

A Ramsey Test Analysis of Causation for Causal Models

Holger Andreas Mario Günther

University of British Columbia LMU Munich, University of Regensburg

November 27, 2018

Abstract

We aim to devise a Ramsey Test analysis of actual causation. Our method is to define a strengthened Ramsey Test for causal models. Unlike the accounts of Halpern and Pearl (2005) and Halpern (2015), the resulting analysis deals satisfactorily with both overdetermination and conjunctive scenarios.

Keywords. Causal Models, Ramsey Test, Actual Causation, Halpern-Pearl Definitions.

Contents

1	Introduction	2
2	An Extension of Causal Model Semantics	4
2.1	Halpern and Pearl’s Causal Model Semantics	4
2.2	Agnostic Models	9
3	A Strengthened Ramsey Test for Causal Models	10
3.1	The Ramsey Test and Causal Models	11
3.2	A Strengthened Ramsey Test for Causal Models	11
3.3	A Ramsey Test Definition of Actual Causation	15
3.4	Minimality	16

4	Applying the Definition of Actual Causation	16
4.1	Overdetermination and Conjunctive Scenarios	17
4.2	Preemption	19
4.3	Conjunctive and Disjunctive Scenarios Combined	22
5	Comparison to the Halpern-Pearl Definitions	24
5.1	Conjunctive Scenarios	24
5.2	Preemption	28
5.3	Symmetric Overdetermination	29
5.4	Summary	30
5.5	A Note on Causal Models	31
6	Conclusion	32

1 Introduction

We propose a Ramsey Test analysis of actual causation in the framework of causal models. The basic idea is to define a Ramsey Test conditional for causal models. The evaluation recipe of this novel Ramsey Test conditional \gg can be informally stated as follows:

First, suspend judgment about the antecedent and the consequent. Second, add the antecedent (hypothetically) to your stock of beliefs. Finally, consider whether or not the consequent is entailed by your beliefs.¹

In brief, the conditional $A \gg C$ should be believed iff, after suspending judgment on A and C , C is believed as a result of assuming A . The conditional \gg gives rise to an analysis of causation via the following template:

C is a cause of E iff C and E occur, and $C \gg E$.

¹This recipe is motivated in Andreas and Günther (2018).

The logical foundations of the belief changes that define \gg can be explicated using AGM-style belief revision theory, as founded by Alchourrón et al. (1985) and Gärdenfors (1988). Using some concepts of belief revision, we define in this paper the conditional \gg in the framework of causal models. Thereby, the above template becomes a surprisingly powerful analysis of causation.

Why analyze causation by a Ramsey Test conditional in the framework of causal models? After all, Halpern and Pearl (2005) already put forward a formally precise definition of actual causation that captures a wide range of causal scenarios, including troublesome cases of preemption. However, as Halpern and Pearl themselves admit, their definition in (2005) is “complicated” (ibid., p. 880). Moreover, even this highly elaborate definition runs into a problem when it comes to analysing conjunctive scenarios. In such scenarios, two events are necessary for an effect to occur. If lightning and a preceding drought, for example, are necessary factors for a forest fire to occur, the conjunction of these factors should qualify as an – if not ‘the’ – actual cause. However, Halpern and Pearl’s (2005) definition does not count the conjunction of both events to be an actual cause, while both events individually are recognized as actual causes.²

Halpern’s (2015) modification of the Halpern and Pearl (2005) definition is more elegant.³ However, as we shall see later on, the modified HP definition inherits the problem of the above mentioned conjunctive scenario. The modification comes with another cost. In scenarios of overdetermination, the modified HP definition does not recognize the overdetermining causes as actual causes. (Surprisingly, the modified HP definition classifies the conjunction of the overdetermining causes as an actual cause.) We observe that no Halpern-Pearl definition satisfactorily solves both conjunctive scenarios and cases of overdetermination. We show that this predicament can be remedied in the framework of causal models by our Ramsey Test analysis of causation.

The plan of our investigation is straightforward: we work upward from causal models via the strengthened Ramsey Test to an analysis of causation. Section 2 extends Halpern and Pearl’s causal model semantics by the notion of an agnostic model. Based on this extension, Section 3 presents our strengthened Ramsey Test and our

²As proven by Eiter and Lukasiewicz (2002), the predecessor definition of Halpern and Pearl (2005) in Halpern and Pearl (2001) runs into the even more general problem: *any* conjunction of two (or more) events are precluded to be an actual cause. We would like to thank an anonymous referee for pointing out the difference between the definitions of 2001 and 2005.

³Henceforth, we refer to the Halpern and Pearl (2005) definition also as ‘HP definition’ and to Halpern’s (2015) modification as ‘modified HP definition’.

analysis of actual causation. Section 4 applies the analysis to overdetermination, conjunctive scenarios, preemption, and a combination of a conjunctive and disjunctive scenario. Section 5 compares our Ramsey Test analysis to the Halpern-Pearl definitions of actual causation.

2 An Extension of Causal Model Semantics

We extend Halpern and Pearl’s (2005) causal model semantics by the notion of an agnostic model. Such models will represent the suspension of judgment in our strengthened Ramsey Test. Upon a brief introduction of causal models, we shall define the notion of an agnostic model and its corresponding operator of suspension.

2.1 Halpern and Pearl’s Causal Model Semantics

The semantics of conditionals due to Halpern and Pearl (2005, pp. 851-852) is defined with respect to a causal model over a signature.

Definition 1. Signature \mathcal{S}

A signature \mathcal{S} is a triple $\mathcal{S} = \langle \mathcal{U}, \mathcal{V}, \mathcal{R} \rangle$, where \mathcal{U} is a finite set of exogenous variables, \mathcal{V} is a finite set of endogenous variables, and \mathcal{R} maps any variable $Y \in \mathcal{U} \cup \mathcal{V}$ to a non-empty (but finite) set $\mathcal{R}(Y)$ of possible values for Y .

Definition 2. Causal Model M

A causal model over signature \mathcal{S} is a tuple $M = \langle \mathcal{S}, \mathcal{F} \rangle$, where \mathcal{F} maps each endogenous variable $X \in \mathcal{V}$ to a function $F_X : (\times_{U \in \mathcal{U}} \mathcal{R}(U)) \times (\times_{Y \in \mathcal{V} \setminus \{X\}} \mathcal{R}(Y)) \rightarrow \mathcal{R}(X)$.

The mapping \mathcal{F} defines a set of (modifiable) structural equations that model the causal influence of exogenous and endogenous variables on other endogenous variables. The function F_X determines the value of $X \in \mathcal{V}$ given the values of all the other variables in $\mathcal{U} \cup \mathcal{V}$. Note that \mathcal{F} defines no structural equation for any exogenous variable $U \in \mathcal{U}$.

A structural equation such as $x = F_X(\vec{u}, y)$ says (in a context where the exogenous variables take the values \vec{u}), if Y were set to y (by means of an intervention), then X would take on the value x . Notice the difference to a direct intervention on X : an intervention that assigns a value $x' \neq x$ to X (by means external to the model)

overrules the value x assigned by $F_X(\vec{u}, y)$. The difference points to a ‘causal’ asymmetry. In any structural equation, the values of the variables on the right-hand side jointly determine the value of the variable on the left-hand side; if, however, the value of the variable on the left-hand side is changed by means of an external intervention, the values of the variables on the right-hand side remain unaffected.

Halpern and Pearl (2005, p. 849) confine themselves to models of recursive structural equations.⁴ In such models, the causal dependences among the variables in \mathcal{V} can be represented by a directed acyclic graph. Hence, there is a strict partial order such that, if $X < Y$, then the value of X may affect the value of Y , but the value of Y cannot have any effect on the value of X . A consequence of the restriction to models M of recursive equations is that, given a context $\vec{U} = \vec{u}$, there is always a unique solution to the equations in M ; for we can always solve the equations in the order given by $<$.

In the spirit of Halpern and Hitchcock (2010, p. 397), we make two assumptions concerning the variables of a causal model: (i) no value of a variable X logically implies a value of another variable Y , and (ii) the different values of the same variable are mutually exclusive and jointly exhaustive.

