A reductive INUS theory of causation

Luna De Souter

June 2025

1 Introduction

Regularity theories of causation within the so-called *INUS*[1] framework aim to define causation in terms of redundancy-free regularity relations between factor values, typically representing event types. These theories build upon John Mackie's (1974) original INUS account and introduce additional non-redundancy, or *minimality*, constraints. INUS theories capture intuitions about causation found in influential accounts such as those by Hume (1748 (1963)) and Mill (1843), and they serve as a philosophical foundation for *configurational comparative methods* (Rihoux and Ragin, 2009), a family of methods used for discovering causal structures from data in various scientific fields, including political science, social science, and health research.

INUS theories seek to provide a *reductive* definition of causation, meaning that they aim to spell out how causal structures are determined by a non-causal reduction base. This distinguishes these theories from, among others, interventionist accounts of causation (Woodward, 2003), which rely on the causal concept of an intervention to define causation. INUS theories hold that the non-causal reduction base of causal structures consists of regularity patterns—that is, sets of regularity relations between event types. An INUS theory successfully reduces causation to regularity patterns if, and only if (iff), it imposes criteria that are satisfied by all and only those regularity patterns that track genuine causal structures. When this condition is met, we can say that the theory *uniquely identifies* causal structures based on regularity patterns.

This paper demonstrates that INUS theories fail to achieve their reductive goal. I establish this by focusing on the theory developed by Baumgartner and Falk (2023a), a recent theory in the INUS framework that imposes the most stringent criteria on regularity patterns. In the remainder of this paper, I refer to this theory as 'the INUS* theory'. I show that the INUS* theory fails to reductively define causation by presenting an example in which the INUS* theory cannot uniquely identify a causal structure based on regularity patterns, because the criteria this theory imposes are insufficiently strict.

[1]'INUS' is a name originally derived from John Mackie's (1974, 62) technical phrase *Insufficient Non-redundant part of an Unnecessary Sufficient condition.*

While the problem that I present bears some resemblance to Jaegwon Kim's (1971, 434) critique of Mackie's (1974) original INUS account—namely, that there are regularity patterns based on which Mackie's INUS account cannot uniquely identify a causal structure—my critique goes beyond Kim's. Kim's critique has been convincingly refuted by Baumgartner and Falk (2023a), who point out that Kim's examples involve overly simplistic candidate causal structures that lack the complexity required to be uniquely identifiable from regularity patterns. Adding complexity to these candidate structures enables their unique identification by the INUS* theory, and Baumgartner and Falk (2023a) maintain that the INUS* theory analyzes causation not for simple toy worlds but for the real world, in which causal structures are sufficiently complex to be uniquely identified based on regularity patterns.

The example that I present differs from examples like Kim's (1971, 434) in that it possesses sufficient complexity to ensure that the INUS* theory's inability to uniquely identify it cannot be attributed to simplicity. Consequently, Baumgartner and Falk's (2023a) rebuttal of Kim's critique does not apply to my example. The INUS* theory's inability to reduce my example structure to regularity patterns therefore reveals a failure in the theory's reductive aim.

To address this problem, I propose amending the INUS* theory by adding a new criterion that enables the unique identification of my example structure. I demonstrate that the so-called *3^{rd} Non-Redundancy condition (NR3)*, a criterion recently proposed by Zhang and Zhang (2025), is appropriate for this purpose.

While Zhang and Zhang (2025) have argued for integrating NR3 into the INUS* theory, their arguments are unlikely to convince INUS theorists, as Zhang and Zhang largely focus on causal discovery and provide little discussion of how integrating NR3 results in a more adequate theory of causation. Importantly, Zhang and Zhang (2025) leave unaddressed a key objection against integrating NR3, despite acknowledging this objection and conceding that it is reasonable.

In addition to demonstrating that the INUS* theory without NR3 fails as a reductive theory of causation and that adding NR3 would restore the theory's ability to reductively define causation, this paper refutes the unaddressed objection against incorporating NR3 into the INUS* theory. Through these contributions, the paper firmly establishes that NR3 must be integrated into the INUS* theory of causation.

2 Background

2.1 Principles of INUS causation

INUS theories aim to define type-level causation while adhering to a number of central principles. A first of these principles is that of *causal anti-realism*, entailing a rejection of the notion that causation is a power or necessitating relationship governing the behaviours of causes and their effects. Instead, INUS theories assert that causation supervenes on these behaviours (Andreas and Günther, 2021). In this respect, INUS theories align with other anti-realist approaches like counterfactual (Lewis, 1973) and probabilistic (Suppes, 1970) accounts of causation, while differing from realist production accounts (e.g. Mumford and Anjum, 2010; see Hall, 2004 for a discussion of the distinction between realist and anti-realist accounts).

Second, INUS theories take causation to be *deterministic*, in the sense that an effect is a deterministic function of its causes: the same causes are accompanied by the same effects. This principle aligns, for instance, with Hume's famous description of a cause as "*an Object, follow'd by another, and where all the Objects, similar to the first, are follow'd by Objects, similar to the second*" (Hume, 1748 (1963), 81) and conforms to common intuitions about causation (Frosch and Johnson-Laird, 2011; Rothe et al., 2018; Schulz and Sommerville, 2006). In providing a deterministic account of causation, INUS theories contrast with probabilistic theories (Suppes, 1970), which assert that causes only determine the probability distributions of their effects, allowing that the same causes lead to different effects on different occasions.

A third principle of INUS theories, discussed less frequently than anti-realism and determinism but crucial to many regularity accounts since Mill (1843), is the *conjunctivity of causation*: often, an individual cause brings about its effect only when combined with other individual causes (Rothman, 1976). For example, striking a match will only cause a fire when instantiated in combination with other factors, such as the presence of flammable materials. All components of such a conjunctive bundle need to occur together to lead to the bundle's effect. The principle of conjunctivity is essential in configurational comparative methods, as many scientific questions investigated using these methods involve complex causal structures in which multiple individual causes interact to determine an effect. For instance, outreach to one's own patients may cause high hepatitis C virus treatment uptake, though only when combined with another individual cause, such as external medical site visits (Yakovchenko et al., 2020).

A fourth principle, also vital to INUS theories though discussed less frequently than anti-realism and determinism, is the *disjunctivity of causation*: an effect can have multiple alternative conjunctive bundles of causes, with only one such bundle needing to be instantiated to bring about the effect. This principle, also known as *equifinality* (Gerring, 2005, 164), implies that different conjunctions of causes can lead to the same effect. For instance, as an alternative to striking a match, a lightning strike (in conjunction with other individual causes) can cause fire. Like conjunctivity, the disjunctivity of causation is crucial for analyzing many scientific questions. For example, support to Ukraine by a given EU member state may be caused by that member state facing a strong threat from Russia, while alternatively being caused by a conjunction of military investments and public support (Haesebrouck, 2024).

2.2 Conceptual preliminaries

The INUS* theory defines causation in terms of regularity relations between factor values, that is, between factors taking specific values (Baumgartner and Falk, 2023a). An example of such a regularity relation is $A=1 \rightarrow B=0$, meaning

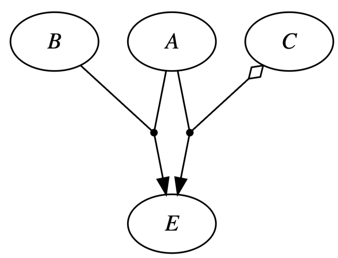

(a)

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

(b)

Figure/Table 1: Figure (a) is a causal hypergraph representing the causal structure $\mathcal{S}_1$. Arrows represent direct causal relevance, '$\bullet$' stands for conjunction, '$\diamond$' means negation, and multiple arrows into the same node form a disjunction. Table (b) represents the so-called *Humean mosaic* of (a), referred to as δ_1, consisting of configurations σ_1 to σ_8.

that factor A taking the value 1 is sufficient for factor B taking the value 0. Here, sufficiency is understood as material implication: whenever A=1 is given, B=0 is given too.

Following Baumgartner and Falk (2023a, 174), I limit my discussion to binary factors, which can take values of either 0 or 1. In this context, I use italicized letters as a shorthand for factors taking a value of 1, and I use the symbol '$\neg$' to indicate the negation of a factor value. For example, 'A' stands for A=1 (e.g. the presence of event type A) and '$\neg A$' stands for A=0 (e.g. the absence of event type A).

Consider the causal structure $\mathcal{S}_1$, depicted by the hypergraph (Falk et al., 2024) in Figure 1a. In this representation, arrows signify direct causal relevance, dots ('$\bullet$') indicate conjunctions of factor values represented at arrow tails, diamonds ('$\diamond$') represent negations, and multiple arrows into the same node form a disjunction. $\mathcal{S}_1$ is a structure over the factor set $\mathbf{F}_1 = \{\text{A, B, C, E}\}$, which I call the *factor frame* of structure $\mathcal{S}_1$. Following Mackie (1974), I discuss causal structures like $\mathcal{S}_1$ relative to a *causal field*, which refers to a set of 'standing conditions' or context factors that have fixed values in the unmeasured causal background and are excluded from the factor frame (Baumgartner and Falk, 2023a, 182).

According to the INUS* theory, the same causal relevance relationships as in Figure 1a can be represented by the following expression:

$$AB \vee A\neg C \leftrightarrow E \tag{1}$$

In this expression, concatenation represents *and*, '$\vee$' means *or*, and '$\leftrightarrow$' denotes *if and only if (iff)*, indicating material equivalence. Both Figure 1a and biconditional (1) imply that there are two alternative conjunctive bundles of causes—henceforth also referred to as *conjunctive causes*—either of which is sufficient for effect E relative to the given causal field: the conjunction of A and B, or the conjunction of A and $\neg C$.

For example, if A, B, and C represent switches in an electric circuit (with, for instance, A meaning that switch A is in position 1 and $\neg A$ meaning that switch A is in position 0), and E represents a lamp being on (E) or off ($\neg E$) in that same circuit, then both the biconditional and the figure imply that the lamp E is turned on by either both A and B being in position 1 or A being in position 1 and C being in position 0. In general, a biconditional that satisfies the criteria of causation imposed by the INUS* theory has a disjunction of conjunctions involving at least two factor values as its (causal) antecedent and a single factor value representing the effect as its consequent (Baumgartner and Falk, 2023a, 177-178, 187).[2]

The sets of regularities to which the INUS* theory aims to reduce causal structures like $\mathcal{S}_1$ are regularities in 'actual distributions of matters of fact' (Baumgartner and Falk, 2023a, 174). To illustrate, consider Table 1b, which represents the distribution of matters of fact over $\mathbf{F}_1$ that determines structure $\mathcal{S}_1$. Each row in this table represents a configuration of the factors in $\mathbf{F}_1$ that complies with the regularities implied by $\mathcal{S}_1$, and every logically possible value combination of the factors in $\mathbf{F}_1$ that complies with the regularities implied by $\mathcal{S}_1$ is featured in Table 1b. For example, as Table 1b lacks any configurations featuring A and B together with $\neg E$, each configuration in this table complies with the regularity $AB \rightarrow E$, which is implied by $\mathcal{S}_1$. Meanwhile, all logically possible combinations of values of A, B, and C appear in at least one configuration in this table, since $\mathcal{S}_1$ does not prohibit any of these combinations: configuration σ_1 features A, B, and C, σ_2 features A, B, and $\neg C$, and so on.

