Factual Difference-Making

Holger Andreas*[1] and Mario Günther†[2]

[1]University of British Columbia
[2]LMU Munich & Carnegie Mellon University

Abstract

In this paper, we analyse causation in terms of factual difference-making. Factual difference-making is an alternative to counterfactual difference-making which does not face the problem of redundant causation. Our analysis defined in a framework of causal models solves many causal scenarios with which accounts in terms of counterfactual difference-making still struggle. The upshot is that causes are, perhaps, better understood as factual difference-makers rather than counterfactual ones.
Keywords. Causation; Difference-Making; Factual Conditionals; Causal Model Semantics; Deviancy

1 Introduction

A cause makes a difference to its effect. This idea of difference-making has been used to analyse token causation. Lewis (1973a) spelled out the idea in terms of counterfactual conditionals: if a cause had not occurred, its effect would not have occurred. However, an effect might still occur even if one of its causes had not occurred. There might simply be another

*Holger.Andreas@ubc.ca.
†Corresponding Author: Mario.Guenther@lmu.de. The authors contributed equally.

event which then brings about the effect. This problem of redundant causation means trouble for accounts of causation in terms of counterfactual difference-making.

In this paper, we spell out the idea of difference-making in terms of factual conditionals: if a cause occurs, so does its effect. This conditional means that the cause implies the effect when it is unsettled whether cause and effect occur. We propose that such conditionals express factual difference-making. And we will show that causation understood as factual difference-making avoids the problem of redundant causation altogether. The upshot is that token causation may be better understood in terms of factual difference-making. Or so we suggest.

We proceed as follows. First, we explain counterfactual difference-making (Section 2) and factual difference-making (Section 3). Then we develop our analysis of token causation in terms of factual difference-making (Section 4). For that purpose, we define factual conditionals for causal models. The gist of the preliminary analysis is that an event C is a cause of another event E only if both events are actual, and C implies E in a state, where C and E are unsettled. We show that our analysis, suitably amended, solves a number of causal scenarios with which counterfactual accounts struggle (Section 5). Our final analysis says that causation is deviant factual difference-making in a forward-directed way. We end by comparing our analysis to accounts in terms of counterfactual difference-making (Section 6).

2 Counterfactual Difference-Making

The idea of counterfactual difference-making motivates the accounts of token causation in the tradition of Lewis (1973a). He defines the idea as follows.[1]

> An event C makes a counterfactual difference to another event E if and only if (iff) two subjunctive conditionals are true:

[1] Lewis (1973a, pp. 561-3) calls his notion of counterfactual difference-making "causal dependence".

(i) if C had occurred, E would have, and

(ii) if C had not occurred, E would not have.

On Lewis's (1973b) semantics of subjunctive conditionals, a subjunctive 'if A had been the case, C would have been the case' is true if both A and C are actually the case.[2] This semantics gives rise to the simple account of counterfactual difference-making.

> An event or absence C is a cause of another event or absence E iff
>
> (i) C and E are actual, and
>
> (ii) the counterfactual 'if C had not been actual, E would not have been' is true.[3]

Given that the distinct events C and E occur, C is a cause of E on this simple account if the counterfactual $\neg C \mathrel{\Box\!\!\rightarrow} \neg E$ is true. This conditional is counterfactual because its antecedent assumes contrary-to-fact that C had *not* occurred. The non-occurrence or absence $\neg C$ of C is represented using a negation. A counterfactual assumes that a counterfact is actual.[4]

The simple account of counterfactual difference-making faces the problem of redundant causation. There are causal scenarios, where there is more than one event that would be sufficient for the effect to occur. In such scenarios, the effect would occur even if one of its genuine causes would have been absent. But if an effect E still occurs even if its cause C had not

[2] Lewis (1973b) calls subjunctive conditionals with factual antecedents also 'counterfactuals'.

[3] Lewis (1986, p. 189) understands an absence $\neg A$ as the non-occurrence of any token event of type A. If the absence $\neg A$ had not been, some token event A would have been. Counterfactual dependence between occurring events is thus only a special case of counterfactual dependence between actual events and absences. The latter is still sufficient for causation, or so says Lewis.

[4] If an event E occurs, E is a fact and the absence $\neg E$ is a counterfact. If an absence $\neg E$ is actual, the counterfact is the occurring event E. We say that all events are contingent and so are all facts and counterfacts. It follows that tautologies and contradictions are neither events nor facts nor counterfacts.

occurred, C does not count as a cause of E on the simple counterfactual account. The problem is that there might be another event which brings about the effect E and yet C is among its genuine causes.

One type of redundant causation is overdetermination. Let's say Suzy and Billy each throw a rock at a window, the rocks impact upon the window at the same time and each rock alone would have been sufficient to break the window. Each rock throwing is arguably a cause of the window's breaking. But the simple account of counterfactual difference-making says that neither Suzy's nor Billy's throwing is a cause of the window's shattering. Had Suzy not thrown her rock, the window would still have shattered—due to Billy's throw. And had Billy not thrown his rock, the window would have shattered anyways—due to Suzy's throw. But then, what caused the shattering of the window? Surely, we do not want to say that the shattering is uncaused.[5]

The overdetermination scenario can be depicted by a neuron diagram (Lewis, 1986). Neuron diagrams represent events and absences and the causal dependences between them. The firing of a neuron indicates the occurrence of some event and the non-firing indicates its non-occurrence. The firing of a neuron is visualized by a gray-shaded node, the non-firing by a white node. The causal dependences between events and absences are represented by arrows. Each arrow with a head represents a stimulatory connection between two neurons, each arrow ending with a black dot an inhibitory connection. There are two laws for neuron diagrams. A neuron does not fire if it is inhibited by at least one. And a neuron fires if it is inhibited by none and stimulated by at least one. Here is the neuron diagram of overdetermination:

[5] Not everyone agrees though, as we will see in Section 6.

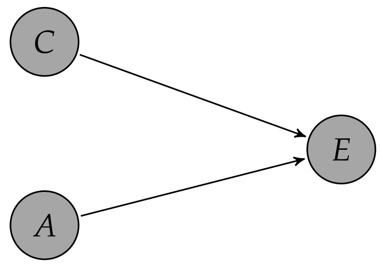

Figure 1: Overdetermination

Neuron C and neuron A fire. The firing of each of C and A alone suffices to excite neuron E. Hence, the common firing of C and A overdetermines E to fire. Suzy's throw and Billy's overdetermine the window to shatter.

What would have happened if C had not fired? Well then, A would still have fired and excited E to fire. And symmetrically, had A not occurred, E would still have occurred due to C. Neither C nor A alone make a counterfactual difference to E. Hence, neither C nor A counts as a cause of E on the simple account of counterfactual difference-making.

We have seen that overdetermination—a type of redundant causation—makes trouble for the simple account of counterfactual difference-making. And no wonder. An overdetermining cause does not make a counterfactual difference to its effect. More sophisticated accounts of causation drop the necessity of counterfactual difference-making for causation in response to the problem of redundant causation. We will explain in Section 6 how they deviate from the simple account of counterfactual difference-making—and so from their very own guiding idea.

3 Factual Difference-Making

We propose another idea of difference-making. In this section, we outline the alternative before we analyse it in terms of causal models in the next section.

> An event or absence C makes a factual difference to another event or absence E iff
>
> (i) C and E are actual, and
>
> (ii) the settling conditional 'if C then E' is true.

A settling conditional 'if A then C' is true iff, in a state unsettled about A and C, if A becomes actual, so does C. An event A is unsettled iff neither A nor $\neg A$ is actual. A settling conditional presupposes a state unsettled about its antecedent A and consequent C. The assumption of the antecedent does not contradict any facts of its unsettled state because there are none with respect to A in this state. The fact A and its counterfact $\neg A$ are both not actual in a state unsettled on A.

The assumption that an event A is unsettled contradicts actuality. In actuality, any event A is settled: it either occurs or else it does not. The unsettling assumption of a true settling conditional 'if A then C' results in a state which is partly undetermined about actuality. This unsettled state contains as many of the actual facts as possible while there is no positive posit about whether $A, \neg A, C$, or $\neg C$ is actual. An unsettled state represents actuality with factual gaps—so to speak. Relative to this gappy or 'less than actual' state, the further assumption of the antecedent merely fills in a factual gap.

The unsettled state of a settling conditional is consistent with all ways the world might be with respect to the antecedent and consequent—it is in particular consistent with the actual facts.[6] The truth of a settling conditional requires an unsettled state with respect to the events and absences expressed in the antecedent and consequent. A true settling conditional 'if A then C' says that the consequent is settled as C upon assuming A as a fact.[7]

[6]Lewis's (1973b) semantics for counterfactuals presupposes possible worlds, where the contingent antecedent of a counterfactual is true. Our semantics only presupposes gaps of fact.

[7]Another way to understand settling conditionals is this: 'if A then C' is true iff, in a state without any information on A and C, C is actual if A is (Andreas and Günther, forthcoming). Perhaps an epistemic reading is most natural for settling conditionals: you believe 'if A then C' iff, after suspending judgment on A and C, you can infer C when

Given that the events or absences C and E are actual, the settling conditional 'if C then E' expresses factual difference-making: in a state unsettled on C and E, the fact C together with the remaining facts implies the fact E. In this case, the settling conditional is *factual*: its antecedent and its consequent express facts.[8] We suggest our factual conditional says that C makes a factual difference to E. Our semantics gives rise to the simple account of factual difference-making.

> An event or absence C is a cause of another event or absence E iff the factual conditional 'If C then E' is true.

The simple account of factual difference-making solves the overdetermination problem without further ado. Assume it is unsettled whether Suzy throws her rock and whether the window shatters. As it is assumed that the window's shattering is unsettled, it must also be unsettled whether Billy throws. For if Billy throws his rock, it is settled that the window shatters in the overdetermination scenario. Now, if Suzy throws her rock, the window shatters. Suzy's throw is a cause of the window's shattering on the simple account of factual difference-making. And so is Billy's. Overdetermining causes make a factual difference to their effects.

We represent unsettled events, or more generally unsettled states, by adding a feature to neuron diagrams: dotted nodes. The state in the overdetermination scenario, where it is unsettled whether Suzy throws, Billy throws, and the window shatters, respectively, is represented by the following neuron diagram.

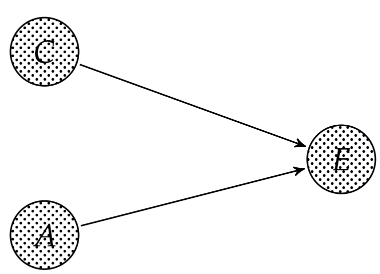

supposing A (Andreas and Günther, 2019, 2020, 2021).

