t distribute.

STRATEGIES FOR SOCIAL INQUIRY

Set-Theoretic Methods for the Social Sciences

A Guide to Qualitative Comparative Analysis

CARSTEN Q. SCHNEIDER
CLAUDIUS WAGEMANN

CAMBRIDGE

Study material. Do not distribute.

These aspects also enable us to analyze **INUS** and **SUIN conditions** with the help of QCA. INUS conditions are defined as insufficient but necessary parts of a condition which is itself unnecessary but sufficient for the result; SUIN conditions refer to sufficient, but unnecessary, parts of a factor that by itself is insufficient, but necessary, for the result.

Causally complex results produced by set-theoretic methods differ from those produced by standard statistical (regression-based) approaches. While more advanced quantitative techniques can mimic some aspects of causal complexity, achieving all of them simultaneously is a challenging and still unresolved task. However, no form of causality induced by the choice of method can be considered superior per se. Instead, researchers should choose the method that rests on assumptions which are most in line with their research question. If hunches about necessary and sufficient conditions exist (and there are many research fields for which this is the case, as Goertz and Starr 2003 and Seawright 2002b: 180f., show), then set-theoretic methods are a plausible choice.

Note that, in this chapter, we have only labeled those conditions as sufficient or necessary for which all empirical evidence was in line with these respective set relations. However, some of the examples have already alluded to the fact that there can be different degrees of deviation from perfect subset relations. In fact, when applying set-theoretic methods to social science data, such observations are the norm. We will discuss this in more detail in Chapter 5.

4 Truth tables

Easy reading guide

In this chapter, we introduce a concept that is at the core of QCA, both in the understanding of it as an approach and as a technique: truth tables. QCA understood as an approach can be perceived as a research phase that aims to construct a truth table. Truth tables contain the empirical information gathered by the researcher, often through years of painstaking work. QCA as a technique, then, consists of the formal analysis of truth tables – the so-called logical minimization – with the aim of identifying sufficient (and necessary) conditions. As such, truth tables become *the* indispensable tool for any QCA, no matter whether we are working with crisp or fuzzy sets. This is one of the primary bases for the argument that crisp-set QCA and fuzzy-set QCA are not fundamentally that different. It also means that most of what we say in this book about truth tables, and their analysis, applies both to csQCA and to fsQCA.[1]

We deem it important to reiterate that, in this and other chapters, we mostly focus on issues related to QCA as a technique for pedagogical reasons, and we therefore take for granted the existence of empirical information upon which the truth table is constructed. However, one integral part of set-theoretic approaches – and the key to their success – consists precisely in the process of collecting this information and constructing truth tables in an iterative process, a process sometimes described as the "back and forth between ideas and evidence" (Ragin 1987). The analysis of the truth table only represents a short "analytic moment" (Ragin 2000) in the process of performing set-theoretic analysis.

In Chapter 3, we engaged in the analysis of necessity and sufficiency without making use of truth tables. One might therefore wonder why we would need truth tables if necessity and sufficiency can also be analyzed simply by screening a standard data matrix. As this chapter will show, truth tables are a much more adequate device for detecting set relations, mainly because they shift the focus from empirical cases to configurations of conditions. This leads to a radically different – and more efficient – approach to the analysis of sufficiency. The analysis of sufficiency based on a data matrix proceeds in a

[1] There are only a few analytically relevant differences in the analysis of a truth table that follow from the difference between crisp and fuzzy sets, such as, for instance, the possibility in fuzzy sets that a given truth table row is a subset of outcome Y but also of its complement ~Y. We will discuss this in section 9.2.2.

Study material. Do not distribute.

bottom-up manner by first focusing on simple sets and then proceeding to more complex sets. In contrast, the analysis of sufficiency based on a truth table proceeds top-down, by first screening all logically possible combinations of conditions and then logically minimizing those conjunctions that have passed the test of sufficiency. Notice, however, that while for sufficiency the truth table approach is (and should be) the dominant strategy, for the analysis of necessity, the bottom-up approach is instead clearly preferable and the top-down approach is meaningless. The reason is simple: a logical AND conjunction of two or more conditions can only be necessary for Y if, and only if, all single conditions involved in the conjunction are necessary on their own.

The organization of this chapter is straightforward: after clarifying what a truth table is (4.1), we show how truth tables are constructed based on empirical information about cases (4.2). In section 4.3 we explain, step by step, how truth tables are analyzed with the help of Boolean algebra. Clearly, this chapter is central to the whole book, simply because truth tables are the *sine qua non* technique for QCA. The chapter should be read in detail and with care. This chapter also provides important information which will supply the main ingredients for the Truth Table Algorithm, the currently accepted minimum standard for a QCA, as it will be introduced in Chapter 7.

4.1 What is a truth table?

The concept of a truth table originates in formal logic. At first glance it might look a lot like a standard data matrix. Just like conventional data matrices, each truth table column denotes a different variable or, better, set. The difference consists in the meaning of rows. In a standard data matrix, each row denotes a different case (or unit of observation). In a truth table, each row instead represents one of the logically possible AND combinations between the conditions. Since each single condition can occur either in its presence or its absence, the total number of truth table rows is calculated by the expression 2^k. The letter k represents the number of conditions used and the number 2 the two different states (presence or absence) in which these conditions can occur. Each row denotes a *qualitatively different* combination of conditions, i.e., the difference between cases in different rows is a difference in kind rather than a difference in degree.

The formula 2^k yields the number of *logically possible* combinations or truth table rows or, slightly misleadingly, logically possible cases. The number of truth table rows increases exponentially with the number of conditions. With three conditions, we end up with eight configurations. With 4 conditions, we already have 16 configurations, with 5 we have 32, and with 10 we have no fewer than 1,024 logically possible cases. In social reality and therefore also in social science research practice, not all of these potential cases materialize empirically. The whole of Chapter 6 is dedicated to the phenomenon of *limited diversity* and provides strategies on how to handle *logical remainders* (Ragin 1987: 104ff.). For the time being, and in order to properly introduce the meaning and analysis of truth tables, in the current chapter we only deal with truth table rows that do not show any such logical remainders.

Venn diagrams are another way to intuitively visualize that k number of conditions produce 2^k logically possible combinations. Figure 4.1 displays three conditions (A, B, C). They all overlap in various ways, creating eight different areas. Each area in the Venn diagram corresponds to one row in a truth table, and each area can be described in the form of a Boolean expression. For example, the area in the middle of the diagram where A, B, and C overlap is the one that contains all the cases where A, B, and C are present. This can be written as A*B*C, or simply ABC (Chapter 2). The upper area of set A is where condition A can be observed (A = 1) and B and C cannot (B = 0, C = 0). This area thus denotes the set A*~B*~C or simply A~B~C. The area outside all three of the circles, but within the rectangle, denotes cases where none of the three conditions is present and can be written as ~A*~B*~C or ~A~B~C, and so on.