Given a signature $\mathcal{S} = \langle \mathcal{U}, \mathcal{V}, \mathcal{R} \rangle$, we say that a formula of the form $X = x$, where $X \in \mathcal{U} \cup \mathcal{V}$ and $x \in \mathcal{R}(X)$, is a primitive value assignment. Halpern and Pearl (2005) and Halpern (2015) interpret primitive value assignments as ‘primitive events’.⁵ We write \vec{X} for a (finite) set $\{X_1, \dots, X_n\}$ of variables, and \vec{x} for a (finite) set $\{x_1, \dots, x_n\}$ of values of the variables in \vec{X} . We abbreviate the set of primitive value assignments $\{X_1 = x_1, \dots, X_n = x_n\}$ by $\vec{X} = \vec{x}$.⁶

The values of the variables in $M = \langle \mathcal{S}, \mathcal{F} \rangle$ are not determined. If the exogenous variables are not set to a value, the structural equations in \mathcal{F} are not triggered and thus no endogenous variable is assigned a value. Only if we set all the exogenous variables $U \in \mathcal{U}$ to exactly one of their values, the structural equations

⁴However, Halpern and Pearl (2005, pp. 883-884) provide a definition of actual causation for non-recursive models in their appendix.

⁵Although we find the interpretation of primitive value assignments as primitive events remarkable, we will relegate the discussion of this issue to another paper. For some challenging observations on this interpretation, see Hall (2007).

⁶We deviate from Halpern and Pearl’s presentation of causal models insofar as \vec{X} denotes a set of variables and not a vector of variables. However, Halpern and Pearl (2005, p. 849, footnote 1) are “implicitly identifying the vector \vec{X} with the subset of \mathcal{V} consisting of the variables in \vec{X} ”. Halpern (2015, footnote 2) uses the vector notation, but sometimes views “ \vec{Z} as a set of variables.” We take the sets as basic and implicitly view them as ordered in the right way.

propagate the causal influence from parent to child variables such that each endogenous variable $V \in \mathcal{V}$ is assigned exactly one value. In other words, if a context $\vec{u} = \{u_1, \dots, u_n\}$ is set for all n variables $U_i \in \mathcal{U}$, the structural equations recursively determine the value of each endogenous variable. The result is that in a contextualized causal model $\langle M, \vec{u} \rangle$ every variable in $\mathcal{U} \cup \mathcal{V}$ is assigned exactly one value.

In Halpern and Pearl’s (2005) semantics, a simple conditional $[Y = y]X = x$ is satisfied in a contextualized causal model $\langle M, \vec{u} \rangle$ if the intervention setting the variable Y to the value y results in X taking on the value x .⁷ Such an intervention is defined by the notion of a submodel $M_{Y=y}$ of M .

Definition 3. Submodel $M_{\vec{X}=\vec{x}}$

Let $M = \langle \mathcal{S}, \mathcal{F} \rangle$ be a causal model, \vec{X} a (possibly empty) set of variables in \mathcal{V} and \vec{x}, \vec{u} sets of values for the variables in \vec{X}, \vec{U} . We call the causal model $M_{\vec{X}=\vec{x}} = \langle \mathcal{S}_{\vec{X}}, \mathcal{F}^{\vec{X}=\vec{x}} \rangle$ over signature $\mathcal{S}_{\vec{X}} = \langle \mathcal{U}, \mathcal{V} \setminus \vec{X}, \mathcal{R} \upharpoonright_{(\mathcal{U} \cup \mathcal{V}) \setminus \vec{X}} \rangle$ a submodel of M .⁸ $\mathcal{F}^{\vec{X}=\vec{x}}$ maps each variable in $\mathcal{V} \setminus \vec{X}$ to a function $F_Y^{\vec{X}=\vec{x}}$ that corresponds to F_Y for the variables in $\mathcal{V} \setminus \vec{X}$, and sets the variables in \vec{X} to \vec{x} .

Let us consider an example of a causal model due to Halpern and Pearl (2005, p. 856):

Example 1. Arsonists Example

Suppose that two arsonists drop lit matches in different parts of a dry forest, and both cause trees to start burning. Consider two scenarios. In the first, called the *disjunctive scenario*, either match by itself suffices to burn down the whole forest. That is, even if only one match were lit, the forest would burn down. In the second scenario, called the *conjunctive scenario*, both matches are necessary to burn down the forest; if only one match were lit, the fire would die down before the forest was consumed.

We can describe the essential structure of these two scenarios using a causal model with four variables:

⁷We follow Fenton-Glynn (2017) in conveying “=” a double role: the sign figures as (i) identity and (ii) an assignment operator such as Halpern and Pearl’s (2005) “ \leftarrow ” or Pearl’s (2009) “do(\cdot)”. Although this double role is mathematically sloppy, it avoids an unnecessary multiplication of notation and is harmless as long as we keep in mind that $[\vec{Y} = \vec{y}] \phi$ expresses a conditional whose antecedent assigns the variables in \vec{Y} the values \vec{y} .

⁸ $\mathcal{R} \upharpoonright_{(\mathcal{U} \cup \mathcal{V}) \setminus \vec{X}}$ is the restriction of \mathcal{R} to the variables in $(\mathcal{U} \cup \mathcal{V}) \setminus \vec{X}$.

- an exogenous variable U that determines, among other things, the motivation and state of mind of the arsonists. For simplicity, assume that $\mathcal{R}(U) = \{u_{00}, u_{10}, u_{01}, u_{11}\}$; if $U = u_{ij}$, then the first arsonist intends to start a fire iff $i = 1$ and the second arsonist intends to start a fire iff $j = 1$. In both scenarios $U = u_{11}$;⁹
- endogenous variables ML_1 and ML_2 , each either 0 or 1, where $ML_i = 0$ if arsonist i does not drop the lit match and $ML_i = 1$ if he does, for $i = 1, 2$;
- an endogenous variable FB for forest burns down, with values 0 (it does not) and 1 (it does).

The two scenarios differ with respect to the structural equation for FB :

- Disjunctive scenario: $FB = F_{FB}(ML_1, ML_2) = \max(ML_1, ML_2)$.
- Conjunctive scenario: $FB = F_{FB}(ML_1, ML_2) = \min(ML_1, ML_2)$.

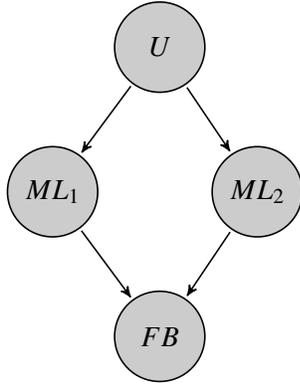


Figure 1: Causal network for the Arsonists Example. The arrows represent the dependences among the variables as encoded by the structural equations.

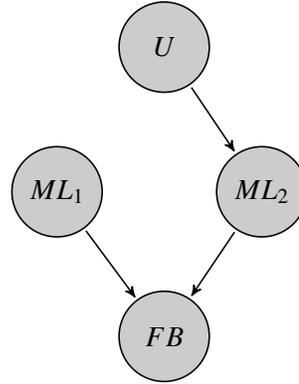


Figure 2: The causal network of the submodel for setting $ML_1 = 0$. The removal of the arrow between U and ML_1 represents that F_{ML_1} does not any more depend on the value of U .

Causal networks can be used to depict the recursive dependences of the structural equations, as in Figure 1 and Figure 2. A causal network is a directed acyclic graph,

⁹We will normally leave the context \vec{u} implicit in the following examples (as is common practice in causal modeling). See, for example, Halpern and Pearl (2005) and Halpern and Hitchcock (2010).

where the nodes correspond to the variables and there is an arrow from a node labeled X to one labeled Y iff X is a parent variable of Y iff F_Y directly depends on the value of X iff $X < Y$ and there is no variable Z such that $X < Z < Y$.¹⁰ The asymmetric dependence of structural equations can be read off a causal network. Each variable only causally affects its descendants. If Y is not a descendant of X , then a change in the value of X has no effect on the value of Y .

