Quality & Quantity
https://doi.org/10.1007/s11135-025-02303-7

QCA vs. CNA: fundamental clarifications

Adrian Dușa[1]

Accepted: 18 July 2025

Abstract
The ongoing debate within Configurational Comparative Methods (CCMs) primarily revolves around how different solution types are derived and interpreted. While Qualitative Comparative Analysis (QCA) generates three types of solutions (conservative, parsimonious, and intermediate), Coincidence Analysis (CNA) produces only one. This difference has fueled discussions regarding their respective methodological strengths and limitations. This paper aims to clarify fundamental misconceptions surrounding QCA, particularly in relation to CNA. It critically examines the role of sufficiency, necessity, and the implications of different minimization approaches. By addressing key misinterpretations (such as the assumption-free nature of certain solutions and the role of counterfactuals) this paper provides a structured comparison of QCA and CNA. Additionally, it highlights the methodological trade-offs involved in prioritizing either robust sufficiency or redundancy-free models. The paper concludes with recommendations for researchers in CCMs, aiming to foster a more precise understanding of these methods and their appropriate applications.

Keywords QCA · Boolean minimization · Robust sufficiency

1 Introduction

Qualitative Comparative Analysis (QCA) and Coincidence Analysis (CNA) are often positioned as competing methodologies within the field of configurational comparative methods (CCMs). While they share a common ambition, such as identifying complex causal structures from empirical data, they diverge substantially in both their formal foundations and their epistemological commitments. Chief among these divergences is the way each method handles solution types, particularly regarding assumptions about unobserved configurations and the minimization strategies employed.

✉ Adrian Dușa
dusa.adrian@unibuc.ro

[1] Department of sociology, University of Bucharest, Șoseaua Panduri nr. 90-92, sector 5, Bucharest 050663, Romania

Published online: 28 July 2025

QCA, in its canonical form, yields three types of solutions: conservative, parsimonious, and intermediate. CNA, by contrast, generates a single solution type, typically equivalent to the parsimonious QCA solution under comparable consistency thresholds. This difference is not merely procedural. It reflects deeper disagreements over how sufficiency, necessity, and empirical relevance should be operationalized within a configurational framework.

This paper addresses several conceptual and procedural ambiguities in the debate between QCA and CNA. It focuses in particular on misunderstandings concerning solution derivation, the interpretation of counterfactuals, and the status of empirical versus theoretical assumptions in minimizing Boolean expressions. The objective is not to promote one method over the other, but to clarify the conceptual terrain on which such comparisons must rest.

To that end, the following sections examine key points of divergence between QCA and CNA, with special attention to the notion of robust sufficiency, the interpretation of conservative solutions, and the risks associated with over-minimization. The aim is to sharpen the conceptual distinctions necessary for meaningful methodological evaluation.

1.1 QCA vs. CNA: key differences

Qualitative Comparative Analysis (QCA) and Coincidence Analysis (CNA) both aim to identify causal structures from configurational data, and find minimally sufficient conditions for a certain outcome. Despite this shared objective, they diverge on central methodological dimensions, some of which thoroughly analyzed by Swiatczak (2021). These divergences pertain to their underlying data structures, the range of solution types they produce, and their treatment of counterfactuals.

- **Data structure.** QCA constructs and analyzes truth tables, i.e., exhaustive representations of observed and logically possible configurations. CNA, by contrast, is based on coincidence lists: collections of co-occurring events without imposing a full configuration space. This fundamental difference constrains CNA to operate within observed data and avoids extrapolating from unobserved combinations.
- **Solution types.** QCA distinguishes three solution types: conservative, parsimonious, and intermediate. These differ in their assumptions about unobserved configurations (remainders) and the extent of minimization applied. CNA offers only one solution type, which under compatible parameters corresponds to QCA's parsimonious solution. Consequently, CNA does not accommodate the distinctions QCA draws between conservative and intermediate solutions.
- **Minimization procedures.** Traditional QCA applies the Quine-McCluskey (QM) algorithm, which minimizes expressions by iteratively introducing counterfactuals. Newer implementations, such as the Consistency Cubes method (Dușa 2018), generate fully QM-compatible solutions while avoiding counterfactual assumptions, relying strictly on observed configurations. CNA, in contrast, rejects the role of counterfactuals entirely and constructs redundancy-free models from empirical co-occurrence patterns alone.
- **Theoretical orientation.** QCA allows for the integration of theoretical expectations, especially in its intermediate solution, via directional assumptions on remainders. CNA avoids this by design. It treats empirical co-occurrence as self-sufficient, with causal claims derived strictly from data-driven regularities. This difference reveals a deeper

divergence: QCA admits a role for theoretical constraints; CNA maintains a strictly observational stance.

- **Empirical equivalence.** Under identical input and parameter settings, CNA and QCA produce an identical solution, typically the parsimonious one (Dușa 2019). However, this equivalence is conditional and does not hold across all cases, particularly when QCA's solution type is chosen based on theoretical or empirical considerations that CNA is unable to model.

In short, QCA and CNA differ less in output (under ideal convergence) than in their methodological framing. Their distinctions are most visible when evaluating the epistemic trade-offs between parsimony, theoretical integration, and empirical robustness.

A central contribution to the debate comes from Baumgartner and Thiem (2017), who evaluate QCA's three solution types (conservative, intermediate, and parsimonious) based on a formal correctness criterion: a solution is deemed correct if it contains only causally relevant conditions. By this criterion, only the parsimonious solution qualifies, leading them to recommend that the conservative and intermediate solutions be "immediately discontinued" with far-reaching implications: several later interpretations, such as Forner et al. (2019, p. 326), extended this conclusion to QCA as a whole, thus conflating a critique of solution types with a critique of the method itself.

