Set-Theoretic Methods and Process Tracing in Multi-Method Designs: Principles of Case Selection after a QCA

Carsten Q. Schneider, Central European University (CEU), Budapest

&

Ingo Rohlfing, University of Cologne[1]

Draft version - Comments are welcome

Please do not quote without the authors' permission

February 2011

Abstract:

We propose rigorous principles for the integration of process tracing with set-theoretic methods, in general, and Qualitative Comparative Analysis (QCA), in particular. We show that set theory-based research produces multiple types of deviant cases. This is a direct consequence of these methods' reliance on set-relational causation in terms of necessary and sufficient conditions. Intentional case selection on the basis of QCA results needs to take into account that each of these types carries a different analytic meaning and therefore their in-depth study must serve different analytic purposes. Using simple 2x2 tables and x-y plots, we address four issues in detail. First, we identify the location of each type of onlier and outlier in a cross-case distribution one obtains based on a QCA solution formula. Second, using formal logic, we specify the reasons for deviance specific for each type of outlier. Third, we show how (not) to compare different pairs of onliers and outliers for within-case comparisons, providing new insights on the classic small N comparison literature. Fourth, we detail the implications of modifying the set-theoretic cross-case model in the light of process tracing evidence. We use data from published research using crisp-set QCA and fuzzy-set QCA studies for exemplifying these four points.

[1] Rotation principle. Both authors contributed equally to the paper.

Table of content

INTRODUCTION

Although multi-method research (MMR) is increasingly discussed in the methodological literature, there is a striking lack of debate on how to combine set-theoretic methods, most notably *Qualitative Comparative Analysis* (QCA),[2] and case studies. If case selection is discussed within the framework of set-theoretic methods, the focus usually is on how to infer a sub-set relation from a distribution of cases (Braumoeller and Goertz 2000; 2002; Clarke 2002; Dion 1998; Goertz 2008; Goertz and Levy 2007; Goertz and Starr 2003; Seawright 2002a). Alternatively, scholars are focusing on the delineation of cases that are then used for a QCA, that is, they discuss the importance of defining the scope conditions as an integral part of the set-theory based research process (Ragin 2000, chap. 2; 2006b; Rihoux and Lobe 2009). Both topics are important but differ from what we are interested in: Which cases can and should one select for process tracing *after* sub-set relations have already been established across a wider set of cases.[3]

It comes as a surprise that there is virtually no elaboration of how to integrate set-theoretic methods and case studies in MMR. Set-theory based approaches, in general, and QCA, in particular are designed for, and promoted as, cross-case method that are strongly connected with the thorough knowledge of cases (Ragin 2000; Rihoux and Lobe 2009). It is a widely agreed claim that the primary purpose of QCA and process tracing is to make sense of and learn more about the cases under study (George and Bennett 2005; Ragin 1987; 2006a).

One reason for this striking gap might stem from the belief that the standard case selection principles elaborated in the extant literature on regression-based case selection (Bäck and Dumont 2007; Lieberman 2005; Rohlfing 2008; Seawright and Gerring 2008) are directly applicable to set-theoretic

[2] Set theoretic approaches in the social sciences subsume a diverse set of approaches including, for example, typological theory (George and Bennett 2005, chap. 11) and Mill’s methods of difference, agreement, and the indirect method of difference (Mahoney 1999; Ragin 1987). Among them, QCA is the most formalized set-theoretic method that uses logical minimization procedures, allows for deviations from perfect subset relations, and incorporates fuzzy-set logic. As will become clear below, these features are at the core of our case selection principles.

[3] Case studies and small-n research, “n” referring to the number of cases, can be qualitative and quantitative (Gerring 2004). In the context of MMR, however, it is most useful to perform small-n qualitative process tracing as a supplement to the large-n cross-case technique.

methods as well. We aim to show that this belief is wrong because set-theoretic and regression-based case selection differs in crucial ways. The main reason for this difference rests on the diametrically opposed perspective on what kind of data pattern provides cues for a potential causal relationship. Most standard statistical techniques are geared towards detecting *symmetric correlations* between independent and the dependent variable. Set-theoretic methods, instead, are looking for *asymmetric set relations* between conditions and the outcome and interpret these relations in terms of necessity and sufficiency (Ragin 1987). These fundamentally different ontological premises have important consequences, such that one cannot directly transfer the insights from the literature on regression-based case selection to the realm of MMR including set-theoretic methods.

We aim to give the first rigorous exposition of how to perform set-theoretic MMR combining QCA and qualitative within-case analyses. This is a valuable enterprise for four reasons related to measurement and causal inference. First, the choice of outliers helps one to reassess the quality of the concepts of interest and the validity of the indicators used for the QCA (Adcock and Collier 2001; Coppedge 1999; Ragin 2000, 317). Second, QCA, just as any other method, may start with the correct set of conditions, but this is of course not necessarily the case. An iterative process between set-theoretic cross-case analyses, on one hand, and insights gathered via process tracing, on the other, promotes the identification of causally irrelevant conditions and conditions that one initially omitted from the analysis. While the back and forth between cross-case level and within-case level has been emphasized in relation with the specification of the population (Ragin 2000, 45) and concept formation and measurement in QCA (Ragin 2000, 317), there is no such discussion the interplay between cases and the QCA model *after* a QCA has been performed. Third, QCA, again like virtually all cross-case analyses, often produces results compatible with different causal processes related to competing explanations for the outcome. Process tracing can help in singling out one explanation for the cross-case pattern or at least narrow down the range of feasible theoretical accounts) (Campbell 1975). Fourth, process tracing following QCA is beneficial when the number of cases is too high for acquiring intimate knowledge of all of them (Ragin 2008). In our view, this is a situation that is more often reached than is usually acknowledged in the literature. We certainly agree that a QCA drawing on the deep knowledge of case is better than one that lacks it. However, one can run a QCA on large datasets with dozens, hundreds, or thousands of observations and this is done more and more often in fact (Cooper and Glaesser 2011; Glaesser and Cooper 2010; Grendstad 2007; Ragin 2008, chap. 11). Our case selection principles are a guide for all those that are confronted with too many cases to be all studied in-depth. Even if sufficiently deep knowledge of all cases exists prior to the QCA, careful post-QCA case selection can still improve causal inference because the set-theoretic results give one a fresh look at the cases (Ragin 2006a).

In order to meet these four goals via the systematic choice of cases, it is mandatory to pay close attention to the core features of set-relational based research (Ragin 1981; Ragin 2006b; Schneider and Wagemann 2010). On a first level, one must take into account whether one is choosing cases after QCA that aimed at necessary or sufficient conditions for an outcome. Second, we will show that for statements of necessity and sufficiency, the case selection procedure hinges on whether one is interested in improving the *consistency* or *coverage* of a set relation (Goertz 2006; Ragin 2006a), two separate measures of fit in QCA. One consequence of the opportunity to focus on different set relations and measures of fit is that one can choose onliers and *different types of outliers*. This contrasts with regression-based MMR where one only has one type of onlier and outlier (Lieberman 2005). In addition to the explication of case selection strategies for various types of outliers in QCA-based MMR, we explain that the nature of each type of case provides strong clues as to what the reasons for its deviance are. Deriving these clues in advance of the case study is an important asset because it allows one to narrow the focus during the within-case analysis. Finally and building on the previous point, we close the circle and detail the implications of different process tracing insights, such as, for instance, an omitted condition or weak measurement, on the cross-case analysis.

In section two of our paper, we first make some general notes about QCA-based MMR and clarify what elements of QCA are important to know for the elaboration of our case selection principles. We start with crisp-set QCA (csQCA) in section three and then extend the principles to fuzzy-set QCA (fsQCA) in section four. In both sections, we first transfer the well-known notions of typical and deviant cases (Eckstein 1975; Levy 2008; Lijphart 1971; Seawright and Gerring 2008) to the realm of QCA-based MMR. We detail where to find which type of typical and deviant case in csQCA and fsQCA, thereby also discussing case selection for single-case studies. Afterwards, we extend the perspective to case selection for *comparative* within-case studies. We show that comparative case selection is not as straightforward as one might think at first sight: some comparisons between typical and deviant cases are viable while others are logically flawed and should thus be avoided. We then exemplify our case selection principles by using data from published QCA research and by showing which cases should be selected on ground of the respective set-relational results. The final section concludes.

QCA AND CASE STUDIES IN MMR: PRELIMINARY NOTES

The distinctive feature of MMR as opposed to single-method designs is the systematic integration of a cross-case and a within-case analysis (Bennett 2002; Lieberman 2005). The sequence in which the cross-case and within-case analyses are performed determines how these two perspectives speak to each other (Rohlfing 2008). If the case study comes first, process tracing serves, among other things, more exploratory purposes with the aim to discern conditions to include in the cross-case model, be this a regression-type model or a QCA (Rihoux and Lobe 2009). The cross-case model then tests the adequacy of the model previously generated in the small-n analysis. Alternatively, if the cross-case analysis goes first and cases are selected for process tracing on the basis of the cross-case results, process tracing results are then used to reconsider the cross-case model.