[8] Goodman (1947, p. 114) likewise calls conditionals whose antecedents and consequents are both true 'factual conditionals'.

This state is fully unsettled with respect to C, A, and E. There is no information on whether any neuron fires or not. But the dependences between the neurons are intact. Hence, if C fires in the fully unsettled state, so does E. C makes a factual difference to E—and so does A.

We have introduced factual difference-making. The simple account of factual difference-making solves the overdetermination problem without further ado—unlike the simple account of counterfactual difference-making. This result is promising and calls for an analysis of causation in terms of factual difference-making.

4 Preliminary Analysis

In this section, we propose our preliminary analysis of causation. The guiding idea is that a cause makes a factual difference to its effect in a forward-directed way. The factual conditional of the previous section is at the heart of our analysis. We will begin by providing a causal model semantics for settling conditionals.

4.1 Causal Models

Causal models are devices to represent causal scenarios. Our causal models $\langle M, V \rangle$ have two components: a set M of structural equations and a set V of literals.[9] The literals tell us which events occur and which do not. $A \in V$ means that the event expressed by A occurs. $\neg A \in V$, by contrast, means that no token event A of the relevant type occurs. In brief, the literals represent which events and absences are actual.

A set of structural equations represents the causal dependences between the events of a causal scenario. A structural equation tells us whether an event occurs given the occurrences and non-occurrences of certain other

[9]Our framework of causal models is inspired but deviates from Halpern's (2000) and Pearl's (2009). Our structural equations are formulas, not functions. Another difference is that our account is confined to binary variables whose values are represented by literals—a propositional variable or its negation.

events. Where A is a propositional variable and ϕ a propositional formula, we say that

$$A = \phi$$

is a structural equation.[10] It tells us whether the event represented by A occurs depending on the combination of events and absences described by ϕ. We may regard ϕ as a truth function. Its arguments represent events and absences. Its truth value determines whether A or $\neg A$. The meaning of this determination is intended to be causal: the combination of events and absences described in ϕ causally determine whether or not A occurs. We call $A = \phi$ the structural equation of A.

The idea behind our causal model semantics is that the propositional variables used in the structural equations are evaluated by the literals. The input values for ϕ are given by the literals in V. In this sense, the literals play a semantic role akin to valuation functions in propositional logic. This will become clear in the next section when we will say what it means for a propositional formula to be satisfied by a set of literals.

Let us have a look at the causal model of overdetermination. This model is given by $\langle \{E = C \vee A\}, \{C, A, E\} \rangle$. The structural equation $E = C \vee A$ says that E fires just in case either C or A fire, or both. The values of the propositional variables are given by the literals C, A, E, which indicate that all neurons fire in the actual situation. For readability, we will depict such causal models $\langle M, V \rangle$ in a two-layered box, where the upper layer shows the set M of structural equations and the lower layer the set V of literals. For the overdetermination scenario, we obtain:

$E = C \vee A$
C, A, E

Let us make the semantic role of literals explicit. If $A \in V$, then A is assigned the truth value true. If $\neg A \in V$, then A is assigned the truth

[10]There are restrictions on causal models. For each propositional variable A of a causal model, there is at most one structural equation $A = \phi$, where ϕ is a propositional formula constructed out of combining propositional variables with the symbols for negation, disjunction, and conjunction. ϕ contains no other symbols, in particular no symbol for implication or bi-implication.

value false. If neither $A \in V$ nor $\neg A \in V$, then the truth value of A is unsettled or indeterminate. We will rely on such values in order to model unsettled events and states—or a lack of information as to which events occur if you prefer.

4.2 Semantics

When does a set V of literals satisfy a propositional formula ϕ? We define a satisfaction relation using the semantics of classical propositional logic. Where $\models_{CL}$ stands for the satisfaction relation of propositional logic, we say:

$$V \models \phi \text{ iff } V \models_{Cl} \phi.$$

In words, V satisfies ϕ iff the set V entails ϕ in the sense of classical propositional logic. Likewise, we define satisfaction for structural equations as follows:

$$V \models A = \phi \text{ iff, either } V \models_{Cl} A \text{ and } V \models_{Cl} \phi, \text{ or } V \models_{Cl} \neg A \text{ and } V \models_{Cl} \neg\phi.$$

In simpler terms, V satisfies the structural equation $A = \phi$ iff V either entails A and ϕ, or else $\neg A$ and $\neg\phi$. Furthermore, we say that V satisfies a set M of structural equations just in case V satisfies each element of M:

$$V \models M \text{ iff } V \models A = \phi \text{ for all } A = \phi \text{ in } M.$$

For sets Γ that contain structural equations and propositional formulas, entailment is understood in the standard way: $\Gamma \models \phi$ iff ϕ is satisfied by any valuation V that satisfies all members of Γ. These concepts at hand, we can define the satisfaction relation for causal models $\langle M, V \rangle$:

$$\langle M, V \rangle \models \phi \text{ iff } M \cup V \models \phi.$$

Our analysis relies on causal models that represent states which are unsettled about whether presumed cause and putative effect occur. We say that a causal model $\langle M, V \rangle$ is *unsettled* about a formula ϕ iff $\langle M, V \rangle$ satisfies none of ϕ and $\neg\phi$. Note that the causal model $\langle M, V \rangle$ is settled on any propositional formula as long as V contains a literal for each propositional variable. However, if $A \notin V$ and $\neg A \notin V$, then $\langle M, V \rangle$ may well be unsettled on certain propositional formulas. The fully unsettled state of the overdetermination scenario is represented by the causal model $\langle \{E = C \lor A\}, \{\} \rangle$.

So far, our causal model semantics parallels classical propositional logic. The semantics of $=$, in particular, does not differ from the semantics of the bi-implication $\leftrightarrow$ of classical logic. This means that our structural equations are symmetric just like bi-implications: we can infer from the left-hand side to the right-hand side and the other way around. The symmetry of structural equations is in our framework of causal models only broken by interventions, as we will see shortly.

The intended meaning of a structural equation is that certain events and absences causally determine the occurrence or non-occurrence of another event. The equality symbol of a structural equation $A = \phi$ is directed in the sense that ϕ describes a combination of events and absences that determines whether or not the event described by A occurs, and not the other way around. This being said, the structural equation carries information that allows inferences from A to ϕ. Such inferences go from an effect to its causal factors. We call them *backward-directed*. If, by contrast, an inference goes from causal factors to an effect, we call it *forward-directed*.

Causation is not symmetric, if not asymmetric outright. If C is a cause of E, it does not follow that E is also a cause of C. The present semantics of structural equations does not account for the non-symmetry of causal relations. To account for the non-symmetry, we introduce interventions. The idea is that interventions isolate the forward-directed consequences of occurrences and non-occurrences of events from the backward-directed ones.

Suppose we want to determine the forward-directed causal consequences of the occurrence of A for a causal model $\langle M, V \rangle$. Intervening on this causal model by A results in a causal model $\langle M', V \cup \{A\} \rangle$. If the equation $A = \phi$ is a member of M, M' is obtained from M by removing this equation. Otherwise $M' = M$. We call $\langle M', V \cup \{A\} \rangle$ the causal submodel of $\langle M, V \rangle$ after the intervention by A. By removing the structural equation

$A = \phi$ from M, backward-directed inferences from A or $\neg A$ are excluded in the causal submodel.

Complex interventions may be represented by a set I of literals. Let us denote interventions by an operator $[\cdot]$ that takes a causal model $\langle M, V \rangle$ and a set I of literals, and returns a causal model—the submodel of $\langle M, V \rangle$ after the intervention by I. The intervention by a set of literals is defined as follows:

$$\langle M, V \rangle [I] = \langle M_I, V \cup I \rangle, \text{ where}$$

$$M_I = \{(A = \phi) \in M \mid A \notin I \text{ and } \neg A \notin I\}.$$

M_I is the subset of M that contains each structural equation $A = \phi$ whose variable A is not evaluated by any member of I. After intervening by I on the causal model $\langle M, V \rangle$, the set I becomes part of the valuation of the resulting submodel. The resulting submodel is again a causal model consisting of a set of structural equations and a set of literals. So iterated interventions are well-defined.

Interventions may well result in inconsistent causal models. Even if the original causal model $\langle M, V \rangle$ is consistent, $M_I \cup V \cup I$ and $V \cup I$ may well be inconsistent. There is, however, no reason for concern. We need interventions only as part of the definition of our settling conditional. And this definition sidesteps such inconsistencies, as we will see below.

Let us illustrate how interventions work. The causal model $\langle \{E = C \vee A\}, \varnothing \rangle$ is unsettled on whether or not, let's say, neurons C and E fire. What happens if neuron C is active? This question is answered by the causal submodel $\langle \{E = C \vee A\}, \varnothing \rangle [C] = \langle \{E = C \vee A\}, \{C\} \rangle$. This causal submodel satisfies E, which means that E is going to fire as well. By contrast, the causal model $\langle \{E = C \vee A\}, \{\neg C, \neg A\} \rangle$ does not satisfy E. This means that neuron E is not active if neither A nor C is. What happens if we intervene by E? $\langle \{E = C \vee A\}, \varnothing \rangle [E] = \langle \varnothing, \{E\} \rangle$. This is not a very interesting causal model since the set of structural equations is empty. Nothing can be inferred about the causal consequences of E. And nothing can be inferred in a backward-directed way about the causal factors of E.

4.3 Settling Conditionals

Our preliminary analysis requires us to define settling conditionals for causal models. The evaluation of a settling conditional can be broken down into two questions. The first question is whether there is a state in which antecedent and consequent are unsettled. If so, the second question is whether assuming the antecedent entails the consequent in this unsettled state. The settling conditional is true iff both questions are answered by 'yes'.

We translate this two step evaluation into our framework of causal models. Unsettled states are represented by unsettled causal models. A settling conditional $C \gg E$ is true in a causal model iff there is a causal model minimally unsettled about C and E, in which C entails E. Here is the first definition of our settling conditional $\gg$.

Definition 1. $\langle M, V \rangle \models C \gg E$
Let $\langle M, V \rangle$ be a causal model. $\langle M, V \rangle \models C \gg E$ iff there is $V' \subseteq V$ such that

(1) $\langle M, V' \rangle$ is unsettled on C and E,

(2) there is no $V'' \subseteq V$ such that $V' \subset V''$ and $\langle M, V'' \rangle$ is unsettled on C and E, and

(3) $\langle M, \emptyset \rangle \, [V'][C] \models E$.

Condition (1) ensures that there is a causal model that is unsettled on C and E. It says that, for $C \gg E$ to be true in $\langle M, V \rangle$, there must be a causal model $\langle M, V' \rangle$ with the same structural equations and possibly less literals that is unsettled on C and E. This means $\langle M, V' \rangle$ does not satisfy C, $\neg C$, E, and $\neg E$.