This position assume a custom interpretation of sufficiency that is neither standard nor universally accepted, namely that a configuration is sufficient only if it contains no causally irrelevant conditions. It is a view that departs from the more commonly held understanding in QCA, where sufficiency is assessed with respect to observed empirical consistency and, in some cases, refined through minimization under the constraint that the outcome must robustly occur in the simplified solution. As will be developed below, this latter idea, central to the notion of robust sufficiency, offers an alternative standard that prioritizes empirical retention over redundancy elimination.

The following sections analyze how differing notions of sufficiency (empirical and counterfactual) underpin these disagreements. In doing so, they clarify the conceptual stakes involved in evaluating QCA solutions and explain why divergent conclusions often stem from distinct, and sometimes incompatible, epistemological commitments.

1.2 The concept of robust sufficiency

The concept of sufficiency is central to QCA but remains subject to multiple interpretations. In its standard empirical form, a configuration is said to be sufficient for an outcome if, in all observed cases where the configuration is present, the outcome is also observed. This form of sufficiency is data-bound and makes no assumptions about what might happen in unobserved or logically possible configurations.

However, QCA does not end with identifying empirical sufficiency. While the conservative solution adheres closely to empirical observations, the parsimonious solution simplifies the configuration space by introducing counterfactual assumptions about unobserved configurations. In doing so, it risks eliminating conditions that may be causally relevant in the observed cases.

To address this concern, the notion of **robust sufficiency** was introduced by Dușa (2019) as an empirical safeguard in the minimization process. Its purpose is to retain the consistency of the outcome in minimized solutions, especially when counterfactuals are used. A configuration is said to be *robustly sufficient* if, after minimization, the outcome still (consistently) occurs in all covered positive cases. The definition is as follows:

"A disjunct in a QCA solution is robustly sufficient if the outcome is guaranteed to occur in its presence." — Dușa (2019), p. 11.

The phrasing "guaranteed to occur" has led to controversy. In particular, Baumgartner (2021) critiques the term as vacuous, arguing that such guarantees cannot be made in real-world, fragmented data. However, this critique rests on a misunderstanding of scope. Robust sufficiency does not claim to provide ontological certainty. It specifies that, within the confines of the observed dataset (or a simulated dataset with a known structure), the minimized configuration continues to cover only those cases where the outcome occurs. The guarantee is internal to the dataset, not a metaphysical claim about causation.

Baumgartner's critique appears to conflate empirical guarantees with ontological certainty. The point of robust sufficiency is not to claim universal truth but to define a solution criterion that maintains consistency with observed positive cases, without discarding potentially relevant conditions in the pursuit of syntactic minimality.

A solution is *robustly sufficient* if, during minimization, no condition that is empirically necessary for the outcome is lost. This definition shifts the emphasis from pure logical reduction to preservation of causal structure.

Consider the following illustrative example. Suppose the following configuration leads to a fire:

$$S \cdot H \cdot O \cdot \sim E \cdot D \rightarrow F$$

where S: short circuit, H: inflammable material, O: oxygen, $\sim E$: absence of extinguisher, D: dry environment, and F: fire. A parsimonious solution derived from this dataset might eliminate S and H, resulting in:

$$O \cdot \sim E \cdot D \rightarrow F$$

which appears sufficient in the data.

However, this simplification omits conditions (e.g., S and H) that are in fact causally necessary. This key issue is especially important in the context of incomplete data. If the dataset lacks cases where D is present but the outcome is absent, then D may appear empirically sufficient, even though we intuitively understand that a mere dry environment cannot produce a fire on its own, knowing from the full causal structure that $S \cdot H \cdot O \sim E \cdot D$ is required. In such cases, D is wrongly identified as sufficient, simply because the data no longer contains the necessary contrast.

The concept of "robust sufficiency" captures this very problem: whenever a relevant condition is over-minimized, the correctness of the QCA solution should be equal to zero, because the outcome fails to happen. Real life researchers never know which conditions are relevant and which are not, but simulated data with an a priori known, full causal structure of the outcome, makes it possible to identify the situations when the relevant conditions are incorrectly eliminated.

Baumgartner (2021) introduces a curious distinction between "robustly sufficient" and "minimally robustly sufficient" configurations. This is analytically artificial in the context of QCA, because QCA is primarily about minimality (the algorithm behind its top-down approach is called "Boolean minimization"). QCA solutions are inherently minimal by design: the true purpose is to find the minimal set of conditions that are sufficient for the outcome, while preserving the property of robust sufficiency. In other words, not minimality for the sake of minimality, but minimality that does not omit necessary conditions in the minimization process.

From this perspective, a proper QCA solution (particularly one derived via conservative or intermediate procedures) should satisfy both requirements simultaneously: it must retain all empirically necessary conjuncts and eliminate all empirically irrelevant ones, such that the final solution must be both minimal and sufficient. If it is not minimal, it has not been fully reduced; if it omits a necessary condition, it ceases to be sufficient. The idea that these two properties should be separated misunderstands the internal logic of QCA minimization, which aims to preserve sufficiency while removing only empirically irrelevant components.

In continuing his arguments, Baumgartner (2021), strongly assert the "guaranteed" aspect is impossible with fuzzy sets data, where the outcome doesn't "*always*" but only partially occurs. Still, it should not be forgotten that QCA does not minimize fuzzy data but a crisp truth table (Ragin 2004). In this fundamental object, configurations display the output either happening or not, regardless of its full or partial consistency with the outcome in the raw / calibrated (fuzzy) data.

Aggregating the conclusions from all of these articles, Haesebrouck and Thomann (2021) were quick to observe that identifying robust sufficient configurations is difficult because it requires an ex-ante knowledge of the full causal structure of an event. While this is true in the world at large, they too overlooked the most important point of the simulation exercise by Baumgartner and Thiem, who evaluated the different QCA solutions against a *known* causal structure.

Robust sufficiency was not designed as a superior definition of sufficiency, it simply formalizes a practical safeguard in QCA: do not eliminate conditions that are necessary for the outcome to occur in the observed data. While it cannot be directly validated in empirical data, it becomes testable and meaningful in simulation-based studies where the full causal model is known.

