

Boolean Difference-Making: A Modern Regularity Theory of Causation

forthcoming in *The British Journal for the Philosophy of Science*, doi: 10.1093/bjps/axz047

Michael Baumgartner and Christoph Falk

Abstract

A regularity theory of causation analyses type-level causation in terms of Boolean difference-making. The essential ingredient that helps this theoretical framework overcome the problems of Hume's and Mill's classical accounts is a principle of non-redundancy: only Boolean dependency structures from which no elements can be eliminated track causation. The first part of this paper argues that the recent regularity theoretic literature has not consistently implemented this principle, for it disregarded an important type of redundancies: structural redundancies. Moreover, it is shown that a regularity theory needs to be underwritten by a hitherto neglected metaphysical background assumption stipulating that the world's causal makeup is not ambiguous. Against that background, the second part then develops a new regularity theory that does justice to all types of redundancies and, thereby, provides the first all-inclusive notion of Boolean difference-making.

1 *Introduction*

2 *Fundamentals*

3 *Structural Redundancies*

4 *Permanence*

5 *Ambiguities*

6 *A New Regularity Theory*

7 *Discussion*

Appendix

1 Introduction

Theories of causation come in many variants, many of which are incompatible. According to some, causation is deterministic, while according to others it is not; some theories take difference-making to be the characteristic feature of causation, others opt for powers or dispositions; some yield that causation is an extrinsic property, according to others it is intrinsic; and so on. Conflicting theories continue to co-exist because they are embedded in, and draw their justification from, incompatible background metaphysics, which are notoriously difficult to reconcile and which, typically, are taken for granted in discussions about causation. Hence, without claiming to be presenting the only or ultimate truth about causation, this paper develops a modern regularity theory of causation.

Regularity theories are embedded in the metaphysical tradition of Humean actualist anti-necessitarianism (Hume [1748], Section 7), according to which there is no causal oomph; rather, causation, possibility, and lawhood supervene on the actual distribution of matters of fact, which itself is a brute fact. Causal laws are convenient summaries of the regularities that happen to emerge from that distribution. Correspondingly, being in accordance with those laws, that is, being empirically possible is a matter of existing (in an atemporal sense) in the actual world. Plainly, as all metaphysical frameworks, actualist anti-necessitarianism is controversial. This paper, however, is not the place to enter that controversy. Its main objective is not metaphysical but pragmatic: to provide a conceptual fundament for the currently spreading *configurational comparative methods* (CCMs) of causal data analysis.¹ We take the anti-necessitarian background to be sufficiently justified if it yields an account of causation that conceptually underwrites CCMs.

CCMs differ from other techniques as regression analytical methods (RAMs; Gelman and Hill [2007]) or Bayes-nets methods (BNMs; Spirtes *et al.* [2000]) in a number of respects (for a discussion of some of these differences see Thiem *et al.* [2016]). Most importantly for our current purposes, while RAMs and BNMs search for causal dependencies among

¹ Qualitative Comparative Analysis (QCA; Rihoux and Ragin [2009]; Thiem [2014]) and Coincidence Analysis (CNA; Baumgartner and Ambühl [2018]) are paradigmatic CCMs. QCA, in particular, has been applied in hundreds of studies, mainly in the social sciences (see the website of the COMPASSS network: <www.compassss.org>).

variables by exploiting their statistical (in-)dependencies, CCMs search for causal dependencies among concrete *values of variables* by exploiting Boolean dependencies as ‘ $A=\alpha_i$ is sufficient/necessary for $B=\beta_i$ ’. To this end, CCMs must be underwritten by a theory of causation that provides a link between Boolean dependencies and causation. This is exactly the field of expertise of regularity theories.

The primary analysandum of regularity theories is causation on the type level, that is, causal relevance relations between variables or factors taking on specific values: ‘ $A=\alpha_i$ is causally relevant to $B=\beta_i$ ’, where $A=\alpha_i$, for instance, stands for (the event type of) malfunctioning traffic lights and $B=\beta_i$ for occurring rear-end collisions. (We will use the terms ‘variable’ and ‘factor’ interchangeably in this paper.) Regularity theories take *difference-making* to be the characteristic feature of causation. What that amounts to can easily be specified in causal terms: $A=\alpha_i$ is a difference-maker of $B=\beta_i$ iff there exist (at least) two scenarios σ_1 and σ_2 such that $A=\alpha_i$ is associated with $B=\beta_i$ in σ_1 and $A\neq\alpha_i$ with $B\neq\beta_i$ in σ_2 while all alternative causes of $B=\beta_i$ are absent in σ_1 and σ_2 , where alternative causes of $B=\beta_i$ are causes located on a causal path to $B=\beta_i$ that does not go through $A=\alpha_i$. However, as regularity theories aim for a *reductive* analysis of causation, they cannot define ‘ $A=\alpha_i$ is causally relevant to $B=\beta_i$ ’ with recourse to the absence of alternative *causes* of $B=\beta_i$. Instead, the difference-making requirement must be captured in terms of non-causal (i.e. Boolean) dependencies, which, as we shall see below, calls for imposing constraints not only on dependence relations between pairs of factor values but also on whole dependency structures. It follows that a causal relation between $A=\alpha_i$ and $B=\beta_i$ does not supervene on intrinsic properties of the (sets of) entities represented by $A=\alpha_i$ and $B=\beta_i$, rather it obtains in virtue of the latter’s function in a whole dependency structure. Finally, regularity theories assume causation to be deterministic.

The primary analysans of regularity theories consists in structures of Boolean dependencies of sufficiency and necessity without redundancies. The principle, originally due to Broad ([1930]) and famously shaped in Mackie’s ([1974]) INUS-theory, that only *redundancy-free* Boolean dependencies track causation, is the essential theoretical ingredient that helped overcome the problems incurred by the classical regularity theories (Hume [1748] and Mill [1843]). To render this principle precise, Graßhoff and May ([2001]) determined

that only *minimally necessary* disjunctions of *minimally sufficient* conditions of scrutinized effects are amenable to a causal interpretation. Syntactically put, causally interpretable Boolean dependencies must be expressible as biconditionals featuring a redundancy-free disjunctive normal form on one side and the scrutinized effect on the other—we shall speak of *RDN-biconditionals*, for short. Baumgartner ([2013]) suggested that this idea could be generalised for the analysis of multi-effect structures by simply conjunctively concatenating atomic RDN-biconditionals with one effect to complex RDN-biconditionals with multiple effects. Baumgartner (implicitly) assumed that by concatenating atomic RDN-biconditionals no new redundancies could be introduced.

The first part of this paper will show that this assumption is false. The non-redundancy principle is not as easily implemented for multi-effect structures as was hoped by Baumgartner ([2013]). Certain RDN-biconditionals in conjunctive sequences of such biconditionals, while internally free of redundancies, can themselves be redundant in the superordinate structure and, as a result, fail to make a difference on the structural level. Hence, what counts as a redundancy-free Boolean dependency structure does not only depend on the minimality of sufficient and necessary conditions but also on the minimality of the conjunctive concatenation of the resulting RDN-biconditionals. That is, the regularity theoretic literature has so far disregarded an important type of redundancies: *structural* redundancies. Furthermore, it will be shown that regularity theories can only consistently capture the difference-making requirement if they are underwritten by a hitherto neglected background assumption stipulating that the causal makeup of the world is not ambiguous.


The second part of the paper then develops a new regularity theory that integrates that metaphysical background, does justice to all types of redundancies, properly generalises the basic idea behind modern regularity theories for multi-effect structures, and, thereby, provides the first all-inclusive notion of Boolean difference-making. To this end, the resulting theory abandons the idea, common to all its regularity theoretic predecessors, that multi-effect structures can be modularly built up from single-effect structures, and replaces it by a form of *causal holism* according to which causation is a holistic property that supervenes on complete distributions of matters of fact and not on proper parts thereof.

2 Fundamentals

A regularity theory assumes that type-level causation is not fundamental but supervenes on actual distributions of matters of fact, that is, on *Humean mosaics* (Lewis [1986], pp. ix-x), which amount to sets of configurations of natural properties coincidentally instantiated by units of observation—events, states of affairs, cases, or whatever other entities the preferred ontology happens to furnish. The problem of rendering the notion of a natural property precise is notoriously difficult. For the purposes of this paper, we bracket it and simply assume that all henceforth analysed properties are natural. Moreover, as is common in the causal modelling literature, we want to remain as non-committal as possible with respect to the ontology of causation and, thus, refer to the causal relata simply as ‘factors taking values’.

Factors represent categorical properties that partition sets of units of observation either into two sets, in case of binary properties, or into more than two (but finitely many) sets, in case of multi-value properties. In the context of CCMs, factors representing binary properties can be crisp-set or fuzzy-set (Thiem [2014]); the former can take on the Boolean identity elements 0 and 1 as possible values, whereas the latter can take on any (continuous) values from the unit interval $[0, 1]$. Factors representing multi-value properties can take on any of an open (but finite) number of non-negative integers as possible values. For simplicity of exposition, we confine ourselves to crisp-set factors in this paper.

