

# A Regularity Theory of Causation

Holger Andreas and Mario Günther\*

Forthcoming in *Pacific Philosophical Quarterly*

## Abstract

In this paper, we propose a regularity theory of causation. The theory aims to be reductive and to align with our pre-theoretic understanding of the causal relation. We show that our theory can account for a wide range of causal scenarios, including isomorphic scenarios, omissions, and scenarios which suggest that causation is not transitive.

**Keywords.** Causation, Regularity Theory, Causal Models.

## 1 Introduction

Hume observed that all our reasoning about matters of particular fact is grounded in the relation of cause and effect. This relation, if any, we understand perfectly. And yet there is no theory of causation at our disposal that tells us reliably what it means that *this* is a cause of *that*. The lack of a satisfying theory of causation is, moreover, not for a lack of effort. There have been many attempts to analyse causation. Hume himself defined a cause to be

an object, followed by another, and where all the objects similar to the first are followed by objects similar to the second. Or in other words where, if the first object had not been, the second never had existed. (Hume, 1748/1975, Sect. VII)

---

\*Mario.Guenther@lmu.de. The authors contributed equally.

The first sentence says that the relation between cause and effect instantiates a regularity (Andreas and Günther, 2021). This idea has been developed further. We learned from Mackie (1965, 1974) and Wright (1985) that the sets of actual conditions which are jointly sufficient for the effect to occur must be minimal. And Baumgartner (2013) shows that the regularities must be non-redundant for the effect. The regularity theory thus counts as a cause each member of any minimal set of actual conditions that are jointly sufficient for the effect to occur in the presence of non-redundant regularities.

The second sentence of the quote says that an effect counterfactually depends on its cause. This idea has been taken up by Lewis (1973) who analyses causation as the transitive closure of counterfactual dependences between actual events and absences. For Hitchcock (2001), a cause is connected by an active causal path to its effect. And Gallow (2021) adds that the active causal path, or rather active causal network, must transmit deviancy from cause to effect.

In this paper, we propose a regularity theory of deterministic token causation.<sup>1</sup> The motivating idea is this: causation is deviant forward-directed inferability along the causal paths of direct non-redundant regularities. Building on Baumgartner's work, we aim to reduce causation to true propositions of particular fact and deviancy from norms. The true propositions of particular facts and minimization procedures give rise to non-redundant regularities. This will allow us to define *causal models* in terms of true propositions of particular fact and direct non-redundant regularities. These causal models are ultimately obtained from the true propositions of particular fact, unlike the causal models employed by the counterfactual theories of Hitchcock (2001), Halpern and Pearl (2005), and others. We will show that our regularity theory aligns with the commonsense understanding of the causal relation—even in causal scenarios where Baumgartner's does not.

We proceed as follows. First, we present a regularity theory due to Lewis (1973) and restrict the regularities to direct and non-redundant ones. We then define causal models as tuples of direct non-redundant regularities and propositions of particular fact. This allows us to state our preliminary theory of token causation. We will show how our theory solves a wide

range of causal scenarios and amend it where necessary. Finally, we aim to establish the reductiveness of our theory by avoiding model-relativity.

## 2 Direct Non-Redundant Regularities

Lewis (1973) authored and rejected a regularity theory. It can be summarized as follows. Let  $A$  be the proposition that is true if and only if (iff) some token event  $a$  of type  $A$  occurs. If no token event of type  $A$  occurs,  $\neg A$  is true. Furthermore, let  $\mathcal{L}$  denote a set of true regularities, and  $\mathcal{F}$  a possibly empty set of true propositions describing particular facts.<sup>2</sup> An event  $c$  is a cause of a distinct event  $e$  iff there is a set  $\mathcal{L}$  of true regularities and a set  $\mathcal{F}$  of true propositions of particular fact such that

- (1)  $C$  and  $E$  are true,
- (2)  $\mathcal{L} \cup \mathcal{F} \models C \rightarrow E$ ,
- (3)  $\mathcal{L} \cup \mathcal{F} \not\models E$ , and
- (4)  $\mathcal{F} \not\models C \rightarrow E$ .

(1) says that cause and effect are actual. (2) says that a cause expressed by  $C$  is sufficient for its effect in the presence of  $\mathcal{L} \cup \mathcal{F}$ . However, (3) says that  $\mathcal{L} \cup \mathcal{F}$  alone does not entail  $E$ . Given  $\mathcal{L} \cup \mathcal{F}$ ,  $C$  is necessary for  $E$ . In this sense,  $\mathcal{L} \cup \mathcal{F} \cup C$  is a *minimal* set sufficient for  $E$ . (4) says that  $\mathcal{F}$  alone does not entail the material implication  $C \rightarrow E$ .  $\mathcal{L}$  is not redundant in the minimal sufficiency of  $\mathcal{L} \cup \mathcal{F} \cup C$  for  $E$ .

Lewis made a case that the regularity theory cannot properly distinguish between genuine causes and both effects and preempted would-be causes. The underlying reason is that some of the true regularities in  $\mathcal{L}$  are spurious. Consider, for illustration, the structure in Figure 1, where  $C$  is a common type-cause of  $D$  and  $E$ . In addition,  $B$  and  $A$  are alternative type-causes of  $D$  and  $E$  respectively.




Figure 1: C is a common type-cause of D and E.

$B$  is sufficient for  $D$ , and so is  $C$ . Likewise, each of  $C$  and  $A$  alone is sufficient for  $E$ . But  $\neg B \wedge D$  is also sufficient for  $E$ . When some token event of  $D$  occurs, and no token event of  $B$ , then some token event of  $C$  occurs. For no effect occurs without any of its causes. And  $C$  is sufficient for  $E$ . We have thus the following true but spurious regularity:

$$(\neg B \wedge D) \vee C \vee A \leftrightarrow E.$$

If a token event of type  $C$  occurs, for example, a token event of type  $E$  occurs. Similarly, if no token event of type  $B$  occurs and a token event of type  $D$  occurs, then a token event of type  $E$  occurs. Furthermore, if a token event of type  $E$  occurs, then a token event of type  $A$ , or of type  $C$  occurs.

Given the above regularity is in  $\mathcal{L}$  and  $\mathcal{F} = \{\neg B\}$ , the occurring event  $d$  is then a cause of the occurring event  $e$  on the above regularity theory. Indeed,  $D$  is then an INUS condition and a NESS condition for  $E$ . If the token events  $d$  of  $D$  and  $e$  of  $E$  occur, Mackie's INUS theory and Wright's NESS theory wrongly count the event  $d$  as a cause of  $e$ .

In the meantime, Baumgartner (2013) has shown that spurious regularities can be pruned from  $\mathcal{L}$  by requiring them to be non-redundant. A non-redundant regularity for  $E$  is a minimally necessary disjunction of minimally sufficient conjunctions for  $E$ . To be precise, let  $\mathbf{C}$  be a set of minimal sets  $\mathcal{A}_i$  of propositions of particular fact that are sufficient for  $E$ .  $\bigwedge \mathcal{A}_i$  is some conjunction of the members of  $\mathcal{A}_i$ , and  $\bigvee_i \bigwedge \mathcal{A}_i$  some disjunction of such conjunctions.<sup>3</sup> We say  $\mathbf{C}$  is *minimally necessary* for  $E$  iff

(a)  $\bigvee_i \bigwedge \mathcal{A}_i \leftrightarrow E$  is true, and

(b) there is no  $\mathbf{C}' \subset \mathbf{C}$  such that  $\bigcup_j \mathcal{A}_j = \bigcup \mathbf{C}'$  and  $\bigvee_j \bigwedge \mathcal{A}_j \leftrightarrow E$  is true.

Whenever  $\mathbf{C}$  is minimally necessary for  $E$ , we say that  $\bigvee_i \bigwedge \mathcal{A}_i \leftrightarrow E$  is a *non-redundant* regularity for  $E$ .

Now,  $\{\{\neg B, D\}, \{C\}, \{A\}\}$  is not minimally necessary for  $E$ , but  $\{\{C\}, \{A\}\}$  is. The non-redundant regularity for  $E$  is thus

$$C \vee A \leftrightarrow E.$$

This solves the problem of joint effects of a common cause under the assumption that any type effect has more than one alternative type cause—more than one minimally sufficient conjunction. The assumption likewise solves the problem of the direction of causation. The non-redundant regularity implies a non-symmetry: each of  $C$  and  $A$  alone is sufficient for  $E$ , but  $E$  is only sufficient for the disjunction  $C \vee A$ . Therefore,  $E$  is not a cause of  $C$  on the above regularity theory (assuming  $\mathcal{L} = \{C \vee A \leftrightarrow E, C \vee B \leftrightarrow D\}$  contains all and only the non-redundant regularities for  $E$  and  $D$ ).

The non-redundant regularities are relative to the set of propositional variables we consider. For this to be seen, consider the above structure restricted to the variable set  $\{C, A, E\}$ . This simple structure of overdetermination is fully described by the non-redundant regularity  $C \vee A \leftrightarrow E$ . Indeed, a token event of type  $E$  occurs only if a token event of type  $C$ , or one of type  $A$  occurs, or both; and for any occurrence of a token event of type  $A$  or  $C$ , there is an occurrence of a token event of type  $E$ .

However, let us consider a structure of preemption enriched by information about the token events. In Figure 2, a grey-shaded node indicates that some token event of the relevant type occurs, and a white node that no token event of the relevant type occurs.




Figure 2: Preemption.