Since any set of configurations has a unique (exhaustive) set of regularities obtaining in it and any set of regularities has a unique (exhaustive) set of configurations satisfying these regularities, there is a one-to-one correspondence between sets of configurations and sets of regularities. I refer to sets of configurations in which the regularity patterns that determine causal structures obtain as *Humean mosaics* (Lewis, 1986, ix-x):

Humean mosaic of a causal structure: A Humean mosaic δ of a causal structure $\mathcal{S}$ over factor frame $\mathbf{F}$ is the exhaustive set of configurations of the factors in $\mathbf{F}$ which are *empirically possible* according to $\mathcal{S}$, that is, which do not violate any regularity relations implied by $\mathcal{S}$.

Similarly, I will refer to the exhaustive set of configurations that do not violate any regularity relations represented by a given logical expression as the 'Humean mosaic' of that expression. For instance, the Humean mosaic in Table 1b, which I will call δ_1, is the mosaic of expression (1) as well as of structure $\mathcal{S}_1$. For simplicity, I will, henceforth, often discuss the INUS* theory's reductive aim in

[2]Requiring that antecedents contain at least two factor values while consequents contain exactly one enables distinguishing causes from effects in a biconditional: the multiple factor values in the antecedent are causes of the single factor value in the consequent (Graßhoff and May, 2001, 97-99).

terms of reducing causal structures to Humean mosaics rather than in terms of reducing causal structures to *regularity relations obtaining in* Humean mosaics.

2.3 Theoretical preliminaries

According to the INUS* theory, a biconditional like (1) accurately tracks causation only if (i) the conjunctions in its antecedent are *minimally sufficient* and (ii) the disjunction formed by these conjunctions is *minimally necessary* for the consequent. A conjunction is minimally sufficient for a consequent iff it is sufficient for the consequent and taking away any of its conjuncts would make it lose its sufficiency (Baumgartner and Falk, 2023a, 176). For instance, if A by itself is sufficient for E, then a conjunction like AB is not minimally sufficient for E because B is redundant for the conjunction's sufficiency. Additionally, a disjunction is minimally necessary for E iff it is necessary for E and removing any of its disjuncts (e.g. AB or $A\neg C$ in $AB \vee A\neg C$) would make it lose its necessity.

While Mackie's (1974) original INUS account already posited that conjunctions of causes must be minimally sufficient and disjunctions must be necessary for the effect—a requirement referred to as the *1st Non-Redundancy condition (NR1)* by Zhang and Zhang (2025)—the requirement that disjunctions must be *minimally* necessary was introduced later (Graßhoff and May, 2001; May, 1999). I follow Zhang and Zhang (2025) in referring to this latter requirement as the *2nd Non-Redundancy condition (NR2)*:

NR2: A biconditional $\Pi \leftrightarrow B$ satisfies NR2 for a Humean mosaic δ iff B is a single factor value and Π is a minimally necessary disjunction of minimally sufficient conjunctions for B in δ.

NR2-biconditional: A biconditional $\Pi \leftrightarrow B$ is an NR2-biconditional for a Humean mosaic δ iff $\Pi \leftrightarrow B$ satisfies NR2 for δ.

If $\Pi \leftrightarrow B$ is an NR2-biconditional, its antecedent Π is called an *NR2-antecedent* (for B) and its consequent B is called an *NR2-consequent*.

Baumgartner and Falk (2023a, 178-179) formally link NR2 to difference-making by proving an equivalence between a true biconditional satisfying NR2 and the existence of difference-making pairs of configurations for this true biconditional. Let $A\phi_1 \vee \phi_2 \vee ... \vee \phi_n$ be a disjunction of conjunctions of factor values, let B be a single factor value, and let δ be a Humean mosaic in which biconditional $A\phi_1 \vee \phi_2 \vee ... \vee \phi_n \leftrightarrow B$ is true (i.e. satisfied in all configurations). ϕ_1 is said to represent A's *accompanying conjunct(s)*, and $\phi_2, ..., \phi_n$ are termed the *alternative conjunctions* of $A\phi_1$. I will also refer to $\phi_2, ..., \phi_n$ as the alternative conjunctions of A in $A\phi_1$ and of ϕ_1 in $A\phi_1$. Then $A\phi_1 \vee \phi_2 \vee ... \vee \phi_n \leftrightarrow B$ satisfies NR2 for δ iff this biconditional satisfies the criterion of *NR2-Difference Making* (DM2) for δ:

NR2-Difference Making (DM2): Biconditional $A\phi_1 \vee \phi_2 \vee ... \vee \phi_n \leftrightarrow B$ satisfies *DM2* for δ iff δ contains at least one *NR2-Difference-Making pair (DM2-pair)* for each factor value appearance in $A\phi_1 \vee \phi_2 \vee ... \vee \phi_n$.

NR2-Difference-Making pair (DM2-pair): A *DM2-pair* for A in $A\phi_1 \vee \phi_2 \vee ... \vee \phi_n \leftrightarrow B$ is a pair of configurations $\{\sigma_i, \sigma_j\}$ such that A and B are given in σ_i and are not given in σ_j, while accompanying conjuncts ϕ_1 are present and alternative conjunctions ϕ_2, ..., ϕ_n are absent in both σ_i and σ_j.

To illustrate the notion of a DM2-pair, let us identify the DM2-pairs in δ_1 for the first appearance of A in biconditional (1) (i.e. the appearance of A in conjunction AB). Each such pair consists of two configurations σ_i and σ_j that must satisfy a number of constraints. First, A and E must appear in σ_i and must not appear in σ_j. Second, A's accompanying conjunct B needs to be given in both σ_i and σ_j. And third, both these configurations must have A's alternative conjunction $A\neg C$ absent.

As any configuration satisfying the constraints on σ_i features A and E, the absence of $A\neg C$ requires that $\neg C$ is not given and thus that C is given in such a configuration. Therefore, the only configuration in δ_1 qualifying as σ_i is σ_1, which has A and E given, B given, and C given. Meanwhile, any configuration satisfying the constraints on σ_j features $\neg A$, automatically ensuring the absence of $A\neg C$ and leaving the value of C irrelevant. Consequently, all configurations in δ_1 with A and E not given and B given, namely σ_5 and σ_6, qualify as σ_j. The DM2-pairs in δ_1 for the first appearance of A in biconditional (1) are thus $\{\sigma_1, \sigma_5\}$ and $\{\sigma_1, \sigma_6\}$.

Note that these pairs differ from the DM2-pairs of A's second appearance, namely $\{\sigma_4,\sigma_6\}$ and $\{\sigma_4,\sigma_8\}$, which feature $\neg C$ and have AB absent. This motivates defining DM2-pairs for factor value *appearances* rather than factor values, which is not explicitly stated but seems to be presupposed in Baumgartner and Falk's (2023a, 179) definition.

The equivalence between NR2 and DM2 shows that the INUS* theory's criterion that a cause A must appear in an NR2-antecedent for its effect B effectively captures the notion that a cause *makes a difference* to its effect while controlling for other causes: the constant presence of accompanying conjuncts and constant absence of alternative conjunctions across configurations in a DM2-pair prevents accompanying conjuncts and alternative conjunctions from accounting for the variation in B across the pair, demonstrating A's indispensability to account for the variation in B (Baumgartner and Falk, 2023a, 178-179).

2.4 Causal models

While an NR2-biconditional can represent a single-effect causal structure, a causal structure involving more than one effect is represented by a conjunction of more than one NR2-biconditional in which each NR2-biconditional has a *distinct consequent*—in the sense that no two biconditionals have a value of the same factor as their consequent. Take the following expression as an example:

$$(AB \vee A\neg C \leftrightarrow E)(\neg B \vee C \leftrightarrow G) \tag{2}$$

Like (1), (2) implies that AB and $A\neg C$ are conjunctive causes of E. Additionally, (2) implies that $\neg B$ and C are (single-conjunct) conjunctive causes of G.

Using our previous electrical circuit interpretation, we could say that G represents a second lamp in the circuit, which turns on when switch B is in position 0 or switch C is in position 1.

Not all conjunctions of NR2-biconditionals with distinct consequents are suitable for representing causal structures. Baumgartner and Falk (2023a, 179-182) note that some such conjunctions contain biconditionals that make no difference to the Humean mosaic of the conjunction. This occurs when a conjunction of NR2-biconditionals (with distinct consequents) for a given mosaic contains an NR2-biconditional $\Pi \leftrightarrow B$ that only expresses regularities already implied by other NR2-biconditionals in the conjunction. In such cases, $\Pi \leftrightarrow B$ is redundant with respect to these other NR2-biconditionals and does not express genuine causal relevance relations in conjunction with them.

Correspondingly, Baumgartner and Falk (2023a, 181-182) add to NR2 an additional criterion of *structural minimality*:

Structural minimality: Let Ψ be a conjunction of NR2-biconditionals for a Humean mosaic δ. Ψ is *structurally minimal* with respect to δ iff the following two conditions are met:

1. Ψ is logically equivalent to the conjunction of *all* NR2-biconditionals for δ; and
2. removing any conjunct from Ψ would result in a conjunction that is no longer logically equivalent to the conjunction of all NR2-biconditionals for δ.

Note that a single NR2-biconditional can be structurally minimal with respect to a Humean mosaic, but only if that biconditional is logically equivalent to the conjunction of all NR2-biconditionals for that mosaic.

I define single- or multiple-effect causal models satisfying structural minimality in addition to NR2 as *NR2-models*:

NR2-model: A conjunction Ψ of biconditionals is an *NR2-model* for a Humean mosaic δ iff the following three conditions are met:

1. each biconditional in Ψ is an NR2-biconditional for δ;
2. each biconditional in Ψ has a distinct consequent; and
3. Ψ is structurally minimal with respect to δ.

2.5 Ambiguities and expansions

Like causal discovery frameworks including the influential approaches developed by Spirtes et al. (2000) and Pearl (2009), regularity theories within the INUS framework face a challenge due to *causal ambiguities*. In the INUS framework, causal ambiguities arise when different NR2-models that cannot represent the same causal structure qualify as NR2-models for the same Humean mosaic. For

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

(a)

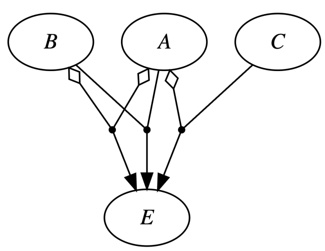

(b)

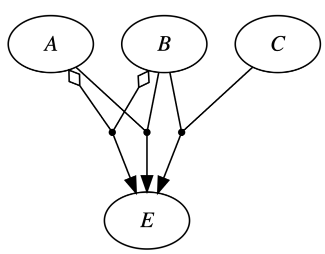

(c)

	A	B	C	E	X_1
σ_1	1	1	1	0	0
σ_2	0	1	1	1	0
σ_3	0	0	1	0	0
σ_4	1	1	0	1	0
σ_5	0	0	0	1	0
σ_6	1	0	1	0	0
σ_7	0	1	0	0	0
σ_8	1	0	0	0	0
σ_9	1	1	1	1	1
σ_{10}	0	1	1	1	1
σ_{11}	0	0	1	1	1
σ_{12}	1	1	0	1	1
σ_{13}	0	0	0	1	1
σ_{14}	1	0	1	0	1
σ_{15}	0	1	0	0	1
σ_{16}	1	0	0	0	1

(d)

Figure/Table 2: Table (a) depicts the Humean mosaic δ_2. Figure (b) plots the candidate causal structure represented by model (3), and Figure (c) plots the candidate causal structure represented by model (4). Table (d) depicts mosaic δ_2', which contains δ_2 (highlighted in grey shading) as its proper part.

instance, consider Humean mosaic δ_2, depicted in Table 2a, over factor frame $\mathbf{F}_2 = \{\text{A, B, C, E}\}$. There are two NR2-models for δ_2:

$$AB \vee \neg A \neg B \vee \neg AC \leftrightarrow E \tag{3}$$

$$AB \vee \neg A \neg B \vee BC \leftrightarrow E \tag{4}$$

As models (3) and (4) are NR2-models for the same Humean mosaic, I say that they are *NR2-equivalent* to each other. Even though models (3) and (4) agree that AB and $\neg A \neg B$ are conjunctive causes of E, these models are *conflicting* causal models, that is, they cannot represent the same causal structure: model (3) implies that $\neg AC$ causes E, which is not implied by model (4), and model (4) implies that BC causes E, which is not implied by model (3). The different candidate causal structures represented by (3) and (4) are plotted by the hypergraphs in Figures 2b and 2c, respectively.