This section has shown that empirical sufficiency, particularly when inferred from minimized solutions, can yield misleading results, can misrepresent causal relevance if necessary conditions are dropped or if the dataset is incomplete. To address this vulnerability, the notion of robust sufficiency introduces a stricter standard: a solution should not only be logically sufficient but also retain all causally necessary components.

Among QCA's three solution types, the conservative solution is the only one that satisfies this requirement by construction: it minimizes solely over observed positive configurations and excludes all remainders from the minimization process. Within the scope of the conservative solution and empirical evidence from the positive output configurations, the outcome is *guaranteed* to occur, because it *always* occurs.

However, this very feature has also led to widespread misconceptions. It is often assumed that the conservative solution treats all unobserved or non-positive configurations as negative: that is, as incompatible with the outcome. This assumption is incorrect. The logic and mechanics of the conservative solution are more subtle, and their clarification is essential for understanding what robust sufficiency entails in practice. The next section addresses these issues in detail.

1.3 The logic of the conservative solution

Among the three QCA solution types, the conservative solution stands out for its methodological restraint: it excludes all counterfactual assumptions and minimizes exclusively over empirically observed configurations where the outcome is present. As such, it preserves all empirically necessary conditions and completely bypasses the risk of over-minimization. For this reason, it is the only solution type that is guaranteed, within the limits of the observed data, to retain robust sufficiency.

However, the conservative solution is also widely misunderstood. A recurring misconception is that it implicitly treats all other configurations (both unobserved and observed negative ones) as incompatible with the outcome. This misunderstanding has led to erroneous claims that the conservative solution "sets all remainders to false," or that it artificially inflates empirical conservatism by assuming universal absence outside the observed configurations.

This section clarifies what the conservative solution does and does not assume. It shows how the conservative solution treats unobserved configurations as unknown (not as negative) and demonstrates that its logic is internally consistent, empirically grounded, and indispensable for evaluating the retention of necessary conditions.

It is often assumed that the conservative solution operates under a strong assumption: that all configurations not associated with the outcome, whether they are unobserved remainders or observed negative cases, should be treated as incompatible with the outcome. This assumption is incorrect. The conservative solution does not assign negative values to unobserved configurations; it simply omits them from the minimization process. In fact, none of the QCA solution types, including the parsimonious, incorporate negative configurations during the Boolean minimization. All solution types rely solely on configurations with outcome value 1, and differ only in how they treat the logical remainders.

The conservative solution proceeds by minimizing over the subset of the truth table where the outcome is present, excluding both remainders and observed negative configurations from the minimization process. This procedure avoids the risk of introducing unverified assumptions about what might happen in unobserved parts of the configuration space.

Confusion has arisen in part due to the terminology used in the QCA literature. For example, Cooper and Glaesser (2016, p. 303) quote Ragin (2008) as saying the conservative solution "effectively sets all the remainders to 'false'." However, Ragin's original phrasing is more precise: "the complex solution defines all remainder combinations as false" (Ragin 2008, p. 173). While the wording is subtle, the difference is substantive. To "set" remainder configurations to false implies that they are explicitly coded with a 0 outcome value, an action the algorithm does not perform. Instead, Ragin's point is that these remainders are excluded from the minimization process, and are thereby treated as unknown.

This distinction is also reflected in the fs/QCA software interface, as depicted in Fig. 1. When specifying which configurations to include in the minimization, users can assign them as *True* (included), *False* (excluded), or *Don't Care* (left to algorithmic discretion). Assigning *False* to a configuration does not mean it has an outcome value of 0, but only that it is not used in the derivation of prime implicants.

In effect, there is no need whatsoever to set all remainder configurations with an output value of 0, neither explicitly in Ragin's fs/QCA software, nor implicitly ("through the back

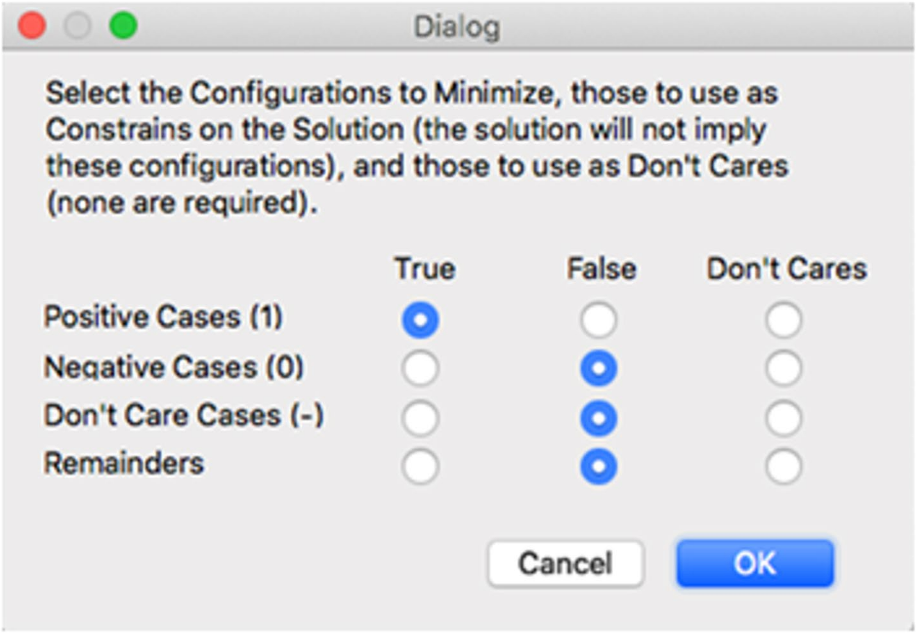

Fig. 1 "Specify Analysis" dialog in the fs/QCA software

door"). This can be verified in any open source minimization software that implements the classical Quine-McCluskey algorithm.