Here, we exclusively deal with the second variant, that is, the scenario in which one starts with QCA. We particularly focus on three aspects of the multi-method design: the choice of suitable onliers and outliers for process tracing; the clues that the cross-case results provide for the within-case analysis; and the effects of altering the cross-case model in light of process tracing insights. These three features deserve particular attention, for they are unique to designs that integrate QCA and case studies as opposed to both performing one method alone or by combining case studies with non-set-theoretic methods. The ways in which QCA and process tracing should be implemented in order to meet their respective standards has been aptly described elsewhere and need not be replicated here.[4]

The fact that there are excellent treatments of QCA also explains why we do not elaborate the elements of QCA that are central for the discussion of case selection. We presume that the reader is familiar with the following issues: the different manifestations of set relations (necessity, sufficiency, equifinality, conjunctural causation (Ragin 1987), INUS causation (Mackie 1965), and SUIN causation (Mahoney, Kimball and Koivu 2009); crisp sets and fuzzy sets (Ragin 2000); standard Boolean notation for analyses of necessity and sufficiency (Ragin 1987); the minimum-scoring rule for determining the membership of a case in a configuration, and the maximum-scoring rule for calculating the membership in an equifinal solution (Ragin 2008); truth tables and their logical minimization (Ragin 1987); consistency and coverage as goodness-of-fit measures (Ragin 2006a); and 2x2 tables and x-y plots as means to present QCA results (Schneider and Grofman 2006). For readers who want to refresh their knowledge about these issues, we prepared an online supplement containing brief discussions of each individual element.[5]

CRISP-SET QCA AND CASE SELECTION

Necessity

Typical cases

A statement of necessity implies that the condition is a superset of the outcome. The implications of this definition of necessity for case selection are exemplified with a 2x2 table that plots each cases' membership in the QCA solution against its membership in the outcome set (Table 1).[6] Typical cases regarding necessity are located in cell 2. Cases in this cell are in line with the set-theoretic statement

[4] On standards of QCA, see for example Ragin (1987; 2000; 2008), Rihoux and Ragin (2008), and Schneider and Wagemann (2010). On process tracing, see for example Hall (2008) and George and Bennett (2005).

[5] The appendix can be found here: xxx.

[6] We follow the definition of necessity widely applied in the literature (Mahoney, Kimball and Koivu 2009) which states that if Y is present, X must be present as well.

(consistency) and show both the condition and the outcome (coverage). Hence, in contrast to outliers (see below), there is no need to distinguish between different types of typical cases.

Table 1: Types of cases in csQCA for necessity

		crisp set membership in QCA solution: 0	crisp set membership in QCA solution: 1
outcome	1	Deviant cases consistency (1)	Typical cases (2)
	0	Irrelevant cases (4)	Irrelevant cases (3)

Deviant cases consistency

Deviant cases regarding consistency are located in cell 1. They display the outcome in the absence of the purported necessary condition, something that is not in line with a perfectly consistent statement of necessity. One reason for the occurrence of the outlier consistency cases can be the omission of a functional equivalent necessary condition of the same higher-order construct (Schneider 2008). Such a measurement strategy is warranted if a condition lacks content validity for all cases under scrutiny and context-specific conditions are to be preferred (Adcock and Collier 2001). Adding another condition through the logical OR operator could turn the deviant cases into typical cases for the new statement of necessity. This would occur if all deviant cases are members of the set of the new condition, because the logical OR requires to assign to cases the maximum set membership across all conditions combined by logical OR (Ragin 1987).

Deviant cases coverage

The second measure of fit in QCA is coverage. Just like consistency, the relevance of a necessary condition is determined by the distribution of cases in different cells of the 2x2 tables in **TABLE 1**. There are two crucial differences for the meaning of outliers, though, which lead us to argue that no meaningful outliers with regard to coverage exist.

A necessary condition is, by definition, a superset of the outcome and thus covers the outcome by definition. Therefore, we deem it potentially misleading to apply the notion of coverage to necessary conditions and prefer to speak of the *relevance* of necessary conditions (Goertz 2006; Ragin 2006a, 90-100) and its *trivialness* (Goertz 2006???).

First, the relevance of a necessary condition decreases as the number of cases in cell 3 of **TABLE 1** increases in relation to the number of cases in cells 2 and 3. However, studying cases in cell 3 cannot contribute to our knowledge of either the modus operandi of necessary condition X, the occurrence of outcome Y, or how to improve the empirical fit of the statement of necessity. In fact, the cases in the lower right corner do not qualify as deviant because we do not expect the outcome to be in place when the necessary condition is present (in this instance, X would be a subset of Y).

Second, a condition is trivial if it is present in all or most of the cases of interest. Because of their omnipresence, such conditions are supersets of any outcome almost by default. For instance, in a study on the conditions for why some EU member states do not violate the Maastricht criteria, it is trivial to state that 'absence of civil war' is a necessary condition for membership in the set of countries that do not violate the Maastricht criteria. Among EU member states, condition 'absence of civil war' is a constant and by virtue of this a superset and a trivial necessary condition of the outcome.[7] Trivialness, thus, is indicated by the *absence* of cases in cell 4 of **TABLE 1**. In other words, if something is wrong with a claim of necessity in terms of trivialness, then there is no case cell 4. For obvious reasons, no within-case studies can be performed on cases that do not exist.

[7] Note that Goertz' (2006) notion of relevance, is substantively (though not mathematically) identical with Ragin's (2006a) notion of coverage. Goertz (2006) suggests a second formula with which to identify conditions that are trivial due to them being close to constant conditions.

Because of this, we conclude that there is no deviant case with regard to the coverage/relevance/trivialness of a necessary condition that lends itself to insightful process tracing. We add, however, that this argument only applies to standalone analyses. For comparative within-case studies, we argue below that cases in cell 3 are useful when being matched with cases in cells 2 or 4.

Sufficiency

Table 2 presents the classification of cases as *typical cases*, *deviant cases for consistency*, and *deviant cases for coverage* in a csQCA on the sufficient condition(s) for an outcome. Note that in the realm of sufficiency, "X" can stand for a single condition or be a placeholder for either a conjunctions of conditions, called a *path*, for multiple paths, pr for the entire solution term one derives from a QCA. As **TABLE 2** highlights, there are two types of deviant cases in QCA-based MMR. This feature sets it apart from regression-based MMR where one can only specify one type of outlier. The cases in cell 4 are not relevant in an analysis of sufficiency because they display neither the cause nor the outcome (Braumoeller and Goertz 2002).[8] In the following, we discuss each of the three types of cases on its own terms.

Table 2: Types of cases in csQCA of sufficiency

membership in outcome (Y)		0	1
	1	Deviant cases coverage (1)	Typical cases (2)
	0	Irrelevant cases (4)	Deviant cases consistency (3)
		membership in QCA solution (X): 0	membership in QCA solution (X): 1

[8] Seawright (2002a; b) argues that one should draw on all four cells of the a 2x2 table in analyses of sufficiency (and necessity). We follow Braumoeller and Goertz's (2002) criticism that Seawright's approach rests on distributional assumptions that are not in line with the distinctive nature of set relations.

Typical cases

A condition X is sufficient for an outcome Y if the set of X is a subset of the set of Y. This means that if X is present, Y should be observed as well. These cases are located in cell 2 of the 2x2 matrix. They are typical cases because they are empirical manifestations of the proclaimed set relation of interest. Any case from cell 2 is suitable for inductive process tracing aiming to discern how exactly the cause is related to the outcome. Similarly, the cases in zone 2 should be selected if one strives to test deductively derived hypothesis on the casual process connecting X to Y (cf. George and Bennett 2005, chaps. 10-11). Cases in cell 1 are formally consistent with a pattern of sufficiency but do not qualify as typical cases. Nothing can be learned about how this path brings about the outcome by studying these cases, simply because they are not members of the sufficient path.

When selecting typical cases (and deviant cases regarding consistency, see below), it is important to do justice to the principle of *diversity*. Diversity manifests itself in equifinal QCA solutions (Ragin 1987, 30), that is, it expresses *causal heterogeneity* by pointing to multiple sufficient paths that all lead to the same outcome. Many, perhaps even most social phenomena are characterized by equifinality, For example, high levels of welfare state spending can be due to an open economy *or* a left-wing government or both (Obinger, Leibfried, Bogedan, Gindulus, Moser and Starke 2010, chap. 1); two countries are at peace with each other because they are both democratic or they maintain close economic ties (Schneider and Gleditsch 2010); governments break down because of external shocks or internal disputes (Laver 2003); and so on. In QCA-based case selection, diversity must be taken into account by selecting one case for each path contained in the solution.[9] For the example of high welfare state spending, this implies that one should select two typical cases: one with high membership in the condition ‘open economy’ and outcome ‘high welfare spending’, and another case with high membership in condition ‘left-wing government’ and outcome ‘high welfare spending’. We summarize the implication of equifinality for case selection in the *principle of diverse case selection*:

Principle of diverse case selection: Choose at least one typical case for each path of the equifinal solution term.

[9] Goertz (2008) calls this “choosing cases diversely”.

This principle has two implications, one related to the breadth of the small-n part in MMR and one to generalization. First, if one aims to get a comprehensive understanding of the outcome of interest, a core justification for QCA as a case-based method (Berg-Schlosser , De Meur, Ragin and Rihoux 2008), one has to select and examine at least as many onliers as there are paths in the solution. Second, the principle of diverse case selection is closely tied to the notion of contingent generalization, a concept discussed in the case study literature (Bunce 2000; George and Bennett 2005, 111-113). In QCA-based MMR, contingent generalization specifically means that the empirical insights one derives through process tracing can and should only be generalized to cases that are members of the same path. This implication follows directly from the idea of equifinality, for it means that there are qualitatively different ways in which the outcome can come about. Consequently, one should refrain from extending the insights derived from process tracing in one onlier to onliers that are members of another path.

In set-theoretic research, it is common that one case is an empirical instance of multiple paths, that is, it is covered by more than one path. As regards the previous example, there are countries with high welfare state spending (outcome) and both a 'high degree of economic openness' and an 'incumbent left-wing government'. We refer to such cases as *jointly covered* and the phenomenon as *joint coverage*.[10] In contrast, we call cases *uniquely covered* if they are a member of only one path of the solution. Cases with simultaneous presence of multiple paths are inadequate for within-case analysis because one should focus on a single path and how it leads to the outcome (George and Bennett 2005, 252). Joint coverage renders this impossible and makes the within-case analysis unnecessarily complicated (Gerring 2007).[11] For this reason, uniquely covered cases create the best context for process tracing and should be the target of case selection. We therefore propose the *principle of unique coverage* for handling this problem and argue that one should always choose cases that are covered by only one path:

[10] Jointly covered cases are related to the difference between the raw coverage and unique coverage of a path. As a rule of thumb, if for a given sufficient path the raw coverage is larger than the unique coverage, then it overlaps with another path identified in the same solution term and some of the cases under analysis must be jointly covered.