The ambiguity about whether $\neg AC$ or BC causes E in δ_2 stems from A and B's multiple appearances within each NR2-biconditional for E. These multiple appearances create specific dependencies among the conjunctions AB, $\neg A \neg B$, $\neg AC$, and BC in δ_2. For instance, since every configuration features either A or $\neg A$, any configuration with both B and C must have either AB or $\neg AC$ present. Consequently, δ_2 contains no configurations with the conjunction BC present while both AB and $\neg AC$ are absent, implying that any configuration in δ_2 with BC and without $\neg AC$ must have AB present.

Note that any mosaic in which BC is insufficient for E contains a configuration with BC and $\neg E$, and, more specifically, that any mosaic in which BC is insufficient for E while $\neg AC$ is sufficient for E contains a configuration with BC, without $\neg AC$, and with $\neg E$. Therefore, as a configuration with BC and without $\neg AC$ must have AB present, any mosaic in which BC is insufficient and $\neg AC$ is sufficient for E contains a configuration with $AB \neg E$. But δ_2 does not contain any configurations with $AB \neg E$: such configurations are empirically impossible because, as both (3) and (4) agree, AB is a conjunctive cause of E (and thus sufficient for E).

In sum, BC being insufficient while $\neg AC$ is sufficient would require a configuration with value combination $AB \neg E$, but this is prohibited by the fact that AB, which shares factors with $\neg AC$ and BC, is sufficient for E (since all NR2-models for δ_2 agree that AB is a conjunctive cause of E). Similarly, the multiple appearances of A and B in AB, $\neg A \neg B$, $\neg AC$, and BC imply that, whenever BC is sufficient for E while AB and $\neg A \neg B$ are conjunctive causes of E, $\neg AC$ must also be sufficient for E, and that whenever one of $\neg AC$ and BC forms a minimally necessary disjunction with AB and $\neg A \neg B$, the other one (i.e. BC or $\neg AC$) must also form a minimally necessary disjunction with AB and $\neg A \neg B$.

The ambiguity between models (3) and (4) is an illustration of Kim's (1971, 434) critique of Mackie's (1974) original INUS account. Kim considers the inability of INUS theories to determine based on δ_2 which one of $\neg AC$ or BC causes E to be a problem, as this inability suggests that these theories cannot uniquely identify the causal structure of mosaic δ_2 based solely on (the regularity relations obtaining in) this mosaic, that is, based solely on the causal structure's non-causal reduction base.

However, Baumgartner and Falk (2023a) argue that causal ambiguities need not be problematic for the INUS* theory, as these ambiguities stem from considering factor frames that include only a limited number of factors. Candidate causal structures over such limited factor frames may be insufficiently complex to be determined by the distribution of matters of fact over the considered factor frame alone. While Baumgartner and Falk (2023a, 187-188) explicitly acknowledge that an INUS theory fails to provide a reductive definition of causation if it cannot uniquely identify the causal structure of the actual world when the factor frame is not limited, they argue that the inability of the INUS* theory to uniquely identify a causal structure based on δ_2 stems not from a flaw in the theory but from the fact that δ_2 reflects an artificially simple toy world which does not *have* a causal structure that is sufficiently complex to be uniquely determined by a Humean mosaic.

When faced with ambiguities, Baumgartner and Falk argue, complexity should be added through a *factor frame expansion*, which amounts to integrating additional factors into the factor frame and allowing them to vary (Baumgartner and Falk, 2023a, 183-184):

Factor frame expansion: In an expansion of a factor frame $\mathbf{F}$ with a Humean mosaic δ to a factor frame $\mathbf{F}'$ with a Humean mosaic δ', the following four conditions are met:

1. $\mathbf{F} \subset \mathbf{F}'$ (i.e. $\mathbf{F}$ is a proper subset of $\mathbf{F}'$);
2. δ is a proper part of δ';
3. at least one factor X which is contained in $\mathbf{F}'$ but not in $\mathbf{F}$ varies in δ'; and
4. all factors in $\mathbf{F}'$ are *modally independent* of each other, meaning that they are 'logically and conceptually independent and not related in terms of metaphysical dependencies such as supervenience, constitution, or grounding' (Baumgartner and Falk, 2023a, 175).

To illustrate the conditions that a factor frame expansion satisfies, consider the expansion of factor frame $\mathbf{F}_2 = \{\text{A, B, C, E}\}$ with mosaic δ_2, by factor X_1, to factor frame $\mathbf{F}'_2 = \{\text{A, B, C, E, X}_1\}$ with mosaic δ'_2. δ'_2 is displayed in Table 2d. This expansion satisfies the first three conditions: $\mathbf{F}_2$ is a proper subset of $\mathbf{F}'_2$, δ_2 is a proper part of δ'_2, and X_1 varies in δ'_2 in the sense that this mosaic contains at least one configuration with X_1 and at least one configuration with $\neg X_1$. Furthermore, in this paper, all factors are assumed to be modally independent, such that the fourth condition is automatically satisfied.

The addition of X_1 to conjunction ABX_1 through this expansion renders it empirically possible for A and B to appear in a configuration together with $\neg E$ (when X_1 is not given), which was required to reveal the insufficiency of BC for E while preserving the sufficiency of $\neg AC$ for E. Model (5) is the only

NR2-model for δ_2':

$$ABX_1 \vee \neg A \neg B \vee \neg AC \leftrightarrow E \tag{5}$$

This model is a *supermodel* of model (3):

Supermodel: An NR2-model $\mathcal{M}$ is a *supermodel* of an NR2-model $\mathcal{M}'$ iff $\mathcal{M}'$ can be derived from $\mathcal{M}$ by removing conjuncts from disjuncts in NR2-antecedents, disjuncts from NR2-antecedents, and/or NR2-biconditionals from the model (Baumgartner and Falk, 2024, 4).

To illustrate, in the case of models (5) and (3), the latter is obtained from the former by removing conjunct X_1 from disjunct ABX_1.

Since model (3) is an NR2-model for the mosaic over $\mathbf{F}_2$ and (3)'s supermodel (5) is an NR2-model for the mosaic over $\mathbf{F}_2'$, the factor frame expansion of $\mathbf{F}_2$ to $\mathbf{F}_2'$ corresponds to a *model expansion* of (3) to (5):

Model expansion: An expansion of a factor frame $\mathbf{F}$ with a Humean mosaic δ to a factor frame $\mathbf{F}'$ with a mosaic δ' is a *model expansion* of NR2-model $\mathcal{M}$ to NR2-model $\mathcal{M}'$ iff:

1. $\mathcal{M}$ is an NR2-model for δ;
2. $\mathcal{M}'$ is an NR2-model for δ'; and
3. $\mathcal{M}'$ is a supermodel of $\mathcal{M}$.

Accordingly, expanding factor frame $\mathbf{F}_2$ to $\mathbf{F}_2'$ corresponds to adding complexity to the causal structure represented by model (3), which resolves the ambiguity between models (3) and (4) and determines that it is $\neg AC$ rather than BC that causes E.

2.6 Defining causal relevance

The INUS* theory's definition of causal relevance (Baumgartner and Falk, 2023a, 192) combines the necessary conditions of causation discussed in Sections 2.3 to 2.5 into two criteria that are necessary and jointly sufficient for causal relevance. The first criterion in this definition posits that causes are contained in antecedents of structurally non-redundant NR2-biconditionals for their effects. Since such containment is not always preserved after factor frame expansions—as exemplified by BC in Section 2.5, which is contained in an NR2-antecedent for E relative to factor frame $\mathbf{F}_2$ but not relative to expanded factor frame $\mathbf{F}_2'$—the second criterion stipulates that a cause's containment in an antecedent of a structurally non-redundant NR2-biconditional is preserved as the factor frame is gradually expanded.

Causal relevance: Let Ω be a disjunction of conjunctions of factor values and B a single factor value. Ω is *causally relevant* to B iff there exists a factor frame $\mathbf{F}$ which is a set of (modally independent) factors including B and all factors featured in Ω, with δ a Humean mosaic over $\mathbf{F}$, such that the following two conditions are satisfied:

1. there exists an NR2-biconditional $\Pi \leftrightarrow B$ for δ that is part of an NR2-model for δ, such that Ω is contained in Π; and
2. for every factor frame expansion $\mathbf{F}' \supset \mathbf{F}$ with mosaic δ': there exists an NR2-biconditional $\Pi' \leftrightarrow B$ for δ' that is part of an NR2-model for δ', such that Ω is contained in Π'.

3 The problem

3.1 Full S-expansions

Section 2.5 has explained how Baumgartner and Falk (2023a) attribute any failure of the INUS* theory to uniquely identify a causal structure based on a given Humean mosaic to insufficient complexity in the candidate causal structures of that mosaic. This response seems to assume that only simple candidate causal structures are unidentifiable by the INUS* theory. However, Baumgartner and Falk (2023a) fail to argue why this assumption would be satisfied. Furthermore, they do not specify criteria for determining when a candidate causal structure is so simple that, if it is unidentifiable by a theory, then this unidentifiability should be attributed to simplicity rather than to a flaw in the theory.

Given this absence of a clear complexity threshold, this section characterizes a complexity level that renders any failure of a theory to uniquely identify structures possessing it attributable to this theory rather than to the simplicity of these structures. Rather than attempting to pinpoint the precise threshold at which a candidate causal structure should become sufficiently complex for unique identification—a likely controversial matter—I define a standard of complexity stringent enough for anyone to agree that, if the INUS* theory fails to uniquely identify a structure with this complexity, then this failure must stem from a problem with the theory, not from the structure's simplicity.

Characterizing this complexity standard requires introducing some terminology. First, while the INUS* theory defines causal relevance as a relation between factor values, the upcoming discussion frequently references factors whose values are causally related. For conciseness, when a value of a factor X is causally relevant to a value of another factor Y, I will say that X is a *causing factor* of Y. Additionally, I define a factor X as an *ancestor* of a factor Y iff X is a causing factor of Y, or a causing factor of a causing factor of Y, and so on.[3] This allows us to define the property of *acyclicity*: a causal structure is *acyclic* iff it does not feature any factor pair {X, Y} where X is an ancestor of Y and Y is an ancestor of X. I limit the upcoming discussion to acyclic causal structures, since the examples in Section 3.2 are acyclic and this limitation in scope will simplify the discussion, particularly in Appendix A.