$C$  is sufficient for  $D$ , which in turn is sufficient for  $E$ . Similarly,  $A$  is sufficient for  $B$ , which in turn is sufficient for  $E$ . But  $C$  is also sufficient to prevent any token event of type  $B$  from occurring. This is represented by the ‘preventive’ arrow between  $C$  and  $B$ . In the scenario under consideration, token events of both  $C$  and  $A$  occur.  $c$  is a cause of  $e$ , but  $a$  is a mere would be cause of  $e$ :  $c$  preempts the causal efficacy of  $a$ .

The regularity  $C \vee A \leftrightarrow E$  is still true and non-redundant for the structure of preemption. However, this regularity does not fully describe the structure relative to the variable set  $\{C, A, D, B, E\}$ . In order to obtain a full description, we require that all non-redundant regularities in  $\mathcal{L}$  are direct. A non-redundant regularity  $\bigvee_i \wedge \mathcal{A}_i \leftrightarrow E$  is *direct* relative to a variable set iff there are no true non-redundant regularities  $\bigvee_j \wedge \mathcal{B}_j \leftrightarrow C$  and  $\bigvee_k \wedge \mathcal{C}_k \leftrightarrow E$ , where the same variable appears as subformula in some  $\wedge \mathcal{A}_i$  and some  $\wedge \mathcal{B}_j$ , while  $C$  appears as subformula in some  $\wedge \mathcal{C}_k$ .

The non-redundant regularity  $C \vee A \leftrightarrow E$  is not direct relative to the extended variable set: there are the true non-redundant regularities  $C \leftrightarrow D$  and  $D \vee B \leftrightarrow E$ , and  $C$  appears as subformula in  $C \wedge A$  and  $C$ , while  $D$  appears in  $D$ . (Any formula is a subformula of itself). The non-redundant regularities  $C \leftrightarrow D$ ,  $A \wedge \neg C \leftrightarrow B$ , and  $D \vee B \leftrightarrow E$ , by contrast, are all direct relative to the extended variable set. There are no other direct non-redundant regularities. In general, a causal structure is fully described by all and only the direct non-redundant regularities.

How does Lewis’s regularity theory amended by direct non-redundant regularities treat preemption? Well,  $c$  is a cause of  $e$ , as it should be. For this to be seen, take  $\mathcal{F} = \emptyset$ . Conditions (1)-(4) are then satisfied. Unfortunately,  $a$  also counts as a cause of  $e$ . For this to be seen, observe

that condition (1) is satisfied, and condition (3) is satisfied for  $\mathcal{F} = \emptyset$  or  $\mathcal{F} = \{\neg B\}$ . Both options also satisfy condition (4). As to condition (2), let's consider the first option  $\mathcal{F} = \emptyset$ . There are only two cases: if  $\neg C$ , then the direct non-redundant regularities in  $\mathcal{L}$  entail  $A \rightarrow E$  via  $B$ . If  $C$ , the direct non-redundant regularities in  $\mathcal{L}$  entail  $E$ , and thus  $A \rightarrow E$ . Hence,  $\mathcal{L} \cup \mathcal{F} \models A \rightarrow E$ . Condition (2) is satisfied, and so  $a$  wrongly counts as a cause of  $e$ .

What went wrong? Well,  $A$  and  $B$  are—just like  $C$  and  $D$ —type-causes of  $E$ . What makes the difference in the preemption scenario between  $C$  and  $A$  is that a token event of type  $D$  occurs, while no token event of type  $B$  occurs. The spirit behind regularity theories is that the true propositions of particular fact should be respected. But the entailment of  $A \rightarrow E$  involves an inference via  $B$ , even though no token event of type  $B$  occurs. This suggests that the regularity theory of Lewis is too liberal as to the choice of the set  $\mathcal{F}$  of true propositions of particular fact. There is a minimality constraint on  $\mathcal{F}$ , but a maximality constraint is lacking which would guarantee that  $\neg B \in \mathcal{F}$ .

Such a maximality constraint alone, however, does not help Lewis's regularity theory. Even if  $\mathcal{F} = \{\neg B\}$ ,  $a$  counts as a cause of  $e$ . For this to be seen, observe first that the direct non-redundant regularities in  $\mathcal{L}$  and  $\neg B$  entail  $\neg A \vee C$ . There are again only two cases: if  $\neg C$ , then  $\mathcal{L}$  and  $\mathcal{F}$  entail  $\neg A$ , and thus  $A \rightarrow E$ . If  $C$ , then  $\mathcal{L}$  and  $\mathcal{F}$  entail  $E$ , and thus  $A \rightarrow E$ . Condition (2) is again satisfied. But observe that this reasoning is artificial. Intuitively,  $a$  is not a cause of  $e$  because  $A$  does not entail  $E$  in a forward-directed way via  $B$ . If you prefer, we may alternatively say that there is no causal path from  $A$  over  $B$  to  $E$  that is active in the sense that a token event of each of the types  $A, B$  and  $E$  occurs.

Baumgartner (2013) develops a theory of token-causation from his theory of type-causation and the notion of an active causal path. His theory of type-causation requires that the regularities remain non-redundant under any suitable extension of the variable set. As a result, there are direct and indirect non-redundant regularities that already tell us which relations of type-causation are transitive, and which are not. He defines, roughly, a token event to be a cause of another iff there is an active causal path from the first token event or absence to the second via the direct non-redundant

regularities. (See p. 98 for the formal details.)

Baumgartner’s regularity theory reduces causation to material implications and minimization procedures. He thereby proposes a theory of causation free of any modal notion. And it is notable that his theory accounts well for many causal scenarios, including overdetermination, preemption, as well as some switching scenarios, and some short-circuits. Only very recently counterfactual theories of causation have been able to account for these scenarios (Andreas and Günther, 2021b; Gallow, 2021).

This being said, Baumgartner’s regularity theory works only under the assumption that each variable set is extendable such that each type effect has at least two alternative type causes. The assumption seems empirically accurate, even though it is contingent: it depends on the actual features of our world. We think it would be better if a theory of causation were even able to come to the right verdicts in all conceptual possibilities—including a scenario where one event is the only cause of another. We aim for such a conceptually accurate regularity theory that aligns with our understanding of the causal relation.

### 3 Causal Models

Causal models represent causal scenarios. In a causal scenario like preemption, certain events occur, others do not, and we have a certain structure that tells us how event types depend on other event types. We can define a causal model  $\langle \mathcal{L}, \mathcal{F} \rangle$  by two components: a set  $\mathcal{L}$  of direct non-redundant regularities and a set  $\mathcal{F}$  of propositions of particular fact.  $A \in \mathcal{F}$  means that some token event  $a$  of type  $A$  occurs.  $\neg A \in \mathcal{F}$ , by contrast, means that no token event of type  $A$  occurs. In other words,  $\neg A$  denotes the absence of any event of type  $A$ , or simply the absence of  $A$ .

The preemption scenario can be represented by a causal model  $\langle \mathcal{L}, \mathcal{F} \rangle$ , where  $\mathcal{L} = \{C \leftrightarrow D, B \leftrightarrow A \wedge \neg C, E \leftrightarrow D \vee B\}$  and  $\mathcal{F} = \{C, A, D, \neg B, E\}$ . For readability, we represent causal models in two-layered boxes. The upper layer shows the set  $\mathcal{L}$  of direct non-redundant regularities ordered in the following way: the single variable is on the left-hand side and the formula



in disjunctive normal form on the right-hand side. The lower layer shows the set  $\mathcal{F}$  of propositions of particular fact. For the preemption scenario, we obtain:

$D \leftrightarrow C$
$B \leftrightarrow A \wedge \neg C$
$E \leftrightarrow D \vee B$
$C, A, D, \neg B, E$

Let  $\Gamma$  be a set that contains propositional formulas. We understand entailment in the standard way:  $\Gamma \models \phi$  iff the propositional formula  $\phi$  is satisfied by any classical valuation that satisfies all members of  $\Gamma$ . We define the entailment relation for causal models as follows:

$$\langle \mathcal{L}, \mathcal{F} \rangle \models \phi \quad \text{iff} \quad \mathcal{L} \cup \mathcal{F} \models \phi.$$

Causation is not symmetric, if not asymmetric outright. If  $C$  is a cause of  $E$ , it is thereby not the case that  $E$  is a cause of  $C$ . Recall that non-redundant regularities of the form  $E \leftrightarrow \bigvee_i \mathcal{A}_i$  are generally not symmetric: the type-effect  $E$  is only sufficient for a disjunction, while each disjunct is sufficient for  $E$ .

There are, however, problematic limiting cases. Non-redundant regularities like  $D \leftrightarrow C$  are symmetric:  $D$  is sufficient for  $C$ , and  $C$  is sufficient for  $D$ . There are two promising ways to determine the direction of causation in these limiting cases. One could reconsider a Humean approach to the direction of causation that relies on the direction of time (Andreas and Günther, 202xc). If, for example, the token events of type  $C$  precede the token events of type  $D$ , we may say that  $D$  is the type-effect. Alternatively, one could extend the variable set until the symmetry is broken (Baumgartner, 2013). If we extend the variable set by  $C'$ , which is another type-cause of  $D$ , the resulting non-redundant regularity  $D \leftrightarrow C \vee C'$  breaks the symmetry. In what follows, we assume that the direction of non-redundant regularities can always be determined. We can then say that any type effect is by convention on the left-hand side of a direct non-redundant regularity and its direct type causes on the right-hand side.

Still, non-redundant regularities allow for inferences in two directions.  $E \leftrightarrow D \vee B$ , for example, allows us to infer from either of the type causes  $D$  and  $B$  to the type effect  $E$ . This inference goes from type causes to some type effect. Let us call any such inference *forward-directed*. By contrast, the type effect  $E$  allows us to infer the disjunction of its type causes  $D \vee B$ . This inference goes from some type effect to its type causes. Let us call any such inference *backward-directed*.

