Interactive Causes: Revising the Markov Condition

Gerhard Schurz*†

This article suggests a revision of the theory of causal nets (TCN). In section 1 we introduce an axiomatization of TCN based on a realistic understanding. It is shown that the causal Markov condition entails three independent principles. In section 2 we analyze *indeterministic decay* as the major counterexample to one of these principles: screening off by common causes (SCC). We call SCC-violating common causes *interactive* causes. In section 3 we develop a revised version of TCN, called TCN*, which accounts for interactive causes. It is shown that there are interactive causal models that admit of no faithful noninteractive reconstruction.

1. Introduction: The Received Theory of Causal Nets. One of the best developed contemporary accounts of causality is the theory of causal nets (TCN). While most earlier accounts attempted to explicate causality by means of definitions, the TCN account characterizes causality via axiomatic principles that explain probabilistic relations of (in)dependence between variables by the existence of directed causal relations between them. This realistic understanding of TCN is informally acknowledged by Spirtes, Glymour, and Scheines (1993, secs. 3.4–5) and Pearl (2009, preface, viii–ix), although both Pearl and Spirtes et al. present TCN's axioms as definitions (Spirtes et al. speak of "axioms" on pp. 4 and 29). An explicitly axiomatic reconstruction of TCN is given in Schurz and Gebharter (2016).

This section gives a brief exposition of the notions and axioms of TCN. A (mathematical) variable is a measurable function X: D → Val(X) from a

Received November 2015; revised August 2016.

*To contact the author, please write to: Department of Philosophy, Heinrich Heine University Duesseldorf, Universitaetsstrasse 1, Geb. 24.5, 40225 Duesseldorf, Germany; e-mail: schurz@phil.uni-duesseldorf.de.

†This work was supported by the Deutsche Forschungsgemeinschaft, research unit FOR 1063. For valuable help I am indebted to Christopher Hitchcock, Paul Näger, Jerry Finkelstein, and an unknown referee.

Philosophy of Science, 84 (July 2017) pp. 456–479. 0031-8248/2017/8403-0004$10.00
Copyright 2017 by the Philosophy of Science Association. All rights reserved.

https://doi.org/10.1086/692143 Published online by Cambridge University Press

domain D of individuals to its value space $Val(X) = \{x_1, x_2, ...\}$, which is a family of properties or a set of numbers. If X denotes color, for example, then $Val(X) = \{red, green, ...\}$, and X assigns a color $X(\alpha)$ to every individual $\alpha \in D$. The domain may also consist of *n*-tuples of individuals, for example, individuals at certain time points. Binary properties are represented by binary variables X_F with value space $\{F,\neg F\}$. Terminological conventions:

- X,Y,Z (possibly indexed) are variables, and U,V,W are sets of variables;
- lowercase x's (or x_i) stand for values of X, and lowercase u's (or u_i), for sets of values of the respective variables in U.

A *causal graph* is a pair (V,E), where $V = \{X_1, X_2, ...\}$ is a set of variables (the "vertices") and $E \subseteq V \times V$ a set of directed arrows $X_i \rightarrow X_j$ (the "edges").[1] A graph (V,E) together with a probability distribution P is called a *causal model* (CM) (V,E,P). When we use causal graphs or models to represent reality, we speak of causal structures or systems, respectively.

The symbol P denotes an objective probability distribution over a suitable algebra AL over the value space over V; that is, $P: AL(\prod_{i\in I} Val(X_i)) \rightarrow [0, 1]$ (with corresponding projections to subspaces). Further terminological conventions:

- P(x) (or P({x})) abbreviates $P(X(\alpha) = x)$, that is, the probability that X takes value x in the underlying domain D (the individual variable α is bound by P).
- Likewise, P(S) stands for $P(X(\alpha) \in S)$, $P(\neg x)$ for $P(X(\alpha) \notin x)$, P(x,y) for $P(X(\alpha) = x \wedge Y(\alpha) = y)$, and P(x|y)for $P(X(\alpha) = x|Y(\alpha) = y)$, that is, the conditional probability of x given y, provided P(y) > 0.

Two variables X,Y are probabilistically dependent (DEP(X,Y)) if and only if (iff) some of their values are dependent; otherwise X, Y are probabilistically independent. More generally, probabilistic (in)dependence between X and Y conditional on (fixed values of) a set of variables U is defined as follows:

- $DEP(X, Y|U)$ iff $\exists$ x, y, u: $P(x|y, u) \neq P(x|u)$ and $P(y, u) > 0$, and
- $INDEP(X, Y|U)$ iff not $DEP(X, Y|U)$, that is, iff $\forall x,y,u$: $P(x|y, u) = P(x|u)$ or $P(y, u) = 0$.

1. The variables in V are ordered by indices $X_1, X_2, ...$ in order to define the Cartesian product of their value space, $\prod_{i\in I} Val(X_i)$; these indices have no substantial meaning. When applying the probability operator P to sets of variables, we assume the variables to be ordered according to their indices.

https://doi.org/10.1086/692143 Published online by Cambridge University Press

If the variables' values are ordered according to size, the notion of *positive/negative* dependence between variables can be defined as follows: X depends positively/negatively on Y iff $P(x_1 > x_2|y_1 > y_2) >/< P(x_1 > x_2)$ (for all x_1, x_2, y_1, y_2).

A causal graph (V,E) is described by using the following notions:

> $X \rightarrow Y$: X is a *direct cause* of Y (Y is a *direct effect* of X).
>
> $X \rightarrow\!\rightarrow Y$: X is a *causal ancestor* of Y (Y is a *causal successor* of X); that is, there is a directed path $X \rightarrow Z_1 \rightarrow ... \rightarrow Z_n \rightarrow Y$ from X to Y (for $n \geq 0$).[2]
>
> $X - Y$: $X \rightarrow Y$ or $X \leftarrow Y$; that is, X and Y are *adjacent*.
>
> $X_1 - - X_n$: a *path* $X_1 - ... - X_n$ between X_1 and X_n that *connects* X_1 and X_n; the variables X_i ($1 \leq i \leq n$) are said to *lie* on this path.

If X_i lies on a path π, then X_i is called (i) a common cause, (ii) an intermediate cause, or (iii) a common effect on π iff (i) $\leftarrow X \rightarrow$, (ii) $\leftarrow X \leftarrow$ or $\rightarrow X \rightarrow$, or (iii) $\rightarrow X \leftarrow$, respectively, is part of π.

The set of all direct causes of a variable X is called the set of its parents, Par(X). Note that TCN is compatible with ontological *indeterminism*; that is, the values of Par(X) need not determine the value of X. *Determinism* in regard to variable X is given when for all possible values $par(X) \in Val(Par(X))$ there exists an $x \in Val(X)$ with $P(x|Par(X)) = 1$.

TCN explains probabilistic dependencies via causal connections by means of two core axioms: causal d-connection (C) and minimality (Min). Variable C is frequently expressed by the equivalent causal Markov condition (M) (Spirtes, Glymour, and Scheines 2000, sec. 3.4.1). We prefer C over M for reasons specified below. The axiom C states that every (conditional) probabilistic dependence between two variables X and Y is the result of a causal "d-connection" between X and Y:

> (1) *Definition*: A CM (V,E,P) satisfies the condition C of *causal d-connection* iff for all $X, Y \in V$ and $U \subseteq V - \{X, Y\}$: if DEP(X, Y|U), then X and Y are d-connected given U in the following sense: X and Y are connected by some path π such that
> (SIC): no intermediate cause on π is in U,
> (SCC): no common cause on π is in U, and
> (CE): every common effect on π is in U or has a causal successor in U.

2. The notion of "causal ancestor" does not imply that causation is transitive (but faithful causation is transitive; see below).

https://doi.org/10.1086/692143 Published online by Cambridge University Press

> *Axiom*: Every physically possible causal system that is [sufficiently] complete satisfies condition C.