[11] This principle ensures that one chooses cases that Gerring (2007) labels as *pathway cases*.

Principle of unique coverage: Choose cases that are covered by just one path.

Deviant cases consistency

If X is deterministically sufficient for Y, one does not observe any cases in zone cell 3 of **TABLE 2**. In practice, however, it is common that the distribution of cases includes some cases that contradict the statement of sufficiency and are located in the lower-right corner of **TABLE 2**. These cases are deviant cases with respect to consistency because they are inconsistent with a perfect set-relation of sufficiency. The puzzling feature of these cases is that they are covered by a specific sufficient path, but do not display the outcome. The two principles of diverse selection and unique coverage extend to deviant cases regarding consistency. One should therefore select at least one uniquely covered outlier for each path of the QCA solution for which one observes deviant cases in zone 3.[12]

Because there are various potential reasons for deviance, outlier consistency cases are worth to examine for multiple purposes. Two of these reasons are related to the conditions that one feeds into the QCA and, consequently, the solution one obtains from a QCA. First, the path that covers the deviant case may be *under-fitted* and lacks a relevant condition. This clue directly follows from the location of the case in cell 3. These cases are deviant because their membership in X exceeds that in Y. One obvious way of lowering its membership in the path is to make the membership in X more demanding by adding a condition to that conjunction. The minimum-scoring rule for the assignment of set memberships to cases comes into play here because the membership of a case in a path is equal to the minimum score of all its constitutive conditions. This follows common sense. More cases will be members of the set of 'open economies' than in the set of 'open economies AND left wing government'. Consequently, exploratory process tracing in an outlier consistency should search for a condition that seems causally relevant in conjunction with the other conditions constituting the path *and* in which the

[12] Readers familiar with QCA will notice that cases in cell 2 and 3 imply a contradictory row in the truth table: one and the same the configuration X leads to different membership scores in Y. Deviant cases consistency are therefore the same as contradictory cases in standard QCA terminology and the joint presence of typical cases and deviant cases consistency indicates the presence of contradictory truth table rows. Below, when discussing the empirical example, it will become clear why the designation of cases as typical and deviant and the implementation of our case selection principles yields added value compared to the usual handling of contradictions in QCA.

deviant case is not a member. When one discerns an omitted condition that fulfills both criteria and adds it to the path, thus creating an expanded path, the outlier ceases to be a member of this new sufficient path and thus converts into an irrelevant case in cell 4 of **TABLE 1**.

The second model-related reason for outliers consistency can be the *over-fitting* of the solution. The membership of a case in the entire QCA solution follows the maximum-scoring rule. Each, case's membership in the overall solution term is calculated by taking the maximum set membership across all paths that form the entire solution. If it turns out in process tracing that there is no causal link between a specific path and the outcome and there are no overwhelming theoretical reasons to suggest otherwise, it may be justified to drop the path from the solution. Although we deem it unlikely that an entire path is spurious, it is a logically possible reason for outliers with respect to consistency and should be taken into account in QCA-based MMR. If the outlier consistency has no membership in any other path constituting the QCA solution - which should be the case due to our case selection principle of unique coverage – then it moves from cell 3 to cell 4 and turns from a deviant case into irrelevant one. In total, then, process tracing that serves to search for model-related sources of deviance should take a two-fold perspective: one should determine whether the path covering an outlier is spurious (model over-fit) and whether a condition has been omitted from the path (model under-fit).

Beyond model over- or under-fit, a case can be deviant for reasons related to the mis-calibration of either the conditions or the outcome or both. Process tracing can also be helpful for gathering information on such issues of concept formation and measurement. If within-case analysis yields credible evidence and arguments for a change of the membership function of one ore more condition, then hitherto deviant cases consistency might also turn into irrelevant cases by shifting from cell 3 to cell 4. Similarly, if good arguments and evidence exist to lower the threshold for being a member of the outcome set, then the deviant case might shift from cell 3 to cell 2 and becomes a typical case.

Deviant cases coverage

Independently of whether there are outliers for consistency or not, one may observe some cases in the upper-left corner of the 2x2 matrix. We argue that these cases are deviant regarding coverage. While this type of deviant cases does not provide counter-evidence against the set-theoretic statement of sufficiency, they nevertheless indicate some empirical inadequacy of the QCA solution because it cannot tell us anything about why the outcome occurs in these cases. From a substantive point of view, they are puzzling, and thus outliers, because QCA aims at being case oriented and, by virtue of this, tries to account for the occurrence of the outcome in all cases in which it occurs. Cases in cell 1 are left unexplained, or uncovered. We therefore label cases in the top-left cell of the 2x2 table as *non-covered cases*.

The model-related reason for the existence of outliers for coverage is the under-fitting of the QCA solution. There is a path missing that, if identified, turns the outlier for coverage into a typical case of the new solution. The within-case study of this type of cases should therefore aim at identifying this missing path. The reason for deviance cannot be the over-fitting of the path for which these cases are outliers as regards coverage and dropping one or more condition from the path has two pitfalls. First, it contradicts our suggestion made for the study of deviant cases consistency where we suggest adding a condition to potentially the same path. Second, if a condition could be dropped and the resulting path would still be consistent enough, then the cross-case QCA would have already identified it.

Our suggestion to find a new path capitalizes on the formal logical necessity that each case in a QCA is a member of one, and only one, *truth table row* (Ragin 1987, 87-89). The location of an outlier for coverage in the upper left cell of **TABLE 2** only tells us that a case it is *not* a member of any row that is implied by the solution.[13] For meaningful case selection, however, we need to know the configuration that the outlier for coverage *is* an instance of. Consequently, we recommend to go back to the truth table and determine the row into which the outlier falls. Given the conditions chosen by the researcher at the outset of the QCA (cf. Amenta and Poulsen 1994; Berg-Schlosser and De Meur 2008), this is the

[13] Presume you run a QCA that includes the conditions A, B, and C and that produces solution AB. We then know that an outlier for coverage cannot be located in the truth table rows ABC and ABc because these rows are implied by path AB.

conjunction that best describes the deviant case. In short, for outliers coverage, we propose the *truth table principle* for the handling of this type of deviant case.

Truth table principle: Do not select deviant cases as regards coverage on the basis of the QCA solution. Instead, determine the truth table row to which the outlier coverage case belongs.

When applying the truth table principle, the principle of diverse case selection is still in force. If two or more deviant cases populate the same row, one can pick any case from them and generalize the insights derived from this case to all other cases that are described by the same configuration. Deviant cases coverage that fall into different truth table rows, however, need to be analyzed separately. In both scenarios, the aim of exploratory process tracing in an outlier for coverage should be the identification of condition that is missing from the truth table row that best describes the case under study.

The rationale is the following: the outlier coverage displays the outcome of interest, whereas a sufficiently high number of cases described by the same configuration lack the outcome, else it would pass the threshold of consistency and thus be included in the QCA solution term. The contradiction between the outlier coverage case, on one side, and the other cases in the same truth table row, on the other side, can be resolved by adding a condition to the original configuration. If the outlier coverage and the other cases have opposite set memberships in the added condition, then they fall into two different rows in the new, expanded truth table and the contradiction is resolved (Ragin 1987, 113-118).

At first sight, one may believe that the number of case studies required by the truth table principle will be prohibitively high because each outlier coverage case may fall into a different truth table row. However, in practice the actual number of different outlier coverage cases is likely to be (much) smaller. First, the number of logically possible outlier coverage cases is likely to be drastically reduced by the fact that they can only fall into those truth table rows that are not implied by (parts of) the QCA solution. For instance, take a truth table consisting of conditions A, B, C, and D and the QCA solution $A*B \rightarrow Y$. Out of the $2^4 = 16$ rows, outlier coverage cases can only occur in 12. The remaining four all contain conjunction A*B and are thus implied by the QCA solution. Second, the number of outlier coverage cases is further limited by the omnipresent phenomenon of limited diversity, which means that not all logically possible outlier coverage cases will actually materialize in the data (Ragin 1987, 106-113).

Contradictions and types of cases

The previous discussion of the truth table principle hints at a more general relationship between typical and deviant cases, on the one hand, and contradictory truth table rows in QCA, on the other. We discuss this relationship in some detail here in order to show that post-QCA case selection for process tracing differs in important ways from the pre-QCA analysis of contradictory rows and yields added value. In QCA, one is confronted with a contradictory truth table row when cases are described by the same configuration but have different outcomes in place. Although one can deal with contradictory rows in the process of logically minimizing a truth table by including rows that are above a given threshold of consistency and exclude all others (Ragin 2008), QCA as a case-based method calls for a more case-centered handling of this problem. The more case-oriented responses to contradictory truth table rows are to reconsider the calibration of the conditions and the outcome, to search for measurement error and incorrectly formed concepts, and to check whether the introduction of a hitherto omitted condition would eliminate the contradiction (Ragin 1987, 113-118). For all these purposes, process tracing is a suitable tool in combination with conceptual and theoretical reasoning (Rihoux and Lobe 2009).

We of course agree that contradictions are best approached prior to the QCA with the means just described. However, this may not always be possible and the only instrument that is left is the imposition of a consistency threshold. If this route is taken, it is certain that one will be confronted with deviant cases in the post-QCA stage of the empirical analysis. Table 3 clarifies how typical and deviant cases are related to contradictory truth tables and is now elaborated step by step.

