceed analogously to the calculation in crisp sets by simply counting the number of cases that are in line with the statement of sufficiency (i.e., those above or on the main diagonal) and then dividing this number by the number of cases that are relevant for the test (i.e., those with membership in X of higher than 0). The plot for X_1 shows no case below the main diagonal. Hence, consistency would be 1. For X_2 there are 10 out of 195 cases with X > 0 below the main diagonal. Hence, consistency for X_2 would be 185/195 = 0.95. For X_3, consistency would be 145/195 = 0.74.

This crisp approach to calculating consistency is deficient, though. Notice that it gives equal weight to all cases below the diagonal. This is not plausible. The distance between cases and the diagonal is clearly of interest because cases that are far below the main diagonal obviously deviate more strongly from the alleged subset relation. For instance, case A in Figure 5.2 has a high degree of membership in the supposedly sufficient condition X, but a relatively low value of Y. It therefore contradicts the sufficiency statement more than cases that fall only slightly below the main diagonal and/or have only weak membership in condition X and outcome Y.

The remedy for these pitfalls is to make use of the more fine-grained information conveyed by each case's fuzzy-set membership in X and Y when calculating consistency (and coverage, see 5.3). This is precisely what Ragin (2006; 2008a: 44–68) suggests with his formula for the consistency of a fuzzy sufficient condition. For each case, the minimum values across the membership scores in X and Y are added up and then divided by the sum of the membership values in X across all cases.

$$Consistency_{Sufficient\ Conditions\ (X_i \leq Y_i)} = \frac{\sum_{i=1}^{I} \min(X_i, Y_i)}{\sum_{i=1}^{I} X_i}.$$

If all cases have smaller or equal membership in X than in Y (as is required for fully consistent sufficiency), then the numerator simply becomes the sum of all X_i, and this formula returns a value of 1. For cases below the main diagonal, their membership in Y provides the minimum. The farther they fall below the diagonal, the bigger the difference between their membership in X and Y and the smaller the sum in the numerator becomes in relation to the sum of X_i in the denominator. Thus, this consistency measure takes into account how far a case falls below the main diagonal, or how far the membership in X exceeds that in Y.

Notice, as well, that this formula is a generalization of the crisp-set consistency formulas above and yields identical consistency values to them (Ragin 2008b: 108n. 5). In case of crisp sets, the X_i in the denominator can only be 0 or 1. Therefore, the sum of all X_i is equal to the number of cases where $X = 1$. In crisp sets, the numerator denotes the number of cases in cell b of our two-by-two table (Table 5.1). This is the only cell where the minimum across X_i and Y_i is 1, since both the X and Y values are 1. In all other cells, the minimum of X_i and Y_i is 0, and the cases contained therein are not added to the numerator.

While being a plausible way of numerically expressing the degree of a subset relation, the consistency formula has one particular shortcoming when applied to fuzzy sets. It does not take into account whether an inconsistent case is above or below the qualitative anchor of 0.5 in X and/or Y. Take, for instance, cases A, B, and C in the XY plot for X_2. Their distances to the main

diagonal are identical. Thus, they equally contribute to the inconsistency of X_2 as a sufficient condition for Y. There is, however, a qualitative difference between cases B and C, on the one hand, and case A, on the other, which is important for evaluating whether X_2 can be interpreted as a sufficient condition for Y. The former two cases display set membership scores in X and Y that are on the same side of the 0.5 qualitative anchor – they are either more in than out of both X and Y (case C) or more out than in of both X and Y (case B). Case A, in contrast, has qualitatively different membership scores in X and Y. Its membership in X_2 is above 0.5, making it a good empirical instance of this condition. Yet, its membership in Y is below 0.5. Hence, case A is a true logically contradictory case while cases B and C are simply inconsistent cases. To summarize this shortcoming, contradictory truth table rows can and do occur both in csQCA and fsQCA and they are, by definition, inconsistent rows. With fuzzy sets, however, not all inconsistent rows are automatically truly logically contradictory.[1] Analytically, inconsistent subset relations that also contain a true logical contradiction are less in line with a statement of sufficiency than simply inconsistent subset relations. They warrant more actions by researchers in terms of the strategies for fixing contradictory rows outlined in section 5.1 before proceeding with the logical minimization.

Which consistency level should researchers impose when identifying single truth table rows as sufficient conditions? For obvious reasons, consistency values close to, or even below, 0.5 should be ruled out, as this indicates that (almost) half of the empirical evidence contradicts the subset relational statement of sufficiency. Even values below 0.75 are often problematic as they have consequences for the subsequent analysis, which we spell out in various places in the remainder of this book (e.g., sections 5.6 and 9.1). As mentioned, with fuzzy sets, not only the consistency score, but also the presence or absence of true logically contradictory cases should be taken into account.[2] In the presence of such cases, researchers should be more reluctant to declare that row as a sufficient condition, independently of its consistency value.

Beyond these rough indications, we would like to make a strong plea for the notion that the exact location of the consistency threshold is heavily dependent on the specific research context. In other words, researchers should not justify their choice of the consistency threshold by making reference to some

[1] The notion of a true logical contradiction also extends to statements of necessity (see section 5.4). Here, a true logically contradictory case is one with $X < Y$ and $X < 0.5$ and $Y > 0.5$.

[2] Perhaps the easiest way to do so is by producing an XY plot and checking whether the lower right area contains cases.

sort of universally accepted consistency threshold, akin to the (largely non-reflected) use of the 95 percent confidence interval in inferential statistics. Instead, researchers should guide their decision by making reference to various research-specific features.

The following guidelines should be used as some rough yardsticks. The more precise and strong the theoretical expectations that can be derived from the literature, the higher the consistency that should be used. The higher the confidence in the precision and validity of the calibration procedure for the conditions and the outcome, the higher the consistency. The lower the number of cases under investigation, the higher the consistency. The more logically contradictory cases, the higher the consistency.[3] In addition, in applied QCA a gap often exists between rows with relatively high and low consistency values that can guide the decision of where to put the consistency threshold. A less often used strategy is to employ the tools of probability theory. Ragin (2000: 109–16) suggests a binomial probability test simple for smaller N (30 or below) and a z test when the N is larger than that. Other authors also combine the assessment of set relations with tools from probability theory.[4] By now, several of the software packages (R and Stata) available for set-theoretic analyses allow for easy use of statistical tests not only of consistency, but also of coverage (see 5.3). Clearly, no precise, universal consistency value can be derived from these guidelines, often not even within a specific project. It is therefore strongly recommended that separate analyses with different thresholds of consistency be run in order to find out how sensitive the results are to the choice of the consistency level. We discuss this issue in further detail under the heading of robustness in section 11.2.