In (1) we distinguish between the definition of the causal d-connection condition and the corresponding axiom, which states that this condition holds for all physically possible systems that are either complete or at least sufficiently complete (as indicated by the square brackets); these two notions of completeness are characterized below.

Two variables X, Y are "d-separated" given U if they are not d-connected given U. The concepts of d-separation and d-connection were developed by Pearl (1988, 117). Axiom C asserts an *implication* from (probabilistic) dependence to d-connection, or in contraposed form, from d-separation to independence (Pearl 1988, 119; 2000, theorem 1.2.4). Axiom C implies the standard causal principle for *unconditional* dependence: if DEP(X,Y), then X and Y are connected by a directed or common cause path (i.e., d-connected by the empty set $\varnothing$).

Our formulation in (1) makes it plain that condition C entails three independent principles:

> SIC (screening off by intermediate causes): Direct causes screen off indirect causes from their effects. Thus, in a [sufficiently] complete CM of the form $X \rightarrow Z \rightarrow Y$, INDEP(X, Y|Z) holds. Condition SIC goes back to Andrei Markov (senior).

> SCC (screening off by common causes): Common causes screen off their effects from each other. So in a [sufficiently] complete CM of the form $X \leftarrow Z \rightarrow Y$, INDEP(X, Y|Z) holds. Condition SCC goes back to Reichenbach (1956).

> CE (common effects): The causes of a common effect are unconditionally independent but may be dependent conditional on a common causal successor. Thus, in a [sufficiently] complete CM of the form $X \rightarrow Z \leftarrow Y$, INDEP(X, Y|Z) holds, and DEP(X, Y|Z) will hold if the CM is faithful (see definition (3)). Condition CE was discovered by Berkson (1946) and is called Berkson's paradox.

The full content of these three principles is expressed as follows: "If DEP(X, Y|U), then there is an X-Y-connecting path π such that (P)," where P is one of the clauses SIC, SCC, or CE in (1), respectively.[3] The modularization of axiom C into three entailed principles is important for our pur-

3. Axiom C is stronger than the conjunction of these three principles, because in C the clauses SIC, SCC, and CE are in the scope of the same existential quantifier over X-Y-connecting paths π.

https://doi.org/10.1086/692143 Published online by Cambridge University Press

pose of revision because it allows us to alter one part of C without changing the others.

A causal graph (model) is *acyclic* iff it contains no cyclic path $X \rightarrow\rightarrow X$. In this article we concentrate on acyclic graphs. For acyclic CMs, condition C is equivalent to the causal Markov condition M (Spirtes et al. 2000, theorem 3.3; Pearl 2009, 16):

> (2) *Definition*: (V,E,P) satisfies the *causal Markov condition* M iff every $X \in V$ is independent of any subset $U \subseteq V - \{X\}$ of nonsuccessors of X, conditional on X's parents: INDEP(X, U|Par(X)).[4]

The equivalence between M and C has been demonstrated by Verma and Pearl (see Pearl 1988, 119–20, theorem 9) and Lauritzen et al. (1990). The latter authors call C the *global* and M the *local* Markov condition, since in its contraposed form, C asserts a (conditional) independence for all d-separation relations of a causal graph, while M asserts such an independence only for the d-separation relations between a variable X and its noneffects conditional on its parents; the other independencies follow from these as probabilistic consequences. For this reason, C expresses the full content of the causal Markov condition in a more transparent way than M. Moreover, the three principles SIC, SCC, and CE are more directly contained in axiom C than in condition M.

The explicit distinction between the definition of C and the axiom C in (1) impels us to critically reflect on the *problem of generality*: do all probabilistic dependencies between analytically independent variables in fact result from causal d-connections? This has been a question of debate for several decades. Two remarks on this debate are important for our purpose: First, although the criticism was often directed against the Markov condition "in general," our reconstruction of axiom C as entailing three independent conditions helps us to locate the problem: what was under attack was not condition SIC or CE but only SCC—screening off by common causes. Second, defenders of TCN usually acknowledge that SCC fails in application to far-distance (EPR/Bell) correlations in quantum mechanics, but they regard SCC as valid for all other domains (Spirtes et al. 2000, 37–38; Pearl 2009, 62). Only a few authors, foremost Cartwright (e.g., Cartwright 2007), have insisted that the causal Markov condition may fail for macrophysical systems as well. Section 2 will support this position, but since the counterexamples only concern subcondition SCC of C, we argue in section 3 that the Markov condition should not be abandoned but should be revised in certain ways.

4. Condition M is explicated in terms of dependence between sets of variables. The latter notion reduces to dependence between variables as follows: DEP(U, V|W) iff DEP(X, Y|W, $U - \{X\}$, $V - \{Y\}$) for some $X \in U$ and $Y \in V$. Similarly, the notion of d-connection can be extended to sets of variables.

https://doi.org/10.1086/692143 Published online by Cambridge University Press

Axiom C asserts an implication from probabilistic dependence to d-connection. The converse implication from d-connection to probabilistic dependence is called the condition of *faithfulness* (Spirtes et al. 2000, 31; Zhang and Spirtes 2008, 24):

> (3) *Definition*: A CM (V,E,P) satisfies the *faithfulness condition* F iff for all X, $Y \in V$ and $U \subseteq V - \{X, Y\}$: if X and Y are d-connected given U, then DEP(X, Y|U).

To avoid confusion, note that Spirtes et al. (2000, 31) define "faithfulness" as the conjunction of conditions C and F. We prefer definition (3) (which is also used in Zhang and Spirtes 2008, 247) because it logically separates F from C. This is essential for our revision, which does not concern F but only one part of C (i.e., SCC).

The probability distributions over the variables X_i of a CM conditional on their parents, $P(X_i|Par(X_i))$, are called the *parameters* of the CM (Pearl 2009, 44). It is well known that the faithfulness condition may fail for CMs whose parameters are fine-tuned in a such way that the influences of two (or more) d-connecting paths are opposite and cancel each other out. An example is presented in figure 1.

Therefore, condition F is not suited as a general axiom of TCN. Proponents of TCN argue that the satisfaction of F is at least highly probable because unfaithful CMs are *parameter instable*, in the sense that their unfaithful independencies can be destroyed by arbitrary small changes of the parameters of the CM (Pearl 2009, def. 2.4.1). This fact implies that unfaithful independencies are statistically rare in all causal systems that are subject to external noise variables that are mutually independent and, thus, produce independent variations of the system's parameters.

However, there is a weakening of F that holds in full generality and constitutes TCN's second core axiom—the condition of minimality (Spirtes et al. 2000, sec. 3.4.2):

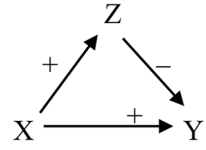


Figure 1. Unfaithful causal graph: DEP(X, Y|Z) but INDEP(X,Y), because the positive influence $X \rightarrow^{+} Y$ and the negative influence $X \rightarrow^{+} Z \rightarrow^{-} Y$ cancel to 0. Here $X \rightarrow^{+/-} Y$ means that the arrow $X \rightarrow Y$ is responsible for a positive/negative dependence of Y on X (as explained), conditional on $par(Y) - \{X\}$. See Schurz and Gebharter (2016), (9), theorem 2.

https://doi.org/10.1086/692143 Published online by Cambridge University Press

(4) *Definition*: A C-satisfying CM (V,E,P) satisfies the *minimality condition* (Min) iff no arrow can be omitted from E without violating condition C.

Axiom: Every physically possible C-satisfying causal system satisfies Min.