Table 3: Relation between contradictions and types of cases

		Membership of case in outcome =1	
		No	Yes
Truth table row included in logical minimization (coded Y=1)	Yes	Deviant case consistency (contradictory case)	Typical case
	No	Deviant case coverage (contradictory case)	Irrelevant

Typical cases fall into a truth table row that is part of the minimization process and have a set membership in the outcome. Deviant cases in relation to consistency are also covered by a row that is included in the minimization, but they are not members of the outcome. These cases are inconsistent with the inference that the configuration in question is sufficient for the outcome. The reason they are covered by the QCA solution is because they are members of a truth table in which the majority of cases do display the outcome. Deviant cases as regards coverage fall into a row that is not part of the minimization but that are members of the outcome set. In other words, the outliers display the outcome that we are interested in, but most of the cases in the same row lack the outcome. As a consequence of this, the truth table row is contradictory and excluded from the actual QCA.

Altogether, this discussion shows that the label of a contradictory truth table row conflates two very different constellations that need to be disentangled when engaging in meaningful cases selection after a QCA. The qualitative differences between outliers for consistency and coverage tends to go unnoticed if one just treats them under the heading' contradictory truth table'. As explained, these two variants of outliers should be approached with a different thrust in process tracing. This underscores the added value of looking at cases in terms of deviant cases for consistency and coverage rather than trying to resolve this issue by simple reference to a less than perfect consistency and coverage score.

Admittedly, our suggestion for the handling of outliers coverage is not different from the strategy to eliminate contradictory truth table rows before a QCA. However, we contend that there is benefit to designating them as outliers coverage. Assigning cases the status as outliers for coverage after a QCA provides important clues that one may not obtain in the pre-QCA stage. More specifically, we argue below that outlier coverage can be fruitfully matched with other cases in a comparative within-case analysis. Obviously, such a matching is only after a QCA.

A similar point applies to the benefits of treating cases as outliers consistency. If handled as cases belonging to a contradictory row that is included in the minimization procedure, we approach all the contradictory rows one after another in an attempt to get rid of the contradiction. Again, one may improve the ground for process tracing by taking a broader picture, which can be clarified by a hypothetical example. Suppose we have one truth table row *A*B* and one row representing the configuration *A*b*. Both rows are part of the minimization process, but for each row there is one or more case that does not display the outcome of interest (i.e., both rows are contradictory). Minimization of the truth table then leads to *A* as one sufficient path. The interesting implication now is that the cases that gave rise to two contradictory rows *prior* to the QCA are now outliers for consistency with regard to the same path. Constellations like this can only be discerned after having performed the minimization and may give a new look at why the cases deviate from the typical cases as regards the same path.

FUZZY-SET QCA AND CASE SELECTION

All of our selection procedures that we detailed so far equally apply when using fuzzy-set QCA (fsQCA). This directly follows from the fact that csQCA is a special case of fsQCA, for the latter not allows cases to have full (non-)membership but also partial membership in sets (Ragin 2000, chap. 6). This more fine-grained measurement approach is also an additional advantage for a meaningful post-QCA case selection. In addition to choosing cases based on their *differences in kind* with fuzzy-sets one can also take into account *differences in degree* among cases that are similar in kind.

So far, all our case selection principles are based on differences and similarities in kind. With fuzzy-sets, two cases are qualitatively identical if their fuzzy-set membership score falls on the same side of the *qualitative anchor* of 0.5 fuzzy-set membership. Accordingly, they are qualitatively different if their membership scores fall on different sides of this anchor. The qualitative anchor at 0.5 therefore is the benchmark establishing differences in kind between cases. With crisp-sets, only qualitative differences in kind are possible, whereas with fuzzy-sets also differences in degree between cases are reflected. This additional information can be used for informed case selection. One implication of more fine-grained set membership scores is that one can distinguish between more or less (and also the most) typical and more or less (and also the most) deviant cases.

This beneficial feature of fuzzy sets can be exemplified with a so called *x-y plot.* It is the functional equivalent to a 2x2 table in csQCA. It plots the fuzzy-set membership of a case in a condition or path (X) against its fuzzy-set membership in the outcome (Y). Following the logic of fuzzy-set relations, a path X is a perfect subset of outcome Y and fully consistent with the statement of sufficiency if all cases are located above the main diagonal. Correspondingly, a condition is fully consistent with the statement of necessity if all cases are below the main diagonal (Ragin 2000, chaps. 8-9). Since the presence of set-relational pattern depends on whether the membership in X is higher (necessity) or lower (sufficiency) than the membership in Y, the main diagonal is a useful aid in x-y plots (Schneider and Grofman 2006).

For case selection purposes, however, the main diagonal alone falls short in providing systematic guidance because it does not offer information about differences in kind between cases. In order to rectify this shortcoming of standard xy plots, and as a means to highlight the important role of differences in kind for purposeful case selection, we superimpose a 2x2 matrix on the x-y plot by adding a dashed horizontal line and vertical line that run through the fuzzy-set membership scores of .5 for X and Y, respectively. As Figure 1 highlights, the result is an enhanced x-y plot with six zones (or areas or cells) into which cases can fall. We now elaborate the role of the zones for case selection and particularly focus on the differences between the choice of cases in csQCA and fsQCA. This implies that we do not discuss all of our principles and insights that are identical in csQCA and fsQCA: these are the principle of diverse case selection, the truth table principle, the potential reasons for the deviance of outliers, and the effects of remedies on the QCA solution.

Figure 1: Enhanced x-y-plot and zones of case selection in fsQCA

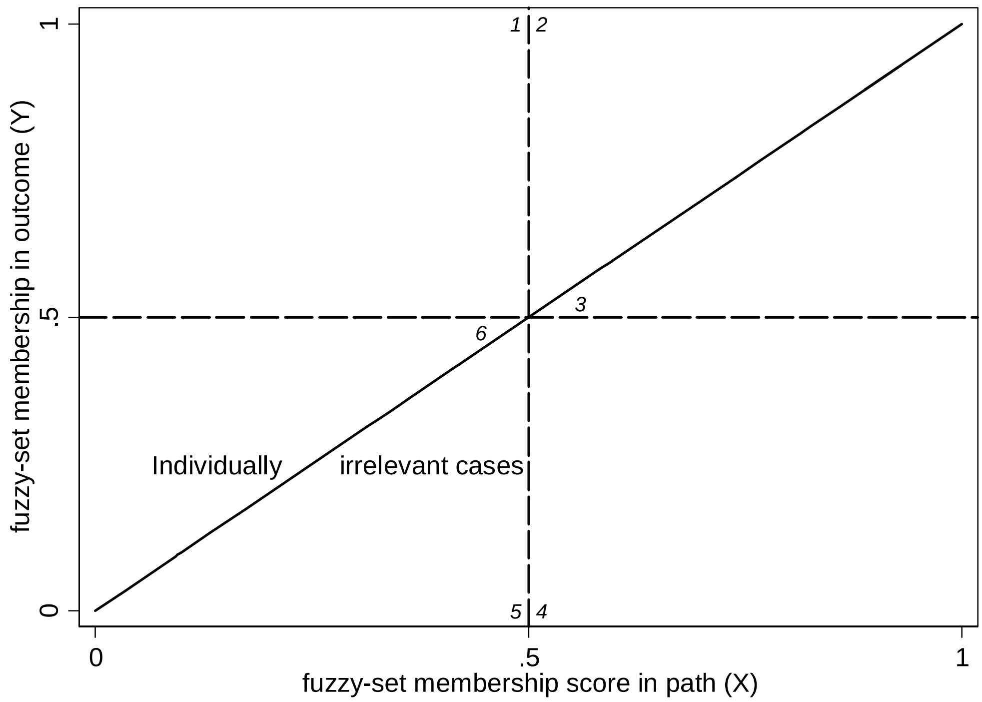

Necessity

Typical cases

Typical cases with regard to necessity are located in zone 3 of Figure 1. The cases are more in than out of both set X and set Y, which is the requirement that typical cases for necessity must meet. Second, among these qualitatively identical cases, their membership in X exceeds the membership in Y. Without this second requirement, one would also count cases in zone 2 as typical, which would be wrong for they are in discord with a set-relational pattern of necessity. This is an important difference to csQCA where all cases in the upper-right quadrant count as typical cases.

A further difference to csQCA-based case selection is that in fsQCA, it possible to determine the most typical case among all cases in area 3. The *ideal-typical onlier* has a membership of 1 in X and Y because the usefulness of a typical case for process tracing increases with the strength of its set membership in X and Y. The rationale is that cases with high membership scores simply are better empirical manifestations of both the path and the outcome, thereby making it easier with process tracing to unravel how the path is connected to the outcome. We capture this salient recommendation in the *principle of high fuzzy-set membership* that only pertains to the choice of onliers in fsQCA:

Principle of high fuzzy-set membership: As a typical case choose the one with the maximum set membership scores in X and Y.

Ideally, then, a typical case has a membership score of 1 in both X and Y. In practice, it is likely that its membership is not the maximum in X and Y. As a consequence of this, the most typical case is the one that lies closest to the upper-right corner of the x-y plot.[14]

Deviant cases consistency

Deviant cases with regard to the consistency are captured by cell 1 of **FIGURE 1**, which corresponds to the upper-left quadrant of a 2x2 table in csQCA. Notice that if one only emphasizes differences in degree among cases, also cases in zone 2 would be deemed good choices for process tracing of deviant cases consistency. They display membership scores in Y that are higher than in X and thus fall above the main diagonal. This, however, ignores that these cases are more in than out of both the set of X and of Y. They are, thus, qualitatively identical to typical cases for necessity. The only puzzling thing about cases in zone 2 is that the membership in Y is larger than in X instead of the other way round. Because of this feature of outliers in area 2, we propose to call them *outliers in degree*. They are qualitatively identical to typical cases and only differ from them in that they are located on the wrong side of the diagonal. Compared to the cases in zone 2, cases in zone 1 are much more puzzling because they are qualitatively different from onliers, being more out than in the set of X and more in than out of the set of Y. Cases in area 1 therefore *outliers in kind* as regards the consistency of necessity. Since outliers in kind are more interesting than outliers in degree, we recommend to first consider cases in zone 1 as

[14] In principle, one can use the fuzzy-set memberships of a case for formalizing case selection in fsQCA. However, this is an involved matter that we discuss in a separate paper.

targets of case selection. Only if there are no cases in this zone – a potential, if not very likely scenario in applied fsQCA - should one turn to outliers in degree in zone 1 for process tracing.