Let us now check whether or not the conditional $[ML_1 = 0]FB = 1$ is true in the contextualized causal model $\langle M, u_{11} \rangle$ of the Arsonists Example. The intervention that sets $ML_1 = 0$ induces a submodel $M_{ML_1=0}$ of M . If the solution to the structural equations of $M_{ML_1=0}$ satisfies $FB = 1$, then $[ML_1 = 0]FB = 1$ is true in the contextualized causal model. If the conditional is indeed satisfied in the contextualized causal model, we write $\langle M, u_{11} \rangle \models [ML_1 = 0]FB = 1$.

In the disjunctive scenario, $\langle M, u_{11} \rangle \models [ML_1 = 0]FB = 1$ iff $\langle M_{ML_1=0}, u_{11} \rangle \models FB = 1$.¹¹ The structural equations for the submodel $M_{ML_1=0}$ are:

- $F_{ML_1}^{ML_1=0} = 0$. Hence, $ML_1 = 0$.
- $F_{ML_2}^{ML_1=0}(u_{11}) = 1$. Hence, $ML_2 = 1$.
- $F_{FB}^{ML_1=0}(ML_1, ML_2) = \max(ML_1, ML_2)$. Since $ML_2 = 1$, $FB = 1$.

The solution of the structural equations of $M_{ML_1=0}$ satisfies $FB = 1$, and thus $\langle M, u_{11} \rangle \models [ML_1 = 0]FB = 1$. Notice the difference between the structural equation $F_{ML_1}(u_{11}) = 1$ and $F_{ML_1}^{ML_1=0} = 0$: the former depends on the value of U , whereas the latter does not. After the intervention that sets $ML_1 = 0$, the variable ML_1 is treated as if it were an exogenous variable, viz. it is assigned a value by its structural equation that does not depend on other variables. The structural equations for the other variables, i. e. $\mathcal{V} \setminus \{ML_1\}$, remain unchanged: $F_{ML_2}(u_{11}) = 1 = F_{ML_2}^{ML_1=0}(u_{11}) = 1$ and $F_{FB}(ML_1, ML_2) = F_{FB}^{ML_1=0}(ML_1, ML_2) = \max(ML_1, ML_2)$.

The conjunctive scenario differs from the disjunctive scenario only in the structural equation for FB : $F_{FB}^{ML_1=0}(ML_1, ML_2) = \min(ML_1, ML_2)$. Since $ML_1 = 0$, $FB = 0$. Hence, $\langle M, u_{11} \rangle \not\models [ML_1 = 0]FB = 1$. We conclude that the conditional $[ML_1 = 0]FB = 1$ is true in the disjunctive scenario and not true in the conjunctive scenario.

¹⁰Note that the acyclicity follows from the assumption that the structural equations are recursive.

¹¹Notice that the operator $[ML_1 = 0]$ relates $\langle M, u_{11} \rangle$ and $\langle M_{ML_1=0}, u_{11} \rangle$. In general, $[\]$ can be seen as both an operator on a (contextualized) causal model and as the antecedent of a conditional.

Before turning to our extension, we observe a correspondence between an arbitrary set of primitive value assignments $\vec{X} = \vec{x}$ and the conjunction of all members of $\vec{X} = \vec{x}$, i.e. $\bigwedge \vec{X} = \vec{x}$. Given a contextualized causal model $\langle M, \vec{u} \rangle$, and an arbitrary Boolean combination of value assignments ϕ , $\langle M, \vec{u} \rangle \models [\vec{X} = \vec{x}] \phi$ iff $\langle M, \vec{u} \rangle \models [\bigwedge \vec{X} = \vec{x}] \phi$, where $\bigwedge \vec{X} = \vec{x} = (X_1 = x_1 \wedge \dots \wedge X_i = x_i)$. Based on this correspondence, we will move back and forth between sets of primitive value assignments and their big conjunctions – without further mentioning. Moreover, by convention we take $\langle M, \vec{u} \rangle \models [\vec{X} = \vec{x}, \vec{Y} = \vec{y}] \phi$ to be synonymous with $\langle M, \vec{u} \rangle \models [\vec{X} = \vec{x} \cup \vec{Y} = \vec{y}] \phi$.

2.2 Agnostic Models

Our Ramsey Test requires a suspension of judgment. In order to define the suspension of judgment, we extend Halpern and Pearl’s (2005) causal model semantics by a new type of model, which we call ‘agnostic’. Agnostic models are meant to represent a suspension of judgment. Intuitively, the actual context is suspended in an agnostic model such that the value assignments are (partially) indeterminate: the values of some variables are suspended, while the values of other variables remain unchanged. Indeterminacy of a variable Y is expressed by assigning Y the set of all of its possible values. That is, an agnostic value assignment is a function \mathcal{A} that assigns each indeterminate variable Y the set $\mathcal{R}(Y) = \{y_1, \dots, y_n\}$ of its mutually exclusive and jointly exhaustive values. Notice that \mathcal{A} assigns sets of values to endogenous and, possibly, to exogenous variables.

An agnostic value assignment contrasts with a context of a causal model. A context \vec{u} uniquely determines the value assignment $\vec{u} \cup \vec{v}$ of the variables $\mathcal{U} \cup \mathcal{V}$ with respect to $\langle M, \vec{u} \rangle$. A *partially* agnostic value assignment \mathcal{A} renders indeterminate the value assignment $\vec{u} \cup \vec{v}$ of *some* of the variables in $\mathcal{U} \cup \mathcal{V}$ with respect to $\langle M, \vec{u} \rangle$.

Similar to Halpern and Pearl’s (2005) $[\]$ operator, which is based on the notion of a submodel, we now define a new operator $[\]_s$ based on the notion of an agnostic model.

Definition 4. Agnostic Model $\langle M_{\vec{X}=\vec{x}}^{\vec{y}}, \mathcal{A} \rangle$

Let $\langle M, \vec{u} \rangle$ be a contextualized causal model and $\vec{X} = \vec{x}$ a set of primitive value assignments for a (possibly empty) subset of the variables in $\mathcal{U} \cup \mathcal{V}$. We say that $\langle M_{\vec{X}=\vec{x}}^{\vec{y}}, \mathcal{A} \rangle$ is the agnostic model of $\langle M, \vec{u} \rangle$ with the core $\vec{X} = \vec{x}$ iff

- (1) $M_{\vec{X}=\vec{x}}$ is a submodel of M , and

(2) $\mathcal{A}(Y_i) = \mathcal{R}(Y_i)$ for all $Y_i \in (\mathcal{U} \cup \mathcal{V}) \setminus \vec{X}$.

An agnostic model $\langle M_{\vec{X}=\vec{x}}, \mathcal{A} \rangle$ with the core $\vec{X} = \vec{x}$ is agnostic with respect to the variables in the set $\vec{Y} = (\mathcal{U} \cup \mathcal{V}) \setminus \vec{X}$. The variables in \vec{X} are set to the values \vec{x} independently of the context value(s) for \vec{U} . Intuitively, the core $\vec{X} = \vec{x}$ contains the variable assignments that are protected from the suspension of judgment. In other words, we are not agnostic about the variables in the core. Hence, the variables of the causal model are partitioned into the variables \vec{Y} , on which we suspend judgment, and the variables \vec{X} , on which we do not suspend judgment.

We define satisfaction of a formula ϕ in an agnostic model as follows: $\langle M_{\vec{X}=\vec{x}}, \mathcal{A} \rangle \models \phi$ iff $\langle M_{\vec{X}=\vec{x}}, \vec{u} \cup \vec{v} \rangle \models \phi$ for any value assignment $\vec{u} \cup \vec{v}$ to the variables $\mathcal{U} \cup \mathcal{V}$ that respects the structural equations of $M_{\vec{X}=\vec{x}}$.¹² A value assignment $\vec{u} \cup \vec{v}$ respects the structural equations of $M_{\vec{X}=\vec{x}}$ just in case it assigns the variables in \vec{X} the values \vec{x} , and does not violate the structural equations of $M_{\vec{X}=\vec{x}}$. Conceptually, a value assignment respects the structural equations of a submodel iff the assignment agrees on the values that are set (by intervention) in the submodel, while being compatible with the remaining structural equations. Note that $M_{\vec{X}=\vec{x}} = M$ in the special case where \vec{X} is the empty set. We define the suspension operator $[\]_s$ as follows:

$$\langle M, \vec{u} \rangle [\vec{X} = \vec{x}]_s = \langle M_{\vec{X}=\vec{x}}, \mathcal{A} \rangle \quad (\text{Def } [\]_s)$$

The suspension operator $[\]_s$ renders the values of the variables $\vec{Y} = (\mathcal{U} \cup \mathcal{V}) \setminus \vec{X}$ indeterminate.