The focus on the crisp-set case allows us, for instance, to conveniently abbreviate the explicit ‘Variable=value’ notation, which generates convoluted syntactic expressions with increasing model complexity. As is conventional in Boolean algebra, we write ‘ A ’ for $A=1$ and ‘ a ’ for $A=0$. While this shorthand simplifies the syntax of causal models, it introduces a risk of misinterpretation, for it yields that the factor A and its taking on the value 1 are both expressed by ‘ A ’. Disambiguation must hence be facilitated by the concrete context in which ‘ A ’ appears. Accordingly, whenever we do not explicitly characterise italicized Roman letters as ‘factors’, we use them in terms of the shorthand notation. Moreover, we write ‘ $A*B$ ’ for the conjunction ‘ $A=1$ and $B=1$ ’, ‘ $A + B$ ’ for the disjunction ‘ $A=1$ or $B=1$ ’,



Figure/Table 1: A causal structure (a) (where ‘•’ symbolises conjunction and ‘◊’ expresses negation) with a corresponding complete Humean mosaic (b) and an incomplete one (c).

‘ $A \rightarrow B$ ’ for the conditional ‘If $A=1$, then $B=1$ ’ ($a + B$), and ‘ $A \leftrightarrow B$ ’ for the biconditional ‘ $A=1$ iff $B=1$ ’ ($A*B + a*b$).

To have a concrete context for our ensuing discussion, consider the causal structure over the set of crisp-set factors $\mathbf{F}_1 = \{A, B, C, D, E\}$ in the hypergraph of Figure 1a. This graph has two non-standard elements that require introduction: arrows merged by ‘•’ symbolise conjunctive relevance, and ‘◊’ expresses that the negation of the factor value at the tail of the arrow is relevant. That is, in Figure 1a, $A*B$ and C are two alternative causes of D and a and c are two alternative causes of E . A possible interpretation of this structure might be the following. Suppose a city has two power stations: a wind farm and a nuclear plant. Let A express that the wind farm is operational and C that the nuclear plant is operational and let operationality be sufficient for a nuclear plant to produce electricity, while a wind farm produces electricity provided it is operational and there is wind (B). Hence, the wind farm being operational while it is windy or the nuclear plant being operational ($A*B + C$) are two alternative causes of the city being power supplied (D). Whereas the wind farm or the nuclear plant not being operational ($a + c$) are two alternative causes of an alarm being triggered (E).

We assume that the structure is deterministic and, for simplicity, that there are no causal paths leading to D and E other than the ones through A , B , and C (meaning that there are no latent paths). It then follows that the elements of \mathbf{F}_1 can be co-instantiated in exactly the 8 types of configurations σ_1 to σ_8 in Table 1b. Type σ_1 , for instance, represents a configuration where factors A , C , and D take the value 1 (the wind farm and the nuclear

plant are operational and the city is power supplied) while B and E take value 0 (there is no wind and the alarm is not triggered); type σ_2 represents a configuration where all factors but E take value 1, and so on. Most logically possible configurations of the factors in \mathbf{F}_1 are determined to be inexistent by Figure 1a. For example, C cannot be combined with d , for C causally determines D . Overall, if the behaviour of the factors in \mathbf{F}_1 is underwritten by Figure 1a, Table 1b lists all and only their empirically possible configurations, which, according to the metaphysical embedding of regularity theories, are the configurations that exist in the actual world. As there are no latent causal paths, Table 1b contains a complete distribution of possible matters of fact for the underlying causal structure. We shall say that Table 1b is the *complete Humean mosaic* for the structure in Figure 1a. By the lights of a regularity theory, that the causal dependencies in Figure 1a obtain, essentially means nothing over and above Table 1b being a complete Humean mosaic.

A regularity theory defines causation in terms of sufficiency and necessity relations among values of factors representing different natural properties that are logically and conceptually independent and not related in terms of metaphysical dependencies such as supervenience, constitution, or grounding—we shall speak of *modally independent* properties, for short. Subject to the nature of the involved factors, sufficiency and necessity relations can be given a classical or a fuzzy-logic rendering (see Baumgartner and Ambühl [2018]). In the context of this paper, we can confine ourselves to the classical rendering in terms of material implication: A is sufficient for B iff $A \rightarrow B$, and A is necessary for B iff $B \rightarrow A$. Clearly, most of these Boolean dependencies have nothing to do with causation. For example, the configuration $A*b*C*e$ is sufficient for D in Table 1b, for this table does not feature the combination of $A*b*C*e$ and d . The same holds for $A*B*C*e$, $A*B*C*E$, etc. Moreover, the disjunction of all sufficient conditions of D is necessary for D ; that is, the following relations of sufficiency and necessity obtain among D and the other factors in \mathbf{F}_1 :

$$A*b*C*e + A*B*C*e + A*B*c*E + a*b*C*E + a*B*C*E \leftrightarrow D \quad (1)$$

(1) obviously does not track causation, as the factor E , for example, is part of every sufficient condition of D , but neither E nor e are causally relevant for D in Figure 1a (whether

the alarm is triggered has no influence on the city's power supply). Still, some relations of sufficiency and necessity in fact reflect causation: in Table 1b, $A*B$ and C are individually sufficient and their disjunction is necessary for D and they are the two alternative causes of D . Accordingly, the crucial problem to be solved by a regularity theory is to filter out those Boolean dependencies that track causation.

The main reason why most structures of Boolean dependencies do not reflect causation is that they tend to contain redundant elements, which are dispensable for those dependencies to obtain. Structures of causal dependencies, by contrast, do not feature redundancies. All components of a causal structure make their own distinctive difference to the behaviour of the factors in that structure. Accordingly, the regularity theoretic analyses must be required to be redundancy-free.

Non-Redundancy (NR): A Boolean dependency structure over a set of factors \mathbf{F} tracks causation only if every component of that structure is indispensable to account for the behaviour of the elements of \mathbf{F} .

When applied to sufficient and necessary conditions, (NR) entails that all factor values that can be removed from such conditions without affecting the latter's sufficiency and necessity are not difference-makers and, hence, not causally relevant. Only *minimally* sufficient and *minimally* necessary conditions possibly track causation (Graßhoff and May [2001]).

Minimal Sufficiency: Let Σ be a conjunction of factor values $Z_1 * \dots * Z_n$ with $1 \leq n$. Σ is a minimally sufficient condition of B iff

- (a) the factors in Σ and B represent different natural and modally independent properties,
- (b) $\Sigma \rightarrow B$, and
- (c) for no proper part Σ' of Σ : $\Sigma' \rightarrow B$ (where a proper part of a conjunction is that conjunction reduced by at least one conjunct).

Minimal Necessity: Let Π be a disjunction (in disjunctive normal form)² of factor

² As DNFs allow for rigorous redundancy elimination, regularity theories are, since Mackie ([1974]), stan-

values $Z_1 * \dots * Z_k + \dots + Z_m * \dots * Z_n$ with $1 \leq n$. Π is a minimally necessary condition of B iff

- (a) the factors in Π and B represent different natural and modally independent properties,
- (b) $B \rightarrow \Pi$, and
- (c) for no proper part Π' of Π : $B \rightarrow \Pi'$ (where a proper part of a disjunction is that disjunction reduced by at least one disjunct).

To illustrate, the first disjunct of (1), $A * b * C * e$, is not a minimally sufficient condition of D because it contains sufficient proper parts. For instance, $b * C * e$ is itself sufficient for D in Table 1b. But $b * C * e$ is likewise not minimally sufficient, as it also contains sufficient proper parts. Overall, D has three minimally sufficient conditions in Table 1b: $A * B$, C , and e .³ Their disjunction is necessary for D , that is, $D \rightarrow A * B + C + e$. That necessary condition, however, still contains the spurious dependence between e and D (the alarm not being triggered is minimally sufficient for the city to be power supplied).⁴ The reason is that it does not amount to a minimally necessary condition, as it contains a necessary proper part, $A * B + C$. Whenever D is given, $A * B + C$ is given. The same does not hold for any other proper part of $A * B + C + e$. Or differently, e is dispensable to account for D because, whenever e is given, so is $A * B + C$. But the reverse does not hold: in configurations σ_6 to σ_8 , $A * B + C$ is given but e is not. In sum, the redundancy-free Boolean dependency structure behind the behaviour of D in Table 1b is this one:

$$A * B + C \leftrightarrow D \quad (2)$$

Plainly, these are exactly those sufficiency and necessity relations that reflect the causes of

hardly formulated in terms of DNFs. Not all critics of regularity theories have taken note of that important syntactic restriction; Maslen ([2012]), for instance, attempts to construct a counterexample to Mackie's INUS-theory featuring a disjunctive expression as part of a conjunctive sufficient condition.

³ All calculations can be replicated using the R script in the online supplementary material.

⁴ e is (at least) an INUS condition of D as defined by Mackie ([1974], p. 62), whose INUS-theory is therefore forced to interpret e as a cause of D . This is an instance of the so-called 'Manchester Factory Hooters' problem.

D in Figure 1a.

