



Published in final edited form as:
Ann Stat. 2012 ; 40(4): 2128–2161.

GENERAL THEORY FOR INTERACTIONS IN SUFFICIENT CAUSE MODELS WITH DICHOTOMOUS EXPOSURES

Tyler J. VanderWeele and

Harvard School of Public Health, Department of Epidemiology, 677 Huntington Avenue, Boston, MA 02115, tvanderw@hsph.harvard.edu, URL: <http://www.hsph.harvard.edu/faculty/tyler-vanderweele/>

Thomas S. Richardson

University of Washington, Department of Statistics, Box 354322, Seattle, WA 98195, thomasr@uw.edu

Abstract

The sufficient-component cause framework assumes the existence of sets of sufficient causes that bring about an event. For a binary outcome and an arbitrary number of binary causes any set of potential outcomes can be replicated by positing a set of sufficient causes; typically this representation is not unique. A sufficient cause interaction is said to be present if within all representations there exists a sufficient cause in which two or more particular causes are all present. A singular interaction is said to be present if for some subset of individuals there is a unique minimal sufficient cause. Empirical and counterfactual conditions are given for sufficient cause interactions and singular interactions between an arbitrary number of causes. Conditions are given for cases in which none, some or all of a given set of causes affect the outcome monotonically. The relations between these results, interactions in linear statistical models and Pearl's probability of causation are discussed.

Keywords

causal inference; counterfactual; epistasis; interaction include keywords that are in title; potential outcomes; synergism

1. Introduction

Rothman's sufficient-component cause model [21] postulates a set of different causal mechanisms each sufficient to bring about the outcome under consideration. Rothman refers to these hypothesized causal mechanisms as "sufficient causes", conceiving of them as minimal sets of actions, events or states of nature which together initiate a process resulting in the outcome.

Thus each sufficient cause is hypothesized to consist of a set of "component causes". Whenever all components of a particular sufficient cause are present, the outcome occurs;

within every sufficient cause, each component would be necessary for that sufficient cause to lead to the outcome. Models of this kind have a long history: a simple version is considered by Cayley [3]; it also corresponds to the INUS model introduced by Mackie [11] in the philosophical literature. Much recent work has sought to relate the model to other causal modelling frameworks [7, 8, 29, 31, 32].

In traditional sufficient-component cause [SCC] models the outcome and all the component causes are events, or equivalently, binary random variables. An SCC model with k component causes, implies a set of 2^k potential outcomes. Conversely, in section 2 we show that for any given list of potential outcomes there is at least one SCC model which represents this set. However, in general there may be many such SCC models.

One question concerns whether, given a set of potential outcomes implied by some (unknown) SCC model, one may infer that two component causes are present within some sufficient cause in the unknown SCC model. In general, it is possible that two SCC models both imply the same set of potential outcomes yet although A and B occur together in some sufficient component cause in the first model, A and B are not present together in any sufficient component cause in the second. In [32] two sufficient component causes are said to form a ‘sufficient cause interaction’ (or to be ‘irreducible’) if they are both present within at least one sufficient cause in *every* SCC model for a given set of potential outcomes. Of course, in general, the distribution of potential outcomes for a given population is also unknown, though it is constrained (marginally) by the observed data from a randomized experiment. In [32] empirical conditions are given which are sufficient to ensure that for any set of potential outcomes compatible with experimental data, all compatible SCC models will contain a sufficient cause involving A and B . These results are an improvement upon earlier empirical tests for the existence of a two-way interaction in an SCC model [22], which required the assumption of monotonicity; see also [1, 9, 10, 15, 30]. The new results are able to establish the existence of an interaction in situations where monotonicity does not hold. In this paper we develop empirical conditions that are sufficient for the existence of a sufficient cause containing a given subset of an arbitrary number of variables both with and without monotonicity assumptions.

As illustrative motivation for the theoretical development we will consider data presented in a study by [26], summarized in Table 1, from a case-control study of bladder cancer examining possible three-way interaction between smoking (1 = present), and genetic variants on NAT2 (0 = R , 1 = S genotype) and NAT1 (1 for the *10 allele) for Caucasian individuals. We return to this example at the end of this paper to examine the evidence for a sufficient cause containing all three of smoking, the S genotype on NAT2 and the *10 allele on NAT1.

The remainder of this paper is organized as follows: Section 2 presents the sufficient-component cause framework as formalized by VanderWeele and Robins [32]. Section 3 describes general n -way irreducible interactions (aka ‘sufficient cause interactions’) and characterizes these in terms of potential outcomes. Section 4 derives empirical conditions for the existence of irreducible interactions both with and without monotonicity assumptions. Section 5 describes ‘singular’ interactions which arise in genetic contexts, provides a

characterization, derives empirical conditions that are sufficient for their existence, and relates this notion to Pearl’s probability of causation. Section 6 discusses the relation between singular and sufficient cause interactions and linear statistical models. Section 7 provides some comments regarding stronger interpretations of sufficient cause models, and returns to the data presented in Table 1. Finally section 8 offers some possible extensions to the present work.

2. Notation and Basic Concepts

We will use the following notation: An *event* is a binary random variable taking values in $\{0, 1\}$. We use uppercase roman to indicate events (X), boldface to indicate sets of events (\mathbf{C}), and lowercase to indicate specific values both for single random variables ($X = x$), and, with slight abuse of notation, for sets $\{\mathbf{C} = \mathbf{c}\} \equiv \{\forall i, (\mathbf{C})_i = (\mathbf{c})_i\}$, and $\{\mathbf{a} = \mathbf{b}\} \equiv \{\forall i, (\mathbf{a})_i = (\mathbf{b})_i\}$; $\mathbf{1}$ and $\mathbf{0}$ are vectors of 1’s and 0’s; the cardinality of a set is denoted $|\mathbf{C}|$. We use fraktur (\mathfrak{B}) to denote collections of sets of events.

The complement of some event X is denoted by $\bar{X} \equiv 1 - X$. A *literal* event associated with X , is either X or \bar{X} . For a given set of events \mathbf{C} , $\mathbb{L}(\mathbf{C})$ is the associated set of literal events:

$$\mathbb{L}(\mathbf{C}) \equiv \mathbf{C} \cup \{\bar{X} | X \in \mathbf{C}\}.$$

For a literal $L \in \mathbb{L}(\mathbf{C})$, and an assignment \mathbf{c} to \mathbf{C} , $(L)_{\mathbf{c}}$ denotes the value assigned to L by \mathbf{c} . The *conjunction* of a set of literal events $\mathbf{B} = \{F_1, \dots, F_m\} \subseteq \mathbb{L}(\mathbf{C})$ is defined as:

$$\bigwedge(\mathbf{B}) \equiv \prod_{i=1}^m F_i = \min\{F_1, \dots, F_m\};$$

note that $\bigwedge(\mathbf{B}) = 1$ iff for *all* i , $F_i = 1$. We also define $B_1 \wedge B_2 \equiv \bigwedge\{B_1, B_2\}$. We will use $\mathbb{I}(A)$ to denote the indicator function for event A . There is a simple correspondence between conjunctions of literals and indicator functions: let $\mathbf{B} = \{X_1, \dots, X_s\}$ and $\mathbf{C} = \{Y_1, \dots, Y_t\}$, then

$$\bigwedge(\{X_1, \dots, X_s, \bar{Y}_1, \dots, \bar{Y}_t\}) = 1 \iff \mathbb{I}(\{\mathbf{B}=1, \mathbf{C}=0\}) = 1. \quad (2.1)$$

Similarly, the *set of literals corresponding to an assignment \mathbf{c} to \mathbf{C}* is defined:

$$\mathbf{B}^{[\mathbf{c}]} \equiv \{L | L \in \mathbb{L}(\mathbf{C}), (L)_{\mathbf{c}} = 1\}$$

so that $\bigwedge(\mathbf{B}^{[\mathbf{c}]}) = \mathbb{I}(\mathbf{C}=\mathbf{c})$; note that $|\mathbf{B}^{[\mathbf{c}]}| = |\mathbf{C}|$. The *disjunction* of a set of binary random variables is defined as:

$$\bigvee(\{Z_1, \dots, Z_p\}) \equiv \max\{Z_1, \dots, Z_p\};$$

note that $\vee(\{Z_1, \dots, Z_p\}) = 1$ iff for *some* j , $Z_j = 1$. Similarly we let $B_1 \vee B_2 \equiv \vee\{B_1, B_2\}$. Given a collection of sets of literals $\mathfrak{B} = \{C_1, \dots, C_q\}$ we define:

$$\bigvee \bigwedge (\mathfrak{B}) \equiv \bigvee_i \left(\bigwedge (C_i) \right).$$

We use $\dot{\mathbb{P}}(\mathbb{L}(\mathbf{C}))$ to denote the set of subsets of $\mathbb{L}(\mathbf{C})$ that do not contain both X and \overline{X} for any $X \in \mathbf{C}$, more formally:

$$\dot{\mathbb{P}}(\mathbb{L}(\mathbf{C})) \equiv \left\{ \mathbf{B} \mid \mathbf{B} \subset \mathbb{L}(\mathbf{C}), \text{ for all } X \in \mathbf{C}, \{X, \overline{X}\} \not\subset \mathbf{B} \right\}.$$

Note that if $\mathbf{B} \in \dot{\mathbb{P}}(\mathbb{L}(\mathbf{C}))$, and $|\mathbf{B}| = |\mathbf{C}|$, so that for all $C \in \mathbf{C}$, exactly one of C or \overline{C} is in \mathbf{B} , then an assignment of values \mathbf{b} to \mathbf{B} induces a unique assignment \mathbf{c} to \mathbf{C} and vice versa.

2.1. Potential Outcomes Models

Consider a potential outcome model [14, 23, 24] with s binary factors, X_1, \dots, X_s , which represent hypothetical interventions or causes and let D denote some binary outcome of interest. We use Ω to denote the sample space of individuals in the population and use ω for a particular sample point. Let $D_{x_1 \dots x_s}(\omega)$ denote the counterfactual value of D for individual ω if the cause X_j were set to the value x_j for $j = 1, \dots, s$. The potential outcomes framework we employ makes two assumptions: first, that for a given individual these counterfactual variables are deterministic; second, in asserting that the counterfactual $D_{x_1 \dots x_s}(\omega)$ is well-defined, it is implicitly assumed that the value that D would take on for individual ω is determined solely by the values that X_1, \dots, X_s are assigned for this individual, and not the assignments made to these variables for other individuals ω' . This latter assumption is often called ‘no interference’ [6], or the Stable Unit Treatment Value Assumption (SUTVA) [25]. An example of a situation where this assumption might fail is a vaccine trial where there is ‘herd’ immunity.

We will use $D_{x_1 \dots x_s}(\omega)$, $D_{X_1=x_1, \dots, X_s=x_s}(\omega)$, $D_{\mathbf{c}}$ and $D_{C=c}(\omega)$, with $\mathbf{C} = \{X_1, \dots, X_s\}$ interchangeably. In this setting there will be 2^s potential outcomes for each individual ω in the population, one potential outcome for each possible value of (X_1, \dots, X_s) ; we use $\mathcal{D}(\mathbf{C}; \omega)$ to denote the set of all such potential outcomes for an individual, and $\mathcal{D}(\mathbf{C}; \Omega)$ for the population. Note that if $G = g(\mathbf{C})$ is some deterministic function of \mathbf{C} then $G_{C=c}(\omega) = g(\mathbf{c})$, and hence is constant; thus our usage is consistent with the definition of $(L)_{\mathbf{c}}$ in the previous section.

The actual observed value of D for individual ω will be denoted by $D(\omega)$ and similarly the actual value of X_1, \dots, X_s by $X_1(\omega), \dots, X_s(\omega)$. Actual and counterfactual outcomes are linked by the *Consistency Axiom* which requires that

$$D_{X_1=X_1(\omega), \dots, X_s=X_s(\omega)}(\omega) = D(\omega), \quad (2.2)$$

i.e. that the value of D which would have been observed if X_1, \dots, X_s had been set to the values they actually took is equal to the value of D which was in fact observed [?]. It follows

from this axiom that $D_{X_1(\omega), \dots, X_5(\omega)}(\omega) = D$ is observed, but it is the only potential outcome for individual ω that is observed.

EXAMPLE 1. Consider a binary outcome D with three binary causes of interest, X_1 , X_2 and X_3 . Suppose that the population consists of two individuals. The potential outcomes (LHS) and actual outcomes (RHS) are shown in Table 2.

We use the notation $A \perp\!\!\!\perp B \mid C$ to indicate that A is independent of B conditional on C in the population distribution.

2.2. Definitions for sufficient cause models

The following definitions generalize those in [32] to sub-populations, $\emptyset \neq \Omega^* \subseteq \Omega$:

DEFINITION 2.1 (Sufficient Cause). A subset \mathbf{B} of the putative (binary) causes $\mathbb{L}(\mathbf{C})$ for D forms a sufficient cause for D (relative to \mathbf{C}) in sub-population Ω^* if for all $\mathbf{c} \in \{0, 1\}^{|\mathbf{C}|}$ such that $(\Lambda(\mathbf{B}))_{\mathbf{c}} = 1$, $D_{\mathbf{c}}(\omega) = 1$ for all $\omega \in \Omega^* \subseteq \Omega$. (We assume that there exists a \mathbf{c}^* such that $(\Lambda(\mathbf{B}))_{\mathbf{c}^*} = 1$.)

Observe that if \mathbf{B} is a sufficient cause for D then any intervention setting the variables \mathbf{C} to \mathbf{c} with $(\Lambda(\mathbf{B}))_{\mathbf{c}} = 1$ will ensure that $D_{\mathbf{c}}(\omega) = 1$ for all $\omega \in \Omega^*$. We restrict the definition to non-empty sets Ω^* , to preclude every set \mathbf{B} being a sufficient cause in an empty sub-population. Likewise we require that there exists some \mathbf{c}^* such that $(\Lambda(\mathbf{B}))_{\mathbf{c}^*} = 1$ in order to avoid logically inconsistent conjunctions, e.g. $X_1 \wedge \overline{X_1}$, being classified (vacuously) as a sufficient cause. As a direct consequence, for any binary random variable X , at most one of X and \overline{X} appear in any sufficient cause \mathbf{B} .

PROPOSITION 2.2. In Ω^* if \mathbf{B} is a sufficient cause for D relative to \mathbf{C} then \mathbf{B} is sufficient for D in any set \mathbf{C}^* with $\mathbf{B} \subseteq \mathbf{C}^* \subseteq \mathbf{C}$.

\mathbf{B} may be sufficient for D relative to \mathbf{C} in Ω^* , but not relative to $\mathbf{C}' \supset \mathbf{C}$.

PROPOSITION 2.3. If \mathbf{B} is a sufficient cause for D relative to \mathbf{C} in Ω^* then \mathbf{B} is sufficient for D relative to \mathbf{C} in any subset $\emptyset \neq \Omega^{**} \subseteq \Omega^*$.

\mathbf{B} may be sufficient for D relative to \mathbf{C} in Ω^* , but not in $\Omega' \supset \Omega^*$.

DEFINITION 2.4 (Minimal Sufficient Cause). A set $\mathbf{B} \subset \mathbb{L}(\mathbf{C})$ forms a minimal sufficient cause for D (relative to \mathbf{C}) in sub-population Ω^* if \mathbf{B} constitutes a sufficient cause for D in Ω^* but no proper subset $\mathbf{B}^* \subset \mathbf{B}$ also forms a sufficient cause for D in Ω^* .

Note that (in some Ω^*) \mathbf{B} may be a minimal sufficient cause for D relative to \mathbf{C} , but not relative to $\mathbf{C}^* \subset \mathbf{C}$, so the analog of Proposition 2.2 does not hold. For individual 2 in Table 2 $\{X_1, X_3\}$ is a minimal sufficient cause relative to $\{X_1, X_2, X_3\}$. However, if we suppose that for $\omega = 2$, X_2 is not caused by X_1 and X_3 , so for all x_1, x_3 , $X_{2X_1=x_1, X_3=x_3}(\omega = 2) = X_2(\omega = 2)$, then $\{X_1, X_3\}$ is not a minimal sufficient cause relative to $\{X_1, X_3\}$:

$$\begin{aligned} D_{X_1=0, X_3=1}(\omega=2) &= D_{X_1=0, X_2=0, X_3=1}(\omega=2) = 1, \\ D_{X_1=1, X_3=1}(\omega=2) &= D_{X_1=1, X_2=0, X_3=1}(\omega=2) = 1, \end{aligned}$$

(since $X_2(\omega=2) = 0$) hence X_3 is a sufficient cause of D relative to $\{X_1, X_3\}$, hence $\{X_1, X_3\}$ is not minimal relative to $\{X_1, X_3\}$ for $\omega = 2$.