To summarize, the suspension of judgment on the variables $(\mathcal{U} \cup \mathcal{V}) \setminus \vec{X}$ with respect to a contextualized causal model $\langle M, \vec{u} \rangle$ results in an agnostic model $\langle M_{\vec{X}=\vec{x}}, \mathcal{A} \rangle$. After this ‘agnostic move’, the model $\langle M_{\vec{X}=\vec{x}}, \mathcal{A} \rangle$ satisfies ϕ iff any value assignment $\vec{u} \cup \vec{v}$ respecting the structural equations of the submodel $M_{\vec{X}=\vec{x}}$ satisfies ϕ .

3 A Strengthened Ramsey Test for Causal Models

We briefly introduce the idea of the Ramsey Test and uncover its relation to Halpern and Pearl’s (2005) causal model semantics. Subsequently, we define our strengthened Ramsey Test. By the template presented in the Introduction, this test gives us an analysis of causation.

¹²Henceforth, we simply write $\vec{u} \cup \vec{v}$ for a value assignment $\mathcal{U} \cup \mathcal{V} = \vec{u} \cup \vec{v}$.

3.1 The Ramsey Test and Causal Models

Ramsey (1929/1990) proposed a test to evaluate conditionals. When an agent has no opinion on a proposition A , she should believe a conditional ‘if A then C ’ if she believes C as a result of adding A to her stock of beliefs. Stalnaker (1968, p. 102) extended Ramsey’s test to cover the remaining epistemic attitudes with respect to the antecedent:

First, add the antecedent (hypothetically) to your stock of beliefs; second, make whatever adjustments are required to maintain consistency (without modifying the hypothetical belief in the antecedent); finally, consider whether or not the consequent is then true.

In brief, the recipe says that the conditional $A > C$ should be believed just in case C is believed as a result of assuming A .

Structurally, Stalnaker’s Ramsey Test resembles the evaluation of conditionals in Halpern and Pearl’s causal model semantics. For this to be seen, recall that the conditional $[Y = y]X = x$ is true if X takes on the value x as a result of setting the variable Y to the value y . Of course, the conditional $[Y = y]X = x$ is only true or false with respect to a contextualized causal model – very much like the Ramsey Test is relative to the beliefs of an epistemic state. The intervention in a causal model plays a role similar to the assumption of the antecedent in the Ramsey Test. The value of the variable X after the intervention corresponds to the belief in the consequent after the Ramsey Test revision. To be more precise, accepting the consequent after carrying out the Ramsey Test is comparable to the variable X taking on the value x according to the structural equations of the submodel $M_{Y=y}$. Notice that the structural equations of the submodel provide a precise meaning to the phrase ‘as a result of’. Hence, Halpern and Pearl’s causal model semantics may be regarded as containing a fully worked out version of the Ramsey Test. In the next section, we devise a strengthened Ramsey Test for causal models.

3.2 A Strengthened Ramsey Test for Causal Models

Let us recall our exposition of the strengthened Ramsey Test:

First, suspend judgment about the antecedent and the consequent. Second, add the antecedent (hypothetically) to your stock of beliefs. Fi-

nally, consider whether or not the consequent is entailed by your beliefs.

This strengthening of Ramsey’s (1929/1990) test is put forth by Andreas and Günther (2018). The idea is to test whether or not the antecedent may serve as a reason for the consequent. One motivation for strengthening the Ramsey Test has been to build a proper relation of relevance into the semantics of conditionals. That is, a conditional is verified only if the antecedent is relevant for the consequent. In contrast, the Ramsey Test semantics by Stalnaker (1968) and Gärdenfors (1988) already validate conditionals when the antecedent and consequent are believed, irrespective of whether or not the antecedent is relevant for the consequent. Moreover, following Rott (1986), Andreas and Günther (2018) use their strengthened Ramsey Test to analyze the conjunction ‘because’ of natural language.

We translate now this strengthened Ramsey Test into the framework of causal models. Let $Prim\langle M, \vec{u} \rangle$ denote the set of primitive value assignments that are satisfied in the contextualized causal model $\langle M, \vec{u} \rangle$. In symbols,

$$Prim\langle M, \vec{u} \rangle = \{X = x \mid \langle M, \vec{u} \rangle \models X = x\}.$$

Having this shorthand for the actual value assignments, how do we determine which beliefs or judgments are to be suspended in our strengthened Ramsey Test? Inspired by AGM belief revision, we define a contraction operator that turns a contextualized causal model into a partially agnostic model such that ϕ is not believed (any more) in the latter. We define the contraction operator via a remainder set of primitive value assignments:

Definition 5. Remainder Set $\langle M, \vec{u} \rangle \perp \phi$

Let $\langle M, \vec{u} \rangle$ be a contextualized causal model and ϕ a non-tautological Boolean combination of value assignments. We call $\langle M, \vec{u} \rangle \perp \phi$ the remainder set of $\langle M, \vec{u} \rangle$ with respect to ϕ :

$P \in \langle M, \vec{u} \rangle \perp \phi$ iff

- (1) $P \subseteq Prim\langle M, \vec{u} \rangle$, and
- (2) $\langle M, \vec{u} \rangle [P]_s \not\models \phi$, and
- (3) there is no P' such that $P \subset P' \subseteq Prim\langle M, \vec{u} \rangle$ and $\langle M, \vec{u} \rangle [P']_s \not\models \phi$.

Condition (1) simply demands that the beliefs that need not be suspended are actual beliefs. Intuitively, condition (2) demands that the contextualized causal model

$\langle M, \vec{u} \rangle$ does not satisfy ϕ after suspending judgment on all the variables of M , except the variables in P . Condition (3) says that each member P of the remainder set is maximal in the sense that adding an actual belief would revoke the suspension of judgment with respect to ϕ . A remainder set $\langle M, \vec{u} \rangle \perp \phi$ contains thus all maximal subsets P of $\text{Prim}\langle M, \vec{u} \rangle$ that do not satisfy ϕ with respect to the respective agnostic model $\langle M_P, \mathcal{A} \rangle$. Notice that each agnostic model $\langle M_P, \mathcal{A} \rangle$ is induced by a member P of the remainder set $\langle M, \vec{u} \rangle \perp \phi$, and thus represents one way to retain as many as possible actual variable assignments, while suspending judgment on ϕ . Moreover, note that any remainder set contains at least one element, which may well be the empty set.

Using the suspension operator and the definition of a remainder set, we define now the contraction operator $[\]_{\sigma}^-$. Let σ be a function that selects exactly one member P_{σ} of the remainder set $\langle M, \vec{u} \rangle \perp \phi$. A contraction of the contextualized causal model $\langle M, \vec{u} \rangle$ by the Boolean combination ϕ is then given by:

$$\langle M, \vec{u} \rangle [\phi]_{\sigma}^- = \langle M, \vec{u} \rangle [\sigma \langle M, \vec{u} \rangle \perp \phi]_s, \quad (\text{Def } [\]_{\sigma}^-)$$

where $\sigma \langle M, \vec{u} \rangle \perp \phi$ designates the member P_{σ} of the remainder set $\langle M, \vec{u} \rangle \perp \phi$ selected by σ . P_{σ} represents the set of primitive value assignments, as chosen by σ , that survived the suspension of judgment. For the sake of brevity, we will henceforth often refer directly to P_{σ} instead of $\sigma \langle M, \vec{u} \rangle \perp \phi$. In the language of AGM, we can say that an agnostic model is obtained by a maxichoice contraction operation (see Gärdenfors (1988, Ch. 4.2)).