(2) is a biconditional featuring a minimally necessary disjunction of minimally sufficient conditions of D , in disjunctive normal form. Although the main operator of (2) is symmetric, meaning that D is likewise minimally necessary and sufficient for $A*B + C$, (2) can only be causally interpreted from left to right. The reason is that, in deterministic systems, every configuration of exogenous factors determines one configuration of endogenous factors, and subject to this principle, factor D cannot be exogenous. While every configuration of A , B , and C determines D to be either present or absent, the presence and absence of D do not determine specific configurations of A , B and C . More concretely, $A*B$ and C each determine D and can, thus, be interpreted as deterministic causes of D , but D neither determines $A*B$ nor C (the power supply of the city does not determine what plant the electricity is coming from). Therefore, D cannot be interpreted as deterministic cause of $A*B$ or C (Baumgartner [2013], pp. 95-6). But clearly, if the left-hand side of (2) would only contain a single factor value, say C , determination would be symmetric and, as a result, a causal interpretation in both directions would be feasible. In Section 5, we will introduce a background assumption ensuring that the complexity of our world is high enough to avoid such ambiguities.

In the regularity theoretic literature (for example, Graßhoff and May [2001]), expressions of the form of (2) are commonly taken to be causally interpretable and, hence, furnished with a label: minimal theories.⁵ As will become clear in Section 3, however, minimally necessary disjunctions of minimally sufficient conditions may—unlike (2)—fail to do justice to (NR) and, correspondingly, to track causation. We prefer to reserve the label of a minimal theory to expressions that are guaranteed to comply with (NR), and, thus, refer to all expressions of type (2) as *RDN-biconditionals* (redundancy-free disjunctive normal form biconditionals):

RDN-Biconditional: A true biconditional $\Pi \leftrightarrow B$ is an RDN-biconditional for B iff Π is a minimally necessary disjunction, in disjunctive normal form, of minimally sufficient conditions of B . (Π is the antecedent and B the consequent of the RDN-biconditional.)

No elements can be eliminated from the antecedent of an RDN-biconditional without break-

⁵ Beirlaen *et al.* ([2018]) use the label MINUS-formulas.

ing a sufficiency or necessity relation expressed by that biconditional, that is, without rendering that biconditional false. Every factor value in an RDN-biconditional's antecedent is indispensable to account for the behaviour of its consequent.

While the idea that the minimality of sufficient and necessary conditions is a precondition of their causal interpretability has been present in the literature at least since (Graßhoff and May [2001]), it has so far not been explicitly connected to the intuition that causes are difference-makers of their effects. To render that connection (formally) precise, reconsider the RDN-biconditional (2) entailed by Table 1b and compare configurations σ_5 and σ_6 in that table. In both of them, factor C takes the value 0 and B the value 1, while A and D change from 0 in σ_5 to 1 in σ_6 . In other words, all the disjuncts in (2) not containing A are not instantiated in the pair $\{\sigma_5, \sigma_6\}$ (i.e. all alternative sufficient conditions of D are absent), whereas the contextual factor value B in combination with which A is sufficient for D is constant. It follows that the change from 0 to 1 in D can only be accounted for by the corresponding change from 0 to 1 in A . This is what it means for A to be indispensable in (2) to account for the behaviour of D . Configurations as σ_5 and σ_6 constitute evidence that A is a difference-maker of D . We shall, hence, say that $\{\sigma_5, \sigma_6\}$ is a *difference-making pair* for A with respect to D . To explicitly define that notion, we follow Mackie ([1974], pp. 66-71) in using X as a placeholder for a (possibly empty) conjunction of factor values $Z_1^* \dots^* Z_i$ and Y as a placeholder for a (possibly empty) disjunction $Z_j^* \dots^* Z_k + \dots + Z_m^* \dots^* Z_n$:

Difference-Making Pair: Let $A^*X + Y \leftrightarrow B$ be true. A difference-making pair for A w.r.t. B relative to $A^*X + Y \leftrightarrow B$ is a pair of configurations $\{\sigma_i, \sigma_j\}$ such that A and B are given in σ_i and not given in σ_j , while $X^* \neg Y$ holds in both σ_i and σ_j .

The (tight) connection between the minimality of sufficient and necessary conditions and the difference-making intuition can now be rendered precise:

Theorem 1. $A^*X + Y \leftrightarrow B$ is an RDN-biconditional iff there exist difference-making pairs for all factor values in $A^*X + Y$.

As to Theorem 1, which is proven in the appendix, eliminating redundancies from a true necessary disjunction of sufficient conditions is a means to ascertain the existence of difference-

making pairs. If all component factor values have difference-making pairs, the disjunction is internally redundancy-free. Internal redundancy-freeness is necessary but, as the next section will show, not sufficient for causal interpretability.

3 Structural Redundancies

The reason why the existence of difference-making pairs for all factor values in an RDN-biconditional is not sufficient for causal interpretability is that the notion of a difference-making pair is not defined in terms of the absence of alternative causes (see p. 3) but of alternative sufficient conditions—which are not guaranteed to be causes. To see this, reconsider the structure in Figure 1a over the set of crisp-set factors $\mathbf{F}_1 = \{A, B, C, D, E\}$ and the corresponding complete mosaic in Table 1b. That mosaic not only entails (2) but also an RDN-biconditional for E and one for C :

$$A*B + C \leftrightarrow D \quad (2)$$

$$a + c \leftrightarrow E \quad (3)$$

$$a*D + e \leftrightarrow C \quad (4)$$

While E is the other effect in Figure 1a, C in fact is exogenous. Nonetheless, C can be expressed as an internally redundancy-free Boolean function of its effects D and e (the nuclear plant is operational iff the city is power supplied and the wind farm is not operational or the alarm is not triggered). In other words, there exist difference-making pairs for all factor values in the antecedent of (4) without any of them actually being causes of C . Even though (4) expresses upstream dependencies, all currently existing regularity theories are forced to causally interpret (4) because they take internal redundancy-freeness to be sufficient for causation. As a result, they cannot reliably distinguish between downstream and upstream dependencies and, thus, fall prey to a standard objection against regularity theories (Mackie [1974], pp. 160-1).⁶

⁶ The example in Figure 1a was deliberately chosen for its simplicity. But this problem can arise in structures of arbitrary complexity. To substantiate this, the replication script in the online supplementary material provides a test loop that randomly draws causal structures, simulates Humean mosaics from those structures,

To avoid that consequence, we not only have to impose (non-causal) difference-making constraints on individual Boolean dependencies but also on whole dependency structures. Each substructure \mathcal{X} of a complex causal structure \mathcal{S} makes its own distinctive difference to the overall behaviour of the factors in \mathcal{S} , that is, for every \mathcal{X} in \mathcal{S} it holds that \mathcal{S} (with \mathcal{X}) and \mathcal{S}' , which results from \mathcal{S} by removing \mathcal{X} , have different ramifications for the behaviour of (some of) the involved factors. Correspondingly, when RDN-biconditionals are conjunctively combined to a complex expression Ψ , Ψ tracks causation only if each conjunct in Ψ makes its own distinctive difference to the behaviour of the factors in Ψ . This, indeed, is a hitherto neglected source of (NR)-violations: RDN-biconditionals, although internally redundancy free, can—as a whole—be redundant in superordinate structures and, hence, fail to make a difference due to a higher-order violation of (NR).

(4) is a case in point. It makes no difference to the behaviour of the factors in \mathbf{F}_1 beyond (2) and (3). To show this, we conjunctively concatenate these RDN-biconditionals:

$$(A*B + C \leftrightarrow D) * (a + c \leftrightarrow E) * (a*D + e \leftrightarrow C) \quad (5)$$

For convenience, let us call the conjunction of all RDN-biconditionals entailed by a Humean mosaic δ the *RDNB-conjunction* of δ . It is a transparent and unambiguous representation of all internally redundancy-free regularities inherent in δ and, as such, will be of central relevance to our regularity theory of causation. (5) is the RDNB-conjunction of Table 1b. (5) is true iff the factors in \mathbf{F}_1 take one of the value configurations in Table 1b. If a mosaic coincides with the truth conditions of a Boolean dependency structure, we shall say that the latter *returns* the former. That is, (5) returns Table 1b and, thereby, accounts for the behaviour of the factors in \mathbf{F}_1 . But (5) has a proper substructure that returns the exact same mosaic:

$$(A*B + C \leftrightarrow D) * (a + c \leftrightarrow E) \quad (6)$$

and checks whether these mosaics entail RDN-biconditionals affected by this problem. It turns out that the check is positive in about 35% of the draws.

(6), which results from (5) by eliminating (4), has precisely the same ramifications for the behaviour of the factors in \mathbf{F}_1 as (5). (6) logically entails (4). The RDNB-conjunction (5) is logically equivalent to its proper substructure (6). It follows that, although (4) expresses a regularity entailed by Table 1b, it is dispensable to account for the behaviour of the factors in \mathbf{F}_1 and, thus, violates (NR); it is spurious. By contrast, neither (2) nor (3) are redundant in (5), for neither the conjunction of (2) and (4) nor the conjunction of (3) and (4) are logically equivalent to (5). Both (2) and (3) make their own distinctive difference to the behaviour of the factors in \mathbf{F}_1 .