One idea behind our theory of causation is that a token cause allows us to infer its token effect in a purely forward-directed way. To implement this idea, we need to isolate the forward-directed consequences of actual events and absences from the backward-directed ones. In general, we can isolate the forward-directed causal consequences of some occurring token event  $a$  of type  $A$  for a causal model  $\langle \mathcal{L}, \mathcal{F}' \rangle$  by a *setting*. Roughly speaking, a setting removes a non-redundant regularity  $A \leftrightarrow \bigvee_i \mathcal{A}_i$  from  $\mathcal{L}$  and replaces it by a true proposition, either  $A$  or  $\neg A$ . Thereby backward-directed inferences from  $A$  or  $\neg A$  are excluded.

Suppose we want to determine the forward-directed causal consequences of the occurring token event  $a$  of type  $A$  for a causal model  $\langle \mathcal{L}, \mathcal{F}' \rangle$ , where  $\mathcal{F}'$  is a subset of the set  $\mathcal{F}$  of true propositions particular facts. The setting of  $A$  in this causal model results in a causal model  $\langle \mathcal{L}_A, \mathcal{F}' \cup \{A\} \rangle$ . If  $A \leftrightarrow \bigvee_i \mathcal{A}_i$  is a member of  $\mathcal{L}$ ,  $\mathcal{L}_A$  is obtained from  $\mathcal{L}$  by removing this non-redundant regularity. Otherwise  $\mathcal{L}_A = \mathcal{L}$ . We call  $\langle \mathcal{L}_A, \mathcal{F}' \cup \{A\} \rangle$  the causal submodel of  $\langle \mathcal{L}, \mathcal{F}' \rangle$  after the setting of  $A$ . By removing the non-redundant regularity for  $A$  from  $\mathcal{L}$ , backward-directed inferences from  $A$  or  $\neg A$  are excluded in the causal submodel.

Complex settings may be represented by a set  $S$  of true propositions of particular fact. Let us denote settings by an operator  $[\cdot]$  that takes a causal model  $\langle \mathcal{L}, \mathcal{F}' \rangle$  and a set  $S$  of true propositions of particular fact, where both  $\mathcal{F}'$  and  $S$  are subsets of the true propositions  $\mathcal{F}$  of particular fact, and returns a causal model—the submodel of  $\langle \mathcal{L}, \mathcal{F}' \rangle$  after the setting of  $S$ . The setting by a set of true propositions of particular fact is defined as follows:

$$\langle \mathcal{L}, \mathcal{F}' \rangle[S] = \langle \mathcal{L}_S, \mathcal{F}' \cup S \rangle \quad (\langle \mathcal{L}, \mathcal{F} \rangle[S])$$

where

$$\mathcal{L}_S = \{(A \leftrightarrow \bigvee_i \bigwedge \mathcal{A}_i \in \mathcal{L}) \mid A \notin S \text{ and } \neg A \notin S\}.$$

$\mathcal{L}_S$  is the subset of  $\mathcal{L}$  that contains each direct non-redundant regularity  $A \leftrightarrow \bigvee_i \bigwedge \mathcal{A}_i$  whose variable  $A$  does not appear in  $S$ . After setting  $S$  in the causal model  $\langle \mathcal{L}, \mathcal{F}' \rangle$ , the set  $S$  becomes part of the propositions of particular fact of the resulting submodel. Note that the resulting submodel is again a causal model consisting of a set of non-redundant regularities and a set of propositions of particular fact.

Settings will always only set true propositions of particular fact. No propositions contrary to the true facts are ever set, unlike the interventions employed by Halpern and Pearl (2005) for example. As a consequence, the resulting submodels are not inconsistent provided the original causal models were not.

We are now in a position to state our preliminary theory of causation.

**Definition 1.** Let  $\langle \mathcal{L}, \mathcal{F} \rangle$  be a causal model such that  $\mathcal{F} \models \mathcal{L}$ .  $c$  is a cause of  $e$  relative to  $\langle \mathcal{L}, \mathcal{F} \rangle$  iff there is a set  $\mathcal{F}' \subseteq \mathcal{F}$  such that all of the following conditions are satisfied:

- (1)  $\langle \mathcal{L}, \mathcal{F} \rangle \models C \wedge E$ .
- (2)  $\langle \mathcal{L}, \mathcal{F}' \rangle \not\models E$ , and there is no  $\mathcal{F}''$  so that  $\mathcal{F}' \subset \mathcal{F}'' \subseteq \mathcal{F}$  and  $\langle \mathcal{L}, \mathcal{F}'' \rangle \not\models E$ .
- (3)  $\langle \mathcal{L}, \emptyset \rangle[\mathcal{F}'][[C]] \models E$ .


(1) says that cause and effect are actual. (2) says that there is a set  $\mathcal{F}'$  of propositions of particular facts that *does not* entail the effect  $E$  in the presence of the direct non-redundant regularities, and  $\mathcal{F}'$  is maximal: any strict superset of  $\mathcal{F}'$  would entail  $E$  in the presence of the direct non-redundant regularities. (3) says that  $\mathcal{F}'$  and the direct non-redundant regularities *do* entail  $E$  together with  $C$  in a forward-directed way.

On our preliminary theory, a cause is each member of any maximised minimal set of actual conditions which are jointly sufficient in a forward-directed way for the effect to occur in the presence of direct non-redundant

regularities. Causation so understood amounts to forward-directed inferability along direct non-redundant regularities.

## 4 Scenarios

Let us revisit the scenario of preemption.



Relative to  $\langle \mathcal{L}, \mathcal{F} \rangle$ ,  $c$  is a cause of  $e$ . Condition (1) is satisfied, and will be satisfied in all of the scenarios to come. Take  $\mathcal{F}' = \{\neg B\}$ . Then  $\langle \mathcal{L}, \mathcal{F}' \rangle \not\models E$ , but any strict superset of  $\mathcal{F}'$  would entail  $E$  in the presence of  $\mathcal{L}$ . Condition (2) is satisfied. And condition (3) is satisfied as well:  $\{D \leftrightarrow C, E \leftrightarrow D \vee B\} \cup \{\neg B\} \cup \{C\} \models E$ .

Relative to  $\langle \mathcal{L}, \mathcal{F} \rangle$ ,  $a$  is not a cause of  $e$ . Condition (2) is only satisfied for  $\mathcal{F}' = \{\neg B\}$ . But then the direct non-redundant regularity of  $B$  is removed from  $\mathcal{L}$  by the setting of  $\neg B$ . Condition (3) is then violated:  $\{D \leftrightarrow C, E \leftrightarrow D \vee B\} \cup \{\neg B\} \cup \{A\} \not\models E$ . Our preliminary theory solves preemption.

We leave it to the reader to verify that our preliminary theory accounts for other classic scenarios like overdetermination, prevention, double prevention, and more. Instead we turn towards the problem of isomorphic causal models.

### 4.1 Isomorphisms

The problem of isomorphic causal models is that there are pairs of scenarios which are structurally indistinguishable for simple causal model accounts,

and yet our causal judgments differ (Hall, 2007, p. 44). We call a causal model account simple if it only factors in structural equations—or our direct non-redundant regularities—together with values of variables—or our propositions of particular fact.

Let us illustrate an instance of the problem. Here is the causal model of a scenario of overdetermination.

$E \leftrightarrow C \vee A$
$C, A, E$

We transform this causal model into a structurally indistinguishable or isomorphic causal model. To do this, negate both sides of the direct non-redundant regularity. Then substitute  $C$  by  $F$ ,  $A$  by  $\neg D$ , and  $E$  by  $\neg E$ . The result is the isomorphic causal model:

$E \leftrightarrow \neg F \wedge D$
$F, \neg D, \neg E$

And indeed,  $\neg E$  is ‘overdetermined’ by  $F$  and  $\neg D$ . This causal model can be represented by Figure 3.




Figure 3: Bogus Prevention.

Here is an example for this structure. Poisonings  $D$  are type causes of deaths  $E$ , and antidotes  $F$  are type preventers of deaths  $E$ . In the scenario under consideration, let’s say, there is an assassin, a potential target, and her bodyguard. The assassin refrains from poisoning target’s coffee  $\neg D$ , and her bodyguard puts antidote in her coffee  $F$ . Target survives, of course. Since target’s coffee is not poisoned in the first place, there is no danger at

all that she dies. The prevention by bodyguard's antidote is *bogus*. And so bodyguard's putting the antidote in her coffee is arguably no cause of her survival (Hiddleston, 2005; Hitchcock, 2007).

Recall that there is no structural difference for simple causal model accounts between  $F$  in the scenario of bogus prevention and  $C$  in the scenario of overdetermination. However, our causal judgments differ. We judge  $C$  to be a cause of  $E$  in the overdetermination scenario, while we do not judge  $F$  to be a cause of  $\neg E$  in the bogus prevention scenario.

Simple causal model accounts of causation—like the accounts of Hitchcock (2001), Halpern and Pearl (2005), and Andreas and Günther (2021b, 2021c) for example—cannot distinguish between  $F$  and  $C$  in the isomorphic causal models:  $C$  counts as a cause iff  $F$  does. This means simple causal model accounts must incorrectly classify  $F$  as a cause in the bogus prevention scenario if they correctly classify  $C$  as a cause in the overdetermination scenario. This is a problem indeed if we take our pre-theoretic understanding of the causal relation seriously.