The ideal-typical deviant case for consistency of necessity is located in the upper-left corner of an x-y plot. In the ideal case, then, the outlier for consistency of necessity has a membership of 1 in the outcome and a membership of 0 in the path because this is the constellation of set memberships that strikes one as most puzzling. More generally, when choosing the outlier consistency for necessity, one should strive for a case with maximum difference between the set memberships in X and Y. We put this line of reasoning as a general principle that applies to the selection of outliers and that we coin the *principle of maximum difference* in fuzzy-set membership scores:

Principle of maximum difference: *Choose the outlier with the maximum difference in its set membership scores in X and Y.*

Sufficiency

Typical cases

In csQCA, typical cases for sufficiency are located in the upper right cell, which corresponds to cell 2 and 3 in the enhanced x-y plot. These are the areas that contain all cases that are more in than out of both X and Y. because of its more fine-grained membership measure, in fsQCA, we can subdivide this group of cases. As typical cases can count only those in zone 2 because cases in zone 3 fall below the main diagonal, thus violating the subset relation of X and Y. Similarly to onliers for necessity, the ideal-typical onlier has a membership of 1 in the path and the outcome. It follows that the best onlier for sufficiency is located as closely as possible to the upper-right corner of the x-y plot. Hence, the p*rinciple of high fuzzy-set membership* equally applies.

Deviant cases consistency

Similarly to deviant cases with respect to consistency of necessity, it is important to distinguish between outliers in kind and degree in the context of consistency for sufficiency. Deviant cases in kind are located in cell 4 of Figure 1. Cases in zone 3 are also inconsistent with a pattern of sufficiency, but they are less puzzling because they are more in than out of both the set of X and Y. In other words, cases in area 3 are outliers in degree. As a consequence of this, one should first determine whether there are outliers in kind available for process tracing. If this is not the case, one should turn to outliers in degree. Following the idea of ideal-typical outliers, the most deviant case for the consistency of sufficiency is located in the lower right corner of Figure 1 or, if this case does not exist empirically, the case that is closest to it. The *principle of maximum difference* thus also applies to outliers for the consistency of sufficiency.

Deviant cases coverage

Deviant cases with regard to coverage are located in the upper left corner of an x-y plot. Their membership in the path is below 0.5 and in the outcome above 0.5. This is exactly what turns cases in zone 1 into deviant cases for coverage. Among these cases, the principle of maximum difference mandates the choice of the case that is placed in the upper-left corner of the plot or that is located closest to this corner.

The discussion of case selection after an fsQCA highlights the usefulness of taking into account both the differences and similarities of kind and of degree among cases. An exclusive emphasis ion differences in kind is not meaningful as it would fail to take advantage of the more fine-grained information contained in fuzzy-set membership scores. At the same time, an exclusive focus on differences in degree would also be flawed as it ignores important qualitative differences among cases on the same side of the main diagonal of an x-y plot. We summarize our line of reasoning that qualitative differences trump differences in degree in two general principles for case selection in fsQCA.

Principle of differences in kind: Case located on different sides of a qualitative anchor differ in kind and should not *be compared with respect to differences in degree.*

Principle of differences in degree: Differences in degree should only be established among cases that are similar in kind and located on the same kind of the main diagonal.

Relating the two principles to our arguments on case selection, one can say that the principle of differences in degree makes a plea for keeping separate cases that are located on the same side of the main diagonal, but different sides of a qualitative anchor. For instance, this pertains to typical cases for the consistency of sufficiency and outliers for coverage of sufficiency. The principle of differences in kind reinforces the argument that one can only compare cases with respect to the degree of their typicalness or deviance if they are similar in kind. The requirement that they are located on the same side of the main diagonal specifically aims at cases in the upper-right quadrant of an x-y plot. They are all similar in kind, but it is of course important to take into account whether they are on the same side of the main diagonal.

SELECTING CASES FOR COMPARATIVE STUDIES AFTER A QCA

So far, our discussion of QCA-based case selection has been concerned with locating individual types of cases in 2x2 tables and x-y plots. Implicitly, we therefore discussed the choice of cases for single-case studies after a QCA. In practice, though, there are at least two good reasons for performing comparative within-case studies (Lees 2006). First, insights gathered on one typical case can be strengthened by selecting another typical case for process tracing (Goertz 2008; Lieberman 2005). Second, one can hardly discern the reason(s) for the deviance of an outlier without knowledge of the causal processes operative in a typical case (George and Bennett 2005, chap. 8). In other words, implicitly, even single case studies do make comparisons to onliers. We think that is more fruitful to make these comparisons explicit. We therefore propose procedures for the choice of cases for comparative process tracing (see Shively 2006). In fact, we argue that the crucial added value of comparative process tracing vis-à-vis single within-case studies is the enhanced ability to search for model-related reasons *why* an outlier is deviant. It is to note that the above-mentioned principle of diverse case selection extends to within-case comparisons. The cases that one compares must be covered by the same path or, as we explain below, by paths that differ from each other in only one condition. Nothing can be learned by studying typical cases that follow different paths, for they influence the outcome in strikingly different ways. Because of this, the number of cross-case comparisons also equals the number of paths in the QCA solution if one aims to get a comprehensive understanding of it.

The following discussion of cross-case comparisons is exemplified with a fsQCA for the sufficient conditions for the outcome. The additional elaboration of comparisons for necessity is beyond the scope of a single paper. We decide for a discussion of comparisons for sufficiency because they are more intricate due to equifinality and conjunctural causation in particular. We argue that one can distinguish two types of basic comparative designs. *Similar-outcome designs* match cases that are qualitatively identical with regard to their membership in the set of the outcome, that is, the membership in Y is above .5 for each case. *Dissimilar-outcome designs* match cases with qualitative differences with regard to their membership in Y. We argue that, taken together, these two designs cover all possible meaningful comparisons in QCA-based MMR. No matter which typical and deviant cases are taken for comparison, at least one of the two or more cases must be a member of the set for the outcome of interest. In order to emphasize this point and the fact that not all logically possible comparisons are analytically meaningful, we formulate the *positive outcome principle* for comparisons:

Positive outcome principle: At least one case involved in comparative process tracing must be a member of the outcome.

Our distinction between similar-outcome and dissimilar-outcome designs seems to mimic the well-known method of agreement and method of difference (Mahoney 1999), also referred to as the most-dissimilar and most-similar comparisons (De Meur and Berg-Schlosser 1996). As we explain in the following, this impression is wrong. We show that there is no legitimate place for the method of agreement in QCA-based small-n comparisons because the QCA does what the method of agreement is used in standalone cross-case case studies. With respect to the method of difference, we argue that it should be applied in two different variants. Thus, performing a QCA in advance of comparative process tracing not only allows one to formalize case selection. It also lays the ground for comparative designs that differ from the established comparisons known from the small-n literature (Tarrow 2010).

We illustrate the logic of comparative case selection with data used by Vis (2009). Her research interest is to explain the occurrence of unpopular social policy reforms (UR). The three conditions are a weak socio-economic situation (WSE), a weak political position of the government (WPP), and rightist government (RIGHT). The 25 cases are social policy reforms of cabinets in Denmark, Germany, The Netherlands, and the United Kingdom. The fuzzy sets memberships used in the fsQCA are presented in the appendix to this paper.[15] The complex solution and the corresponding measures of fit can be found in Table 4.

Table 4: Solution and measures of fit for complex solution

Number of path	Path	Consistency	Raw Coverage	Unique coverage
1	WSE*WPP	0.90	0.71	0.24
2	WSE*RIGHT	0.91	0.62	0.15
Solution consistency: 0.90 Solution coverage: 0.86				

The entire solution as well as both paths are above the plausible threshold for consistency of 0.80. The solution coverage and the raw coverage of both paths are relatively high. However, the unique coverage of each path is relatively low, indicating that many cases are covered by both paths. In the following, we

15 A detailed description of the coding of the outcome and the conditions can be found in Vis's article (2009, 37-43, 50).

take this solution as the basis for elaborating our case selection principles for comparative process tracing. This also requires the identification of typical case and of deviant case for consistency and coverage.

Similar-outcome comparisons

Logically speaking, three different designs are available within the realm of similar-outcome comparisons: a comparison of several typical cases, a comparison of typical cases with outliers in relation to coverage, and a comparison of an outlier consistency from zone 1 with a consistent individually irrelevant case from zone 6. We argue that only the latter makes sense.

The comparison of an outlier consistency with an individually irrelevant case design does not have any inferential merit in an MMR aiming at sufficient conditions because the outcome is *absent* in both cases. This violates the positive outcome principle, which stipulates that *at least one case* in the comparative design displays the outcome of interest. Nothing can be learned empirically about sufficient condition for the outcome by exclusively looking at cases that lack the outcome. As regards the study by Vis, it is apparent that one can hardly learn anything about the conditions under which unpopular social policy reforms occur if one only examines cases of non-unpopular reform.

The arguments for why a comparison of an onlier with an outlier coverage case is not feasible are more intricate and we discuss them in detail below.

Matching multiple onliers

The comparison of two onliers should be guided by the goal to select typical cases spanning a maximum range of membership scores in the condition and the outcome. One can significantly increase the confidence in the theoretical account at hand if it can be shown that similar causal processes are operative in onliers with different membership in the condition and the outcome (Goertz 2008; Lieberman 2005). Ideally, one would compare two cases on or closely above the main diagonal that separates zone 2 and 3, with one case being located on the upper and the other on the lower end of the diagonal. Applied to the example by Vis, one should select the two cases on the main diagonal for the within-case comparison of onliers. These are the first government of Luebbert (upper case on the diagonal) and the fourth cabinet of Helmut Kohl (lower case on the diagonal). Figure 2 presents the distribution of onliers for the path *WSE*WPP*.