It is not a peculiarity of our example that the non-redundant RDN-biconditionals are the ones that correspond to causal (downstream) dependencies. As we have seen above (p. 10), in deterministic systems (of some minimal complexity), every configuration of the exogenous factors determines exactly one configuration of the endogenous factors, but not every configuration of the endogenous factors determines exactly one configuration of the exogenous factors. A complete Humean mosaic δ is composed of x^n logically possible configurations of n exogenous factors, each of which can take x values. In each of these x^n configurations, the values of the endogenous factors are determined by the values of the exogenous factors in accordance with the causal structure behind δ . It follows that the values of the endogenous factors in δ can be expressed as functions of the exogenous factors, which, in our framework, correspond to downstream RDN-biconditionals. As (4) demonstrates, this does not entail that no exogenous factors are expressible as RDN-biconditionals of endogenous factors, but it does entail that the downstream RDN-biconditionals always suffice to account for the behaviour of all the factors in δ , to the effect that upstream RDN-biconditionals are redundant for that purpose (because they are recoverable from the downstream RDN-biconditionals). More generally, let Γ be δ 's RDNB-conjunction; the conjunction Ψ of all downstream RDN-biconditionals entailed by δ determines the exact same behaviour of the factors in δ as Γ , meaning that Γ and Ψ are equivalent. Our example is a mere instance of that general principle.

To preclude a causal interpretation of redundant substructures of Boolean dependency structures, not only sufficient and necessary conditions must be minimised but also the struc-

tures as a whole. More formally, conjunctions of RDN-biconditionals entailed by a Humean mosaic δ only track causation if they are structurally minimal:

Structural Minimality: Let δ be a Humean mosaic over the factor set \mathbf{F}_δ and let $\Gamma = \Phi_1 * \dots * \Phi_n$, $n \geq 1$, be δ 's RDNB-conjunction. A conjunction $\Psi = \Phi_k * \dots * \Phi_m$, $1 \leq k \leq m \leq n$, of RDN-biconditionals from Γ is structurally minimal relative to δ iff

- (a) Ψ is logically equivalent to Γ ;
- (b) there does not exist a Ψ' that results from Ψ by eliminating at least one conjunct such that Ψ and Ψ' are logically equivalent.

That a conjunction of RDN-biconditionals Ψ is structurally minimal entails that it states the same as δ 's RDNB-conjunction and that it does not contain an equivalent proper substructure, which, in turn, means that Ψ and all of its substructures return different mosaics. It follows that each conjunct Φ_i in Ψ has some ramification for the behaviour of the involved factors not shared by any other conjunct, and that, as a whole, Ψ accounts for the behaviour of all factors in \mathbf{F}_δ .

Contrary to our example in Table 1b, the RDNB-conjunction of many mosaics can be broken down into multiple structurally minimal conjunctions (for illustrations see Section 5). In consequence, that a particular RDN-biconditional Φ_1 is not contained in a particular structurally minimal Ψ_1 does not exclude that Φ_1 is contained in another structurally minimal conjunction Ψ_2 . Or differently, that Φ_1 is not part of Ψ_1 only means that Φ_1 does not have ramifications for the behaviour of the involved factors over and above the other RDN-biconditionals in Ψ_1 , it does not, however, mean that Φ_1 is structurally redundant simpliciter, for it might be non-redundant relative to Ψ_2 . Φ_1 is structurally redundant only if Φ_1 is not contained in *any* structurally minimal conjunction.

Structural Redundancy: An RDN-biconditional Φ_i entailed by a Humean mosaic δ is structurally redundant relative to δ iff Φ_i is not contained in any conjunction of RDN-biconditionals that is structurally minimal relative to δ .

A structurally redundant RDN-biconditional is redundant (simpliciter) to account for the behaviour of the factors in \mathbf{F}_δ , meaning it does not make a difference on the structural level,

that is, it does not track causation. By contrast, every RDN-biconditional that is contained in some structurally minimal conjunction is *structurally indispensable* as it satisfies the difference-making requirement on the structural level. Structural indispensability is necessary but, as the next section will show, still not sufficient for causal interpretability.

4 Permanence

Real-life causal structures commonly are not as simple as the one in Figure 1a. Causes amount to very complex conjunctions of factor values and, on the type level, there normally exist more than two alternative paths to one effect. To do justice to real-life causal complexities while, at the same time, ensuring that Boolean dependency structures remain manageable, Mackie ([1974], pp. 34-5, 63) relativizes regularities to what he calls a *causal field*, that is, to a fixed configuration of context factors. A more realistic scenario than the one in (6), thus, is that A , a , B , C , and c are mere parts of alternative causes of D and E within a field \mathcal{F} (where X_1, X_2, \dots and Y_1, Y_2 are placeholders for conjunctions and disjunctions, respectively, of additional factor values):

$$\text{in } \mathcal{F} : (A * B * X_1 + C * X_2 + Y_1 \leftrightarrow D) * (a * X_3 + c * X_4 + Y_2 \leftrightarrow E) \quad (7)$$

In scientific discovery contexts, the constancy of the field, of course, is difficult to ensure, which is why real-life data tend not to be as noise-free as Table 1b. Hence, when causally analysing data, strict Boolean dependencies can typically only be approximated. To this end, CCMs provide various parameters of model fit (Ragin [2006]). But since the focus of this paper is conceptual, we will not further discuss these methodological issues here. Likewise, we abstain from making the field-relativity explicit and from using placeholders for additional conjunctions and disjunctions. Instead, we do justice to the complexity of causal structures by assuming Boolean dependency structures to be *open for expansions*, that is, for the integration of further factors.

The remainder of this section will show that expanding Boolean dependency structures provides an important additional handle to constrain their causal interpretability. What

counts as an RDN-biconditional is relative to the analysed factor set. That is, factors contained in an RDN-biconditional relative to a set \mathbf{F}_i may not be contained in an RDN-biconditional relative to a superset $\mathbf{F}_j \supset \mathbf{F}_i$; and some sets faithfully reflect causation, while others do not. To see this, reconsider the structure in Figure 1a and assume that it is analysed without measuring the factor B , that is, relative to $\mathbf{F}_2 = \{A, C, D, E\}$. Consequently, one causal path to D is missing from the analysis. The resulting list of empirically possible configurations in Table 1c, thus, amounts to an incomplete Humean mosaic. It does not allow for expressing the behaviour of D as a function of $\mathbf{F}_2 \setminus \{D\}$, because in the configurations σ_4 and σ_6 all factors in $\mathbf{F}_2 \setminus \{D\}$ are constant while D changes. The RDNB-conjunction of Table 1c is structurally minimal and only features (3) and (4):

$$(a + c \leftrightarrow E) * (a * D + e \leftrightarrow C) \quad (8)$$

Despite its structural minimality, (8) does not track causation, for C is not actually endogenous in Figure 1a. The reason why (4) is not identified as spurious is that \mathbf{F}_2 is *underspecified*, meaning that there exists a latent causal path to an endogenous factor that is not constant in the corresponding causal field and, thus, induces a variation in the endogenous factor that cannot be accounted for based on the factors in \mathbf{F}_2 . In consequence, the Boolean dependencies among the elements of \mathbf{F}_2 cannot be completely freed of redundancies.

Plainly, whether a factor set \mathbf{F}_δ is underspecified depends on the underlying causal structure. Accordingly, in the conceptual context of analysing causation or in the epistemic context of searching for the causal structure behind \mathbf{F}_δ , \mathbf{F}_δ cannot be assumed to be free of underspecification (for this would presuppose clarity on causation and the causal structure behind \mathbf{F}_δ). Fortunately, neither context requires such an assumption because by *gradually expanding* factor sets spurious regularities are identified. When \mathbf{F}_2 is expanded to \mathbf{F}_1 , there no longer are any varying latent paths. Thus, D becomes expressible as a function of $\mathbf{F}_1 \setminus \{D\}$, meaning that (2) follows, which, as we have seen in the previous section, reveals the spuriousness of (4). Generally, regularities that appear to be of a difference-making type relative to a set \mathbf{F}_δ , but in fact are spurious, are identified as such in the course of gradual expansions of \mathbf{F}_δ .

But in order to reliably reveal the spuriousness of Boolean dependencies, expansions of factor sets must be suitable for causal modelling. A *suitable expansion* \mathbf{F}'_{δ} of a factor set \mathbf{F}_{δ} is a superset of \mathbf{F}_{δ} , which is the result of introducing factors into \mathbf{F}_{δ} representing natural properties that are modally independent of one another and of the properties represented by the elements of \mathbf{F}_{δ} . A suitable expansion \mathbf{F}'_{δ} of \mathbf{F}_{δ} reveals that an RDN-biconditional $\Pi_i \leftrightarrow B$, which is structurally indispensable relative to \mathbf{F}_{δ} , features redundancies or is itself redundant if there does not exist a structurally indispensable RDN-biconditional $\Pi_j \leftrightarrow B$ over \mathbf{F}'_{δ} such that all components of Π_i are also components of Π_j . If there does not exist a suitable expansion \mathbf{F}'_{δ} revealing redundancies in (or of) $\Pi_i \leftrightarrow B$, $\Pi_i \leftrightarrow B$ is *permanently redundancy-free*. A structurally indispensable RDN-biconditional tracks causation only if it is permanently redundancy-free.