Similarly, if \mathbf{B} is a minimal sufficient cause for D relative to \mathbf{C} in Ω^* , it does not follow that \mathbf{B} is a *minimal* sufficient cause for D relative to \mathbf{C} in subsets $\Omega^{**} \subseteq \Omega^*$, so the analog to Proposition 2.3 does not hold. In particular, it may be the case that for all $\omega \in \Omega^*$, \mathbf{B} is not a minimal sufficient cause for D in $\{\omega\}$.

In the language of digital circuit theory [12], sufficient causes are termed ‘implicants’, and minimal sufficient causes are ‘prime implicants’.

DEFINITION 2.5 (Determinative Set of Sufficient Causes). A set of sufficient causes for D , $\mathfrak{B} = \{\mathbf{B}_1, \dots, \mathbf{B}_n\} \subseteq \mathbb{P}(\mathbb{L}(\mathbf{C}))$, is said to be determinative for D (relative to \mathbf{C}) in sub-population Ω^* if for all $\omega \in \Omega^*$ and for all \mathbf{c} , $D_{\mathbf{c}}(\omega) = 1$ if and only if $(\bigvee \bigwedge (\mathfrak{B}))_{\mathbf{c}} = 1$.

We will refer to a determinative set of sufficient causes for D as a *sufficient cause model*. Observe that in any sub-population Ω^* for which there exists a determinative set of sufficient causes, the vectors of potential outcomes for D are identical, so $\mathcal{D}(\mathbf{C}, \omega) = \mathcal{D}(\mathbf{C}, \omega')$ for all $\omega, \omega' \in \Omega^*$.

DEFINITION 2.6 (Non-redundant Set of Sufficient Causes). A determinative set of sufficient causes \mathfrak{B} , for D , is said to be non-redundant (in Ω^* , relative to \mathbf{C}) if there is no proper subset $\mathfrak{B}^* \subset \mathfrak{B}$ that is also determinative for D .

Note that sufficient causes are conjunctions, while sets of sufficient causes form disjunctions of conjunctions; minimality refers to the components in a particular conjunction, that each component is required for the conjunction to be sufficient for D ; non-redundancy implies that each conjunction is required for the disjunction of the set of conjunctions to be determinative. If for some set of sufficient causes $\mathfrak{B} \subseteq \mathbb{P}(\mathbb{L}(\mathbf{C}))$, for all $X \in \mathbf{C}$, and all $\mathbf{B} \in \mathfrak{B}$, either $X \in \mathbf{B}$ or $\bar{X} \in \mathbf{B}$ then \mathfrak{B} is a non-redundant set of sufficient causes.

EXAMPLE 1 (Revisited). The set $\mathfrak{B}_1 = \{\{X_1, X_2\}, \{X_2, \bar{X}_3\}, \{\bar{X}_2, X_3\}\}$ forms a determinative set of sufficient causes for the individual $\omega = 2$, since:

$$D_{\mathbf{c}}(\omega=2) = ((X_1 \wedge X_2) \vee (X_2 \wedge \bar{X}_3) \vee (\bar{X}_2 \wedge X_3))_{\mathbf{c}} \quad (2.3)$$

as does $\mathfrak{B}_2 = \{\{X_1, X_3\}, \{X_2, \bar{X}_3\}, \{\bar{X}_2, X_3\}\}$:

$$D_{\mathbf{c}}(\omega=2) = ((X_1 \wedge X_3) \vee (X_2 \wedge \bar{X}_3) \vee (\bar{X}_2 \wedge X_3))_{\mathbf{c}}. \quad (2.4)$$

As this example shows, determinative sets of sufficient causes are not, in general, unique.

2.3. Sufficient cause representations for a population

As noted, if \mathbf{B} is a sufficient cause for D in Ω^* , then all the units in Ω^* will have $D = 1$ for any assignment \mathbf{c} to \mathbf{C} , such that $(\Lambda(\mathbf{B}))_{\mathbf{c}} = 1$. In most realistic settings it is unlikely that any set \mathbf{B} will be sufficient to ensure $D = 1$ in an entire population. Consequently different sets of sufficient causes will be required within different sub-populations. A sufficient cause representation is a set of sub-populations each with its own determinative sufficient cause representation:

DEFINITION 2.7. A sufficient cause representation $(\mathbf{A}, \mathfrak{B})$ for $\mathcal{D}(\mathbf{C}; \Omega)$ is an ordered set $\mathbf{A} = \langle A_1, \dots, A_p \rangle$ of binary random variables, with $(A_i)_{\mathbf{c}} = A_i$ for all i, \mathbf{c} , and a set $\mathfrak{B} = \langle \mathbf{B}_1, \dots, \mathbf{B}_p \rangle$, with $\mathbf{B}_i \in \mathbb{P}(\mathbb{L}(\mathbf{C}))$, such that for all ω, \mathbf{c} , $D_{\mathbf{c}}(\omega) = 1 \Leftrightarrow$ for some j , $A_j(\omega) = 1$ and $(\Lambda(\mathbf{B}_j))_{\mathbf{c}} = 1$.

Note that the binary random variables A_i and the sets \mathbf{B}_i are naturally paired via the orderings of \mathbf{A} and \mathfrak{B} ; we will refer to a pair (A_i, \mathbf{B}_i) as *occurring* in the representation. The requirement that $(A_i)_{\mathbf{c}} = A_i$ for all i, \mathbf{c} implies that $\mathbf{A} \cap \mathbf{C} = \emptyset$, and further that the A_i are unaffected by interventions on the X_i ; this is in keeping with the interpretation of the A_i as defining pre-existing sub-populations with particular sets of potential outcomes for D .

PROPOSITION 2.8. If $(\mathbf{A}, \mathfrak{B})$ is a sufficient cause representation for $\mathcal{D}(\mathbf{C}; \Omega)$ then \mathbf{B}_i is a sufficient cause of D in the sub-population in which $A_i(\omega) = 1$.

PROPOSITION 2.9. If $(\mathbf{A}, \mathfrak{B})$ is a sufficient cause representation for $\mathcal{D}(\mathbf{C}; \Omega)$, then for all $\mathbf{A}^* \subseteq \mathbf{A}$, if

$$\emptyset \neq \Omega_{\mathbf{A} \setminus \mathbf{A}^*}^{\mathbf{A}^*} \equiv \{\omega \mid \text{for all } A_i \in \mathbf{A}, \quad A_i(\omega) = 1 \text{ iff } A_i \in \mathbf{A}^*\},$$

then

$$\mathfrak{B}^{\mathbf{A}^*} \equiv \{\mathbf{B}_i \mid \mathbf{B}_i \in \mathfrak{B}; A_i \in \mathbf{A}^*\}$$

forms a determinative set of sufficient causes (relative to \mathbf{C}) for $\Omega_{\mathbf{A} \setminus \mathbf{A}^*}^{\mathbf{A}^*}$.

Note that $\Omega_{\mathbf{A} \setminus \mathbf{A}^*}^{\mathbf{A}^*}$ consists of the sub-population in which $A_i(\omega) = 1$ for all $A_i \in \mathbf{A}^*$ and $A_j(\omega) = 0$ for all $A_j \in \mathbf{A} \setminus \mathbf{A}^*$.

Proof: Suppose for some $\omega \in \Omega_{\mathbf{A} \setminus \mathbf{A}^*}^{\mathbf{A}^*}$, $\mathbf{B}_j \in \mathfrak{B}^{\mathbf{A}^*}$, and \mathbf{c} we have $(\Lambda(\mathbf{B}_j))_{\mathbf{c}} = 1$. Since $\omega \in \Omega_{\mathbf{A} \setminus \mathbf{A}^*}^{\mathbf{A}^*}$, $A_j(\omega) = 1$. It then follows from the definition of a sufficient cause representation that $D_{\mathbf{c}}(\omega) = 1$. Conversely, suppose $D_{\mathbf{c}}(\omega) = 1$. As $(\mathbf{A}, \mathfrak{B})$ is a sufficient cause

representation, for some j , $A_j(\omega) = 1$ and $(\Lambda(\mathbf{B}_j))_{\mathbf{c}} = 1$. Since, by hypothesis, $\omega \in \Omega_{\mathbf{A} \setminus \mathbf{A}^*}^{\mathbf{A}^*}$, it follows that $A_j \in \mathbf{A}^*$, hence $\mathbf{B}_j \in \mathfrak{B}^{\mathbf{A}^*}$.

THEOREM 2.10. *For any $\mathcal{D}(\mathbf{C}; \Omega)$ there exists a sufficient cause representation $(\mathbf{A}, \mathfrak{B})$.*

Proof: Let $p = 2^{|\mathbf{C}|}$, and define $\mathfrak{B} \equiv \{ \mathbf{B} | \mathbf{B} \subseteq \mathbb{P}(\mathbb{L}(\mathbf{C})), |\mathbf{B}| = |\mathbf{C}| \} \equiv \langle \mathbf{B}_1, \dots, \mathbf{B}_p \rangle$, ordered arbitrarily. Further define $A_i(\omega) \equiv D_{\mathbf{B}_i=1}(\omega)$. Given an arbitrary \mathbf{c} , for some j , $\mathbf{B}^{[\mathbf{c}]} = \mathbf{B}_j$, by construction of \mathfrak{B} . We then have

$$D_{\mathbf{c}}(\omega) = 1 \iff D_{\mathbf{B}_j=1}(\omega) = 1 \iff A_j(\omega) = 1 \text{ and } \left(\bigwedge (\mathbf{B}_j) \right)_{\mathbf{c}} = 1,$$

as required. The last step follows since by definition $\mathbf{B}^{[\mathbf{c}]} = \mathbf{1}$ iff $\mathbf{C} = \mathbf{c}$.

[32] proves this for the case of $|\mathbf{C}| = 2$; see also [7] and [29] for discussion of the case $|\mathbf{C}| = 1$.

EXAMPLE 1 (Revisited). The construction given in the proof of Theorem 2.10 would yield the following sets of sufficient causes to represent $\mathcal{D}(\mathbf{C}; \Omega)$ shown in Table 2:

$$\begin{aligned} \mathfrak{B} &= \langle \mathbf{B}_1, \dots, \mathbf{B}_8 \rangle \\ &= \left\langle \{X_1, X_2, X_3\}, \{X_1, X_2, \bar{X}_3\}, \{X_1, \bar{X}_2, X_3\}, \{X_1, \bar{X}_2, \bar{X}_3\}, \right. \\ &\quad \left. \{\bar{X}_1, X_2, X_3\}, \{\bar{X}_1, X_2, \bar{X}_3\}, \{\bar{X}_1, \bar{X}_2, X_3\}, \{\bar{X}_1, \bar{X}_2, \bar{X}_3\} \right\rangle, \end{aligned} \tag{2.5}$$

with $A_1 = A_4 = A_5 = 0, A_2 = A_3 = A_6 = A_7 = 1, A_8 = \mathbb{I}(\{\omega = 2\})$.

3. Irreducible conjunctions

We saw in Example 1 above with $\omega = 2$ that an individual’s potential outcomes may be such that there are two determinative sets of common causes \mathfrak{B} and \mathfrak{B}' and $\{X_1, X_2\}$ is in \mathfrak{B} , but not in \mathfrak{B}' . However, certain conjunctions are such that in every representation either the conjunction is present or it is contained in some larger conjunction; such conjunctions are said to be ‘irreducible’:

DEFINITION 3.1. $\mathbf{B} \in \mathbb{P}(\mathbb{L}(\mathbf{C}))$ is said to be irreducible for $\mathcal{D}(\mathbf{C}, \Omega)$ if in every representation $(\mathbf{A}, \mathfrak{B})$ for $\mathcal{D}(\mathbf{C}, \Omega)$ there exists $\mathbf{B}_i \in \mathfrak{B}$, with $\mathbf{B} \subseteq \mathbf{B}_i$.

[32] also refer to irreducibility of \mathbf{B} for $\mathcal{D}(\mathbf{C}, \Omega)$ as a ‘sufficient cause interaction’ between the components of \mathbf{B} . (Note, however, that if \mathbf{B} is irreducible, this does not in general imply that \mathbf{B} is either a minimal sufficient cause, or even a sufficient cause, only that there is a sufficient cause that contains \mathbf{B} .) It can be shown (via Theorem 3.2 below) that $\{X_2, \bar{X}_3\}$ and $\{\bar{X}_2, X_3\}$ are irreducible for $\mathcal{D}(\mathbf{C}; \Omega)$ in Table 2. In §7 we provide an interpretation of irreducibility in terms of the existence of a mechanism involving the variables in \mathbf{B} . We now characterize irreducibility:

THEOREM 3.2. Let $\mathbf{C} = \mathbf{C}_1 \dot{\cup} \mathbf{C}_2$, $\mathbf{B} \in \dot{\mathbb{P}}(\mathbb{L}(\mathbf{C}_1))$, $|\mathbf{B}| = |\mathbf{C}_1|$, then \mathbf{B} is irreducible for $\mathcal{D}(\mathbf{C}, \Omega)$ iff there exists $\omega^* \in \Omega$ and values \mathbf{c}_2^* for \mathbf{C}_2 such that: (i) $D_{\mathbf{B}=1, \mathbf{C}_2=\mathbf{c}_2^*}(\omega^*) = 1$; (ii) for all $L \in \mathbf{B}$, $D_{\mathbf{B} \setminus \{L\}=1, L=0, \mathbf{C}_2=\mathbf{c}_2^*}(\omega^*) = 0$.

Here $\mathbf{C}_1 \dot{\cup} \mathbf{C}_2$ indicates the disjoint union of \mathbf{C}_1 and \mathbf{C}_2 . Thus \mathbf{B} is irreducible if and only if there exists an individual in Ω , who would have response $D = 1$ if every literal in \mathbf{B} is set to 1, but $D = 0$ whenever one literal is set to 0 and the rest to 1 (in some context $\mathbf{C}_2 = \mathbf{c}_2^*$). Note that conditions (i) and (ii) are equivalent to:

$$D_{\mathbf{B}=1, \mathbf{C}_2=\mathbf{c}_2^*}(\omega^*) - \sum_{L \in \mathbf{B}} D_{\mathbf{B} \setminus \{L\}=1, L=0, \mathbf{C}_2=\mathbf{c}_2^*}(\omega^*) > 0. \quad (3.1)$$

Proof: (\Rightarrow) We adapt the proof of Theorem 2.10, to show that if for all $\omega \in \Omega$ and assignments \mathbf{c}_2^* to \mathbf{C}_2 , at least one of (i) or (ii) does not hold, then there exists a representation $(\mathbf{A}, \mathfrak{B})$ for $\mathcal{D}(\mathbf{C}, \Omega)$ such that for all $\mathbf{B}_i \in \mathfrak{B}$, $\mathbf{B} \subseteq \mathbf{B}_i$.

Define:

$$\begin{aligned} \mathfrak{B}^\dagger &\equiv \langle \mathbf{B}_i^\dagger \rangle \equiv \{ \mathbf{B}^* | \mathbf{B}^* \in \dot{\mathbb{P}}(\mathbb{L}(\mathbf{C})), |\mathbf{B}^*| = |\mathbf{C}|, \mathbf{B} \not\subseteq \mathbf{B}^* \}, \\ \mathfrak{B}^\ddagger &\equiv \langle \mathbf{B}_i^\ddagger \rangle \equiv \{ \mathbf{B}^* | \mathbf{B}^* \in \dot{\mathbb{P}}(\mathbb{L}(\mathbf{C})), |\mathbf{B}^*| = |\mathbf{C}| - 1, \mathbf{B} \setminus \mathbf{B}^* = \{L\}, L \in \mathbf{B} \}, \end{aligned}$$

under arbitrary orderings. Thus \mathfrak{B}^\dagger is the set of subsets of exactly $|\mathbf{C}|$ literals that do not include \mathbf{B} as a subset, while \mathfrak{B}^\ddagger contains those subsets of size $|\mathbf{C}| - 1$ that contain all but one literals in \mathbf{B} .

For $\mathbf{B}_i^\dagger \in \mathfrak{B}^\dagger$ define the corresponding $\mathbf{A}_i^\dagger(\omega) \equiv D_{\mathbf{B}_i^\dagger=1}(\omega)$,

For $\mathbf{B}_i^\ddagger \in \mathfrak{B}^\ddagger$ define $\mathbf{A}_i^\ddagger(\omega) \equiv D_{\mathbf{B}_i^\ddagger=1, L_i=0}(\omega) D_{\mathbf{B}_i^\ddagger=1, L_i=1}(\omega)$, where $\{L_i\} \equiv \mathbf{B} \setminus \mathbf{B}_i^\ddagger$.