Figure 2: Typical cases for path WSE*WPP

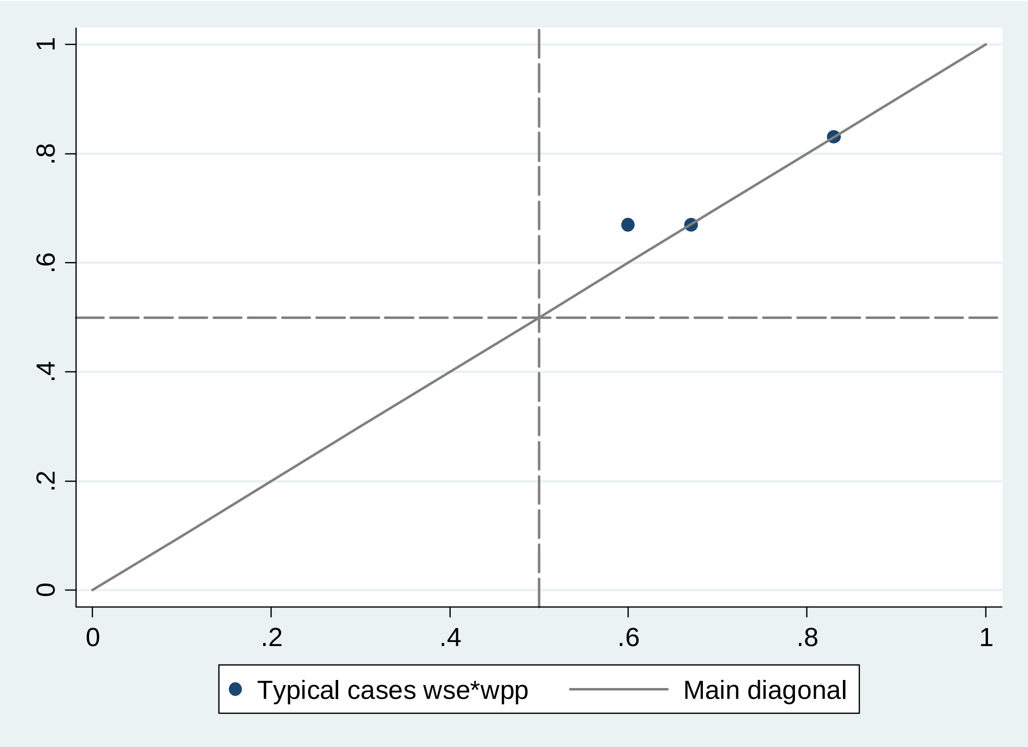

The recommendation to maximize the range of set memberships in QCA-based case selection is, in fact, an adaptation of the argument to maximize the variation of the independent variable of interest in pairwise comparisons (Lijphart 1971). The caveat to consider in the context of set-theoretic methods is that the variation in X should be limited to cases in zone 6 rather than covering the whole variation on the conditions involved so as to pay justice to the qualitative anchors. In set-theoretic methods it would be outright wrong to use the full range of set membership score as if there were no qualitative differences.

The impossibility of matching onliers with outlier coverage cases

A second variant of a similar-outcome comparison that seems to suggest itself includes an onlier and outlier for coverage. However, such a comparison is not feasible. This claim might sound counterintuitive, but a reconsideration of the reasons why a case is deviant with respect to coverage shows that intuition is misleading in this instance. A case is an outlier for coverage because it is not covered by any path of the solution, but displays the outcome. There must be another sufficient path that is not a subset of any paths included in the QCA solution. In other words, the path that covers the outlier coverage is fundamentally different from the paths in the solution. Consequently, nothing can be learned from comparing a typical case and an outlier for coverage.

The discussion of similar-outcome comparisons after a QCA provides two interesting insights into the more general small-n literature on comparisons. First, the comparison of a typical case and a deviant case is commonly advertised as one of the key benefits of case studies as it sheds light on potentially omitted causes that, once taken into the picture, remove the puzzle (George and Bennett 2005, 20). Our discussion has revealed that in the case of QCA-based MMR this statement is under-specified. It only applies to one type of deviant case – that of consistency. It makes little sense, however, with regard to deviant cases with regard to coverage.

Second, if, as we claim, a comparison of onliers with outlier coverage cases is not a good analytic strategy, then also Mill's method of agreement is not. Both comparisons consist in matching two cases with the same membership in the outcome and at least one condition and different membership scores in all other conditions. The usual inference from this that the condition(s) in which they share the same membership is the cause for the occurrence of the outcome. As (Ragin 1987, 36-38) argues, this conclusion is flawed not only when the outcome of interest is the result of equifinal processes, but also when there is conjunctural causation. To this we can add two related reasons why Mill's method of agreement should not be used. First, because the two cases that are matched differ so much, they most likely belong to a different causal universe and require different explanations. Second, if one attributes a causal role to one single condition based on Mill's method of agreement, then one is making untested, and often implausible, assumptions about all the logically possible but empirically unobserved cases that can be created based on the conditions specified for defining the difference between cases (Schneider/Wagemann 2007: 73-77???). Both pitfalls remain hidden in the standard application of Mill.

Dissimilar-outcome comparisons

Logically speaking, there are three different pairs of onlier and outlier cases that would fit into this scheme: onlier – outlier consistency; outlier coverage – individually irrelevant case; and outlier coverage – outlier consistency. The latter comparison is meaningless. These two types of cases not only differ in their membership in Y but also Y. there is nothing puzzling about these cases and nothing can be learned by studying them in a comparative within-case analysis. The two other types of comparisons, however, make analytic sense.

Matching a typical case with an outlier consistency

A typical case and an outlier for consistency are both members of the same path, but differ with respect to the membership of the outcome. The similarity on X and dissimilarity on Y is what makes this form of dissimilar-outcome comparisons attractive. The aim of this comparison is to identify a condition that is causally irrelevant. By dropping such a condition from a path, the hitherto outlier consistency case turns into an irrelevant case.

A comparison of an onlier and an outlier for consistency should meet two criteria. First, following the maximum membership principle, the membership scores in the path should be as high as possible for both cases. Second, the membership in X should be as similar as possible. It is true that the typical case and deviant case are qualitative identical with respect to the path. However, one should additionally take advantage of fuzzy set measurement and make sure that the fuzzy-set membership in X is minimal. Third, the two cases should have a maximum difference in their membership in Y because a large difference in the membership in the outcome represents a puzzle in light of the similar membership in X. In terms of the location of the cases in an x-y plot, the joint application of all three criteria imply that the onlier and outlier are located on the far right, exhibit a large vertical distance and are located on the same vertical line that runs through a high membership score in X.

Figure 3 presents the onliers and outliers for consistency that are uniquely covered by path *WSE*WPP*. As the plot shows, there is not much of a choice between deviant cases because there is only one. In our empirical example, the outlier for consistency is the fourth Danish government that was led by Schlüter between May 1988 and December 1990. The scarcity of outliers consistency cases should not come as a surprise when the consistency threshold was set reasonably high. One A consistency score 0.8 (recommended (Ragin 2006a) and widely used) leaves some room for inconsistency, but not much. The consequence is that the number of outliers in kind in zone 4 is rather limited because they are likely to decrease the consistency figure substantially. The number of outliers in degree in zone 3 can be somewhat higher because they are likely to be closer to the diagonal than the outlier in kind.

Figure 3: Typical cases and deviant case for consistency for path WSE*WPP

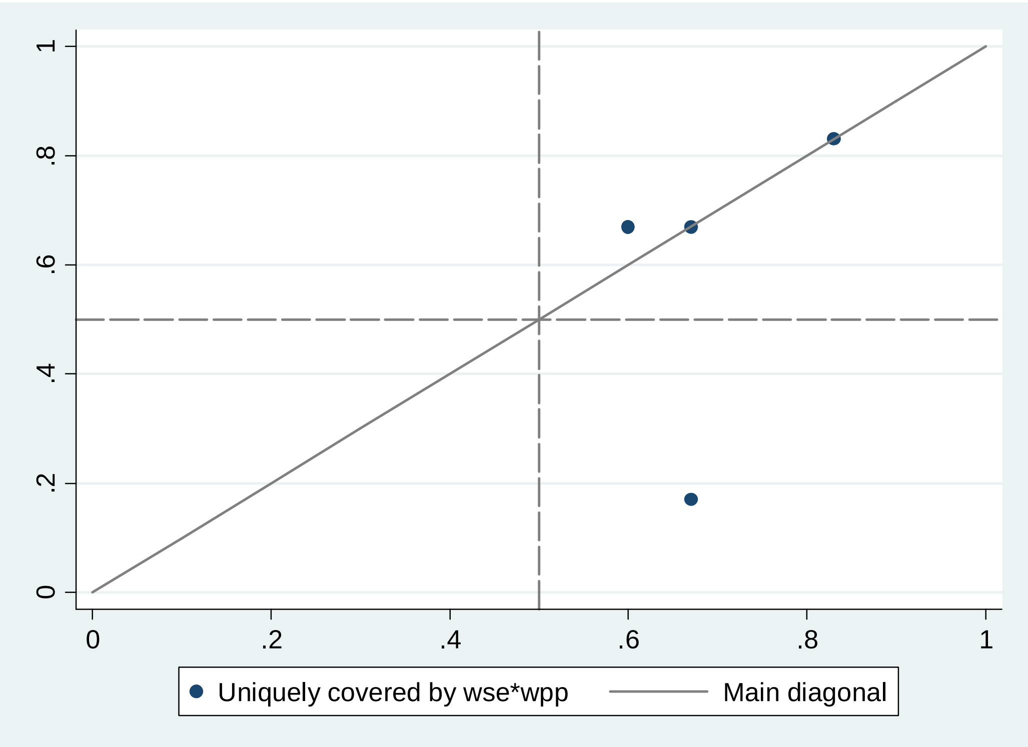

In this instance, the outlier should be compared with the typical case that has exactly the same membership in X, which is the third cabinet led by Thatcher between June 1987 and April 1992. Figure 3 highlights that there is one other onlier with a higher membership in Y and, thus, a larger distance to the onlier. However, the increase in the gap in the set membership for Y comes at the expense of an even larger decrease in the similarity as regards the set membership in the path. For this reason, a comparison of the fourth Schlüter government and the third government by Thatcher is the best match for comparative process tracing.