Furthermore, consider again the example model (5). I call X_1 a *specific factor* and its value X_1 a *specific factor value* with respect to this model:

[3] More precisely, the relation 'is an ancestor of' is the transitive closure of the relation 'is a causing factor of'.

Specific factor: A factor A is a *specific factor* with respect to an NR2-model $\mathcal{M}$ (and with respect to a structure $\mathcal{S}$ represented by $\mathcal{M}$) iff:

1. exactly one value of A appears in $\mathcal{M}$;
2. this value appears exactly once in $\mathcal{M}$; and
3. this appearance occurs in the antecedent of an NR2-biconditional of $\mathcal{M}$.

Specific factor value: A factor value A is a *specific factor value* in an NR2-model $\mathcal{M}$ (and in a structure $\mathcal{S}$ represented by $\mathcal{M}$) iff A appears in $\mathcal{M}$; and the factor of which A is a value is a specific factor with respect to $\mathcal{M}$.

This means that a specific factor value in a model appears only once in this model, that no other values of the same factor appear in this model, and that this factor value is *exogenous* (i.e. not *endogenous*) with respect to this model:

Endogeneity: A factor A, and all values of A, are *endogenous* with respect to an NR2-model $\mathcal{M}$ (and with respect to a structure $\mathcal{S}$ represented by $\mathcal{M}$) iff a value of A appears as an NR2-consequent in $\mathcal{M}$ (i.e. iff a value of A is an effect in $\mathcal{S}$).

Exogeneity: A factor A, and all values of A, are *exogenous* with respect to an NR2-model $\mathcal{M}$ (and with respect to a structure $\mathcal{S}$ represented by $\mathcal{M}$) iff A appears in $\mathcal{M}$ and is *not* endogenous with respect to $\mathcal{M}$ (i.e. iff *no* value of A is an effect in $\mathcal{S}$).

Additionally, if a specific factor value A appears in an NR2-antecedent for a factor value B in an NR2-model, then I will refer to A as a *specific cause* of factor value B, and to factor A as a *specific causing factor* of factor B, with respect to that model. Similarly, if a factor value A which is *not* a specific factor value appears in an NR2-antecedent for factor value B in an NR2-model, then A is a *non-specific cause* of B, and A a *non-specific causing factor* of B, with respect to that model.

We have seen how the presented expansion of model (3) (with mosaic δ_2), by specific factor X_1, to model (5) (with mosaic δ_2') resolves the ambiguity between models (3) and (4). X_1 acted as a kind of switch in (5) that could be turned on (i.e. X_1 is given) or off (i.e. X_1 is not given) irrespective of the values of A and B. Among the empirically possible configurations with $\neg X_1$, there are configurations with factor value combination $AB\neg E$ that make the distribution of matters of fact in (5)'s Humean mosaic sufficiently diverse to reveal that $\neg AC$, and not BC, is a conjunctive cause of E. So, making empirically possible additional factor value combinations of endogenous factor E with E's non-specific causing factors A and B, whose multiple appearances in (3) gave rise to ambiguities, here results in the unique determination of the causal structure represented by (5).

Unlike δ_2, mosaic δ_2' contains both a configuration featuring value combination AB together with E and a configuration featuring AB together with $\neg E$.

An even more diverse distribution of matters of fact than the one in δ_2' could contain, for *each* empirically possible value combination of A and B, a configuration with E as well as one with $\neg E$ together with that combination. In such a case, I say that E *varies independently* of A and B:

Independent variation (factor): A factor X *varies independently* of a factor set $\mathbf{F}$ in a Humean mosaic δ iff, for any combination of values of factors in $\mathbf{F}$ that appears in at least one configuration in δ, δ contains at least one configuration featuring this value combination together with X and at least one configuration featuring this value combination together with $\neg X$.[4]

Independent variation (factor set): A factor set $\mathbf{F}$ *varies independently* in a Humean mosaic δ iff every logically possible value combination of the factors in $\mathbf{F}$ appears in at least one configuration in δ.

The introduced terminology allows to motivate and define the complexity standard that prevents attributing identification failures to the simplicity of candidate causal structures. I propose that such a complexity standard should have the property of guaranteeing that any endogenous factor (such as E in (5)) varies independently of its non-specific causing factors (such as A and B in (5)). As I will prove below, such independent variation can be guaranteed by adding 'switches' similar to X_1 in (5) until a *full S-expansion* is achieved:

Full S-expansion: A model $\mathcal{M}$ (and a structure $\mathcal{S}$ represented by $\mathcal{M}$) is a *full S-expansion* iff:

1. $\mathcal{M}$ contains a specific factor value in each of its disjuncts; and
2. $\mathcal{M}$ contains a disjunct comprising only specific factor values in each of its NR2-antecedents.

As an example of a full S-expansion, consider the following expansion of model (3) with a Humean mosaic called δ_2^* (not displayed in this paper due to space constraints):

$$ABX_1 \vee \neg A \neg B X_2 \vee \neg AC \vee Y_1 \leftrightarrow E \tag{6}$$

One of the specific factor values X_1, X_2, C, and Y_1 appears in each disjunct in model (6), and Y_1 forms a disjunct comprising only specific factor values in this model's NR2-antecedent.

To demonstrate that each endogenous factor varies independently of its non-specific causing factors in the Humean mosaic of any (acyclic) full S-expansion, I begin with the example model (6). In this model, such independent variation requires that E varies independently of A and B in δ_2^*, which is proven as follows:

[4] 'Independent variation' as used in this paper is not to be confused with probabilistic independence.

Proof. In general, the set of all factors that are exogenous with respect to a causal structure varies independently in the Humean mosaic of that structure, because nothing in the causal structure implies that any combination of values of these exogenous factors is empirically impossible, such that all of these combinations are contained in the Humean mosaic of the structure (Baumgartner and Falk, 2023a, 175).

As non-specific factors A and B and specific factors X_1, X_2, C, and Y_1 are exogenous with respect to (6), the factor set $\{\mathrm{A}, \mathrm{B}, \mathrm{X}_1, \mathrm{X}_2, \mathrm{C}, \mathrm{Y}_1\}$ varies independently in δ_2^*. So, for any combination of values of A and B that appears in δ_2^*, δ_2^* contains a configuration featuring this combination while having all specific factor values absent. Since each disjunct in E's antecedent contains a specific factor value, and since a disjunct is absent when any of its conjuncts is absent, all disjuncts in E's antecedent are absent in this configuration, resulting in the absence of E's antecedent. Given that E's antecedent is necessary for E, E cannot be featured in this configuration. Therefore, for each value combination of A and B appearing in δ_2^*, δ_2^* contains at least one configuration featuring this combination together with $\neg E$.

Furthermore, the independent variation of exogenous factors ensures that, for any value combination of A and B that appears in δ_2^*, δ_2^* contains at least one configuration featuring this combination while having the disjunct in E's antecedent comprising only specific factor values (here, Y_1) present. Since any disjunct in E's antecedent is sufficient for E, E is given in this configuration. Hence, for each value combination of A and B featured in δ_2^*, δ_2^* contains at least one configuration featuring that combination together with E. Consequently, as, for any value combination of A and B appearing in δ_2^*, δ_2^* contains at least one configuration featuring this combination together with E and at least one configuration featuring this combination together with $\neg E$, E varies independently of A and B in δ_2^*. □

Appendix A demonstrates that this independent variation of endogenous factors from their non-specific causing factors generalizes to all acyclic full S-expansions. It follows that adding complexity to a causal structure as is done in a full S-expansion leads to a highly diverse distribution of matters of fact. The specific factor values act like switches that can determine the value of a given endogenous factor regardless of the values of that endogenous factor's non-specific causing factors. This is reminiscent of intervention variables in accounts like Woodward's (2003, 98), which are factors that can change the value of a factor X on which they intervene while breaking the causal dependence between X and X's other causing factors.

While qualifying as a full S-expansion imposes a very stringent demand on the complexity of causal structures, what matters for my argument is that, *if* a causal structure with this complexity level is not uniquely identified by the INUS* theory, then this unidentifiability cannot be attributed to the structure's simplicity but instead demonstrates a problem with the theory itself.

3.2 The triangle problem

Consider the following NR2-model with a Humean mosaic referred to as δ_3 (a mosaic not displayed due to space constraints):[5]

$$(Y_1 \vee AX_1 \leftrightarrow B)(Y_2 \vee ABX_2 \leftrightarrow C) \tag{7}$$

Model (7) implies that Y_1 causes B; that A in conjunction with X_1 causes B; that Y_2 causes C; and that the conjunction of A, B, and X_2 causes C. I refer to this model as a 'triangle' model, because it represents a causal triangle structure with non-specific factor values A, B, and C as its vertices.

This triangle model is fully S-expanded: it has a disjunct exclusively comprising specific factor values in each of its NR2-antecedents (Y_1 in $Y_1 \vee AX_1$ and Y_2 in $Y_2 \vee ABX_2$) and a specific factor value in each of its other disjuncts (X_1 in AX_1 and X_2 in ABX_2). Therefore, it should be complex enough to be uniquely determined by its Humean mosaic δ_3.

However, two other models are NR2-equivalent to model (7):

$$(Y_1 \vee AX_1 \leftrightarrow B)(Y_2 \vee AY_1X_2 \vee AX_1X_2 \leftrightarrow C) \tag{8}$$

$$(Y_1 \vee AX_1 \leftrightarrow B)(Y_2 \vee AY_1X_2 \vee \neg Y_1BX_2 \leftrightarrow C) \tag{9}$$

The ambiguity between models (7) and (8) is not a *causal* ambiguity but instead exemplifies what Baumgartner and Falk (2023a, 184) call a *functional ambiguity*. The distinctive feature of functional ambiguities is that they are ambiguities between non-conflicting models, that is, models which can be interpreted as representing the same causal structure. Specifically, if model (7) represents the direct causal relevance relations in a causal structure (relative to the factor frame {Y$_1$, A, X$_1$, B, Y$_2$, X$_2$, C}), then model (8) represents indirect causal relevance relations in that same structure. This is because (8) can be obtained from (7) by replacing B in the NR2-antecedent $Y_2 \vee ABX_2$ by B's causes $Y_1 \vee AX_1$; rewriting the resulting expression $Y_2 \vee A(Y_1 \vee AX_1)X_2$ as a disjunction of conjunctions $Y_2 \vee AY_1X_2 \vee AAX_1X_2$; and removing a redundant appearance of A from AAX_1X_2. Model (8) thus 'skips over' B in explicitly representing the causal relevance of $Y_1 \vee AX_1$ to C implied by (7) and can be viewed as a less fine-grained representation of the structure represented by (7). So, the NR2-equivalence between models (7) and (8) does not jeopardize the INUS* theory's ability to reduce the structure represented by (7) to a Humean mosaic (Baumgartner and Falk, 2023a, 184-185).

Model (9), in contrast, cannot represent the same structure as model (7). This is because (9) implies that $\neg Y_1BX_2$ causes C—which is not implied by model (7)—while model (7) indicates that ABX_2 causes C—which is not implied by model (9). Consequently, the ambiguity between models (7) and (9) is not a functional ambiguity but a genuine causal ambiguity. Therefore, the

[5] I thank Timm Lampert for bringing to my attention a similar example structure, which inspired this one.