Besides, *setting* all remainders to false is most of the times not even possible. In the engineering field, it is customary to have dozens and sometimes hundreds of concomitant electrical inputs (the equivalents of the QCA explanatory conditions). It is literally impossible to actually produce a complete truth table for 50 conditions having over 10^{15} rows. Even for 30 conditions the truth table contains over 1 billion rows, and no software requires such a huge truth table to produce a solution. Such a procedure would very soon run out of available computer memory, given the exponential growth of the truth table.

This logic is consistent with both classical and modern implementations of Boolean minimization. Engineering tools like Espresso (Brayton et al. 1984) distinguish explicitly between observed positive configurations, observed negative configurations, and don't care conditions (remainders). The remainders are marked with a dash and excluded from influencing the outcome logic. QCA software does the same.

Even if one were to assign a value of 0 to all remainder configurations, this would only produce the same conservative solution. But doing so is not required, and in practical applications, often infeasible.

More recently, Haesebrouck and Thomann (2021) describe the conservative solution under the simple assumption that unobserved configurations are insufficient for the outcome. While this interpretation is more cautious, it still misrepresents the actual mechanism. The conservative solution makes no assumption about the sufficiency or insufficiency of unobserved configurations; rather, it brackets them entirely. The algorithm treats them as unknown, thereby refraining from assigning any causal meaning to them. This point is crucial because it distinguishes omission from negation: the former is a deliberate methodological choice aimed at avoiding unfounded generalizations, while the latter would require a substantive empirical or theoretical claim.

A clearer formulation is offered by Schneider and Wagemann (2012, p. 188), who write: “The Standard Analysis procedure allows for different solution terms to be produced depending on whether assumptions on logical remainders are made. If no assumptions are made, the conservative solution is produced”. This description aligns precisely with the intent of the conservative solution: it does not rely on assumptions about unknown configurations and thereby maximizes empirical caution. Still, this clarity is not always reflected in secondary discussions, which often conflate excluding a configuration from minimization with assigning it a 0 outcome value.

This section concludes that the conservative solution is not based on hidden assumptions or rigid negations. It is a product of empirical minimalism: it minimizes over what is observed, assumes nothing about what is not, and yields a solution that (by construction) retains all necessary conditions that co-occur with the outcome. This is precisely what makes it robustly sufficient, and uniquely valuable when empirical integrity is prioritized over logical elegance.

1.4 Robustness vs. redundancy in solution criteria

The contrast between QCA and CNA is often portrayed in terms of algorithmic differences, but a more fundamental divergence lies in what each method treats as its primary evaluative criterion. QCA, particularly in its conservative and intermediate forms, prioritizes robust sufficiency: that is, preserving configurations that reliably instantiate the outcome across all observed cases. CNA, by contrast, emphasizes redundancy-free modeling, seeking to identify only those conditions that are not logically or causally superfluous.

These divergent priorities result in different optimization goals. While QCA tolerates some redundancy if it ensures that no necessary condition is lost, CNA minimizes aggressively to eliminate all irrelevance, even at the risk of discarding necessary elements. This trade-off can be formalized through three stylized scenarios. Considering a known causal structure consisting of two disjunctive configurations: $A \cdot B + C \cdot D$, these scenarios are:

1. A solution that includes $A \cdot B \cdot E + C \cdot D \cdot F$: robustly sufficient but not fully minimal, where irrelevant conditions (E, F) remain.
2. A solution that recovers exactly $A \cdot B + C \cdot D$: both robustly sufficient and redundancy-free, ideal for both QCA and CNA.
3. A solution that reports $A + C \cdot D$: redundancy-free, but over-minimized since B is a necessary component omitted from the conjunction with A.

Scenario (2) is ideal, but often unattainable in practice. Scenario (1) is typically preferred in QCA, as it guarantees outcome consistency even if parsimony is compromised. Scenario (3) is favored by CNA, which considers it “correct” under its redundancy criterion, even though the outcome is not instantiated when A occurs without B.

The difference becomes critical when interpreting partial solutions. In QCA, dropping a condition like B from $A \cdot B$ results in loss of robust sufficiency: A may no longer align consistently with the outcome. In CNA, however, A alone may still appear in the model if B's necessity is not formally enforced.

Statements such as those by Haesebrouck and Thomann (2021, p. 12) that "causally relevant conditions might be removed from sufficient combinations" are conceptually problematic and self contradicting. If a condition is causally relevant, then removing it renders the combination no longer sufficient. Conversely, if sufficiency is preserved, the removed condition was not causally relevant to begin with. Both cannot hold simultaneously.

In real-world applications, the full causal structure is rarely known. This makes it difficult to adjudicate between these approaches using empirical data alone. Simulation studies, where the true model is known, reveal that redundancy-free models often over-minimize and fail to retain all necessary conditions. In contrast, robust sufficiency, even when not fully minimal, tends to preserve the causal architecture.

The preference for robustness over redundancy is grounded in the principle that causal claims should not exceed what the data supports. In complex configurational settings, this often means tolerating some irrelevance in order to avoid losing what is essential.

The following section examines how this trade-off manifests in logic-based minimization and explores the potential pitfalls of over-minimization, including the so-called "empty set" paradox.

1.5 The empty set paradox and the limits of parsimony

A recurring theme in the debate between QCA and CNA concerns the degree to which minimization should be pursued. CNA favors maximal parsimony, aiming to eliminate all redundancy. QCA, by contrast, accepts a trade-off: some redundancy may be tolerable if it ensures the preservation of necessary conditions. This section examines the risks associated with excessive minimization, including a rarely acknowledged consequence: the so-called "empty set" paradox.

In QCA, the parsimonious solution is always a logical reduction of the conservative solution. The conservative solution is derived strictly from observed configurations where the outcome is present; the parsimonious solution introduces counterfactuals to simplify those configurations. Since every condition in the parsimonious solution must already appear in the conservative one, the latter is always a logical superset of the former.