The representation is given by $(\mathbf{A}, \mathfrak{B}) \equiv (\mathbf{A}^\dagger \cup \mathbf{A}^\ddagger, \mathfrak{B}^\dagger \cup \mathfrak{B}^\ddagger)$, where

$\mathbf{A}^\dagger \equiv \langle \mathbf{A}_i^\dagger \rangle$, $\mathbf{A}^\ddagger \equiv \langle \mathbf{A}_i^\ddagger \rangle$. To see this, first note that if for some ω and \mathbf{c} , there is a pair (A_j, \mathbf{B}_j) in $(\mathbf{A}, \mathfrak{B})$ such that $A_j(\omega) = 1$ and $(\wedge \mathbf{B}_j)_\mathbf{c} = 1$ then by construction of \mathbf{A}^\dagger and \mathbf{A}^\ddagger it follows that $D_\mathbf{c}(\omega) = 1$. For the converse, suppose that for some \mathbf{c} and ω , $D_\mathbf{c}(\omega) = 1$. There are two cases to consider:

$(\wedge \mathbf{B})_\mathbf{c} = 0$. In this case $\mathbf{B} \not\subseteq \mathbf{B}^{[\mathbf{c}]}$, so for some j , $\mathbf{B}_j^\dagger = \mathbf{B}^{[\mathbf{c}]}$, hence

$$A_j^\dagger(\omega) \equiv D_{\mathbf{B}_j^\dagger=1}(\omega) = D_\mathbf{c}(\omega) = 1, \text{ as required.}$$

$(\wedge \mathbf{B})_\mathbf{c} = 1$. Let \mathbf{c} be partitioned as $(\mathbf{c}_1, \mathbf{c}_2)$. Since (i) holds with $\mathbf{c}_2^* = \mathbf{c}_2$, (ii) does not. Thus for some $L \in \mathbf{B}$, $D_{\mathbf{B} \setminus \{L\}=1, L=0, \mathbf{C}_2=\mathbf{c}_2}(\omega) = 1$. By construction of \mathfrak{B}^\ddagger , for some j , $\mathbf{B}_j^\ddagger = \mathbf{B}^{[\mathbf{c}]} \setminus \{L\}$, so $(\wedge \mathbf{B}_j^\ddagger)_\mathbf{c} = 1$. Since $1 = D_\mathbf{c}(\omega) = D_{\mathbf{B} \setminus \{L\}=1, L=0, \mathbf{C}_2=\mathbf{c}_2}(\omega)$, we have $A_j^\ddagger(\omega) = 1$, as required.

(\Leftarrow) Suppose for a contradiction, that for some ω^* and \mathbf{c}_2^* , (i) and (ii) hold, but \mathbf{B} is not irreducible. Then there exists a representation $(\mathbf{A}, \mathfrak{B})$ such that for all $\mathbf{B}_i \in \mathfrak{B}$, $\mathbf{B} \not\subseteq \mathbf{B}_i$. By (i), $D_{\mathbf{B}=1, \mathbf{c}_2^*}(\omega^*) = 1$. Thus for some pair (A_j, \mathbf{B}_j) , $A_j(\omega^*) = 1$, and $\mathbf{B}_j \subseteq \mathbf{B} \cup \mathbf{B}[\mathbf{c}_2^*]$. Since $\mathbf{B} \not\subseteq \mathbf{B}_j$, there exists some $L \in \mathbf{B} \setminus \mathbf{B}_j$, but then since $A_j(\omega^*) = 1$ and

$\left(\bigwedge_{\mathbf{B} \setminus \{L\}=1, L=0, \mathbf{C}_2=\mathbf{c}_2^*} (\mathbf{B}_j)\right) = 1$, we have $D_{\mathbf{B} \setminus \{L\}=1, L=0, \mathbf{C}_2=\mathbf{c}_2^*}(\omega^*) = 1$, which is a contradiction.

COROLLARY 3.3. *If \mathbf{B} is irreducible for $\mathcal{D}(\mathbf{C}, \Omega)$ then for any $\Omega^* \supset \Omega$, \mathbf{B} is irreducible for $\mathcal{D}(\mathbf{C}, \Omega^*)$.*

Proof: By Theorem 3.2, since if Ω satisfies (i) and (ii) then so does Ω^* .

3.1. \mathbf{B} irreducible for $\mathcal{D}(\mathbf{C}, \Omega)$ with $|\mathbf{B}| = |\mathbf{C}|$

In the special case where $|\mathbf{B}| = |\mathbf{C}|$, the concepts of minimal sufficient cause for some ω^* and irreducibility coincide.

PROPOSITION 3.4. *If $\mathbf{B} \in \dot{\mathbb{P}}(\mathbb{L}(\mathbf{C}))$ and $|\mathbf{B}| = |\mathbf{C}|$ then \mathbf{B} is a minimal sufficient cause for some $\omega^* \in \Omega$ relative to \mathbf{C} i \mathbf{B} is irreducible for $\mathcal{D}(\mathbf{C}, \omega)$.*

Proof: If $|\mathbf{B}| = |\mathbf{C}|$ then condition (i) in Theorem 3.2 (taking $\mathbf{C}_2 = \emptyset$) holds iff \mathbf{B} is a sufficient cause for D for ω^* , and similarly condition (ii) holds iff \mathbf{B} is a minimal sufficient cause for D (for ω^*).

Thus we have the following:

COROLLARY 3.5. *If $\mathbf{B} \in \dot{\mathbb{P}}(\mathbb{L}(\mathbf{C}))$, $|\mathbf{B}| = |\mathbf{C}|$ and \mathbf{B} is a minimal sufficient cause for D for some $\omega^* \in \Omega$ then $\mathbf{B} \in \mathfrak{B}$ for every representation $(\mathbf{A}, \mathfrak{B})$ for $\mathcal{D}(\mathbf{C}, \Omega)$.*

Proof: Immediate from Proposition 3.4.

3.2. \mathbf{B} irreducible for $\mathcal{D}(\mathbf{C}, \Omega)$ with $|\mathbf{B}| > |\mathbf{C}|$

When $|\mathbf{B}| < |\mathbf{C}|$ the conditions for irreducibility and for being a minimal sufficient cause are logically distinct. Condition (i) in Theorem 3.2 requires $D_{\mathbf{B}=1, \mathbf{c}_2^*}(\omega^*) = 1$ for one assignment \mathbf{c}_2^* (and some ω^*), while if \mathbf{B} is a sufficient cause (for ω^*) then this condition is required to hold for all assignments \mathbf{c}_2^* ; in contrast condition (ii) in Theorem 3.2 requires that there exists a single \mathbf{c}_2^* (and some ω^*) such that for all $L \in \mathbf{B}$, $D_{\mathbf{B} \setminus \{L\}=1, L=0, \mathbf{C}_2=\mathbf{c}_2^*}(\omega^*) = 0$, while for \mathbf{B} to be a *minimal* sufficient cause for ω^* merely requires that for all $L \in \mathbf{B}$, there exists an assignment \mathbf{c}_2^L such that $D_{\mathbf{B} \setminus \{L\}=1, L=0, \mathbf{C}_2=\mathbf{c}_2^L}(\omega^*) = 0$.

EXAMPLE 1 (Revisited). Let $\mathbf{C} = \{X_1, X_2, X_3\}$, $\omega = \{2\}$. Relative to \mathbf{C} , $\{X_1, X_2\}$ is a minimal sufficient cause for $\omega = 2$ since $D_{111}(2) = D_{110}(2) = 1$, and $D_{011}(2) = D_{100}(2) = 0$. However $\{X_1, X_2\}$ is not irreducible for $\mathcal{D}(\mathbf{C}, \Omega)$ because we have $D_{101}(2) = D_{010}(2) = 1$, hence condition (ii) in Theorem 3.2 is not satisfied for either $X_3 = 0$, or $X_3 = 1$. Conversely $\{X_1\}$ is

irreducible for $\mathbf{C} = \{2\}$ since $D_{111}(2) = 1$, while $D_{011}(2) = 0$, but $\{X_1\}$ is not a sufficient cause because $D_{100}(2) = 0$.

Though irreducibility of \mathbf{B} for $\mathcal{D}(\mathbf{C}, \Omega)$ neither implies, nor is implied by \mathbf{B} being a minimal sufficient cause for some $\omega \in \Omega$, it does imply that every sufficient cause representation for $\mathcal{D}(\mathbf{C}, \Omega)$ contains at least one conjunction \mathbf{B}_j of which \mathbf{B} is a (possibly proper) subset. However, *prima facie* this still leaves open the possibility that, for example, every representation either includes $\mathbf{B} \cup \{L\}$ or $\mathbf{B} \cup \{\bar{L}\}$, for some L , but no representation includes both. However, this cannot occur:

COROLLARY 3.6. *Let $\mathbf{C} = \mathbf{C}_1 \cup \mathbf{C}_2$, $\mathbf{B} \in \dot{\mathbb{P}}(\mathbb{L}(\mathbf{C}_1))$ if \mathbf{B} is irreducible for $\mathcal{D}(\mathbf{C}, \Omega)$ then there exists a set $\mathbf{B}^* \in \dot{\mathbb{P}}(\mathbb{L}(\mathbf{C}))$, with $|\mathbf{B}^*| = |\mathbf{C}|$ such that in every representation $(\mathbf{A}, \mathfrak{B})$ for $\mathcal{D}(\mathbf{C}, \Omega)$ there exists $\mathbf{B}_j \in \mathfrak{B}$, with $\mathbf{B} \subseteq \mathbf{B}_j \subseteq \mathbf{B}^*$.*

Thus irreducibility of \mathbf{B} further implies that there is a set \mathbf{B}^* of size $|\mathbf{C}|$ such that in every representation there is at least one conjunct containing \mathbf{B} that is itself contained in \mathbf{B}^* . However, it should be noted that, in general, there may be more than one conjunct \mathbf{B}_j with $\mathbf{B} \subseteq \mathbf{B}_j \subseteq \mathbf{B}^*$.

Proof: Immediate from Theorem 3.2, taking $\mathbf{B}^* = \mathbf{B} \cup \mathbf{B}^{[c_2^*]}$.

Finally, we note that a conjunction that is both irreducible and a minimal sufficient cause corresponds to an ‘essential prime implicant’ in digital circuit theory [12]. The Quine-McCluskey algorithm [13, 18, 19] finds the set of essential prime implicants for a given Boolean function, which here corresponds to the potential outcomes $\mathcal{D}(\mathbf{C}, \omega)$ for an individual.

3.3. Enlarging the set of potential causes

As noted in §2.2 a set \mathbf{B} may be a minimal sufficient cause for \mathbf{C} but not a superset \mathbf{C}' . Irreducibility is also not preserved without further conditions. To state these conditions we require the following:

DEFINITION 3.7. *X' is said to be not causally influenced by a set \mathbf{C} if for all $\omega \in \Omega$, the potential outcomes $X'_{\mathbf{C}=\mathbf{c}}(\omega)$ are constant as \mathbf{c} varies.*

We will also assume that if every $X' \in \mathbf{C}'$ is not causally influenced by \mathbf{C} then the *Relativized Consistency Axiom* holds:

$$D_{\mathbf{C}=\mathbf{c}, \mathbf{C}'=\mathbf{c}'(\omega)}(\omega) = D_{\mathbf{C}=\mathbf{c}}(\omega), \quad (3.2)$$

i.e. that if variables in \mathbf{C}' are not causally influenced by the variables in \mathbf{C} then the counterfactual value of D intervening to set \mathbf{C} to \mathbf{c} is the same as the counterfactual value of D intervening to set \mathbf{C} to \mathbf{c} and the variables in \mathbf{C}' to the values they actually took on.

We now have the following Corollary to Theorem 3.2:

COROLLARY 3.8. *Let $\mathbf{C} = \mathbf{C}_1 \cup \mathbf{C}_2$, $\mathbf{B} \in \dot{\mathbb{P}}(\mathbb{L}(\mathbf{C}_1))$, $|\mathbf{B}| = |\mathbf{C}_1|$. If \mathbf{B} is irreducible for $\mathcal{D}(\mathbf{C}, \Omega)$, $\mathbf{C}' \cap \mathbf{C} = \emptyset$ and for all $X' \in \mathbf{C}'$, X' is not causally influenced by \mathbf{C} , then \mathbf{B} is irreducible for $\mathcal{D}(\mathbf{C} \cup \mathbf{C}', \Omega)$.*

Proof: By Theorem 3.2 there exists $\omega^* \in \Omega$ and an assignment \mathbf{c}_2^* to \mathbf{C}_2 such that (i) and (ii) hold. Let $\mathbf{c}' = \mathbf{C}'(\omega^*)$. Since variables in \mathbf{C}' are not causally influenced by \mathbf{C} , for all assignments \mathbf{b} :

$$D_{\mathbf{B}=\mathbf{b}, \mathbf{C}_2=\mathbf{c}_2^*, \mathbf{C}'=\mathbf{c}'}(\omega^*) = D_{\mathbf{B}=\mathbf{b}, \mathbf{C}_2=\mathbf{c}_2^*, \mathbf{C}'=\mathbf{C}'(\omega^*)}(\omega^*) = D_{\mathbf{B}=\mathbf{b}, \mathbf{C}_2=\mathbf{c}_2^*}(\omega^*),$$

the second equality here follows from (3.2). It follows that ω^* and $(\mathbf{c}_2^*, \mathbf{c}')$ obey (i) and (ii) in Theorem 3.2 with respect to $\mathbf{C} \cup \mathbf{C}'$.

The assumption that every variable in \mathbf{C}' is not causally influenced by \mathbf{C} , is required because otherwise we may have $\mathbf{C}'(\omega^*) = (\mathbf{C}')_{\mathbf{B}=\mathbf{b}^*}(\omega^*)$ for some assignment \mathbf{b}^* to \mathbf{B} . For example, let $\mathbf{C} = \{X_1, X_2, X_3\}$, and suppose that

$$\begin{aligned} D_{X_1=x_1, X_2=x_2, X_3=x_3}(\omega) &= x_3 \\ (X_3)_{X_1=x_1, X_2=x_2}(\omega) &= x_1 \wedge x_2 \end{aligned}$$

for all $\omega \in \Omega$. In this case $\{X_1, X_2\}$ is irreducible for $\mathcal{D}(\{X_1, X_2\}, \Omega)$, but not for $\mathcal{D}(\{X_1, X_2, X_3\}, \Omega)$. We saw earlier that if \mathbf{B} is a minimal sufficient cause for \mathbf{C} then this does not imply that \mathbf{B} is a minimal sufficient cause with respect to subsets of \mathbf{C} . Here we see that if \mathbf{B} is irreducible with respect for $\mathcal{D}(\mathbf{C}, \Omega)$ then this does not imply irreducibility for supersets $\mathbf{C}^* \supset \mathbf{C}$, unless every variable in $\mathbf{C}^* \setminus \mathbf{C}$ is not causally influenced by a variable in \mathbf{C} .

4. Tests for irreducibility

In this section we derive empirical conditions which imply that a given conjunction \mathbf{B} is irreducible for $\mathcal{D}(\mathbf{C}, \Omega)$. Our first approach is via condition (3.1).

4.1. Adjusting for Measured Confounders

To detect that (3.1) holds requires us to identify the mean of potential outcomes in certain subpopulations. This is only possible if we have no unmeasured confounders [??]:

DEFINITION 4.1. *A set of covariates \mathbf{W} suffice to adjust for confounding of (the effect of) \mathbf{C} on D if:*

$$D_{\mathbf{C}=\mathbf{c}} \perp\!\!\!\perp \mathbf{C} \mid \mathbf{W}=\mathbf{w} \quad (4.1)$$

for all \mathbf{c}, \mathbf{w} .

PROPOSITION 4.2. *If a set \mathbf{W} suffices to adjust for confounding of \mathbf{C} on D and $P(\mathbf{C} = \mathbf{c}, \mathbf{W} = \mathbf{w}) > 0$ then*

$$E[D_{\mathbf{C}=\mathbf{c}} | \mathbf{W}=\mathbf{w}] = E[D | \mathbf{C}=\mathbf{c}, \mathbf{W}=\mathbf{w}].$$

The proof of this is standard and hence omitted.