While Venn diagrams are generally a very useful tool for the graphical representation of set-theoretic statements, two caveats need to be made. First, as the number of conditions grows beyond four or five, it becomes difficult to draw and interpret Venn diagrams. Second, note that Venn diagrams such as the one displayed in Figure 4.1 display only sets and their intersections. In Chapter 3, however, we used Venn diagrams to visualize subset relations of sufficiency and necessity between conditions and an outcome. Of course, Venn diagrams can do both simultaneously, i.e., show the subset relation of an intersection of conditions and an outcome.

4.2 How to get from a data matrix to a truth table

4.2.1 Crisp sets

In order to show how to construct a truth table based on information on cases stored in a data matrix, let us go back to our data matrix from section 3.1.1.2. How do we get from here to a truth table? While most of the relevant software

Study material. Do not distribute.

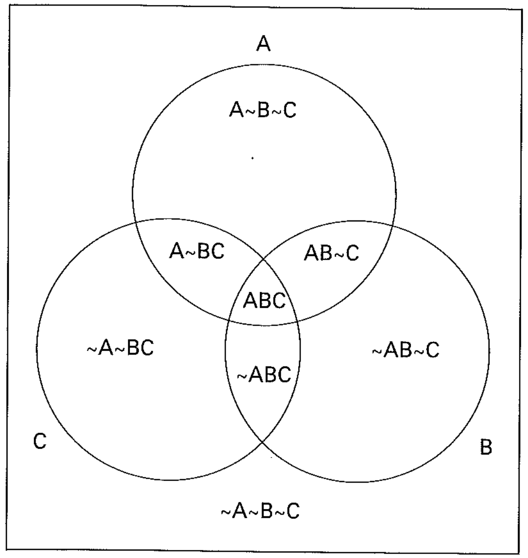

Figure 4.1 Venn diagram with three conditions

packages are able to produce a truth table based on a data matrix representing set-membership scores, it is worth spelling out the three simple steps that are needed.

First, we write down all 2^k logically possible combinations of the k conditions, leaving the column for the outcome value empty. Second, we assign each case from our data matrix to the truth table row that corresponds with its values in the k conditions. Each case can belong to only one truth table row, but individual truth table rows might contain more than one case. In our example, we observe that Argentina and Venezuela display identical values on all three conditions – they had a violent upheaval, have an ethnically homogeneous population and a pluralistic party system. They therefore belong to the same truth table row labeled A*B*C. The same holds true for Peru and Ecuador, which are both assigned to the truth table row A*~B*~C (violent upheaval, no ethnically homogeneous population, no pluralistic party system). In this way, we assign each case to one of the eight logically possible truth table rows.

Third, an outcome value has to be attributed to every truth table row. It is determined by the outcome values of the empirical cases that fall into the respective row. For instance, Colombia falls into row ~A*~B*~C and shows outcome Y. No other case falls into this row. Hence, the outcome value of row ~A*~B*~C is Y = 1. Likewise, neither Argentina nor Venezuela shows a stable

Table 4.1 Data matrix with ten cases, three conditions, and outcome

		Conditions			Outcome
Row	Cases	A	B	C	Y
1	ARG	1	1	1	0
2	PER	1	0	0	0
3	BOL	1	1	0	0
4	CHI	0	1	0	1
5	ECU	1	0	0	0
6	BRZ	0	1	1	1
7	URU	1	0	1	1
8	PAR	0	0	1	1
9	COL	0	0	0	1
10	VEN	1	1	1	0

Y = set of countries with stable democracies
A = set of countries with violent upheavals in the past
B = set of countries with ethnically homogeneous population
C = set of countries with pluralistic party system

democracy, and the outcome value of truth table row A*B*C is Y = 0.[2] Based on this procedure, the data matrix in Table 3.2 yields the truth table displayed in Table 4.2.

The truth table consists of $2^3 = 8$ rows. Strictly speaking, the columns "Row," "~Y," and "Cases" do not belong to the truth table but are included for illustrative purposes. It is important to understand the information contained in the "Outcome" column. From a case perspective, the value of 1 indicates that cases with the given characteristics also show the outcome of interest. For instance, from row 1 in Table 4.2 we learn that cases that did not have a violent upheaval and have no ethnically homogeneous population and have no pluralistic party system are stable democracies. If we shift perspective from cases to configurations, we can say that conjunction ~A~B~C (row 1) is sufficient for Y. A truth table row with outcome Y = 1 is explicitly linked (Ragin and Rihoux 2004) to this outcome. In essence, each truth table row is a statement of sufficiency (Ragin 2008a).

[2] Of course, when applied to real data, it is common that cases attributed to the same truth table row display different membership scores in the outcome. Such rows are called *contradictory rows* (Ragin 2000). Chapter 5 is dedicated to discussing this crucially important issue. For the time being, in order to present the logic of truth tables and their analysis, we present examples of truth tables that are contradiction-free.

Study material. Do not distribute.

Table 4.2 Hypothetical truth table with three conditions

	Conditions			Outcome		
Row	A	B	C	Y	~Y	Cases
1	0	0	0	1	0	COL
2	0	0	1	1	0	PAR
3	0	1	0	1	0	CHI
4	0	1	1	1	0	BRZ
5	1	0	0	0	1	PER, EC
6	1	0	1	1	0	URU
7	1	1	0	0	1	BOL
8	1	1	1	0	1	AR, VEN

See Table 3.2
~Y = set of countries with non-stable democracies

4.2.2 Fuzzy sets

The three steps for converting a data matrix into a truth table also apply when the underlying data are not crisp but fuzzy sets. We first create the truth table, then assign each case to one of these rows, and then determine the outcome value for each row. Since fuzzy sets allow for any set-membership score between 0 and 1, whereas truth tables consists of only 0s and 1s, this might seem puzzling.

The creation of the truth table is the least problematic step. Just as with crisp sets, the number of truth table rows based on fuzzy sets is given by the formula 2^k. This is because, just like crisp sets, fuzzy sets establish a *qualitative* difference between cases above the 0.5 qualitative anchor (more in than out of the set) vis-à-vis cases below that anchor (more out than in). This is why k fuzzy-set conditions yield 2^k truth table rows.[3]

The attribution of cases to specific truth table rows, a rather straightforward exercise based on crisp sets, requires more explanation when dealing with fuzzy sets. With crisp sets, in order to identify the truth table row to which a case belongs, we simply need to find the *exact* match between the case's crisp-set membership scores and the truth table rows. With fuzzy sets, however, cases with fuzzy-set membership scores in the k conditions do not

[3] The situation is different in multi-value QCA (mvQCA). In section 10.2, we discuss the consequences of the fact that with k-number of multi-value "sets," the number of truth table rows is (much) higher than 2^k.