In sum, the consistency formula indicates the degree to which the statement of sufficiency is in line with the empirical evidence at hand. The more cases that deviate from the subset pattern and the stronger their deviation, the lower the consistency value. Of course, consistency can be calculated for any statement of sufficiency of arbitrary complexity. Put differently, X in the consistency formula is simply a placeholder for a set that might consist of the logical AND and OR combination of several sets. Regardless of how many sets are combined with different logical operators, each case has only *one* set membership score in that complex set. This implies that consistency values can be calculated for single truth table rows; single paths which have been identified as sufficient; or even an entire solution formula. When the consistency statement is on truth table rows, it is called "raw consistency," while the consistency value for the entire solution is called "solution consistency."

[3] As explained, with fuzzy sets cases can be inconsistent with a postulated subset relation without, however, being logical contradictory cases. A further guideline for choosing the consistency threshold for X as a sufficient condition for Y applies in fuzzy sets: with fuzzy sets, X can be a subset of both Y and ~Y. Since declaring X as sufficient for both Y and ~Y amounts to a logical contradiction, only those rows should be declared as sufficient for Y that display high consistency as sufficient conditions for Y and low values for ~Y. We return in greater detail to the more general issue of simultaneous subset relations in section 9.2.

[4] See, for instance, Braumoeller and Goertz (2003); Dion (2003); Caramani (2009); Eliason and Stryker (2009).

At-a-glance: consistency of sufficient conditions

Consistency provides a numerical expression for the degree to which the empirical information deviates from a perfect subset relation. This information plays a crucial role when deciding which **truth table rows** can be interpreted as **sufficient conditions** and can thus be included in the **logical minimization** process.

With **crisp sets**, inconsistency by default stems from logically contradictory cases. With **fuzzy sets**, it does not have to. Therefore, researchers are advised to check for the presence of true logically contradictory cases, in addition to the consistency value, before attributing the status of a sufficient condition to a truth table row.

Consistency can be calculated for single conditions as well as for more complex statements.

Researchers should justify their consistency threshold by making reference to research-specific features, such as the strengths of theoretical expectation and the quality of the data. The consistency value for sufficient conditions should preferably be higher than 0.75.

5.3 Coverage of sufficient conditions

Once a subset relation has been established via the use of the consistency parameter, another question can be asked: what is the relation in size between the subset (X) and the superset (Y)? The answer to this question expresses how much of outcome Y is covered by condition X, thus expressing the empirical importance of X for explaining Y.

Consider the three situations depicted in Figure 5.3. It displays three different conditions (X_1 to X_3) for the same outcome Y. All three conditions are identical with regard to their slight inconsistency as sufficient conditions (the ratio of area d over areas b and d is equal in all three Venn diagrams). What differs between the three is the size of the set of X in relation to the set of Y. Condition X_1 is larger than condition X_2 is larger than X_3. Since the set of Y is constant and the ratio between areas b and d remains the same, the varying

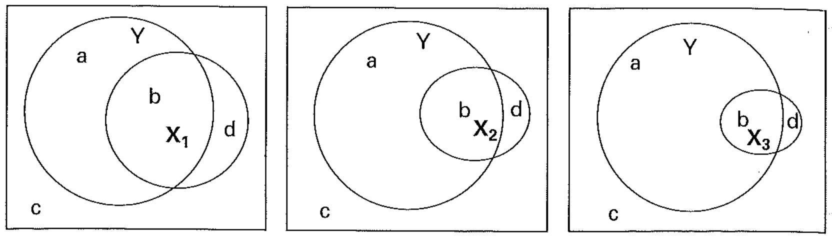

Figure 5.3 Venn diagrams – different levels of coverage sufficiency

size of the set of X means a variation in the amount of cases with Y = 1 that are *covered* by X_1, X_2, and X_3, respectively. In other words: X_1, X_2, and X_3 have different *coverages*. The coverage measure expresses the degree to which the consistent part of sufficient condition X overlaps with outcome Y.

For illustration, let Y be the set of pupils with high test scores, X_1 be the set of students who study hard, X_2 be the set of students who study hard and are talented, and X_3 be the set of students who study hard, are talented, and cheat. Of course, membership in X_3 is more difficult to obtain because we require the joint presence of various characteristics of pupils. This is why less X_3 is smaller than X_2 and X_1. This also implies that fewer members of outcome Y share the features denoted by set X_3.

Table 5.2 represents the same empirical information as the Venn diagrams in Figure 5.3. In each of the three scenarios, the number of cases that are members of the outcome (Y = 1) remains the same (210) and the consistency scores are identical.[5]

What makes the three scenarios different is that as we move from condition X_1 to X_2, and then further to X_3, the number of cases that have membership in the respective condition X decreases. At the same time, the number of cases with membership in Y is constant. In terms of cells in the two-by-two tables, this means that, as more cases move from cell b into cell a, the ratio of the consistent part of X over the total number of cases with Y decreases. The consistent part of X accounts for an increasingly smaller portion of Y. The formula for calculating the coverage of X for Y can then be written as follows (Ragin 2006, 2008a: 44–68).

$$\text{Coverage of } X \text{ as a sufficient condition for } Y = \frac{\text{Number of cases where } X=1 \text{ and } Y=1}{\text{Number of cases where } Y=1}.$$

[5] $X_1 = 200 / 208 = 0.96$; $X_2 = 120 / 125 = 0.96$; $X_3 = 24 / 25 = 0.96$.

Table 5.2 Two-by-two tables – different levels of coverage sufficiency

Outcome Y		X_1 = 0	X_1 = 1	X_2 = 0	X_2 = 1	X_3 = 0	X_3 = 1
	1	10 a	200 b	90 a	120 b	186 a	24 b
	0	20 c	8 d	20 c	5 d	20 c	1 d

Or, in terms of our cells in the two-by-two table:

$$\text{Coverage of } X \text{ as a sufficient condition for } Y = \frac{\text{Number of cases cell b}}{\text{Number of cases cells } a+b}.$$

Based on these formulas, the coverage values for the three sufficient conditions are as follows:

$X_1 = 200 / 210 = 0.95$
$X_2 = 120 / 210 = 0.57$
$X_3 = 24 / 210 = 0.11$.

The graphical intuition gained by looking at the Venn diagrams in Figure 5.3 is corroborated by these coverage values: the coverage of X_1 is higher than that of X_2 is higher than that of X_3. With crisp sets, full coverage is achieved when cell a is empty of cases.

With fuzzy sets, we have to use XY plots. Figure 5.4 displays three conditions, all with identical consistency values (0.91), but different coverage. As we go from X_1 to X_2 and further to X_3, we see that cases tend to fall closer and closer to the Y-axis, i.e., where X is close to 0 – just as with the two-by-two tables above.