Condition (2) says that the causal model must be unsettled in a minimal way. This is implemented by imposing maximality on the set V' of actual literals featuring in condition (1). There is no causal model $\langle M, V'' \rangle$ unsettled on C and E that is more settled than $\langle M, V' \rangle$. The condition requires that the factual gaps are minimal. The condition is motivated by factual difference-making which holds the actual facts in high regard. The deviation from the facts is therefore meant to be as small as possible. This means the factual gaps are to be minimal while ensuring that the causal model is properly unsettled.[11]

Condition (3) says that the consequent E follows from the antecedent C in the presence of the structural equations M and the literals V'. One might wonder why we do not implement the condition as follows: $\langle M, V' \rangle\, [C] \models E$. The reason is that we want to exclude backward-directed inferences from V' and C to E. Our conception of causation is forward-directed: an event is only a cause if it allows to infer the effect in a purely forward-directed way.

We are now in a position to state our preliminary analysis of causation in terms of factual difference-making.

Definition 2. Cause (preliminary)
Let $\langle M, V \rangle$ be a causal model such that $V \models M$. C is a cause of E relative to $\langle M, V \rangle$ iff

(C1) $\langle M, V \rangle \models C \wedge E$, and

(C2) $\langle M, V \rangle \models C \gg E$.

Our analysis presupposes a consistent representation of a causal scenario: a causal model $\langle M, V \rangle$ whose literals V satisfy the structural equations in M. Indeed, an inconsistent causal model seems to be a non-starter for an analysis of causation.

Condition (C1) says that a cause C and its effect E must be actual. Under this condition, condition (C2) says that a cause C must make a factual difference to its effect E in a forward-directed way. The preliminary analysis reconstructs this forward-directed difference-making by means of a causal model that contains all the information about the dependences of a causal scenario, but no information as to whether the cause and the effect are actual. It is time to apply our preliminary analysis.

[11] Condition (2) mirrors a constraint Lewis (1973b) imposes on counterfactuals. The non-actual worlds satisfying the antecedent of a counterfactual are meant to be as similar as possible to the actual world. The idea is to look only at worlds, which minimally deviate from the actual while satisfying the counterfact in the antecedent.

5 Causal Scenarios

In this section, we show that our analysis of causation solves a set of causal scenarios which make trouble for counterfactual accounts. We refine our analysis twice along the way. The first refinement is meant to deal with causal relations that are non-transitive. The second is meant to deal with scenarios where causation seems to depend on what is deviant from normality. We begin by considering a type of redundant causation which differs from overdetermination.

5.1 Early Preemption

Early preemption is another type of redundant causation. An effect E is caused by a genuine cause C. But even if C had not occurred, E would have been brought about by another event A. The simple account of counterfactual difference-making therefore does not count C as a cause of E. As it is, however, C caused E.

An example of early preemption runs as follows. Suzy (C) and Billy (A) each throw a rock at a window. Suzy's rock deflects Billy's mid-flight so that Billy's does not touch the window ($\neg B$). Only Suzy's rock impacts upon the window (D) and so it shatters (E). Had Suzy not thrown, however, Billy's rock would not have been deflected mid-air and would have shattered the window.

The following neuron diagram is canonical for the structure of early preemption.

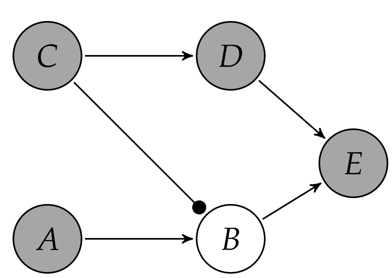

Figure 2: Early Preemption

C's firing excites neuron D, which in turn leads to an excitation of neuron E. At the same time, C's firing inhibits the excitation of B. Had C not fired, however, A would have excited B, which in turn would have led to an excitation of E. The actual cause C preempts the mere would-be cause A. The neuron diagram of early preemption translates into the following causal model $\langle M, V \rangle$:

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

Relative to the model $\langle M, V \rangle$, C is a cause of E. For this to be seen, consider the following causal model $\langle M, V' \rangle$ which is unsettled on C and E:

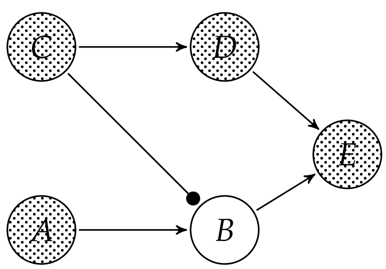

$D = C$ $B = A \wedge \neg C$ $E = D \vee B$
$\neg B$

$\langle M, V' \rangle$ is minimally unsettled among the causal models unsettled on C and E. In particular, if we were to add A to the set V' of literals, the resulting causal model would satisfy C and thus E. Intervening by V' and C yields:

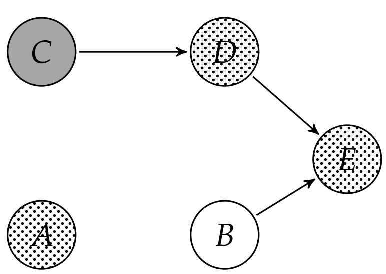

$D = C$ $E = D \vee B$
$C, \neg B$

This causal model determines E to be true. In more formal terms, $\langle M, \emptyset\rangle[V'][\{C\}] \models E$. We have shown that C makes a factual difference to E in a forward-directed way.

It remains to show that A is not a cause of E relative to $\langle M, V\rangle$. And indeed, there is no causal model $\langle M, V'\rangle$ which is minimally unsettled on A and E while $\langle M, \emptyset\rangle[V'][\{A\}] \models E$. $\langle M, V'\rangle$ is only minimally unsettled for $V' = \{\neg B\}$. Intervening by V' and A yields:

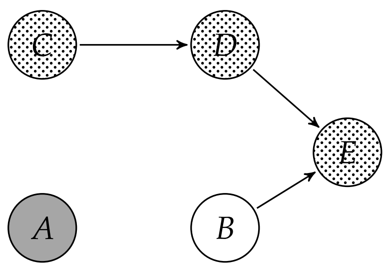

$D = C$
$E = D \vee B$
$A, \neg B$

This causal model does not determine E to be true. Hence, the factual conditional 'if A then E' is false: $\langle M, V\rangle \not\models A \gg E$. A does not make a factual difference to E in a forward-directed way.[12]

5.2 Late Preemption

Lewis (1986, p. 200) subdivides preemption into early and late. We have discussed early preemption in the previous section: a backup process is cut off before the process started by the preempting cause brings about the effect. In scenarios of late preemption, by contrast, the backup process is cut off only because the genuine cause brings about the effect before the preempted cause could do so. Lewis (2000, p. 184) provides the following story for late preemption:

> Billy and Suzy throw rocks at a bottle. Suzy throws first, or maybe she throws harder. Her rock arrives first. The bottle

[12]The intuitive reason for this verdict is this: even if it is unsettled whether or not neurons A and E fire, neuron B may still not fire—as is actual. Recall: factual difference-making respects the actual facts to a maximal extent while ensuring unsettledness.

shatters. When Billy's rock gets to where the bottle used to be, there is nothing there but flying shards of glass. Without Suzy's throw, the impact of Billy's rock on the intact bottle would have been one of the final steps in the causal chain from Billy's throw to the shattering of the bottle. But, thanks to Suzy's preempting throw, that impact never happens.

Crucially, the backup process initiated by Billy's throw is cut off only by Suzy's rock impacting the bottle. Until her rock impacts the bottle, there is always a backup process that would bring about the shattering of the bottle an instant later.

Late preemption can be represented by the following neuron diagram.[13]

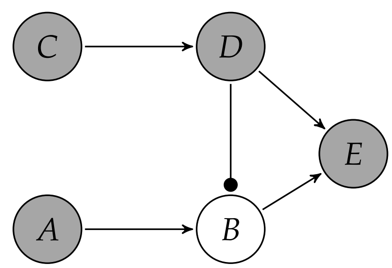

Figure 3

Suzy throws her rock (C) and Billy his (A). Suzy's rock impacts the bottle (D), and so the bottle shatters (E). Suzy's rock impacting the bottle (D) prevents Billy's rock from impacting the bottle ($\neg B$). The neuron diagram of late preemption translates into the following causal model $\langle M, V \rangle$:

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

[13]How to best represent late preemption in neuron diagrams and causal models is somewhat controversial (Hall, 2007; Hitchcock, 2007b; Paul and Hall, 2013). We follow Halpern and Pearl (2005, pp. 861-2).

Only the equation for B differs from the causal model of early preemption: the occurrence of B requires A to occur and the absence of D instead of the absence of C. This difference seems negligible given that D occurs just in case C occurs. It is thus unsurprising that our analysis treats late preemption analogous to early preemption.

Relative to $\langle M, V \rangle$, C is a cause of E. For this to be seen, consider the following causal model $\langle M, V' \rangle$ which is minimally unsettled on C and E:

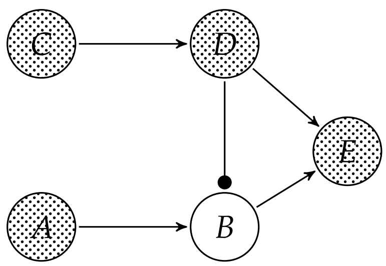

$D = C$ $B = A \wedge \neg D$ $E = D \vee B$
$\neg B$

Intervening by V' and C yields:

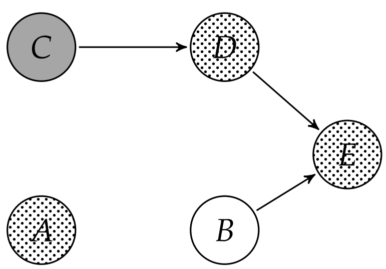

$D = C$ $E = D \vee B$
$C, \neg B$

This causal model determines E to be true. In more formal terms, $\langle M, \emptyset \rangle [V'][\{C\}] \models E$. We have shown that C makes a factual difference to E in a forward-directed way.

It remains to show that A is not a cause of E relative to $\langle M, V \rangle$. And indeed, there is no causal model $\langle M, V' \rangle$ which is minimally unsettled on A and E while $\langle M, \emptyset \rangle [V'][\{A\}] \models E$. $\langle M, V' \rangle$ is only minimally unsettled for $V' = \{\neg B\}$. Intervening by V' and A yields:

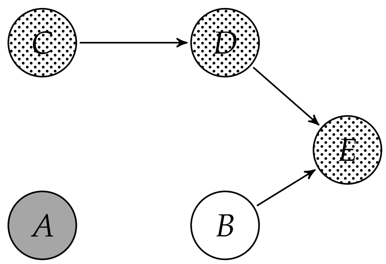

$$D = C$$
$$E = D \lor B$$
$$A, \neg B$$

This causal model does not determine E to be true. Hence, A is not a cause of E. We have shown that our analysis of causation as forward-directed factual difference-making solves the problem of preemption—another type of redundant causation.