This relationship holds unless the minimization process eliminates conditions that are actually necessary. When that happens, the resulting solution no longer guarantees the presence of the outcome. An extreme illustration of this occurs in Baumgartner and Thiem (2017)'s simulation study, where one of the derived solutions is the empty set (i.e., a solution with no causal conditions at all). Under their correctness criterion, based on redundancy-free sufficiency, the empty set is considered "correct" because it does not contain any causally irrelevant conditions. But this result is counterintuitive: the empty set cannot explain anything. It neither contains causal information nor distinguishes between presence and absence of the outcome.

Formally, the empty set is a subset of every set. But sufficiency in QCA is not just a matter of logical subset relations, it also requires that the outcome consistently occurs when the condition (or configuration) is present. Since the empty set contains no conditions, it cannot satisfy this requirement. It cannot be deemed sufficient for anything unless the outcome occurs in all possible configurations, a case of trivial sufficiency that rarely applies in practice.

Moreover, while it is logically correct that the empty set is a subset of any other set, this fact is irrelevant in the context of QCA minimization. In QCA, a valid solution must cover all observed positive configurations, it must be a superset of those configurations. The paradox lies precisely here: the empty set, by definition, contains no conditions and therefore cannot be a superset of any non-empty configuration. To regard the empty set as a valid solution is to mistake syntactic permissibility for substantive adequacy. A solution that fails to cover the cases it is meant to explain cannot be deemed sufficient in any empirical sense.

This paradox highlights the risks of purely formal criteria. A redundancy-free model may be logically valid but empirically vacuous. It may satisfy a minimalist standard while failing to preserve any meaningful link to the outcome. Such outcomes arise from applying propositional logic mechanically, without anchoring it in observed empirical structure.

If the principles of propositional logic were applied in legal reasoning, the consequences would be absurd: individuals could be declared guilty merely because their innocence could not be proven. But legal systems based on the rule of law require positive evidence of wrongdoing, not merely the absence of disproof. The burden is on demonstrating the presence of causal responsibility, not assuming it by default.

Propositional logic imposes a different standard. It defines sufficiency in purely syntactic terms and treats absence of contradiction as tantamount to truth. But conclusions derived from data alone, without theoretical guidance, are highly susceptible to error. This concern is well known in the quantitative social sciences, where "fishing for significance" is a longstanding methodological warning. Sound research practice insists that hypotheses be formulated prior to data collection, and that data be used to test them, not generate them post hoc.

This problem is analogous to debates in other fields, such as machine learning, where purely data-driven models may find spurious patterns that do not generalize. The distinction is between internal consistency and external validity. Logical sufficiency is not enough; it must be constrained by the empirical regularities it purports to explain.

In sum, the empty set paradox exemplifies the danger of over-minimization. Eliminating all redundancy at the expense of causal relevance defeats the explanatory purpose of the method. QCA's allowance for empirical redundancy is not a flaw but a safeguard. It ensures that the outcome is grounded in actual causal structures (that are empirically instantiated and meaningfully covered) not merely in the formal possibility of contradiction avoidance.

1.6 The role of counterfactuals in QCA solutions

The intermediate solution in QCA was introduced to address the tension between preserving empirical robustness and achieving parsimony. While the conservative solution avoids all counterfactuals and the parsimonious solution embraces all plausible ones, the intermediate solution allows for the selective inclusion of remainders based on theoretical expectations. First proposed by Ragin and Sonnett (2005) and further formalized by Schneider and Wagemann (2012) through their Enhanced Standard Analysis (ESA), the intermediate solution provides a middle ground that aims to balance empirical caution with theoretical relevance.

Counterfactuals in QCA refer to unobserved configurations (remainder rows) that could, in principle, occur but are not represented in the data. The inclusion of such configurations in the minimization process requires assumptions about how the outcome would behave if they were present. The conservative solution avoids such assumptions entirely; the parsimonious solution includes all remainders without restriction. The intermediate solution filters remainders, including only those that align with theoretical expectations, generating the so-called "easy counterfactuals."

A common misconception is that the intermediate solution is derived by simply omitting problematic remainders from the minimization process. In fact, the ESA modifies the solution by reclassifying certain remainders as negative cases based on theoretical reasoning, when there is proof the outcome *cannot* happen. This changes the truth table and thus the resulting set of prime implicants. It is not a matter of partial minimization over a filtered set, but rather a transformation of the input based on theory-informed decisions.

Certain causal configurations are logically or physically impossible and therefore cannot be observed. In such cases, the outcome cannot be instantiated either. For example, reading in total darkness is not feasible under normal circumstances. If light is a trivial but necessary condition for reading, then any unobserved configuration in which light is absent can legitimately be assigned an outcome value of zero. These are not speculative assumptions but grounded in background knowledge about causal impossibility.

Incorporating such knowledge enhances the precision of the resulting solution, not by forcing it toward absolute minimality but by ensuring it is as minimal as logically possible, given what is known. The goal is to combine observed data with theoretically justified inferences, thereby producing solutions that are both logically parsimonious and robustly sufficient.

In their formulation of an atomic minimal theory, Baumgartner and Ambühl (2018) argue for the superiority of the parsimonious solution as the only one that yields a minimally necessary disjunction of minimally sufficient conditions. However, this view is incomplete. As Baumgartner (2013, p. 90) himself notes, a minimally sufficient configuration must instantiate the outcome it is meant to explain: "…E is instantiated in the same situation as its minimally sufficient conditions…".

This condition is not optional. If the outcome E fails to occur when its purported sufficient conditions are present, then the configuration, however minimal or redundancy-free, cannot be deemed sufficient. Sufficiency entails instantiation. A logically clean solution that fails to instantiate the outcome is, by definition, insufficient.

This procedure of identifying difficult, impossible or untenable counterfactuals, while more demanding, offers a more disciplined approach to incorporating theoretical knowledge. It neither presumes universal sufficiency (as in the parsimonious solution) nor limits itself to observed configurations (as in the conservative one). Instead, it tests which assumptions can be made on theoretical grounds and recalibrates the minimization accordingly.