We are now in a position to define our strengthened Ramsey Test conditional \gg . Let $P_{\sigma} = \sigma \langle M, \vec{u} \rangle \perp (\vec{X} = \vec{x} \vee \phi)$ for a σ that picks out exactly one member of the remainder set $\langle M, \vec{u} \rangle \perp (\vec{X} = \vec{x} \vee \phi)$. Then, $\vec{X} = \vec{x} \gg \phi$ iff there is a σ such that ϕ is entailed by the solution to the structural equations of the agnostic model $\langle M_{P_{\sigma}}, \mathcal{A} \rangle$, once \vec{X} is set to \vec{x} . The evaluation of $\vec{X} = \vec{x} \gg \phi$ comprises two steps: (i) contracting the set $\text{Prim}\langle M, \vec{u} \rangle$ of primitive value assignments such that we become agnostic as to whether or not $(\vec{X} = \vec{x}) \vee \phi$. This operation yields a set of agnostic models. (ii) Checking whether or not ϕ is entailed by the structural equations of some agnostic model $\langle M_{P_{\sigma}}, \mathcal{A} \rangle$ after setting \vec{X} to its original values \vec{x} . We define the strengthened Ramsey Test for the contextualized causal model $\langle M, \vec{u} \rangle$ as follows:

$$\begin{aligned} \langle M, \vec{u} \rangle \models \vec{X} = \vec{x} \gg \phi \text{ iff there is some } \sigma \text{ such that} \\ \langle M, \vec{u} \rangle [\vec{X} = \vec{x} \vee \phi]_{\sigma}^- \models [\vec{X} = \vec{x}] \phi \end{aligned} \quad (\text{SRT}_{SE})$$

Our strengthened Ramsey Test says that $\vec{X} = \vec{x} \gg \phi$ is satisfied in $\langle M, \vec{u} \rangle$ iff there is

a member P of the remainder set $\langle M, \vec{u} \rangle_{\perp} (\vec{X} = \vec{x} \vee \phi)$ such that $\langle M_P, \mathcal{A} \rangle \models [\vec{X} = \vec{x}] \phi$. In all but one of the examples below, the remainder set is a singleton. In such a case, there is only one selection function σ , which we may then leave implicit.

To further understand SRT_{SE} , it is instructive to note the following proposition and corollary.

Proposition 1. Let $\langle M, \vec{u} \rangle$ be a contextualized causal model, $\vec{X} = \vec{x}$ a set of primitive value assignments, and ϕ a Boolean combination of value assignments. Then $\langle M, \vec{u} \rangle \models \vec{X} = \vec{x} \gg \phi$ iff $\langle M_{P_{\sigma, \vec{X}=\vec{x}}}, \vec{u} \cup \vec{v} \rangle \models \phi$ for some σ and all value assignments $\vec{u} \cup \vec{v}$ that respect the structural equations of $M_{P_{\sigma, \vec{X}=\vec{x}}}$.

Proof. By SRT_{SE} , $\langle M, \vec{u} \rangle \models \vec{X} = \vec{x} \gg \phi$ iff there is some σ such that $\langle M, \vec{u} \rangle [\vec{X} = \vec{x} \vee \phi]_{\sigma} \models [\vec{X} = \vec{x}] \phi$. Abbreviating $\sigma \langle M, \vec{u} \rangle_{\perp} (\vec{X} = \vec{x} \vee \phi)$ by P_{σ} , the definition of $[\]_s$ gives us $\langle M_{P_{\sigma}}, \mathcal{A} \rangle \models [\vec{X} = \vec{x}] \phi$. By the definition of $[\]$, we obtain $\langle M_{P_{\sigma, \vec{X}=\vec{x}}}, \mathcal{A} \rangle \models \phi$. This means, by the definition of satisfaction in an agnostic model, nothing but $\langle M_{P_{\sigma, \vec{X}=\vec{x}}}, \vec{u} \cup \vec{v} \rangle \models \phi$ for all value assignments $\vec{u} \cup \vec{v}$ that respect the structural equations of $M_{P_{\sigma, \vec{X}=\vec{x}}}$. \square

Corollary 1. Let $\langle M, \vec{u} \rangle$ be a contextualized causal model, P_{σ} and $\vec{X} = \vec{x}$ sets of primitive value assignments, and ϕ a Boolean combination of value assignments. Then $\langle M, \vec{u} \rangle [P_{\sigma}]_s \models [\vec{X} = \vec{x}] \phi$ iff $\langle M, \vec{u} \rangle [P_{\sigma} \cup \vec{X} = \vec{x}]_s \models \phi$.

Proof. Both sides of the equivalence are equivalent to $\langle M_{P_{\sigma, \vec{X}=\vec{x}}}, \mathcal{A} \rangle \models \phi$. \square

In light of this proof, we adopt the following notational convention for any operator $[\]^*$ on causal models:

$$\langle M, \vec{u} \rangle [\dots]^* \models \phi \text{ iff } \langle M, \vec{u} \rangle \models [\dots]^* \phi.$$

In words, $[\]^*$ can be seen as both an operator on a contextualized causal model and as the antecedent of a conditional. In particular, we can present the strengthened Ramsey Test as a nested conditional:

$$\langle M, \vec{u} \rangle \models \vec{X} = \vec{x} \gg \phi \text{ iff there is some } \sigma \text{ such that } \langle M, \vec{u} \rangle \models [\vec{X} = \vec{x} \vee \phi]_{\sigma} [\vec{X} = \vec{x}] \phi.$$

3.3 A Ramsey Test Definition of Actual Causation

Based on the template of the Introduction, we define actual causation with respect to a contextualized causal model as follows:

Definition 6. $\vec{X} = \vec{x}$ is an Actual Cause of ϕ relative to $\langle M, \vec{u} \rangle$

Let $\langle M, \vec{u} \rangle$ be a causal model, $\vec{X} = \vec{x}$ a set of primitive value assignments, and ϕ a Boolean combination of value assignments. $\vec{X} = \vec{x}$ is an actual cause of ϕ relative to $\langle M, \vec{u} \rangle$ iff

$$\text{C1 } \langle M, \vec{u} \rangle \models \vec{X} = \vec{x} \wedge \phi, \text{ and}$$

$$\text{C2 } \langle M, \vec{u} \rangle \models \vec{X} = \vec{x} \gg \phi.$$

In essence, our definition says: (i) $\vec{X} = \vec{x}$ is an actual cause of ϕ iff $\vec{X} = \vec{x}$ and ϕ actually occur, and (ii) there is a model agnostic on $\vec{X} = \vec{x} \vee \phi$ such that assuming $\vec{X} = \vec{x}$ determines ϕ to be true.

Let us go a bit further into the details of our definition. According to C1, for $\vec{X} = \vec{x}$ to be an actual cause of ϕ , both $\vec{X} = \vec{x}$ and ϕ must occur in the actual context. C2 translates to $\langle M, \vec{u} \rangle [\vec{X} = \vec{x} \vee \phi]_{\sigma}^{-} \models [\vec{X} = \vec{x}] \phi$ for some σ . The suspension of judgment on the actual belief $\vec{X} = \vec{x} \vee \phi$ induces one agnostic model $\langle M_{P_{\sigma}}, \mathcal{A} \rangle$ for each member P_{σ} of the remainder set $\langle M, \vec{u} \rangle \perp (\vec{X} = \vec{x} \vee \phi)$. The beliefs in P_{σ} , respectively, are protected from the suspension of judgment.¹³ This protection is implemented by setting the structural equations of the variables in P_{σ} to their actual values, as expressed by $M_{P_{\sigma}}$. Each agnostic model $\langle M_{P_{\sigma}}, \mathcal{A} \rangle$ represents one way to retain as many as possible actual beliefs, while providing no information on the presumed cause and the putative effect. In any agnostic model, the actual primitive value assignments of $P_{\sigma} \subset \text{Prim}\langle M, \vec{u} \rangle$ are still believed, that is, they survive the suspension of judgment. These actual beliefs may thus influence our judgments on relations of actual causation.