Our preliminary theory of causation is a simple causal model account, and so is likewise susceptible to the problem of isomorphic causal models. Hitchcock (2007), Hall (2007), Halpern and Hitchcock (2015), and Halpern (2015) all aim to solve the problem by taking into account default or normality considerations. The underlying idea is that the status of genuine causes depends on being deviant from what is normal (Beebe, 2004; McGrath, 2005). On this view, genuine effects are brought about by causes that are more deviant from normality than its non-actual alternatives. Gallow (2021) goes even further by elevating the transmission of deviancy to the mark of causation: an event is a cause in virtue of transmitting its deviancy to its effect.

We also aim to resolve the problem of isomorphic causal models by a condition of deviancy. But what is deviancy? For now, we follow Gallow (2021, p. 54) in saying that *prima facie* an occurring event is more deviant than its absence. But this is just a first approximation. The question of what constitutes deviancy is more intricate. We will discuss it later.

To amend our theory by a deviancy condition, we need the notion of a root proposition. We say a proposition  $F$  of particular fact is a root proposition

in a causal model  $\langle \mathcal{L}, \mathcal{F} \rangle$  iff there is no direct non-redundant regularity for  $F$  or  $\neg F$  in  $\mathcal{L}$ . Finally, let  $\mathcal{F}_r$  denote the set of root propositions in  $\mathcal{F}$ .

We amend now our theory by a condition of deviancy. The condition applies to all genuine causes for the effect under consideration and to all root propositions which are not in the maximised minimal set  $\mathcal{F}'$  of propositions of particular fact that are jointly sufficient for the effect in the presence of the direct non-redundant regularities. It says that all such genuine causes and root propositions have to be more deviant than their negation. To be precise, we require the following condition in addition to conditions (1)-(3):

- (4) for all  $F \in \mathcal{F} \setminus \mathcal{F}'$ , if  $F \equiv C$  or  $F \in \mathcal{F}_r$ , then  $F$  is more deviant than  $\neg F$ .

The deviancy condition (4) says that, for  $c$  to be a cause of  $e$ , the proposition  $C$  and each root proposition  $F$  that is not in the maximised minimal set  $\mathcal{F}'$  of propositions of particular fact must be more deviant than its respective negation. On the amended theory, causation is thus understood as deviant forward-directed inferability along direct non-redundant regularities.

How does our amended theory treat bogus prevention? Well, we know from the isomorphism between overdetermination and bogus prevention that  $f$  counts as a cause of the non-occurrence of any event  $e$  of type  $E$ —or simply  $e$ 's absence—on our preliminary theory for  $\mathcal{F}' = \emptyset$ . However, the deviancy condition is violated.  $\neg D$  is a root proposition of particular fact that is not in  $\mathcal{F}' = \emptyset$ , but  $\neg D$  is less deviant than  $D$ . Therefore,  $f$  is not a cause of  $e$ 's absence.  $d$ 's absence is normal and so is likewise no cause of  $e$ 's absence. Or so says our amended theory.

Bogus prevention illustrates the underlying rationale of condition (4). The root propositions in  $\mathcal{F} \setminus \mathcal{F}'$  are causes or backup causes if they are more deviant than their negations. Otherwise they would remain in  $\mathcal{F}'$  in virtue of its maximality. In the bogus prevention scenario,  $d$ 's absence and  $f$ 's occurrence are no causes because of the normality of the root proposition  $\neg D$ . The normality of non-occurrence of any event of type  $D$  entails that  $d$ 's absence and  $f$ 's occurrence are no causes.

Like the simple causal model accounts, Baumgartner's (2013) regularity theory wrongly says that  $d$ 's absence and  $f$ 's occurrence are causes of  $e$ 's

absence (p. 105). He prefers this result over a relativization of causation to normality or typicality (p. 106). This is where we fundamentally part ways. We take our pre-theoretic understanding of the causal relation in the bogus prevention scenario at face value.

This being said, Baumgartner aims for a proof of concept that his regularity theory could be enriched by a notion of normality or typicality—much like simple causal model accounts (p. 106). And indeed, he shows how his regularity theory so enriched can account for the bogus preventer  $f$  by disregarding an atypical but empirically possible situation—possible according to the type-level structure of bogus prevention.

He assumes that the least typical empirical possibility is the one in which some token events of type  $D$  and of type  $F$  occur. By disregarding this empirical possibility, a token event of type  $E$  occurs just in case one of type  $D$  occurs. The non-redundant regularity simplifies to  $E \leftrightarrow D$ .  $\neg F$  does not figure any longer in the non-redundant regularity. And so  $f$  ceases to be a cause of  $e$ 's absence on his enriched theory. But  $d$ 's absence remains a cause of  $e$ 's absence. And this goes against our commonsense judgment: the normal absence of poison does not cause target's survival.

## 4.2 Omissions

Omissions pose another problem for many theories of causation. In a scenario of omission, an event fails to occur and so another event occurs. However, had the event occurred, it would have prevented the other event from occurring. The basic structure of omissions can be represented by Figure 4.




Figure 4: Omission.



An event  $c$  of type  $C$  occurs and brings about an event  $e$  of type  $E$ . No event  $f$  of type  $F$  occurs. However, had an event  $f$  of type  $F$  occurred, it would have prevented  $e$  from occurring. Here is the causal model for the scenario of omission.

$E \leftrightarrow \neg F \wedge C$
$\neg F, C, E$

Relative to this causal model,  $f$ 's absence is not a cause of  $e$ , given the above convention about deviancy. For this to be seen, observe that conditions (2) and (3) are only satisfied for  $\mathcal{F}' = \{C\}$ . But then condition (4) is violated:  $\neg F \in \mathcal{F} \setminus \mathcal{F}'$  and  $\neg F$  is a root proposition, and yet  $\neg F$  is less deviant than  $F$ .<sup>4</sup>

Indeed, many omissions are no causes. Putin's failure to water my plant, for example, did not cause it to dry up and die. However, some omissions intuitively *do* count as causes. My neighbour promised me to water my plant, but she didn't and it died. Here my neighbour's failure to water my plant should count as a cause of its death (McGrath, 2005). Our theory can capture this phenomenon if we refine our notion of deviancy.

We have said that *prima facie* an occurring event is more deviant than its absence. We say now in addition that the absence  $\neg A$  of any event of type  $A$  is more deviant than an event of type  $A$  if  $\neg A$  violates a norm that is active in the scenario under consideration (Beebe, 2004; Andreas et al., 2022). My neighbour's omission to water my plant is an absence that violates the active norm of promise-keeping. My neighbour deviated from this norm and so her omission is more deviant than its negation. Our theory says then that my neighbour's failure to water my plant is a cause of the plant's death. Putin, by contrast, did not promise to water my plant. His omission is thus less deviant than his watering my plant, and so does not count as a cause of my plant's death. Or so says our theory.

We have illustrated how condition (4) helps overcome the problem of isomorphic causal models and how it accounts for simple scenarios of omission. As to the latter, deviant omissions are genuine causes, non-deviant ones are not. In general, a genuine cause is deviant and allows

to infer its effect in a forward-directed way along direct non-redundant regularities.

Baumgartner's (2013) theory says that all omissions of  $f$  in the above scenario are causes of  $e$ . There is only one set of actual conditions that are jointly sufficient for  $e$  to occur in the presence of the non-redundant regularity:  $\{\neg F, C\}$ . What does his theory say when we enrich it by some typicality ranking?

Observe that the non-redundant regularity of the omission scenario is of the same form as the one of bogus prevention. There Baumgartner assumed that the empirical possibility in which token events of type  $F$  and type  $C$  occur is less typical than any other. By disregarding this empirical possibility, an event of type  $E$  occurs just in case one of type  $C$  occurs. The non-redundant regularity of the scenario simplifies to  $E \leftrightarrow C$ . And so the default omission of  $f$  ceases to be a cause of  $e$  on his enriched theory. So far so good.

There is, however, work left. Baumgartner has not yet said how normality, typicality, or deviancy are supposed to figure in the confines of his theory in general. It remains unclear how his enriched regularity theory applies to arbitrary causal scenarios.

### 4.3 Entanglements

Another problem comes from causes that are entangled with one another. We first illustrate the problem and then refine our theory in response. Consider a structure, where two type causes  $A$  and  $C$  are necessary for an effect  $E$  to occur. Moreover,  $A$  and  $C$  are entangled: one of them is a type cause of the other. And let's say in the actual scenario a token event of each type occurs. Figure 5 graphically represents this scenario, where the thick circle around  $E$  means that both  $A$  and  $C$  are necessary for  $E$ .




Figure 5: Subcause.

We say that  $C$  is here a *subcause* of  $A$ . Here is the causal model for this subcause scenario.

$C \leftrightarrow A$
$E \leftrightarrow C \wedge A$
$A, C, E$

Entangled causes are tightly related. Here, the subcause  $C$  depends directly and exclusively on the ‘supercause’  $A$ . Given the direct non-redundant regularities, the occurrence of some token event of type  $E$  is determined by whether or not some token event of type  $A$  occurs. In this sense, the type-cause  $C$  is subordinate to the type-cause  $A$ .

Our present theory of causation does not count the token event of type  $C$  as a cause of the token event of type  $E$ . For this to be seen, note that  $\langle \mathcal{L}, \mathcal{F}' \rangle \not\models E$  only for  $\mathcal{F}' = \emptyset$ . But then  $\langle \mathcal{L}, \emptyset \rangle[\mathcal{F}'][\{C\}] \not\models E$ . However, this seems to be the wrong verdict.

Why does our present theory of causation fail for entangled causes? Well, it seems that the tight connection between entangled causes like  $A$  and  $C$  is the culprit. A solution is thus to cut this tight connection between entangled causes by removing non-redundant regularities from the causal model. In the subcause scenario, this means we could cut the tight connection between the subcause  $C$  and its supercause  $A$  by removing the direct non-redundant regularity  $C \leftrightarrow A$ . We thus generalize our theory of causation as follows.