Note that if \mathbf{W} is sufficient to adjust for confounding of \mathbf{C} on D , then \mathbf{W} is also sufficient to adjust for confounding of \mathbf{B} on D , where $\mathbf{B} \in \dot{\mathbb{P}}(\mathbb{L}(\mathbf{C}))$, $|\mathbf{B}|=|\mathbf{C}|$.

4.2. Tests for irreducibility without monotonicity

THEOREM 4.3. *Let $\mathbf{C}=\mathbf{C}_1 \cup \mathbf{C}_2$, $\mathbf{B} \in \dot{\mathbb{P}}(\mathbb{L}(\mathbf{C}_1))$, $|\mathbf{B}| = |\mathbf{C}_1|$. If \mathbf{W} is sufficient to adjust for confounding of \mathbf{C} on D , and for some \mathbf{c}_2, \mathbf{w} ,*

$$0 < E[D | \mathbf{B}=1, \mathbf{C}_2=\mathbf{c}_2, \mathbf{W}=\mathbf{w}] - \sum_{L \in \mathbf{B}} E[D | \mathbf{B} \setminus \{L\}=1, L=0, \mathbf{C}_2=\mathbf{c}_2, \mathbf{W}=\mathbf{w}] \quad (4.2)$$

then \mathbf{B} is irreducible for $\mathcal{D}(\mathbf{C}, \Omega)$.

Proof: We prove the contrapositive. Suppose that \mathbf{B} is not irreducible for $\mathcal{D}(\mathbf{C}, \Omega)$. Then by Theorem 3.2, for all $\omega \in \Omega$, and all \mathbf{c}_2 ,

$$D_{\mathbf{B}=1, \mathbf{C}_2=\mathbf{c}_2}(\omega) - \sum_{L \in \mathbf{B}} D_{\mathbf{B} \setminus \{L\}=1, L=0, \mathbf{C}_2=\mathbf{c}_2}(\omega) \leq 0.$$

Hence for any \mathbf{w} ,

$$E \left[D_{\mathbf{B}=1, \mathbf{C}_2=\mathbf{c}_2} - \sum_{L \in \mathbf{B}} D_{\mathbf{B} \setminus \{L\}=1, L=0, \mathbf{C}_2=\mathbf{c}_2} \mid \mathbf{W}=\mathbf{w} \right] \leq 0.$$

Applying Proposition 4.2 to each term implies the negation of (4.2).

The condition provided in Theorem 4.3 can be empirically tested with t -test type statistics if \mathbf{W} consists of a small number of categorical or binary variables or using regression or inverse probability of treatment weighting methods [20, 34] otherwise.

It follows from Corollary 3.8 that condition (4.2) further establishes that \mathbf{B} is irreducible for $\mathcal{D}(\mathbf{C} \cup \mathbf{C}', \Omega)$ so long as every variable in \mathbf{C}' is not causally influenced by variables in \mathbf{C} .

It may be shown that condition (4.2) is the sole restriction on the law of $(D, \mathbf{C}, \mathbf{W})$ implied by the negation of irreducibility.

4.3. Graphs

In the next section we develop more powerful tests under monotonicity assumptions. However, to state these conditions we first introduce some concepts from graph theory:

DEFINITION 4.4. A graph \mathfrak{G} defined on a set \mathbf{B} is a collection of pairs of elements in \mathbf{B} , $\mathfrak{G} \equiv \{\mathbf{E} \mid \mathbf{E} = \{B_1, B_2\} \subseteq \mathbf{B}, B_1 \neq B_2\}$.

This is the usual definition of a graph, except that the vertex set here is a set of literals. We will refer to sets in \mathfrak{G} as *edges*, which we will represent pictorially as $B_1 - B_2$.

DEFINITION 4.5. Two elements $L, L^* \in \mathbf{B}$ are said to be connected in \mathfrak{G} if there exists a sequence $L = L_1, \dots, L_p = L^*$ of distinct elements in \mathbf{B} such that $\{L_i, L_{i+1}\} \in \mathfrak{G}$ for $i = 1, \dots, p - 1$.

The sequence of edges joining L and L^* is said to form a *path* in \mathfrak{G} .

Definition 4.6. A graph \mathfrak{G} on \mathbf{B} is said to form a tree if $|\mathfrak{G}| = |\mathbf{B}| - 1$, and any pair of distinct elements in \mathbf{B} are connected in \mathfrak{G} .

In a tree there is a unique path between any two elements.

PROPOSITION 4.7. Let \mathfrak{T} be a tree on \mathbf{B} . For each element $R \in \mathbf{B}$ there is a natural bijection:

$$\phi_R^{\mathfrak{T}}: \mathbf{B} \setminus \{R\} \leftrightarrow \mathfrak{T}$$

given by $\phi_R^{\mathfrak{T}}(L) = \mathbf{E} = \{L', L\}$ where $\mathbf{E} \in \mathfrak{T}$ is the last edge on a path from R to L .

It is not hard to show that for a graph \mathfrak{G} , if the bijections described in Proposition 4.7 exist then \mathfrak{G} is a tree.

THEOREM 4.8 (Cayley [4]). On a set \mathbf{B} there are $|\mathbf{B}|^{|\mathbf{B}|-2}$ different trees.

4.4. Monotonicity

Sometimes it may be known that a certain cause has an effect on an outcome that is either always positive or always negative.

DEFINITION 4.9. B_i has a positive monotonic effect on D relative to a set \mathbf{B} (with $B_i \in \mathbf{B}$) in a population Ω if for all $\omega \in \Omega$ and all values \mathbf{b}_{-i} for the variables in $\mathbf{B} \setminus \{B_i\}$,

$$D_{\mathbf{B}\{B_i\}=\mathbf{b}_{-i}, B_i=1}(\omega) \geq D_{\mathbf{B}\{B_i\}=\mathbf{b}_{-i}, B_i=0}(\omega).$$

Similarly we say that L has a *negative monotonic effect* relative to $\mathbf{B} \cup \{L\}$ if \overline{L} has a

positive monotonic effect relative to $\mathbf{B} \cup \{\overline{L}\}$. Note that the case in which

$D_{\mathbf{B}\{B_i\}=\mathbf{b}_{-i}, B_i=1}(\omega) = D_{\mathbf{B}\{B_i\}=\mathbf{b}_{-i}, B_i=0}(\omega)$ for all ω , and hence B_i has no effect on D relative to \mathbf{B} , is included as a degenerate case.

The definition of a positive monotonic effect requires that an intervention does not decrease D for every individual, not simply on average, regardless of the other interventions taken. This is thus a strong assumption; see [33] for further discussion.

PROPOSITION 4.10. The set of potential outcomes that are compatible with monotonicity is given by the Sperner number $S(n)$.

*** Cite Dedekind again?

4.5. Tests for irreducibility with monotonicity

Knowledge of the monotonicity of certain potential causes allows for the construction of more powerful statistical tests for irreducibility than those given by Theorem 4.3.

THEOREM 4.11. Let $\mathbf{C} = \mathbf{C}_1 \dot{\cup} \mathbf{C}_2$, $\mathbf{B} = (\mathbf{B}_+ \dot{\cup} \mathbf{B}')$ $\in \dot{\mathcal{P}}(\mathbb{L}(\mathbf{C}_1))$, $|\mathbf{B}| = |\mathbf{C}_1|$ and suppose that every $L \in \mathbf{B}_+$ has a positive monotonic effect on D relative to \mathbf{C} . If for some tree \mathfrak{T} on \mathbf{B}_+ , $\omega^* \in \Omega$ and some \mathbf{c}_2 we have:

$$0 < D_{\mathbf{B}=1, \mathbf{C}_2=\mathbf{c}_2}(\omega^*) - \sum_{L \in \mathbf{B}} D_{\mathbf{B} \setminus \{L\}=1, L=0, \mathbf{C}_2=\mathbf{c}_2}(\omega^*) + \sum_{\mathbf{E} \in \mathfrak{T}} D_{\mathbf{B} \setminus \mathbf{E}=1, \mathbf{E}=0, \mathbf{C}_2=\mathbf{c}_2}(\omega^*). \quad (4.3)$$

then \mathbf{B} is irreducible for $\mathcal{D}(\mathbf{C}, \Omega)$.

If we know that X has a negative monotonic effect on D , then we may use this Theorem to construct more powerful tests of the irreducibility of sets containing \bar{X} with respect to $\mathcal{D}(\mathbf{C}, \Omega)$. Under the assumption that every $L \in \mathbf{C}$ has a monotonic effect on D we have shown via direct calculation using cddlib [?] that for $|\mathbf{C}| = 4$, the conditions (4.3) are the sole restrictions on the law of $(D, \mathbf{C}, \mathbf{W})$ implied by the negation of irreducibility. We conjecture that this holds in general.

Proof: By Theorem 3.2 it is sufficient to prove that under the monotonicity assumption on \mathbf{B}_+ , (4.3) implies (3.1). Suppose that (3.1) does not hold, so that for all values \mathbf{c}_2 , and all $\omega^* \in \Omega$,

$$D_{\mathbf{B}=1, \mathbf{C}_2=\mathbf{c}_2}(\omega^*) - \sum_{L \in \mathbf{B}} D_{\mathbf{B} \setminus \{L\}=1, L=0, \mathbf{C}_2=\mathbf{c}_2}(\omega^*) \leq 0.$$

Then for each $\omega^* \in \Omega$ there exists $R \in \mathbf{B}_+$ such that

$$D_{\mathbf{B}=1, \mathbf{C}_2=\mathbf{c}_2}(\omega^*) - \sum_{L \in \mathbf{B}' \cup \{R\}} D_{\mathbf{B} \setminus \{L\}=1, L=0, \mathbf{C}_2=\mathbf{c}_2}(\omega^*) \leq 0.$$

For a given tree \mathfrak{T} , the remaining terms on the RHS of (4.3) are:

$$\begin{aligned}
 & - \sum_{L \in \mathbf{B}_+ \setminus \{R\}} D_{\mathbf{B} \setminus \{L\}=1, L=0, \mathbf{C}_2=\mathbf{c}_2}(\omega^*) + \sum_{\mathbf{E} \in \mathfrak{T}} D_{\mathbf{B} \setminus \mathbf{E}=1, \mathbf{E}=0, \mathbf{C}_2=\mathbf{c}_2}(\omega^*) \\
 & = \sum_{L \in \mathbf{B}_+ \setminus \{R\}} \left(D_{\mathbf{B} \setminus \phi_{\mathfrak{T}}^{\tilde{L}}(L)=1, \phi_{\mathfrak{T}}^{\tilde{L}}(L)=0, \mathbf{C}_2=\mathbf{c}_2}(\omega^*) - D_{\mathbf{B} \setminus \{L\}=1, L=0, \mathbf{C}_2=\mathbf{c}_2}(\omega^*) \right),
 \end{aligned}$$

by Proposition 4.7. Finally since $\phi_{\mathfrak{T}}^{\tilde{L}}(L) = \{L, L'\} \subseteq \mathbf{B}_+$, L' has a positive monotonic effect on D relative to \mathbf{C} , thus no term in the sum is positive. Consequently for all $\omega^* \in \Omega$, (4.3) does not hold for any tree \mathfrak{T} .

EXAMPLE 2. In the case $\mathbf{B} = \{X_1, X_2\} = \mathbf{B}_+ = \mathbf{C}$, where there is only one tree on \mathbf{B}_+ , consisting of a single edge $X_1 - X_2$. Thus if X_1 and X_2 have a positive monotonic effect on D (relative to \mathbf{C}) then Theorem 4.11 implies that if the following inequality holds for some $\omega \in \Omega$ then

$$D_{11}(\omega) - (D_{10}(\omega) + D_{01}(\omega)) + D_{00}(\omega) > 0$$

then $\{X_1, X_2\}$ is irreducible for $\mathcal{D}(\mathbf{C}, \Omega)$.

EXAMPLE 3. If $\mathbf{B} = \{X_1, X_2, X_3\} = \mathbf{B}_+ = \mathbf{C}$, then there are three trees on \mathbf{B}_+ , see Figure 1. These correspond to the following conditions:

- a. $D_{111}(\omega) - (D_{110}(\omega) + D_{101}(\omega) + D_{011}(\omega)) + (D_{010}(\omega) + D_{001}(\omega)) > 0$;
- b. $D_{111}(\omega) - (D_{110}(\omega) + D_{101}(\omega) + D_{011}(\omega)) + (D_{100}(\omega) + D_{001}(\omega)) > 0$;
- c. $D_{111}(\omega) - (D_{110}(\omega) + D_{101}(\omega) + D_{011}(\omega)) + (D_{100}(\omega) + D_{010}(\omega)) > 0$.

Thus \mathbf{B} is irreducible for $\mathcal{D}(\mathbf{C}, \Omega)$ if at least one holds for some $\omega \in \Omega$.

COROLLARY 4.12. Let $\mathbf{C} = \mathbf{C}_1 \dot{\cup} \mathbf{C}_2$, $\mathbf{B} = (\mathbf{B}_+ \dot{\cup} \mathbf{B}')$ $\in \dot{\mathbb{P}}(\mathbb{L}(\mathbf{C}_1))$ $|\mathbf{B}| = |\mathbf{C}_1|$. Suppose that every $L \in \mathbf{B}_+$ has a positive monotonic effect on D relative to \mathbf{C} , and \mathbf{W} is sufficient to adjust for confounding of \mathbf{C} on D . If for some tree \mathfrak{T} on \mathbf{B}_+ , and some \mathbf{c}_2, \mathbf{w} we have:

$$\begin{aligned}
 & 0 < E[D | \mathbf{B} = 1, \mathbf{C}_2 = \mathbf{c}_2, \mathbf{W} = \mathbf{w}] \\
 & - \sum_{L \in \mathbf{B}} E[D | \mathbf{B} \setminus \{L\} = 1, L = 0, \mathbf{C}_2 = \mathbf{c}_2, \mathbf{W} = \mathbf{w}] \\
 & + \sum_{\mathbf{E} \in \mathfrak{T}} E[D | \mathbf{B} \setminus \mathbf{E} = 1, \mathbf{E} = 0, \mathbf{C}_2 = \mathbf{c}_2, \mathbf{W} = \mathbf{w}],
 \end{aligned} \tag{4.4}$$

then \mathbf{B} is irreducible for $\mathcal{D}(\mathbf{C}, \Omega)$.

Proof: Directly analogous to the proof of Theorem 4.3.

The special case of the previous Corollary where $|\mathbf{B}_+| = |\mathbf{C}| = 2$, and $\mathbf{W} = \emptyset$, appears in Rothman and Greenland [22]; see also Koopman [10].

Theorem 4.8 implies that if every literal in \mathbf{B} has a positive monotonic effect on D then we will have $|\mathbf{B}|^{|\mathbf{B}|-2}$ conditions to test, each of which is sufficient to establish the irreducibility

of \mathbf{B} for $\mathcal{D}(\mathbf{C}, \Omega)$. As before, the conditions (4.4) may be tested via t -test type statistics or using various statistical models.

As with the results in §4.2, we may apply Corollary 3.8 to establish that \mathbf{B} is irreducible for $D(\mathbf{C} \cup \mathbf{C}')$ if every variable in \mathbf{C}' is not causally influenced by variables in \mathbf{C} .

4.6. Tests for a minimal sufficient cause under monotonicity

As noted in section §3.1 if $|\mathbf{B}| = |\mathbf{C}|$ then irreducible conjunctions are also minimal sufficient causes. Thus in this special case the tests of irreducibility given in Theorem 4.3 and Corollary 4.12 also establish that \mathbf{B} is a minimal sufficient cause relative to \mathbf{C} . When $|\mathbf{B}| < |\mathbf{C}|$ these tests do not in general establish this. However, under positive monotonicity assumptions on \mathbf{C}_2 , such tests may be obtained by taking $\mathbf{c}_2 = \mathbf{0}$:

PROPOSITION 4.13. *Let $\mathbf{C} = \mathbf{C}_1 \dot{\cup} \mathbf{C}_2$, $\mathbf{B} \in \dot{\mathbb{P}}(\mathbb{L}(\mathbf{C}_1))$, $|\mathbf{B}| = |\mathbf{C}_1|$. Suppose every $L \in \mathbf{C}_2$ has a positive monotonic effect on D relative to \mathbf{C} . If (i) $D_{\mathbf{B}=\mathbf{1}, \mathbf{C}_2=\mathbf{0}}(\omega^*) = 1$ and (ii) for all $L \in \mathbf{B}$, $D_{\mathbf{B} \setminus \{L\}=\mathbf{1}, L=\mathbf{0}, \mathbf{C}_2=\mathbf{0}}(\omega^*) = 0$, then \mathbf{B} is a minimal sufficient cause for D relative to \mathbf{C} for ω^* .*

Proof: For any \mathbf{c}_2 , $D_{\mathbf{B}=\mathbf{1}, \mathbf{C}_2=\mathbf{c}_2}(\omega) = D_{\mathbf{B}=\mathbf{1}, \mathbf{C}_2=\mathbf{0}}(\omega)$ by the monotonicity assumption. Hence \mathbf{B} is a sufficient cause for D relative to \mathbf{C} for ω^* . Minimality follows directly from condition (ii).