Although the comparison of an onlier and an outlier for consistency may remind one of Mill's method of difference (Lijphart 1975), it is fallacious to treat the two comparison as equivalent. The method of difference is applied when two cases differ on the outcome and when all conditions but one are identical. When comparing an onlier and an outlier for consistency, in contrast, the two cases are covered by exactly the same path. This design hence mirrors the classic comparison of a typical and deviant case where two cases are exactly the same but display a different outcome nevertheless. In the empirical example, the Schlüter government and Thatcher government were both confronted with a weak socio-economic situation and a weak political situation. The QCA solution tells us that governments then undertake unpopular reforms, which is what Thatcher did but Schlüter not. As explained before, the puzzle is solved if and when one is able to identify a condition that has been omitted from the path in which the typical case (Thatcher) has a high set membership and the deviant case (Schlüter) has not. *After* this condition has been added, the former outlier turns into an individually irrelevant case.

Matching an outlier coverage with an individually irrelevant case

Above, we explain why outlier coverage cases should not be matched with onlier cases for a within-case analysis. Instead, a subset of the individually irrelevant cases provides adequate cases for such a comparison. Our argument is this: an outlier coverage case needs to be compared with an individually irrelevant case that falls into the same truth table row. This is an ideal setting for within-case analysis: both cases display the same combination of conditions, yet differ in their membership in the outcome. Notice that this is identical to the pairwise comparison of an onlier with an outlier consistency described above. Hence, also here the aim of the within-case analysis is to identify a condition from the conjunction that might be causally irrelevant. Eliminating this condition from the conjunction would leave the individually irrelevant case as it is but turn the outlier coverage case into an onlier in the newly created sufficient path.

Due to the principle of diverse case selection, separate pairwise comparisons are required for each outlier coverage that alls into a separate truth table row. The maximum number of comparisons is a function of the number of conditions involved in the truth table and the number of conditions involved in the sufficient paths in the QCA solution. The formula for calculating the maximum number of comparisons required is $2*k – (k-n)^2$, where *k* indicates the number of conditions in the truth table and *n* the number of conditions involved the sufficient path.[16] For instance, the maximum number of pairwise comparisons in a study with a truth table of 5 conditions and a sufficient path consisting of 2 conditions would be 25 – (5-2)2 = 15.

The potential number of required comparison seems prohibitively high to be feasible in research practice. There are, however, two general features of observational that both contribute to (drastically) reducing this number. First, as our example indicates, many of the logically possible combinations do not occur in the data due to limited diversity (Ragin 1987). In addition, even if they do, they sometimes only occur among the outliers coverage cases or the individually irrelevant cases, but not both. Obviously, a comparative within-case study is only feasible if a specific conjunction occurs in both sets of cases. Second, the number of eligible cases for pairwise comparisons is further limited when equifinality is at play. This is so because one should only examine outliers coverage that are not covered by *any* path of the QCA solution. Their number is usually not all too high. For these reasons, the potentially high number of pairwise comparisons between outliers coverage and irrelevant cases will be much lower in practice and probably often in a range that is manageable. If the number of required pairwise comparisons is not feasible or simply not performed, one should explicate the possible comparison that have not been analyzed and integrate the heightened level of uncertainty into the interpretation of the cross-sectional QCA solution formula.

[16] Truth table rows that are a subset of the sufficient path should be excluded. This is achieved by the expression $(k-n)^2$. In the presence of equifinality, n represents the number of conditions that are involved in any of the multiple sufficient conjunctions.

Figure 4: Plot of non-covered and individually irrelevant cases

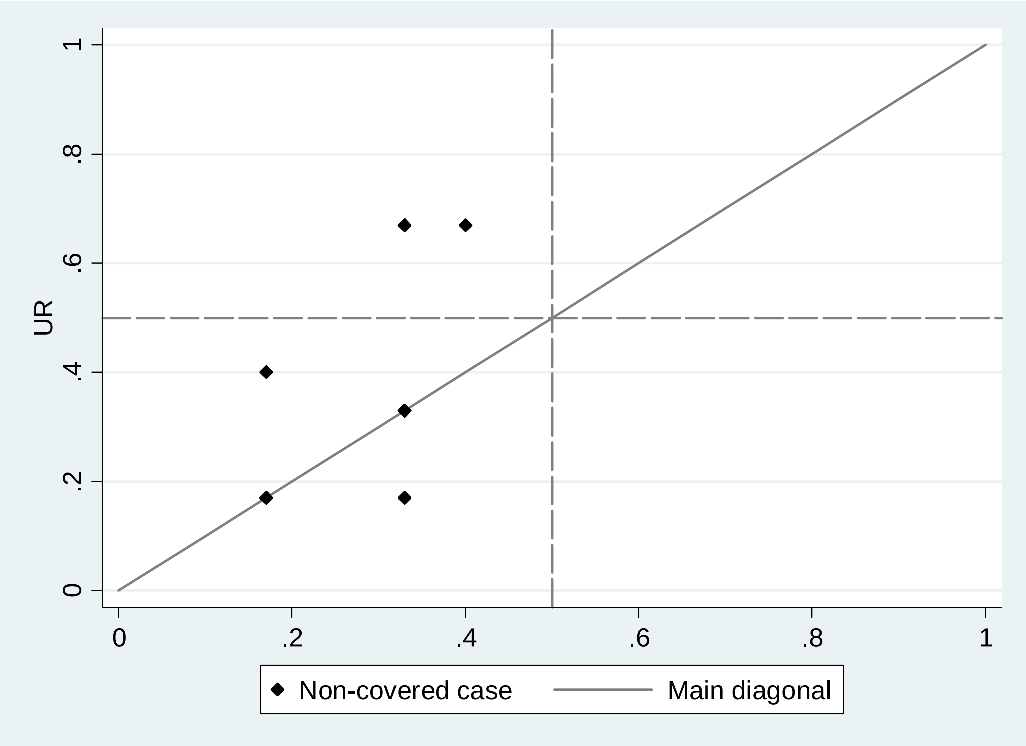

Summary

The configurational view on causation renders QCA a highly suitable cross-case method for combination with process tracing without having to reinvent the wheel. Some of the arguments we make in the context of QCA-based MMR are firmly grounded in general principles developed by the established literature on single case and comparative case studies. By staying true to QCA's insistence on set theory, configurational causation, and causal heterogeneity, our strategies are able to vastly diminish some of the pertinent problems of small-n research (Tarrow 2010). However, it also became apparent that a simple transfer of all established procedures is neither meaningful nor viable. QCA-based pairwise comparisons mitigate three risks. First, drawing inferences based on strong assumptions about counterfactuals (Ragin 2008, chap. 8-9), that often remain hidden not only from the reader but also the researcher herself. Second, the erroneous attribution of causal effects to single variables rather than configurations, as most prominently done in Mill's method of difference and agreement (Lieberson 1991). Third, the neglect of potential equifinality and implications for causal inference (George and Bennett 2005, chap. 8).

While formalizing established comparative designs in the context of QCA-based MMR, we further expanded the perspective by discussing additional, viable cross-case comparisons that hitherto have not been highlighted enough in the existing literature. To our knowledge, no research has engaged in an explicit comparison of outlier coverage cases with individually irrelevant cases. Cases that are not members of a path and the outcome are usually seen as analytically irrelevant. By matching them with outlier coverage cases, they can yield meaningful insights, however. In addition, we have shown how QCA results can be used to derive clues about the most likely sources of deviance that should guide process tracing and the insights of which are in turn fed back into the cross-case model. By introducing different basic forms of reasons for deviance and model misfit, researchers performing within-case analysis have more precise guidance as to what they ought to be looking for. Especially the insight that one should differentiate between the under-fitting of a path (outlier consistency) as opposed to the under-fitting of an entire QCA solution term (outlier coverage) does not feature in MMR based on statistical techniques. The under-fitting of a path as indicated by an outlier consistency requires to acknowledge that these cases need to be explained by dding a hitherto missing condition to an already identified sufficient path. The under-fitting of a solution as indicted by an outlier coverage, instead, implies that the solution in solving this type of misfit consists in finding an entirely new path for explaining this type of case. In sum, the comparative within-case analysis that involves outlier coverage cases leads to a new path whereas those involving outlier consistency cases produces a new path that is a subset of an already existing path.

Table 5: Fuzzy-set memberships for Vis (2009)

Government	WPP	WSE	RIGHT	UR
Lubbers I	0.33	0.83	1.00	0.83
Lubbers II	0.17	0.33	1.00	0.33
Lubbers III	0.33	0.67	0.60	0.67
Kok I	0.17	0.40	0.40	0.67
Kok II	0.33	0.33	0.40	0.17
Balkenende II	0.67	0.67	1.00	0.83
Kohl I	0.17	0.33	1.00	0.33
Kohl II	0.33	0.17	1.00	0.17
Kohl III	0.17	0.33	1.00	0.33
Kohl IV	0.67	0.67	1.00	0.67
Schröder I	0.33	0.40	0.00	0.17
Schröder II	0.83	0.83	0.00	0.83
Schlüter I	0.33	0.33	1.00	0.33
Schlüter II	0.33	0.60	1.00	0.67
Schlüter IV	0.33	0.67	1.00	0.17
Schlüter V	0.60	0.67	1.00	0.33
N.Rasmussen I	0.17	0.17	0.40	0.17
N.Rasmussen II (& III)	0.60	0.60	0.25	0.83
N.Rasmussen IV	0.33	0.33	0.25	0.67
Thatcher I	0.17	0.83	1.00	0.83
Thatcher II	0.33	0.33	1.00	0.67
Thatcher III	0.33	0.67	1.00	0.67
Major I	0.33	0.60	1.00	0.67
Blair I	0.17	0.33	0.00	0.40
Blair II	0.33	0.33	0.00	0.33

References

Adcock, Robert and David Collier (2001): Measurement Validity: A Shared Standard for Qualitative and Quantitative Research. *American Political Science Review* 95 (3): 529-546.