**Definition 2.** Let  $\langle \mathcal{L}, \mathcal{F} \rangle$  be a causal model such that  $\mathcal{F} \models \mathcal{L}$ .  $c$  is a cause of  $e$  relative to  $\langle \mathcal{L}, \mathcal{F} \rangle$  iff there is a set  $\mathcal{L}' \subseteq \mathcal{L}$  and a set  $\mathcal{F}' \subseteq \mathcal{F}$  such that all of the following conditions are satisfied:

- (1)  $\langle \mathcal{L}, \mathcal{F} \rangle \models C \wedge E$ .
- (2)  $\langle \mathcal{L}', \mathcal{F}' \rangle \not\models E$ , and there is no  $\mathcal{F}''$  so that  $\mathcal{F}' \subset \mathcal{F}'' \subseteq \mathcal{F}$  and  $\langle \mathcal{L}', \mathcal{F}'' \rangle \not\models E$ .
- (3)  $\langle \mathcal{L}', \emptyset \rangle[\mathcal{F}'][\{C\}] \models E$ .
- (4) For all  $F \in \mathcal{F} \setminus \mathcal{F}'$ , if  $F \equiv C$  or  $F \in \mathcal{F}_r$ , then  $F$  is more deviant than  $\neg F$ .

This proper generalization of our amended theory solves the subcause scenario. For this to be seen, take  $\mathcal{L}' = \{E \leftrightarrow C \wedge A\}$  and  $\mathcal{F}' = \{A\}$ . Well then, condition (2) is satisfied, and so are conditions (3) and (4). Hence, the occurring token event  $c$  of type  $C$  is a cause of the occurring token event  $e$  of type  $E$ , as desired.<sup>5</sup>

#### 4.4 Non-Transitivity

A final challenge for the current theory arises from causal scenarios which suggest that causation is not transitive. The transitivity of causation means this: whenever a token event  $c$  of type  $C$  is a cause of one of type  $A$  and the one of type  $A$  is a cause of a token event  $e$  of type  $E$ , then  $c$  is a cause of  $e$ . It seems often plausible to judge  $c$  a cause of  $e$  if you judge  $c$  a cause of  $a$  and  $a$  a cause of  $e$ . However, several scenarios have been put forth which suggest that our causal judgments are not transitive (McDermott, 1995; Lewis, 2000; Paul, 2000).

One of the examples against transitivity is due to Hitchcock (2001, p. 276). A boulder is dislodged and rolls toward a hiker. The hiker sees the boulder coming and ducks, so that she does not get hit by the boulder. If the hiker had not ducked, however, the boulder would have hit her.

The boulder scenario seems to show that there are cases where causation is not transitive: the dislodged boulder causes the ducking of the hiker, which in turn causes the hiker to remain untouched by the boulder. But it is counterintuitive to say that the dislodging of the boulder causes the hiker to remain unscathed. Unlike other accounts, our theory does not rely on transitivity to handle certain causal scenarios. We are thus free to deny that our causal judgment is invariably transitive.

The structure of the boulder scenario can be represented by Figure 6 (Gallow, 2021, p. 53).




Figure 6: Short Circuit.

Hall (2007, p. 36) calls this structure a *short circuit*: the boulder’s dislodgement ( $f$ ) threatens to hit the hiker by a rolling boulder ( $b$ ), and at the same time provokes an action—the ducking ( $d$ )—that prevents this threat from being effective: no token event  $e$  of type  $E$  occurs.

The token event  $f$  should not count as a cause of the absence of any event  $e$  of type  $E$ , because  $f$  creates *and* cancels the threat to bring about some token event  $e$  of type  $E$  (Paul and Hall, 2013, p. 216). Our current analysis, however, says that  $f$  is a cause of  $e$ ’s absence. To see this, consider the causal model of the boulder scenario:

$B \leftrightarrow F$
$D \leftrightarrow F$
$E \leftrightarrow B \wedge \neg D$
$F, B, D, \neg E$

Take  $\mathcal{L}' = \{D \leftrightarrow F, E \leftrightarrow B \wedge \neg D\}$  and  $\mathcal{F}' = \{B\}$ . Then condition (2) is satisfied, and so are the conditions (3) and (4).

What went wrong? Well, it seems that  $f$  comes out as a cause of  $e$ ’s absence because we disregard that  $F$  is type-causally connected to  $E$  via  $B$ . By removing the direct non-redundant regularity  $B \leftrightarrow F$ , we lose the information that  $f$  creates a threat to bring about  $e$ ’s occurrence. But, intuitively, this information seems to be crucial to determine whether  $f$  is a cause of  $e$ ’s absence. And so it seems to be unwarranted to neglect one of the type-causal paths from  $F$  to  $E$ .

The generalization of our theory to handle scenarios of entangled causes overshoots. Removing arbitrary direct non-redundant regularities from  $\mathcal{L}$  can lead to counterintuitive results, as the boulder scenario illustrates. In response, we put a constraint on which direct non-redundant regularities can be removed. The idea is this: when testing whether a token event  $c$  of type  $C$  is a cause of a token event  $e$  of type  $E$ , the causal paths from  $C$  to  $E$  need to remain intact.

To implement the idea that constrains the removal of direct non-redundant regularities, we introduce some terminology. Recall that  $A \leftrightarrow \bigvee_i \bigwedge \mathcal{A}_i$  is the direct non-redundant regularity of  $A$ . We say  $A$  is a *child* variable of the *parent* variables appearing as subformulas in some  $\mathcal{A}_i$ . Let  $B$  be one of the parent variables appearing in some  $\mathcal{A}_i$ .  $A$  is then one of its first descendants. The child variables of  $A$  are the child variables of one of  $B$ 's child variables and so are among  $B$ 's second descendants. And so on. The descendants of  $B$  are the variables you can reach by following the child relation. In general, the descendants of some variable  $B$  are the variables in the transitive closure of the child relation starting from  $B$ . Finally, let  $B$  be a proposition of the form  $D$  or  $\neg D$ , where  $D$  is a propositional letter. We then say that the descendants of  $B$  are all the variables (of the causal model under consideration) which are descendants of the variable  $D$  of which  $B$  is a proposition.

For our final theory of causation, we require the following condition in addition to conditions (1)-(4):

- (5) For all descendants  $A$  of  $C$ , the direct non-redundant regularity of  $A$  is in  $\mathcal{L}'$ .

Condition (5) ensures that the causal paths starting from a candidate cause  $C$  remain in  $\mathcal{L}'$ . On our final theory, causation is understood as deviant forward-directed inferability along the causal paths of direct non-redundant regularities from cause to effect. We show now how this condition helps to handle scenarios that challenge the transitivity of causation.

#### 4.4.1 Short Circuit

Let us, first, reconsider the boulder scenario of Figure 6.  $f$  does no longer count as a cause of  $e$ 's absence. For this to be seen, observe that all variables are descendants of  $F$ . Condition (5) thus prohibits to remove any direct non-redundant regularity from  $\mathcal{L}$ . As a consequence, there is no  $\mathcal{L}'$  and  $\mathcal{F}'$  so that condition (2) is satisfied. Even for  $\mathcal{F}' = \emptyset$ , the direct non-redundant regularities entail  $\neg E$ .

By contrast,  $d$  counts as a cause of  $e$ 's absence. For this to be seen, observe that  $F$  is not a descendant of  $D$ . Hence, the direct non-redundant regularity  $D \leftrightarrow F$  can be removed from  $\mathcal{L}$ . Take  $\mathcal{L}' = \{B \leftrightarrow F, E \leftrightarrow B \wedge \neg D\}$  and  $\mathcal{F}' = \{F, B\}$ .  $\langle \mathcal{L}', \mathcal{F}' \rangle$  does not entail  $\neg E$ , and  $\mathcal{F}'$  is maximal: any strict superset of  $\mathcal{F}'$  would entail  $\neg E$  in the presence of the direct non-redundant regularities in  $\mathcal{L}'$ . But, of course,  $\langle \mathcal{L}', \emptyset \rangle[\mathcal{F}'][\{D\}] \models \neg E$ .

On Baumgartner's (2013) theory, the type-level structure of the boulder scenario is *empirically equivalent* to the type level structure  $\mathcal{L} = \{B \leftrightarrow F, D \leftrightarrow F, \neg E \leftrightarrow F \vee \neg F\}$ . A situation in which  $B$  and  $\neg D$  are true is impossible according to this and the original type-level structure of the boulder scenario. In both, there are only two empirically possible situations over the four variables:  $\{F, B, D, \neg E\}$  and  $\{\neg F, \neg B, \neg D, \neg E\}$ . The type-level structure of the boulder scenario only allows for empirical possibilities, in which no token event of type  $E$  occurs. And so it is empirically equivalent to a type-level structure including the indirect non-redundant regularity  $\neg E \leftrightarrow F \vee \neg F$  but without the regularity  $E \leftrightarrow B \wedge \neg D$ . This shows, so argues Baumgartner, that the latter regularity is *empirically redundant* (pp. 101-4). After all, the indirect non-redundant regularity says that a token event of type  $E$  never occurs.

The two type-level structures are empirically equivalent: the situations compatible with the respective non-redundant regularities coincide. However, we think that the type-level structure of the boulder scenario is a conceptual possibility for which a theory of causation should be able to account for. And while Baumgartner's theory correctly says that  $f$  is not a cause of  $e$ 's absence, it also says that  $d$  is not a cause of  $e$ 's absence. The latter is wrong on our understanding of the causal relation. After all, we judge the hiker's ducking to be a cause of her remaining untouched by the

boulder.