5.3 Boulder Scenario

There are other causal scenarios which mean trouble for accounts in terms of counterfactual difference-making. One of these scenarios goes as follows. 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 (Hitchcock, 2001, cf. p. 276).

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 unscathed. But it is odd to say that the dislodging of the boulder causes the hiker to remain unscathed.

The structure of the boulder scenario can be represented by the following neuron diagram (Gallow, 2021, p. 53).

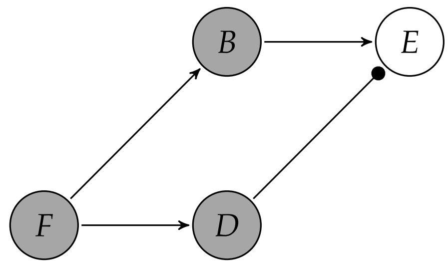

Figure 4: Boulder

Hall (2007, p. 36) calls the network of Figure 4 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 ($\neg E$).

In neuron speak, F fires and thereby excites neuron B to fire, which in turn threatens to excite neuron E. At the same time, F's firing excites neuron D, whose firing prevents E from firing. So F's firing creates a process via B that threatens to bring about E and at the same time initiates another process via D that prevents the threat. F cancels its *own* threat—the threat via B—to prevent E. F should not count as a cause of $\neg E$ because F creates *and* cancels the threat to bring about E (Paul and Hall, 2013, p. 216). The neuron diagram of the boulder scenario translates into the following causal model $\langle M, V \rangle$:

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

Relative to $\langle M, V \rangle$, our preliminary analysis does not count F as a cause of $\neg E$. The reason is that there is no causal model $\langle M, V' \rangle$ which is unsettled about F and $\neg E$. Even the causal model $\langle M, \emptyset \rangle$ satisfies $\neg E$, and so is settled about $\neg E$. There are only two cases. If F is actual, so is B and D, and thus $\neg E$. If $\neg F$ is actual, so is $\neg B$ and $\neg D$, and thus likewise $\neg E$. As there is no causal model $\langle M, V' \rangle$ unsettled on $\neg E$, $\langle M, V \rangle$ does not satisfy $F \gg \neg E$. F does not make a factual difference to $\neg E$ in a forward-directed way. We see here that factual difference-making requires there to be a causal possibility consistent with the structural equations in which the effect is not actual. Otherwise there can be no difference-making.

Interestingly, the simple account of counterfactual difference-making has no trouble with the boulder scenario. The dislodgement of the boulder does not make a counterfactual difference to the hiker's remaining unscathed. Furthermore, the hiker's ducking makes a counterfactual difference to her remaining unscathed. If she had not ducked, the boulder would have hit her.

Our preliminary analysis has the problem that the ducking of the hiker does not count as a cause of her remaining unscathed. There is still no causal model which is unsettled about the ducking and the remaining unscathed. But this is the wrong verdict. Intuitively, the ducking makes a factual difference to her remaining unscathed in a forward-directed way.

What went wrong? Well, the structural equations allow for backward-directed inferences. There are two cases. If the hiker ducks, the boulder has been dislodged, and so rolls toward the hiker. If the hiker does not duck, the boulder has not been dislodged, and so does not roll toward her. This suggests a solution motivated by forward-directedness. We can simply block the backward-directed inference from whether she ducks to whether the boulder is dislodged. We implement this idea by removing the structural equation $D = F$ from the causal model.

Our implementation requires some terminology. Recall that $A = \phi$ is the structural equation of A. We say A is a *child* variable of the *parent* variables occurring in ϕ. Let B be one of the parent variables in ϕ. 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 its 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. Similarly, the ancestors of some variable B are the variables in the transitive closure of the inverse child relation—the parent relation—starting from B. Finally, we say that the descendants (ancestors) of a literal L are all the variables (of the causal model under consideration) which are descendants (ancestors) of the variable of which L is a literal.

We amend the definition of our conditional by condition (4).

Definition 3. $\langle M, V \rangle \models C \gg E$
Let $\langle M, V \rangle$ be a causal model. $\langle M, V \rangle \models C \gg E$ iff there is $V' \subseteq V$ and $M' \subseteq M$ such that

(1) $\langle M', V' \rangle$ is unsettled on C and E,

(2) there is no $V'' \subseteq V$ such that $V' \subset V''$ and $\langle M', V'' \rangle$ is unsettled on C and E,

(3) $\langle M', \emptyset \rangle \, [V'][C] \models E$, and

(4) the structural equation of each descendant of C is in M'.[14]

Condition (4) ensures all and only the forward-directed inferences from a candidate cause C. Our amended analysis allows for more unsettled causal models by disregarding certain backward-directed inferences. This helps to solve the boulder scenario and further scenarios which challenge the transitivity of causation.

On the amended analysis, F is still not a cause of $\neg E$ relative to $\langle M, V \rangle$. For this to be seen, observe that all variables are descendants of F. Condition (4) thus prohibits to remove any structural equation. As a consequence, there is no causal model $\langle M', V' \rangle$ unsettled on F and $\neg E$.

By contrast, D is now a cause of $\neg E$. For this to be seen, observe that D is not a descendant of itself. Hence, the structural equation $D = F$ can be removed. Consider the following causal model $\langle M', V' \rangle$ which is minimally unsettled on D and $\neg E$:

[14] Definition 3 in conjunction with the condition that C and E must be actual results in an analysis of causation similar to Definition 2 in Andreas and Günther (2024). The structure of condition (4) is first introduced and further motivated in Andreas and Günther (2024).

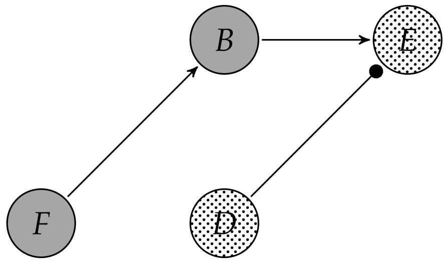

$B = F$ $E = B \wedge \neg D$
F, B

Intervening by V' and D yields:

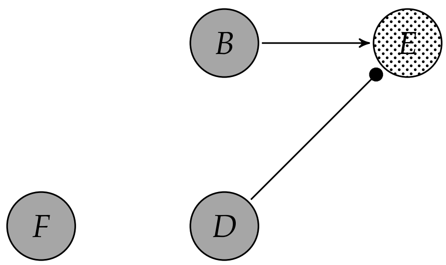

$E = B \wedge \neg D$
F, B, D

This causal model determines $\neg E$ to be true. In more formal terms, $\langle M', \emptyset \rangle [V'][\{D\}] \models \neg E$. We have shown that D makes a factual difference to $\neg E$ in a forward-directed way.[15]

5.4 Bogus Prevention

There is yet another problem which befalls any simple causal model account. We call a causal model account simple if it only factors in structural equations together with values of variables, or alternatively our sets of literals. The problem 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). This is the problem of isomorphic causal models.

[15] Relative to $\langle M, V \rangle$, B is not a cause of $\neg E$. For this to be seen, note that D is in the actual circumstances necessary to infer $\neg E$ and E is a descendant of B. Any causal model which contains E's structural equation and is unsettled on B and $\neg E$ must be unsettled on D as well. An intervention by B in any such causal model does not bring about $\neg E$.

Let us illustrate an instance of the problem. Recall the causal model of the scenario of overdetermination.

$E = C \lor A$
C, A, E

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

$E = \neg F \land D$
$F, \neg D, \neg E$

And indeed, $\neg E$ is 'overdetermined' by F and $\neg D$. The structure of this causal model can be represented by the following neuron diagram.

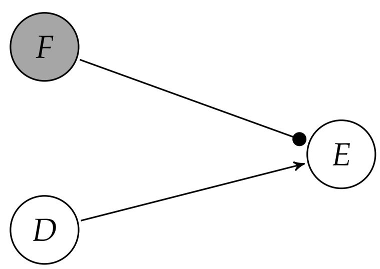

Figure 5: Bogus Prevention

Neuron F fires and thereby would inhibit that neuron E gets excited. However, since neuron D is not firing in the first place, there is no danger at all that neuron E gets excited. Even if F had not fired, E would still not fire. The prevention of E by F is *bogus*. And so F is arguably no cause of $\neg E$—as the simple account of counterfactual difference making says.

A story for the neuron diagram of bogus prevention goes as follows. 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 $\neg E$, 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, 2007a).

As we said above, 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, C is a cause of E in the overdetermination scenario, while F is not a cause of $\neg E$ in the bogus prevention scenario.

Simple causal model accounts of causation—like the accounts of Hitchcock (2001) and Halpern and Pearl (2005) for example—only factor in structural equations and values of variables. As a consequence, they 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.

Our current analysis of causation is a simple causal model account, and so is likewise susceptible to the problem of isomorphic causal models. Hitchcock (2007a), Hall (2007), Halpern (2008), 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 (Beebee, 2004; McGrath, 2005). On this view, genuine effects are brought about by causes that are more deviant from normality than its counterfacts.

We follow Andreas and Günther (2024) to solve the problem of isomorphic causal models by a condition of deviancy. On this strategy, it suffices for now to say that *prima facie* an occurring event is more deviant than a non-occurring event. In the neuron diagrams to come, a firing neuron is thus more deviant than a non-firing neuron. But this is just a first approximation. The question of what constitutes deviancy is more intricate, as we will see below.

Our condition is motivated by the idea that any factual difference-maker of an effect must be deviant. We implement the idea as follows: any C'

potentially making a factual difference to the effect E under consideration must be deviant if it is neither a descendant nor an ancestor of the candidate cause C. Note that C is neither a descendant nor an ancestor of itself. Any one $C' \in V \setminus V'$ is possibly a factual difference-maker of E because it entails the effect together with the literals V' and the structural equations M'. Otherwise C' would remain in V' in virtue of its maximality. To be precise, we add the following condition to conditions (1)-(4) of our settling conditional:

(5) for any literal $C' \in V \setminus V'$ whose variable is neither a descendant nor an ancestor of C, C' is more deviant than $\neg C'$.[16]

The deviancy condition (5) says this: for C to be a cause of E, each possible factual difference-maker C' whose variable is neither a descendant nor an ancestor of C must be more deviant than its respective counterfact. There may be non-deviant events and absences connecting a genuine cause to its effect, but each cause must be deviant. This concludes our analysis: causation is deviant factual difference-making in a forward-directed way.