INUS* theory fails to uniquely identify model (7) based on its Humean mosaic. As model (7) is fully S-expanded, this failure cannot be attributed to simplicity. Correspondingly, this failure, which I term 'the triangle problem', poses a challenge to the theory's reductive aim.

An INUS* theorist might try to dismiss this challenge by simply positing that model (7) does not represent any part of the real world's causal structure—a response that aligns with Baumgartner and Falk's (2023a, 187) 'Causal Uniqueness assumption', which states that the actual world has one determinate causal structure. However, such a response would be *ad hoc*, and positing that model (7) does not represent any part of the real world's causal structure would be highly implausible. Moreover, model (7) is not the only fully S-expanded structure unidentifiable by the INUS* theory. Instead, it is a representative of a broader family of structures—which I call *stemmatic structures*—that the INUS* theory fails to uniquely identify based on Humean mosaics.

I define stemmatic structures as those generated by one of two procedures. The first procedure begins with model (7) and extends the causal chain from A, through B, to C by adding intermediate links to it. For example, incorporating a link R between A and B yields:

$$(Y_1 \vee AX_1 \leftrightarrow R)(Y_3 \vee RX_3 \leftrightarrow B)(Y_2 \vee ABX_2 \leftrightarrow C) \tag{10}$$

This procedure generates additional stemmatic structures by adding more links to the causal chain from A, through R and B, to C. A common feature of these structures is that a factor appears alongside one of its ancestors in the same conjunctive cause. For instance, in the structure represented by model (10), the conjunctive cause ABX_2 features B alongside its ancestor A.

The second procedure also starts with model (7) but adds a link between A and C that is *not* on the chain through B, as in the following model:

$$(Y_1 \vee AX_1 \leftrightarrow B)(Y_3 \vee AX_3 \leftrightarrow T)(Y_2 \vee BTX_2 \leftrightarrow C) \tag{11}$$

This procedure generates further stemmatic structures by adding more links to either of the chains from A to C. In all structures generated using this second procedure, two factors that share a common ancestor appear together in a conjunctive cause. For example, in the structure represented by model (11), the two factors B and T, which share a common ancestor A, appear together in the conjunctive cause BTX_2.

These two procedures can generate infinitely many stemmatic structures, all fully S-expanded yet unidentifiable by the INUS* theory. Model (10), for instance, is a full S-expansion with specific factor values Y_1, X_1, Y_3, X_3, Y_2, and X_2, while being NR2-equivalent to the following model:

$$(Y_1 \vee AX_1 \leftrightarrow R)(Y_3 \vee RX_3 \leftrightarrow B)(Y_2 \vee AY_3X_2 \vee AY_1X_3X_2 \vee \neg Y_1RX_3X_2 \leftrightarrow C) \tag{12}$$

Since models (10) and (12) cannot represent the same causal structure, the structure represented by (10) is unidentifiable even though it is fully S-expanded.

While stemmatic structures represent only a subset of the fully S-expanded structures that the INUS* theory fails to uniquely identify, they demonstrate that infinitely many such structures can be generated. This highlights the severity of the challenge to the INUS* theory's reductive aim. While positing that model (7) does not represent any part of the real world's causal structure was already *ad hoc* and implausible, denying that any stemmatic structure is part of the causal structure of the real world would be entirely untenable. The INUS* theory must be amended to restore its reductive aim.

4 The solution

4.1 The 3rd Non-Redundancy condition (NR3)

Fortunately, recent work by Zhang and Zhang (2025) has provided a promising candidate for a criterion that, when added to the INUS* theory, enables the unique identification of the triangle structure and other stemmatic structures. Zhang and Zhang have proposed this criterion, called the 3rd Non-Redundancy condition (NR3), independently of the triangle problem, with the goal of removing redundancies from NR2-models that are not yet eliminated by the non-redundancy criteria imposed by the INUS* theory (i.e. NR2 and structural minimality).

NR2 stipulated that conjuncts in conjunctive causes must be non-redundant for the sufficiency of the conjunction to which they belong and that conjunctions in causal antecedents must be non-redundant for the necessity of the disjunction to which they belong. Note that, since the antecedent of an NR2-biconditional is sufficient and necessary for the consequent of that biconditional, the set of factors in the antecedent *determines* the factor featured in the consequent:

Determination of a factor: A factor set $\mathbf{F}$ *determines* a factor B in a Humean mosaic δ iff, for every value combination of the factors in $\mathbf{F}$ that appears in at least one configuration in δ, all configurations in δ with this value combination have the same value for B (cf. Zhang and Zhang's, 2025, 7, definition of B being a partial function of $\mathbf{F}$).

NR3 adds to the stipulations of NR2 the requirement that, for any causal antecedent for a given consequent, no proper subset of the factors appearing in that antecedent determines the factor featured in that consequent. In other words, NR3 stipulates that the set of factors in a causal antecedent must be *functionally minimal*:

Functional minimality: A factor set $\mathbf{F}$ is *functionally minimal* with respect to a factor B in a Humean mosaic δ iff the following two conditions are met (Zhang and Zhang, 2025, 7):

1. $\mathbf{F}$ determines B in δ; and
2. there is no $\mathbf{F}' \subset \mathbf{F}$ such that $\mathbf{F}'$ determines B in δ.

NR3: An NR2-biconditional $\Pi \leftrightarrow B$ for a Humean mosaic δ, with $\mathbf{F}_\Pi$ the set of factors featured in Π, satisfies NR3 for δ iff $\mathbf{F}_\Pi$ is functionally minimal with respect to B in δ.

NR3-biconditional: An NR2-biconditional $\Pi \leftrightarrow B$ for a Humean mosaic δ is an *NR3-biconditional* for δ iff $\Pi \leftrightarrow B$ satisfies NR3 for δ.

NR3-model: An NR2-model $\mathcal{M}$ for a Humean mosaic δ is an *NR3-model* for δ iff all of $\mathcal{M}$'s NR2-biconditionals satisfy NR3 for δ.

To illustrate NR3, consider models (7) and (9). The exhaustive set of factors in model (7)'s NR2-antecedent for C is {Y_2, A, B, X_2}. As model (7)'s antecedent for C is sufficient and necessary for C, the factor set {Y_2, A, B, X_2} determines C in δ_3. Meanwhile, model (9)'s NR2-antecedent for C features factors {Y_2, A, B, X_2, Y_1}. Since {Y_2, A, B, X_2}, which determines C, is a proper subset of {Y_2, A, B, X_2, Y_1}, {Y_2, A, B, X_2, Y_1} is *not* functionally minimal (Y_1 is redundant in it for determining C). This implies that model (9)'s NR2-antecedent for C violates NR3, and thereby that model (9) itself violates NR3.

As model (9) violates NR3, adding Zhang and Zhang's (2025) criterion of NR3 to the INUS* theory would enable the unique identification of the triangle structure. Moreover, this addition would enable the unique identification of stemmatic structures and other fully S-expanded structures that cannot be uniquely identified by the INUS* theory.[6]

Nevertheless, before NR3 can be incorporated into the INUS* theory to uniquely identify stemmatic structures and thereby solve the triangle problem, the adequacy of NR3 as a criterion of INUS causation must be assessed. Zhang and Zhang (2025) argue for the incorporation of NR3 mainly through its capacity to eliminate ambiguities, and this alone is insufficient justification for regularity theorists: ambiguity reduction is only beneficial if the eliminated models fail to satisfy appropriate difference-making criteria, while—as I explain in the following section—a significant objection against NR3 is that its equivalent difference-making criterion appears to conflict with the principle of conjunctivity.

4.2 NR3 and difference-making

Recall that DM2, the difference-making criterion equivalent to NR2, requires that for any appearance of a factor value A in the NR2-antecedent for a consequent B, there is a pair of configurations $\{\sigma_i, \sigma_j\}$ such that A and B are given in σ_i and are not given in σ_j while any accompanying conjuncts are present and any alternative conjunctions are absent in both σ_i and σ_j (i.e. a DM2-pair). For example, consider Table 3, which displays all the configurations in δ_3 that qualify as members of DM2-pairs for Y_1 in $Y_2 \vee AY_1X_2 \vee \neg Y_1BX_2 \leftrightarrow C$ in model (9). In each of these configurations, accompanying conjuncts A and X_2 are

[6]While a formal proof that NR3 allows for the unique identification of all fully S-expanded structures remains outstanding, extensive theoretical analysis and computational simulations have not yielded any counterexamples to this conjecture.

	Y_1	C	Y_2	A	X_2	B	X_1
σ_1	1	1	0	1	1	1	1
σ_2	1	1	0	1	1	1	0
σ_3	0	0	0	1	1	0	0

Table 3: All the configurations in δ_3 that qualify as members of DM2-pairs for Y_1 in $Y_2 \vee AY_1X_2 \vee \neg Y_1BX_2 \leftrightarrow C$ in model (9).

present and alternative conjunctions Y_2 and $\neg Y_1BX_2$ are absent. Furthermore, configurations σ_1 and σ_2 both feature Y_1 and C, and configuration σ_3 features neither Y_1 nor C. Therefore, $\{\sigma_1, \sigma_3\}$ and $\{\sigma_2, \sigma_3\}$ are the DM2-pairs for Y_1 in $Y_2 \vee AY_1X_2 \vee \neg Y_1BX_2 \leftrightarrow C$.

The difference-making requirement equivalent to NR3, referred to as the criterion of *NR3-Difference Making (DM3)* (Zhang and Zhang, 2025), involves a stricter notion of a difference-making pair than DM2 does: DM3 requires that all factors besides Y_1 in antecedent $Y_2 \vee AY_1X_2 \vee \neg Y_1BX_2$ (i.e. Y_2, A, X_2, and B) are constant across difference-making pairs. Note that neither $\{\sigma_1, \sigma_3\}$ nor $\{\sigma_2, \sigma_3\}$ satisfies this stricter notion of a difference-making pair imposed by DM3, because B varies across both these pairs.

To express DM3 more precisely, let $A\phi_1 \vee \phi_2 \vee ... \vee \phi_n$ be a disjunction of conjunctions of factor values, let $\mathbf{F}_\phi$ be the set of factors in $A\phi_1 \vee \phi_2 \vee ... \vee \phi_n$, and let B be a single factor value. Additionally, I refer to the factors in $\mathbf{F}_\phi \backslash \{\mathrm{A}\}$ as the *fellow causing factors* of A. An NR2-biconditional satisfies NR3 for a Humean mosaic iff that biconditional satisfies DM3 for that mosaic (Zhang and Zhang, 2025)[7]:

NR3-Difference Making (DM3): An NR2-biconditional $A\phi_1 \vee \phi_2 \vee ... \vee \phi_n \leftrightarrow B$ for a Humean mosaic δ satisfies DM3 for δ iff δ contains at least one *NR3-Difference-Making pair (DM3-pair)* for at least one value of each factor in $\mathbf{F}_\phi$.[8]

NR3-Difference-Making pair (DM3-pair): An *NR3-Difference-Making pair (DM3-pair)* for A in biconditional $A\phi_1 \vee \phi_2 \vee ... \vee \phi_n \leftrightarrow B$ is a pair of configurations $\{\sigma_i, \sigma_j\}$ such that:

1. A and B are given in σ_i and are not given in σ_j; and
2. any factor in $\mathbf{F}_\phi \backslash \{\mathrm{A}\}$ (i.e. each of A's fellow causing factors) has the same value in σ_i and σ_j.