The suspension of judgment is designed to retain maximal sets P_{σ} of actual beliefs that do not allow us to infer any information on presumed cause and effect in the respective agnostic model $\langle M_{P_{\sigma}}, \mathcal{A} \rangle$. Our test for a relation of actual causation is: can we infer ϕ by the structural equations of some agnostic model $\langle M_{P_{\sigma}}, \mathcal{A} \rangle$ when we assume $\vec{X} = \vec{x}$? Each such model partitions the variables $\mathcal{U} \cup \mathcal{V}$ of the original contextualized model $\langle M, \vec{u} \rangle$ into the protected variables in P_{σ} and the

¹³However, note that the set P_{σ} of protected beliefs is not always a strict superset of the empty set.

indeterminate ones. P_σ may therefore be interpreted as a set of background beliefs against which we test whether assuming $\vec{X} = \vec{x}$ lets us infer ϕ . If we succeed, $(\mathcal{U} \cup \mathcal{V}) \setminus P_\sigma$ contains the variables of a causal path from $\vec{X} = \vec{x}$ to ϕ , and, possibly, some remaining variables. Hence, we may conceive of an agnostic model as dividing the variables into (i) background, (ii) causal path, and (iii) remaining variables. The difference between causal path and remaining variables is that only the former receive determinate values when $\vec{X} = \vec{x}$ is assumed in the agnostic model.

Figuratively speaking, our definition searches for a suitable division of the variables in order to reproduce relations of actual causation. For the reproduction, some agnostic model must provide a suitable division of the variables into types (i), (ii), and (iii). A division is suitable iff the effect can be inferred by assuming the cause in the presence of the background beliefs. There is no relation of actual causation when for all divisions, either there is no causal path between presumed cause and putative effect, or there is such a path, but the cause does not determine the effect to be true.

3.4 Minimality

As we shall see below, we do not need any minimality condition to treat the canonical examples of overdetermination, preemption, and conjunctive scenarios. Nor do we need such a condition to treat the combination of a conjunctive with a disjunctive scenario discussed in Section 4.3. A minimality condition is here superfluous due to the simplicity of the examples and standard translations into causal models. Notably, causal models usually do not contain irrelevant variables. However, we can implement a minimality condition by refining C2 as follows:

$$\text{C2}' \quad \langle M, \vec{u} \rangle \models \vec{X}' = \vec{x}' \gg \phi \text{ for all } \vec{X}' = \vec{x}' \subseteq \vec{X} = \vec{x}.$$

C2' says that each subset of an actual cause needs to be an actual cause. If a subset $\vec{X}' = \vec{x}'$ is irrelevant, it follows that $\langle M, \vec{u} \rangle \not\models \vec{X}' = \vec{x}' \gg \phi$. Hence, C2' prevents that actual causes contain irrelevant conjuncts. A more thorough motivation and presentation of C2' will be given in a follow-up paper.

4 Applying the Definition of Actual Causation

We apply now our definition of actual causation to overdetermination, conjunctive scenarios, preemption, and a combination of a disjunctive and a conjunctive

scenario.

4.1 Overdetermination and Conjunctive Scenarios

Recall the Arsonists Example from Section 2.1. The causal model comprises three endogenous variables ML_1 , ML_2 , and FB , each taking either the value 0 or the value 1. In the actual context u_{11} , all three variables take the value 1. In the disjunctive scenario, each of ML_1 and ML_2 are individually sufficient for $FB = 1$. The disjunctive scenario is thus a scenario of symmetric overdetermination. The set of actual value assignments is $Prim\langle M, u_{11} \rangle = \{ML_1 = 1, ML_2 = 1, FB = 1\}$.

By Definition 6, $ML_1 = 1$ is an actual cause of $FB = 1$ with respect to the contextualized causal model $\langle M, u_{11} \rangle$ iff

C1 $\langle M, u_{11} \rangle \models ML_1 = 1 \wedge FB = 1$, and

C2 $\langle M, u_{11} \rangle \models ML_1 = 1 \gg FB = 1$.

C1 is obviously satisfied. We verify condition C2. By SRT_{SE} ,

$\langle M, u_{11} \rangle \models ML_1 = 1 \gg FB = 1$ iff there is σ such that
 $\langle M, u_{11} \rangle [ML_1 = 1 \vee FB = 1]_{\sigma}^- \models [ML_1 = 1]FB = 1$.

Abbreviating $\sigma \langle M, u_{11} \rangle \perp (ML_1 = 1 \vee FB = 1)$ by P_{σ} , the definition of $[\]_s$ gives us $\langle M_{P_{\sigma}}, \mathcal{A} \rangle \models [ML_1 = 1]FB = 1$. Notice that the remainder set $\langle M, u_{11} \rangle \perp (ML_1 = 1 \vee FB = 1)$ contains only one set P_{σ} of value assignments, viz. the empty set. For any member of $Prim\langle M, \vec{u} \rangle$, it is easy to show that $\langle M, u_{11} \rangle [P_{\sigma}]_s \models ML_1 = 1 \vee FB = 1$ – in violation of the suspension of judgment. This claim is entirely trivial for $P_{\sigma} = \{ML_1 = 1\}$ and $P_{\sigma} = \{FB = 1\}$. In the case of $P_{\sigma} = \{ML_2 = 1\}$, it is the structural equation for FB that determines $FB = 1$. Hence, in the agnostic model $\langle M_{P_{\sigma}}, \mathcal{A} \rangle$ of $\langle M, u_{11} \rangle$ no actual value assignment is retained. In other words, no beliefs are protected from the suspension of judgment in the disjunctive scenario.

It remains to check whether or not $\langle M_{P_{\sigma}}, \mathcal{A} \rangle \models [ML_1 = 1]FB = 1$ for $P_{\sigma} = \emptyset$. If we set ML_1 back to its original value 1, FB takes on the value 1 under each value assignment respecting the structural equations of $M_{\emptyset, ML_1=1}$. In symbols, $\langle M_{\emptyset, ML_1=1}, \mathcal{A} \rangle \models FB = 1$. We conclude that ML_1 is an actual cause of $FB = 1$ in the disjunctive scenario.

It should be noted that the disjunctive scenario is isomorphic to a scenario due to Hiddleston (2005, p. 32). This scenario is called ‘bogus prevention’ in Hitchcock

(2007, p. 523). In a case of bogus prevention, there is an event P that would prevent the event C to bring about a third event E , but C does not occur in the first place. Hence, P only vacuously prevents C from bringing about E . Yet, P counts as a cause of the non-occurrence of E on our analysis, as well as on the accounts of Hitchcock (2001) and Halpern and Pearl (2005). Counting P to be a cause of E , however, does not seem to be correct.¹⁴

As observed by Hiddleston (2005) and Hall (2007), the isomorphism between the disjunctive scenario and bogus prevention is an instance of a more general problem: there are pairs of scenarios that can be represented by isomorphic causal models, even though there is no isomorphism at the level of our causal judgments in the two scenarios. We shall deal with the problem of isomorphism and scenarios of prevention in a follow-up paper, exploiting there the distinction between the occurrence and the non-occurrence of an event. This distinction, however, is not explicitly captured by the purely formal semantics of causal models. For, it is only a matter of convention that the value 1 stands for an occurrent event, while 0 stands for the absence of an event.

Let us now move on to the conjunctive scenario of the Arsonists example. Here, both ML_1 and ML_2 need to take the value 1 for FB taking on the value 1. Accordingly, the structural equation for FB is $F_{FB} = \min(ML_1, ML_2)$, while the scenario remains otherwise unchanged. Now, the suspension of judgment on $ML_1 = 1 \vee FB = 1$ does not require to retract the belief in ML_2 since ML_1 alone is not sufficient to entail $FB = 1$ via the structural equations. However, there is still only one element in the remainder set $\langle M, u_{11} \rangle_{\perp}(ML_1 = 1 \vee FB = 1)$, viz. $P_{\sigma} = \{ML_2 = 1\}$. The reason is that there is a possible value assignment $\vec{u} \cup \vec{v} = \{ML_1 = 0, ML_2 = 1, FB = 0\}$ respecting the structural equations of $M_{P_{\sigma}}$ such that $\langle M_{P_{\sigma}}, \vec{u} \cup \vec{v} \rangle \not\models FB = 1$. Hence, $\langle M, u_{11} \rangle [P_{\sigma}]_s \not\models FB = 1$. $ML_2 = 1$ is thus a protected background belief in the conjunctive scenario.