Table 4.3 Hypothetical data matrix with fuzzy-set membership scores

Row	Cases	A	B	C
1	ARG	0.8	0.9	1
2	PER	0.7	0	0
3	BOL	0.6	1	0.1
4	CHI	0.3	0.9	0.2
5	ECU	0.9	0.1	0.3
6	BRZ	0.2	0.8	0.9
7	URU	0.9	0.2	0.8
8	PAR	0.2	0.3	0.7
9	COL	0.2	0.4	0.4
10	VEN	0.9	0.7	0.6

A = set of countries with violent upheavals in the past
B = set of countries with ethnically homogeneous population
C = set of countries with pluralistic party system

exactly match any of the truth table rows. For instance, to which truth table row does Chile in row 4 of Table 4.3 belong, with its set membership scores of A = 0.3, B = 0.9, and C = 0.2?

In order to shed light on this, Ragin (2008a: ch. 7) refers to the concept of a *property space*, going back to Paul Lazarsfeld's (1937) initial ideas. Each set constitutes one dimension of the property space (Barton 1955). The three fuzzy-set conditions in our example thus yield a three-dimensional space as displayed in Figure 4.2. There are several important features of this property space.

First, regardless of a case's membership in conditions A, B, and C, it falls, by definition, inside the property space. This is because both set membership and the dimensions of the property space have their minimum at 0 and their maximum at 1. Second, based on the set membership in A, B, and C, each case has one precise location inside the cube. Third, each corner of the property space directly corresponds to a specific combination of values in A, B, and C. More precisely, each corner represents one specific combination of the two extreme values that are possible in fuzzy sets – full membership (1) and full non-membership (0). For example, the corner in the bottom left front of Figure 4.2 denotes the situation in which all three fuzzy sets take on the value of 0. This corner can therefore be labeled the "0,0,0" or the ~A~B~C corner. Following this logic, we can describe the lower right corner in the front as "1,0,0," the top right rear corner as "1,1,1," and so on. Fourth, because each

Study material. Do not distribute.

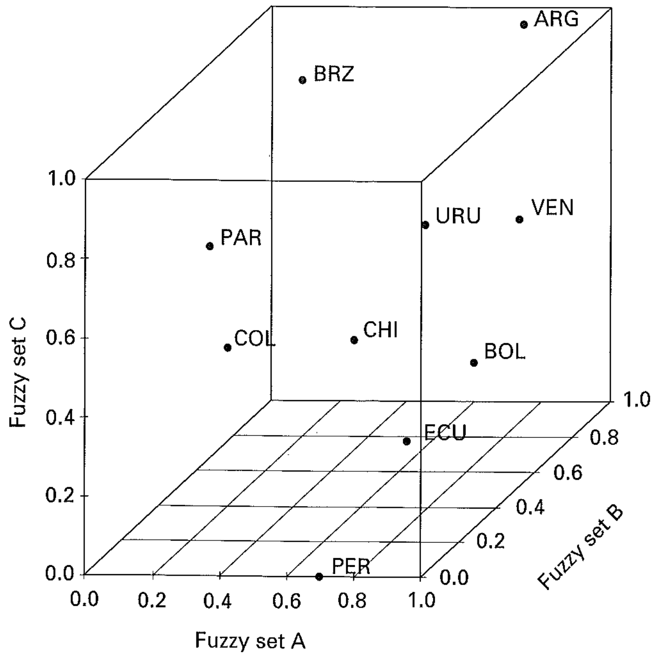

Figure 4.2 Three-dimensional property space

corner denotes a specific combination of extreme membership scores in the conditions, we can perceive of these corners as ideal-typical situations (Weber 1906). Cases that fall exactly in one of the corners are empirical instances of the ideal type denoted by that corner. Unless a case has full (non-)membership in all conditions that constitute the property space, in other words, unless a case exclusively displays crisp-set membership scores, it will not be located directly in one of the corners. Thus, most of the time in fuzzy-set analyses, many, if not all, cases get close to these ideal types only to some (varying) degree. Below, we will explain how the distance to the ideal types can be calculated.

Fifth, a property space with three dimensions has eight corners. This number should ring a bell. A truth table based on three conditions has eight rows. This is no coincidence, but directly follows from the fact that the corners of a property space, spanned by fuzzy sets, are equivalent to the rows of a truth table.[4] This equivalence exists because the corners of a property space defined by fuzzy sets denote the situation where the values of these fuzzy conditions take on the extreme values 0 or 1. In other words, the corners are where the fuzzy sets show crisp-set membership scores.

[4] The metaphor of the cube only works for three conditions. Other geometrical objects would be needed for the representation of other numbers of conditions. However, the basic principle remains the same.

We can summarize the insights gained so far in the following manner. With k number of conditions, we create a property space with 2^k corners and these corners correspond to one of the 2^k (a) ideal types; (b) truth table rows; (c) logical AND conjunctions between the k conditions.

As mentioned, with fuzzy sets cases usually have membership values between 0 and 1. Consequently, they can be located anywhere in the property space, as Figure 4.2 indicates. Some might be closer to one of the corners than to the others. We therefore have to find a way to establish two things: first, to which corner a given case most belongs, and second, how far this case is a member of this ideal type (aka truth table row).

In order to explain the principle by which the membership of cases in each corner is calculated, let us focus on two cases from Table 4.3, Venezuela and Ecuador, reproduced in Table 4.4.[5]

Looking at Venezuela, we see that its membership in all three conditions is above the qualitative anchor 0.5. If asked which of the 2^k ideal types this country resembles most, it is plausible to say that this country is more of an *ethnically relatively homogenous state with a pluralistic party system that experienced a violent upheaval* than any other logically possible type. In other words, intuitively, we would locate Venezuela closest to the ABC row of a truth table or the "1,1,1" corner of a property space, an intuition visually supported by Venezuela's location in Figure 4.2. The same logic of locating a case applies to Ecuador. It is more in than out of set A and more out than in sets B and C, respectively. This makes Ecuador closer to the "1,0,0" corner than any other and an instance of an *ethnically heterogeneous population without a pluralistic party system and without a violent upheaval* (row A~B~C).