Just as with consistency, so also the coverage formula for sufficient conditions suggested by Ragin (2006, 2008a: ch. 3) makes use of the more fine-grained information contained in fuzzy sets looks as follows:

$$\text{Coverage}_{\text{Sufficient Conditions } (X_i \le Y_i)} = \frac{\sum_{i=1}^{I} \min(X_i, Y_i)}{\sum_{i=1}^{I} Y_i}.$$

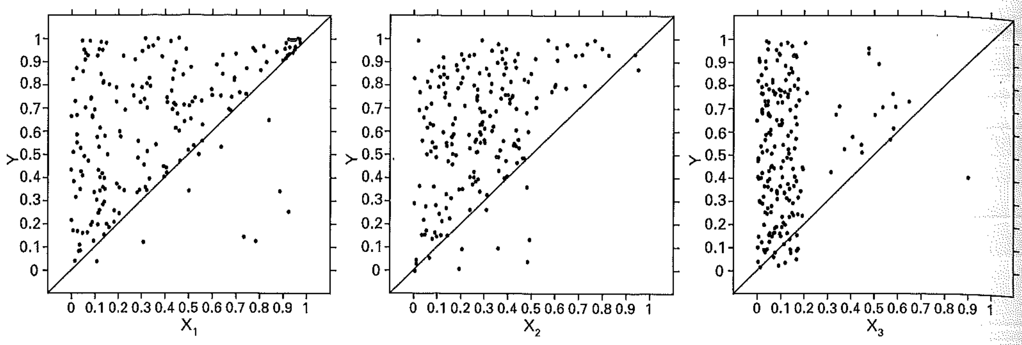

Figure 5.4 XY plot – different levels of coverage sufficiency

Applied to our three XY plots in Figure 5.3, the following coverage values are obtained:

$X_1 = 0.81; X_2 = 0.6; X_3 = 0.19,$

confirming our visual impression of X_1 being empirically more important than X_2, which, in turn, is empirically more important than X_3.

The more cases that are located in the upper left corner, and the farther away from the main diagonal these cases are, the lower the coverage. Those cases are good empirical instances of the outcome (high membership in Y) for which we lack, however, an adequate explanation because they are weak empirical instances of the sufficient condition (low membership in X).[6] The coverage formula takes into account how far above the main diagonal cases are located and, hence, how much of their fuzzy-set membership in Y is not covered by their membership in X. Cases in the upper left corner contribute little to the sum in the numerator (only their small X value) and much to the denominator (their high Y value).

As can be seen, with fuzzy sets, the calculation of coverage also takes into account that part of each *inconsistent* case's membership in Y that is covered by X.[7] As a consequence, coverage also increases due to cases that are inconsistent with the statement of sufficiency. This is an unfortunate property of the coverage parameter. It can be argued, though, that its effect is bound to be marginal and usually does not trigger substantive changes in the interpretation of

[6] The upper left corner corresponds to cell a of our Venn diagram (Figure 5.3) and our two-by-two table (Table 5.2).

[7] The problem does not affect csQCA because each inconsistent case has a membership value of 0 in Y. Thus, for each inconsistent case 0 is added to both the numerator and the denominator in the coverage formula.

the results.[8] Several features reduce the effect. First, it is not the entire inconsistent case that is counted into the coverage formula but only that part of its membership in Y that is actually covered by X. Second, coverage is only calculated for conditions that have passed a threshold of consistency. This ensures that the number of cases (far) below the main diagonal is small and therefore their distorting effect on the coverage formula low. This, incidentally, provides another argument against choosing too low levels of consistency, for it would unduly boost the coverage values of (too inconsistent) sufficient conditions. From all this follows a clear rule of thumb: the consistency of a sufficient condition must always be calculated *before* its degree of coverage, and coverage should only be calculated for conditions that passed the test of consistency (Ragin 2006).[9] It makes no sense to calculate and interpret the coverage of a condition that is not sufficient.

Recall that equifinality is an important part of the epistemological foundation of set-theoretic methods, in general, and QCA, in particular (section 3.3). Different conditions (or combinations thereof) can lead to the same outcome. As a consequence of this, we can and should calculate the coverage of these different parts separately (Ragin 2006, 2008a: 54–68). It should be established how much of the outcome is covered by each of these paths. This is called *raw coverage*. We also might want to know how much of the outcome is covered only by a specific path – the *unique coverage*. The distinction between raw and unique coverage is important because different sufficient paths can overlap, i.e., the same case can follow multiple paths toward the outcome.[10] In these cases, the outcome occurs for more than one reason. Note that if no logically redundant path (section 4.3.2) is included in the solution, then all paths have a unique coverage higher than 0. We are also interested in finding out how much of the outcome is covered by the entire solution term – the so-called *solution coverage*. For instance, consider the equifinal and conjunctural solution term $\sim A \sim C + \sim BC + F \sim D \rightarrow Y$. We can calculate the raw coverage and the unique coverage of sufficient paths $\sim A \sim C$, $\sim BC$, and $F \sim D$, respectively. In addition, we can calculate the solution coverage of the term $\sim A \sim C + \sim BC + F \sim D$.

[8] In section 9.2.1, we demonstrate under which circumstances this feature of the coverage formula produces misleading results and suggest alternative coverage formulas.

[9] This is analogous to that in multivariate regression, where beta-coefficients should be interpreted only for significant variables.

[10] Recall from section 2.2 that the logical OR operator used in QCA solution formulas is a non-exclusionary logical OR. One and the same case is allowed to be a member of more than just one sufficient condition or path.

For all three types of coverage, the coverage formula for sufficiency reported above directly applies. All that needs to be changed is what the placeholder X in this formula stands for: each case's membership in the path of interest, e.g., in the term ~A~C (raw coverage) or in the entire solution term (solution coverage). Unique coverage is calculated by subtracting from the solution coverage the amount of coverage that is obtained by all paths except the one whose unique coverage we are interested in. For instance, the unique coverage of path ~A~C would be calculated in the following way:

Unique coverage ~A~C = solution coverage – coverage (~BC + F~D).

A Venn diagram might help convey an intuitive representation of the different types of coverage. Figure 5.5 displays a Venn diagram for the solution term:

$X_1 + X_2 + X_3 \rightarrow Y$.

The rectangular box denotes all cases in the study. The largest set is outcome Y and each circle represents one of the three sufficient paths X_1–X_3. Needless to stay, X can stand for a conjunction of conditions. Furthermore, since the circles for X_1–X_3 are fully contained within the set of Y, we know that each single path and the entire solution term are fully consistent as sufficient conditions. What varies between paths is their raw and unique coverage.

The raw coverage of a single path is represented by the size of its set in relation to the size of set Y. We see that the raw coverage of X_2 is higher than that of X_1, which is higher than that of X_3, simply because area (IV) is bigger than areas (I) and (II), which are bigger than areas (II) and (III). The unique coverage is that area of a condition that does not overlap with another sufficient condition. As Figure 5.5 shows, paths X_1 and X_3 partially overlap. Therefore, the unique coverage of X_1 is equal to area (I) whereas that of X_3 is equal to area (III). Since condition X_2 does not overlap with any of the other paths, its unique coverage is the same as its raw coverage. Finally, the solution coverage of the term $X_1 + X_2 + X_3 \rightarrow Y$ is the sum of the areas (I)–(IV) in relation to the area of Y. Since the three paths jointly do not fill the entire circle for outcome Y, we can also see that the solution coverage is lower than 1.