The condition of deviancy applies to the above scenario of bogus prevention as follows. On our analysis, the absence of poison $\neg D$ and the presence of antidote F in the coffee are no causes because it is normal that coffees are not poisoned. $\neg D \in V \setminus V'$ is neither a descendant nor an ancestor of the respective candidate cause, and it is less deviant than D. It is the normality of the absence $\neg D$ which implies that it and the event F are no causes of target's survival. Or so says our final analysis which thereby solves both overdetermination and bogus prevention.

5.5 Omissions

Omissions pose another problem for many accounts 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

[16]The structure of condition (5) is first introduced and further motivated in Andreas and Günther (2024).

from occurring. The basic structure of omissions can be represented as follows.

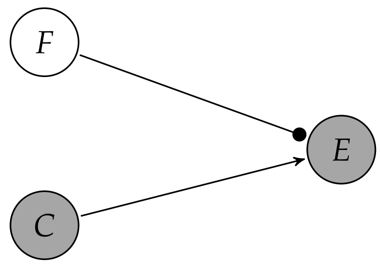

Figure 6: Omission

An event C occurs and brings about an event E. F fails to occur. However, had F occurred, it would have prevented E from occurring. Hence, the simple account of counterfactual difference-making says that $\neg F$ is a cause of E—an unfortunate verdict in its generality.

Here is the causal model for the scenario of omission.

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

On our analysis, the absence $\neg F$ is not a cause of E, given the above convention about deviancy. For then, condition (5) of $\neg F \gg E$ is violated. To see this, consider the only causal model $\langle M, V' \rangle$ which is minimally unsettled on $\neg F$ and E:

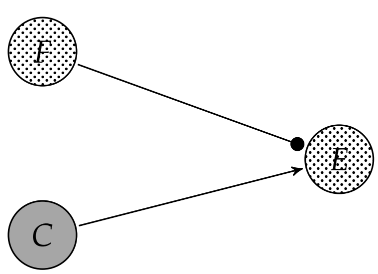

$E = \neg F \wedge C$
C

We see that $\neg F$ is in $V \setminus V'$. And by the above convention $\neg F$ is less deviant than F. We have shown that $\neg F$ is not a cause of E, if $\neg F$ is less deviant than F.[17]

Indeed, many omissions are not 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 $\neg A$ is more deviant than its counterfact A if $\neg A$ violates a norm that is active in the scenario under consideration (Beebee, 2004; Andreas et al., 2023). 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 counterfact. Our final analysis 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.

We have illustrated how condition (5) of deviancy can help to overcome the problem of isomorphic causal models and how it can account for simple scenarios of omission. As to the latter, deviant omissions are genuine causes, non-deviant ones are not. Or so says our analysis on the refined understanding of deviancy.

6 Comparisons

We have analysed causation in terms of factual difference-making. Causation is deviant factual difference-making in a forward-directed way. We have shown that our analysis provides the intuitive verdicts in certain causal scenarios: overdetermination, early and late preemption, the boulder scenario, bogus prevention, and omissions. In this section, we compare our analysis to accounts in terms of counterfactual difference-making.

[17] It is easy to check that C is a cause of E. Just take $V' = \{\neg F\}$.

Let us revisit the simple account of counterfactual difference-making. The simple account does not recognise overdetermining causes and preempted would-be causes. The failure to recognise genuine causes in cases of redundant causation has been taken to show that the simple account of counterfactual difference-making is false. Counterfactual dependence between actual events and absences should not be elevated to be necessary and sufficient for causation. Still, accounts of causation in terms of counterfactual difference-making retain that counterfactual dependence between actual events and absences is sufficient for causation.[18] The problem posed by redundant causation motivates dropping counterfactual difference-making as necessary for causation: there are genuine causes which do not make a counterfactual difference. In other words, the problem provides motivation for deviating from the guiding idea that causation *is* counterfactual difference-making. Our analysis, by contrast, does not deviate from its guiding idea: each cause makes a factual difference to its effect in a forward-directed way. Factual difference-making remains necessary for causation. Or so says our analysis.

Deviating from the guiding idea of counterfactual difference-making is presumably justified if and only if the resulting account of causation preserves the good-making features and does not lead to new problems. One such feature of the simple account of counterfactual difference-making is a solution to the boulder scenario. Had the boulder not been dislodged, the hiker would still have remained unscathed. The boulder does not make a counterfactual difference to the hiker's remaining unscathed. Furthermore, had the hiker not ducked, the boulder would have hit her. The ducking makes a counterfactual difference to her remaining unscathed. This is remarkable because many sophisticated accounts in terms of counterfactual difference-making face troubles in the boulder scenario, for example Lewis's (1973a), de facto accounts, and causal model accounts—as we will see below.

[18] See Lewis (1973a, 2000); Ramachandran (1997); Hitchcock (2001); Yablo (2002, 2004); Woodward (2003); Hall (2004, 2007); Halpern and Pearl (2005); Halpern (2015) among many others.

6.1 Chains of Counterfactual Difference-Making

Lewis (1973a, p. 557) proclaims that we "think of a cause as something that makes a difference, and the difference it makes must be a difference from what would have happened without it". This sounds just like the simple account of counterfactual difference-making. And yet, Lewis does not say that an event is a cause *in virtue of* making a counterfactual difference to another. He rather analyses causation as the transitive closure of such difference-making.

> An event or absence C is a cause of another event or absence E iff there is a counterfactual difference-making chain running from C to E.[19]

A counterfactual difference-making chain is a finite sequence of distinct actual events and absences such that each element in the sequence makes a counterfactual difference to its successor. In symbols, a finite sequence $\langle C, D_1, ..., D_n, E\rangle$ of distinct actual events and absences is a difference-making chain running from C to E iff C makes a counterfactual difference to D_1, D_1 makes a counterfactual difference to ..., and D_n makes a counterfactual difference to E. Lewis thinks of causes as initiators of chains of counterfactual difference-making.[20]

On Lewis's analysis, causation does not suffice for counterfactual difference-making. Suppose an event C makes a counterfactual difference to another event D, which in turn makes such a difference to a third event E. Then C is a cause of E—even if C does not make a counterfactual difference to E. This is the case in the early preemption scenario depicted in Figure 2. Suzy's throw of a rock makes a counterfactual difference as to whether or not her rock impacts upon the window. And given that Suzy's

[19]Lewis (1973a, p. 563) uses the term "causal chain" for counterfactual difference-making chain.

[20]Counterfactual difference-making, of course, still suffices for causation. Suppose the event C makes a counterfactual difference to the distinct event E. Then there is a finite sequence $\langle C, E\rangle$ of distinct actuals such that each element in the sequence makes a counterfactual difference to its successor. Hence, C is a cause of E on Lewis's (1973a) analysis.

rock deflected Billy's, her rock impacting upon the window makes a difference as to whether or not the window shatters. Still, Suzy's throw does not make a counterfactual difference as to the window's shattering. Had she not thrown her rock, Billy's rock would not have been deflected and so would have shattered the window. Difference-making is not transitive. But Lewis thinks causation is.

Lewis's analysis solves the early preemption scenario. There is no sequence of distinct actual events and absences from A to E such that each element makes a counterfactual difference to its successor. If A had not occurred, B still would not have occurred. A does not initiate a counterfactual difference-making chain to E. As we have just seen, C counts as a cause of E. C and E occur, and there is the sequence $\langle C, D, E \rangle$ of distinct events such that the counterfactuals $\neg C \mathrel{\Box\!\!\rightarrow} \neg D$ and $\neg D \mathrel{\Box\!\!\rightarrow} \neg E$ are true.

One might object that the latter counterfactual is only true if the counterfactual $\neg D \mathrel{\Box\!\!\rightarrow} \neg C$ is false. Otherwise C would not have occurred had D not occurred, and so C would not have preempted the process started by A, which then would have caused E. Lewis (1986, p. 201) blocks this argument by arguing that the counterfactual $\neg D \mathrel{\Box\!\!\rightarrow} \neg C$ is *backtracking*: if D had not occurred, its past cause C would still have occurred but somehow failed to cause D; so C would still have interferred with the process started by A and E would not have occurred. We give all the accounts in terms of counterfactual difference-making that they can identify which counterfactuals are backtracking so that backtracking counterfactuals can either be stipulated to be false or else can simply be neglected when determining causation. After all, establishing the direction of causation is a tough problem which seems to be an unsolved problem on any account.[21] Just as counterfactuals are assumed to be non-backtracking, we assume our settling conditionals to be forward-directed.

Lewis's solution to early preemption works only because backtracking is barred and there is an intermediate event D between C and E. Had this intermediate event D not occurred, $\neg B$ would still have been actual due to C's occurrence. In other words, C is a cause of E only because the back-

[21] But see Andreas and Günther (2024); Andreas and Günther (2024) for some first steps to establish the direction of causation.

tracker $\neg D \mathrel{\Box\!\!\rightarrow} \neg C$ is false, and so is $\neg D \mathrel{\Box\!\!\rightarrow} B$. But this means the solution requires that there is some event intermediate between cause and effect such that had this event not occurred, the preempted process would still have been preempted. Hence, Lewis's solution does not apply to late preemption in Figure 3. If Suzy's rock had not impacted upon the window $\neg D$, the bottle would have shattered anyways E. The reason is that the counterfactual $\neg D \mathrel{\Box\!\!\rightarrow} B$ is true: if Suzy's rock had not impacted upon the window, Billy's rock would have. So there is no sequence of stepwise counterfactual dependences that links Suzy's throw to the bottle's shattering. This explains why Lewis subdivides preemption into early and late—even though the scenarios seem to be all too similar. A uniform solution for early and late preemption would be desirable.

As is well-known, Lewis's analysis does not count symmetric overdeterminers as causes. Unlike preemption, there is no causally relevant difference between the overdeterminers depicted in Figure 1. Lewis's analysis delivers the same verdict as the simple account of counterfactual difference-making: neither Suzy's throw nor Billy's throw counts as a cause of the window's shattering—a verdict which strikes many as wrong.

Lewis (1986, pp. 199&200), however, thinks it is only clear that the symmetric overdeterminers are on a par:

> It may or may not be clear whether either [overdeterminer] is a cause; but it is clear at least that their claims are equal. There is nothing to choose between them. Both or neither must count as causes.