Beyond this intuitive attribution of cases to property space corners, aka ideal types, is there a standardized way to precisely define the membership of cases in truth table rows? Yes, there is. Remember that each of the 2^k corners corresponds to one of the 2^k logically possible AND combination of conditions. Remember also (see section 2.1) that the membership of cases in an intersection is determined by their minimum set membership across the single conditions. It is therefore easy to calculate a case's membership in *all* logically possible combinations of conditions, aka corners of the property space. Table 4.5 contains this information for our two cases displayed in Table 4.4.

Venezuela has a fuzzy-set membership of 0.6 in ideal type ABC. This is the minimum across conditions A (0.9), B (0.7), and C (0.6). Ecuador has

[5] We do not report each case's membership in outcome Y, because it is irrelevant for identifying the truth table row a case belongs to. When performing the three steps of converting a data matrix into a truth table, the outcome is added only in the third step.

Study material. Do not distribute.

Table 4.4 Fuzzy-set data matrix with two cases

Case	Conditions A	B	C
VEN	0.9	0.7	0.6
ECU	0.9	0.1	0.3

See Table 4.3

a membership value of 0.7 in the ideal type A~B~C, which is the minimum across A (0.9), ~B (0.9), and ~C (0.7). Both cases are not full instances of their respective ideal types, as indicated by their membership score of less than 1.

As Table 4.5 also shows, each case has a partial membership not only in its own ideal type, but also in all of the other corners of the property space. These membership scores are, however, quite low, a direct consequence of the minimum scoring rule that governs the calculation of set membership scores for conjunctions (section 2.1). The crucial point is that, while each case has a partial membership in all rows, there is only one row in which its membership exceeds the qualitative anchor of 0.5. This is a golden rule for fuzzy sets: no matter how many fuzzy sets are combined, any given case has a membership of higher than 0.5 in one and only one of the 2^k logically possible combinations.

This important mathematical property of fuzzy sets is crucial for our task at hand – identifying the truth table row to which a case best belongs, which turns out to be that truth table row in which its partial set membership is higher than 0.5.

There is one exception to the rule that each case is more in than out of one and only one logically possible combination. Whenever a case holds a membership of exactly 0.5 in one or more of the constitutive conditions, then its membership will not exceed 0.5 in any of the truth table rows. To demonstrate this, we add a third hypothetical case to our data matrix, which has a set membership of 0.5 in condition C. Since both C and ~C take on the value of 0.5, no single ideal type out of the eight possibilities can arrive at a value of greater than 0.5. No minimum from the three single conditions and their complements can be greater than 0.5. Furthermore, there are two ideal types for which the minimum is exactly 0.5. The 0.5 anchor is sometimes referred to as the point of maximum ambiguity (Ragin 2000). It expresses the fact that a case's empirical attributes are such that it cannot be decided whether the case

Table 4.5 Fuzzy-set membership in ideal types for hypothetical data matrix

Case	Conditions			Property space corners/Ideal types/Truth table rows (logically possible combinations of conditions)							
	A	B	C	ABC	AB~C	A~BC	A~B~C	~ABC	~AB~C	~A~BC	~A~B~C
VEN	0.9	0.7	0.6	**0.6**	0.4	0.3	0.3	0.1	0.1	0.1	0.1
ECU	0.9	0.1	0.3	0.1	0.1	0.3	**0.7**	0.1	0.1	0.1	0.1

Table 4.6 Fuzzy-set ideal types for hypothetical data matrix

Case	Conditions			Property space corners/Ideal types/Truth table rows (logically possible combinations of conditions)							
	A	B	C	ABC	AB~C	A~BC	A~B~C	~ABC	~AB~C	~A~BC	~A~B~C
VEN	0.9	0.7	0.6	**0.6**	0.4	0.3	0.3	0.1	0.1	0.1	0.1
ECU	0.9	0.1	0.3	0.1	0.1	0.3	**0.7**	0.1	0.1	0.1	0.1
HYP	0.8	0.1	0.5	0.1	0.1	**0.5**	**0.5**	0.1	0.1	0.2	0.2

is more a member of the set being studied or more a member of the complement of that set. It is because of this ambiguous status that such a case cannot be attributed to any of the 2^k logically possible ideal types that involve this set or its complement.

One practical lesson from this is to be careful about assigning the fuzzy-set membership score of 0.5 to cases. Doing so not only prevents the attribution of such a case to any of the truth table rows, but also represents the weakest possible conceptual statement about that case.

Getting back to our task of representing fuzzy-set data in a truth table, we now know that such a truth table has 2^k rows and that each case is more in than out of one, and only one, of these rows while holding partial membership scores in most, if not all, other rows as well. What remains to be resolved is to determine the outcome value with which each of the 2^k rows is connected. In order to answer this question, remember that each truth table row is a statement of sufficiency. This means that each truth table row should be considered a sufficient conjunction for the outcome if each case's membership in this row is smaller than or equal to its membership in the outcome (see section 3.1.2.1).

Study material. Do not distribute.

Table 4.7 Fuzzy-set membership in rows and outcome

	Conditions			Truth table rows								Outcome
Cases	A	B	C	ABC	AB~C	A~BC	A~B~C	~ABC	~AB~C	~A~BC	~A~B~C	Y
ARG	0.8	0.9	1	0.8	0	0.1	0	0.2	0	0.1	0	0.1
PER	0.7	0	0	0	0	0	0.7	0	0	0	0.3	0.4
BOL	0.6	1	0.1	0.1	0.6	0	0	0.1	0.4	0	0	0.3
CHI	0.3	0.9	0.2	0.2	0.3	0.1	0.1	0.2	0.7	0.1	0.1	0.6
ECU	0.9	0.1	0.3	0.1	0.1	0.3	0.7	0.1	0.1	0.1	0.1	0.4
BRZ	0.2	0.8	0.9	0.2	0.1	0.2	0.1	0.8	0.1	0.2	0.1	0.7
URU	0.9	0.2	0.8	0.2	0.2	0.8	0.2	0.1	0.1	0.1	0.1	0.8
PAR	0.2	0.3	0.7	0.2	0.2	0.2	0.2	0.3	0.3	0.7	0.3	0.9
COL	0.2	0.4	0.4	0.2	0.2	0.2	0.2	0.4	0.4	0.4	0.6	1
VEN	0.9	0.7	0.6	0.6	0.4	0.3	0.3	0.1	0.1	0.1	0.1	0.3
Membership in row ≤ Membership in Y				0	0	1	0	0	0	1	1	

Table 4.7 displays the fuzzy-set membership scores of our ten hypothetical cases in the three conditions, the eight truth table rows, and the outcome stable democracy (Y). For each truth table row, we asses whether each case's membership in it is smaller than or equal to its membership in Y. If so, the respective row is a subset of the outcome, thus fulfills the criterion of a sufficient condition and therefore receives a score of 1. If, however, one or more case's membership in the row exceeds that in the outcome, then the respective row is not a perfect subset of Y and receives a score of 0. As the last row of Table 4.7 shows, three conjunctions – A~BC, ~A~BC, and ~A~B~C – are perfect subsets of Y. For all other truth table rows, one or more cases deviate from the subset pattern of sufficiency and these rows are therefore not considered as sufficient for Y.[6]

We now have all the relevant information at hand to represent a fuzzy-set data matrix in a standard crisp truth table format. For each row, we know which cases belong to it and whether it is a subset of the outcome. The truth table that results from our hypothetical fuzzy-set data is shown in Table 4.8.