[7] My formulation of DM3 differs slightly from Zhang and Zhang's (2025, 13) by explicitly presenting DM3 as an additional requirement on biconditionals that already satisfy NR2. This modification prevents cases in which biconditionals could satisfy DM3 while violating DM2—an unintended possibility under Zhang and Zhang's original formulation, exemplified by biconditionals like $AB \vee \neg A \neg B \vee \neg AC \vee BC \leftrightarrow E$. While requiring biconditionals to already satisfy NR2 represents a technical departure from Zhang and Zhang's formulation, my formulation better serves Zhang and Zhang's (2025, 13) expressed purpose of formulating a difference-making criterion strictly stronger than DM2.

[8] As Zhang and Zhang (2025, 12-14) point out, it is possible that, for some factor in $\mathbf{F}_\phi$, δ contains a DM3-pair for one but not both values of that factor.

As Zhang and Zhang (2025, 14) themselves acknowledge, an objection to DM3 is that DM2 is already sufficiently strict as a difference-making criterion in the INUS framework because, according to this framework, it suffices to have accompanying conjuncts present and alternative conjunctions absent across a difference-making pair (as DM2 requires) to ensure that the variation in the effect across a difference-making pair is attributable to the variation in the considered candidate cause (A in our definition).

Additionally requiring that fellow causing factors are constant (as DM3 does) is only motivated when presupposing that variation in an individual fellow causing factor such as B in model (9) can account for variation in the effect without changing the truth value of any alternative conjunctions like $\neg Y_1 B X_2$ in model (9). But this presupposition seems to conflict with the principle of conjunctivity, which states that all conjuncts in a conjunction of causes must occur together to bring about the effect, implying that variation in an effect can only be attributed to a difference in value of an individual fellow causing factor if this difference changes the truth value of a conjunction that features this factor.

Although Zhang and Zhang (2025, 14) acknowledge this objection, even granting that it is a 'reasonable response' to the proposal of adding the requirement that individual causing factors are kept constant across a difference-making pair (as DM3 does), they neither provide a convincing refutation nor conclude that DM3 is less appropriate than DM2. Instead, they suggest that the need for a causing factor to make a difference to its effect while fellow causing factors are constant—which they claim to be highly intuitive but which an INUS theorist might find unappealing based precisely on the principle of conjunctivity—simply outweighs the need for a difference-making criterion to align with the principle of conjunctivity. By neglecting to refute this objection convincingly, Zhang and Zhang leave unresolved the question of whether adopting NR3 (and thus the equivalent DM3) in the INUS* theory can lead to an adequate theory of causation.

4.3 DM3 reformulated

Fortunately, the apparent conflict between DM3 and the principle of conjunctivity dissolves upon closer examination of DM3. While DM3 does allow individual causing factors to account for variation in their effects independently of the alternative conjunctions in which they appear *in a given model*, deriving from this that DM3 conflicts with conjunctivity misses a crucial point: DM3 still requires that variation in an individual causing factor can account for variation in an effect only through variation in *some* conjunction that features this individual causing factor, just not in an alternative conjunction that appears in the given model.

What DM3 actually requires by holding all fellow causing factors constant across difference-making pairs is that *all* sufficient conjunctions composed exclusively of these fellow causing factors remain constant across these pairs. To make this explicit, I formulate a difference-making criterion called *NR3-Difference-Making′ (DM3′)* which, as proven in Appendix B, is equivalent to DM3, and which differs from DM3 in the second criterion of a DM3′-pair. Again, let $A\phi_1 \vee \phi_2 \vee ... \vee \phi_n$ be a disjunction of conjunctions of factor values, let $\mathbf{F}_\phi$ be the set of factors in $A\phi_1 \vee \phi_2 \vee ... \vee \phi_n$, and let B be a single factor value.

NR3-Difference Making′ (DM3′): An NR2-biconditional $A\phi_1 \vee \phi_2 \vee ... \vee \phi_n \leftrightarrow B$ for a Humean mosaic δ satisfies DM3′ for δ iff δ contains at least one *NR3-Difference-Making′ pair (DM3′-pair)* for at least one value of each factor in $\mathbf{F}_\phi$.

NR3-Difference-Making′ pair (DM3′-pair): An *NR3-Difference-Making′ pair (DM3′-pair)* for A in biconditional $A\phi_1 \vee \phi_2 \vee ... \vee \phi_n \leftrightarrow B$ is a pair of configurations $\{\sigma_i, \sigma_j\}$ such that:

1. A and B are given in σ_i and are not given in σ_j; and
2. any (minimally) sufficient conjunctions for B consisting exclusively of factors in $\mathbf{F}_\phi \backslash \{\text{A}\}$ are absent in σ_i and σ_j.

The equivalence of DM3′ to DM3, and therefore to NR3, reveals that motivating NR3/DM3 is fully compatible with the principle of conjunctivity. Like DM2, DM3 requires certain (minimally) sufficient conjunctions to be constantly absent across difference-making pairs, adhering to the principle that individual causes only bring about their effects when each component in a conjunction to which they belong is instantiated.

The key difference between DM2 and DM3 is that DM3 considers a broader range of conjunctions when evaluating whether changes in the effect can be attributed to A rather than to A's fellow causing factors. While DM2 only considers the minimally sufficient conjunctions that are explicitly present in the biconditional under consideration, DM3 considers all minimally sufficient conjunctions that can be formed exclusively from fellow causing factors.

This broadening of the range of considered conjunctions constitutes a meaningful improvement to the INUS* theory rather than a departure from its principles. Since the theory aims to characterize how causal structures are determined by Humean mosaics rather than defining these structures in isolation, it is appropriate to consider *every* conjunction formed exclusively from fellow causing factors that is minimally sufficient within a given Humean mosaic when determining whether a factor's values are causally relevant in that mosaic.

4.4 A new definition of causal relevance

Having refuted the objection to NR3, it is now clear how NR3 combines eliminating redundant elements from causal structures with imposing an adequate difference-making criterion. This makes NR3 an effective solution to the triangle problem, warranting that it is integrated into the INUS* theory's definition of causal relevance. This integration results in the following definition:

Causal relevance′: Let Ω be a disjunction of conjunctions of factor values and B a single factor value. Ω is *causally relevant* to B iff there exists a factor frame **F** which is a set of (modally independent) factors including B and all factors featured in Ω, with δ a Humean mosaic over **F**, such that the following two conditions are satisfied:

1. there exists an NR3-biconditional $\Pi \leftrightarrow B$ for δ that is part of an NR3-model for δ, such that Ω is contained in Π; and
2. for every factor frame expansion $\mathbf{F}' \supset \mathbf{F}$ with mosaic δ': there exists an NR3-biconditional $\Pi' \leftrightarrow B$ for δ' that is part of an NR3-model for δ', such that Ω is contained in Π'.

5 Conclusion

This paper has demonstrated a detrimental limitation of the INUS* theory of causation as presented by Baumgartner and Falk (2023a): its inability to uniquely identify stemmatic causal structures based on Humean mosaics, while this failure cannot be attributed to the simplicity of these structures. This limitation reveals that the theory fails in its reductive aim of defining causation purely in terms of regularities. However, I have shown that incorporating Zhang and Zhang's (2025) NR3 criterion alongside the INUS* theory's existing criteria enables the unique identification of stemmatic structures. Furthermore, I have refuted a significant objection against NR3—this criterion's alleged conflict with the principle of conjunctivity. By demonstrating how NR3's difference-making criterion can be reformulated to explicitly align with conjunctivity, I have established NR3 as an adequate addition to the INUS* theory, thereby restoring this theory's reductive aim.

Acknowledgements

I thank Michael Baumgartner for his invaluable feedback during the development of this article and the underlying research. Furthermore, I am very grateful to Timm Lampert for inspiring this research by sharing the example causal structure that motivated this work. I also thank Sander Beckers for extensive comments on an earlier draft, and Jiji Zhang for helpful discussions. Finally, I am grateful to the participants in the second annual conference of the Society for Philosophy of Causation in Göttingen, the Bergen Philosophy of Science Workshop 2024, and the Departmental Seminar at the University of Bergen Philosophy Department for helpful feedback.

References

H. Andreas and M. Günther. Regularity and inferential theories of causation. In E. N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Fall 2021 edition,

2021. URL `https://plato.stanford.edu/archives/fall2021/entries/causation-regularity/`.

M. Baumgartner and C. Falk. Boolean difference-making: A modern regularity theory of causation. *The British Journal for the Philosophy of Science*, 74(1): 171–197, 2023a. doi: 10.1093/bjps/axz047.

M. Baumgartner and C. Falk. Quantifying the quality of configurational causal models. *Journal of Causal Inference*, 12(1), 2024. doi: 10.1515/jci-2023-0032.

C. Falk, M. Ambühl, and M. Baumgartner. ***causalHyperGraph**: Drawing Causal Hypergraphs*. R package version 0.1.0, 2024. `https://cran.r-project.org/package=causalHyperGraph`.

C. A. Frosch and P. N. Johnson-Laird. Is everyday causation deterministic or probabilistic? *Acta Psychologica*, 137(3):280–291, 2011. doi: 10.1016/j.actpsy.2011.01.015.

J. Gerring. Causation: A unified framework for the social sciences. *Journal of Theoretical Politics*, 17(2):163–198, 2005. doi: 10.1177/0951629805050859.

G. Graßhoff and M. May. Causal regularities. In W. Spohn, M. Ledwig, and M. Esfeld, editors, *Current Issues in Causation*, pages 85–114. Mentis, 2001.

T. Haesebrouck. EU member state support to Ukraine. *Foreign Policy Analysis*, 20(2), 2024. doi: 10.1093/fpa/orae005.

N. Hall. Two concepts of causation. In J. Collins, N. Hall, and L. A. Paul, editors, *Causation and Counterfactuals*, pages 225–276. The MIT Press, 2004.

D. Hume. *An Enquiry Concerning Human Understanding and Other Essays*. Washington Square Press, New York, 1748 (1963).

J. Kim. Causes and events: Mackie on causation. *The Journal of Philosophy*, 68(14):426–441, 1971. doi: 10.2307/2025175.

D. Lewis. Causation. *The Journal of Philosophy*, 70(17):556–567, 1973. doi: 10.2307/2025310.

D. Lewis. *Philosophical Papers II*. Oxford University Press, Oxford, 1986.

J. L. Mackie. *The Cement of the Universe: A Study of Causation*. Clarendon Press, Oxford, 1974.

M. May. *Kausales Schliessen: Eine Untersuchung über kausale Erklärungen und Theorienbildung*. PhD, Universität Hamburg, Hamburg, 1999.

J. S. Mill. *A System of Logic*. John W. Parker, London, 1843.

S. Mumford and R. Anjum. A powerful theory of causation. In A. Marmodoro, editor, *The Metaphysics of Powers: Their Grounding and Their Manifestations*, pages 143–159. Routledge, 2010.

J. Pearl. *Causality: Models, Reasoning, and Inference.* Cambridge University Press, New York, 2009.

B. Rihoux and C. C. Ragin, editors. *Configurational Comparative Methods: Qualitative Comparative Analysis (QCA) and Related Techniques.* Applied Social Research Methods. Sage, Los Angeles, 2009.