#### 4.4.2 Extended Double Prevention

In a scenario of double prevention, an event prevents another event which—had it occurred—would have prevented a third event. Hall (2004, p. 247) presents an extension of a simple scenario of double prevention, which we call extended double prevention. The structure of this scenario can be represented by Figure 7:




Figure 7: Extended Double Prevention.

Here the event of type *C* double prevents the event of type *E* because *C* prevents an event of type *D* from occurring that—had it occurred—would have prevented the occurring event of type *E*. In the present scenario, however, the occurring events of type *B* and of type *C* have an occurring common cause of type *F*, and the subgraph *F – B – C – D* is a short circuit. *F* starts a process via *B* that threatens to prevent *E*. At the same time, *F* initiates another process via *C* that prevents the threat. *F* cancels its *own* threat—the threat via *B*—to prevent *E*. Paul and Hall (2013, p. 216) argue, among others, that *F* is not a cause of *E* in extended double prevention.

Here is the causal model for the scenario of extended double prevention.



$B \leftrightarrow F$ $C \leftrightarrow F$ $D \leftrightarrow B \wedge \neg C$ $E \leftrightarrow A \wedge \neg D$
$A, F, B, C, \neg D, E$

Relative to this causal model,  $f$  is not a cause of  $e$ . For this to be seen, observe first that all non-root variables are descendants of  $F$ . Condition (5) prohibits to remove any direct non-redundant regularity from  $\mathcal{L}$ . Any set  $\mathcal{F}'$  of particular facts that is supposed to satisfy condition (3) needs to contain  $A$ ; otherwise  $\langle \mathcal{L}', \emptyset \rangle[\mathcal{F}'][\{F\}] \not\models E$ . However, if  $A \in \mathcal{F}$ , then condition (2) is violated:  $\langle \mathcal{L}', \mathcal{F}' \rangle \models E$ .<sup>6</sup>

On Baumgartner's (2013) theory, the type-level structure of extended double prevention is empirically equivalent to the type-level structure  $\mathcal{L} = \{B \leftrightarrow F, C \leftrightarrow F, E \leftrightarrow A\}$  (see pp. 101-2). As a consequence,  $a$  is the only cause of  $e$ . The double preventer  $c$ , in particular, does not count as a cause of  $e$ .

#### 4.4.3 Modified Extended Double Prevention

In Figure 8, Paul and Hall (2013, p. 198-9) ask us to consider a slight modification of extended double prevention by 'adding a non-occurring event of type  $G'$ '.




Figure 8: Modified Extended Double Prevention.

Paul and Hall argue that ‘the original verdict stands:’  $f$  is not a cause of  $e$  because no event of type  $G$  occurs and this non-occurrence does not make any difference (p. 199). Here is the causal model for the scenario of modified extended double prevention.

$B \leftrightarrow G \vee F$
$C \leftrightarrow F$
$D \leftrightarrow B \wedge \neg C$
$E \leftrightarrow A \wedge \neg D$
$A, \neg G, F, B, C, \neg D, E$

Relative to this causal model,  $f$  is not a cause of  $e$ . For this to be seen, observe first that all non-root variables are descendants of  $F$ . Condition (5) prohibits to remove any direct non-redundant regularity from  $\mathcal{L}$ . There are only two cases in which  $\langle \mathcal{L}', \mathcal{F}' \rangle$  meets condition (2). First, take  $\mathcal{F}' = \{\neg G, \neg D\}$ . But then condition (3) is violated:  $\langle \mathcal{L}', \emptyset \rangle[\mathcal{F}'][\{F\}] \not\models E$ .

Second, take  $\mathcal{F}' = \{A, B\}$ . In this case,  $\neg G$  is in  $\mathcal{F} \setminus \mathcal{F}'$  and  $\neg G \in \mathcal{F}_r$ . Since  $\neg G$  is less deviant than  $G$ , condition (4) is violated. We have shown that  $f$  is not a cause of  $e$  in the scenario of modified extended double prevention.

On Baumgartner’s (2013) theory,  $f$  counts as a cause of  $e$ . The type-level structure of modified extended double prevention is empirically accurate,

and yet his theory delivers the wrong result. This being said, his enriched theory could obtain the correct result, if the empirical possibility, in which  $a$  and  $f$  occur but  $g$  does not, is less typical than the other empirical possibilities. But this condition is hard to assess because Baumgartner did not yet tell us what makes one empirical possibility more or less typical than another.

#### 4.4.4 Isomorphic Modified Extended Double Prevention

Paul and Hall (2013, pp.198-9) consider a scenario that is isomorphic to modified extended double prevention. They represent it by Figure 9.




Figure 9: Isomorphic Modified Extended Double Prevention.

In this scenario, the occurring event of type  $E$  requires that both an event of type  $A$  and one of type  $D$  occurs. Paul and Hall (2013, p. 199) argue that this time  $c$  clearly is a cause of  $e$ . Here is the causal model for the scenario of isomorphic modified extended double prevention.

$B \leftrightarrow G \wedge \neg C$
$H \leftrightarrow C$
$D \leftrightarrow B \vee H$
$E \leftrightarrow A \wedge D$
$A, G, C, \neg B, H, D, E$

Relative to this causal model,  $c$  is a cause of  $e$ . Take  $\mathcal{L}' = \mathcal{L}$  and  $\mathcal{F}' = \{A, \neg B\}$ . Then condition (2) is satisfied:  $\langle \mathcal{L}', \mathcal{F}' \rangle \not\models E$ , and  $\mathcal{F}'$  is maximal: any strict superset of  $\mathcal{F}'$  would entail  $E$  in the presence of the direct non-redundant regularities. And condition (3) is satisfied as well:  $\langle \mathcal{L}', \emptyset \rangle [\mathcal{F}'] [\{C\}] \models E$ . Condition (4) is satisfied as well:  $G$  is in  $\mathcal{F} \setminus \mathcal{F}'$  and  $G \in \mathcal{F}_r$ , but  $G$  is also more deviant than  $\neg G$ . And condition (5) is trivially satisfied since we did not remove any direct non-redundant regularity. We have shown that  $c$  is a cause of  $e$  in the scenario of isomorphic modified extended double prevention.<sup>7</sup>

#### 4.4.5 Switch

Switching scenarios are paradigmatic for causal scenarios where our causal judgments are not transitive. In switching scenarios, some occurring event of type  $F$  helps determine the causal path by which another event is brought about. Crucially, the other event would also occur via an alternative causal path if no event of type  $F$  had occurred.

To make it more concrete, consider a story provided by Hall (2000, p. 205). Flipper is standing by a switch in the railroad tracks. A train approaches in the distance. She flips the switch, so that the train travels down the right track, instead of the left. Since the tracks reconverge up ahead, the train arrives at its destination all the same. The commonsense judgment is that flipping the switch is not a cause of the train's arrival—even though flipping the switch is a cause of the train's travelling on the right track, and the train's travelling on the right track is a cause of the train's arrival (Paul and Hall, 2013, p. 232).

The structure of this switch scenario can be represented by Figure 10:




Figure 10: Switch.

The flipping of the switch  $f$  causes the train to travel on the right track  $r$  and prevents the train from travelling on the left track  $l$ . And the travelling on the right track  $r$  causes the train to arrive at its destination  $e$ . However, the flipping of the switch  $f$  arguably is not a cause of the train's arrival  $e$ . Here is the causal model for the switch scenario.

$L \leftrightarrow \neg F$
$R \leftrightarrow F$
$E \leftrightarrow L \vee R$
$F, \neg L, R, E$

Relative to this causal model,  $f$  is not a cause of  $e$ . For this to be seen, observe first that all variables are descendants of  $F$ . Condition (5) thus prohibits to remove any direct non-redundant regularity from  $\mathcal{L}$ . But then there is no  $\mathcal{F}'$  so that condition (2) is satisfied. Even for  $\mathcal{F}' = \emptyset$ , the direct non-redundant regularities entail  $E$ .

By contrast,  $f$  is a cause of  $r$ . Take  $\mathcal{L}' = \mathcal{L}$  and  $\mathcal{F}' = \{E\}$ . Condition (2) is then satisfied:  $\langle \mathcal{L}', \mathcal{F}' \rangle \not\models R$ , and  $\mathcal{F}'$  is maximal: any strict superset of  $\mathcal{F}'$  would entail  $R$  in the presence of the direct non-redundant regularities. The other conditions are trivially satisfied.

Likewise,  $r$  is a cause of  $e$ . Take  $\mathcal{L}' = \{L \leftrightarrow \neg F, E \leftrightarrow L \vee R\}$  and  $\mathcal{F}' = \{F, \neg L\}$ . Condition (2) is then satisfied because  $\langle \mathcal{L}', \mathcal{F}' \rangle \not\models E$ , and  $\mathcal{F}'$  is maximal: any strict superset of  $\mathcal{F}'$  would entail  $E$  in the presence of the direct non-redundant regularities. The other conditions are trivially satisfied.

The representation of switching scenarios is somewhat controversial. Our theory of causation delivers the desired results for the ‘basic’ switch discussed by Paul and Hall (2013, p.232), for the more realistic switches discussed by Hitchcock (2009, p.395-6), and also a switch discussed by Halpern (2016, p.72). We leave it to the reader to verify that our final theory of causation delivers also the desired results for all of the previously mentioned scenarios.

On Baumgartner’s (2013) theory, the type-level structure of our simple switch is empirically equivalent to  $\mathcal{L} = \{L \leftrightarrow \neg F, R \leftrightarrow F, E \leftrightarrow F \vee \neg F\}$ . And so there is no cause of  $e$ .  $r$  does, in particular, not count as a cause of  $e$ . This being said, Baumgartner’s theory delivers the correct verdicts for Paul and Hall’s basic switch.