Before we continue and explain how a truth table is analyzed using the tools of formal logic, several important points should be underlined. First, regardless of whether crisp or fuzzy sets are used, a truth table is at the core of QCA. Second, when representing fuzzy sets in a crisp truth table, the more fine-grained information contained in fuzzy sets is crucial and remains available at all times. In other words, the procedure that leads to a truth table like that in Table 4.8 does not involve any conversion of fuzzy sets into crisp sets. The information conveyed by fuzzy-set membership scores is used both when assigning cases to rows and when assessing whether a row is a subset of the outcome. Third, when producing a truth table based on fuzzy sets, the value (1 or 0) in the outcome column does not mean that all cases in that row have a membership of 1 or 0, respectively, in the outcome. Instead, the outcome column values express that the row can be considered a sufficient condition for the outcome. This is why in Table 4.8 we label the outcome column "Sufficient for Y." Fourth, when assessing the subset relation between a row and the outcome set, all cases are taken into account, not just those that are good instances of the particular row (i.e., those with a membership score above 0.5). The 0.5 qualitative anchor is thus crucial for attributing a case to a row but inconsequential when assessing the subset relation between two fuzzy sets.[7]

At-a-glance: what is a truth table? How to get from a data matrix to a truth table

Truth tables are an important tool in QCA. Although they look similar to **crisp-set** data matrices, they express a different type of information. While the single rows in data matrices correspond to actual cases (or units of observation), in truth tables, single rows denote logically possible **configurations** of **conditions**.

Three steps are needed in order to construct a truth table: First, all 2^k logically possible AND combinations of conditions are written down, with k being the number of conditions. Second, each case is assigned to the **truth table row** in which it has the highest membership. This is straightforward in **crisp-set QCA** because each case is a full member of one row and a full non-member of all the other rows. In **fuzzy-set QCA,** cases usually have partial membership in all rows but they can have a membership of higher than 0.5 in only one row. Cases are therefore attributed to this one row to which they fit best. (Exception: if one or more conditions are given a fuzzy value of 0.5, then the case will not have a membership value of greater than 0.5 in any ideal type.) Third, for each row the **outcome** value has to be defined. It is 1 for all rows that are a subset of, and thus **sufficient** for, the outcome and 0 otherwise.

These three steps yield a truth table that can be subjected to analysis, regardless of whether the underlying data consist of crisp or fuzzy sets.

[6] As mentioned, in Chapter 5 we will deal with the question of how much deviation one can or should allow for before dismissing a subset relation.

[7] In section 5.2, we qualify this statement and argue that researchers should pay attention to whether the cases that contradict the statement of sufficiency (or necessity) are located on different sides of the 0.5 qualitative anchor in the condition and the outcome, respectively. We will label these cases "logically contradictory cases."

Study material. Do not distribute.

Table 4.8 Truth table derived from hypothetical fuzzy-set data

	Conditions			Sufficient for	Cases with
Row	A	B	C	Y	membership ≤ 0.5 in row*
1	0	0	0	1	COL (0.6)
2	0	0	1	1	PAR (0.7)
3	0	1	0	0	CHI (0.7)
4	0	1	1	0	BRZ (0.8)
5	1	0	0	0	PER (0.7), ECU (0.7)
6	1	0	1	1	URU (0.8)
7	1	1	0	0	BOL (0.6)
8	1	1	1	0	AR (0.8), VEN (0.6)

* Numbers in parentheses = fuzzy-set membership of case in row

4.3 Analyzing truth tables

Truth tables can be created from both crisp-set data and from fuzzy-set data. The outcome column indicates whether the specific truth table row, or conjunction of conditions, is sufficient for the outcome of interest. If so, this is indicated by the value of 1 in the outcome column.[8] Hence, if we started our research asking which conditions are sufficient for our outcome of interest, the truth table provides a first answer: all rows that are linked to the outcome value of 1 are the sufficient conditions. This answer, however, is often not very informative and difficult to handle, simply because there might be many such rows in a truth table. Almost always, we would like to obtain a more succinct and parsimonious answer. For this, in QCA we apply the rules of Boolean algebra. The so-called Quine–McCluskey algorithm is used for logically minimizing the various sufficiency statements contained in a truth table (Klir *et al.* 1997: 61). It is important to point out that this form of truth table analysis is applicable only to the analysis of sufficiency. For the analysis of necessity, the bottom-up procedure presented in sections 3.2.1.2 and 3.2.2.2 has to be used. In fact, in section 9.1 we show that any inference about the presence or absence of necessary conditions based on the top-down logical minimization of truth tables is prone to produce flawed results. A truth table, thus, does not play an important role in the analysis of necessity. In the following, we present the steps involved in the Quine–McCluskey algorithm (see also Ragin 1987: ch. 6).

[8] Only in csQCA and only if there are no contradictory truth table rows (see Chapter 5), does the value of 1 in the outcome column indicate that *all* cases in that row are, in fact, members of the outcome. In all other scenarios – i.e., in fsQCA and/or when there are contradictory rows – a value of 1 in the outcome column of a truth table does not necessarily mean that all cases in that row are members of the outcome of interest.

4.3.1 Matching similar conjunctions

We return to the truth table already used in the section on crisp sets (4.2.1). Note that such a truth table could also be the result of converting a fuzzy-set data matrix into a truth table. Therefore, although we are now working with the example derived from the demonstration for crisp sets, truth table analysis is identical regardless of whether the underlying data consists of crisp or fuzzy sets.

The first step is to create a Boolean expression of all those truth table rows that are connected to the outcome to be explained. In our case, these are the rows with Y = 1 (rows 1, 2, 3, 4, and 6). Row 1 can be written as ~A~B~C, row 2 as ~A~BC, and so on. Conjunctions representing a truth table row are also called *primitive expressions*. The information contained in Table 4.9 can be expressed as follows:

row 1 + row 2 + row 3 + row 4 + row 6

~A~B~C + ~A~BC + ~AB~C + ~ABC + A~BC → Y.