To decide whether both throws or neither should count as a cause is for Lewis up to our best theory of causation. It is "spoils to the victor for lack of firm common-sense judgements."[22] (p. 208) Woodward (2003, p. 85) counters: "My guess is that Lewis is wrong about common sense." Or in the words of Paul and Hall (2013, p. 152): "It seems perfectly commonsensical to say that both overdeterminers are causes, and perfectly

[22]Lewis (1986, p. 212) writes: "I used to think that all cases of overdetermination, as opposed to preemption, could be left as spoils to the victor; and that is what I still think about these residual cases." Our example of overdetermination is such a residual case.

puzzling to say that neither are." Indeed, Lewis's later analyses both say that the individual overdeterminers are causes (Lewis, 1986, 2000).[23]

One might reply on behalf of Lewis's (1973a) analysis and the simple account like this: even though the individual rock throws of Suzy and Billy do not count as causes, the disjunction or mereological sum of both throws does.[24] But even proponents of disjunctive causes are hesitant to endorse this solution for overdetermination (Sartorio, 2006). Suppose there is a disjunctive event 'Suzy or Billy throws a rock' and this event is a cause of the window's shattering.[25] Did then Suzy's throw cause the window to shatter? If so, why is this event not recognised as a cause on its own? And why would it be a cause? After all, Suzy's rock does not make a counterfactual difference to the window's shattering—nor is there a difference-making chain between the 'non-disjunctive' events.

If Suzy's throw does not cause the window to shatter but the disjunctive event *at least one throw* does, it seems that Suzy can truthfully say 'I did not cause the window to shatter'. This is a good defense in any legal system, where a person can only be held responsible for an event if the person caused it. If Suzy did not cause the window's shattering, she cannot be held responsible for her vandalism in many legal systems. Perhaps the legal systems, which require causation for responsibility, are wrong. But it is on the defender of counterfactual difference-making to explain why we should find Suzy guilty of vandalism even though she didn't do it—or

[23]These analyses overshoot in early and late preemption: the preempted would-be cause wrongly counts as a cause, respectively.

[24]Coady (2004) defends the simple account of counterfactual difference-making by arguing that redundant causes individually are no causes but their "combination" is. (p. 326ff.) He furthermore tries to explain away the existence of preemption scenarios: they are either scenarios of overdetermination, or else the 'preempting cause' is a genuine non-redundant cause. We remain unconvinced on both fronts as long as there is no convincing account of "ordinary standards of fragility" for events. We think it is unlikely that suitable standards can be found and that Coady's defense collapses without these standards.

[25]It is not so clear to us whether there are disjunctive events. One would need to come up with an ontology of events, where disjunctive events have a proper place and time. We don't know how to do this. Furthermore, one would need to explain away the reductio that the mereological sum of the overdeterminers is a cause provided by Paul and Hall (2013, p. 151, fn. 21).

else why we should not find her guilty. Notably, our analysis in terms of factual difference-making has no such further explanatory need.

Lewis's analysis, somewhat surprisingly, misclassifies the dislodged boulder as a cause of the hiker's remaining unscathed—even though there is no counterfactual difference-making. Barring backtracking, there is a counterfactual difference-making chain: had the boulder not been dislodged, the hiker would not have ducked; and had the hiker not ducked, the boulder would have hit her. This means Lewis's solution to early preemption overshoots. There, Suzy's throw is a cause but no counterfactual difference-maker. In the boulder scenario by contrast, the dislodged boulder is no cause but initiates a chain of counterfactual difference-making. So Lewis's analysis is forced to count the dislodged boulder as a cause. To be clear, defining causation as the transitive closure of counterfactual dependence between actual events and absences solved early preemption. But the transitivity imposed on causation is of no help for late preemption and backfires in the boulder scenario. This result questions whether the deviation from counterfactual difference-making by imposing transitivity is warranted—or only motivated by solving early preemption. Be that as it may, the accounts to follow do not impose transitivity on causation.[26]

6.2 De Facto Dependence

Redundant causation shows that counterfactual difference-making is not necessary for causation. In early preemption, Suzy's throw is a genuine cause of the window's shattering but does not make a counterfactual difference. Billy's throw, by contrast, is not a cause: his rock does not touch the window. There is, however, concealed counterfactual difference-making: given that Billy's rock does not touch the window, the window would not have shattered if Suzy had not thrown her rock. Suzy's throw is perhaps a cause of the window's shattering in virtue of making a counterfactual difference when holding fixed that Billy's rock does not touch

[26]Maslen (2004) analyses the non-transitivity of causation by an account of causation in terms of counterfactual difference-making which relies on making *contrast situations* explicit. Her account explains the intuition that causation is not transitive in certain scenarios and so merits further attention.

the window. This is the idea of de facto dependence (Yablo, 2002).

The simple de facto account goes as follows.

> An event or absence C is a cause of another event or absence E iff
>
> (i) C and E are actual, and
>
> (ii) there is a set F of 'non-disjunctive' actual events and absences such that the counterfactual $(\neg C \wedge \bigwedge F) \;\Box\!\!\rightarrow \neg E$ is true.[27]

The de facto counterfactual $(\neg C \wedge \bigwedge F) \;\Box\!\!\rightarrow \neg E$ says 'if C had not been actual but the events and absences in F had still been actual, E would not have been actual'. The general idea is this: effects depend de facto on their causes—they counterfactually depend on their causes when the right surrounding events and absences are held fixed. The idea immediately poses the question: what are the 'right events and absences' to be held fixed? While it seems clear in the preemption case, a general answer is difficult to give.[28]

The simple de facto account provides a straightforward and uniform solution to early and late preemption. However, the simple de facto account fails for the boulder scenario—it identifies the dislodged boulder as a cause of the hiker's remaining unscathed. If the boulder had not been dislodged but it still had been rolling toward the hiker, the hiker would not have ducked and so would have been hit by the boulder. This de facto counterfactual is true. And yet, it is strange. How could the boulder not have been dislodged and still have rolled toward the hiker and hit her? This seems causally impossible. But wait. How could Billy's rock not have impacted upon the window if Suzy had not thrown? This seems causally

[27] $\bigwedge F$ denotes some conjunction of the members of F.

[28] The simple de facto account does not restrict the set F of events and absences over and above imposing actuality on its members. As a result, counterfactual difference-making between actuals is sufficient for causation. For $F = \emptyset$, the simple de facto account reduces to the simple account of counterfactual difference-making. Hence, the simple de facto account—like Lewis's analysis—recognises more causes than the simple account of counterfactual difference-making.

impossible as well in the scenarios of early and late preemption. A defender of a de facto account should explain why we can hold fixed that Billy's rock does not touch the window in preemption, while we cannot hold fixed that the boulder is rolling toward the hiker.

Hitchcock (2001, pp. 297-8.) thinks the answer is that we are not "willing to take seriously" certain far-fetched and contrary-to-fact combinations of events and absences. But why is holding fixed that Billy's rock does not touch the window if Suzy had not thrown a less far-fetched combination than the boulder rolling toward the hiker if it had not been dislodged? After all, both combinations involve counterfacts and violate the causal dependences between the events and absences. And no wonder. Accounts in terms of counterfactual difference-making rely on counterfacts and 'miracles' to explain causation. A dependence or law violation in point are true backtracking counterfactuals: if Suzy's rock had not touched the window, she would somehow still have thrown the rock—'with unfailing accuracy' we may add to the description of the preemption scenarios. Our analysis in terms of factual difference-making stays clear of these problems because it has no need for assuming causally impossible contrary-to-fact combinations of events and absences.

The simple de facto account succumbs to overdetermination. Suppose C and A overdetermine E. Then there is no set F of 'non-disjunctive' actual events and absences such that $(\neg C \wedge \bigwedge F) \mathrel{\Box\!\!\rightarrow} \neg E$ is true. And similarly, A is not a cause of E. Yablo (2002) writes about overdetermination:

> But then what does cause the window to break? Not the conjunction of the two throws, since the effect could too easily have occurred without it. Not the disjunction, because we are hard put to regard the disjunction as a genuine event. Could it be that nothing causes the window to break? This goes somewhat against the grain. An event that was caused (the breaking was not a miracle!) should, one feels, have causes. (p. 139)

Yablo agrees with "Lewis that the case is *intuitively* undecidable" and "can be left as 'spoils to the victor.'" (ibid.) His own de facto account aims to capture our putative indecision as to whether the individual overdeterminers count as causes. This depends on whether there is a *right* set of actual events and absences which is *more natural* than any *wrong* alternative. Yablo assumes that the following two sets of actuals are most natural for determining whether C is a cause of E.

$\{\neg A \vee C\}$: Billy's rock does not hit the window without Suzy's.

$\{A \vee \neg C\}$: Suzy's rock does not hit the window without Billy's.

'Disjunctive events' usually count as less natural—if not unnatural outright. But Yablo's account requires of the right F not that it is natural—only that it is more natural as compared to the wrong alternatives. As it turns out, $\{\neg A \vee C\}$ is right for determining whether C is a cause of E, and $\{A \vee \neg C\}$ is wrong. Symmetrically, $\{A \vee \neg C\}$ is right for determining whether A is a cause of E, and $\{\neg A \vee C\}$ is wrong. If the two sets are indeed the only most natural ones, as Yablo assumes, then each of the two throws C and A counts as a cause if 'more natural' is understood as *at least as natural*; they do not count as causes if 'more natural' is understood as *strictly more natural*. This is an elegant explanation of our putative indecision. But what if you are decided that the individual overdeterminers are causes? Well then, you should firmly understand 'more natural' as 'at least as natural'.

We have seen that what causes what depends on Yablo's de facto account on what is comparatively more natural. The problem is that he "postpone[s] (= ignore[s]) the question of what is the best thing to mean by 'natural'." (p. 133, fn. 11) A lack of a general account of comparative naturalness does not mean that his de facto account is beyond repair. Indeed, perhaps a worked-out account of comparative naturalness could explain the difference between the de facto counterfactuals in the boulder scenario and preemption. But as it stands, Yablo's account only provides clear verdicts when the comparative naturalness of sets of actual events and absences is beyond doubt—not so often.

6.3 Causal Model Accounts

There are sophisticated de facto accounts of causation which—unlike Yablo's—have no need for comparative naturalness. Like our analysis, they provide clear verdicts by relying on an interventionist semantics of conditionals.[29] These accounts rely on the causal models of Pearl (2009) to define counterfactuals in terms of interventions. The interventionist counterfactuals, in turn, are suitable to spell out de facto and *de counterfacto* dependence. Relative to the causal model of overdetermination, for example, Suzy's throw is a cause of the window's shattering because the window's shattering counterfactually depends on Suzy's throw if we hold the counterfact that Billy does not throw fixed by intervention. Holding fixed the variable A at its counterfactual value $\neg A$ reveals a hidden counterfactual dependence of the effect E on its cause C. This is a straightforward solution to the problem of overdetermination.

The causal model accounts can solve overdetermination if they lift the restriction of the simple de facto account that only actual events and absences can be held fixed. But this move from de facto to de counterfacto dependence opens the door for a plethora of new problems. Billy's throw, for example, comes out as a cause of the window's shattering in early preemption on the simple de counterfacto account: A and E is actual and the de counterfacto conditional $(\neg A \wedge \bigwedge CF) \mathrel{\Box\!\!\rightarrow} \neg E$ is true for the set $CF = \{\neg C, \neg D\}$ of facts and counterfacts.[30] What is direly needed is a restriction on which facts and counterfacts can be held fixed. And in the best case this restriction should be clear and well-motivated.