In contrast to axiom C, axiom Min can be justified by a methodological requirement that is a variant of Ockham's razor: causal arrows are only assumed if they are responsible for at least some (conditional) probabilistic dependence, for values that either are factually realized or can be brought about by interventions.[5]

Schurz and Gebharter (2016, sec. 3) investigate whether the theory TCN has *empirical content* by which it can be independently tested. Since causal arrows are regarded as theoretical (unobservable) relations, the TCN would have empirical content if it excluded logically possible probability distributions, that is, if there were probability models (V,P) that cannot be extended to TCN-satisfying CMs. The authors prove that the core of TCN, axiomatized by C + Min, is without empirical content. If an axiom of cautious faithfulness is added to C + Min—"cautious" in the sense that it asserts faithfulness merely with high probability—this extended version of TCN acquires weak probabilistic content. More empirical content is obtained if one adds to TCN the axiom of *temporal forward-directedness* T. This axiom applies to CMs whose variables are event variables, that is, variables X_i whose values are possible events having a location in time, $t(X_i)$. For these CMs condition T states that (for all $X, Y \in V$) $X \rightarrow Y$ implies $t(X) < t(Y)$; that is, causes precede their effects in time. The extension of C + F by T excludes a variety of logically possible probability distributions (2016, sec. 3.3).

Several objections against the truth of the causal Markov condition have been successfully defeated by proponents of TCN (see Spirtes et al. 2000, 59–63; Pearl 2009, 62; Hitchcock 2010). In particular, it has been shown that C can only be true if

a) the variables in V are analytically independent; otherwise one obtains "arrows" on purely conceptual grounds (Hitchcock 2010, secs. 2.3, 2.10), and
b) the description of the causal system is sufficiently *refined* (cf. Haussman and Woodward 1999, 527–29).

5. To ensure the compatibility of Min with the interventionist explication of a "direct cause" (see Zhang and Spirtes 2011) we assume that values that can be brought about by intervention have a positive probability.

https://doi.org/10.1086/692143 Published online by Cambridge University Press

Since the omission of certain variables or edges from a C-satisfying CM can lead to a violation of C, axiom C refers prima facie to physically possible CMs that are *complete* in the following sense:

> (5) *Definition*: (V,E,P) is a complete physically possible CM iff (V,E,P) is true in a physically possible world w and every variable in w that is a direct cause or effect of a variable in V is itself in V and its attached edges are in E.

Unfortunately, complete causal systems are typically not epistemically accessible. What does condition C entail for coarse-grained descriptions of causal systems that omit some detail and, thus, are incomplete?

A *structural coarsening* of a CM removes some variables and their edges but preserves cause-effect relations between the remaining variables. It is well known that the satisfaction of condition C is not preserved by every structural coarsening of a CM but only by structural coarsenings that are causally *sufficient*, that is, do not remove a common cause (Spirtes et al. 2000, 22, 27). However, there is a second kind of coarsening that plays a role in the next section and is called *conceptual coarsening*. Here a CM is described with the help of a compound variable Z that is definable in terms of more fine-grained variables $X_1, \ldots, X_n$; we write $Z = f(X_1, ..., X_n)$. When a model is conceptually coarsened by a compound variable Z, then the set of defining variables $X_1, \ldots, X_n$ is replaced in V by the compound variable Z and an arrow between a variable Y and Z is added whenever the original model contained a corresponding arrow between Y and some defining variable of Z. This operation is illustrated in figure 2.

Assume that in the refined model M_1 of figure 2 axiom C is violated because $DEP(X_1, X_3|X_2)$ holds (since X_2 is an interactive cause; see sec. 2). Nevertheless $INDEP(Y_1, Y_2|X_1, X_2, X_3)$ holds in M_1; thus, by assuming f to be injective, $INDEP(Y_1, Y_2|Z)$ will hold in M_2 (since each value triple $[x_1,x_2,x_3]$ is uniquely translated into one Z-value $z = f[x_1,x_2,x_3]$). So M_2 satisfies C, although its conceptual refinement M_1 fails to satisfy C. This shows that a conceptual coarsening of a CM may hide a violation of C.

To sum up, the satisfaction of C by a complete CM M* does not imply that a coarsening of it will satisfy C, nor does the violation of C by M* imply that a coarsening of M* will violate C. What can be guaranteed is that a CM is "sufficiently complete," in the sense of not distorting the true laws of causality, if its variables describe reality at the most fine-grained level and do not omit true common causes. Turning this into a definition would not be

$$G_1: \; Y_1 \leftarrow X_1 \leftarrow X_2 \rightarrow X_3 \rightarrow Y_2 \qquad G_2: \; Y_1 \leftarrow Z \rightarrow Y_2$$

Figure 2. Coarsening of causal graph G_1 by the compound variable $Z = f(X_1, X_2, X_3)$.

https://doi.org/10.1086/692143 Published online by Cambridge University Press

very useful, however, since most descriptions of reality are coarse-grained. All that can be said of coarse-grained descriptions is this:

> (6) If C is satisfied by every complete causal system of a world w, then every coarse-grained CM that is true in w has a true refinement in w that satisfies C.

Although condition 6 is analytically true, it gives us a heuristic for testing C:

> (7) *Test heuristic for C*: If a true causal model M is found such that not only M but every true refinement of M that is consistent with the established background knowledge violates C, then the truth of C in the actual world is strongly disconfirmed. On the other hand, if all CMs considered so far satisfy C for all plausible refinements of them, then C is confirmed.

In the next section we will make use of this heuristic.

2. Interactive Causes: Violation of SCC in Indeterministic Decay Processes. Given that axiom C is empirically empty in isolation, how can the truth of C be tested by empirical means at all? Following from the results concerning the empirical content of TCN mentioned in section 1, there are two ways of testing axiom C: either one strengthens C by condition F or by condition T. It is natural to assume condition T in the counterexample to which we now turn, namely, the example of spontaneous decay. However, our arguments do not presuppose this condition: we will see that even without assuming T certain causal structures involving spontaneous decay violate axiom C with high plausibility, because they cannot satisfy C in a faithful way.

The example of spontaneous decay has been given in the literature several times, but it was often presented as if it were a specific aspect of quantum mechanics (cf. van Fraassen 1980, 29–30; Arntzenius 2010, sec. 2.1). In contrast, Cartwright (2007, 122) has emphasized that the counterexample only requires an indeterministic decay process. Indeterministic processes are found not only in quantum physics but also in classical physics.[6]

Consider a particle (of type) q that decays with a certain probability p into two particles (of types) q_1 and q_2. The decay is objectively indetermined; that is, the probability of q's decay conditional on all common causes in the past is (significantly) smaller than 1; however, if the particle decays, then it always (or almost always) decays into q_1 and q_2. Or at least: If q decays and one decay product is q_1, then the other decay product is (almost) always q_2. Examples are radioactive decay (e.g., the α-decay of uranium $U^{238} \rightarrow Th^{234} + He^4$),

6. Certain classical differential equations admit of trajectory branching (e.g., Earman 1986, chap. 3; Norton 2008).

https://doi.org/10.1086/692143 Published online by Cambridge University Press

spontaneous chemical decay (e.g., the decay of ethene $C_2H_4 \rightarrow CH_4 + C$), or in everyday life the spontaneous breaking of an eroded mountain rock into two pieces, and so on.

What leads to a violation of SCC is the combination of two factors, (i) indeterminism plus (ii) the law describing a decay. Physically speaking the latter is an instance of a *conservation* law. Generally speaking, a conservation law for a quantity X says that the sum of the X-values of all parts of a closed system remains constant through time. Classical macrophysical or chemical decay is described by the conservation of *mass* (m): if q decays into q_1 and q_2, then $m(q) = m(q_1) + m(q_2)$. This explains *qualitative* decay laws as follows: if q decays into two parts and one part is q_1, then the other one must be q_2, because q_2 is the only physically possible decay component for that type of decay that conserves mass, via $m(q_2) = m(q) - m(q_1)$. More general conservation laws of physics, such as the conservation of energy, momentum, and charge, give rise to a more general notion of indeterministic decay that is introduced at the end of this section. In the next couple of paragraphs we discuss simple qualitative decay.