Computationally, the calculation of coverage (and also of consistency) is not very demanding. Although all relevant software packages except Tosmana 1.3.2 automatically provide these parameters of fit, it might be helpful to demonstrate how coverage (and consistency) are calculated by hand. We do this with an example from Vis (2009),[11] who aims at explaining why some

[11] In the remainder of the book, we repeatedly refer to data from this and other published examples, and we will introduce the study in further detail below (section 8.2).

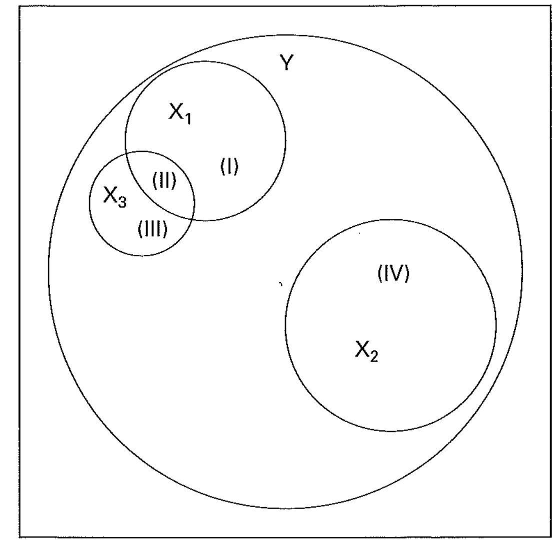

Figure 5.5 Venn diagram – equifinal solution term and types of coverage

governments in Western Europe engage in unpopular reforms (U). She finds that a weak political position of the governments (P) combined with a weak socio-economic situation (S) or a right-wing government (R) combined with a weak socio-economic situation (S) are the sufficient conditions for unpopular reform. Formally:

$P{*}S + R{*}S \rightarrow U$.

Table 5.3 contains the membership scores for the 25 cases, in both the solution and the outcome. In addition, the last column indicates for each case the minimum membership score across the solution and the outcome. This is a crucial quantity as it is used in the numerator for calculating both consistency and coverage. Calculating the solution coverage is straightforward. We simply add up the values in column "min(PS+RS,U)" and divide it by the sum of the membership scores in the column for outcome U. This yields a coverage for PS+RS of

Solution Coverage (PS+RS) = 10.96 / 12.74 = 0.86.

Each case with a higher membership in U than in PS+RS contributes to less-than-perfect coverage. Among them, cases like Kok I or N. Rasmussen IV are particularly striking, for they are more in than out of the outcome set U but more out than in both sufficient paths and thus remain uncovered and therefore unexplained by solution term PS+SR.

Table 5.3 Fuzzy-set membership in solution and outcome (Vis 2009)

	Solution	Outcome	
Government	PS + RS	U	min(PS+RS,U)
Lubbers I	0.83	0.83	0.83
Lubbers II	0.33	0.33	0.33
Lubbers III	0.60	0.67	0.60
Kok I	0.40	0.67	0.40
Kok II	0.33	0.17	0.17
Balkenende II	0.67	0.83	0.67
Kohl I	0.33	0.33	0.33
Kohl II	0.17	0.17	0.17
Kohl III	0.33	0.33	0.33
Kohl IV	0.67	0.67	0.67
Schröder I	0.33	0.17	0.17
Schröder II	0.83	0.83	0.83
Schlüter I	0.33	0.33	0.33
Schlüter II	0.60	0.67	0.60
Schlüter IV	0.67	0.17	0.17
Schlüter V	0.67	0.33	0.33
N. Rasmussen I	0.17	0.17	0.17
N. Rasmussen II (& III)	0.60	0.83	0.60
N. Rasmussen IV	0.33	0.67	0.33
Thatcher I	0.83	0.83	0.83
Thatcher II	0.33	0.67	0.33
Thatcher III	0.67	0.67	0.67
Major I	0.60	0.67	0.60
Blair I	0.17	0.40	0.17
Blair II	0.33	0.33	0.33
SUM fuzzy membership	(a) 12.12	(b) 12.74	(c) 10.96
Coverage sufficiency (c/b)		0.86	
Consistency sufficiency (c/a)		0.90	

As a matter of fact, the calculation of consistency for the solution PS+RS is equally straightforward. Simply add up the scores in the last column once again, but this time divide it by the sum of scores in the path PS+RS:

Solution Consistency = 10.96 / 12.12 = 0.90.

Less-than-perfect consistency is caused by cases whose membership in PS+RS exceeds their membership in U such as, for instance, Kok II, Schröder I, Schlüter IV and V.

The calculation of raw and unique coverage is equally simple. We demonstrate it for path PS. Table 5.4 displays each case's fuzzy set membership scores in path PS, outcome U, and the minimum score across these two sets.

By adding up the scores in the last column and dividing the total by the sum of column PS, we obtain a consistency value of 7.94 / 8.69 = 0.91. The raw coverage of PS is 7.94 / 12.74 = 0.62.

In order to calculate the unique coverage of path PS, we need to subtract from the solution coverage all of what can be covered by any other path in the solution except PS. Since, in our example, there is only one other path (RS),[12] we have to calculate the coverage of RS (0.71) and then subtract it from the solution coverage (0.86) in order to obtain the unique coverage of path PS:

Unique coverage PS: 0.86 – 0.71 = 0.15.

The calculation of the unique coverage of path RS (not displayed in Table 5.4) is equally simple. From the solution coverage we subtract the raw coverage of path PS:

Unique coverage RS: 0.86 – 0.62 = 0.24.

The unique coverage scores reveal that each path has some unique contributions to covering the outcome. In general, marginal differences in the coverage level should not be over-interpreted. Of equal, if not more, interest should be the cases that are uniquely covered. A case is uniquely covered if it holds a membership value higher than 0.5 in only one sufficient path (Schneider and Rohlfing in press and section 11.4, below). In our example, it turns out that out of the five cases that are more in than out of path PS, only two – Schröder II and Rasmussen II (&III) – are uniquely covered. The other three also have a membership greater than 0.5 in path RS. Path RS, in turn, has ten cases with membership greater than 0.5, and seven of them are uniquely covered by that path. Path RS is therefore empirically more important than path PS to an extent beyond what is reflected by comparing only their unique coverage formulas. The practical suggestion is that researchers should not only calculate, report, and interpret the raw and unique coverage scores, but also should go back to the cases and identify the uniquely covered cases.

Notice that for consistency, we argued that a lower threshold exists in principle, even if its precise location is subject to judgment (section 5.2). For

[12] If the solution term consists of more than two paths, then we must calculate the joint coverage of all paths except the one we are interested in. It is not correct to simply add up the raw coverage of all these paths, because paths might partially overlap.