We have the following corollaries which provide conditions under which \mathbf{B} is a minimal sufficient cause for D relative to \mathbf{C} for some $\omega \in \Omega$, in addition to being irreducible for $\mathcal{D}(\mathbf{C}, \Omega)$:

COROLLARY 4.14. *Let $\mathbf{C} = \mathbf{C}_1 \dot{\cup} \mathbf{C}_2$, $\mathbf{B} \in \dot{\mathbb{P}}(\mathbb{L}(\mathbf{C}_1))$, $|\mathbf{B}| = |\mathbf{C}_1|$. Suppose every $L \in \mathbf{C}_2$ has a positive monotonic effect on D relative to \mathbf{C} and \mathbf{W} is sufficient to adjust for confounding of \mathbf{C} on D . If (4.2) holds with $\mathbf{c}_2 = \mathbf{0}$ for some \mathbf{w} then \mathbf{B} is a minimal sufficient cause of D relative to \mathbf{C} for some $\omega \in \Omega$.*

Proof: Immediate from Proposition 4.13 and Theorem 4.3.

COROLLARY 4.15. *Let $\mathbf{C} = \mathbf{C}_1 \dot{\cup} \mathbf{C}_2$, $\mathbf{B} = (\mathbf{B}_+ \dot{\cup} \mathbf{B}') \in \dot{\mathbb{P}}(\mathbb{L}(\mathbf{C}_1))$, $|\mathbf{B}| = |\mathbf{C}_1|$. Suppose that every $L \in \mathbf{B}_+ \cup \mathbf{C}_2$ has a positive monotonic effect on D relative to \mathbf{C} , and \mathbf{W} is sufficient to adjust for confounding of \mathbf{C} on D .*

If (4.4) holds with $\mathbf{c}_2 = \mathbf{0}$ for some \mathbf{w} and some tree \mathfrak{T} on \mathbf{B}_+ then \mathbf{B} is a minimal sufficient cause of D relative to \mathbf{C} for some $\omega \in \Omega$.

Proof: Immediate from Proposition 4.13 and Corollary 4.12.

5. Singular interactions

In the genetics literature, in the context of two binary genetic factors, X_1 and X_2 , ‘biologic’ or ‘compositional’ *epistasis* [2, 5, 17] is said to be present if for some ω^* , $D_{11}(\omega^*) = 1$ but $D_{10}(\omega^*) = D_{01}(\omega^*) = D_{00}(\omega^*) = 0$; in this case the effect of one genetic factor is effectively

masked when the other genetic factor is absent. If $\{X_1, X_2\}$ is a minimal sufficient cause of D relative to $\{X_1, X_2\}$ for ω^* then although this implies $D_{11}(\omega^*) = 1$ and $D_{10}(\omega^*) = D_{01}(\omega^*) = 0$, it does *not* imply $D_{00}(\omega^*) = 0$. This motivates the following:

DEFINITION 5.1. A minimal sufficient cause \mathbf{B} for D relative to \mathbf{C} for ω^* is said to be singular if there is no $\mathbf{B}' \in \dot{\mathbb{P}}(\mathbb{L}(\mathbf{C}))$, $\mathbf{B}' \subsetneq \mathbf{B}$, forming a minimal sufficient cause for D relative to \mathbf{C} for ω^* . \mathbf{B} is singular for $\mathcal{D}(\mathbf{C}, \Omega)$ if \mathbf{B} is singular relative to \mathbf{C} for some $\omega^* \in \Omega$.

If \mathbf{B} is singular for $\mathcal{D}(\mathbf{C}, \Omega)$ then we will also refer to a *singular interaction* between the components of \mathbf{B} . We now characterize singularity in terms of potential outcomes:

THEOREM 5.2. Let $\mathbf{C} = \mathbf{C}_1 \cup \mathbf{C}_2$, $\mathbf{B} \in \dot{\mathbb{P}}(\mathbb{L}(\mathbf{C}_1))$, $|\mathbf{B}| = |\mathbf{C}_1|$ then \mathbf{B} is singular for $\mathcal{D}(\mathbf{C}, \Omega)$ iff there exists $\omega^* \in \Omega$ such that

$$\text{for all values } \mathbf{c}_2^*, \mathbf{b}: D_{\mathbf{B}=\mathbf{b}, \mathbf{C}_2=\mathbf{c}_2^*}(\omega^*) = 1 \iff \mathbf{b} = \mathbf{1}. \quad (5.1)$$

Note that (5.1) is equivalent to:

$$D_{\mathbf{C}=\mathbf{c}}(\omega^*) = \left(\bigwedge (\mathbf{B}) \right)_{\mathbf{c}} \text{ for all } \mathbf{c}. \quad (5.2)$$

Thus if \mathbf{B} is singular for $\mathcal{D}(\mathbf{C}, \Omega)$ then there is some individual ω^* whose potential outcomes are given by the single conjunction \mathbf{B} .

Proof: By definition, \mathbf{B} is a sufficient cause for D for ω^* iff $(\mathbf{b} = \mathbf{1} \Rightarrow D_{\mathbf{B}=\mathbf{b}, \mathbf{C}_2=\mathbf{c}_2^*}(\omega^*) = 1)$. Thus it is sufficient to show that, assuming \mathbf{B} is a minimal sufficient cause for D for ω^* , there are no other minimal sufficient causes of D for ω^* iff $(D_{\mathbf{B}=\mathbf{b}, \mathbf{C}_2=\mathbf{c}_2^*}(\omega^*) = 1 \Rightarrow \mathbf{b} = \mathbf{1})$.

Suppose \mathbf{B} is the only minimal sufficient cause for D for ω^* , but that for some $\mathbf{b}^* \neq \mathbf{1}$, $D_{\mathbf{B}=\mathbf{b}^*, \mathbf{C}_2=\mathbf{c}_2^*}(\omega^*) = 1$. Let $\mathbf{B}^\dagger \equiv \mathbf{B}[\mathbf{B}=\mathbf{b}^*, \mathbf{C}_2=\mathbf{c}_2^*]$, \mathbf{B}^\dagger forms a sufficient cause for D for ω^* , and $\mathbf{B} \not\subseteq \mathbf{B}^\dagger$. Hence there is some $\mathbf{B}' \subsetneq \mathbf{B}^\dagger$ that is a minimal sufficient cause for D for ω^* , and $\mathbf{B} \not\subseteq \mathbf{B}'$, a contradiction.

Conversely suppose $(D_{\mathbf{B}=\mathbf{b}, \mathbf{C}_2=\mathbf{c}_2^*}(\omega^*) = 1 \Rightarrow \mathbf{b} = \mathbf{1})$ but there exists another minimal sufficient cause \mathbf{B}' for D for ω^* , and $\mathbf{B} \not\subseteq \mathbf{B}'$. Since \mathbf{B}' is minimal, $\mathbf{B} \not\subseteq \mathbf{B}'$. Thus there exists a $\tilde{\mathbf{c}}$ such that $(\mathbf{B})_{\tilde{\mathbf{c}}} \neq 1$, but $(\mathbf{B}')_{\tilde{\mathbf{c}}} = 1$ and hence $D_{\mathbf{C}=\tilde{\mathbf{c}}}(\omega^*) = 1$, a contradiction.

COROLLARY 5.3. For $\mathcal{D}(\mathbf{C}, \Omega)$, if \mathbf{B} is singular then \mathbf{B} is irreducible.

Proof: Immediate from (5.1) and the definition of irreducibility.

Theorem 5.4 relates singular interactions to properties of the set of sufficient cause representations for $\mathcal{D}(\mathbf{C}, \Omega)$.

THEOREM 5.4. Let $\mathbf{B} \in \dot{\mathbb{P}}(\mathbb{L}(\mathbf{C}))$. \mathbf{B} is singular for $\mathcal{D}(\mathbf{C}, \Omega)$ iff there exists $\omega^* \in \Omega$ such that in every representation $(\mathbf{A}, \mathfrak{B})$ for $\mathcal{D}(\mathbf{C}, \Omega)$, (i) for all $\mathbf{B}^* \in \dot{\mathbb{P}}(\mathbb{L}(\mathbf{C}))$, with $|\mathbf{B}^*| = |\mathbf{C}|$ and $\mathbf{B} \subseteq \mathbf{B}^*$ there exists $\mathbf{B}_i \in \mathfrak{B}$ with $\mathbf{B}_i \subseteq \mathbf{B}^*$ and $A_i(\omega^*) = 1$; (ii) for all $\mathbf{B}_i \in \mathfrak{B}$ such that $\mathbf{B} \not\subseteq \mathbf{B}_i$, $A_i(\omega^*) = 0$.

Proof: Let $\mathbf{C} = \mathbf{C}_1 \dot{\cup} \mathbf{C}_2$, where $\mathbf{B} \in \dot{\mathbb{P}}(\mathbb{L}(\mathbf{C}_1))$, and $|\mathbf{B}| = |\mathbf{C}_1|$.

(\Rightarrow) Suppose \mathbf{B} is singular for $\mathcal{D}(\mathbf{C}, \Omega)$, so that some $\omega^* \in \Omega$ satisfies (5.2). Then for any representation $(\mathbf{A}, \mathfrak{B})$ for $\mathcal{D}(\mathbf{C}, \Omega)$ and any \mathbf{B}^* such that $|\mathbf{B}^*| = |\mathbf{C}|$ and $\mathbf{B} \subseteq \mathbf{B}^*$ we can select values \mathbf{c}_2^* so that $\mathbf{B}^* = [\mathbf{B}=1, \mathbf{C}_2=\mathbf{c}_2^*]$. Since $D_{\mathbf{B}=1, \mathbf{C}_2=\mathbf{c}_2^*}(\omega^*) = 1$ there exists $A_i \in \mathbf{A}$, $\mathbf{B}_i \in \mathfrak{B}$ with $A_i(\omega^*) = 1$ and $\left(\bigwedge_{\mathbf{B}=1, \mathbf{C}_2=\mathbf{c}_2^*} (\mathbf{B}_i)\right) = 1$. Thus $\mathbf{B}_i \subseteq \mathbf{B}^*$, so (i) holds. For all $\mathbf{B}_i \in \mathfrak{B}$ such that $\mathbf{B} \not\subseteq \mathbf{B}_i$, we can choose $\tilde{\mathbf{B}} \in \dot{\mathbb{P}}(\mathbb{L}(\mathbf{C}_1))$, $|\tilde{\mathbf{B}}| = |\mathbf{C}_1|$ with $\tilde{\mathbf{B}} \neq \mathbf{B}$ and values $\tilde{\mathbf{c}}_2$ so that $\mathbf{B}_i \subseteq \mathbf{B}[\tilde{\mathbf{B}}=1, \mathbf{C}_2=\tilde{\mathbf{c}}_2]$. Since $D_{\tilde{\mathbf{B}}=1, \mathbf{C}_2=\tilde{\mathbf{c}}_2}(\omega^*) = 0$ we have $A_i(\omega^*) = 0$ since $\left(\bigwedge_{\tilde{\mathbf{B}}=1, \mathbf{C}_2=\tilde{\mathbf{c}}_2} (\mathbf{B}_i)\right) = 1$, so (ii) holds as required.

(\Leftarrow) Suppose there exists $\omega^* \in \Omega$ such that every representation $(\mathbf{A}, \mathfrak{B})$ satisfies (i) and (ii). We will show that (5.1) holds. For any values \mathbf{c}_2^* let $\mathbf{B}^* \equiv \mathbf{B}[\mathbf{B}=1, \mathbf{C}_2=\mathbf{c}_2^*]$, so $|\mathbf{B}^*| = |\mathbf{C}|$ and $\mathbf{B} \subseteq \mathbf{B}^*$. Thus by (i) there exists $\mathbf{B}_i \in \mathfrak{B}$ with $\mathbf{B}_i \subseteq \mathbf{B}^*$ and $A_i(\omega^*) = 1$. Hence $D_{\mathbf{B}=1, \mathbf{C}_2=\mathbf{c}_2^*}(\omega^*) = 1$ since $A_i(\omega^*) = 1$ and $\left(\bigwedge_{\mathbf{B}=1, \mathbf{C}_2=\mathbf{c}_2^*} (\mathbf{B}_i)\right) = 1$. Conversely for any $\mathbf{b}' \in \mathbb{L}$, let $\mathbf{B}' \equiv \mathbf{B}[\mathbf{B}=\mathbf{b}']$, so $|\mathbf{B}'| = |\mathbf{C}_1|$ with $\mathbf{B}' \subseteq \mathbf{B}$. Thus for all $\mathbf{B}_i \in \mathfrak{B}$ such that $\left(\bigwedge_{\mathbf{B}'=1, \mathbf{C}_2=\mathbf{c}_2^*} (\mathbf{B}_i)\right) = 1$, $\mathbf{B} \not\subseteq \mathbf{B}_i$ and thus by (ii) $A_i(\omega^*) = 0$. Hence $D_{\mathbf{B}'=1, \mathbf{C}_2=\mathbf{c}_2^*}(\omega^*) = 0$.

We now consider results relevant for testing for singular interactions with or without monotonicity assumptions.

THEOREM 5.5. Let $\mathbf{B} = \mathbf{B}_+ \dot{\cup} \mathbf{B}' \in \dot{\mathbb{P}}(\mathbb{L}(\mathbf{C}))$, $|\mathbf{B}| = |\mathbf{C}|$ and suppose that every $L \in \mathbf{B}_+$ has a positive monotonic effect on D relative to \mathbf{C} . If for some tree \mathfrak{T} on \mathbf{B}_+ and some $\omega^* \in \Omega$ we have

$$\begin{aligned} D_{\mathbf{B}=1}(\omega^*) &= \sum_{L \in \mathbf{B}_+} D_{\mathbf{B} \setminus \{L\}=1, L=0}(\omega^*) \\ &= \sum_{\tilde{\mathbf{B}}: \emptyset \neq \tilde{\mathbf{B}} \subseteq \mathbf{B}'} D_{\mathbf{B} \setminus \tilde{\mathbf{B}}=1, \tilde{\mathbf{B}}=0}(\omega^*) + \sum_{\mathbf{E} \in \mathfrak{T}} D_{\mathbf{B} \setminus \mathbf{E}=1, \mathbf{E}=0}(\omega^*) > 0 \end{aligned} \quad (5.3)$$

then \mathbf{B} is singular for $\mathcal{D}(\mathbf{C}, \Omega)$.

Proof: By Theorem 5.4 \mathbf{B} is singular for $\mathcal{D}(\mathbf{C}, \Omega)$ iff

$$\text{for some } \omega^* \in \Omega, \quad D_{\mathbf{B}=1}(\omega^*) - \sum_{\tilde{\mathbf{B}} \subseteq \mathbf{B}} D_{\mathbf{B} \setminus \tilde{\mathbf{B}}=1, \tilde{\mathbf{B}}=0}(\omega^*) > 0. \quad (5.4)$$

Suppose for a contradiction that (5.4) does not hold but (5.3) holds for some $\omega^* \in \Omega$. Since \mathbf{B}_+ has a positive monotonic effect on D relative to \mathbf{C} , if $\tilde{\mathbf{B}} \subseteq \mathbf{B}$ is such that $\tilde{\mathbf{B}} \cap \mathbf{B}_+ \neq \emptyset$ then $D_{\tilde{\mathbf{B}}|\tilde{\mathbf{B}}=1, \tilde{\mathbf{B}}=0}(\omega^*)=1$ implies $D_{\mathbf{B} \setminus \{L\}=1, L=0}(\omega^*)=1$ for some $L \in \tilde{\mathbf{B}}$. Hence for all $\omega \in \Omega$

$$D_{\mathbf{B}=1}(\omega) - \sum_{L \in \mathbf{B}} D_{\mathbf{B} \setminus \{L\}=1, L=0}(\omega) - \sum_{\substack{\tilde{\mathbf{B}} \subseteq \mathbf{B}' \\ |\tilde{\mathbf{B}}| \geq 2}} D_{\tilde{\mathbf{B}}|\tilde{\mathbf{B}}=1, \tilde{\mathbf{B}}=0}(\omega) \leq 0. \tag{5.5}$$

By applying the same argument to the first two terms on the LHS of (5.5) as was applied in the proof of Theorem 4.11 we have that (5.3) does not hold for all $\omega \in \Omega$, which is a contradiction.