## 5 Reductiveness

We have proposed a regularity theory of causation. On our preliminary theory, a cause is each member of any maximised minimal set of actual conditions which are jointly sufficient in a forward-directed way for the effect to occur in the presence of the direct non-redundant regularities. Causation so understood is forward-directed inferability along direct non-redundant regularities. The direct non-redundant regularities are material bi-implications obtained from true propositions of particular fact and minimization procedures. Causation is thus reduced to true propositions of particular fact.

However, our causal relation is so far relative to a set of selected variables. For the true non-redundant regularities are direct relative to a variable set (see Section 2). Extending the variable set under consideration may render direct non-redundant regularities indirect. Even worse, an extension of the variable set may render a non-redundant regularity redundant (Baumgartner, 2013, pp.93-4). This poses the metaphysical question how we can identify causation in the world, as opposed to causation in a model containing a limited amount of variables.

Our answer has two parts. First, causation in the world requires each non-

redundant regularity to be *stable*—to remain non-redundant under any apt extension of the variable set. An extension of a variable set is *apt* just in case the additional variables do not introduce dependences among the variables that are stronger than causation, such as logical or mereological relations, supervenience, or grounding. Stable regularities are not relative to a particular variable set because they remain non-redundant under *any* apt extension.

Second, we say that a stable regularity  $\bigvee_i \mathcal{A}_i \leftrightarrow E$  is *metaphysically direct* iff there are no stable regularities  $\bigvee_j \mathcal{B}_j \leftrightarrow C$  and  $\bigvee_k \mathcal{C}_k \leftrightarrow E$ , where the same variable appears as subformula in some  $\mathcal{A}_i$  and some  $\mathcal{B}_j$ , while  $C$  appears as subformula in some  $\mathcal{C}_k$ . Metaphysically direct regularities inherit from stable regularities that they are not relative to a particular variable set. And they remain metaphysically direct under each apt extension of the variable set.

Metaphysically direct regularities allow us to define causation without model relativity. We say  $c$  is a cause of  $e$  iff there is a set  $\mathcal{L}' \subseteq \mathcal{L}$  of metaphysically direct regularities and a set  $\mathcal{F}' \subseteq \mathcal{F}$  of propositions of particular fact such that our conditions for causation are satisfied. This way to avoid model relativity may well lead to very fine-grained metaphysically direct regularities. But it is also a way to reduce causation in the world to true propositions of particular fact.

We have amended our preliminary theory by a condition of deviancy. Causation so understood is deviant forward-directed inferability along direct non-redundant regularities. We have said that occurring events are *prima facie* more deviant than their absences (Section 4.1). We could have said instead that occurring events have a metaphysically different status from their absences. This metaphysical postulate would have saved that our amended theory reduces causation to true propositions of particular fact. And it still solves the problem of isomorphic causal models and accounts for the non-causal status of many omissions. However, some omissions are judged to be causes.<sup>8</sup> The reason seems to be that the latter, but not the former, violate some norm. This is why we said in addition that an absence is more deviant than an event if the absence violates a norm active in the scenario under consideration. Causation is thus reduced to true propositions of particular fact and deviancy from norms.

Blanchard and Schaffer (2017) criticise the distinction between default and deviant events. Consider the causal model for bogus prevention in Section 4.1. They claim that this causal model is not apt for the scenario where no poison is administered because it lacks essential structure: a variable for whether or not the administered antidote neutralizes some poison. The causal model so enriched is not isomorphic to overdetermination and the administration of antidote is no cause of target's survival. This solution, so they conclude, shows that we do not need a default-deviant distinction but apt causal models.

Indeed, the metaphysically direct regularities presumably entail that a neutralization event occurs only if a poisoning event and an administration of antidote event occurs. One must wonder, however, whether there are no causal scenarios that share the simple structure of bogus prevention. It seems at least conceptually possible that  $E \leftrightarrow \neg F \wedge D$  is a true metaphysically direct regularity of some scenario, in which an event of type  $F$  occurs, but no events of types  $D$  and  $E$ . But then the bogus preventer of type  $F$  is not a cause of the absence of type  $E$ . Blanchard and Schaffer still need to explain this verdict.

There is, of course, more to say about deviancy, norms, and the distinction between occurring events and absences. Hitchcock and Knobe (2009) argue that our judgement of causation is influenced by statistical and moral norms, as well as norms of proper functioning. However, they dodge the questions of what norms are in the first place and when they are active in a certain scenario. These are open problems. Another is that the distinction between occurring events and absences is not always clear cut. Is, for example, the death of a plant an occurring event, or rather the absence of continued living? It seems that this depends on the involved norms. Is it normal that plants continue to live? Or is being alive the deviant state to being dead? The distinction between occurring events and absences seems to be related to the deviancy from norms. But an investigation of this relation must await another occasion.

We have further amended our theory by a condition which ensures that the causal paths from a cause to its effect remain intact. Causation so understood is deviant forward-directed inferability along the causal paths of direct non-redundant regularities from cause to effect. The further



amendment helps to handle scenarios that challenge the transitivity of causation, like short circuits and switches. Our final regularity theory thus accounts for our judgments on a wide range of causal scenarios—wider than Baumgartner’s (2013) theory. We summarize the results in the Conclusion.

Except for perhaps norms, our regularity theory has no need for any modal notions. It does neither rely on a notion of nomic sufficiency (Hausman, 1998), nor on a notion of epistemic inferability (Andreas and Günther, 2019; Andreas and Günther, 2021a, 2021b). It does likewise not rely on counterfactuals. The latter distinguishes our theory from many contemporary theories of causation (Ramachandran, 1997; Lewis, 2000; Yablo, 2002; Woodward, 2003; Hall, 2004, 2007). These counterfactual theories are not reductive if they rely on antecedently given causal structures—as they often do in the form of given causal models (Hitchcock, 2001, 2007; Hiddleston, 2005; Halpern and Pearl, 2005; Halpern, 2016; Andreas and Günther, 2021b; Gallow, 2021). Hitchcock (2001) and Halpern and Hitchcock (2010) offer some guidelines on what constitutes an apt causal model. But the question is not fully answered (Blanchard and Schaffer, 2017).

Our regularity theory, by contrast, has no need for discerning apt from non-apt causal models. It tells us what causes what relative to a particular causal model understood as a tuple of true direct non-redundant regularities and true propositions of particular fact. And we have proposed a way how we could avoid this model-relativity if we had the metaphysically direct regularities at our disposal.

## 6 Conclusion

We have put forth a regularity theory of causation. The theory says, in essence, that causation is deviant forward-directed inferability along the causal paths of direct non-redundant regularities. The verdicts of our theory aligns with our pre-theoretic understanding of the causal relation in the considered set of scenarios. In virtue of its deviancy condition, the theory can account for the problem of isomorphic causal models and the problem of omissions. It can, furthermore, account for entangled causes

and scenarios that challenge the transitivity of causation. The following table summarizes the results of our regularity theory in comparison to Baumgartner's (2013).

Causes of $e$ or $\neg e$	Baumgartner (2013)	Our Regularity Theory
Overdetermination	$c, a$	$c, a$
Preemption	$c, d$	$c, d$
Bogus Prevention	$(f), \neg d$	–
Default Omission	$(\neg f), c$	$c$
Deviant Omission	$(\neg f), c$	$\neg f, c$
Subcause	$c, a$	$c, a$
Short Circuit	–	$d$
Extended DP	$a$	$c, a$
Mod. Extended DP	$(f), c, \neg d, a$	$c, a$
Iso. Mod. Extended DP	$c, h, d, a$	$c, h, d, a$
Simple Switch	–	$r$

The table shows the causes of  $e$  or  $e$ 's absence—here abbreviated by  $\neg e$ —in the respective causal scenario. Expressions like  $(f)$  mean that a suitable notion of typicality or deviancy could undo the causal status of the token event  $f$  on Baumgartner's theory. 'DP' stands for Double Prevention.

In the scenario of bogus prevention, there is no poisoning and target survives. The absence of poison in target's coffee, here  $\neg d$ , is as much a cause of her survival as the absence of someone stabbing her or the absence of someone shooting her: the absences are not deviant and so are no causes. Bodyguard's putting antidote into her coffee  $f$  is in this scenario likewise no cause of her survival. Baumgartner can secure the latter but not the former verdict by imposing a certain typicality ranking on bogus prevention.

We have discussed a short circuit scenario. A boulder is dislodged and rolls toward a hiker. The hiker sees the boulder coming and ducks, so that the boulder does not hit her. Had she not ducked, however, the boulder would have hit her. The dislodgement of the boulder  $f$  is no cause of the hiker's remaining unscathed. Our theories agree. The hiker's ducking  $d$  is a cause of the hiker's remaining untouched by the boulder. Our theory says so, but Baumgartner's does not. A similar point applies to our simple

switch. The train travelling on the right track  $r$  causes the train to arrive at its destination. Our theory says so, unlike Baumgartner's. The table shows further differences.

We think our theory aligns with our pre-theoretic judgments about causation, and more so than Baumgartner's. This is no wonder in a sense: our theory aims to be conceptually accurate, whereas Baumgartner's seems to aim for empirical accurateness. And it may well turn out that his theory is empirically more accurate—perhaps scenarios like simple short circuits and simple switches have no empirical causal structure. But even if it should turn out this way, these scenarios seem to be conceptually possible. And so a conceptually accurate theory should still align with our commonsense understanding of causation.