Table 5.4 Fuzzy-set membership in path PS and outcome (Vis 2009)

	Path	Outcome	
Government	PS	U	min(PS,U)
Lubbers I	0.33	0.83	0.33
Lubbers II	0.17	0.33	0.17
Lubbers III	0.33	0.67	0.33
Kok I	0.17	0.67	0.17
Kok II	0.33	0.17	0.17
Balkenende II	0.67	0.83	0.67
Kohl I	0.17	0.33	0.17
Kohl II	0.17	0.17	0.17
Kohl III	0.17	0.33	0.17
Kohl IV	0.67	0.67	0.67
Schröder I	0.33	0.17	0.17
Schröder II	0.83	0.83	0.83
Schlüter I	0.33	0.33	0.33
Schlüter II	0.33	0.67	0.33
Schlüter IV	0.33	0.17	0.17
Schlüter V	0.6	0.33	0.33
N. Rasmussen I	0.17	0.17	0.17
N. RasmussenII (& III)	0.6	0.83	0.6
N. Rasmussen IV	0.33	0.67	0.33
Thatcher I	0.17	0.83	0.17
Thatcher II	0.33	0.67	0.33
Thatcher III	0.33	0.67	0.33
Major I	0.33	0.67	0.33
Blair I	0.17	0.4	0.17
Blair II	0.33	0.33	0.33
SUM fuzzy membership	(a) 8.69	(b) 12.74	(c) 7.94
Coverage sufficiency (c/b)		0.62	
Consistency sufficiency (c/a)		0.91	

coverage, no lower threshold exists. The reason for this is that consistency establishes whether a subset relation exists, whereas coverage expresses how empirically important a subset relation is. Conditions with low coverage cover only a little of the outcome of interest, but that little might be of huge theoretical or substantive importance. Of course, conditions with zero unique coverage should be either disregarded or interpreted with care. Such zero coverage will always happen when logically redundant prime implicants (section 4.3.2) are included in the solution term.

At-a-glance: coverage of sufficient conditions

Coverage sufficiency expresses how much of the outcome is covered (explained) by the condition in question. The formula sums all minima of X and Y in the numerator and divides it by the sum of all Y values.

Raw coverage indicates how much of the membership in the outcome is covered by the membership in a single path; the **unique coverage** instead indicates how much a single path *uniquely* covers. The **solution coverage** expresses how much is covered by the entire solution term.

The empirical importance expressed by coverage is not the same as the theoretical or substantive relevance of a sufficient condition. Thus, low-coverage paths might still be of great substantive interest.

Uniquely covered cases are those that hold a membership value higher than 0.5 in only one sufficient path. When substantively interpreting sufficient paths and assessing their importance, researchers should make reference to these uniquely covered cases.

Unlike the case with **consistency**, there is no lower threshold for coverage.

5.4 Consistency of necessary conditions

The notions of consistency and, with some qualification, coverage can be applied to necessary conditions. If X is necessary for Y, then X is a superset of Y, whereas if X is sufficient for Y, then it is a subset of Y (Chapter 3). One consequence of this mirror-image relation between necessity and sufficiency is that the formulas for the parameter of fit are closely related. As a matter of fact, as we will see now, the formula for consistency sufficiency is mathematically identical to the standard formula for coverage necessity, and the formula for coverage sufficiency is mathematically identical to that for consistency necessity. In the following, we explain the rationale for these formulas.

Let us start by having a look at the following three two-by-two tables. Each of them displays the same outcome Y but three different conditions (X_4, X_5, and X_6). The numbers in the cells indicate the number of cases.

If a condition is necessary for the outcome, then no case may show the outcome without the condition. This means that cell a of our two-by-two table must be empty. When making a statement of necessity for outcome Y, then, cases that do not show the outcome are irrelevant (section 3.2.1.1). In two-by-two tables, this means that cells c and d are irrelevant for the assessment of necessity. The degree to which a condition is consistent with the statement of necessity thus depends on the ratio of cases in cells a and b. If all of these cases are located in cell b, the condition is fully consistent. The more of these cases fall into cell a, the lower consistency becomes.

Table 5.5 Two-by-two tables – consistent and inconsistent necessary conditions

Outcome Y		X_4 = 0	X_4 = 1	X_5 = 0	X_5 = 1	X_6 = 0	X_6 = 1
	1	0 a	100 b	10 a	90 b	80 a	20 b
	0	10 c	100 d	10 c	100 d	10 c	100 d

For condition X_4, cell a in Table 5.5 is empty. Thus it is a fully consistent necessary condition for Y. What about conditions X_5 and X_6? For both conditions, cell a contains cases. Certainly, X_5 and X_6 are not fully consistent necessary conditions for Y. X_6 is less consistent with the statement of necessity than X_5, because of those cases that matter (cells a and b) more cases (50) are located in the forbidden cell than for X_5 (10). Ragin (2006) suggests the following formula for calculating consistency of a necessary condition:

$$\textit{Consistency of X as a necessary condition for } Y = \frac{\textit{Number of cases where } X=1 \textit{ and } Y=1}{\textit{Number of cases where } Y=1}.$$

In the numerator we add up all cases that are members of both the outcome and the necessary condition and in the denominator we add up all cases that are members of the outcome. Applied to Table 5.5, the formula can be rewritten as:

$$\textit{Consistency of X as a necessary condition for } Y = \frac{\textit{Number of cases cell b}}{\textit{Number of cases cells } a+b}.$$

Plugging in the values for the conditions X_4, X_5, and X_6, we obtain the following consistency scores as necessary conditions:

$X_4 = 100/(0 + 100) = 1$
$X_5 = 90/(10 + 90) = 0.9$
$X_6 = 20/(20 + 80) = 0.2$.

When discussing consistency and coverage for sufficiency, we have already pointed out that with fuzzy sets, one should make use of the more fine-grained information contained in fuzzy-set membership scores. The same holds true

when dealing with necessity. With fuzzy sets, the consistency of a necessary condition is given by the degree to which each case's membership in X is equal to or greater than their membership in Y. When calculating consistency necessity, we therefore relate each case's membership in X that is consistent with the statement of necessity to the sum of each case's membership in Y. This logic can be expressed by the following formula (Ragin 2006):

$$Consistency_{\textit{Necessary Conditions } (X_i \geq Y_i)} = \frac{\sum_{i=1}^{I} \min(X_i, Y_i)}{\sum_{i=1}^{I} Y_i}.$$

If for all cases the X values are equal to or greater than their Y values, then they are all below or on the main diagonal and the formula takes on the consistency value of 1, since the minimum of X and Y is in all cases the Y value. The more cases that display a membership in Y that exceeds their membership in X (and the greater the amount by which Y exceeds X in these cases), the more cases lie above the diagonal (and the farther above the diagonal they lie). In this scenario, the consistency value for the necessary condition deviates ever more from a value of 1, since smaller values go into the numerator than into the denominator.