To reconstruct qualitative decay processes, we assume binary variables Q_i that apply to spatiotemporal locations α of a small but finite size as their domain. The value of $Q_i(\alpha)$ indicates whether a given type of particle q_i is present ($Q_i(\alpha) = q_i$) or absent ($Q_i(\alpha) = \text{off}$) at the location α or in a certain neighborhood $f(\alpha)$ of α that is determined by a spatiotemporal function f; in the latter case the variable carries "f" as a second index: $Q_{f,i}$. In what follows, we consider a symmetric decay process in which a particle q decays into two type-identical particles q_l and q_r that are moving away from each other in opposite directions (see fig. 3). We let $Q_l(\alpha)$ be the binary variable denoting the (non)occurrence of a particle q_l shortly after the decay time at the *left* spatial side of α and $Q_r(\alpha)$ the (non)occurrence of q_r at the *right* side of α.

Since neither the state of the particle Q nor any further variable before the decay screens off Q_l from Q_r, condition SCC of C is violated: $P(q_l|q_r, q) >$

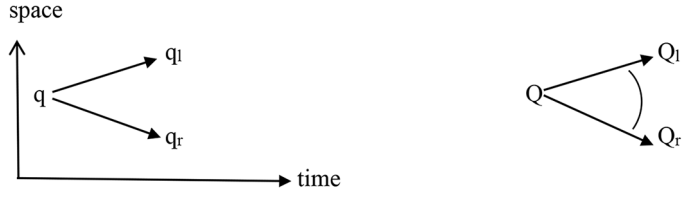


Figure 3. Indeterministic decay: Spatiotemporal diagram (*left*) and causal graph (*right*). Particle q (described by variable Q) decays into decay products q_l and q_r (described by Q_l and Q_r, respectively). SCC fails since neither Q nor any extended set of common causes in the past screens off Q_l from Q_r. "Q" is called an interactive (common) cause and is graphically represented by the arc.

https://doi.org/10.1086/692143 Published online by Cambridge University Press

$P(q_l|q)$; that is, $DEP(Q_l, Q_r|Q)$. Moreover, $DEP(Q_l, Q_r|Past)$ holds for any set "Past" of variables describing events before the decay time.

Note that for the violation of SCC one need not assume that the conservation law holds with probability 1; it suffices that it holds with a probability p' higher than the decay probability p. We follow Salmon (1984, 158–77) and call a common cause that does not screen off its effects an *interactive* (common) *cause* and a common cause that does screen off its effects a *conjunctive cause*.

There are four prominent defense strategies for SCC against this type of counterexample:

(I) *Value refinement of the interactive cause*: Salmon's example of an interactive cause was one of deterministic physics. Spirtes et al. (1993, 63) argued that if in that example the state of particle q immediately before the decay were described with maximal precision, it would imply the values of Q_l and Q_r with probability 1 and, thus, screen them off. However, in our example particle q decays only with a certain probability, so this defense does not apply.

(II) *Linking the effects*: One way to reconcile interactive forks with condition SCC is to add a direct causal link between the two effects (see fig. 4*a*). This possibility is typically excluded in the given examples of indeterministic decay because the two effects occur simultaneously; so a causal link between them would violate condition T, which holds in classical and relativistic physics because of their assumption of locality. In contrast, in quantum mechanics it is debated whether nonlocal causal links should be admitted (see Wood and Spekkens 2015, 14ff.; Näger 2016).

Even without the assumption of condition T, the causal reconstruction in figure 4*a* leads to the problem of unexplained violations of faithfulness. There is a crucial difference between the conditional dependence between effects of an interactive cause and that between effects of a common cause with an additional causal link between them. In a faithful reconstruction of the latter model this causal link can be used to transmit the influence of an intervention on one effect to the other effect conditional on the common cause, but this is impossible in the case of an interactive fork. In our example of decay, such an intervention could be realized by positioning a device

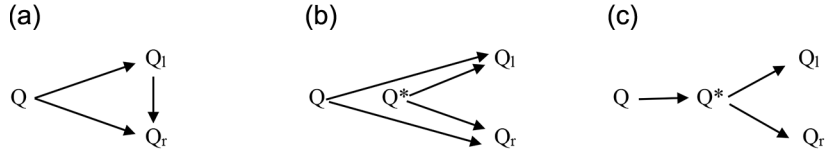


Figure 4. Defense strategies against the counterexample in figure 3. *a*, Linking the effects; *b*, hidden common cause; *c*, hidden intermediate cause.

https://doi.org/10.1086/692143 Published online by Cambridge University Press

at the left side of α that emits particles of type q_l so that the probability of finding p_1 conditional on the presence of q is increased, compared to the distribution without this device. Formally, $P_I(q_l|q) > P(q_l|q)$, where P is the preintervention and P_I the postintervention distribution. While in the case of indeterministic decay, the postintervention probability of the other decay product q_r would be unchanged ($P_I(q_r|q) = P(q_r|q)$), it would be increased in a faithful version of figure 4*a* ($P_I(q_r|q) > P(q_r|q)$). In what follows we call this characteristic feature of interactive forks the *nontransmission* condition. In section 3 we demonstrate that a C-satisfying reconstruction of this feature is not possible without postulating mysterious unfaithful independencies.

(III) *Introduction of a hidden common cause*: Pearl (2000, 62, fig. 2.6) suggested explaining $DEP(Q_l, Q_r|Q)$ by assuming an additional hidden common cause Q* that together with Q screens off Q_l from Q_r ($INDEP(Q_l, Q_r|Q, Q^*)$; see fig. 4*b*). However, this defense strategy cannot explain a second characteristic feature: Interactive cause variables screen off their common effects when their value is "off," that is, when the decaying particle or conserved quantity is absent. To see this, assume that even if there is no particle q in the spatiotemporal region α there is a small probability that a subparticle $q_{l/r}$ spontaneously appears in the left/right neighborhood of α, because it has moved in from the environment. Since the spontaneous appearances of q_l and q_r are mutually independent, $INDEP(Q_l, Q_r|Q = \text{off})$ must hold. However, this feature is not explained by the graph in figure 4*b*: Since the two common causes Q and Q* are causally unconnected, Q = off cannot entirely screen off Q_l from Q_r because of the influence of Q*. The common cause scenario could only explain this feature if a directed arrow led from Q to Q* that inactivates Q* whenever Q is absent, $P(Q^* = \text{off}|Q = \text{off}) = 1$. But then Q* would play the role of an intermediate cause, which is the fourth defense strategy to be discussed below. We call this second characteristic feature of interactive causes the *absence condition*. It becomes important in a revised search procedure for interactive CMs.

(IV) *Introduction of a hidden intermediate cause*: Even if the particle decays with a probability of less than 1, many authors have argued for the existence of some intermediate cause q* (described by variable Q*) in the form of a *breaking event* or *state* of the particle immediately before the decay, so that q* causes the decay with certainty and, hence, Q* screens off Q_l from Q_r (see fig. 4*c*). This defense strategy is, for example, used by Haussman and Woodward (1999, 563) against Cartwright's counterexample (2007, 122), in which a chemical process C succeeds 80% of the time, and when it succeeds it invariably produces two end products, a chemical X and a pollutant Y. These authors assume a certain intermediate event, the "operation of firing F," that deterministically produces X and Y and screens off X from Y.

In this defense strategy the idea is to try to find a description of the decay process that makes it deterministic. The refutation of this strategy is more involved. We will present three arguments to show that this strategy either vi-

https://doi.org/10.1086/692143 Published online by Cambridge University Press

olates the assumption of indeterminism or rests on the artificial effects of a too coarse-grained description.

Argument 1: The first argument says that a *breaking event* is nothing but a temporal succession of two consecutive events or states: at time t, the particle q is still in the unbroken state, while at the next time point, it is in the "broken state," which means that instead of q we now have the two decay products q_l and q_r, possibly very close together but already separated. If we reconstruct this in a discrete time, the picture is as in figure 5.