The following corollary to Theorem 5.5 generalizes the discussion in [27,28] to an arbitrary number of dichotomous factors:

Corollary 5.6. *Let $\mathbf{B} = \mathbf{B}_+ \dot{\cup} \mathbf{B}' \in \mathbb{P}(\mathbb{L}(\mathbf{C}))$, $|\mathbf{B}| = |\mathbf{C}|$. Suppose that every $L \in \mathbf{B}_+$ has a positive monotonic effect on D relative to \mathbf{B} , and \mathbf{W} is sufficient to adjust for confounding of \mathbf{C} on D . If for some tree \mathfrak{T} on \mathbf{B}_+ , and some \mathbf{w} we have:*

$$\begin{aligned} 0 < & E [D|\mathbf{B}=1, \mathbf{W}=\mathbf{w}] - \sum_{L \in \mathbf{B}_+} E [D|\mathbf{B} \setminus \{L\}=1, L=0, \mathbf{W}=\mathbf{w}] \\ & - \sum_{\substack{\tilde{\mathbf{B}}: \emptyset \neq \tilde{\mathbf{B}} \subseteq \mathbf{B}' \\ \mathbf{E} \in \mathfrak{T}}} E [D|\tilde{\mathbf{B}}|\tilde{\mathbf{B}}=1, \tilde{\mathbf{B}}=0, \mathbf{W}=\mathbf{w}] \\ & + \sum_{\mathbf{E} \in \mathfrak{T}} E [D|\mathbf{B} \setminus \mathbf{E}=1, \mathbf{E}=0, \mathbf{W}=\mathbf{w}] \end{aligned} \tag{5.6}$$

then \mathbf{B} is singular for $\mathcal{D}(\mathbf{C}, \Omega)$.

Proof: By applying Proposition 4.2 to each term in (5.3).

Condition (5.6) leads directly to a statistical test of compositional epistasis. This is notable since some claims in the genetics literature appear to suggest that such tests did not exist [5].

As stated in the next corollary, from Theorem 5.5 if all or all but one of the elements of \mathbf{B} have positive monotonic effects on D then singularity and irreducibility coincide:

COROLLARY 5.7. *Suppose $|\mathbf{B}| = |\mathbf{C}|$ and that for all or all but one of $B_i \in \mathbf{B}$, B_i has a positive monotonic effect on D relative to \mathbf{B} then \mathbf{B} is singular for $\mathcal{D}(\mathbf{C}, \Omega)$ if \mathbf{B} is irreducible for $\mathcal{D}(\mathbf{C}, \Omega)$.*

An important consequence of this corollary is that when there is at most one variable that does not have a positive monotonic effect condition (4.4) establishes that \mathbf{B} is singular in addition to being irreducible for $\mathcal{D}(\mathbf{C}, \Omega)$.

Proof: Let \mathbf{B}' denote the one or zero elements of \mathbf{B} that do not have a monotonic effect on D relative to \mathbf{C} . If \mathbf{B} is irreducible for $\mathcal{D}(\mathbf{C}, \Omega)$ then by Theorem 4.11,

$$D_{\mathbf{B}=1}(\omega^*) - \sum_{L \in \mathbf{B}} D_{\mathbf{B} \setminus \{L\}=1, L=0}(\omega^*) + \sum_{\mathbf{E} \in \mathfrak{T}} D_{\mathbf{B} \setminus \mathbf{E}=1, \mathbf{E}=0}(\omega^*) > 0.$$

Since the third term on the LHS of (5.3) vanishes when $|\mathbf{B}'| = 1$, it follows that \mathbf{B} is singular for $\mathcal{D}(\mathbf{C}, \Omega)$. The converse is given in Corollary 5.3.

COROLLARY 5.8. *Suppose $|\mathbf{B}| = |\mathbf{C}|$ and that for all or all but one of $B_i \in \mathbf{B} \cup \mathbf{C}'$, B_i has a positive monotonic effect on D relative to $\mathbf{B} \cup \mathbf{C}$, for all $X' \in \mathbf{C}'$, X' is not causally influenced by \mathbf{C} and \mathbf{B} is singular for $\mathcal{D}(\mathbf{C}, \Omega)$ then \mathbf{B} is singular for $\mathcal{D}(\mathbf{C} \cup \mathbf{C}', \Omega)$.*

Proof: By Corollary 5.7, \mathbf{B} is irreducible relative to $\mathcal{D}(\mathbf{C}, \Omega)$. Hence by Corollary 3.8 \mathbf{B} is singular relative to $\mathcal{D}(\mathbf{C} \cup \mathbf{C}', \Omega)$. The conclusion then follows from a further application of Corollary 5.7.

5.1. Relation to Pearl's Probability of Causation

Pearl [16, chapter 9] defined the probability of necessity and sufficiency (PNS) of cause X for outcome D to be $P(D_{X=1}(\omega) = 1, D_{X=0}(\omega) = 0)$. In other words $\text{PNS}(D, X)$ is the probability that D would occur if X occurred and would not have done so had X not occurred. We generalize this to the setting in which there are multiple causes \mathbf{B} :

DEFINITION 5.9. *For $\mathbf{B} \subseteq \dot{\mathbb{P}}(\mathbb{L}(\mathbf{C}))$, the probability of necessity and sufficiency of \mathbf{B} causing D , is:*

$$\text{PNS}(D, \mathbf{B}) \equiv P(D_{\mathbf{B}=1} = 1 \text{ and for all } \mathbf{b} \neq 1, D_{\mathbf{B}=\mathbf{b}} = 0).$$

Thus $\text{PNS}(D, \mathbf{B})$ is the probability that D would occur if every literal $L \in \mathbf{B}$ occurred and would not have done so had at least one literal in \mathbf{B} not occurred.

PROPOSITION 5.10. *If $|\mathbf{B}| = |\mathbf{C}|$ then $\text{PNS}(D, \mathbf{B}) > 0$ if \mathbf{B} is singular for $\mathcal{D}(\mathbf{C}, \Omega)$.*

Proof: Follows directly from Theorem 5.4 and Definition 5.9.

This connection also provides an interpretation for condition (5.6).

PROPOSITION 5.11. *Under the conditions of Corollary 5.6, with $\mathbf{W} = \emptyset = \mathbf{B}_+$ $\text{PNS}(\mathbf{B}, D)$ is bounded below by the RHS of (5.6).*

Proof:

$$\begin{aligned} \text{PNS}(D, \mathbf{B}) &= P(D_{\mathbf{B}=1} = 1 \text{ and for all } \mathbf{b} \neq 1, D_{\mathbf{B}=\mathbf{b}} = 0) \\ &\geq P(D_{\mathbf{B}=1} = 1) + P(\text{for all } \mathbf{b} \neq 1, D_{\mathbf{B}=\mathbf{b}} = 0) - 1 \\ &= P(D_{\mathbf{B}=1} = 1) - P(\text{for some } \mathbf{b} \neq 1, D_{\mathbf{B}=\mathbf{b}} = 1) \\ &\geq P(D_{\mathbf{B}=1} = 1) - \sum_{\mathbf{b} \neq 1} P(D_{\mathbf{B}=\mathbf{b}} = 1) \\ &= E[D=1|\mathbf{B}=1] - \sum_{\mathbf{b} \neq 1} E[D|\mathbf{B}=\mathbf{b}] \end{aligned}$$

which is the RHS of (5.6) with $\mathbf{W} = \emptyset = \mathbf{B}_+$.

This generalizes some of the lower bounds on $\text{PNS}(D, X)$ given by Pearl [16, §9.2]. The assumption in Proposition 5.11 that \mathbf{B} and D are unconfounded, so $\mathbf{W} = \emptyset$, and the absence of monotonicity assumptions so $\mathbf{B}_+ = \emptyset$, are for expositional convenience.

6. Relation to Statistical Models with Linear Links

In related work [32] it is noted that the presence of interaction terms in statistical models do not in general correspond to sufficient conditions for irreducibility. Consider for example a saturated Bernoulli regression model for D with identity link and binary regressors $\mathbf{C} = \{X_1, \dots, X_p\}$:

$$E[D|\mathbf{C}=\mathbf{c}] = \sum_{\tilde{\mathbf{B}} \subseteq \mathbf{C}} \beta_{\tilde{\mathbf{B}}} \left(\bigwedge_{\tilde{\mathbf{B}}} (\tilde{\mathbf{B}}) \right)_{\mathbf{c}}. \tag{6.1}$$

Note that with $\mathbf{c} = (x_1, \dots, x_p)$ then $\left(\bigwedge_{\tilde{\mathbf{B}}} (\tilde{\mathbf{B}}) \right)_{\mathbf{c}} = \prod_{x_i \in \tilde{\mathbf{B}}} x_i$, the usual multiplicative interaction term. The conditions, given earlier, for detecting the presence of irreducibility and singularity lead to linear restrictions on the regression coefficients $\beta_{\tilde{\mathbf{B}}}$:

$$\sum_{\tilde{\mathbf{B}} \subseteq \mathbf{C}} m_{\tilde{\mathbf{B}}} \beta_{\tilde{\mathbf{B}}} > 0. \tag{6.2}$$

Note that (6.2) includes an intercept β_{\emptyset} . First we define

$$\text{deg}_{\mathfrak{T}}(L) \equiv \{ \mathbf{E} | \mathbf{E} \in \mathfrak{T}, L \in \mathbf{E} \},$$

the *degree* of L in a tree \mathfrak{T} , to be the number of edges in \mathfrak{T} that contain L .

PROPOSITION 6.1. *Under the conditions of Theorem 4.11, with $\mathbf{B} = \mathbf{C}$, condition (4.3) is equivalent to restriction (6.2) with $m_{\tilde{\mathbf{B}}} = m_{\tilde{\mathbf{B}}}^{\text{irred}}$ where:*

$$m_{\tilde{\mathbf{B}}}^{\text{irred}} \equiv 1 - |\mathbf{B} \setminus \tilde{\mathbf{B}}| + |\mathfrak{T}| - \sum_{L \in \tilde{\mathbf{B}} \cap \mathbf{B}_+} \text{deg}_{\mathfrak{T}}(L) + \left| \{ \mathbf{E} | \mathbf{E} \in \mathfrak{T}, \mathbf{E} \subseteq \tilde{\mathbf{B}} \cap \mathbf{B}_+ \} \right|. \tag{6.3}$$

Note that since \mathfrak{T} is a tree on \mathbf{B}_+ , the last term in (6.3) has a natural graphical interpretation as the number of edges in the ‘induced subgraph’ of \mathfrak{T} on the subset $\tilde{\mathbf{B}}$. Definition (6.3) also subsumes condition (4.2) given in Theorem 4.3 (without monotonicity), in which case the last three terms in (6.3) vanish. Though Proposition 6.1 assumes that $\mathbf{C}_2 = \emptyset$, the condition given by (6.2) and (6.3) continues to apply in the case where $\mathbf{c}_2 = \mathbf{0}$ as in Corollaries 4.14 and 4.15; obvious extensions apply to the case where $\mathbf{c}_2 = \mathbf{0}$.

Proof: This follows by counting the number (and sign) of expectations in (4.3) for which $\tilde{\mathbf{B}}$ is a subset of the variables assigned the value 1 in the conditioning event. The first two terms in (4.3) correspond, respectively, to the first two terms in (6.3). The remaining three terms in (6.3) correspond to the last sum in (4.3): $|\mathfrak{T}|$, the number of edges in \mathfrak{T} , is the total

number of terms in the sum. The sum over degrees subtracts the number of terms in which some $L \in \tilde{\mathbf{B}}$ is assigned zero. Since this double counts terms corresponding to edges contained in $\tilde{\mathbf{B}}$ the last term corrects for this.

PROPOSITION 6.2. *Under the conditions of Theorem 5.5, with $\mathbf{B} = \mathbf{C}$, condition (5.3) is equivalent to restriction (6.2) with $m_{\tilde{\mathbf{B}}} = m_{\tilde{\mathbf{B}}}^{sing}$ where:*

$$m_{\tilde{\mathbf{B}}}^{sing} \equiv m_{\tilde{\mathbf{B}}}^{irred} + \left(|\mathbf{B}' \setminus \tilde{\mathbf{B}}| \right) - \left(2^{|\mathbf{B}' \setminus \tilde{\mathbf{B}}|} - 1 \right). \quad (6.4)$$

Proof: Expression (6.4) follows from another counting argument similar to the proof of Proposition 6.1, together with the observation that conditions (4.3) and (5.3) only differ in that the terms in the sum over L in (4.3) for $L \in \mathbf{B}'$ are replaced by a sum over all subsets of \mathbf{B}' .

EXAMPLE 4 (Two-way interactions). Consider the saturated Bernoulli regression with with identity link with $\mathbf{C} = \{X_1, X_2\}$:

$$E [D | X_1 = x_1, X_2 = x_2] = \beta_{\emptyset} + \beta_1 x_1 + \beta_2 x_2 + \beta_{12} x_1 x_2.$$

Suppose that X_1 and X_2 are unconfounded with respect to D , so (4.1) holds with $\mathbf{W} = \emptyset$. Proposition 6.1 implies that $\{X_1, X_2\}$ is irreducible relative to \mathbf{C} if $\beta_{12} > \beta_{\emptyset}$; Proposition 6.2 implies that $\{X_1, X_2\}$ is singular relative to \mathbf{C} if $\beta_{12} > 2\beta_{\emptyset}$. If one of X_1 or X_2 have positive monotonic effects on D relative to \mathbf{C} then Proposition 6.1 and Corollary 5.7 imply that $\{X_1, X_2\}$ is both irreducible and singular relative to \mathbf{C} if $\beta_{12} > \beta_{\emptyset}$. If X_1 and X_2 have positive monotonic effects on D relative to \mathbf{C} then Proposition 6.1 and Corollary 5.7 imply that $\{X_1, X_2\}$ is both irreducible and singular relative to \mathbf{C} if $\beta_{12} > 0$.

Thus only under the assumption of positive monotonic effects for both X_1 and X_2 does the sufficient condition for the irreducibility and singularity of $\{X_1, X_2\}$ coincide with the classical two-way interaction term β_{12} being positive.

It also follows from Proposition 3.4 that if $\{X_1, X_2\}$ is irreducible relative to \mathbf{C} then there exists some $\omega \in \mathbf{C}$ for whom $\{X_1, X_2\}$ is a minimal sufficient cause relative to \mathbf{C} (since $|\mathbf{B}| = |\mathbf{C}|$).

EXAMPLE 5 (Three-way interactions). The saturated Bernoulli regression with three binary variables and a identity link can be written as:

$$E [D = 1 | X_1 = x_1, X_2 = x_2, X_3 = x_3] = \beta_{\emptyset} + \beta_1 x_1 + \beta_2 x_2 + \beta_3 x_3 + \beta_{12} x_1 x_2 + \beta_{13} x_1 x_3 + \beta_{23} x_2 x_3 + \beta_{123} x_1 x_2 x_3.$$

Suppose that $\mathbf{C} = \{X_1, X_2, X_3\}$ is unconfounded for D . Proposition 6.1 implies that $\{X_1, X_2, X_3\}$ is irreducible relative to \mathbf{C} if

$$\beta_{123} > 2\beta_{\emptyset} + \beta_1 + \beta_2 + \beta_3. \quad (6.5)$$

It follows from Proposition 3.4 that if $\{X_1, X_2, X_3\}$ is irreducible relative to \mathbf{C} then there exists some $\omega \in \Omega$ for whom $\{X_1, X_2, X_3\}$ is a minimal sufficient cause relative to \mathbf{C} (since $|\mathbf{B}| = |\mathbf{C}|$).

Proposition 6.2 implies $\{X_1, X_2, X_3\}$ is singular relative to \mathbf{C} if

$$\beta_{123} > 6\beta_{\emptyset} + 2\beta_1 + 2\beta_2 + 2\beta_3.$$

However, if X_1, X_2 and X_3 have positive monotonic effects on D (relative to \mathbf{C}) then Proposition 6.1 implies $\{X_1, X_2, X_3\}$ is irreducible relative to \mathbf{C} if any of the following hold:

$$\beta_{123} > \beta_1, \quad \beta_{123} > \beta_2, \quad \beta_{123} > \beta_3; \quad (6.6)$$

equivalently, $\beta_{123} > \min\{\beta_1, \beta_2, \beta_3\}$. By Corollary 5.7 this also establishes that $\{X_1, X_2, X_3\}$ is singular relative to \mathbf{C} .