A. Rothe, B. Deverett, R. Mayrhofer, and C. Kemp. Successful structure learning from observational data. *Cognition*, 179:266–297, 2018. doi: 10.1016/j.cognition.2018.06.003.

K. J. Rothman. Causes. *American Journal of Epidemiology*, 104(6):587–592, 1976. doi: 10.1093/oxfordjournals.aje.a112335.

L. E. Schulz and J. Sommerville. God does not play dice: Causal determinism and preschoolers' causal inferences. *Child Development*, 77(2):427–442, 2006. doi: 10.1111/j.1467-8624.2006.00880.x.

P. Spirtes, C. Glymour, and R. Scheines. *Causation, Prediction, and Search.* The MIT Press, Cambridge, 2000.

P. Suppes. *A Probabilistic Theory of Causality.* Acta philosophica Fennica. North Holland Publishing Company, Amsterdam, 1970.

J. Woodward. *Making Things Happen: A Theory of Causal Explanation.* Oxford studies in philosophy of science. Oxford University Press, New York, 2003.

V. Yakovchenko, E. J. Miech, M. J. Chinman, M. Chartier, et al. and S. S. Strategy configurations directly linked to higher hepatitis C virus treatment starts: An applied use of configurational comparative methods. *Medical Care*, 58(5):e31–e38, 2020. doi: 10.1097/MLR.0000000000001319.

J. Zhang and K. Zhang. A new minimality condition for Boolean accounts of causal regularities. *Erkenntnis*, 90:67–86, 2025. doi: 10.1007/s10670-023-00685-4.

A Independent variation in full S-expansions

This appendix establishes that, in any acyclic full S-expansion, each endogenous factor varies independently of its non-specific causing factors in the Humean mosaic of that full S-expansion.

Section 3.1 explained why endogenous factor E varies independently of its non-specific causing factors A and B in the mosaic of model (6). This independence arises provided that, for every empirically possible value combination of E's non-specific causing factors and every logically possible value combination of E's specific causing factors, this mosaic contains a configuration featuring these two value combinations together. The logically possible value combinations of

E's specific causing factors include the combination in which no specific factor value in E's antecedent is given (such that E is not given), as well as combinations in which the disjunct comprising only specific factors in this antecedent is present (such that E is given). Therefore, E's independent variation from its non-specific causing factors is guaranteed when the set comprising E's specific factors *varies independently* of the set comprising E's non-specific factors:

Independent variation (factor set of factor set): A factor set $\mathbf{F}_1$ *varies independently* of a factor set $\mathbf{F}_2$ in a Humean mosaic δ iff, for every combination of values of the factors in $\mathbf{F}_2$ that appears in δ and every logically possible value combination of the factors in $\mathbf{F}_1$, δ contains at least one configuration featuring these two value combinations together.

In model (6), such independent variation is guaranteed because all of E's causing factors are exogenous. However, this is not always the case in a full S-expansion. Therefore, what remains to be established in order to demonstrate that any endogenous factor in a full S-expansion varies independently of its non-specific causing factors in the mosaic of that full S-expansion is that, for any endogenous factor in any full S-expansion (even endogenous factors that have endogenous causing factors), the set comprising this endogenous factor's specific causing factors varies independently of the set comprising this endogenous factor's non-specific causing factors.

To establish that this holds, consider any full S-expansion $\mathcal{M}$ with Humean mosaic δ over factor frame $\mathbf{F}_{\mathcal{M}}$. Let factor W $\in \mathbf{F}_{\mathcal{M}}$ be an endogenous factor with respect to $\mathcal{M}$. Furthermore, let $\mathbf{F}_{\text{sp}} = \{\text{C}_1, ..., \text{C}_n\} \subset \mathbf{F}_{\mathcal{M}}$ be the factor set comprising all specific causing factors of W with respect to $\mathcal{M}$, and let $\mathbf{F}_{\text{nsp}} \subset \mathbf{F}_{\mathcal{M}}$ refer to the factor set comprising all non-specific causing factors of W with respect to $\mathcal{M}$. The goal is then to demonstrate that $\mathbf{F}_{\text{sp}}$ varies independently of $\mathbf{F}_{\text{nsp}}$ in δ. I achieve this in two steps. First, I show that each factor in $\mathbf{F}_{\text{sp}}$ varies independently of the set comprising all other factors in $\mathbf{F}_{\text{sp}}$ together with all factors in $\mathbf{F}_{\text{nsp}}$, for instance, C_1 varies independently of $(\mathbf{F}_{\text{sp}} \backslash \{\text{C}_1\}) \cup \mathbf{F}_{\text{nsp}}$. Second, I establish that, given this first result, $\mathbf{F}_{\text{sp}}$ varies independently of $\mathbf{F}_{\text{nsp}}$.

Recall from Section 3.1 that all exogenous factors with respect to a causal structure vary independently of each other in the Humean mosaic of that structure because nothing in the causal structure implies that any combination of their values is empirically impossible. More generally, any factor X varies independently of any factor set *not* containing a factor Y such that any of the following relationships holds:

1. X is an ancestor of Y;
2. Y is an ancestor of X; or
3. X and Y have a common ancestor.

Unless one of these relationships holds between X and a factor in a given factor set, a causal structure imposes no restrictions preventing any value of X from occurring together with any empirically possible combination of values from this factor set.

Therefore, C_1 varies independently of $(\mathbf{F}_{sp}\backslash\{C_1\})\cup\mathbf{F}_{nsp}$ in δ if $(\mathbf{F}_{sp}\backslash\{C_1\})\cup\mathbf{F}_{nsp}$ does not contain any factor D such that any of the following relationships holds:

1. C_1 is an ancestor of D;
2. D is an ancestor of C_1; or
3. C_1 and D have a common ancestor.

To show that none of these three relationships holds between C_1 and any factor in $(\mathbf{F}_{sp}\backslash\{C_1\})\cup\mathbf{F}_{nsp}$, I start by establishing that (1.) C_1 cannot be an ancestor of any factor in $(\mathbf{F}_{sp}\backslash\{C_1\})\cup\mathbf{F}_{nsp}$. As factors in $\mathbf{F}_{sp}\backslash\{C_1\}$ are specific factors and therefore, by definition, exogenous, such that they do not have any ancestors, C_1 cannot be an ancestor of any factor in $\mathbf{F}_{sp}\backslash\{C_1\}$. Furthermore, since C_1 is a specific factor, C_1 appears in only one causal antecedent in $\mathcal{M}$ (here: that of a value of W). Hence, C_1 can be an ancestor of a factor in $\mathbf{F}_{nsp}$ only through W being an ancestor of that factor in $\mathbf{F}_{nsp}$ (not through C_1 causing some factor other than W which is an ancestor of that factor in $\mathbf{F}_{nsp}$). Furthermore, factors in $\mathbf{F}_{nsp}$ are non-specific causing factors of W and therefore, by definition, ancestors of W. Therefore, C_1 can be an ancestor of a factor in $\mathbf{F}_{nsp}$ only if W is an ancestor of that factor in $\mathbf{F}_{nsp}$ while that factor in $\mathbf{F}_{nsp}$ is also an ancestor of W, which is prohibited by the assumption of acyclicity. As C_1 can neither be an ancestor of a factor in $\mathbf{F}_{sp}\backslash\{C_1\}$ nor be an ancestor of a factor in $\mathbf{F}_{nsp}$, C_1 cannot be an ancestor of any factor in $(\mathbf{F}_{sp}\backslash\{C_1\})\cup\mathbf{F}_{nsp}$. Furthermore, as C_1 is a specific factor and therefore, by definition, does not have any ancestors, no factor in $(\mathbf{F}_{sp}\backslash\{C_1\})\cup\mathbf{F}_{nsp}$ can be (2.) an ancestor of C_1 or (3.) share an ancestor with C_1.

So, as there cannot be a factor D in $(\mathbf{F}_{sp}\backslash\{C_1\})\cup\mathbf{F}_{nsp}$ such that C_1 is an ancestor of D, D is an ancestor of C_1, or C_1 and D have a common ancestor, C_1 varies independently of $(\mathbf{F}_{sp}\backslash\{C_1\})\cup\mathbf{F}_{nsp}$ in δ. Analogously, C_2 varies independently of $(\mathbf{F}_{sp}\backslash\{C_2\})\cup\mathbf{F}_{nsp}$ in δ, and so on. In sum, each specific causing factor of an endogenous factor varies independently of the set comprising all other specific causing factors and all non-specific causing factors of that endogenous factor in the mosaic of a full S-expansion.

This fact allows to demonstrate that $\mathbf{F}_{sp}$ varies independently of $\mathbf{F}_{nsp}$ in δ. Consider an arbitrary configuration σ_1 in δ. For illustrative purposes, let us represent the values that the factors C_1, C_2, and C_3 take in σ_1 by c_1, c_2, and c_3, respectively. So, $C_1 = c_1$, $C_2 = c_2$, and $C_3 = c_3$ in σ_1, where c_1, c_2, and c_3 each represent either the value 1 or the value 0. Given that C_1 varies independently of $(\mathbf{F}_{sp}\backslash\{C_1\})\cup\mathbf{F}_{nsp}$ in δ, δ contains another observation—let us call it σ_2—with the values of the factors in $(\mathbf{F}_{sp}\backslash\{C_1\})\cup\mathbf{F}_{nsp}$ the same as in σ_1, but in which $C_1 \neq c_1$ (as per the definition of a factor varying independently of a factor set). Thus, σ_2 has $C_1 \neq c_1$ and $C_2 = c_2$ while featuring the same value combination of the factors in $(\mathbf{F}_{sp}\backslash\{C_1, C_2\})\cup\mathbf{F}_{nsp}$ as σ_1 does.

Furthermore, given that δ contains σ_1 and σ_2, C_2's independent variation of $(\mathbf{F}_{sp}\backslash\{C_2\})\cup\mathbf{F}_{nsp}$ in δ implies that δ contains at least two more configurations in addition to σ_1 and σ_2. First, δ must contain a configuration—let us call it σ_3—with the same value combination of the factors in $(\mathbf{F}_{\mathrm{sp}}\backslash\{\mathrm{C}_2\})\cup\mathbf{F}_{\mathrm{nsp}}$ as σ_1 but with a different value of C_2 than the one in σ_1. This means that σ_3 has $\mathrm{C}_1 = c_1$ and $\mathrm{C}_2 \neq c_2$ while featuring the same value combination of the factors in $(\mathbf{F}_{\mathrm{sp}}\backslash\{\mathrm{C}_1, \mathrm{C}_2\})\cup\mathbf{F}_{\mathrm{nsp}}$ as σ_1 does.

Second, δ must contain a configuration—let us call it σ_4—with the same value combination of the factors in $(\mathbf{F}_{\mathrm{sp}}\backslash\{\mathrm{C}_2\})\cup\mathbf{F}_{\mathrm{nsp}}$ as σ_2 but with a different value of C_2 than the one in σ_2. So, σ_4 has $\mathrm{C}_1 \neq c_1$ and $\mathrm{C}_2 \neq c_2$ while featuring the same value combination of the factors in $(\mathbf{F}_{\mathrm{sp}}\backslash\{\mathrm{C}_1, \mathrm{C}_2\})\cup\mathbf{F}_{\mathrm{nsp}}$ as σ_2 does. Importantly, this means that σ_1, σ_2, σ_3, and σ_4 each have the same value combination of $(\mathbf{F}_{\mathrm{sp}}\backslash\{\mathrm{C}_1, \mathrm{C}_2\})\cup\mathbf{F}_{\mathrm{nsp}}$.