Since the breaking event b(t, t + 1) is a compound of both the cause and its direct effects, it screens off its direct effects trivially; that is, INDEP(Q_l(t + 1), Q_r(t + 1)|b(t, t + 1)) holds for analytical reasons; the screening off of the indirect effects follows from this by condition SIC. As soon as we refine the description and discriminate between the events at time t and at time t + 1, the screening off disappears.

Proponents of this defense strategy tend to reply to this argument with the objection that time is continuous, so there are no such two consecutive time points t and t + 1. A closer analysis shows that the continuity assumption cannot rectify the violation of the causal Markov condition. If time is continuous, two cases are possible: in case 1, there is a last time point $t \in \mathbb{R}$ at which q is unbroken, while for all times $t' > t$ the particle is broken. Thus, the set of times for which the particle is unbroken forms a right-closed interval and the set of times for which it is broken an adjacent left-open interval (see fig. 6*a*). Then for all times $t' < t$ Q(t) does not screen off $Q_l(t')$ from $Q_r(t')$, which is a violation of condition SCC.

Alternatively, in case 2 there is a first time point $t \in \mathbb{R}$ at which q is replaced by its decay products q_l and q_r. Then the set of times for which the particle is unbroken forms a right-open interval, and the set of times for which it is broken, an adjacent left-closed interval (see fig. 6*b*). Then for all times $t' > t$ Q(t') does not screen off Q_l(t) from Q_r(t), which is again a violation of SCC.

Argument 2: The first argument shows that for every analysis of a breaking event as a succession of fine-grained events condition SCC fails. Some proponents of SCC reply to this argument as follows: although this may be

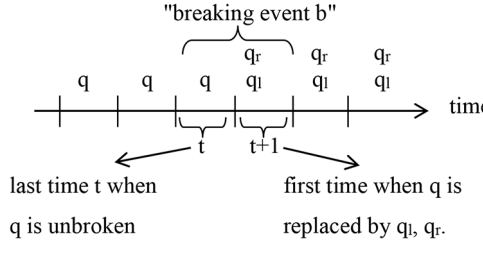


Figure 5. Reconstruction of a "breaking event" in discrete time.

https://doi.org/10.1086/692143 Published online by Cambridge University Press

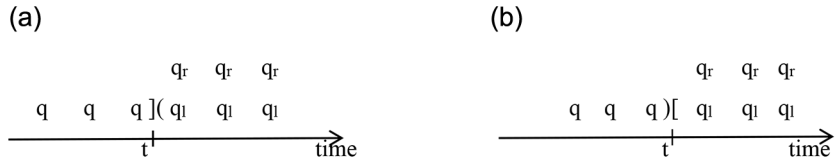


Figure 6. Two possible reconstructions of a "breaking event" in continuous time. *a*, t = last time at which q is unbroken; unbroken state: right-closed interval "]"; broken state: left-open interval "(." *b*, t = first time at which q is broken; unbroken state: right-open interval ")"; broken state: left-closed interval "[."

so, it is at least possible to describe the causal structure in such a way that SCC holds, namely, by using the compound variable of a "breaking event." At this point our test heuristic of section 1, condition 7, comes into play. What the defender of SCC says with this objection is, in effect, that SCC holds merely because the description is coarse-grained and hides the failure of SCC, which appears in all refined models that decompose the breaking event into a succession of fine-grained events. By test condition 7 this is strong reason to suppose that condition SCC is violated in the actual world.

Argument 3: The objection remains that the "breaking event" is more than an instantaneous event consisting of a mere succession of an unbroken and a broken state. Rather, it refers to a specific *prebreaking stage* q* that the particle enters shortly before it decays (see fig. 7).

Even if this assumption is made, the analysis given in the previous argument applies: either there is a first time point t at which the particle is broken (i.e., q* is replaced by q_1 and q_2), or there is a last time point at which the particle is unbroken (i.e., q* exists). If the decay is indeterministic, then in both cases condition SCC is violated. The only way to save condition SCC is to assume (as in argument 1) that the transition from the prebreaking state q* to the decay products q_l and q_r is deterministic. But this is implausible, since according to contemporary scientific understanding, radioactive decay or the breaking of chemical bonds (etc.) are indeterministic processes. Moreover, why should we assume that although the beginning (the left end) of the particle's prebreaking stage is an indeterministic process ($P(q^*|q) < 1$), its ending (the right end) should be a deterministic process? This seems too strange to be plausible: if the decay process is indeterministic, then it is indetermin-

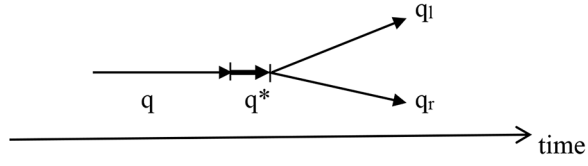


Figure 7. Particle q enters a prebreaking state q* before it decays.

https://doi.org/10.1086/692143 Published online by Cambridge University Press

istic "on both sides." To summarize, in all plausible reconstructions of indeterministic decay processes, the particle's state before the decay does not screen off the states of its decay products.

So far we have only analyzed *simple* (qualitative) indeterministic decay, in which the decay itself occurs spontaneously. However, the same reasoning applies in cases of *generalized indeterministic decay*: here the decay event may have a deterministic cause, but a conserved quantity X of the unbroken particle q is indeterministically distributed over its decay products, in the sense that (i) $P(X(q_r) = x_2|X(q) = x_1)$ is small but (ii) $P(X(q_r) = x_2|X(q) = x_1, X(q_l) = x_1 - x_2)$ is (close to) 1. Conditions i and ii imply the interactive fork condition $DEP(X(q_r)|X(q), X(q_l))$. As an example, consider the generalized decay of a high-energy gamma particle (γ) into a positron (e^+) and an electron (e^-), which is also called "pair creation." It is caused by the clash of a γ particle against an atomic nucleus a. Since the momentum $\vec{p}$ of the γ particle is absorbed by a ($\vec{p}(a) = \vec{p}(\gamma)$), the total momentum of p and e must be 0 ($\vec{p}(e^+) + \vec{p}(e^-) = 0$). This means that e^+ and e^- fly apart in opposite directions, but the particular direction (and momentum) is undetermined. This leads to the interactive fork conditions i and ii above, if we substitute γ for q, e^+ for q_l, e^- for q_r, and $\vec{p}$ for X.

The notion of generalized indeterministic decay also applies to far-distance correlations in quantum-mechanical EPR/Bell experiments. Here a source emits a pair of entangled particles (e.g., electrons) flying in opposite directions. The interactive cause is the state of the entangled two-electron system q, whose total spin in the direction of the *z*-axis is conserved. Upon measurement at one of the two "sides" of the spatially extended two-electron system, q's wave function collapses into a statistical ensemble of the states of the two separate electrons q_l and q_r. This collapse is described as a generalized decay. While conditional on the measurement this decay obtains with certainty, the total spin is indeterministically distributed over the two electrons, satisfying the conservation law $S(q_l) + S(q_r) = S(q)$. Depending on the states (polarization angles) of the measurement devices M_l and M_r, measurement of $S(q_l)$ and $S(q_r)$ yields the results R_l and R_r, respectively (with $Val(R_i) = \{up, down\}$), whose distribution satisfies Bell's famous inequality-violating correlations: $DEP(R(q_l), R(q_r)|M_1, M_2, S(q))$ (see Wood and Spekkens 2015). In section 3 we argue that this correlation can be faithfully explained by an interactive cause but not by a common cause model. The difference from classical indeterministic decay does not concern the basic structure of an interactive fork but rather the nonlocal nature of the entangled system whose wave function stretches over a huge distance, although it is instantaneously influenced by measurements on either side of the electron pair.