The enhanced parsimonious solution represents a knowledge-augmented variant of the standard parsimonious solution, with direct implications over the intermediate solution. This middle ground solution type, introduced by Ragin and Sonnett (2005), is constructed by comparing the conservative and parsimonious solutions and integrating theoretically justified counterfactuals. The process is formalized in the Enhanced Standard Analysis, which refines the minimization procedure in four steps:

1. Identify the set of Simplifying Assumptions, remainders that are actually used during minimization.
2. Evaluate these assumptions against directional expectations derived from theory.

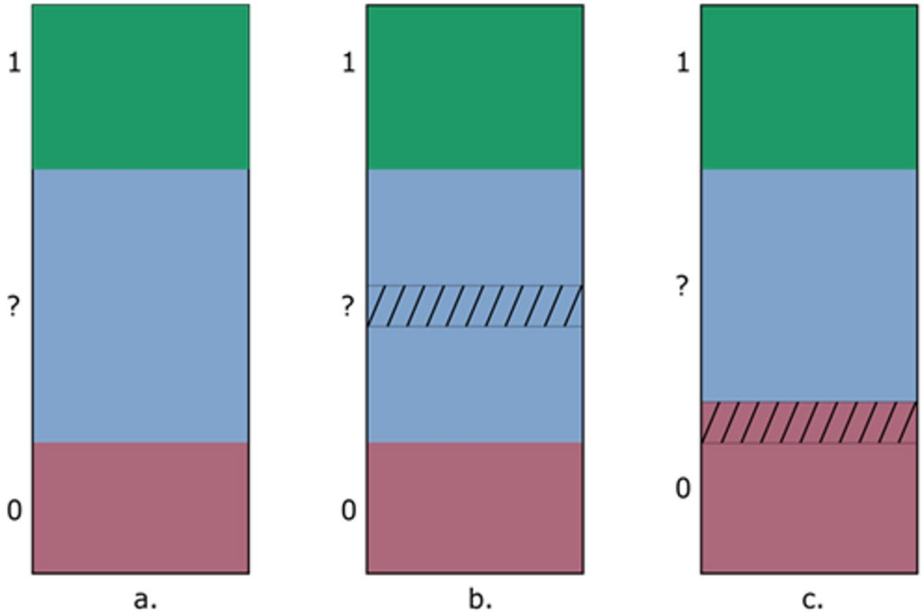

Fig. 2 Truth table configurations

3. Retain only those that are consistent with the theoretical expectations, thus generating the Easy Counterfactuals.
4. Add the Easy Counterfactuals to the set of observed positive configurations and apply a classical Quine-McCluskey minimization to the augmented truth table.

This approach allows the intermediate solution to incorporate theoretical expectations into the solution logic, thereby distinguishing plausible remainders from those that lack empirical or conceptual justification.

Figure 2 illustrates the structure of a standard truth table, including three types of configurations: observed positive cases (coded 1), observed negative cases (coded 0), and remainders (?). This classification is essential to understanding how the ESA operates. By selectively reclassifying remainders as negative cases (based on theory rather than assumptions) the ESA modifies the set of configurations used for minimization.

Critics of the ESA, such as Cooper and Glaesser (2016), have argued that it introduces additional complexity and opacity into the analysis. While this is not unfounded, it misunderstands the purpose of the ESA. The added complexity is not arbitrary; it reflects the epistemological challenge of balancing parsimony with theoretical plausibility. The ESA clarifies rather than obscures, provided its procedural logic is properly understood.

What actually occurs is illustrated in Fig. 2c: the remainders deemed incompatible with theoretical expectations are not simply discarded, but reassigned as negative configurations. This reclassification modifies the structure of the truth table and yields a different, enhanced parsimonious solution. The change in the intermediate solution does not arise from excluding remainders from the pool of easy counterfactuals, but rather from altering the logic of the superset-subset relation. The new set of easy counterfactuals is derived from a revised parsimonious solution, one that now reflects both empirical data and theoretically grounded reclassifications. This updated relation to the conservative solution is what ultimately defines the intermediate solution.

Despite its rigor, this procedure is surprisingly often misunderstood. One persistent misconception is that the ESA merely excludes certain remainders from the minimization process. In fact, it alters the truth table structure itself by modifying which configurations are treated as positive or negative inputs. This misunderstanding contributes to persistent confusion not only about how the intermediate solution is derived, but also about its epistemological position between conservative caution and theoretical ambition.

The intermediate solution exemplifies a broader point: QCA, unlike CNA, accommodates both empirical and theoretical input. It recognizes that causal inference in the social sciences cannot be fully automated through formal logic. Theory matters, not only in interpreting results, but in shaping the inferential process itself.

The following section addresses this broader point more directly by discussing the role of prior knowledge and its implications for within-case understanding and causal inference.

1.7 The epistemic role of prior knowledge

Causal inference in configurational methods relies not only on identifying patterns in the data but also on understanding the context in which those patterns emerge. As Ragin (1987) describes, this process is a “dialogue between ideas and evidence”, a dynamic interplay where provisional insights guide further engagement with the data, including calibration refinements and reinterpretations of case knowledge.

In contrast, approaches that privilege redundancy-free models often discard the data once a solution is found, as if no further interaction with the empirical material is necessary. This view assumes both that the dataset perfectly reflects reality and that no interpretive depth is gained by re-examining cases. Yet, the essence of causal understanding often lies precisely at the point where conclusions invite further scrutiny.

A condition is not deemed necessary simply because it shows high consistency. It requires theoretical justification. The search for necessary conditions without reference to theory risks degenerating into data mining, a concern long noted in critical realism and quantitative methodology alike. Identifying empirical regularities alone is insufficient when the goal is to explain why outcomes occur.

QCA’s framework is consistent with regularity theories of causation but not reducible to them. Its minimization algorithm can indeed identify Boolean difference makers, much like an unsupervised machine learning routine. But stronger results are obtained through a supervised logic: theory must precede data exploration and continue to inform it throughout. Theory is not an afterthought but a precondition for meaningful inference.