If only X_1 and X_2 have positive monotonic effects on D relative to \mathbf{C} then Proposition 6.1 implies that $\{X_1, X_2, X_3\}$ is irreducible relative to \mathbf{C} if

$$\beta_{123} > \beta_{\emptyset} + \beta_1 + \beta_2. \quad (6.7)$$

By Corollary 5.7 condition (6.7) also implies that $\{X_1, X_2, X_3\}$ is singular relative to \mathbf{C} (since only X_3 does not have a positive monotonic effect on D). As we would expect condition (6.7) is weaker than (6.5) but stronger than any of the conditions (6.6). If only one potential cause has a monotonic effect on D relative to \mathbf{C} then we can only use (6.5) to establish irreducibility.

Thus for three-way interactions $\beta_{123} > 0$ does not correspond to any of the sufficient conditions for irreducibility or singularity of $\{X_1, X_2, X_3\}$ relative to \mathbf{C} , regardless of whether or not monotonicity assumptions hold.

7. Interpretation of sufficient cause models

As mentioned in §2.1 the observed data for an individual $(\mathbf{C}(\omega), D(\omega))$ represents a strict subset of the potential outcomes $\mathcal{D}(\mathbf{C}, \omega)$; this is the ‘fundamental problem of causal inference’. Further, as we have seen, for a given set of potential outcomes there can exist different determinative sets of minimal sufficient causes \mathbf{B} for the same set of potential outcomes; see (2.3) and (2.4). Thus we have the following for an individual:

$$\begin{array}{ccc}
 \vdots & \searrow & \vdots \\
 \mathbf{B} & \rightarrow & \mathcal{D}(\mathbf{C}, \omega) \\
 \vdots & \nearrow & \vdots
 \end{array}
 \quad
 \begin{array}{ccc}
 \vdots & \searrow & \\
 & \rightarrow & (\mathbf{C}(\omega), D(\omega)) \\
 \vdots & \nearrow &
 \end{array}
 \quad (7.1)$$

many-one many-one

It is typically impossible to know the set of potential outcomes for an individual $\mathcal{D}(\mathbf{C}, \omega)$, even when $\mathbf{C} = \{X\}$, even from randomized experiments. However possession of this knowledge would permit one to predict how a given individual would respond under any given pattern of exposures $\mathbf{C} = \mathbf{c}$.

The results in this paper demonstrate that, given data from a randomized experiment (or when sufficient variables are measured to adjust for confounding) it is possible to infer the existence of an individual for whom all sets of minimal sufficient causes \mathfrak{B} have certain features in common. However, given the double many-one relationship (7.1), and the fact that the set of potential outcomes $\mathcal{D}(\mathbf{C}, \omega)$, if they were known, apparently address all potential counterfactual queries, it is natural to ask what is to be gained by considering such representations. We now motivate our results by presenting several different interpretations of sufficient cause representations.

7.1. The descriptive interpretation

Under this view sets of minimal sufficient causes are merely a way to describe the set of potential outcomes $\mathcal{D}(\mathbf{C}, \Omega)$. The representation $(\mathbf{A}, \mathfrak{B})$ may be more compact; compare Table 2 and (2.5). Extending this to a population Ω , the variables \mathbf{A} in a representation $(\mathbf{A}, \mathfrak{B})$ merely describe subpopulations with particular patterns of potential outcomes. Knowing that there exists an individual for whom all representations \mathfrak{B} have certain features in common provides qualitative information about the set of potential outcomes.

For two binary causes $\{X_1, X_2\}$, Theorem 3.2 implies that $\{X_1, X_2\}$ is irreducible relative to \mathbf{C} for ω^* if $D_{11}(\omega^*) = 1$ and $D_{10}(\omega^*) = D_{01}(\omega^*) = 0$. Such a pattern is of interest insofar as it indicates that the causal process resulting in this individual's potential outcomes $\mathcal{D}(\mathbf{C}, \omega^*)$ is such that (for some setting of the variables in $\mathbf{C} \setminus \{X_1, X_2\}$), $D = 1$ if both $X_1 = 1$ and $X_2 = 1$, but not when $X_1 = 1$ and $X_2 = 0$ or vice versa.

Similarly it follows from Theorem 5.2 that if $\{X_1, X_2\}$ is singular relative to \mathbf{C} for ω^* if $D_{11}(\omega^*) = 1$ and $D_{10}(\omega^*) = D_{01}(\omega^*) = D_{00}(\omega^*) = 0$. Hence the causal process producing $\mathcal{D}(\mathbf{C}, \omega^*)$ is such that, for some setting of the variables in $\mathbf{C} \setminus \{X_1, X_2\}$, $D = 1$ if both $X_1 = 1$ and $X_2 = 1$, but not when either $X_1 = 0$ or $X_2 = 0$.

In contrast to the classical notions of interaction arising in linear models (see §6), irreducibility and singularity are causal in that they relate to the potential outcomes. §4 and §5 contain empirical tests for the presence of irreducible or singular interactions.

7.2. Generative mechanism interpretations

A minimal sufficient cause representation may be interpreted in terms of a 'generative mechanism':

DEFINITION 7.1. A mechanism $M(\omega)$ relative to \mathbf{C} takes as input an assignment \mathbf{c} to \mathbf{C} , and outputs a 'state' $M_{\mathbf{c}}(\omega)$ which is either 'active' (1) or 'inactive' (0). A mechanism is said to be generative for D if whenever it is active, the event $D = 1$ is caused, so that $M_{\mathbf{c}}(\omega) = 1$

implies $D_{\mathbf{c}}(\omega) = 1$. Conversely, a mechanism is said to be preventive for D if whenever $M_{\mathbf{c}}(\omega) = 1$, $D_{\mathbf{c}}(\omega) = 0$ is caused.

Though this definition refers to a mechanism ‘causing’ $D = 1$ or $D = 0$, we abstain from defining this more formally in terms of potential outcomes until the next section. Our reason for proceeding in this way is that there may be circumstances in which an investigator is able to posit the existence of a causal mechanism causing $D = 1$ or $D = 0$, e.g. based on experiments manipulating the inputs \mathbf{C} and output D , but lacks sufficiently detailed information to posit well-defined counterfactual outcomes involving interventions on these (hypothesized) mechanisms.

DEFINITION 7.2. A set of generative mechanisms $\mathbf{M} = \langle M^1, \dots, M^P \rangle$ will be said to be exhaustive for a given set of potential outcomes $\mathcal{D}(\mathbf{C}, \Omega)$ if for all $\omega \in \Omega$, and all \mathbf{c} , if $D_{\mathbf{c}}(\omega) = 1$ then for some $M^i \in \mathbf{M}$, $M_{\mathbf{c}}^i(\omega) = 1$.

Note that if \mathbf{M} forms an exhaustive set of mechanisms for $\mathcal{D}(\mathbf{C}, \Omega)$, then it follows that in a context in which no mechanism M^i is active, $D = 0$.

PROPOSITION 7.3. If $\mathbf{M} = \langle M^1, \dots, M^P \rangle$ forms an exhaustive set of generative mechanisms for $\mathcal{D}(\mathbf{C}, \Omega)$ then $\mathcal{D} = \bigvee(\mathbf{M})$ and $D_{\mathbf{c}}(\omega) = \bigvee(\mathbf{M}_{\mathbf{c}}(\omega))$.

Proof: Follows Definitions 7.1 and 7.2. 2

PROPOSITION 7.4. Suppose \mathbf{M} forms an exhaustive set of generative mechanisms for $\mathcal{D}(\mathbf{C}, \Omega)$. If $\mathbf{B} \in \dot{\mathbb{P}}(\mathbb{L}(\mathbf{C}))$, $|\mathbf{B}| = |\mathbf{C}|$ and \mathbf{B} is irreducible for $\mathcal{D}(\mathbf{C}; \Omega)$ then there exists an individual ω^* and a mechanism M^i such that $M_{\mathbf{B}=1}^i(\omega^*) = 1$ but for all $L \in \mathbf{B}$, $M_{\mathbf{B} \setminus \{L\}=1, L=0}^i(\omega^*) = 0$.

Thus if there exists an exhaustive set of generative mechanisms for $\mathcal{D}(\mathbf{C}, \Omega)$ and \mathbf{B} is irreducible then there is an individual ω^* and a mechanism M^i such that M^i is active when all the literals in \mathbf{B} take the value 1, and is inactive when any one literal is 0 and the rest continue to take the value 1.

Proof: By Theorem 3.2, since \mathbf{B} is irreducible for $\mathcal{D}(\mathbf{C}; \Omega)$, there exists $\omega^* \in \Omega$ such that $D_{\mathbf{B}=1}(\omega^*) = 1$ and for all $L \in \mathbf{B}$, $D_{\mathbf{B} \setminus \{L\}=1, L=0}(\omega)$. Since \mathbf{M} is an exhaustive set of generative mechanisms for $\mathcal{D}(\mathbf{C}, \Omega)$, we have that for all \mathbf{c} , $D_{\mathbf{c}}(\omega^*) = \bigvee(\mathbf{M}_{\mathbf{c}}(\omega))$. Since $D_{\mathbf{B}=1}(\omega^*) = 1$, for some $M^i \in \mathbf{M}$, $M_{\mathbf{B}=1}^i(\omega^*) = 1$. Since for all $L \in \mathbf{B}$, $D_{\mathbf{B} \setminus \{L\}=1, L=0}(\omega^*) = 0$ we have that $M_{\mathbf{B} \setminus \{L\}=1, L=0}^i(\omega^*) = 0$.

PROPOSITION 7.5. Suppose \mathbf{M} forms an exhaustive set of generative mechanisms for $\mathcal{D}(\mathbf{C}, \Omega)$. If $\mathbf{B} \in \dot{\mathbb{P}}(\mathbb{L}(\mathbf{C}))$, $|\mathbf{B}| = |\mathbf{C}|$ and \mathbf{B} is singular for $\mathcal{D}(\mathbf{C}; \Omega)$ then there exists an individual ω^* and a mechanism M^i such that $M_{\mathbf{B}=\mathbf{b}}^i(\omega^*) = 1$ iff $\mathbf{b} = \mathbf{1}$.

Hence under the conditions of Proposition 7.5, if \mathbf{B} is singular then there is an individual ω^* and a mechanism M^i such that M^i is active if and only if all the literals in \mathbf{B} take the value 1.

Proof: Similar to the proof of Proposition 7.4, replacing Theorem 3.2 by Theorem 5.2.

As the next example shows, the assumption that there exists an exhaustive set of generative mechanisms is substantive, and does not hold in all cases.

EXAMPLE 6. Suppose $\mathbf{C} = \{X_1, X_2\}$ where X_1 and X_2 denote the presence of a variant allele at two particular loci. Let M^1 and M^2 denote two different proteins such that M^i is produced iff $X_i = 0$, i.e. the associated allele is not present. Finally let D denote some characteristic whose occurrence is blocked by the presence of either M^1 or M^2 (or both). In this example

$$\begin{aligned} M_{x_1x_2}^i(\omega) &= (1 - x_i) \\ D_{x_1x_2}(\omega) &= (1 - M_{x_1x_2}^1) \vee (1 - M_{x_1x_2}^2) = x_1x_2. \end{aligned}$$

By De Morgan's Law, the second equation here may also be expressed as:

$$1 - D_{x_1x_2}(\omega) = M_{x_1x_2}^1(\omega) M_{x_1x_2}^2(\omega) = 1 - x_1x_2.$$

The mechanisms M^1 and M^2 are preventive for D , so that $D = 1$ only occurs when both mechanisms are inactive. An exhaustive set of generative mechanisms does not exist because in this example there are no generative mechanisms.

It is natural to suppose that mechanisms are 'modular' and thus may be isolated or rendered inactive without affecting other such mechanisms. This may be formalized via potential outcomes:

DEFINITION 7.6. An exhaustive set of generative mechanisms \mathbf{M} for $\mathcal{D}(\mathbf{C}, \Omega)$ are said to support counterfactuals if there exist well-defined potential outcomes $D_{\mathbf{C}=\mathbf{c}, \mathbf{M}=\mathbf{m}}(\omega)$ and $D_{\mathbf{M}=\mathbf{m}}(\omega)$ such that

$$D_{\mathbf{C}=\mathbf{c}, \mathbf{M}=\mathbf{m}}(\omega) = D_{\mathbf{M}=\mathbf{m}}(\omega) = \left(\bigvee (\mathbf{M}) \right)_{\mathbf{m}}.$$

The important assumption here is the existence of the potential outcomes $D_{\mathbf{m}}(\omega)$ and $D_{\mathbf{c}, \mathbf{m}}(\omega)$. Note that if \mathbf{M} supports counterfactuals then interventions on \mathbf{C} do not affect D if interventions are also made on \mathbf{M} .

PROPOSITION 7.7. If the exhaustive set of generative mechanisms \mathbf{M} support counterfactuals then

$$D_{\mathbf{M}=\mathbf{M}(\omega)}(\omega) = D_{\mathbf{C}=\mathbf{C}(\omega), \mathbf{M}=\mathbf{M}(\omega)}(\omega) = D(\omega)$$

so that consistency holds for the potential outcomes $D_{\mathbf{m}}(\omega)$ and $D_{\mathbf{c}, \mathbf{m}}(\omega)$.

Proof: This follows because

$$D_{\mathbf{C}=\mathbf{C}(\omega), \mathbf{M}=\mathbf{M}(\omega)}(\omega) = D_{\mathbf{M}=\mathbf{M}(\omega)}(\omega) = \bigvee (\mathbf{M}(\omega)) = D(\omega).$$

7.3. Counterfactual interpretation of a minimal sufficient cause representation

If we have an exhaustive set of generative mechanisms which supports counterfactuals, and further each mechanism is a conjunction of literals, then there will be a minimal sufficient cause representation that itself supports counterfactuals.

DEFINITION 7.8. A representation $(\mathbf{A}, \mathfrak{B})$ for $\mathcal{D}(\mathbf{C}, \Omega)$ will be said to be structural if for each pair (A_i, \mathbf{B}_i) , $A_i \in \mathbf{A}$, $\mathbf{B}_i \in \mathfrak{B}$ there exists a generative mechanism (or mechanisms) M^i such that $M^i = A_i \wedge (\bigwedge (\mathbf{B}_i))$ and

$$M^i_{\mathbf{C}=\mathbf{c}}(\omega) = A_i(\omega) \wedge \left(\bigwedge (\mathbf{B}_i) \right)_{\mathbf{c}}.$$

Thus if $(\mathbf{A}, \mathfrak{B})$ is structural then each pair (A_i, \mathbf{B}_i) , $A_i \in \mathbf{A}$, $\mathbf{B}_i \in \mathfrak{B}$ corresponds to a mechanism M^i . Thus in this case the variables $A_i(\omega)$ may be interpreted as indicating whether the corresponding mechanism(s) M^i is ‘present’ in individual ω . We may thus associate potential outcomes with the A_i , corresponding to removing (or inserting) the corresponding mechanism(s). This interpretation of the A_i ’s is consistent with the notion of ‘co-cause’ which arises in the literature on minimal sufficient causes.

We note that ‘structural’ is often used as a synonym for ‘causal’. However, even under the weak interpretation, a sufficient cause representation is causal in that it represents a set of potential outcomes. The word is used in Definition 7.8 to connote that the *structure* of the representation itself represents (additional) potential outcomes arising from intervention on a set of mechanisms \mathbf{M} that correspond with the pairs (A_i, \mathbf{B}_i) , $A_i \in \mathbf{A}$, $\mathbf{B}_i \in \mathfrak{B}$. Note that there need not be a unique structural representation $(\mathbf{A}, \mathfrak{B})$ for $\mathcal{D}(\mathbf{C}, \Omega)$. There might be several functionally equivalent generative mechanisms corresponding to a given pair (A_i, \mathbf{B}_i) .

PROPOSITION 7.9. If a representation $(\mathbf{A}, \mathfrak{B})$ for $\mathcal{D}(\mathbf{C}, \Omega)$, where $\mathbf{A} = \langle A_1, \dots, A_p \rangle$, is structural then the associated set of generative mechanisms $\langle M^1, \dots, M^p \rangle$ is exhaustive.

Proof: This follows from Definitions 2.7 and 7.2. 2.