5 Ambiguities


Before we can assemble the analytical tools developed above in a new regularity theory, we have to introduce a metaphysical background assumption that, although needed to consistently implement (NR), has not been made transparent in the literature so far. The need for that assumption arises from the problem of model ambiguities, which is a widespread phenomenon in all causal modelling frameworks (see Spirtes *et al.* [2000], pp. 59-72; or Eberhardt 2013). A regularity theory is confronted with a model ambiguity when a mosaic entails more than one RDN-biconditional for at least one effect Z . Two cases must be distinguished: either (i) it is possible to interpret the different RDN-biconditionals of Z as representing (distinct aspects/levels of) one and the same causal structure or (ii) that is not possible. In case (i), we shall speak of a mere *functional ambiguity*, whereas case (ii) amounts to a genuine *causal ambiguity*.

We illustrate case (i) with the configurations in Table 2a, again over the factor set $\mathbf{F}_1 = \{A, B, C, D, E\}$. The only factors in that table whose behaviour can be expressed as a function of other factors in \mathbf{F}_1 are C and E . (9) is the corresponding RDNB-conjunction:


$$(A + B \leftrightarrow C) * (C + D \leftrightarrow E) * (A + B + D \leftrightarrow E) \quad (9)$$

#	A	B	C	D	E
σ_1	0	0	0	0	0
σ_2	1	0	1	0	1
σ_3	0	1	1	0	1
σ_4	1	1	1	0	1
σ_5	0	0	0	1	1
σ_6	1	0	1	1	1
σ_7	0	1	1	1	1
σ_8	1	1	1	1	1

(a)



(b)



(c)

Figure/Table 2: Table (a) entails two non-equivalent RDN-biconditionals for E . Structures (b) and (c) both return Table (a).

That is, two RDN-biconditionals for E are entailed; one expressing E as a function of C (and D) and another one expressing it as a function $A + B$ (and D). However, in light of the first conjunct of (9), which states the equivalence of C and $A + B$, these two RDN-biconditionals can be transformed into one another by interchanging C and $A + B$. Hence, they have exactly the same ramifications for the behaviour of the factors in \mathbf{F}_1 , which, in turn, entails that (9) contains redundant proper parts, meaning it is not structurally minimal. It can be broken down into two structurally minimal conjunctions:

$$(A + B \leftrightarrow C) * (C + D \leftrightarrow E) \quad (10)$$

$$(A + B \leftrightarrow C) * (A + B + D \leftrightarrow E) \quad (11)$$

When causally interpreted, (10) expresses the causal chain in Figure 2b and (11) the common-cause structure in Figure 2c—the core difference being that C is a cause of E in the former but not in the latter structure.

This type of ambiguity is ubiquitous in deterministic causation. Baumgartner ([2008a]) has dubbed it the *causal chain problem*: to every deterministic chain there exists an empirically indistinguishable common-cause structure. In a nutshell, the reason is that the behaviour of an outcome Z in a deterministic chain can be expressed as a function f_1 of Z 's direct causes, which, in turn, are functions of their own direct causes; it follows that Z can also be expressed as a function f_2 that is the result of replacing some of Z 's direct causes in f_1 by their direct causes (i.e. Z 's indirect causes) and eliminating redundancies. Baumgart-

ner ([2008a]) proposes to solve the causal chain problem by expanding factor sets to check whether the ambiguities disappear. In our example, if the dependence between C and E vanishes in mosaics over supersets of \mathbf{F}_1 , C is not a cause of E and the underlying structure is of common-cause form. If, by contrast, that dependence is permanent across factor set expansions, the underlying structure is a chain. In that case, both RDN-biconditionals of E track causation: one expresses direct causation, the other indirect causation. That means both structurally minimal conjunctions (10) and (11) can be causally interpreted jointly, which is why they constitute a mere functional ambiguity.

It does not hold generally, however, that multiple RDN-biconditionals with identical consequents can be causally interpreted jointly. To illustrate case (ii), consider Table 3a over $\mathbf{F}_3 = \{A, B, C, D\}$. The RDNB-conjunction of Table 3a consists of two RDN-biconditionals for D , (12) and (13), which are logically equivalent. When causally interpreted, they both identify the following set of causally relevant factor values $\{A, a, B, b, c\}$. However, they place a different Boolean ordering over these causes: according to (12), the set of alternative causes of D is $\{A*b, a*B, A*c\}$; according to (13), it is $\{A*b, a*B, B*c\}$. If (12) and (13) are causally interpreted jointly, it follows that D has four alternative causes: $A*b$, $a*B$, $A*c$, and $B*c$. Such an interpretation, however, violates the difference-making requirement subject to which a cause must make a difference to its effect when all alternative causes are absent: there does not exist a difference-making pair for c w.r.t. D in Table 3a such that the background constantly features $A * \neg(A*b + a*B + B*c)$ or $B * \neg(A*b + a*B + A*c)$. That is, (12) and (13) cannot both be causally interpreted; only one of them possibly tracks causation. But in light of their logical equivalence it is completely undetermined which one. Table 3a thus yields a proper causal ambiguity.

If Table 3a records the possible configurations of the factors in \mathbf{F}_3 relative to some field, in which further relevant factors are constant, the ambiguity between (12) and (13) can be resolved by suitably expanding \mathbf{F}_3 . To make this concrete, suppose that integrating the factor E into \mathbf{F}_3 yields the mosaic in Table 3b, which contains Table 3a as a proper part (highlighted with grey shading). Whenever the added factor E takes the value 1, the factors in \mathbf{F}_3 are instantiated in the configurations recorded in Table 3a; but when E takes the value

#	A	B	C	D
σ_1	0	0	0	0
σ_2	0	0	1	0
σ_3	1	1	1	0
σ_4	1	0	0	1
σ_5	0	1	0	1
σ_6	1	1	0	1
σ_7	1	0	1	1
σ_8	0	1	1	1

(a)

#	A	B	C	D	E
σ_1	0	0	0	0	0
σ_2	1	0	0	0	0
σ_3	0	0	1	0	0
σ_4	1	0	1	0	0
σ_5	1	1	1	0	0
σ_6	0	1	0	1	0
σ_7	1	1	0	1	0
σ_8	0	1	1	1	0
σ_9	0	0	0	0	1
σ_{10}	0	0	1	0	1
σ_{11}	1	1	1	0	1
σ_{12}	1	0	0	1	1
σ_{13}	0	1	0	1	1
σ_{14}	1	1	0	1	1
σ_{15}	1	0	1	1	1
σ_{16}	0	1	1	1	1

(b)

$A*b + a*B + A*c \leftrightarrow D \quad (12)$
 $A*b + a*B + B*c \leftrightarrow D \quad (13)$
 $A*b*E + a*B + B*c \leftrightarrow D \quad (14)$

Table 3: Table (a) entails two logically equivalent RDN-biconditionals for D . Table (b) results from (a) by expansion and entails only one RDN-biconditional for D .

0, further configurations are possible. Table 3b only entails one RDN-biconditional for D : (14). That is, while it is impossible to determine whether D is caused by $A*c$ or $B*c$ relative to Table 3a, Table 3b resolves that ambiguity in favour of $B*c$.

Whereas functional ambiguities may or may not be resolved by factor set expansions, the resolvability of causal ambiguities is crucial for a regularity theory aiming to spell out causation in difference-making terms. If Table 3a could not be expanded (say, because it is complete), (12) and (13) would be permanently redundancy-free, both internally and structurally, which would entail that they both identify difference-makers of D . As we have seen above, however, that cannot be true. Therefore, in order to consistently exploit the idea that causes make a difference to their effects when all alternative causes are absent, Table 3a must be expandable such that the ambiguity between (12) and (13) is resolved.

Although it is easy to devise artificial toy worlds (for instance, in thought experiments targeting the adequacy of theories of causation) without determinate causal structures, we take it as a given that our world is not of this kind. Its causal makeup may be beyond our epistemic reach, but it is ultimately *one* determinate makeup. A regularity theory, therefore, needs to be underwritten by the metaphysical background assumption that causal ambiguities are always due to an insufficient evidential basis, rather than to the ultimate causal

indeterminateness of the world. In principle, causal ambiguities can always be resolved by expanding factor sets. In other words, we assume causal uniqueness for complete mosaics:

Causal Uniqueness (CU): Every complete Humean mosaic corresponds to one determinate causal structure.

One corollary of (CU) deserves separate mention: complete mosaics entail dependency structures with a minimal complexity sufficient to distinguish between causes and effects. The reason is that, as we have seen in Section 2), dependency structures as $Z \leftrightarrow B$ induce a symmetry of determination leaving the direction of causation ambiguous. Subject to (CU), all RDN-biconditionals entailed by complete mosaics have a minimal complexity of $Z_1 + Z_2 \leftrightarrow B$ or, equivalently, $z_1 * z_2 \leftrightarrow b$, which can only be causally interpreted from left to right in line with the principle that every configuration of exogenous factors determines exactly one configuration of endogenous factors.

Being a background assumption, (CU) specifies a precondition for a regularity theory to apply. If there exist worlds whose complete mosaics give rise to causal ambiguities, a regularity theory does not apply to them. That either means that there is no causation in such worlds or, if there is, that another theoretical framework (a non-difference-making theory) needs to be invoked. We do not want to speculate about the existence of worlds with indeterminate difference-making relations. What matters for our purposes is merely that our world is not of this kind. In the end, a regularity theory achieves its aim if it succeeds in analysing causation in the actual world.