Hitchcock (2001) restricts the set of facts and counterfacts by the idea that a specific *weakly active* path from cause to effect must remain intact. For C to be a cause of E in a causal model, he effectively requires that C makes a counterfactual difference to E when the variables between C and E—which are *not* on a specific directed path from C to E—are held fixed at certain actual or non-actual values; the values, which may be held fixed, are restricted to those which do not change the values of the variables lying on the specific path. This account solves early and late preemption as well as overdetermination. However, it fails for the boulder scenario: there is a directed path on which the dislodged boulder makes a counterfactual dif-

[29] Hitchcock (2001); Woodward (2003); Halpern and Pearl (2005); Hall (2007); Halpern (2015)

[30] Another among the many problems for the simple de counterfacto account is the scenario discussed by Paul and Hall (2013, pp. 198-9).

ference to the hiker's remaining unscathed via the hiker's ducking when holding fixed that the boulder is rolling towards the hiker off the path.

The accounts of Halpern and Pearl (2005) and Halpern (2015) succumb to the boulder scenario for similar reasons. The former solves overdetermination by relying on de counterfacto dependence. But it is difficult to see an intuitive motivation for their restrictions on what facts and counterfacts may be held fixed (Andreas et al., 2023, Sect. 7). On both accounts, candidates for causes and effects are not only single variable assignments. Causes may be sets of variable assignments and effects are propositional formulas.

Halpern's (2015) modified account is very similar to the simple de facto account: it tests for counterfactual dependence when certain variables are held fixed by intervention at their actual values. Just like the simple de facto account, the modified account solves early and late preemption but fails for overdetermination. There is simply no set of actual value assignments such that the effect would counterfactually depend on one of the overdeterminers. This being said, the set $\{C, A\}$ of variable assignments counts as a cause of the overdetermined effect E. Holding nothing fixed, if $\neg C$ and $\neg A$ were the case, $\neg E$ would be the case. While the only cause of the effect is the set, its members are *parts* of the cause. And parts of causes are "what we think of as causes"—or so says Halpern (2016, p. 25).[31] Perhaps a cause is nothing but an element of a minimal set which makes a counterfactual difference (Andreas and Günther, 2021). This move would pose the mereological problems we mentioned above and would require a suitable notion of minimality. It is up to the defenders of counterfactual difference-making to explore this avenue.

[31] Halpern's (2015) account does not count the set $\{C, A\}$ as a cause in a conjunctive scenario, where the occurrence of both events C and A is necessary and sufficient for the effect to occur. But, as Andreas and Günther (2021) pointed out, why is $\{C, A\}$ not 'the'—or at least a—cause of the effect?

6.4 Normality

All the considered accounts of causation in terms of counterfactual difference-making have in common that counterfactual dependence between actual events and absences is sufficient for causation. As a consequence, they all count too many omissions as causes. If Putin had watered my plant, it would not have died. Indeed, if the Queen of England had watered my plant, it would not have died. The same is true of anyone who didn't water my plant.[32] This is an unwelcome result which questions the sufficiency of counterfactual dependence for causation.

One might block the unwelcome result by restricting causation to counterfactual difference-making between occurring events.[33] This solution requires an ontological distinction between events and absences. Cases of prevention are then non-causal. Preventing an accident, for example, would not be causing the accident to be absent. And omissions would be non-causes in general, as defended by Beebee (2004). But it seems that preventing an event from happening is nothing but causing it to be absent. And some omissions seem to be causes while others are not.[34]

Another way to avoid causal omissions is to say that only occurring events can be causes while both events and absences can be effects. Then preventers may be causes, but omissions are always non-causes. My neighbour's failure to water my plant would then not cause it to die—even though she promised to water it. We are convinced by the argument of McGrath (2005) which establishes that the causal status of omissions depends on normality considerations.

[32] Sartorio (2010) puts forth the Prince of Wales problem—a successor to the Queen of England problem. Like its predecessor, it is designed to show that, if counterfactual difference-making suffices for causation, there is too much of it. Maslen (2020) pushes back against Sartorio's argument that there are too many unwanted positive and negative causes.

[33] One should presumably not require that all elements in a chain of counterfactual difference-making must be occurring events; at least not without defusing Schaffer's (2000) argument that there are cases of genuine causation, where a cause is related to its effect via absences.

[34] The view that preventers and omissions are no genuine causes but may still figure in true counterfactual claims about genuine causation has been defended by Dowe (2001).

The accounts of Halpern and Pearl (2005) and Halpern (2015) can be amended by a condition of normality (Halpern, 2016, pp. 79-81&90-1). A normality order over possible worlds allows them to represent the different views about the causal efficacy of omissions to be found in the literature. The amended accounts understand causation roughly as de facto—or de counterfacto—dependence *witnessed by a possible world which is at least as normal as the actual one*. If Putin had watered my plant, it would not have died. True—but the world witnessing the counterfactual is less normal than the actual world, where Putin doesn't water my plant. By contrast, the world in which my neighbour waters the plant is at least as normal as the actual world, where she does not. Hence, my neighbour's failure to water the plant is a cause of its dying, whereas Putin's failure is not. The latter shows that counterfactual dependence between actual events and absences is not sufficient for causation on the amended accounts—and it is not necessary as we knew before. On the positive side, the omission scenario of Figure 6 poses no longer a problem due to the normality condition.

The amended accounts can also solve bogus prevention. Recall that this scenario is isomorphic to overdetermination. Hence, Halpern and Pearl's (2005) account wrongly identifies bodyguard's administering of antidote and assassin's refraining from poisoning target's coffee as individual causes of target's survival. Crucial is the de counterfacto conditional 'if Bodyguard had not administered the antidote and assassin had put in the poison, target would have died'. The world witnessing the de counterfacto dependence is the one where Bodyguard does not put in the antidote, assassin puts in the poison and target dies. We have no clear intuition whether this world is more or less normal than the actual world. Halpern (2016, pp. 88-9) uses this lack of clarity: under the assumption that not putting anything in target's coffee is most normal, he declares the actual world *incomparable* to the witness world, and so the bodyguard's administration of antidote is no cause of target's survival on the amended account.

The result is similar for Halpern's (2015) modified account. The above de counterfacto conditional expresses at the same time the simple counterfactual dependence of target's survival on the set containing both bodyguard's and assassin's actions. The witness world is the same and it may still be incomparable to the actual world. Hence, bodyguard's administering of antidote is not part of a cause of target's survival on the amended account. One must wonder, however, why the witness world is not just as normal as the actual world and so at least as normal as the actual world. This would also explain the lack of clarity whether the witness world is more or less normal than the actual world. Besides this worry, the strategy to compare the normality of worlds looks promising so far.

On both Halpern-Pearl accounts, the dislodged boulder counts as a cause of the hiker's remaining unscathed. The reason is the de facto counterfactual 'if the boulder had not been dislodged but it would still roll toward the hiker, the hiker would have been hit'. It seems that the witness world—where the boulder has not been dislodged, but still rolls toward the hiker, the hiker does not duck, and so the boulder hits her—is less normal than the actual world. After all, it is a causally impossible world. Halpern (2016, p. 80) concurs by assuming that worlds which satisfy the structural equations are more normal than worlds which do not. If so, the dislodged boulder does not count as a cause of the hiker's remaining unscathed. This seems like a win for the amended accounts!

Let's not be too hasty, however. Consider a modification of the causal model of the boulder scenario by replacing the structural equation of B with $B = F \vee F'$ and adding $\neg F'$ to the set V of literals—or set of value assignments if you prefer. $\neg F'$ stands for the absence of another boulder being dislodged. This modification just adds an absence to the causal model and so does not change the actual scenario. Hence, the dislodged boulder F is still not a cause of the hiker's remaining unscathed $\neg E$—as our analysis says. And yet, F is now a cause of $\neg E$ on Halpern and Pearl's (2005) amended account. Crucial is the de counterfacto conditional $\neg F \wedge F' \mathrel{\Box\!\!\rightarrow} E$: 'if the boulder had not been dislodged but another boulder had been dislodged, the hiker would have been hit'. In the witness world, the boulder is not dislodged, the other boulder is dislodged, a boulder rolls toward the hiker, the hiker does not duck, and so is hit by the boulder. This witness world is causally possible and seems at least as normal as the actual world. It follows from the same conditional and witness world that the dislodged boulder F is part of a cause of the hiker's remaining unscathed $\neg E$ on Halpern's (2015) amended account—the other part is the absence of another boulder $\neg F'$. This seems odd.

The modified account amended by a normality condition also faces troubles in the preemption scenarios. Recall that de facto accounts identify Suzy's throw as a cause of the window's shattering because her throw makes a counterfactual difference to the window's shattering when holding fixed that Billy's throw does not touch the window. In the witness world, Suzy does not throw, Billy does with unfailing accuracy, but his rock somehow does not touch the window, neither does Suzy's, and so the window does not shatter. It seems that this witness world is less normal than the actual world because it violates the structural equations. A natural solution to preemption is lost.

Suzy's throw still counts as part of a cause of the window's shattering—the other part being Billy's throw. The counterfactual 'if neither Suzy nor Billy had thrown a rock, the window would not have shattered' is true. Nothing happens in the causally possible witness world. This world seems to be at least as normal as the actual world. Hence, Billy's throw is part of a cause of the window's shattering on Halpern's (2015) amended account. Halpern and Pearl's (2005) earlier account amended by the normality condition has no such devastating consequence.

We have looked at accounts motivated by the idea that causes are counterfactual difference-makers. The simple account of counterfactual difference-making succumbs to the problem of redundant causation. This motivates dropping the necessity of counterfactual difference-making between actual events and absences for causation. The simple de facto and de counterfacto accounts do so but retain its sufficiency for causation. As a consequence, all omissions count as causes. This motivates also dropping the sufficiency of counterfactual difference-making. The amended causal model accounts, which rely on comparing the normality between possible worlds, drop both necessity and sufficiency of counterfactual difference-making. One must wonder what remains of the guiding idea that causes *are* counterfactual difference-makers.

We have seen that sophisticated accounts in terms of counterfactual difference-making solve the problem posed by redundant causation.

However, the extant solutions create new problems. Lewis's (1973a) imposition of transitivity on counterfactual difference-making to solve early preemption creates the problem that the dislodged boulder counts as a cause of the hiker's remaining unscathed. The sophisticated de facto and de counterfacto accounts of Yablo (2002), Hitchcock (2001), Halpern and Pearl (2005) and Halpern (2015) provide the same problematic verdict among others. Amending the Halpern-Pearl accounts by a condition of normality solves the boulder scenario but not a factual equivalent thereof. The extant solutions to the problem of redundant causation consistently lead to new problems. In this sense, redundant causation still haunts accounts in terms of counterfactual difference-making.