The conclusion of this section is that in indeterministic decay processes (in the simple or generalized sense) condition SCC is violated. Philosophers

https://doi.org/10.1086/692143 Published online by Cambridge University Press

who accept this conclusion may nevertheless defend the causal Markov condition by pointing out that indeterministic decays are exceptions that do not arise in everyday life or in the special (life or social) sciences. This does not seem to be true, either. An everyday life example of a violation of SCC is a rock q in the Alps from which a piece q_1 breaks off spontaneously, leaving the now smaller rock q_2 behind, with $\text{mass}(q) = \text{mass}(q_1) + \text{mass}(q_2)$. Given rock q at time t, the probability that in the next moment of time $t + 1$ a piece q_1 breaks off is low, but given q_2 at time $t + 1$, with near certainty the missing piece q_1 is in the (dangerous) state of falling downward at time $t + 1$.

Concerning the social sciences, we conjecture that quantities possessed by social groups (national wealth, etc.) may undergo processes of generalized decay: often the total amount of such a quantity remains fairly constant, while its social distribution undergoes unpredictable changes. An elaboration of this conjecture is left to the future.

In conclusion, it appears that violations of condition SCC due to interactive causes may occur in all domains. In the next section we develop a revision of condition C that accounts for interactive causes.

3. Revision of TCN. The minimal revision of TCN that accounts for interactive causes is to exempt interactive causes from the d-connection condition C, or more precisely, to exempt them from the set of nodes whose presence in the conditioning set blocks a d-connection. In what follows we call the minimal revisions of the notions of d-connection, condition C, and faithfulness "d*-connection," "condition C*," and "d*-faithfulness," respectively. The unasterisked and asterisked notions coincide for CMs without interactive causes but differ in regard to interactive causes. We will argue that the minimal revision is the right revision, except for an additional constraint that accounts for the absence condition (recall sec. 2).

We display an interactive cause graphically by an arc that connects the arrows pointing toward its direct effects.[7] We call such an arc an *interactive arc* and assume that effects connected by such an arc are d*-connected conditional on their interactive cause (see fig. 8*a*); we also say they are "interactively d*-connected." If the interactive CM is d*-faithful, then two interactively d*-connected effects will be probabilistically dependent conditional on their interactive cause; we say in this case that they are "interactively (probabilistically) dependent." Moreover, an interactive cause may have more than two effects, as in figure 8*b*: an example is the β-decay of a neutron into proton, electron, and anti-neutrino.

7. Arcs are more specific than double-headed arrows, which are used in the literature to represent unexplained correlations of different sorts, e.g., due to hidden conjunctive causes. See Richardson (2003).

https://doi.org/10.1086/692143 Published online by Cambridge University Press

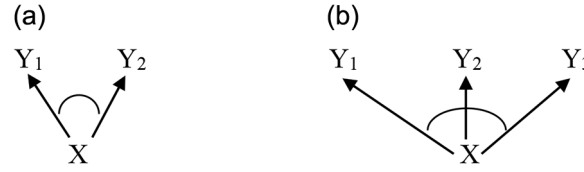


Figure 8. Interactive forks are displayed by arcs connecting interacting arrows.

In section 2 we characterized the correlations between interactively d*-connected effects by the nontransmission condition, which distinguishes them from the conditional correlations between effects of a common cause that result from of a causal link between them. In the latter case, the conditional dependence between the effects can be used to propagate the influence of an intervention on the causing effect to the other effect (provided the CM is faithful). This is depicted in figure 9*a*, where Z_1 figures as an intervention variable on the causing effect Y_1: the result is DEP(Y_2, Z_1|X). In contrast, the conditional correlation between the effects of interactive causes cannot be used to propagate the influence of an intervention: in figure 9*b* INDEP(Y_2, Z_1|X) holds.

Condition C* (the minimal revision of C) is sufficient to account for the behavior displayed in figure 9. It implies for figure 9*a* that Y_2 is d*-connected (as well as d-connected) with Z_1 given ∅ and given X, along the path $Z_1 \rightarrow Y_1 \rightarrow Y_2$. For figure 9*b* it yields that Y_2 is not d*-connected (or d-connected) with Z_1 given ∅ or given X. To ensure d*-connection, the common effect Y_1 has to be included in the conditioning set; that is, in figure 9*b* Y_2 is d*-connected (and d-connected) with Z_1 given Y_1.

Condition C* has the further consequence that in figure 9*b* Y_2 is d*-connected with Z_1 given {X,Y_1}, which is not the case for d-connection in figures 9*b* or 9*a*. Do we have a corresponding probabilistic dependence DEP(Y_2, Z_1|X, Y_1) in figure 9*b*? If the CM is d*-faithful, the answer is yes. Here is an intuitive example that illustrates this dependence: Assume that the interactive cause $Y_1 \leftarrow \overset{\frown}{X} \rightarrow Y_2$ in figure 9*b* describes the spontaneous decay of particle q into q_l and q_r, but the conservation law is not entirely

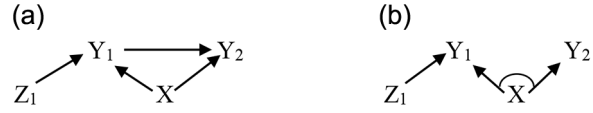


Figure 9. Distinguishing correlations produced by directed arrows (*a*) from correlations produced by interactive causes (*b*). In both CMs (which are assumed to be faithful) we have DEP(Y_2, Y_1|X), but only in *a* do we have DEP(Y_2, Z_1|X); that is, Z_1 can be used as intervention variable on Y_1 with respect to Y_2. In *b* INDEP(Y_2, Z_1|X) holds (nontransmission condition).

https://doi.org/10.1086/692143 Published online by Cambridge University Press

deterministic ($P(q_l|q_r, q) < 1$), and Z_1 is an additional factor whose presence increases the probability of q_l's presence ($Y_1 = q_l$) after the decay. Factor $Z_1 =$ on may represent a field that attracts q_l particles and increases the probability that they fly into region α from outside it. Assume the particle decays ($X = q$) and the right decay product is absent ($Y_2 =$ off). Then the absence of q_r in spite of the presence of q increases the probability that the presence of q_l was caused by $Z_1 =$ on. This means that $P(Z_1 = \text{on}|Y_2 = \text{off}, X = q, Y_1 = q_l) > P(Z_1 = \text{on}|X = q, Y_1 = q_l)$; that is, $DEP(Z_1, Y_2|X, Y_1)$. Thus, the revised condition C* yields the correct result.

A noninteractive and d-faithful reconstruction of the (in)dependences characterizing figure 9*b* would be possible by inverting the direction of the causal arrow $Y_1 \rightarrow Y_2$ in the CM of figure 9*a*; the resulting structure would explain $DEP(Y_1, Y_2|X)$, $INDEP(Z_1, Y_2|X)$, and $DEP(Z_1, Y_2|X, Y_1)$. This is no longer possible, however, in the CM in figure 10*a*, which extends the CM in figure 9*b* by adding a second cause Z_2 to Y_2: $Z_2 \rightarrow Y_2$. The (in)dependences characterizing this CM cannot be described by a noninteractive CM in a faithful way. To explain $DEP(Y_1, Y_2|X)$, $INDEP(Z_1, Y_2|X)$ and $DEP(Z_1, Y_2|X, Y_1)$ in a d-faithful way, we would have to add the arrow $Y_2 \rightarrow Y_1$ to the CM without the arc above X. But this would turn $INDEP(Z_2, Y_1|X)$ into an d-unfaithful independence because then Z_2 is d-connected with Y_1 given X. Moreover, we would additionally need the arrow $Z_2 \rightarrow Y_1$ to explain $DEP(Z_2, Y_1|X, Y_2)$. This unfaithful CM is depicted in figure 10*b*; the bold arrows generate unfaithfulness.