We have set out to propose a regularity theory which reduces causation to true propositions of particular fact and deviancy from norms. And we have come a long way. We have defined direct non-redundant regularities in terms of true propositions of particular fact and minimization procedures. In turn, we have defined causal models in terms of direct non-redundant regularities and propositions of particular fact. The regularity theory we proposed says what causes what relative to such a causal model. Finally, we have made a metaphysical proposal of how to avoid the model-relativity.

Our regularity theory is, however, still incomplete. We haven't said much on what norms are and when events deviate from norms in a given scenario. We have also bracketed the question whether or not norms can be reduced to propositions of particular matter of fact. It seems to us that the relation between the deviancy from norms and the distinction between events and absences, in particular, deserves more attention.

We generally obtain the direction of causation via the direct non-redundant regularities. However, this strategy does not work for scenarios where some type effect has only a single type cause. We have pointed toward solutions for these simplistic but conceptually possible scenarios due to Baumgartner (2013) and Andreas and Günther (202xc). But a comprehensive treatment in the confines of our theory is left for future work.<sup>9</sup>

There is no doubt: more work is to be done. But for now, we may hope to be

one step closer to a regularity theory that grounds the type relation of cause and effect in matters of particular fact and aligns with our understanding of token causation.

Department of Economics, Philosophy and Political Science,  
University of British Columbia

Munich Center for Mathematical Philosophy,  
LMU Munich

## Notes

<sup>1</sup>We set the issue of indeterministic causation aside.

<sup>2</sup>True propositions of particular fact describe actual events and absences—events and absences of the actual world ranging from the most distant past to the farthest future.

<sup>3</sup> $\bigvee_i \wedge \mathcal{A}_i$  is short for  $\bigvee_{i \leq n} \wedge \mathcal{A}_i$ , where  $n$  is the cardinality of  $\mathbf{C}$ . Likewise,  $\bigcup_i \mathcal{A}_i$  is short for  $\bigcup_{i \leq n} \mathcal{A}_i$ .

<sup>4</sup>By contrast,  $c$  is a cause of  $e$ . Take  $\mathcal{F}' = \{\neg F\}$ . Conditions (2) and (3) are then satisfied. And since  $\neg F$  is not in  $\mathcal{F} \setminus \mathcal{F}'$ , condition (4) is trivially met.

<sup>5</sup>Baumgartner's (2013) theory correctly says that both token events of type  $A$  and  $C$ , respectively, are causes of the token event of type  $E$ . Beckers (2021, pp. 1361-3) puts forth a series of six scenarios in order to support his causal model account of causation and to challenge others. The latter four scenarios contain entangled causes. We leave it to the reader to verify that our theory delivers the commonsense results Beckers desires.

<sup>6</sup>By contrast,  $c$  is a cause of  $e$ . Take  $\mathcal{F}' = \{F, B, A\}$  and  $\mathcal{L}' = \mathcal{L} \setminus \{C \leftrightarrow F\}$ . Conditions (1)-(5) are then met.

<sup>7</sup>Baumgartner's (2013) theory correctly says that  $c$  is a cause of  $e$  (see pp. 101-2).

<sup>8</sup>This claim is somewhat controversial. Beebe (2004) argues that omissions are never causes. Dowe (2000, Ch. 6) argues that omissions are no genuine causes, but may figure in true counterfactual claims about genuine causation. We are convinced by McGrath (2005) who argues that the causal status of omissions depends on norms.

<sup>9</sup>In Andreas and Günther (202xa), we further develop the present regularity theory but without aiming for reductivity. We show in this sequel paper that there is another challenge to the reductivity of Baumgartner’s theory and ours. We do not think that this challenge is insurmountable. Unfortunately, we discovered it too late for dealing with it here. A solution must await another occasion.

## References

- Andreas, H., Armgardt, M., and Günther, M. (2022). ‘Counterfactuals for Causal Responsibility in Legal Contexts’. *Artificial Intelligence and Law* .
- Andreas, H. and Günther, M. (2019). ‘On the Ramsey Test Analysis of “Because”’. *Erkenntnis* **84**(6): 1229–1262.
- Andreas, H. and Günther, M. (2021a). ‘A Ramsey Test Analysis of Causation for Causal Models’. *The British Journal for the Philosophy of Science* **72**(2): 587–615.
- Andreas, H. and Günther, M. (2021b). ‘Difference-Making Causation’. *Journal of Philosophy* **118**(12): 680–701.
- Andreas, H. and Günther, M. (2021). ‘Regularity and Inferential Theories of Causation’. In *The Stanford Encyclopedia of Philosophy*, edited by E. N. Zalta, Metaphysics Research Lab, Stanford University. Fall 2021 edn.
- Andreas, H. and Günther, M. (202xa). ‘A Lewisian Regularity Theory’. *Philosophical Studies* .
- Andreas, H. and Günther, M. (202xb). ‘Actual Causation’. *dialectica* .
- Andreas, H. and Günther, M. (202xc). *From Reasons to Causes. An Epochetic Theory of Causation*. Book Manuscript.
- Baumgartner, M. (2013). ‘A Regularity Theoretic Approach to Actual Causation’. *Erkenntnis* **78**(1): 85–109.
- Beckers, S. (2021). ‘Causal Sufficiency and Actual Causation’. *Journal of Philosophical Logic* **50**(6): 1341–1374.

- Beebe, H. (2004). 'Causing and Nothingness'. In *Causation and Counterfactuals*, edited by L. A. Paul, E. J. Hall, and J. Collins, Cambridge, MA, USA: MIT Press. 291–308.
- Blanchard, T. and Schaffer, J. (2017). 'Cause Without Default'. In *Making a Difference*, edited by H. Beebe, C. Hitchcock, and H. Price, Oxford: Oxford University Press. 175–214.
- Dowe, P. (2000). *Physical Causation*. Cambridge Studies in Probability, Induction and Decision Theory, Cambridge University Press.
- Gallow, D. J. (2021). 'A Model-Invariant Theory of Causation'. *The Philosophical Review* **130**(1): 45–96.
- Hall, N. (2000). 'Causation and the Price of Transitivity'. *Journal of Philosophy* **97**(4): 198.
- Hall, N. (2004). 'Two Concepts of Causation'. In *Causation and Counterfactuals*, edited by J. Collins, N. Hall, and L. Paul, Cambridge, MA.: MIT Press. 225–276.
- Hall, N. (2007). 'Structural Equations and Causation'. *Philosophical Studies* **132**(1): 109–136.
- Halpern, J. (2016). *Actual Causality*. Cambridge, MA: MIT Press.
- Halpern, J. Y. (2015). 'A Modification of the Halpern-Pearl Definition of Causality'. *Proc. 24th International Joint Conference on Artificial Intelligence (IJCAI 2015)* : 3022–3033.
- Halpern, J. Y. and Hitchcock, C. (2010). 'Actual Causation and the Art of Modeling'. In *Heuristics, Probability, and Causality: a Tribute to Judea Pearl*, edited by R. Dechter, H. Geffner, and J. Y. Halpern, London: College Publications. 383–406.
- Halpern, J. Y. and Hitchcock, C. (2015). 'Graded Causation and Defaults'. *British Journal for the Philosophy of Science* **66**(2): 413–457.
- Halpern, J. Y. and Pearl, J. (2005). 'Causes and Explanations: A Structural-Model Approach. Part I: Causes'. *British Journal for the Philosophy of Science* **56**(4): 843–887.

- Hausman, D. M. (1998). *Causal Asymmetries*. Cambridge Studies in Probability, Induction and Decision Theory, Cambridge University Press.
- Hiddleston, E. (2005). 'Causal Powers'. *The British Journal for the Philosophy of Science* **56**(1): 27–59.
- Hitchcock, C. (2001). 'The Intransitivity of Causation Revealed in Equations and Graphs'. *The Journal of Philosophy* **98**(6): 273–299.
- Hitchcock, C. (2007). 'Prevention, Preemption, and the Principle of Sufficient Reason'. *Philosophical Review* **116**(4): 495–532.
- Hitchcock, C. (2009). 'Structural Equations and Causation: Six Counterexamples'. *Philosophical Studies* **144**(3): 391–401.
- Hitchcock, C. and Knobe, J. (2009). 'Cause and Norm'. *The Journal of Philosophy* **106**(11): 587–612.
- Hume, D. (1748/1975). *An Enquiry Concerning Human Understanding*. Oxford: Clarendon, 3rd edn.
- Lewis, D. (1973). 'Causation'. *Journal of Philosophy* **70**(17): 556–567.
- Lewis, D. (2000). 'Causation as Influence'. *Journal of Philosophy* **97**(4): 182–197.
- Mackie, J. L. (1965). 'Causes and Conditions'. *American Philosophical Quarterly* **2**(4): 245–264.
- Mackie, J. L. (1974). *The Cement of the Universe: A Study of Causation*. Clarendon Library of Logic and Philosophy, Clarendon Press.
- McDermott, M. (1995). 'Redundant Causation'. *The British Journal for the Philosophy of Science* **46**(4): 523–544.
- McGrath, S. (2005). 'Causation By Omission: A Dilemma'. *Philosophical Studies* **123**(1-2): 125–148.
- Paul, L. A. (2000). 'Aspect Causation'. *Journal of Philosophy* **97**(4): 235.
- Paul, L. A. and Hall, N. (2013). *Causation: A User's Guide*. Oxford.

- Ramachandran, M. (1997). 'A Counterfactual Analysis of Causation'. *Mind* **106**(422): 263–277.
- Woodward, J. (2003). *Making Things Happen : A Theory of Causal Explanation*. Oxford: Oxford University Press.
- Wright, R. W. (1985). 'Causation in Tort Law'. *California Law Review* **73**(6): 1735–1828.
- Yablo, S. (2002). 'De Facto Dependence'. *The Journal of Philosophy* **99**(3): 130–148.