It remains to check whether or not $\langle M_{P_{\sigma}}, \mathcal{A} \rangle \models [ML_1 = 1]FB = 1$. If we set ML_1 back to 1 in the agnostic model $\langle M_{P_{\sigma}}, \mathcal{A} \rangle$, FB takes on the value 1 under each value assignment respecting the structural equations of $M_{P_{\sigma}, ML_1=1}$. We conclude that $ML_1 = 1$ is an actual cause of $FB = 1$ in the conjunctive scenario.

By the same reasoning pattern, ML_2 is an actual cause of $FB = 1$ in both the disjunctive and conjunctive scenario due to the symmetry between ML_1 and ML_2 .

¹⁴For an interesting step towards a solution of bogus prevention, see Hitchcock (2007). For another, see Halpern (2008) and Halpern and Hitchcock (2015). The problem in question does not arise in Halpern (2015), at the price of the old overdetermination problem.

We leave it to the reader to verify that the conjunction $ML_1 = 1 \wedge ML_2 = 1$ is also an actual cause of $FB = 1$ in both scenarios.

4.2 Preemption

Let us consider the following case of preemption analyzed by Halpern and Pearl (2005, p. 861):

Example 2. Suzy & Billy Throw Rocks at a Bottle

Suzy and Billy both pick up rocks and throw them at a bottle. Suzy's rock gets there first, shattering the bottle. Since both throws are perfectly accurate, Billy's would have shattered the bottle had it not been preempted by Suzy's throw.

In such scenarios, the actual cause preempts a mere potential cause. Here, Suzy's throw is the actual cause of the bottle's shattering, which preempts Billy's throw. The key characteristic in this preemption case is: if Billy throws, the bottle will shatter – no matter what else is *actually* going on. Even if we counterfactually assume that Suzy is not throwing her rock at the bottle, knowing that Billy throws still allows us to infer that the bottle shatters. Nevertheless, Billy's throw does intuitively not count as an *actual* cause of the bottle's shattering; it is a mere potential cause of the shattering. Other scenarios of preemption have a similar structure.

We formalize the example following Halpern and Pearl (2005, pp. 861-864). The example comprises five endogenous variables:

- ST for 'Suzy throws' with two values: 0 (Suzy does not throw), 1 (she does).
- BT for 'Billy throws' with two values: 0 (Billy does not throw), 1 (he does).
- SH for 'Suzy's rock hits the bottle' with two values: 0 (it does not), 1 (it does).
- BH for 'Billy's rock hits the bottle' with two values: 0 (it does not), 1 (it does).
- BS for 'the bottle shatters' with two values: 0 (it does not), 1 (it does).

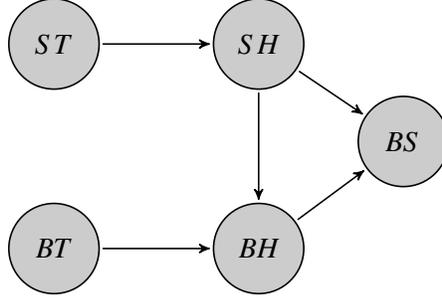


Figure 3: Causal network for the Suzy & Billy Throw Rocks at a Bottle Example.

We have the following structural equations:

- $F_{SH} = ST$.
- $F_{BH} = \min(BT, \neg SH)$. Note in particular $BH = 1$ iff $BT = 1$ and $SH = 0$.
- $F_{BS} = \max(SH, BH)$.

The actual context \vec{u} determines that $ST = BT = 1$. The ensuing set of actual value assignments is $\text{Prim}\langle M, \vec{u} \rangle = \{ST = 1, SH = 1, BS = 1, BT = 1, BH = 0\}$.

We check first whether Suzy's throw ($ST = 1$) is an actual cause of the bottle's shattering ($BS = 1$). By Definition 6, $ST = 1$ is an actual cause of $BS = 1$ with respect to the contextualized causal model $\langle M, \vec{u} \rangle$ iff

C1 $\langle M, \vec{u} \rangle \models ST = 1 \wedge BS = 1$, and

C2 $\langle M, \vec{u} \rangle \models ST = 1 \gg BS = 1$.

Condition C1 is obviously satisfied. We verify C2. By SRT_{SE} ,

$\langle M, \vec{u} \rangle \models ST = 1 \gg BS = 1$ iff there is some σ such that
 $\langle M, \vec{u} \rangle [ST = 1 \vee BS = 1]_{\sigma} \models [ST = 1]BS = 1$.

Abbreviating $\sigma \langle M, \vec{u} \rangle \perp (ST = 1 \vee BS = 1)$ by P_{σ} , the definition of $[\]_s$ gives us $\langle M_{P_{\sigma}}, \mathcal{A} \rangle \models [ST = 1]BS = 1$. By Definition 5, $P_{\sigma} = \{BT = 1, BH = 0\}$ is the only member of the remainder set $\langle M, \vec{u} \rangle \perp (ST = 1 \vee BS = 1)$. It is instructive to take a closer look at why this is so. Clearly, P_{σ} is a subset of $\text{Prim}\langle M, \vec{u} \rangle$. Furthermore, $\langle M, \vec{u} \rangle [P_{\sigma}]_s \not\models ST = 1 \vee BS = 1$. The reason is that there is a

possible value assignment $\vec{u} \cup \vec{v} = \{ST = 0, SH = 0, BS = 0, BT = 1, BH = 0\}$ that respects the structural equations of M_{P_σ} such that, in this value assignment, $ST = 1 \vee BS = 1$ is false. (Respecting the structural equations means here that the values assigned to BT and BH are 1 and 0, respectively, and the other values of the assignment do not violate the structural equations of the submodel M_{P_σ} .) Hence, $\langle M_{P_\sigma}, \mathcal{A} \rangle \not\models ST = 1 \vee BS = 1$. However, if an actual value assignment to one or more of the other variables would remain in P_σ , the agnostic model $\langle M_{P_\sigma}, \mathcal{A} \rangle$ would satisfy $ST = 1 \vee BS = 1$. Notice also that the set $\{BH = 0\}$ is a subset of $\{BT = 1, BH = 0\}$, and so it is excluded as a member of the remainder set $\langle M, \vec{u} \rangle \perp (ST = 1 \vee BS = 1)$ by the maximality condition of Definition 5. After suspending judgment, our agent still believes that Billy throws a rock but does not hit the bottle.

It remains to check whether or not $\langle M_{P_\sigma}, \mathcal{A} \rangle \models [ST = 1]BS = 1$, which is the case iff $\langle M_{P_\sigma, ST=1}, \vec{u} \cup \vec{v} \rangle \models BS = 1$ for each possible value assignment $\vec{u} \cup \vec{v}$ respecting the structural equations of $M_{P_\sigma, ST=1}$. The structural equations of $M_{P_\sigma, ST=1}$ ensure that $ST = 1$ and so $SH = 1$, which in turn leads to $BS = 1$. We obtain $\langle M_{P_\sigma, ST=1}, \mathcal{A} \rangle \models BS = 1$. Therefore, Suzy's throw is recognized as an actual cause of the bottle's shattering.

We demonstrate now that Billy's throw ($BT = 1$) is not an actual cause of the bottle's shattering ($BS = 1$). The reasoning is analogous to the above case, where Suzy throws. By Definition 6, $BT = 1$ is an actual cause of $BS = 1$ with respect to the contextualized causal model $\langle M, \vec{u} \rangle$ iff

C1 $\langle M, \vec{u} \rangle \models BT = 1 \wedge BS = 1$, and

C2 $\langle M, \vec{u} \rangle \models BT = 1 \gg BS = 1$.