Let us briefly demonstrate this formula with an example. Schneider, Schulze-Bentrop, and Paunescu (2010) are interested in, among other things, the necessary conditions for the high share of high-tech sector exports in proportion to all exports (EXPORT) in 19 OECD countries from 1990 to 2003 (N = 76). They identify high unemployment protection (EMP); high coverage of collective bargaining (BARGAIN); high share of university-trained citizens (UNI); high share of occupation-trained citizens (OCCUP); high share of stock market capitalized indigenous firms (STOCK); and high share of cross-border mergers and acquisitions as a measure of institutional arbitrage (MA).[13] The results of applying the formula to all conditions and their complements, ordered by consistency values, are shown in Table 5.6.[14]

The condition STOCK has the highest consistency value (0.89) and researchers might see good reasons to interpret it as a necessary condition. However, an inspection of the XY plot (Figure 5.6) reveals that this is not so clear, in the end. First, a considerable number of cases fall above the diagonal. Second,

13 The data matrix can be found in the online appendix (www.cambridge.org/schneider-wagemann).
14 See the online How-to section for Chapter 5 on using the software packages for calculating the parameters of fit (www.cambridge.org/schneider-wagemann).

Table 5.6 Analysis necessity, single conditions (Schneider *et al.* 2010: 255)

Condition	Consistency
STOCK	0.89
UNI	0.81
MA	0.72
~OCCUP	0.71
BARGAIN	0.68
~EMP	0.64
OCCUP	0.58
~BARGAIN	0.50
~MA	0.50
~UNI	0.31
~STOCK	0.24

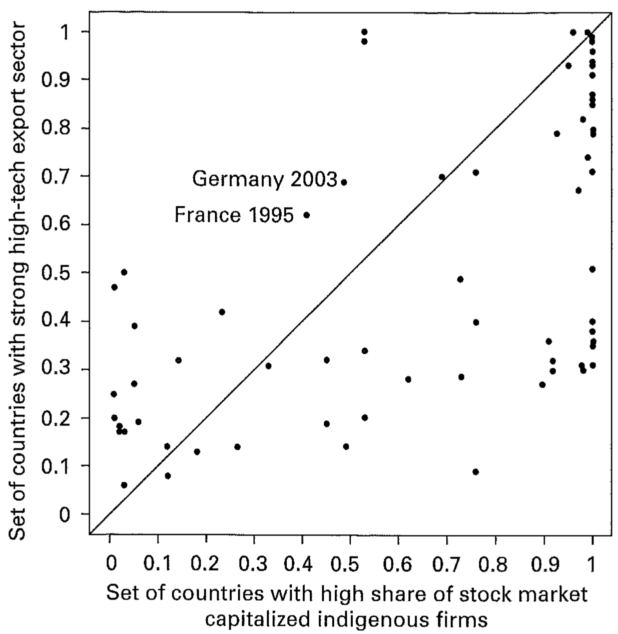

Figure 5.6 XY plot – condition STOCK, outcome EXPORT

among the inconsistent cases there are two true logical contradictory cases (section 5.2; see footnote 1, above): France in 1995 (STOCK = 0.41; EXPORT = 0.62), and Germany in 2003 (STOCK = 0.49; EXPORT = 0.69). Both cases are more out of than in the alleged necessary condition while being more in than out of the outcome. Therefore, interpreting STOCK as a necessary condition for EXPORT does not seem warranted.

More generally, just as with the assessment of sufficiency, so also with necessity it is important that researchers not only use the consistency level, but also check if true logical contradictory cases exist. For necessary conditions, a consistency threshold of at least 0.9 seems advisable (Ragin 2006). One obvious rationale behind this is that higher consistency values reduce the likelihood of true logical contradictions. In section 9.1 we provide further reasons for high consistency levels for necessary conditions.[15]

The formula for the consistency of a necessary condition should look familiar to the reader. In fact, it is mathematically identical to the formula for calculating the coverage of a sufficient condition. However, the two have very different substantive interpretations. The point of a consistency test for a necessary condition is to determine the degree to which an outcome Y is a *subset* of a condition X. We expect from the very beginning that many, if not most, cases display membership values in Y that are smaller than their respective membership in X. In contrast, the purpose of a test for the coverage of a sufficient condition is to find out the portion of an outcome Y that is covered by a consistent sufficient condition X. This means that we will already know that Y is a consistent (enough) *superset* of X, such that the majority of cases will have larger Y values than X values. One practical implication for research is that the calculation of consistency must always precede that of coverage. To start with, is it meaningless to interpret the coverage of a non-consistent necessary or sufficient condition. Moreover, this procedure avoids confusion when interpreting the results obtained from the consistency and coverage formulas for necessity and sufficiency (Ragin 2008a: 63).

At-a-glance: consistency of necessary conditions

The **consistency** measure for **necessary conditions** assesses the degree to which the empirical information at hand is in line with the statement of necessity, i.e., how far the outcome can be considered a subset of the condition. As in the case of **sufficiency**, with **fuzzy sets**, the parameter takes into account both how many cases deviate from the pattern of necessity and how strongly they deviate.

The formulas for consistency of necessity, on the one hand, and **coverage** of **sufficiency**, on the other, are mathematically identical but have different substantive interpretations.

[15] To anticipate the arguments: high consistency thresholds are also conducive to avoiding the pitfalls of (a) necessary conditions disappearing from sufficiency solution terms (*hidden necessary conditions*) and (b) false necessary conditions appearing in sufficiency solutions (*false necessary conditions*).

5.5 Coverage of necessary conditions

The reader might already suspect that the mutual, formal equivalence of coverage and consistency between necessary and sufficient conditions might also be extended to the *coverage* of necessary conditions. Following this reasoning, the formula for the coverage of necessary conditions should be equal to the formula for the consistency of sufficient conditions and should therefore read for crisp sets:

$$\textit{Coverage of X as a necessary condition for } Y = \frac{\textit{Number of cases where } X = 1 \textit{ and } Y = 1}{\textit{Number of cases where } X = 1}$$

and, applied to a two-by-two table, as:

$$\textit{Coverage of X as a necessary condition for } Y = \frac{\textit{Number of cases cell } b}{\textit{Number of cases cells } b + d}$$

and, for both fuzzy and crisp sets:

$$\textit{Coverage of X as a necessary condition for } Y = \frac{\sum_{i=1}^{I} \min(X_i, Y_i)}{\sum_{i=1}^{I} X_i}.$$

And, indeed, these are the formulas for the coverage of a necessary condition as suggested by Ragin (2006a, 2008a: 61) and currently implemented in the relevant software.

The formula for the coverage of a necessary condition expresses how much smaller the outcome set Y is in relation to set X. According to this formula, if X and Y are of roughly equal size, then the coverage of X as a necessary condition is high. Put differently, the more the size of X exceeds that of Y, the lower the coverage of X as a necessary condition.

The label *coverage* is misleading, though. If X has passed our consistency test as a necessary condition, then, by definition, X is a superset of Y and thus X fully covers Y. In other words, by virtue of being necessary, X always fully covers all cases of membership in Y. Ragin (2008a: 60–63) and Goertz (2006a) therefore point out that, next to consistency, the issue at stake when dealing with necessary conditions, is that of *relevance* (Ragin) or *trivialness* (Goertz). Thus, despite all symmetry in these parameters, the interpretation of the coverage value for necessity and that for sufficiency are fundamentally different.