PROPOSITION 7.10. Suppose that \mathbf{M} forms an exhaustive set of generative mechanisms for $\mathcal{D}(\mathbf{C}, \Omega)$ and \mathbf{M} supports counterfactuals. If for all $M^i \in \mathbf{M}$ there exists $\mathbf{B}_i \in \dot{\mathbb{P}}(\mathbb{L}(\mathbf{C}))$, and an A_i such that for all \mathbf{c} , and $\omega \in \Omega$, if $A_i(\omega) = 1$ then $(M^i)_{\mathbf{c}}(\omega) = \bigwedge (\mathbf{B}_i)_{\mathbf{c}}$, then $(\mathbf{A} = \langle A_1, \dots, A_p \rangle, \mathfrak{B} = \langle \mathbf{B}_1, \dots, \mathbf{B}_p \rangle)$ forms a representation for $\mathcal{D}(\mathbf{C}, \Omega)$ that is structural.

Proof: This follows from Definitions 2.7 and 7.8.

PROPOSITION 7.11. If there is some representation $(\mathbf{A}, \mathfrak{B})$ that is structural and $\mathbf{B} \in \dot{\mathbb{P}}(\mathbb{L}(\mathbf{C}))$ is irreducible for $\mathcal{D}(\mathbf{C}; \Omega)$ then there exists a mechanism M_i that is active only if $\mathbf{B} = \mathbf{1}$.

Proof: If \mathbf{B} is irreducible for $\mathcal{D}(\mathbf{C};\Omega)$ then there exists $\mathbf{B}_i \in \mathfrak{B}$ with $\mathbf{B} \subseteq \mathbf{B}_i$; the mechanism $M^i = A_i \wedge (\Lambda(\mathbf{B}_i))$ is such that $M^i = 1$ only if $\mathbf{B} = \mathbf{1}$.

Note that the conclusion of Proposition 7.11, unlike those of Propositions 7.4 and 7.5, does not make reference to an individual ω^* . This is because Proposition 7.11 assumes that there is a representation $(\mathbf{A}, \mathfrak{B})$ that is structural: in this representation the A_i variables may be seen as a constituent part of the corresponding mechanism M_i .

Note that there may exist a set of exhaustive generative mechanisms, but these mechanisms may not themselves be conjunction of literals so that there is no sufficient cause representation for $\mathcal{D}(\mathbf{C};\Omega)$ that is structural:

EXAMPLE 7. Suppose $\mathbf{C} = \{X_1, X_2\}$ where X_1 and X_2 again denote the presence of variant alleles, acquired by maternal and paternal inheritance respectively, at a particular locus. Let M denote a protein that is produced iff either $X_1 = X_2 = 1$ or $X_1 = X_2 = 0$ and let D denote some characteristic that occurs iff $M = 1$. Suppose we can intervene to remove or add the protein. We then have that

$$\begin{aligned} M_{x_1x_2}(\omega) &= x_1x_2 \vee (1-x_1)(1-x_2) \\ D_{x_1x_2}(\omega) &= M_{x_1x_2}(\omega) \\ D_{x_1x_2m}(\omega) &= D_m(\omega) = m. \end{aligned}$$

Thus $\{M\}$ constitutes an exhaustive set of generative mechanisms for $\mathcal{D}(\mathbf{C}, \Omega)$. We have a representation for $\mathcal{D}(\mathbf{C}, \Omega)$:

$$D_{x_1x_2}(\omega) = (A_1(\omega) X_1 X_2 \vee A_2(\omega) (1-X_1)(1-X_2))_{x_1x_2}$$

with $A_1(\omega) = A_2(\omega) = 1$ for all $\omega \in \Omega$. However, this representation is not structural because $A_1(\omega)X_1X_2$ and $A_2(\omega)(1-X_1)(1-X_2)$ do not constitute separate mechanisms for which interventions are conceivable; there is only one mechanism M , the protein. Since for any $\omega \in \Omega$, $D_{11}(\omega) = 1$ and $D_{10}(\omega) = D_{01}(\omega) = 0$, $\{X_1, X_2\}$ is irreducible relative to \mathbf{C} ; however it is not the case that there is a mechanism M_i that is active only if $X_1X_2 = 1$ since for the only mechanism M it is the case that $M = 1$ if $X_1 = X_2 = 0$. Note, however, in this example there is still a mechanism, namely M , that will be ‘active’ if $X_1 = X_2 = 1$ but will be ‘inactive’ if only one of X_1 or X_2 is 1.

EXAMPLE 8. To illustrate the results in the paper we consider again the data presented in Table 1. We let D denote bladder cancer, X_1 smoking, X_2 the S NAT2 genotype, and X_3 the *10 allele on NAT1. As discussed in Example 5, if the effect of $C = \{X_1, X_2, X_3\}$ is unconfounded for D and we fit the model

$$E[D=1|X_1=x_1, X_2=x_2, X_3=x_3] = \beta_0 + \beta_1x_1 + \beta_2x_2 + \beta_3x_3 + \beta_{12}x_1x_2 + \beta_{13}x_1x_3 + \beta_{23}x_2x_3 + \beta_{123}x_1x_2x_3 \quad (7.2)$$

then if X_1, X_2 and X_3 have positive monotonic effects on D (relative to C) then $\{X_1, X_2, X_3\}$ is irreducible relative to C if any of the following hold:

$$\beta_{123} > \beta_1, \quad \beta_{123} > \beta_2, \quad \beta_{123} > \beta_3.$$

We cannot fit the model (7.2) directly with case control data. However, under the assumption that the outcome is rare (reasonable with bladder cancer) so that odds ratios approximate risk ratios, we can fit the model:

$$\begin{aligned} E[D=1|X_1=x_1, X_2 \\ =x_2, X_3 \\ =x_3] &= E[D \\ =1|X_1 \\ =0, X_2 \\ =0, X_3 \\ =0] &= \theta_1 x_1 + \theta_2 x_2 + \theta_3 x_3 + \theta_{12} x_1 x_2 \\ &+ \theta_{13} x_1 x_3 \\ &+ \theta_{23} x_2 x_3 \\ &+ \theta_{123} x_1 x_2 x_3 \end{aligned} \quad (7.3)$$

and the conditions for the irreducibility of $\{X_1, X_2, X_3\}$ relative to C under monotonicity of $\{X_1, X_2, X_3\}$ become:

$$\theta_{123} > \theta_1, \quad \theta_{123} > \theta_2, \quad \theta_{123} > \theta_3.$$

If we fit model (7.3) using maximum likelihood, we find that

$$\begin{aligned} \theta_{123} - \theta_1 &= 1.21 \quad (9.5\% \quad CI: -3.83, 6.26), \\ \theta_{123} - \theta_2 &= 2.93 \quad (9.5\% \quad CI: -2.85, 8.72), \\ \theta_{123} - \theta_3 &= 2.97 \quad (9.5\% \quad CI: -2.80, 8.74). \end{aligned}$$

In each case the point estimate suggests evidence of irreducibility, under monotonicity of $\{X_1, X_2, X_3\}$, but the sample size is not sufficiently large to draw this conclusion confidently. With monotonicity of $\{X_1, X_2, X_3\}$, irreducibility also implies a singular interaction for $\{X_1, X_2, X_3\}$. If we assume that only $\{X_1, X_2\}$ or $\{X_1, X_3\}$ or $\{X_2, X_3\}$ are monotonic relative to C then the conditions for irreducibility in Example 5 can be expressed respectively as:

$$\theta_{123} > 1 + \theta_1 + \theta_2, \quad \theta_{123} > 1 + \theta_1 + \theta_3, \quad \theta_{123} > 1 + \theta_2 + \theta_3.$$

From model (7.3) we have that

$$\begin{aligned} \theta_{123} - (1 + \theta_1 + \theta_2) &= 0.09 \quad (9.5\% \quad CI: -4.77, 4.96), \\ \theta_{123} - (1 + \theta_1 + \theta_3) &= 0.13 \quad (9.5\% \quad CI: -4.69, 4.95), \\ \theta_{123} - (1 + \theta_2 + \theta_3) &= 1.86 \quad (9.5\% \quad CI: -3.41, 7.12). \end{aligned}$$

Again, in each case the point estimate suggests evidence of irreducibility, under monotonicity of just two of the three exposures, but the sample size is not sufficiently large to draw this conclusion confidently. With monotonicity of two of the three exposures, irreducibility also implies a singular interaction for $\{X_1, X_2, X_3\}$. The test for irreducibility in Example 5 without assumptions about monotonicity can be expressed as:

$$\theta_{123} > 2 + \theta_1 + \theta_2 + \theta_3.$$

From model (7.3) we have that

$$\theta_{123} - (2 + \theta_1 + \theta_2 + \theta_3) = -0.99 \quad (95\% \text{ CI: } -5.86, 3.88).$$

In this case, not even the point estimate is positive.

If $\{X_1, X_2, X_3\}$ is in fact irreducible and if there exists a representation that is structural then it follows by Proportion 7.11 that there exists a mechanism that is active only if $X_1 = 1, X_2 = 1, X_3 = 1$.

8. Concluding Remarks

In this paper we have developed general results for notions of interaction that we referred to as ‘irreducibility’ (aka ‘a sufficient cause interaction’) and ‘singularity’ for sufficient cause models with an arbitrary number of dichotomous causes. The theory could be extended by developing notions of sufficient cause, irreducibility and singularity for causes and outcomes that are categorical and/or ordinal in nature.

Acknowledgments

We thank Stephen Stigler for pointing us to earlier work by Cayley on minimal sufficient cause models. We thank James Robins for helpful conversations. Tyler VanderWeele was supported by the National Institutes of Health (R01 ES017876). Thomas Richardson was supported by the National Science Foundation (CRI 0855230), the National Institutes of Health (R01 AI032475), and The Institute of Advanced Studies, University of Bologna.

References

- [1]. Aickin, M. Causal Analysis in Biomedicine and Epidemiology Based on Minimal Sufficient Causation. Marcel Dekker; New York: 2002.
- [2]. Bateson, W. Mendel’s Principles of Heredity. Cambridge University Press;
- [3]. Cayley A. Note on a question in the theory of probabilities. London, Edinburgh and Dublin Philosophical Magazine. 1853; VI:259.
- [4]. Cayley A. A theorem on trees. Quart. J. Math. 1889; 23:376–378.
- [5]. Cordell H. Epistasis: what it means, what it doesn’t mean, and statistical methods to detect it in humans. Hum. Mol. Genet. 2002; 11:2463–2468. [PubMed: 12351582]
- [6]. Cox, DR. The Planning of Experiments. Wiley; New York: 1958.
- [7]. Flanders D. Sufficient-component cause and potential outcome models. Eur. J. Epidemiol. 2006; 21:847–853. [PubMed: 17048084]
- [8]. Greenland S, Brumback B. An overview of relations among causal modelling methods. Int. J. Epidemiol. 2002; 31:1030–1037. [PubMed: 12435780]
- [9]. Greenland S, Poole C. Invariants and noninvariants in the concept of interdependent effects. Scand. J. Work Environ. Health. 1988; 14:125–129. [PubMed: 3387960]

- [10]. Koopman JS. Interaction between discrete causes. *Am. J. Epidemiol.* 1981; 113:716–724. [PubMed: 7234861]
- [11]. Mackie JL. Causes and conditions. *American Philosophical Quarterly.* 1965; 2:245–255.
- [12]. Marcovitz, AB. *Introduction to Logic Design.* McGraw-Hill; 2001.
- [13]. McCluskey EJ. Minimization of Boolean functions. *The Bell System Technical Journal.* 1956; 35:1417–1444.
- [14]. Neyman, J. Sur les applications de la thar des probabilités aux expériences agaricales: Essay des principe. In: Dabrowska, D.; Speed, T., editors. *Statistical Science* 5. 1923. p. 463-472.Excerpts in English
- [15]. Novick LR, Cheng PW. Assessing interactive causal influence. *Psychological Review.* 2004; 111:455–485. [PubMed: 15065918]
- [16]. Pearl, J. *Causality: Models, Reasoning, and Inference.* Cambridge University Press; Cambridge: 2000.
- [17]. Phillips P. Epistasis – the essential role of gene interactions in the structure and evolution of genetic systems. *Nature Reviews Genetics.* 2008; 9:855–867.
- [18]. Quine WV. The problem of simplifying truth functions. *American Math. Monthly.* 1952; 59:521–531.
- [19]. Quine WV. A way to simplify truth functions. *American Math. Monthly.* 1955; 62:627–631. Robins JM. A new approach to causal inference in mortality studies with sustained exposure period - application to control of the healthy worker survivor effect. *Mathematical Modelling.* 1986; 7:1393–1512.
- [20]. Robins, JM. Marginal structural models versus structural nested models as tools for causal inference. In: Halloran, E.; Berry, D., editors. *Statistical Models in Epidemiology: the Environment and Clinical Trials.* Springer; New York: 1999. p. 95-134. Rosenbaum PR, Rubin DB. The central role of the propensity score in observational studies for causal effects. *Biometrika.* 1983; 70:41–55.
- [21]. Rothman KJ. Causes. *Am. J. Epidemiol.* 1976; 104:587–592. [PubMed: 998606]
- [22]. Rothman, KJ.; Greenland, S. *Modern Epidemiology.* Lippincott-Raven; Philadelphia: 1998.
- [23]. Rubin DB. Estimating causal effects of treatments in randomized and nonrandomized studies. *J. Educ. Psychol.* 1974; 66:688–701.
- [24]. Rubin DB. Bayesian inference for causal effects: The role of randomization. *Ann. Statist.* 1978; 6:34–58.
- [25]. Rubin DB. Comment on neyman 1923 and causal inference in experiments and observational studies. *Statistical Science.* 1990; 5:472–480.
- [26]. Taylor JA, Umbach DM, Stephens E, Castranio T, Paulson D, Robertson G, Mohler JL, Bell DA. The role of N-acetylation polymorphisms in smoking-associated bladder cancer: Evidence of a gene-gene-exposure three-way interaction. *Cancer Research.* 1998; 58:3603–3610. [PubMed: 9721868]
- [27]. VanderWeele T. Empirical tests for compositional epistasis. *Nature Reviews Genetics.* 2010; 11:166.
- [28]. VanderWeele T. Epistatic interactions. *Statistical Applications in Genetics and Molecular Biology.* 2010; 9:1–22.
- [29]. VanderWeele TJ, Hernán MA. From counterfactuals to sufficient component causes, and vice versa. *Eur. J. Epidemiol.* 2006; 21:855–858. [PubMed: 17225959]
- [30]. VanderWeele TJ, Robins JM. Directed acyclic graphs, sufficient causes and the properties of conditioning on a common effect. *Am. J. Epidemiol.* 2007; 166:1096–1104. [PubMed: 17702973]
- [31]. VanderWeele TJ, Robins JM. The identification of synergism in the sufficient-component cause framework. *Epidemiol.* 2007; 18:329–339.
- [32]. VanderWeele TJ, Robins JM. Empirical and counterfactual conditions for sufficient cause interactions. *Biometrika.* 2008; 95:49–61.
- [33]. VanderWeele TJ, Robins JM. Signed directed acyclic graphs for causal inference. *J. Roy. Statist. Soc., Series B.* 2010; 72:111–127.

- [34]. Vansteelandt S, Goetghebeur E. Causal inference with generalized structural mean models. *J. Roy. Statist. Soc., Series B.* 2003; 65:817–835.




Fig 1.
The three trees on $\{X_1, X_2, X_3\}$.

Table 1

Smoking	NAT2	NAT1*10	Cases (<i>n</i> = 215)	Controls (<i>n</i> = 191)	Odds Ratio (95% CI)	
0	0	0	6	13	1	
0	0	1	8	16	1.1	(0.3, 3.9)
0	1	0	16	31	1.1	(0.4, 3.5)
0	1	1	6	10	1.3	(0.3, 5.3)
1	0	0	42	32	2.8	(1.1, 8.3)
1	0	1	41	26	3.4	(1.2, 10.1)
1	1	0	61	51	2.6	(0.9, 7.3)
1	1	1	35	12	6.3	(2.0, 20.3)

All potential outcomes and actual outcomes for three binary causes X_1 , X_2 and X_3 , in a population with two individuals.

Table 2

Individual	D_{000}	D_{001}	D_{010}	D_{011}	D_{100}	D_{101}	D_{110}	D_{111}	D	(X_1, X_2, X_3)
1	0	1	1	0	0	1	1	0	1	(1, 0, 1)
2	0	1	1	0	0	1	1	1	0	(0, 0, 0)