This logic parallels insights from quantitative research. In analyzing large samples, researchers do not generalize from thousands of discrete cases but from the single population they represent. In comparative designs, cases are often entire populations, and differences among them require deep knowledge of their structure and history. The same holds in CCMs, where the explanatory power of a model depends not just on its coverage or consistency, but on the researcher’s understanding of each case’s internal logic.

Within-case knowledge is essential. A regularity-theoretic algorithm may flag Boolean difference makers, but it is only by situating those differences within a case-specific context that one can determine whether they reflect genuine causal mechanisms. This is why post-QCA process tracing, as proposed by Rohlfing and Schneider (2013); Rohlfing and Schneider (2018); Rohlfing (2014), is an important complement to cross-case analysis.

Prior knowledge also informs how we treat unobserved configurations. In QCA, remainders can be included or excluded from minimization depending on whether they are considered theoretically plausible. Those included are treated as if they were observed with a positive outcome, paving the way for the Standard Analysis. Those excluded are treated as if they were observed with a negative one, leading to the Enhanced Standard Analysis, with a different (enhanced) parsimonious solution. This practice reflects a dual reliance on empirical regularities and counterfactual reasoning, an integration that QCA supports explicitly and that CNA, by design, avoids.

In sum, robust configurational inference depends not only on the data, but on the researchers' ability to link patterns to plausible causal narratives. The regularities we observe are meaningful only when anchored in what we already know, or are willing to investigate further, about the cases they describe.

2 Concluding remarks

This paper has examined key differences between QCA and CNA, not only in terms of algorithmic procedures but more fundamentally in their assumptions about sufficiency, minimization, and the role of theory. It has argued that much of the disagreement between the two approaches stems from incompatible standards: QCA emphasizes empirical robustness and theoretical integration (as advocated by Ragin 2000, 2008; Schneider and Wagemann 2012), while CNA prioritizes redundancy elimination and artificial parsimony, extensively using Mackie's (1974) regularity theory, which Schneider (2018) categorised as ideal in theory but almost never achieved in reality.

A central theme has been the status of the conservative solution. Contrary to claims that it is methodologically inferior, the conservative solution is the only one that strictly avoids counterfactual assumptions and minimizes over the observed, outcome present configurations. As such, it offers the most empirically cautious strategy and is the only solution that, by construction, retains all necessary conditions observed in the data. Mischaracterizations of this solution, such as the assumption that all remainders are treated as negative, have obscured its role in ensuring robust sufficiency.

The intermediate solution, introduced to reconcile empirical caution with theoretical insight, also plays a crucial role. Through the Enhanced Standard Analysis, it incorporates theory-informed remainders without resorting to blanket counterfactuals. This allows QCA to navigate between overfitting and over-minimization, a capacity absent in CNA, which lacks an equivalent mechanism for integrating prior knowledge.

The discussion of the empty set paradox has highlighted the risks of formal correctness criteria that disregard empirical grounding. Redundancy-free models may be logically coherent yet fail to preserve any causal meaning when necessary conditions are discarded. In such cases, the gain in parsimony is offset by a loss in explanatory adequacy.

Baumgartner (2021) contends that the conservative and intermediate solutions are methodologically inferior to the parsimonious solution because they depend on substantive interpretability and robust sufficiency standards that, in his view, are ill-suited to real-world data, which is inherently incomplete and fragmented.

It is true that empirical data rarely, if ever, offer a complete representation of the underlying causal structure. But this observation does not discredit robust sufficiency; rather, it explains its purpose. Robust sufficiency was never meant to guarantee that the outcome always occurs in all conceivable contexts. Instead, it asserts that the outcome consistently occurs within the boundaries of the observed configurations: that is, within the empirical scope of the truth table.

The validity of this standard is not metaphysical but methodological. When working with real, noisy, or partial data, the analyst cannot know in advance which conditions are causally necessary. However, in simulated contexts (where the true causal structure is known) robust sufficiency becomes testable. It provides a criterion for evaluating whether minimization procedures preserve or eliminate essential causal components.

As demonstrated in Dușa (2019), the parsimonious solution promoted by CNA adheres to a stricter, propositional logic-based definition of sufficiency. It values syntactic minimality and redundancy elimination, even when this entails removing conditions that are, in fact, causally relevant. The consequence is a form of over-minimization that may yield solutions that are formally valid but empirically inadequate. Robust sufficiency, by contrast, is designed to avoid this outcome. It ensures that simplification does not come at the cost of losing what matters most: the capacity to reproduce the outcome under observed causal configurations.

It offers an alternative evaluative standard. It does not demand universal guarantees, but requires that minimized solutions preserve the empirical relationships observed in the data, especially those that instantiate the outcome. This concept is not an abstract ideal but a practical response to the challenges of working with incomplete, noisy, or fragmentary data.

Ultimately, the choice between QCA and CNA is not just a matter of software or syntax, but of epistemological stance. CNA's approach is efficient but risks over-minimization. QCA's approach is more cautious, aiming to eliminate redundancy only insofar as it does not compromise empirical adequacy or theoretical plausibility.

Social phenomena rarely occur in controlled environments free of irrelevant conditions. Real-world outcomes emerge from complex interactions, many of which include background noise. A method that discards relevant information for the sake of syntactic elegance risks misrepresenting the very complexity it aims to explain.

QCA was designed not only to detect causal patterns but, perhaps more importantly, to support causal understanding. That aim requires more than identifying regularities. It demands interpretive judgment, theoretical insight and, above all, a willingness to engage with the cases themselves.

Purely data-driven solutions grounded in empirical regularities can be informative, but they remain theoretically incomplete. While such models may identify surface-level patterns, they cannot by themselves distinguish spurious correlations from meaningful causal relationships. This is why QCA's long-standing integration of both regularity and counterfactual reasoning is not an ad hoc addition, but a methodological strength. Combining these perspectives allows researchers to approximate more closely the underlying processes that generate the observed outcomes.