In sum, if C_1 varies independently of $(\mathbf{F}_{\mathrm{sp}}\backslash\{\mathrm{C}_1\})\cup\mathbf{F}_{\mathrm{nsp}}$ and C_2 varies independently of $(\mathbf{F}_{\mathrm{sp}}\backslash\{\mathrm{C}_2\})\cup\mathbf{F}_{\mathrm{nsp}}$ in δ, then δ must contain $2\cdot 2 = 4$ configurations with the same value combination of $(\mathbf{F}_{\mathrm{sp}}\backslash\{\mathrm{C}_1, \mathrm{C}_2\})\cup\mathbf{F}_{\mathrm{nsp}}$ and with distinct value combinations of C_1 and C_2, that is, $\mathrm{C}_1 = c_1$, $\mathrm{C}_2 = c_2$ (in σ_1); $\mathrm{C}_1 \neq c_1$, $\mathrm{C}_2 = c_2$ (in σ_2); $\mathrm{C}_1 = c_1$, $\mathrm{C}_2 \neq c_2$ (in σ_3); and $\mathrm{C}_1 \neq c_1$, $\mathrm{C}_2 \neq c_2$ (in σ_4).

Similarly, since C_3 varies independently of $(\mathbf{F}_{\mathrm{sp}}\backslash\{\mathrm{C}_3\})\cup\mathbf{F}_{\mathrm{nsp}}$, δ must contain four configurations in addition to the four configurations σ_1, σ_2, σ_3, and σ_4. These new configurations must match σ_1, σ_2, σ_3, and σ_4 in their values of the factors in $(\mathbf{F}_{\mathrm{sp}}\backslash\{\mathrm{C}_1, \mathrm{C}_2, \mathrm{C}_3\})\cup\mathbf{F}_{\mathrm{nsp}}$ but must have $\mathrm{C}_3 \neq c_3$. Consequently, δ must contain $(2\cdot 2)\cdot 2 = 2^3 = 8$ configurations with the same value combination of the factors in $(\mathbf{F}_{\mathrm{sp}}\backslash\{\mathrm{C}_1, \mathrm{C}_2, \mathrm{C}_3\})\cup\mathbf{F}_{\mathrm{nsp}}$ and with distinct value combinations of C_1, C_2, and C_3.

The same reasoning applies to C_4, C_5, and so on, until C_n. Overall, when each factor in $\mathbf{F}_{\mathrm{sp}} = \{\mathrm{C}_1, \mathrm{C}_2, ..., \mathrm{C}_n\}$ varies independently of the set comprising the other factors in $\mathbf{F}_{\mathrm{sp}}$ together with the factors in $\mathbf{F}_{\mathrm{nsp}}$, then δ must contain 2^n configurations with distinct value combinations of the factors in $\mathbf{F}_{\mathrm{sp}}$ and with the same value combination of the factors in $\mathbf{F}_{\mathrm{nsp}}$.

Since $\mathbf{F}_{\mathrm{sp}} = \{\mathrm{C}_1, \mathrm{C}_2, ..., \mathrm{C}_n\}$ contains n binary factors, there exist 2^n distinct logically possible value combinations of the factors in $\mathbf{F}_{\mathrm{sp}}$. Therefore, if δ contains 2^n configurations with distinct value combinations of the factors in $\mathbf{F}_{\mathrm{sp}}$ and with the same value combination of the factors in $\mathbf{F}_{\mathrm{nsp}}$, then δ contains, for each logically possible value combination of the factors in $\mathbf{F}_{\mathrm{sp}}$, at least one configuration with this value combination of the factors in $\mathbf{F}_{\mathrm{sp}}$ together with this same value combination of the factors in $\mathbf{F}_{\mathrm{nsp}}$.

This reasoning generalizes to every value combination of the factors in $\mathbf{F}_{\mathrm{nsp}}$ that appears in δ. This demonstrates that, for every value combination of the factors in $\mathbf{F}_{\mathrm{nsp}}$ that appears in δ and every logically possible value combination of the factors in $\mathbf{F}_{\mathrm{sp}}$, δ contains at least one configuration featuring these two value combinations together. Therefore, $\mathbf{F}_{\mathrm{sp}}$ varies independently of $\mathbf{F}_{\mathrm{nsp}}$ in this mosaic.

As established above, this result implies that any endogenous factor in an acyclic full S-expansion varies independently of its non-specific causing factors in the mosaic of that full S-expansion.

B Equivalence of DM3 and DM3′

This appendix provides a two-part proof of the equivalence between DM3 and DM3′.[9] First, I show that any NR2-biconditional satisfying DM3 in a Humean mosaic must also satisfy DM3′ in that mosaic:

Proof. Consider the NR2-biconditional $A\phi_1 \vee \phi_2 \vee ... \vee \phi_n \leftrightarrow B$ for a Humean mosaic δ, with $\mathbf{F}_\phi$ the exhaustive set of factors appearing in $A\phi_1 \vee \phi_2 \vee ... \vee \phi_n$. Consider DM3-pair $\{\sigma_i, \sigma_j\}$ for A and B in δ. Since B is not given in σ_j, any sufficient conjunctions for B, and therefore any minimally sufficient conjunctions for B consisting exclusively of factors in $\mathbf{F}_\phi \backslash \{\text{A}\}$, are absent in σ_j. Furthermore, as all factors in $\mathbf{F}_\phi \backslash \{\text{A}\}$ have the same value in σ_i and σ_j, all conjunctions consisting exclusively of factors in $\mathbf{F}_\phi \backslash \{\text{A}\}$ have the same value in σ_i and σ_j. Since all minimally sufficient conjunctions for B consisting exclusively of factors in $\mathbf{F}_\phi \backslash \{\text{A}\}$ are absent in σ_j and have the same value in σ_i and σ_j, all minimally sufficient conjunctions for B consisting exclusively of factors in $\mathbf{F}_\phi \backslash \{\text{A}\}$ are absent in σ_i and σ_j. This implies that every DM3-pair is also a DM3′-pair, such that any NR2-biconditional that has a DM3-pair for at least one value of each factor in its antecedent also has a DM3′-pair for at least one value of each factor in its antecedent. □

Second, and crucially for the INUS framework, I will show that any NR2-biconditional that violates DM3 in a Humean mosaic must also violate DM3′ in that mosaic. This establishes that when an NR2-biconditional fails to satisfy DM3, it contains at least one factor for which all DM2-pairs involve variation in the truth value of a (minimally) sufficient conjunction made up exclusively of that factor's fellow causing factors.

Proof. Consider the NR2-biconditional $A\phi_1 \vee \phi_2 \vee ... \vee \phi_n \leftrightarrow B$ for a Humean mosaic δ, with $\mathbf{F}_\phi$ the exhaustive set of factors appearing in $A\phi_1 \vee \phi_2 \vee ... \vee \phi_n$.

Suppose that $A\phi_1 \vee \phi_2 \vee ... \vee \phi_n \leftrightarrow B$ violates DM3. Then there is at least one factor appearing in the antecedent of this biconditional—let us take A as an example—such that δ contains no DM3-pairs for either of A's values with respect to this biconditional. This implies, first, that δ contains no pair of configurations with all factors in $\mathbf{F}_\phi \backslash \{\text{A}\}$ constant and A and B both given in one configuration and both not given in the other (which would constitute a DM3-pair for A). Second, it implies that δ contains no pair of configurations with all factors in $\mathbf{F}_\phi \backslash \{\text{A}\}$ constant and $\neg A$ and B both featured in one and both not featured in the other (which would constitute a DM3-pair for $\neg A$). Therefore,

(i) δ does not contain a pair of configurations across which all factors in $\mathbf{F}_\phi \backslash \{\text{A}\}$ are constant while A and B both vary.

Furthermore, as $A\phi_1 \vee \phi_2 \vee ... \vee \phi_n \leftrightarrow B$ is an NR2-biconditional and therefore, by definition, true in δ, δ does not contain a pair of configurations across which

[9] Parts of the upcoming proof closely follow strategies used by Zhang and Zhang (2025).

all factors in $\mathbf{F}_\phi$ are constant (such that the truth value of $A\phi_1 \vee \phi_2 \vee ... \vee \phi_n$ is constant) while B varies, since the truth value of the consequent of the biconditional would differ from that of the antecedent in one configuration of such a pair. As $\mathbf{F}_\phi$ can be rewritten as $(\mathbf{F}_\phi \backslash \{A\}) \cup \{A\}$, this can be reformulated as follows:

(ii) δ does not contain a pair of configurations across which all factors in $\mathbf{F}_\phi \backslash \{A\}$ are constant and A is constant while B varies.

Since any pair of configurations must have either A varying or A constant, statements (i) and (ii) above imply the following:

(iii) δ does not contain a pair of configurations across which all factors in $\mathbf{F}_\phi \backslash \{A\}$ are constant while B varies.

This entails that $\mathbf{F}_\phi \backslash \{A\}$ determines B in δ.

I now prove that, if $\mathbf{F}_\phi \backslash \{A\}$ determines B in δ, then δ cannot contain a DM3′-pair for any of A's values for $A\phi_1 \vee \phi_2 \vee ... \vee \phi_n \leftrightarrow B$. First, since $\mathbf{F}_\phi \backslash \{A\}$ determines B in δ, there exists an expression involving only factors in $\mathbf{F}_\phi \backslash \{A\}$ that is both sufficient and necessary for B in δ. This expression can be equivalently rewritten as a disjunction of conjunctions Π, for which the following three claims hold:

1. each disjunct in Π is a conjunction featuring only factors from $\mathbf{F}_\phi \backslash \{A\}$ (because Π is a disjunction of conjunctions featuring only factors from $\mathbf{F}_\phi \backslash \{A\}$);
2. each disjunct in Π is sufficient for B in δ (since Π is sufficient for B in δ); and
3. every configuration in δ in which B is given must have at least one disjunct of Π present (as Π is necessary for B in δ).

From these three points, it follows that every configuration in δ in which B is given must contain at least one sufficient conjunction for B featuring only factors from $\mathbf{F}_\phi \backslash \{A\}$.

Furthermore, whenever a sufficient conjunction for B featuring only factors from $\mathbf{F}_\phi \backslash \{A\}$ is present, a *minimally* sufficient conjunction for B featuring only factors from $\mathbf{F}_\phi \backslash \{A\}$ is present, as any non-minimally sufficient conjunction yields a minimally sufficient conjunction when conjuncts redundant for its sufficiency are removed.

Therefore, if δ contains no DM3-pairs for (any value of) A, then any configuration in δ in which B is given must have at least one minimally sufficient conjunction for B featuring only factors from $\mathbf{F}_\phi \backslash \{A\}$ present. This fact implies that δ contains no DM3′-pair for A, since, by definition, any DM3′-pair for A contains a configuration with B and without any minimally sufficient conjunction for B featuring only factors from $\mathbf{F}_\phi \backslash \{A\}$. As δ cannot contain any DM3′-pairs for A if it contains no DM3-pairs for A, it follows that any NR2-biconditional that violates DM3 also violates DM3′. □