Each of these five primitive expressions has been defined as a sufficient condition for Y in the process of creating the truth table. This formula is the most complex way in which we can express the information about sufficiency contained in the truth table. The task now consists in reformulating the same logical truth in a less complex manner.

This process is called *logical minimization*. It is guided by the following first principle of logical minimization: if two truth table rows, which are both linked to the outcome, differ in only one condition – with that condition being present in one row and absent in the other – then this condition can be considered logically redundant and irrelevant for producing the outcome in the presence of the remaining conditions involved in these rows. The logically redundant condition can be omitted, and the two rows can be merged into a simpler sufficient conjunction of conditions.

Let us apply this principle to our example. Row 1 (~A~B~C) and row 2 (~A~BC) are identical except for the value condition C takes on: it is absent in row 1 and present in row 2. Thus, this information can be summarized in the

Study material. Do not distribute.

Table 4.9 Example of hypothetical truth table

	Conditions			Outcome
Row	A	B	C	Y
1	0	0	0	1
2	0	0	1	1
3	0	1	0	1
4	0	1	1	1
5	1	0	0	0
6	1	0	1	1
7	1	1	0	0
8	1	1	1	0

See Table 4.2

logically identical expression ~A~B. In other words, we can write the information about sufficiency in Table 4.9 like this:

rows 1 and 2 + row 3 + row 4 + row 6
~A~B + ~AB~C + ~ABC + A~BC → Y.

With reference to our example, this means that the absence of a violent upheaval in the past combined with an ethnically non-homogenous society (~A~B) is a sufficient condition for a stable democracy (Y), regardless of whether a pluralistic party system is in place (C) or not (~C).

Let us apply the same logical minimization principle to the primitive expressions ~AB~C (row 3) and ~ABC (row 4). They differ only with regard to the value of condition C, which therefore can be dropped with the two rows rewritten as ~AB. Together with the previous minimization of rows 1 and 2, we can now write:

rows 1 and 2 + rows 3 and 4 + row 6
~A~B + ~AB +A~BC → Y.

The same principle of logical minimization, matching a pair of primitive expressions that differ in the value of only one condition, can be equally applied to any two conjunctions that lead to the same outcome. In our example, conjunctions ~A~B and ~AB differ only in the value of condition B, which can be dropped, and the two expressions can be simplified to ~A. This means that condition ~A is sufficient for Y regardless of the values conditions B or C take. Our simplified solution formula now looks like this:

rows 1 to 4 + row 6
~A + A~BC → Y.

This formula is logically equivalent to the most complex formula and to all intermediate formulas.

Since it is based on the same data as our example in section 3.1.1.2, we note a difference in the solution term. In section 3.1.1.2, the very same data resulted in the solution:

~A + ~BC → Y.

The difference consists of the role attributed to condition A when it is combined with ~BC. The question is whether the inclusion of condition A is required when our aim is to find the most parsimonious solution term for the information contained in Table 4.9. The answer is that it is not required. Why? Conjunction ~BC includes both primitive expressions A~BC (row 6) and ~A~BC (row 2), i.e., by saying that ~BC is sufficient for Y, we also say that both A~BC and ~A~BC are sufficient for Y. Since these two primitive expressions differ only in the value of A, condition A can be dropped. Notice that the process of logical minimization allows for using the same primitive expression for more than one logical minimization. In our example, the primitive expression ~A~BC in row 2 can be matched with both the primitive expression in row 1 (~A~B~C, leading to ~A~B) and the one in row 6 (A~BC, leading to ~BC). This simply means that this primitive expression of row 2 is covered by more than one *prime implicant*, an issue that we address in more detail in section 4.3.2. For the moment, we can confirm the solution term:

~A + ~BC → Y.

We reiterate that this formula is one of several ways of summarizing the information on sufficiency contained in Table 4.9. All of the different solution formulas that we have reported here, as well as the intermediate steps of the minimization process, (a) are logically equivalent; (b) express the same information contained in the truth table; (c) do not contradict each other, nor do they contradict the information contained in the truth table; and (d) are acceptable summaries of the empirical information at hand.

The principle that more than one solution term is an acceptable and logically correct representation of the data in the truth table is a general feature of QCA. The decision on which solution formula to choose as the basis for the substantive interpretation of the available information depends on many research-specific issues that have nothing to do with formal logic. There are several potential reasons that we might prefer the formula

Study material. Do not distribute.

~A + A~BC → Y

over

~A + ~BC → Y.

Imagine, for instance, that the literature on the emergence of stable democracies (Y) makes a strong point that a democracy cannot stabilize in the presence of violent upheavals (A). However, as solution term A~BC (i.e., row 6 in Table 4.9) demonstrates, there is empirical evidence that warrants a qualification of this claim: if combined with ~BC, A can indeed be a causally relevant INUS condition for Y. Following our hypothetical example, and contrary to the hypothetical claim in the literature, stable democracies occur in the presence of violent upheavals – but only when these countries additionally have an ethnically non-homogeneous population (~B) and a pluralistic party system (C).[9] While it is true that the formula

~A + ~BC → Y

also contains this information, the role of condition A remains less visible. The formula that includes the term A~BC is simply more helpful in connecting the empirical findings with pre-existing theoretical knowledge and expectations on this particular topic.

A related argument to this is that more complex solution formulas help to direct attention to hitherto unexplained cases. Imagine that the literature on the stability of democracy has thus far failed to find an explanation for why a certain country, which we will call X, is a stable democracy. Further assume that country X can best be described by conjunction A~BC. By preferring the solution term that explicitly includes this conjunction as a sufficient path towards Y, we are able to demonstrate why country X displays a stable democracy in a more straightforward manner than with the more parsimonious solution term.

4.3.2 Logically redundant prime implicants

The Quine–McCluskey algorithm consists of more than the elimination of single conditions from a pairwise matching of similar conjunctions. There are situations in which this procedure yields a solution formula that can be

[9] Caution: this should be read to mean that an upheaval *can* be violent in such circumstances but does not necessarily *have* to be. Remember that the component ~A~BC is also implicitly contained in the solution. In the scenario of a heterogeneous society with a pluralistic party system, this allows for the absence of a violent upheaval.

further minimized, but not by using the rule we have used so far. Another minimization principle is therefore needed (Ragin 1987: 95–98).