The example of figure 10*a* shows that there are d*-faithful interactive CMs that admit of no faithful noninteractive reconstruction. Thus, misidentifying an interactive as a conjunctive cause may lead to unfaithfulness. This fact may explain recent findings of Näger (2016) and Wood and Spekkens (2015, 2) who applied TCN to quantum-mechanical EPR/Bell correlations and conclude, in the words of Näger (2016, 12), that "the causal faithfulness condition and the causal Markov condition cannot both hold in the quantum realm." Recall the EPR/Bell experiment described in section 2: if we identify the variable S(q) of this experiment (describing the spin state of the entangled electron pair q before measurement) with the variable X in figure 10*a*, the variables M_1 and M_2 describing the state of the measurement devices with Z_1 and Z_2 in figure 10*a*, and the variables $R(q_l)$ and $R(q_r)$ describing the mea-



Figure 10. *a*, d*-faithful interactive CM whose (in)dependencies have no faithful noninteractive modeling; *b*, d-unfaithful CM of the (in)dependencies characterizing *a* (bold arrows generate unfaithfulness).

https://doi.org/10.1086/692143 Published online by Cambridge University Press

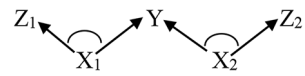


Figure 11. Two interactive causes may share an effect: Z_1 and Z_2 are d*-connected and dependent only if Y is in the conditioning set.

surement results with Y_1 and Y_2 in figure 10*a*, then we obtain the causal structure describing the EPR/Bell experiment. The nontransmission condition that characterizes interactive causes, INDEP(Z_1, $Y_2|X$), corresponds to the famous "no signaling" condition. As we explained above, a C-satisfying reconstruction of the probabilistic (in)dependences in this experiment without assuming a hidden common cause is only possible if one introduces the additional arrows $Y_2 \rightarrow Y_1$ and $Z_2 \rightarrow Y_1$ that imply "superluminal causation" and generate unfaithfulness (Wood and Spekkens 2015, fig. 23).[8] In contrast, the reconstruction of the EPR/Bell experiment as an interactive CM can explain all (in)dependencies of this experiment in a d*-faithful way. These remarks are intended as promissory notes concerning the applicability of the proposed account to quantum mechanics; working out the details would require a paper of its own.

So far we have considered interactive effects having a further cause. In figure 11 we consider two interactive forks that are connected via a common effect. Here Z_1 and Z_2 are not d*-connected given $\{X_1,X_2\}$ but are d*-connected given $\{X_1,X_2,Y\}$, that is, conditional on their common effect. By C* we get INDEP($Z_1, Z_2|X_1, X_2$), and by d*-faithfulness, DEP($Z_1, Z_2|X_1, X_2, Y$). As an example, assume that two particles x_1 and x_2 decay into the partially overlapping decay products z_1,y and y,z_2, respectively. The presence of z_2 conditional on x_2 makes y's presence more probable, but this probability increase is not propagated to z_1 conditional on x_1 because the decay of x_1 is independent from that of x_2. However, conditional on the presence of y, x_1, and x_2, the absence of z_2 increases the probability that z_1 is present because x_1 decayed. Again the minimal revision gives the right result.

We finally turn to interactive causes with more than two effects, as in figure 8*b* above. We require that if more than two effects are connected by one arc, then each effect must be interactively dependent on at least one other effect. We consider this as the extension of the minimality axiom for d*-connections: Each d*-connection by an interactive arc must produce at least one interactive probabilistic dependence, so that removing this arc causes a violation of condition C*. But, interactive correlations within one arc need

8. Even the assumption of a hidden common cause of Y_1 and Y_2 cannot explain the distribution over $\{Z_1,Y_1,X,Y_2,Z_2\}$ because it violates Bell's inequalities, but it opens two further possibilities of unfaithful C-satisfying reconstructions, "superdeterminism" and "retrocausation" (Wood and Spekkens 2015, 14, figs. 25–27).

https://doi.org/10.1086/692143 Published online by Cambridge University Press

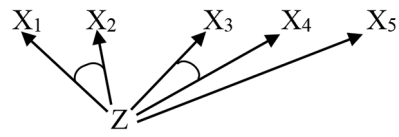


Figure 12. One interactive cause causes two interactive pairs of effects and one ordinary effect. Variables not connected by an interactive arc are conjunctively related.

not always be transitive, so that all pairs of effects within one arc become interactively correlated; this is only the case if the interactive CM is d*-faithful.

We admit the possibility that the same interactive cause produces two (or more) sets of interactive effects in separate interactive events. This is displayed in figure 12, in which $\{X_1,X_2\}$ and $\{X_3,X_4\}$ are sets of interactively d*-connected effects of Z; moreover, X_5 is an additional conjunctive effect of Z. As an example, consider high-energy γ-rays (X) that create two different kinds of particle pairs, electron positrons (X_1,X_2) and muon antimuons (X_3,X_4), and in addition produce heat (X_5). Following from the above requirement on arcs, different sets of interactive effects of one interactive cause are *disjoint*; that is, a variable in one set is merely conjunctively correlated with the variables in another set.

We have seen that the proposed minimal revision of TCN handles the discussed examples in a satisfactory way. Now we are ready to state the exact formulation of our revision of TCN. Recall that the absence of the interactive cause is expressed by letting the interactive cause variable X take the value "off," $X(\alpha) = \text{off}$, meaning that the conserved quantity represented by X is not present in the region α. If the conserved quantity is positive (unsigned)—as in the case of mass, energy, or momentum—the value "off" is identified with the value 0. In application to signed magnitudes such as charge, this identification fails since a particle may have a total charge of 0 but consist of two subparticles having a positive and negative charge of equal amount; in this case the conserved quantity is present, although its total value is 0.

On the basis of the preceding considerations we formally define the notion of interactive CMs as follows ("Pow" for "power set"):

(8) *Definition*: An *interactive causal model* (ICM) is a quadruple (V,E,I,P) such that:

1. (V,E,P) is an ordinary CM (as defined in sec. 1).
2. $I \subseteq V \times \text{Pow}(V)$ is a nonempty set of pairs of the form (X,U), so-called interactive forks, satisfying:
   - *a*) $X \in V$ and $\text{off} \in \text{Val}(X)$ (possession of an absence value),
   - *b*) for all $Y \in U$: $X \rightarrow Y \in E$, and
   - *c*) for all (X,U_i), $(X,U_j) \in I$: $U_i \cap U_j = \varnothing$ (disjointness of arcs).

https://doi.org/10.1086/692143 Published online by Cambridge University Press

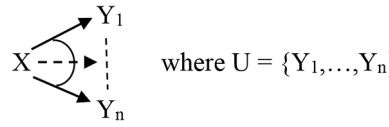


Figure 13. Interactive fork (X,U).

The node X of (X,U) is (called) an *interactive cause* with respect to U, the set U is an *interactive arc*, and the interactive fork (X,U) is graphically displayed as in figure 13.

3. (V,E,I) is called the interactive causal graph of the ICM.

Interactive (as well as conjunctive) causes are now relativized to sets U_i of effects. Variable X is called an interactive (conjunctive) cause simpliciter if X is an interactive (conjunctive) cause with respect to some U_i (these latter notions are no longer disjoint). For $U \subseteq V$ a subset of variables of an interactive causal graph (V,E,I), $I(U) \subseteq U$ denotes U's subset of interactive causes in (V,E,I). The notion $I(U) = \text{off}$ abbreviates the conjunction of the conditions $X = \text{off}$ for all $X \in I(U)$ (should I(U) be empty, the condition $I(U) = \text{off}$ becomes vacuous). An interactive cause Z lies on a path π iff π carries $Y \leftarrow \overset{\frown}{Z} \rightarrow X$; that is, Z's outgoing arrows on this path are connected by an arc (which may connect other effects), or formally, if for some $U \subseteq V$, $(Z, U) \in I$ with $\{X, Y\} \subseteq U$. A conjunctive cause Z lies on path π iff π carries $Y \leftarrow Z \rightarrow X$ and Z carries no arc; that is, there exists no $U \subseteq V$ with $(Z, U) \in I$ and $\{X, Y\} \subseteq U$. The other graph-theoretic notions remain unchanged.