Moreover, the analytic task does not end with the derivation of a formal solution. As Pattyn et al. (2020) rightly emphasize, causal explanation requires understanding not just how something happens, but why it happens. This explanatory depth cannot be achieved by algorithms alone, no matter how efficient.

Causal inference ultimately requires interpretive judgment. Algorithms can identify difference-makers, but they cannot explain their meaning. This part of the research process remains firmly rooted in human reasoning, drawing on theoretical knowledge, contextual insight, and methodological flexibility to make sense of complexity in ways that computational procedures cannot.

Funding The author declares that no funds, grants, or other support were received during the preparation of this manuscript.

Declarations

Competing interests The author has no relevant financial or non-financial interests to disclose.

References

Baumgartner, M.: 'Qualitative Comparative Analysis and robust sufficiency', *Quality* & *Quantity*. Springer Science; Business Media LLC. (2021). https://doi.org/10.1007/s11135-021-01157-z

Baumgartner, M., Ambühl, M.: Causal modeling with multi-value and fuzzy-set coincidence analysis. In: Political Science Research and Methods, vol. 8, pp. 526–542. Cambridge University Press (CUP) (2018). 310.1017/psrm.2018.45

Baumgartner, M.: A Regularity Theoretic Approach To Actual Causation', *Erkenntnis*, vol. 78, pp. 85–109. Springer Science (2013). S110.1007/s10670-013-9438-3 Business Media LLC

Baumgartner, M., Thiem, A.: Often trusted but never (Properly) tested: Evaluating qualitative comparative analysis. In: *Sociological Methods* & *Research*, vol. 49, pp. 279–311. SAGE (2017). 210.1177/0049124117701487

Brayton, R.K., Hachtel, G.D., McMullen, C.T., Sangiovanni-Vincentelli, A.L.: *Logic minimization algorithms for VLSI synthesis–1984*. Edited by R. K. Brayton, G. D. Hachtel, C. T. McMullen, and A. L. Sangiovanni-Vincentelli. Springer US (The kluwer international series in engineering and computer science). (1984). https://doi.org/10.1007/978-1-4613-2821-6

Cooper, B., Glaesser, J.: Qualitative comparative analysis, necessary conditions, and limited diversity: Some problematic consequences of Schneider and wagemann's enhanced standard analysis. Field Methods. **28**(3), 300–315 (2016). https://doi.org/10.1177/1525822X15598974

Dușa, A.: Consistency cubes. Fast, A., efficient method for exact boolean minimization. R J. **10**(2), 357–370 (2018). https://doi.org/10.32614/RJ-2018-080

Dușa, A.: *QCA with R. A Comprehensive Resource*. Cham, Switzerland: Springer International Publishing. (2019). https://doi.org/10.1007/978-3-319-75668-4

Forner, D., Ozcan, S., Bacon, D.: 'Machine learning approach for national innovation performance data analysis', in *Proceedings of the 8th International Conference on Data Science, Technology and Applications*. SCITEPRESS - Science; Technology Publications. (2019). https://doi.org/10.5220/0007953603250331

Haesebrouck, T., Thomann, E.: 'Introduction: Causation, inferences, and solution types in configurational comparative methods', *Quality* & *Quantity*. Springer Science; Business Media LLC. (2021). https://doi.org/10.1007/s11135-021-01209-4

Mackie, J.L.: *The Cement of the Universe: A Study of Causation*. 2002 reprinted. Oxford, UK: Clarendon Press. (1974)

Pattyn, V., Álamos-Concha, P., Cambre, B., Rihoux, B., Schalembier, B.: 'Policy Effectiveness through Configurational and Mechanistic Lenses: Lessons for Concept Development', *Journal of Comparative Policy Analysis: Research and Practice*. Informa UK Limited. (2020). https://doi.org/10.1080/13876988.2020.1773263

Ragin, C.C.: The Comparative Method. University of California Press, Berkeley, Los Angeles, London (1987)

Ragin, C.C.: *Fuzzy-Set Social Science*. 1st edn. Edited by C. C. C. Ragin. University of Chicago Press. (2000)

Ragin, C.C.: 'From Fuzzy Sets to Crisp Truth Tables', *COMPASSS Working Papers Series*. (2004). Available at: http://www.compasss.org/wpseries/Ragin2004.pdf

Ragin, C.C.: Redesigning Social Inquiry. Fuzzy Sets and Beyond. University of Chicago Press, Chicago; London (2008)

Ragin, C.C., Sonnett, J.: Between complexity and parsimony: Limited diversity, counterfactual cases, and comparative analysis. In: Kropp, S., Minkenberg, M. (eds.) Vergleichen in Der Politikwissenschaft, pp. 180–197. VS Verlag für Sozialwissenschaften, Wiesbaden (2005). https://doi.org/10.1007/978-3-322-80441-9

Rohlfing, I.: Comparative hypothesis testing via process tracing. Sociological Methods Research. **43**(4), 606–642 (2014). https://doi.org/10.1177/0049124113503142

Rohlfing, I., Schneider, C.Q.: Improving research on necessary conditions: Formalized case selection for process tracing after QCA. Polit. Res. Q. **66**(1), 220–235 (2013). https://doi.org/10.1177/1065912912468269i

Rohlfing, I., Schneider, C.Q.: A unifying framework for causal analysis in Set-Theoretic multimethod research. Sociological Methods Research. **47**(1), 37–63 (2018). https://doi.org/10.1177/0049124115626170

Schneider, C.Q.: Realists and idealists in QCA. Political Anal. (2018). https://doi.org/10.1017/pan.2017.45

Schneider, C., Wagemann, C.: Set-Theoretic Methods for the Social Sciences. A Guide To Qualitative Comparative Analysis. Cambridge University Press, Cambridge (2012)

Swiatczak, M.D.: 'Different algorithms, different models', *Quality* & *Quantity*. Springer Science; Business Media LLC. (2021). https://doi.org/10.1007/s11135-021-01193-9