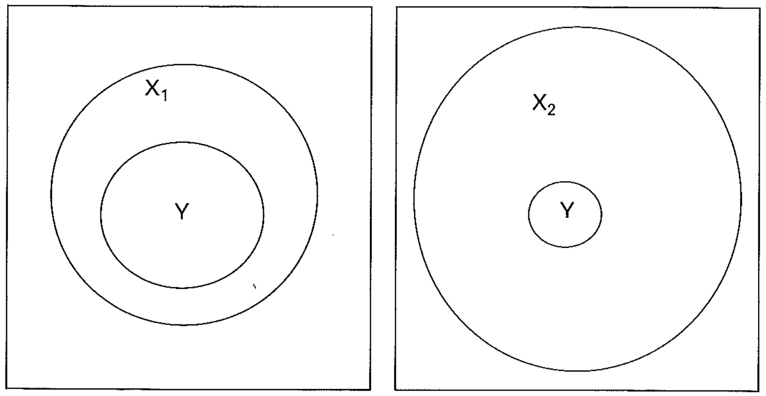

Figure 5.7 Venn diagrams – trivial and non-trivial necessary conditions

In order to understand what is meant with relevance and trivialness, consider the two Venn diagrams in Figure 5.7. Let Y be the set of speeches in a country's parliament during which parliamentarians curse. X_1 is the set of male members of parliament and X_2 the set of parliamentarians born in that country. Clearly, both conditions are fully consistent supersets of the outcome and thus pass the formal requirement as necessary conditions. The relation in size of sets X_1 and Y is more in proportion than that between X_2 and Y. Hence, if we applied the coverage formula to these two empirical scenarios, X_1 (male persons) would receive a higher score and thus be deemed more relevant as a necessary condition for cursing than X_2 (being born in the country). X_2 is a trivially necessary condition for Y, simply because so many more members in parliament are born in the country (X_2) than curse during parliamentary debate (Y). The coverage formula suggested by Ragin (2006) and described here adequately captures this form of trivialness.

Let us apply Ragin's coverage formula to the example by Schneider *et al.* (2010). Calculating coverage only makes sense for those conditions that have passed the consistency threshold. Schneider *et al.* (2010: 255) convincingly argue that condition MA can be interpreted as a functional equivalent (see section 3.2.1.2) to condition STOCK. As Table 5.7 shows, the consistency value of the term MA+STOCK is above the 0.9 threshold.

Table 5.7 Analysis necessity, functional equivalents (Schneider *et al.* 2010: 255)

Condition	Consistency	Coverage
STOCK	0.89	0.72
MA+STOCK	0.92	0.68

The coverage value for the disjunction is 0.68. This is lower than the value of STOCK alone.[16] This suggests that the size of the logical OR set, compared to the outcome set, has increased, which should come as no surprise, because combining sets with the logical OR requires taking the maximum value of each case across the combined sets (see section 2.4). Since membership in outcome Y remains the same, the relation in size between sets Y and STOCK, on the one hand, and Y and MA+STOCK, on the other, increases. In an XY plot, this is graphically displayed by more cases falling further to the right-hand side of the plot. Just compare the XY plots in Figure 5.6 and Figure 5.8.

Two points are worth mentioning about the coverage formula. First, values for coverage necessity tend to be rather high. Unlike coverage sufficiency, in research practice, values far below 0.5 are rare and those close to 0 hardly ever seen. This suggests that when assessing the trivialness of necessary conditions, researchers should not be misled by seemingly high coverage values. In addition, the XY plot should always be carefully examined to ascertain whether most cases are clustering close to the vertical right axis thus suggesting trivialness.

The second issue related to the correct interpretation of the coverage formula for assessing trivialness is this: a condition X can be trivially necessary even when it is of roughly equal size to outcome Y. This happens when not only X, but also Y are very big in size and thus close to being constants (Goertz 2006a). In such a scenario, the formula for coverage necessity will yield a high value and researchers might be inclined to interpret X as a relevant necessary condition. This seems odd, though. Because of their size, both X and Y cover almost all cases and come thus very close to the universal set. Indeed, there are two sources of trivialness of a necessary condition: first, X is much bigger than Y; second, X and Y are close to being constants. Both sources of trivialness need to be taken into account and no condition interpreted as necessary in either of the two situations. The currently predominant formula

[16] However, in both analyses, the same two true logical contradictory cases occur (France in 1995 and Germany in 2003), thus providing further illustration that high consistency values alone are often not enough for a definite statement on a set relation.

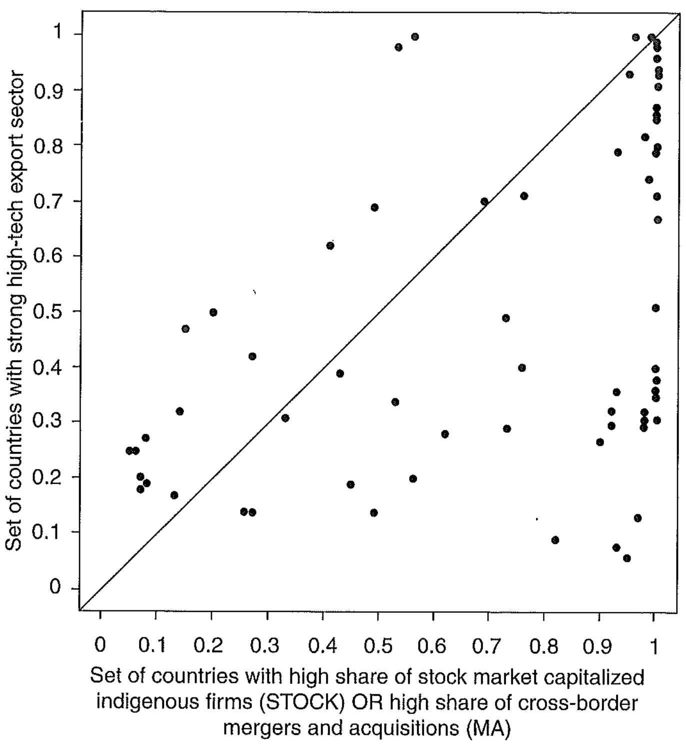

Figure 5.8 XY plot – condition MA+STOCK, outcome EXPORT

for coverage necessity, however, which we have presented here handles only the first source of trivialness well. In section 9.2.1, we provide a detailed discussion of this issue and suggest an alternative formula for calculating the relevance of a necessary condition which also takes into account the second source of trivialness.

At-a-glance: coverage of necessary conditions

The standard **coverage** measure for **necessary conditions** is better interpreted as a measure of the **relevance** of a necessary condition.

High values indicate relevance, whereas low values indicate **trivialness**.

Conditions that pass the **consistency** test as a necessary condition should not be deemed to be relevant necessary conditions unless they also obtain a high value in the relevance measure.

The coverage measure for necessity captures only one source of trivialness, though. It detects whether the outcome set is much smaller than the condition set but is not capable of capturing whether both the condition and the outcome are (close to) universal sets.