Our notion of d*-connection exempts interactive causes from clause *a*:

(9) *Definition*: Two nodes X, Y of an interactive graph (V,E,I) are *d*-connected* given a set of nodes U iff they are connected by a path π in (V,E) such that

a) no intermediate cause or conjunctive common cause lying on π is in U, and

b) every common effect on π is in U or has an effect in U.

The next definition explicates the revised causal d*-connection condition. It consists of the minimal revision (i) together with the absence condition (ii):

(10) *Definition*: An ICM (V,E,I,P) satisfies the *causal d*-connection condition* C* iff for all $X, Y \in V$ and $U \subseteq V - \{X, Y\}$:

i) if $\text{DEP}(X, Y|U)$, then X and Y are d*-connected given U, and

ii) if $\text{DEP}(X, Y|U, I(U) = \text{off})$, then X and Y are d-connected given U.

Axiom: Every physically possible system (V,E,P) that is [sufficiently] complete satisfies condition C*.

https://doi.org/10.1086/692143 Published online by Cambridge University Press

Observe that d-connection implies d*-connection, so if a CM satisfies C, it satisfies C*.

Condition 10.ii is equivalent to "if X and Y are d-separated given U, then INDEP(X, Y|U, I(U) = off)." This condition breaks the symmetry between C and C*, but it is useful for the purpose of constructing a discovery algorithm for ICMs.

Condition C* is the revised global Markov condition. For finite acyclic ICMs C* is equivalent to a correspondingly revised local Markov condition M*. We say that Y is an *interactive nonsuccessor* of X if Y is neither an ordinary X successor nor a successor of a variable Z that is connected with X by an interactive arc:

(11) *Definition*: An ICM (V,E,I,P) satisfies the *revised Markov condition* M* iff every $X \in V$ is

i) independent of any subset of interactive nonsuccessors of X conditional on Par(X), and

ii) independent of any subset of (ordinary) nonsuccessors of X conditional on Par(X) and the condition I(Par(X)) = off.

Theorem 1. For acyclic finite ICMs, conditions C* and M* are equivalent. (Proof omitted.)

The conditions of d*-faithfulness and d*-minimality are defined as follows:

(12) *Definition*:

1. An ICM (V,E,I,P) satisfies the condition of *d*-faithfulness* (F*) iff for all $X,Y \in V$ and $U \subseteq V - \{X, Y\}$: if X and Y are d*-connected given U, then DEP(X, Y|U).
2. A C*-satisfying ICM (V,E,I,P) satisfies the condition of *d*-minimality* (Min*) iff no arrow can be removed from E and no interactive fork from I without violating condition C*.

Axiom: Every physically possible C*-satisfying causal system satisfies Min*.

Axioms C* and Min* form the core of TCN*. Possible extensions are obtained by adding a cautious version of F* or condition T.

4. Conclusion and Outlook: The Discovery Problem for TCN*. In this article we argued that processes of indeterministic decay involve interactive causes. They violate the principle of "screening-off by common causes," which is entailed by causal Markov condition, the core axiom of the theory of causal nets (TCN). We showed that there are interactive CMs that admit of

https://doi.org/10.1086/692143 Published online by Cambridge University Press



Figure 14. Additional discovery step for ICMs: Can an undirected triangle (*a*) be reconstructed as an interactive fork (*b*)?

no faithful noninteractive reconstruction. We developed a revised version the TCN, TCN*, that correctly predicts the probabilistic (in)dependencies of interactive CMs and coincides with TCN for CMs without interactive causes.

We cannot discuss the discovery problem for TCN*. Instead, we confine ourselves to brief observations. Recall the Spirtes et al. discovery algorithm for noninteractive CMs (Spirtes et al. 1993, 114; Pearl 2000, 50): it first finds the unique undirected causal graph (assuming conditions C and F) and then orients some edges as common effects. We suggest that a revised discovery algorithm for ICMs can be obtained from the Spirtes et al. algorithm by adding a third step, roughly described as follows. For every undirected triangle (X,Y,Z) in the undirected graph and node X in this triangle, as in figure 14*a*, the algorithm checks whether an interactive cause interpretation of X is possible, as in figure 14*b*. This presupposes that (i) the absence condition holds ($\text{off} \in \text{Val}(X)$ and $\text{INDEP}(Y, Z|X = \text{off}, U)$ for some $U \subseteq V$) and (ii) edges between X, Y and neighboring nodes that have already been oriented are coherent with the nontransmission condition. The design of a revised discovery algorithm along these lines is a task for the future.

REFERENCES

Arntzenius, Frank. 2010. "Reichenbach's Common Cause Principle." In *Stanford Encyclopedia of Philosophy*, ed. Edwin N. Zalta. Stanford, CA: Stanford University. http://plato.stanford.edu/archives/fall2010/entries/physics-Rpcc/.

Berkson, Joseph. 1946. "Limitations of the Application of Fourfold Table Analysis to Hospital Data." *Biometrics Bulletin* 2 (3): 47–53.

Cartwright, Nancy. 2007. *Hunting Causes and Using Them*. Cambridge: Cambridge University Press.

Earman, John. 1986. *A Primer on Determinism*. Dordrecht: Reidel.

Haussman, Daniel, and James Woodward. 1999. "Independence, Invariance, and the Causal Markov Condition." *British Journal for the Philosophy of Science* 50:1–63.

Hitchcock, Christopher. 2010. "Probabilistic Causation." In *Stanford Encyclopedia of Philosophy*, ed. Edwin N. Zalta. Stanford, CA: Stanford University. http://plato.stanford.edu/archives/win2011/entries/causation-probabilistic/.

Lauritzen, Steffen L., Philip A. Dawid, B. N. Larsen, and H.-G. Leimer. 1990. "Independence Properties of Directed Markov-Fields." *Networks* 20:491–505.

Näger, Paul. 2016. "The Causal Problem of Entanglement." *Synthese* 193:1127–55.

Norton, John. 2008. "The Dome: An Unexpectedly Simple Failure of Determinism." *Philosophy of Science* 75:786–98.

https://doi.org/10.1086/692143 Published online by Cambridge University Press

Pearl, Judea. 1988. *Probabilistic Reasoning in Intelligent Systems*. San Francisco: Kaufmann.
———. 2000. *Causality: Models, Reasoning and Inference*. Cambridge: Cambridge University Press.
———. 2009. *Causality: Models, Reasoning and Inference*. 2nd ed. Cambridge: Cambridge University Press.
Reichenbach, Hans. 1956. *The Direction of Time*. Berkeley: University of California Press.
Richardson, T. 2003. "Markov Properties for Acyclic Directed Mixed Graphs." *Scandinavian Journal of Statistics* 30:145–57.
Salmon, Wesley. 1984. *Scientific Explanation and the Causal Structure of the World*. Princeton, NJ: Princeton University Press.
Schurz, Gerhard, and Alexander Gebharter. 2016. "Causality as a Theoretical Concept." *Synthese* 193:1073–1103.
Spirtes, Peter, Clark Glymour, and Richard Scheines. 1993. *Causation, Prediction, and Search*. Cambridge, MA: MIT Press.
———. 2000. *Causation, Prediction, and Search*. 2nd ed. Cambridge, MA: MIT Press.
Van Fraassen, Bas. 1980. *The Scientific Image*. Oxford: Clarendon.
Wood, Christopher J., and Robert W. Spekkens. 2015. "The Lesson of Causal Discovery Algorithms for Quantum Correlations." *New Journal of Physics* 17. doi:10.1088/1367-2630/17/3/033002.
Zhang, Jiji, and Peter Spirtes. 2008. "Detection of Unfaithfulness and Robust Causal Inference." *Minds and Machines* 18:239–71.
———. 2011. "Intervention, Determinism, and the Causal Minimality Condition." *Synthese* 182 (3): 335–47.

https://doi.org/10.1086/692143 Published online by Cambridge University Press