6 A New Regularity Theory

We have now collected all ingredients for a new regularity theory of causation. To present that theory, we will proceed in two steps. First, we introduce the notions of a *minimal theory* and of an *atomic minimal theory*, and second, we define causal relevance in terms of containment in permanently redundancy-free atomic minimal theories.

Roughly, a minimal theory is a structurally minimal conjunction of (one or more) RDN-biconditionals. As shown above, the minimality of such a conjunction hinges on the superordinate dependency structure in which it is embedded, which, in turn, depends on the Humean

mosaic over an analysed factor set. A Humean mosaic δ over a set \mathbf{F}_δ is a set of the empirically possible value configurations of the factors in \mathbf{F}_δ . The anti-necessitarian background of regularity theories provides an actualist rendering of the notion of an empirically possible configuration. Causation then supervenes on the actually existing distribution of matters of fact, which, in turn, is a brute fact of our world. If \mathbf{F}_δ contains exogenous factors on all causal paths in the structure Δ behind δ , the corresponding mosaic δ is complete. Subject to (CU), every complete mosaic is underwritten by exactly one causal structure Δ . Complete mosaics allow for complete redundancy elimination. Hence, a minimal theory entailed by a complete δ is free of all redundancies and, thereby, identifies Boolean difference-makers—it is guaranteed to truthfully reflect Δ .

However, causal relevance cannot simply be defined in terms of minimal theories entailed by complete mosaics. The reason was anticipated in Section 4: clarity on mosaic completeness presupposes clarity on causal paths, which is exactly what a theory of causal relevance is supposed to supply and thus, on pain of circularity, cannot presuppose. This problem could be avoided by resorting to all-encompassing *world-mosaics* featuring the empirically possible value configurations of all (modally independent) factors throughout spacetime. World-mosaics do not presuppose clarity on causation and still, as no causal paths can be latent in world-mosaics, allow for complete redundancy elimination. Yet, analysing causation in terms of minimal theories entailed by world-mosaics would yield a theory according to which causation between any pair of factor values depends on the distribution of matters of fact throughout spacetime. Such a theory would not be methodologically implementable, as it would induce infeasible demands on data collection and processing. In fact, however, mosaics as the ones in Tables 1b, 2a, and 3b, which fall far short of world-mosaics, provide reliable evidence on causal relations; and indeed, configurational comparative methods (CCMs) exploit that evidence. Since a core purpose of the theory developed here is to conceptually underwrite CCMs, avoiding the circularity threat by defining causation in terms of world-mosaics is not an option for us.

Accordingly, in the first step of our analysis, we neither confine the notion of a minimal theory to complete mosaics nor to world-mosaics. A minimal theory inferred from a mosaic

δ over any factor set—whether underspecified or not—amounts to a transparent representation of the difference-making evidence contained in δ , which is the chief characteristic of causation for a regularity theory. If δ is complete, that evidence is faithful to the underlying causal structure Δ , but if δ is incomplete, it may misleadingly suggest the causal nature of some dependencies which in fact are spurious. As shown in Section 4, however, factor set expansions gradually rectify minimal theories entailed by a misleading δ by eliminating spurious dependencies and, thereby, ‘zooming in’ on the true Δ —thus the aforementioned second step in our analysis.

Building on the conceptual inventory previously introduced, the following is our definition of the notion of a minimal theory (simpliciter).

Minimal Theory: Let δ be a Humean mosaic over the factor set \mathbf{F}_δ . A minimal theory for δ over \mathbf{F}_δ is a conjunction $\Psi = \Phi_1 * \dots * \Phi_n$, $1 \leq n$, of RDN-biconditionals such that the following conditions hold:

- (a) Ψ is structurally minimal relative to δ ,
- (b) any two Φ_i and Φ_j in Ψ have different consequents.

Condition (a) entails that Ψ is logically equivalent to the RDNB-conjunction of δ (i.e. the conjunction of all RDN-biconditionals entailed by δ) and that Ψ does not contain a logically equivalent proper part. While the purpose of that condition is clear (see Section 3), condition (b) requires explication.

As shown in the previous section, mosaics sometimes entail multiple RDN-biconditionals with identical consequents, some of which—the causally ambiguous ones—cannot be causally interpreted jointly. In order for a minimal theory Ψ to exhibit the difference-making evidence in δ , Ψ must be a candidate representation of the causal structure behind δ . To this end, it must not comprise any RDN-biconditionals that cannot be causally interpreted jointly. The structural minimality restriction in (a) prohibits the concatenation of some causally ambiguous RDN-biconditionals—for instance, of (12) and (13)—but not of all of them. There exist structurally minimal conjunctions comprising RDN-biconditionals with identical consequents. Such conjunctions cannot be interpreted as one causal structure because, in a

structurally minimal conjunction, two RDN-biconditionals with an identical consequent Z have non-equivalent ramifications for the behaviour of Z . In deterministic causal structures, where the behaviour of no outcome follows multiple non-equivalent functional patterns, this amounts to a causal ambiguity with respect to Z .⁷ The purpose of condition (b), hence, is to ensure that minimal theories do not comprise causally ambiguous RDN-biconditionals.


To make this concrete, consider Table 4a, which entails various RDN-biconditionals for various outcomes (for an overview see the online supplementary material). One conjunction of these RDN-biconditionals that satisfies condition (a) and, nonetheless, comprises two RDN-biconditionals with identical consequents is the following:

$$(a + b + c \leftrightarrow D) * (A + C \leftrightarrow E) * (A + B \leftrightarrow E) \quad (15)$$

(15) cannot be interpreted as one causal structure. The two RDN-biconditionals of E cannot be seen as expressing direct and indirect causal relevance relations in a chain because the only conceivable chain-interpretation of (15) would be that B (resp. C) causes E via C (resp. B),⁸ but Table 4a entails no RDN-biconditionals for either B or C , which hence are exogenous. Neither can A , B , and C be interpreted as alternative causes of E because that interpretation violates the difference-making requirement: there is no pair of configurations in Table 4a such that both A and B are absent and a variation of C is associated with a variation of E . Subject to condition (b), the two RDN-biconditionals of E cannot be combined in one minimal theory but must be allocated to different theories. More generally, whenever a mosaic δ yields a causal ambiguity, the RDNB-conjunction of δ must be broken down in as many minimal theories as there are causal structures possibly underwriting δ . The overall causal inference to be drawn from a δ entailing multiple minimal theories Ψ_1 to Ψ_n is (inclusively) *disjunctive*: the evidence in δ is such that Ψ_1 or Ψ_2 or ... or Ψ_n corresponds to the underlying causal structure Δ .

⁷ As (9) illustrates, not all conjunctions of RDN-biconditionals with identical consequents yield causal ambiguities. (9), however, is not structurally minimal.

⁸ There cannot be a chain from A (B) via B (A) to E because A and B are disjuncts in the same RDN-biconditional, and disjuncts in an RDN-biconditional cannot be directly causally related, otherwise that antecedent would contain redundancies.



Figure/Table 4: Table (a) entails two RDN-biconditionals for E . Figure (b) is a switching structure with switch F , and Table (c) the corresponding mosaic.

A minimal theory entailed by a mosaic δ rigorously implements (NR) relative to δ . To properly read relations of Boolean difference-making off of minimal theories, a further specification is needed. To see this, consider the switching structure in Figure 4b with the mosaic in Table 4c. The ultimate effect, G , has two alternative causes, $D + E$, which themselves have two alternative causes each, $A + B^*F$ for D and $C + B^*f$ for E . Importantly, factor F functions as a switch for the causal impact of B on D and E . The combination of B and F causes D , and the combination of B and f causes E . But independently of F , B is sufficient for G . Hence, factor F only makes a difference to whether the influence of B on G is mediated by D or by E but not to G itself. Nonetheless, F and f appear in minimal theories with G as ultimate outcome. In total, Table 4c entails four minimal theories (which constitute a functional ambiguity):

$$(A + B^*F \leftrightarrow D) * (C + B^*f \leftrightarrow E) * (D + E \leftrightarrow G) \quad (16)$$

$$(A + B^*F \leftrightarrow D) * (C + B^*f \leftrightarrow E) * (A + B + C \leftrightarrow G) \quad (17)$$

$$(A + B^*F \leftrightarrow D) * (C + B^*f \leftrightarrow E) * (A + B + E \leftrightarrow G) \quad (18)$$

$$(A + B^*F \leftrightarrow D) * (C + B^*f \leftrightarrow E) * (B + C + D \leftrightarrow G) \quad (19)$$

That the factor F is contained in these minimal theories apparently must not be taken to entail that F is a difference-maker of G . F is not contained in minimally sufficient conditions of G and, correspondingly, it does not appear in an RDN-biconditional for G in any of the theories (16) to (19). That, in turn, shows that it is not membership in minimal theories (simpliciter) that tracks difference-making relations, but membership in RDN-biconditionals contained in minimal theories, which we label *atomic minimal theories*:

Atomic Minimal Theory: An atomic minimal theory Φ of B for a Humean mosaic δ over the set of factors \mathbf{F}_δ is an RDN-biconditional for B contained (as a conjunct) in a minimal theory Ψ for δ over \mathbf{F}_δ .