Logical equivalence can often be detected quite easily, such as in one example presented in section 4.3.1, where we show that ~A + ~BC → Y ≙ ~A + A~BC → Y. However, logical equivalence is not always so easy to detect. For this reason, we introduce a more general procedure for arriving at solution formulas that cannot be further minimized.

In order to understand this further step in the minimization procedure, we introduce the notion of prime implicants. Prime implicants can be defined as the end products of the logical minimization process through pairwise comparisons of conjunctions introduced in section 4.3.1. In other words, the solution term that we achieve through the pairwise comparison of conjunctions consists of prime implicants that are combined through logical OR. Under certain circumstances, though, it happens that one or more of those prime implicants are logically redundant. They can be dropped from the solution term in order to obtain the most parsimonious formula.

How can we identify logically redundant prime implicants? In order to answer this, we introduce some new hypothetical data. Suppose that the outcome to be explained is the presence of a consolidated democracy (C). As potential conditions, we choose whether a country is rich (R), is ethnically homogeneous (E), and has a parliamentary government (P). Suppose the empirical information contained in a truth table can be written using the following primitive expressions, aka truth table rows:

REP + RE~P + ~REP + ~R~EP → C.

Figure 4.3 shows the prime implicants obtained by applying the minimization strategy just introduced.

The complexity of the logical statement has been reduced from four paths (each composed of three individual conditions) to three paths (each composed of two individual conditions). These three new paths (RE, EP, and ~RP) are the prime implicants. They logically contain all the primitive expressions and cannot be further minimized with the minimization procedure we have described up to now.

It is nevertheless possible that this solution term contains logically redundant prime implicants. Therefore, we introduce a second rule for the minimization of solution formula: a prime implicant is logically redundant if all of the primitive expressions are covered without it being included in the solution formula. Hence, a solution formula without such a prime implicant

Study material. Do not distribute.

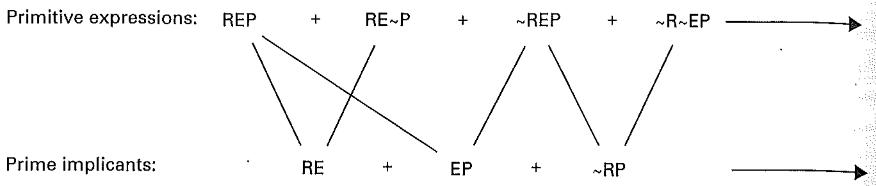

Figure 4.3 Logical minimization of primitive expressions to prime implicants

does not violate the truth value of the information contained in the truth table. Remember, the guiding principle of logical minimization is to express the same logical statement in a more parsimonious manner. The overarching requirement is that the truth value contained in the original truth table is not violated. The same logical statement can be expressed without the prime implicant in question and still adhere to this requirement.

We present the identification of logically redundant prime implicants in two different ways. First, we use the tool of a prime implicant chart, and then we use a Venn diagram. A prime implicant chart displays the primitive expressions in the columns and the prime implicants in the rows. Table 4.10 displays the prime implicant chart for our example on consolidated democracies (C). Crosses in the cells indicate which primitive expression is covered by which prime implicant(s). Each prime implicant covers at least one, but usually more, primitive expressions. In order to preserve the truth value contained in the truth table, each primitive expression must be covered by at least one prime implicant. Sometimes, there are primitive expressions that are covered by more than one prime implicant. It is here where the key to logically redundant prime implicant lies: a prime implicant is logically redundant if, and only if, all primitive expressions are covered even without it.

Take the situation displayed in Table 4.10. This table can be read as follows: the prime implicant RE covers the primitive expressions REP and RE~P, since RE is the result of the logical minimization of REP and RE~P by dropping condition P. This is indicated by the two crosses in the row for RE. The other two prime implicants, EP and ~RP, also both cover two prime implicants each, as can be seen by the two Xs in their rows.

For the truth value to be preserved, there must be at least one X per column in a prime implicant chart like that in Table 4.10. RE cannot be dropped, for it would leave primitive expression RE~P uncovered. It is therefore *not* logically redundant. ~RP cannot be dropped and is thus not logically redundant either, for it would leave primitive expression ~R~EP uncovered. Prime implicant EP, however, is logically redundant. It can be removed from the table without

Table 4.10 Prime implicant chart

	Primitive expressions/Truth table rows			
Prime implicants	REP	RE~P	~REP	~R~EP
RE	X	X		
~RP			X	X
EP	X		X	

any of the four primitive expressions remaining uncovered. REP is already covered by the prime implicant RE and ~REP by ~RP. Therefore, we can minimize our solution to:

RE + ~RP → C.

The notion of logically redundant prime implicants can also be explained by invoking the notion of intersecting sets displayed in a Venn diagram. Figure 4.4 displays the Venn diagram of our hypothetical example. In addition to the eight (2^3) different areas, which correspond to the logically possible combinations (aka truth table rows) between R, E, and P, the Venn diagram also indicates the location of the prime implicants (RE, RP, ~RP).

What Figure 4.4 demonstrates is this: the two prime implicants RE and ~RP jointly cover the entire area that is also covered by the third prime implicant EP (highlighted by the dark-gray area). Put differently, EP is logically implied, or is a subset of, the expression RE + ~RP. For this reason, EP is logically redundant and can be removed from the solution term. We say that it *can be removed*, because logically redundant prime implicants might well be of substantive interest. If so, they can and should be left in the solution formula. In these circumstances, such a formula is simply not the most parsimonious expression of the empirical information at hand.

Note that in the above example, there is only one logically redundant prime implicant (EP). This leaves no discretion to the researcher as to which prime implicant needs to be dropped in order to produce the most parsimonious solution. Very often in applied QCA, though, there are several logically redundant prime implicants and some of them are tied. Two logically redundant prime implicants are tied if either one or the other, but not both, can be dropped without violating the truth value of the solution term. This implies that in the presence of tied logically redundant prime implicants, there can be more than one most parsimonious solution term.

Study material. Do not distribute.

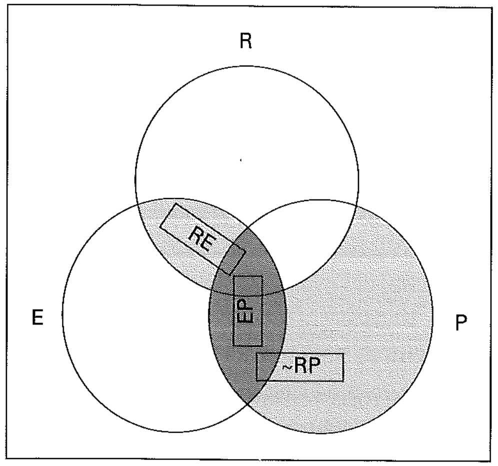

Figure 4.4 Venn diagram with logically redundant prime implicant

4.3.3 Issues related to the analysis of the non-occurrence of the outcome

Set relations are asymmetric (see section 3.3.3). One implication of this asymmetry is that the occurrence and the non-occurrence of an outcome of interest require separate analyses. All the analytic steps described so far that lead from a data matrix to a truth table and the logical minimization of the latter equally apply when the non-occurrence of the outcome is analyzed. Thus, continuing from our example displayed in Table 4.2 above, we now select ~Y (the set of countries with non-stable democracies) as the outcome of interest.