5.6 Issues related to consistency and coverage

The concepts of consistency and coverage contribute in important ways to making set-theoretic methods, in general, and QCA, in particular, a more adequate and useful tool for analyzing social science questions. They allow for the use of set theory and formal logic to find patterns in noisy social science data. Despite – or perhaps precisely because of – their usefulness researchers employing QCA should resist the temptation to reduce this method to a simple hunt for high values of consistency and coverage. This would clearly be against the spirit of set-theoretic methods and would deprive them of their main strength: being grounded in the qualitative research practices of engaging in an iterative dialogue between ideas and evidence. Consistency and coverage are better thought of as numerical summaries that describe the data patterns in the underlying dataset. QCA is above all a *qualitative* data technique, and its primary purpose consists in interpreting and understanding the cases under study. Neither should specific consistency values obtain the status of universally applicable thresholds. Nor should individual cases disappear behind, or be hidden by, consistency and coverage values. Instead, researchers must carefully judge and then explicitly argue which consistency threshold is adequate for their specific research and then also perform several analyses with consistency values that vary within a reasonable range. When using fuzzy sets, we also advise paying close attention to which cases are true logical contradictions (consistency), uniquely covered, and which ones are not covered at all.

Consistency is the parameter which should always be assessed first. The reason is straightforward. It only makes sense to calculate the coverage of a sufficient (or necessary) condition if that condition has already been identified as being consistently sufficient (or necessary). If the consistency value is too low for the condition to be considered sufficient (or necessary), the calculation of coverage is meaningless. Along these lines, while there are consistency levels below which a condition cannot be considered as sufficient (or necessary), such lower-bound thresholds do not exist for coverage. With sufficiency, very low coverage values indicate that only a small portion of the outcome of interest is explained by that condition. However, that little bit might still be of great theoretical and substantive importance. With necessity, low levels of coverage indicate trivialness whereas high levels might or might not indicate relevant conditions, an issue we come back to in section 9.2.1.

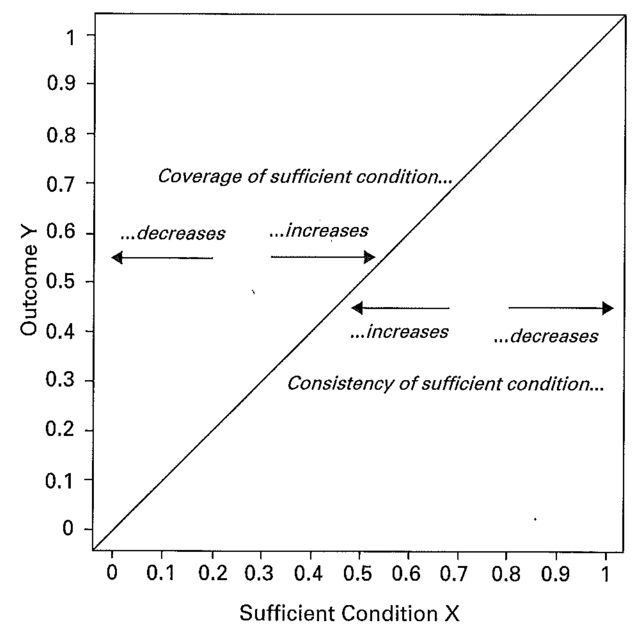

Figure 5.9 XY plot – the tension between consistency and coverage of sufficient conditions

Also note that, in research practice, higher consistency values often come at the price of lower coverage values. In the analysis of sufficiency, this works the following way. We can increase consistency by adding single conditions through logical AND. For instance, we might enlarge conjunction A*B*C to conjunction A*B*C*D*E. The more conditions that are combined, the more difficult membership in it becomes (section 2.1). This makes the set ever smaller, and thus makes it more likely to be a consistent subset of the outcome. At the same time, however, and precisely because membership becomes more and more difficult, long conjunctions cover less and less of the outcome simply because so few cases are members of this conjunction. A similar logic applies to the analysis of necessity. Here we increase consistency by adding conditions through logical OR. For instance, we extend expression A+B+C to A+B+C+D+E. The more conditions that are added, the easier membership in it becomes. This makes the set ever bigger and thus more likely to be a consistent superset of the outcome. But at the same time, and as with the process just described for conjunctions, long OR expressions cover more and more cases of the entire set of cases under study since membership becomes ever easier, and they thus risk becoming trivial necessary conditions (section 9.2.1). The inherent tradeoff between consistency and coverage is graphically depicted in the XY plot in Figure 5.9 for an analysis of sufficiency, but works in the same way in the analysis of necessity.

At-a-glance: issues related to consistency and coverage

Consistency is the central measure for the assessment of set relations. Only if consistency is satisfactory should **coverage** be calculated.

Often it is not possible to achieve high values for the consistency and coverage measures at the same time. Indeed, there is a tradeoff between the two: to increase consistency often means to decrease coverage and vice versa.

Parameters of fit are not an end in themselves. The main focus should always be on the cases under study. Researchers should identify the cases that contribute to inconsistency and to low coverage.

6 Limited diversity and logical remainders

Easy reading guide

As seen in Chapter 4, the analysis of truth tables is at the core of QCA. In Chapter 5 we presented, among other things, the consistency value as a parameter for assessing whether a given truth table row could be considered a subset of, and thus sufficient for, the outcome. What if, however, there is not enough empirical evidence for a given row in order to assess whether it is sufficient? In other words, what if a row consists of a conjunction of properties that is logically possible but not empirically observed? Treating these so-called *logical remainder rows* in a conscious manner is both crucial for, and an asset of, set-theoretic methods. As this chapter shows, assumptions about remainders do have a direct impact on the results obtained and some assumptions are more plausible than others.

The presence of logical remainders is called *limited diversity*. This can be defined as the set of all logically possible combinations of conditions for which either no or not enough empirical evidence is at hand. It is a universal phenomenon in comparative social science research. The effect that these logically possible, yet empirically unobserved, "cases" have upon the possibilities for drawing evidence-based inferences is perhaps among the most understudied topics in social science research methodology. This is why we dedicate extensive space here to the different strategies that researchers facing limited diversity should be aware of. In Chapter 8, we add further strategies that go beyond the current best practice approach.

In this chapter, we first explain how to detect logical remainders (6.1). Second, we discuss why virtually all social science data is limited in its diversity. We do so by differentiating between different sources, and thus different types, of logical remainders (6.2). Since limited diversity afflicts the capacity for drawing inference, regardless of which specific method is applied, be it statistical or not, we then delimit in section 6.3 the phenomenon of logical remainders from other seemingly related notions in the social science methodology literature (such as missing values). In the final, main section of this chapter, we spell out the principles of the so-called Standard Analysis (Ragin 2008b) as the currently predominant procedure in applied QCA for making plausible assumptions about logical remainders (6.4). The aim of this chapter is to formulate set-theoretic strategies that help to keep the impact of logical remainders on inferences under the conscious control of the researcher.

The proper handling of logical remainders is of central importance for QCA. This chapter is certainly a must-read for beginners. Even experienced users will profit from studying