Amenta, Edwin and Jane D. Poulsen (1994): Where to Begin: A Survey of Five Approaches to Selecting Independent Variables for Qualitative Comparative Analysis. *Sociological Methods Research* 23 (1): 22-53.

Bäck, Hanna and Patrick Dumont (2007): Combining Large-*N* and Small-*N* Strategies: The Way Forward in Coalition Research. *West European Politics* 30 (3): 467-501.

Berg-Schlosser , Dirk and Gisèle De Meur (2008): Comparative Research Design: Case and Variable Selection. Rihoux, Benoit and Charles Ragin (eds.): *Configurational Comparative Methods*. Thousand Oaks: Sage: 19-32

Berg-Schlosser , Dirk, Gisèle De Meur, Charles C. Ragin and Benoît Rihoux (2008): Qualitative Comparative Analysis as an Approach. Rihoux, Benoit and Charles Ragin (eds.): *Configurational Comparative Methods*. Thousand Oaks: Sage: 1-18

Braumoeller, Bear F. and Gary Goertz (2000): The Methodology of Necessary Conditions. *American Journal of Political Science* 44 (4): 844-858.

Braumoeller, Bear F. and Gary Goertz (2002): Watching Your Posterior: Comment on Seawright. *Political Analysis* 10 (2): 198-203.

Bunce, Valerie (2000): Comparative Democratization: Big and Bounded Generalizations. *Comparative Political Studies* 33 (6-7): 703-734.

Campbell, Donald T. (1975): Degrees of Freedom and the Case Study. *Comparative Political Studies* 8 (2): 178-193.

Clarke, Kevin A. (2002): The Reverend and the Ravens: Comment on Seawright. *Political Analysis* 10 (2): 194-197.

Cooper, Barry and Judith Glaesser (2011): Using Case-Based Approaches to Analyse Large Datasets: A Comparison of Ragin's Fsqca and Fuzzy Cluster Analysis. *International Journal of Social Research Methodology* 14 (1): 31-48.

Coppedge, Michael (1999): Thickening Thin Concepts and Theories - Combining Large N and Small in Comparative Politics. *Comparative Politics* 31 (4): 465-476.

De Meur, Gisèle and Dirk Berg-Schlosser (1996): Conditions of Authoritarianism, Fascism, and Democracy in Interwar Europe: Systematic Matching and Contrasting of Cases for "Small N" Analysis. *Comparative Political Studies* 29 (4): 423-468.

Dion, Douglas (1998): Evidence and Inference in the Comparative Case Study. *Comparative Politics* 30 (2): 127-145.

Eckstein, Harry (1975): Case Study and Theory in Political Science. Greenstein, Fred I. and Nelson W. Polsby (eds.): *Strategies of Inquiry. Handbook of Political Science, Vol. 7*. Reading, Mass.: Addison-Wesley: 79-137

George, Alexander L. and Andrew Bennett (2005): *Case Studies and Theory Development in the Social Sciences*. Cambridge, Mass.: MIT Press.

Gerring, John (2004): What Is a Case Study and What Is It Good For? *American Political Science Review* 98 (2): 341-354.

Gerring, John (2007): Is There a (Viable) Crucial-Case Method? *Comparative Political Studies* 40 (3): 231-253.

Glaesser, Judith and Barry Cooper (2010): Selectivity and Flexibility in the German Secondary School System: A Configurational Analysis of Recent Data from the German Socio-Economic Panel. *European Sociological Review*:

Goertz, Gary and Harvey Starr (2003): Introduction: Necessary Condition Logics, Research Design, and Theory. Goertz, Gary and Harvey Starr (eds.): *Necessary Conditions. Theory, Methodology, and Applications*. Lanham: Rowman & Littlefield: 1-24

Goertz, Gary (2006): Assessing the Trivialness, Relevance, and Relative Importance of Necessary and Sufficient Conditions in Social Science. *Studies in Comparative International Development* 41 (2): 88-109.

Goertz, Gary and Jack S. Levy (2007): Causal Explanation, Necessary Conditions, and Case Studies. Levy, Jack S. and Gary Goertz (eds.): *Explaining War and Peace: Case Studies and Necessary Condition Counterfactuals*. London: Routledge: 9-45

Goertz, Gary (2008): Choosing Case for Case Studies: A Qualitative Logic. *Newsletter of the APSA Section on Qualitative & Multi-Method Research*: 11-14.

Grendstad, Gunnar (2007): Causal Complexity and Party Preference. *European Journal of Political Research* 46 (1): 121-149.

Hall, Peter A. (2008): Systematic Process Analysis: When and How to Use It. *European Political Science* 7 (3): 304-317.

Laver, Michael (2003): Government Termination. *Annual Review of Political Science* 6: 23-40.

Lees, Charles (2006): We Are All Comparativists Now: Why and How Single-Country Scholarship Must Adapt and Incorporate the Comparative Politics Approach. *Comparative Political Studies* 39 (9): 1084-1108.

Levy, Jack S. (2008): Case Studies: Types, Designs, and Logics of Inference. *Conflict Management and Peace Science* 25 (1): 1-18.

Lieberman, Evan S. (2005): Nested Analysis as a Mixed-Method Strategy for Comparative Research. *American Political Science Review* 99 (3): 435-452.

Lieberson, Stanley (1991): Small Ns and Big Conclusions: An Examination of the Reasoning in Comparative Studies Based on a Small Number of Cases. *Social Forces* 70 (2): 307-320.

Lijphart, Arend (1971): Comparative Politics and the Comparative Method. *American Political Science Review* 65 (3): 682-693.

Lijphart, Arend (1975): Comparable-Cases Strategy in Comparative Research. *Comparative Political Studies* 8 (2): 158-177.

Mackie, John L. (1965): Causes and Conditions. *American Philosophical Quarterly* 2: 245-264.

Mahoney, James (1999): Nominal, Ordinal, and Narrative Appraisal in Macrocausal Analysis. *American Journal of Sociology* 104 (4): 1154-1196.

Mahoney, James, Erin Kimball and Kendra L. Koivu (2009): The Logic of Historical Explanation in the Social Sciences. *Comparative Political Studies* 42 (1): 114-146.

Obinger, Herbert, Stephan Leibfried, Claudia Bogedan, Edith Gindulus, Julia Moser and Peter Starke (2010): *Transformations of the Welfare State: Small States, Big Lessons*. Oxford: Oxford University Press.

Ragin, Charles (1981): Comparative Sociology and the Comparative Method. *International Journal of Comparative Sociology* 22 (1-2): 102-120.

Ragin, Charles C. (1987): *The Comparative Method: Moving Beyond Quantitative and Qualitative Strategies*. Berkeley: University of Berkeley Press.

Ragin, Charles C. (2000): *Fuzzy-Set Social Science*. Chicago: University of Chicago Press.

Ragin, Charles C. (2006a): Set Relations in Social Research: Evaluating Their Consistency and Coverage. *Political Analysis* 14 (3): 291-310.

Ragin, Charles C. (2006b): How to Lure Analytic Social Science out of the Doldrums: Some Lessons from Comparative Research. *International Sociology* 21 (5): 633-646.

Ragin, Charles C. (2008): *Redesigning Social Inquiry*. Chicago: Chicago University Press.

Rihoux, Benoit and Charles Ragin (eds.) (2008): *Configurational Comparative Methods*. Thousand Oaks: Sage.

Rihoux, Benoît and Bojana Lobe (2009): The Case of Qualitative Comparative Analysis (Qca): Adding Leverage for Cross-Case Comparison. Byrne, David and Charles C. Ragin (eds.): *The Sage Handbook of Case-Based Methods*. Thousand Oaks: Sage: 222-242

Rohlfing, Ingo (2008): What You See and What You Get: Pitfalls and Principles of Nested Analysis in Comparative Research. *Comparative Political Studies* 41 (11): 1492-1514.

Schneider, Carsten Q. and Bernard Grofman (2006): It Might Look Like a Regression ... But It's Not! An Intuitive Approach to the Presentation of Qca and Fs/Qca Results. *Compasss Working Paper, WP2006-39*:

Schneider, Carsten Q. (2008): *The Consolidation of Democracy: Comparing Europe and Latin America*. London: Routledge.

Schneider, Carsten Q. and Claudius Wagemann (2010): Standards of Good Practice in Qualitative Comparative Analysis (Qca) and Fuzzy Sets. *Comparative Sociology* 9 (3): 397-418.

Schneider, Gerald and Nils Petter Gleditsch (2010): The Capitalist Peace: The Origins and Prospects of a Liberal Idea. *International Interactions* 36 (2): 107-114.

Seawright, Jason (2002a): Testing for Necessary and/or Sufficient Conditions: Which Cases Are Relevant? *Political Analysis* 10: 178-193.

Seawright, Jason (2002b): What Counts as Evidence? Reply. *Political Analysis* 10 (2): 204-207.

Seawright, Jason and John Gerring (2008): Case Selection Techniques in Case Study Research: A Menu of Qualitative and Quantitative Options. *Political Research Quarterly* 61 (2): 294-308.

Shively, W. Phillips (2006): Case Selection: Insights from *Rethinking Social Inquiry*. *Political Analysis* 14 (3): 344-347.

Tarrow, Sidney (2010): The Strategy of Paired Comparison: Toward a Theory of Practice. *Comparative Political Studies* 43 (2): 230-259.

Vis, Barbara (2009): Governments and Unpopular Social Policy Reform: Biting the Bullet or Steering Clear? *European Journal of Political Research* 48 (1): 31-57.