Starting with the analysis of necessity, we see that whenever ~Y is present, A is also present (see also section 3.2.1.2):

$A \leftarrow \sim Y.$

Having experienced a violent upheaval in the past turns out to be a necessary condition for having an unstable democratic system.

For the analysis of sufficiency, we apply the Quine–McCluskey algorithm based on all rows with ~Y = 1. This yields the following result:

$A{\sim}C + AB \rightarrow \sim Y.$

This can be rewritten by factoring out condition A (section 2.4.1) as:

$A\,(B + {\sim}C) \rightarrow \sim Y.$

Non-stable democracies occur in societies that have experienced a violent upheaval in the past and, at the same time, have an ethnically homogeneous

population and/or do not have a pluralistic party system. Again, this account of outcome ~Y is different from that for Y.

Three important remarks need to be made. One is of a general nature and the second and third stem from the particular simplicity of the example we have chosen. First, if, indeed, as we claim, the occurrence and the non-occurrence of a phenomenon, such as the stability of democracy and the non-stability, constitute two qualitatively different events that warrant separate explanations, then it often makes sense to resort to different theories and hypotheses to explain those outcomes. In other words, rather than just changing the outcome value from Y to ~Y in the *same* truth table, one might have to choose different conditions and thus construct an entirely *new* truth table. This directly follows from conceptual asymmetry, i.e., the fact that the negation of a concept often contains various qualitatively different notions. For instance, the set of non-democracies denotes military regimes, theocracies, and one-party regimes, to mention just a few. Likewise, for example, the set of non-married people comprises singles, widows, etc. In short, asymmetry might not only require different conditions for explaining Y and ~Y respectively. It also might require different conditions for the qualitatively different outcomes captured within ~Y.

The second and third caveats need to be made because our simple example produces two features in the solution term that usually do not hold when set-theoretic methods are applied to observational data. Discussing these features generates some general insights, though, and should help to avoid two mistakes often found in the applied QCA literature.

The second caveat is the following: in the analysis of necessity, we have identified condition A as necessary. At the same time, the analysis of sufficiency has revealed two paths, both of which involve condition A. It therefore seems that whenever a single condition is part of all sufficient paths, then this condition must be necessary for the outcome. Likewise, it might seem that that if no single condition appears in all sufficient paths, then there is no necessary condition. Both conclusions are likely to be wrong in applied QCA. They hold only if the sufficiency analysis is performed on a fully specified truth table, i.e., a truth table in which the outcome value for each of the 2^k logically possible combinations of conditions is either 1 or 0.[10] As we show in detail in sections 5.1, 6.1, and 6.2, this is hardly ever the case when formal logic meets noisy social science data. In applied QCA, truth tables

[10] In fsQCA, even then the conclusions about necessary conditions might be erroneous for reasons that we discuss in detail in section 9.1.

Study material. Do not distribute.

almost invariably contain rows that are contradictory or logical remainders. Whenever these types of rows are present, the sufficiency analysis of a truth table runs the risk of not correctly revealing the presence or absence of necessary conditions. In Chapter 9, we spell out the detailed circumstances under which false necessary conditions appear and when true necessary conditions disappear. For the time being, it suffices to state that it is always recommended that analyses of necessity and sufficiency be kept separate and that statements of necessity and sufficiency, respectively, be based only on analyses of necessity and sufficiency, respectively.

The third caveat has similar roots to the second. Due to the simplicity of our example, it provides an exceptional instance in which it would be possible to derive the sufficiency solution formula for ~Y based on the formula for Y, without performing a separate analysis. Making use of DeMorgan's law[11] we can convert the sufficiency solution term for Y

$$\sim A + \sim B * C \rightarrow Y$$

into

$$A * (B + \sim C) \rightarrow \sim Y.$$

This is identical to the formula derived through empirical analysis of outcome ~Y based on Table 4.2.

However, as mentioned, in social science research practice this procedure is problematic. It works properly only in a fully specified truth table, i.e., when there are no contradictions (section 5.1) or logical remainders (sections 6.1 and 6.2). Otherwise, the results produced by an application of DeMorgan's law imply claims about some truth table rows that either go unnoticed or are untenable, or both (Chapter 8 and section 9.1). Since fully specified truth tables are rare in practice, the meaningful use of the procedure described here and thus of DeMorgan's law is very limited.

For all these reasons, the standard of good practice (section 11.1) should be to always perform separate analyses of the occurrence and the non-occurrence of the outcome and to always analyze necessity and sufficiency in separate steps, not the least, since causal asymmetry also refers to the fact that substantial reasons might require us to use different causal factors for the explanation of the occurrence and the different types of non-occurrence of an outcome, respectively.

[11] In section 3.3.3, we have described how, if we have fully specified truth tables (as in the case under examination), the arrow in the statement of sufficiency can be replaced by an equals sign (=). Consequently, it is possible to negate (e.g., through the application of DeMorgan's law, see section 2.3) both sides of the equation without altering the truth value of the statement.

At-a-glance: analyzing truth tables

The **Quine–McCluskey algorithm** uses the simplification rules of Boolean expressions on the **truth table**. It starts out by listing all **configurations** of **conditions** for which **sufficiency** has been confirmed. Subsequently, the logical expression is minimized, with the help of the rules of Boolean algebra. Examining **prime implicants** often makes possible further simplifications that are not apparent at a first glance.

Factoring out those **INUS conditions** that appear in all sufficient paths does not show that a condition is **necessary** for the outcome. Because of this, necessary and sufficient conditions must be examined separately. It is recommended that analyses of necessary conditions be performed before analyses of sufficient conditions.

The **non-occurrence** of the outcome has to be analyzed separately. Only when there are neither configurations lacking any empirical cases nor **contradictory truth table rows** can DeMorgan's law be applied.

Note that the information contained in any given truth table can be expressed through different solution terms. The principles of **logical minimization** ensure that these formulas are logically equivalent and differ only in the degree of complexity. The decision about which of these solution terms to put at the center of the substantive interpretation needs to be guided by theoretical and substantive considerations.