Now we are in a position to define causal relevance (type-level causation). We define it not just for single factor values but for Boolean expressions in disjunctive normal form (DNF). That, on the one hand, maximises generality but, on the other, bends ordinary speech a bit. By ascribing causal relevance to a DNF as $Z_1 * Z_2 + Z_3$ we mean (i) that Z_1 , Z_2 , and Z_3 are causally relevant factor values, (ii) that Z_1 and Z_2 are jointly relevant, and (iii) that $Z_1 * Z_2$ and Z_3 are alternatively relevant. Moreover, we say that a DNF Ω is *contained* in another DNF Π iff every conjunction in Ω is a conjunct in a conjunction in Π , and any two disjuncts in Ω are conjuncts in two different disjuncts in Π .

Causal Relevance (CR): Let Ω be a Boolean expression in DNF. Ω is causally relevant for B iff there exists a set of (modally independent) factors \mathbf{F}_δ containing B and all factors in Ω , such that δ is a Humean mosaic over \mathbf{F}_δ , and the following conditions hold:

- (a) there exists an atomic minimal theory $\Pi \leftrightarrow B$ for δ over \mathbf{F}_δ such that Ω is contained in Π ,
- (b) for every suitable expansion $\mathbf{F}'_{\delta'} \supset \mathbf{F}_\delta$ and corresponding mosaic δ' : there exists an atomic minimal theory $\Pi' \leftrightarrow B$ for δ' over $\mathbf{F}'_{\delta'}$ such that Ω is contained in Π' .

(CR) can be less formally expressed as follows: Ω is causally relevant for B iff Ω is contained in the antecedent of a permanently redundancy-free atomic minimal theory of

B . Ω can have any complexity. If Ω is a single factor value A , a conjunction $A * C$, or a disjunction $A + C$, (CR) provides the conditions under which, respectively, A is causally relevant, A and C are jointly relevant, and A and C are alternatively relevant for B .

Before discussing some implications of (CR) in the next section, two features of (CR) deserve separate emphasis. First, (CR) is formulated against the background of (CU), which ensures that suitably expanding factor sets will always resolve causal ambiguities. That is, if an incomplete mosaic entails multiple minimal theories that cannot be causally interpreted jointly, only as many will survive after expansion as can be causally interpreted jointly. Second, causal relevance as defined in (CR) is non-transitive. It is possible for Z_1 to be causally relevant for Z_2 , which is causally relevant for Z_3 , without Z_1 being causally relevant for Z_3 . The switching structure in Figure 4b is a case in point. Presuming that the minimal theory (16) is permanently redundancy-free, it follows that F is causally relevant for both D and E , which are relevant for G , but F is not causally relevant for G .

7 Discussion

While (CR) draws on analytical tools from previous regularity theoretic proposals (Mackie [1974], Graßhoff and May [2001]; Baumgartner [2008b], [2013]), it assembles these tools in a way that implicates a departure from an implicit consensus among its predecessors. The latter all entail (or presuppose) that multi-effect structures can be modularly built up from single-effect structures, meaning that, in order to identify the causes of some effect B , it suffices to identify members of permanently redundancy-free sufficient and necessary conditions of B . This paper suggests that this modularity principle cannot be sustained. According to (CR), the redundancy-freeness of Boolean dependency structures and, thus, their causal interpretability cannot be assessed for single-effect structures individually but only for complete multi-effect structures. As a result, (CR) entails a form of *causal holism* according to which causation is a holistic property that supervenes on complete Humean mosaics and not on proper parts thereof.

That holism has a number of notable ramifications. For instance, it yields that (CR) is more restrictive in sanctioning the causal interpretability of Boolean dependency structures

than its predecessor theories: all dependencies that can be causally interpreted according to (CR) can also be causally interpreted according to its predecessors but not vice versa. This is particularly important in light of the fact that most of the classical objections levelled against regularity theories since the times of Hume and Mill contend that these theories overgenerate, meaning they warrant the causal interpretation of regularities which in fact are spurious. It follows that those overgeneration problems that have already been solved by (CR)'s predecessors are solved correspondingly by (CR); this concerns in particular the problems of empty and single-case regularities (see Armstrong [1983]) as well as the problem of spurious regularities due to common causes, as most famously instantiated in Mackie's ([1974], pp. 83-7) 'Manchester Factory Hooters' example. For detailed discussions of these issues the reader is hence referred to (Graßhoff and May [2001]) and (Baumgartner [2008b], [2013]).

Section 3 has shown, however, that one overgeneration problem has not been addressed by (CR)'s predecessors. It can happen that the behaviour of exogenous factors is expressible in terms of RDN-biconditionals featuring their own effects. (CR)'s predecessors cannot reliably distinguish between upstream and downstream regularities because they do not ensure that all substructures of a complex Boolean dependency structure make a difference on the structural level. (CR) solves that problem by prohibiting the causal interpretation of structurally redundant RDN-biconditionals.

Of course, regularity theories have also been objected to on the ground that they undergenerate in case of irreducible indeterminism (Dowe and Noordhof [2004]). While standard interpretations of quantum mechanics advocate the existence of irreducibly indeterministic processes, non-standard interpretations disagree. Hence, there is no consensus on whether our universe is deterministic or not. Moreover, even if irreducibly indeterministic processes exist, there are many open questions—as for instance raised by phenomena of the EPR type—with respect to the *causal* interpretability of these processes (Healey [2010]). In the present context, we can sidestep these foundational questions, for, as indicated in the introduction, regularity theories aim to capture the intuition that causation is a deterministic form of dependence, that is, they analyse deterministic variants of causation (only). The notion of causal relevance spelled out in (CR) must, hence, be understood in terms of *deter-*

ministic causal relevance. If there should turn out to exist irreducibly indeterministic causal relevance relations, other theoretical frameworks would have to be invoked.

Another upshot of the causal holism entailed by (CR) is that unmistakable causal inferences can only be drawn from complete Humean mosaics. It is, of course, questionable whether complete Humean mosaics for other than artificial causal structures are ever available to human reasoners. That is, causal relevance as defined in (CR) can typically only be approximated in scientific practice. Nonetheless, atomic minimal theories inferred from incomplete mosaics transparently represent the causal evidence contained in those mosaics. Even though, in the absence of complete mosaics, causal inferences run a risk of being refuted in the light of factor set expansions, such inferences become increasingly warranted the longer memberships in minimal theories are stable throughout a series of expansions. That is, the inference to causal relevance as defined by (CR) is *inherently inductive*, which—we contend—nicely captures the nature of causal inference in scientific practice.

Finally, as (CR) is the first regularity theory that eliminates structural redundancies and, thus, provides the first notion of Boolean difference-making that rigorously implements (NR), configurational comparative methods, as QCA or CNA, which output Boolean causal models, are well-advised to understand causal relevance in terms of (CR). QCA focuses on single-effect structures and considers building multi-effect structures as optional. (CR) calls for a revision of that approach. Reliable Boolean causal inference not only requires expanding the evidence base on the causes of single effects, but necessitates also aggregating single- to multi-effect structures. While such an aggregation has always been an essential element in the procedural protocol of CNA, CNA has, so far, conceived of this aggregation in too simplistic a manner: it solely conjunctively concatenates minimal biconditionals inferred from processed data. According to (CR), an additional iteration of (structural) redundancy elimination is required.

We end with two caveats. First, note that (CR) does not distinguish between direct and indirect causal relevance. In light of the non-transitivity of (CR)-defined causal relevance, indirect relevance cannot simply be spelled out in terms of the transitive closure of direct relevance, which, in turn, is accounted for in terms of containment in permanently redun-

dancy-free atomic minimal theories. Discriminating between direct and indirect relevance presupposes a notion of a causal chain, which, for reasons of space, we cannot properly introduce here. Second, note again that (CR) provides a notion of type-level causation. Token-level causation or actual causation must be cashed out in terms of a suitable spatiotemporal instantiation of a type-level structure. Building a corresponding token-level account on the basis of (CR) must also await another occasion.

Appendix

In this appendix, we prove Theorem 1, p. 11, which states the following equivalence:

(i) $A * X + Y \leftrightarrow B$ is an RDN-biconditional.

\leftrightarrow

(ii) There exist difference-making pairs for all factor values in $A * X + Y$.

For our proof, recall that X stands for a (possibly empty) conjunction $Z_1 * \dots * Z_i$, and Y for a (possibly empty) disjunction $Z_j * \dots * Z_k + \dots + Z_m * \dots * Z_n$. Moreover, let A denote an arbitrary factor value on the left-hand side of the biconditional in (i). We prove both entailment directions separately.