The de facto and de counterfacto accounts need to answer what facts and counterfacts can be held fixed when testing for counterfactual difference-making. A satisfying answer is wanting. The extant proposals rely on comparative naturalness and comparative normality. But these notions are not yet worked out in sufficient detail. And so the accounts amended by them do not always provide clear verdicts. All of the problems taken together shed doubt on whether accounts of causation in terms of counterfactual difference-making are viable. Our analysis in terms of factual difference-making, by contrast, stays true to its guiding idea, does not create new problems by solving redundant causation, and gives clear verdicts. It solves all of the considered scenarios—and many more.

7 Conclusion

We have analysed token causation in terms of factual difference-making. In a nutshell, causation is deviant factual difference-making in a forward-directed way. Our analysis solves the problem of redundant causation without further problems. The underlying reason is that preempting causes and overdeterminers are factual difference-makers—even though they are not counterfactual difference-makers. Scenarios of redundant causation show that there are causes which do not make a counterfactual difference to their effects. By contrast, all causes make a factual difference to their effects. Factual difference-making is necessary for causation, while counterfactual difference-making is not. This allows accounts of causation to stay true to the guiding idea of factual difference-making but not to its counterfactual cousin.

Our analysis does not require counterfactual dependence for causation. Instead it relies on factual conditionals, which assume facts in their antecedents—unlike counterfactual conditionals. As there is no need for contrary-to-fact events and absences in the antecedent of a factual conditional, our analysis can stay clear of 'miracles'—violations of law or the like—and the accompanying challenges. Our analysis merely requires factual difference-making: in a state unsettled on cause and effect, if the cause becomes actual, so does the effect.

The extant solutions to the problem of redundant causation from accounts of causation in terms of counterfactual difference-making lead to further problems, such as unintuitive or unclear verdicts, or the need for a mereology which allows for disjunctive events or parts of causes. Our analysis faces no such issues. We have, of course, no proof that accounts in terms of counterfactual difference-making are beyond repair. But we have amassed enough reasons for taking another kind of difference-making seriously—and enough reasons for thinking that accounts in terms of factual difference-making may very well advance our understanding of causation. Our analysis can be seen as an example of such an account—an example to be improved upon.

Perhaps we shouldn't think of causes as counterfactual difference-makers. After all, counterfactual difference-making is neither necessary nor sufficient for causation—provided redundant causes and some omissions are genuine causes. Perhaps we should think of causes as factual difference-makers.

Acknowledgments. We would like to thank Cei Maslen, Edwin Mares, Philip Pettit, and Christopher Hitchcock for very helpful comments. We are grateful for the opportunity to present parts of this work at the 11th European Conference for Analytic Philosophy at the University of Vienna in 2023 and the Pacific Division Meeting of the American Philosophical Association in 2022. Mario Günther was supported by a Junior Residency in the Center for Advanced Studies and a Junior Research Fund by LMUexcellent, which is funded by the Federal Ministry of Education and Research (BMBF) and the Free State of Bavaria under the Excellence Strategy of the Federal Government and the Länder.

Disclosure Statement. The authors report there are no competing interests to declare.

References

Andreas, Holger , Armgardt, Matthias , and Günther, Mario (2023). 'Counterfactuals for Causal Responsibility in Legal Contexts'. *Artificial Intelligence and Law* **31**(1): 115–132. doi:10.1007/s10506-021-09307-2.

Andreas, Holger and Günther, Mario (2019). 'On the Ramsey Test Analysis of "Because"'. *Erkenntnis* **84**(6): 1229–1262. doi:10.1007/s10670-018-0006-8.

Andreas, Holger and Günther, Mario (2020). 'Causation in Terms of Production'. *Philosophical Studies* **177**(6): 1565–1591. doi:10.1007/s11098-019-01275-3.

Andreas, Holger and Günther, Mario (2021). 'A Ramsey Test Analysis of Causation for Causal Models'. *The British Journal for the Philosophy of Science* **72**(2): 587–615. doi:10.1093/bjps/axy074.

Andreas, Holger and Günther, Mario (2021). 'Difference-Making Causation'. *Journal of Philosophy* **118**(12): 680–701. doi:10.5840/jphil20211181243.

Andreas, Holger and Günther, Mario (2024). 'A Lewisian Regularity Theory'. *Philosophical Studies* doi:10.1007/s11098-024-02149-z.

Andreas, Holger and Günther, Mario (2024). 'A Regularity Theory of Causation'. *Pacific Philosophical Quarterly* **105**(1): 2–32. doi:10.1111/papq.12447.

Andreas, Holger and Günther, Mario (forthcoming). 'Actual Causation'. *dialectica* .

Beebee, Helen (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.

Coady, David (2004). ‘Preempting Preemption’. In *Causation and Counterfactuals*, edited by J. Collins, N. Hall, and L. Paul, MIT Press. 325–340.

Dowe, P. (2001). ‘A Counterfactual Theory of Prevention and “Causation” by Omission’. *Australasian Journal of Philosophy* **79**(2): 216–226. doi:10.1080/713659223.

Gallow, Dmitri J. (2021). ‘A Model-Invariant Theory of Causation’. *The Philosophical Review* **130**(1): 45–96. doi:10.1215/00318108-8699682.

Goodman, Nelson (1947). ‘The Problem of Counterfactual Conditionals’. *The Journal of Philosophy* **44**(5): 113–128. doi:10.2307/2019988.

Hall, Ned (2004). ‘Two Concepts of Causation’. In *Causation and Counterfactuals*, edited by J. Collins, N. Hall, and L. Paul, MIT Press. pp. 225–276.

Hall, Ned (2007). ‘Structural Equations and Causation’. *Philosophical Studies* **132**(1): 109–136. doi:10.1007/s11098-006-9057-9.

Halpern, Joseph (2016). *Actual Causality*. Cambridge, MA: MIT Press.

Halpern, Joseph Y. (2000). ‘Axiomatizing Causal Reasoning’. *Journal of Artificial Intelligence Research* **12**(1): 317–337. doi:10.1613/jair.648.

Halpern, Joseph Y. (2008). ‘Defaults and Normality in Causal Structures’. In *Proceedings of the Eleventh International Conference on Principles of Knowledge Representation and Reasoning*, edited by G. Brewka and J. Lang. Menlo Park, CA: AAAI Press, 198–208.

Halpern, Joseph Y. (2015). ‘A Modification of the Halpern-Pearl Definition of Causality’. *Proc. 24th International Joint Conference on Artificial Intelligence (IJCAI 2015)* : 3022–3033.

Halpern, Joseph Y. and Hitchcock, Christopher (2015). ‘Graded Causation and Defaults’. *British Journal for the Philosophy of Science* **66**(2): 413–457. doi:10.1093/bjps/axt050.

Halpern, Joseph Y. and Pearl, Judea (2005). 'Causes and Explanations: A Structural-Model Approach. Part I: Causes'. *British Journal for the Philosophy of Science* **56**(4): 843–887. doi:10.1093/bjps/axi147.

Hiddleston, Eric (2005). 'Causal Powers'. *British Journal for the Philosophy of Science* **56**(1): 27–59. doi:10.1093/phisci/axi102.

Hitchcock, Christopher (2001). 'The Intransitivity of Causation Revealed in Equations and Graphs'. *Journal of Philosophy* **98**(6): 273–299. doi:10.2307/2678432.

Hitchcock, Christopher (2007a). 'Prevention, Preemption, and the Principle of Sufficient Reason'. *Philosophical Review* **116**(4): 495–532. doi:10.1215/00318108-2007-012.

Hitchcock, Christopher (2007b). 'What's Wrong with Neuron Diagrams?'. In *Causation and Explanation*, edited by J. K. Campbell, M. O'Rourke, and H. S. Silverstein, MIT Press. pp. 4–69.

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

Lewis, D. (1973b). *Counterfactuals*. Oxford: Blackwell.

Lewis, David (1986). 'Postscripts to "Causation"'. In *Philosophical Papers. Volume II*, edited by D. Lewis, Oxford University Press. pp. 172–213.

Lewis, David (2000). 'Causation as Influence'. *Journal of Philosophy* **97**(4): 182–197. doi:10.2307/2678389.

Maslen, Cei (2004). 'Causes, Contrasts, and the Nontransitivity of Causation'. In *Causation and Counterfactuals*, edited by N. Hall, L. A. Paul, and J. Collins, Mass.: Mit Press. 341–357.

Maslen, Cei (2020). 'Causation, Absences, and the Prince of Wales'. *Synthese* **197**(11): 4783–4794. doi:10.1007/s11229-015-0980-1.

McGrath, Sarah (2005). 'Causation By Omission: A Dilemma'. *Philosophical Studies* **123**(1-2): 125–148. doi:10.1007/s11098-004-5216-z.

Paul, Laurie and Hall, Ned (2013). *Causation: A User's Guide*. Oxford. doi:10.1093/acprof:oso/9780199673445.001.0001.

Pearl, Judea (2009). *Causality: Models, Reasoning and Inference*. New York, NY, USA: Cambridge University Press, 2nd edn. doi:10.1017/CBO9780511803161.

Ramachandran, Murali (1997). 'A Counterfactual Analysis of Causation'. *Mind* **106**(422): 263–277. doi:10.1093/mind/106.422.263.

Sartorio, Carolina (2006). 'Disjunctive Causes'. *Journal of Philosophy* **103**(10): 521–538. doi:10.5840/jphil20061031029.

Sartorio, Carolina (2010). 'The Prince of Wales Problem for Counterfactual Theories of Causation'. In *New Waves in Metaphysics*, edited by A. Hazlett, London: Palgrave Macmillan UK. 259–276. doi:10.1057/9780230297425_14.

Schaffer, Jonathan (2000). 'Trumping Preemption'. *Journal of Philosophy* **97**(4): 165–181. doi:10.2307/2678388.

Woodward, James (2003). *Making Things Happen : A Theory of Causal Explanation*. Oxford University Press. doi:10.1093/0195155270.001.0001.

Yablo, Stephen (2002). 'De Facto Dependence'. *Journal of Philosophy* **99**(3): 130–148. doi:10.2307/3655640.

Yablo, Stephen (2004). 'Advertisement for a Sketch of an Outline of a Proto-Theory of Causation'. In *Causation and Counterfactuals*, edited by N. Hall, L. A. Paul, and J. Collins, Cambridge: MIT Press. 119–137. doi:10.1093/acprof:oso/9780199266487.003.0005.