Proof of (i) \rightarrow (ii):

(i) entails that $A * X + Y$ is minimally necessary for B . It follows that Y alone is not necessary for B , which means that there exists a configuration σ_i featuring both $B = 1$ and $Y = 0$. Still, as to (i), $A * X + Y$ is necessary for B . Since it holds that $B = 1$ while $Y = 0$ in σ_i , σ_i must feature $A * X = 1$ (otherwise B could not take the value 1 in σ_i). In sum, σ_i features $A = B = X = 1$ and $Y = 0$. (i) moreover entails that $A * X$ is minimally sufficient for B . It follows that X alone is not sufficient for B , which means that there exists a configuration σ_j featuring $A = B = 0$ and $X = 1$. Factor B only takes the value 0 if all of its other sufficient conditions in Y are absent, hence, $Y = 0$ in σ_j . In sum, σ_j features $A = B = Y = 0$ and $X = 1$. The pair $\{\sigma_i, \sigma_j\}$ is a difference-making pair for A w.r.t. B . As A denotes an arbitrary factor value, the above argument can be repeated for every element of $A * X + Y$. This proves that if (i) holds, there exists a difference-making pair for every factor value in $A * X + Y$. \square

Proof of (ii) \rightarrow (i):

The notion of a difference-making pair is defined for the elements of antecedents (in DNF) of true biconditionals of the form $A * X + Y \leftrightarrow B$. Thus, (ii) entails that $A * X + Y \leftrightarrow B$ is true, which, in turn, means that $A * X$ and every disjunct in Y is sufficient for B . As to (ii), moreover, there is a configuration σ_j featuring $A = B = Y = 0$ and $X = 1$, meaning that X alone is not sufficient for B . It follows that A is a non-redundant part of the sufficient condition $A * X$. Since A is an arbitrary factor value in $A * X + Y$, this argument can be repeated for every other conjunct of $A * X$ as well as for every conjunct of every disjunct in Y , meaning that $A * X + Y$ is exclusively composed of minimally sufficient conditions of B . Furthermore, (ii) entails that $A * X + Y$ is necessary for B and that there exists a configuration σ_i featuring $A = X = B = 1$ and $Y = 0$, meaning that Y alone is not necessary for B . In other words, $A * X$ is a non-redundant part of the necessary condition $A * X + Y$. By the same token, every disjunct in Y can be shown to be non-redundant, meaning that $A * X + Y$ is a minimally necessary condition of B . In sum, $A * X + Y$ is a minimally necessary disjunction of minimally sufficient conditions of B , that is, an RDN-biconditional. \square

Acknowledgements

We thank Jiji Zhang for helpful discussions and two anonymous reviewers for thoughtful and thorough comments on earlier drafts. Moreover, we are indebted to the University of Bergen and the Trond Mohn Foundation for generous support of this research (grant no. 811886).

Michael Baumgartner

University of Bergen

Department of Philosophy

Bergen, Norway

michael.baumgartner@uib.no

Christoph Falk

University of Geneva

Department of Philosophy

Geneva, Switzerland

chri.falk@gmail.com

References

- Armstrong, D. M. [1983]: *What is a Law of Nature?*, Cambridge: Cambridge University Press.
- Baumgartner, M. [2008a]: ‘The Causal Chain Problem’, *Erkenntnis*, **69**, pp. 201–26.
- Baumgartner, M. [2008b]: ‘Regularity Theories Reassessed’, *Philosophia*, **36**, pp. 327–54.
- Baumgartner, M. [2013]: ‘A Regularity Theoretic Approach to Actual Causation’, *Erkenntnis*, **78**, pp. 85–109.
- Baumgartner, M. and Ambühl, M. [2018]: ‘Causal Modeling with Multi-Value and Fuzzy-Set Coincidence Analysis’, *Political Science Research and Methods*, pp. 1–17. doi: 10.1017/psrm.2018.45.
- Beirlaen, M., Leuridan, B. and Van De Putte, F. [2018]: ‘A Logic for the Discovery of Deterministic Causal Regularities’, *Synthese*, **195**, pp. 367–99.
- Broad, C. D. [1930]: ‘The Principles of Demonstrative Induction (I.)–(II.)’, *Mind*, **39**, pp. 302–17, pp. 426–39.
- Dowe, P. and Noordhof, P. (eds) [2004]: *Cause and Chance: Causation in an Indeterministic World*, London: Routledge.
- Eberhardt, F. [2013]: ‘Experimental Indistinguishability of Causal Structures’, *Philosophy of Science*, **80**, pp. 684–96.
- Gelman, A. and Hill, J. [2007]: *Data Analysis Using Regression and Multi-level/Hierarchical Models*, Cambridge: Cambridge University Press.
- Graßhoff, G. and May, M. [2001]: ‘Causal Regularities’, in W. Spohn, M. Ledwig and M. Esfeld (eds), *Current Issues in Causation*, Paderborn: Mentis, pp. 85–114.

- Healey, R. [2010]: ‘Causation in Quantum Mechanics’, in H. Beebe, C. Hitchcock and P. Menzies (eds), *The Oxford Handbook of Causation*, Oxford University Press.
- Hume, D. [1999 (1748)]: *An Enquiry Concerning Human Understanding*, Oxford: Oxford University Press.
- Lewis, D. [1986]: *Philosophical Papers II*, Oxford: Oxford University Press.
- Mackie, J. L. [1974]: *The Cement of the Universe: A Study of Causation*, Oxford: Clarendon Press.
- Maslen, C. [2012]: ‘Regularity Accounts of Causation and the Problem of Pre-Emption: Dark Prospects Indeed’, *Erkenntnis*, **77**, pp. 419–34.
- Mill, J. S. [1843]: *A System of Logic*, London: John W. Parker.
- Ragin, C. C. [2006]: ‘Set Relations in Social Research: Evaluating Their Consistency and Coverage’, *Political Analysis*, **14**, pp. 291–310.
- Rihoux, B. and Ragin, C. C. (eds) [2009]: *Configurational Comparative Methods: Qualitative Comparative Analysis (QCA) and Related Techniques*, Thousand Oaks: Sage.
- Spirtes, P., Glymour, C. and Scheines, R. [2000]: *Causation, Prediction, and Search* (2 ed.), Cambridge: MIT Press.
- Thiem, A. [2014]: ‘Unifying Configurational Comparative Methods: Generalized-Set Qualitative Comparative Analysis’, *Sociological Methods & Research*, **43**, pp. 313–37.
- Thiem, A., Baumgartner, M. and Bol, D. [2016]: ‘Still Lost in Translation! A Correction of Three Misunderstandings between Configurational Comparativists and Regression Analysts’, *Comparative Political Studies*, **49**, pp. 742–74.

Online Supplementary Material

```

# R replication script
# #####
# R version used 3.6
# Required R package (package version used: 2.2.1)
library(cna)

# Fundamentals
# -----
# Table 1b:
dat1 <- allCombs(c(2,2,2,2,2)) -1
(tab1b <- selectCases("(A*B + C <-> D)*(c + a <-> E)", dat1))
# Minimally sufficient conditions for D:
ana1 <- cna(tab1b, what = "mac")
subset(msc(ana1), outcome == "D")
# RDN-biconditional for D:
subset(asf(ana1), outcome == "D")

# Structural redundancies
# -----
# RDNB-conjunction for Table 1b:
ana1
# Structurally minimal conjunction of RDN-biconditionals:
minimalizeCsf(ana1)

# Test loop estimating the frequency of structural redundancies
# (different runs will generate different results; redundancy
# ratios will vary accordingly):
n <- 100
score <- vector("list", n)
for(i in 1:n){
  cat(i, "\n")
  x <- randomCsf(full.tt(8), n.asf = 3, compl = 2)
  y <- selectCases(x)
  score[[i]] <- csf(cna(y, details = T, rm.dup.factors = F), 1)
}
eval <- Filter(function(x) dim(x)[1] > 0,
               lapply(score, function(x) subset(x, x$redundant == TRUE)))
# Structural redundancy ratio:
length(eval)/n

```

```

# Permanence
# -----
# Table 1c:
(tab1c <- tt2df(tab1b)[, -2])
ana2 <- cna(tab1c)
# Structurally minimal conjunction of RDN-biconditionals:
minimalizeCsf(csf(ana2)$condition, dat1)

# Ambiguities
# -----
# Table 2a:
(tab2a <- selectCases("(A + B <-> C)*(C + D <-> E)", dat1))
cna(tab2a)
# Tables 3a/b:
dat2 <- allCombs(c(2,2,2,2)) -1
(tab3a <- selectCases("A*b + a*B + A*c <-> D", dat2))
# Two structurally minimal RDN-biconditionals:
cna(tab3a)
# Ambiguity resolution through factor set expansion:
(tab3b <- selectCases("A*b*E + a*B + B*c <-> D", dat1))
cna(tab3b)

# A new regularity theory
# -----
# Table 4a:
(tab4a <- selectCases("(a + b + c <-> D)*(A + C <-> E)*(A + B <-> E)", dat1))
ana1 <- cna(tab4a, details = T)
# All RDN-biconditionals entailed by Table 4a:
asf(ana1)$condition
# All minimal theories for Table 4a:
(mt <- as.vector(minimalizeCsf(subset(csf(ana1, Inf),
                                   exhaustiveness == 1)$condition, dat1)$condition))

# Table 4c:
dat3 <- allCombs(rep(2,7)) -1
(tab4c <- selectCases("(A + B*F <-> D)*(C + B*f <-> E)*(D + E <-> G)",
                      dat3))
cna(tab4c)

```