Condition C1 is obviously satisfied. We falsify C2. By SRT_{SE} ,

$\langle M, \vec{u} \rangle \models BT = 1 \gg BS = 1$ iff there is some σ such that
 $\langle M, \vec{u} \rangle [BT = 1 \vee BS = 1]_\sigma^- \models [BT = 1]BS = 1$.

Abbreviating $\sigma \langle M, \vec{u} \rangle \perp (BT = 1 \vee BS = 1)$ by P_σ , the definition of $[\]_s$ gives us $\langle M_{P_\sigma}, \mathcal{A} \rangle \models [BT = 1]BS = 1$. This time $P_\sigma = \{BH = 0\}$ is the only element of the remainder set $\langle M, \vec{u} \rangle \perp (BT = 1 \vee BS = 1)$. $P_\sigma = \{BH = 0\}$ is maximal in the sense that $\langle M, \vec{u} \rangle [P_\sigma]_s \not\models BT = 1 \vee BS = 1$, whereas, for each strict superset $P' \supset P_\sigma$, $\langle M, \vec{u} \rangle [P']_s \models BT = 1 \vee BS = 1$. The reason is that there is the possible value assignment $\vec{u} \cup \vec{v} = \{ST = 0, BT = 0, SH = 0, BH = 0, BS = 0\}$ that respects

the structural equations of M_{P_σ} and does not satisfy $BT = 1 \vee BS = 1$. Hence, $\langle M, \vec{u} \rangle [P_\sigma]_s \not\models BT = 1 \vee BS = 1$. Adding any other primitive value assignment (or assignments) of the set of actual assignments to P_σ would result in an agnostic model that satisfies $BT = 1 \vee BS = 1$. That Billy’s rock does not hit the bottle is thus the only protected background belief.

Let us now determine whether or not $\langle M_{P_\sigma}, \mathcal{A} \rangle \models [BT = 1]BS = 1$. If we set $BT = 1$ in the agnostic model $\langle M_{P_\sigma}, \mathcal{A} \rangle$, the model does not satisfy $BS = 1$. The reason is that the structural equation for BH is set to its actual value 0, and so assuming the value of BT does not change the value of BS . In other words, the intervention that sets BH on its actual value cancels BH ’s dependence on BT (and SH). Consequently, the value of BH cannot be determined by the value of BT (and/or SH).

More formally, $\langle M_{P_\sigma}, \mathcal{A} \rangle \models [BT = 1]BS = 1$ iff $\langle M_{BH=0, BT=1}, \vec{u} \cup \vec{v} \rangle \models BS = 1$ for each possible value assignment $\vec{u} \cup \vec{v}$ respecting the structural equations of $M_{P_\sigma, BT=1}$. Since, however, $\vec{u} \cup \vec{v} = \{ST = 0, BT = 1, SH = 0, BH = 0, BS = 0\}$ respects the structural equations of $M_{P_\sigma, BT=1}$, we have $\langle M_{P_\sigma}, \mathcal{A} \rangle \not\models [BT = 1]BS = 1$. Hence, Billy’s throw is not an actual cause of the bottle’s shattering.

On our account, actuality influences whether or not there is an *actual* causal relation. The actual value assignment $BH = 0$ is exempt from the suspension of judgment, and thus remains in P_σ . In other words, the ‘actuality’ $BH = 0$ survives the suspension of judgment by setting the structural equation for BH to the value 0. Consequently, BH takes on its actual value in the agnostic model $\langle M_{P_\sigma}, \mathcal{A} \rangle$ independent of the values for BT and SH . This is an example of how certain parts of actuality may “intervene” to tell apart actual causes from merely potential causes.

4.3 Conjunctive and Disjunctive Scenarios Combined

So far, we have only considered examples where the relevant remainder set contains exactly one element. If we combine a conjunctive with a disjunctive scenario, however, we may have two elements in the remainder set. (Of course, there are various other scenarios, where the remainder set contains more than one element.) Let us consider such a combination of a conjunctive and a disjunctive scenario.

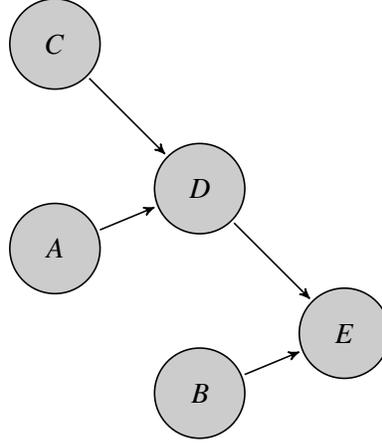


Figure 4: A combination of a conjunctive and disjunctive scenario.

We have the following structural equations:

- $F_D = \max(C, A)$
- $F_E = \min(D, B)$.

The actual context \vec{u} determines that $A = B = C = 1$. The ensuing set of actual value assignments is $\text{Prim}\langle M, \vec{u} \rangle = \{A = 1, B = 1, C = 1, D = 1, E = 1\}$. We check whether $C = 1$ is an actual cause of $E = 1$ relative to $\langle M, \vec{u} \rangle$. By Definition 6, for this to be the case, the following two conditions must be met:

C1 $\langle M, \vec{u} \rangle \models C = 1 \wedge E = 1$, and

C2 $\langle M, \vec{u} \rangle \models C = 1 \gg E = 1$.

Condition C1 is obviously satisfied. We verify C2. By SRT_{SE} ,

$\langle M, \vec{u} \rangle \models C = 1 \gg E = 1$ iff there is some σ such that
 $\langle M, \vec{u} \rangle [C = 1 \vee E = 1]_{\sigma}^- \models [C = 1]E = 1$.

Here, the remainder set $\langle M, \vec{u} \rangle \perp (C = 1 \vee E = 1)$ contains two elements, viz. $\{B = 1\}$ and $\{A = 1, D = 1\}$. Indeed, neither $B = 1$ nor $A = 1 \wedge D = 1$, individually, are sufficient to infer the value of E or C . Given the structural equation for E , if we have no information on the value of D , $B = 1$ does not allow us to infer the value

of E ; similarly, if we have no information on the value of B , $D = 1$ (together with $A = 1$) does not allow us to infer the value of E . Hence, we have two selection functions such that $P_{\sigma_1} = \{B = 1\}$ and $P_{\sigma_2} = \{A = 1, D = 1\}$. Accordingly, $\langle M, \vec{u} \rangle [C = 1 \vee E = 1]_{\sigma}^- = \langle M, \vec{u} \rangle [\sigma \langle M, \vec{u} \rangle \perp (C = 1 \vee E = 1)]_s$ depends on which σ_i (for $i = 1, 2$) we choose.

It remains to check whether or not $\langle M_{P_\sigma}, \mathcal{A} \rangle \models [C = 1]E = 1$ for at least one of the selection functions. Let us have a look at $P_{\sigma_1} = \{B = 1\}$ first. It is easy to see that the structural equations of $M_{B=1, C=1}$ ensure that $E = 1$ for each possible value assignment respecting the equations. We obtain $\langle M_{P_{\sigma_1}, C=1}, \mathcal{A} \rangle \models E = 1$. Therefore, $C = 1$ is an actual cause of $E = 1$ (relative to $\langle M, \vec{u} \rangle$).

Observe, however, that the submodel induced by σ_2 does not allow us to infer $E = 1$ from setting C to the value 1. The structural equations of $M_{A=1, D=1, C=1}$ do not ensure that $E = 1$. In fact, without $B = 1$, we cannot infer that $E = 1$. As we have just seen, not any member P of the remainder set $\langle M, \vec{u} \rangle \perp (C = 1 \vee E = 1)$ satisfies the condition that $\langle M_P, \mathcal{A} \rangle \models [C = 1]E = 1$. The flexibility of selection functions is crucial to obtain the correct results in the present combination of a conjunctive and disjunctive scenario. This is but one of many causal scenarios where the selection functions do substantial work. It can be shown, for instance, that the flexibility coming with the selection functions allows for a successful treatment of typical scenarios of double prevention.

5 Comparison to the Halpern-Pearl Definitions

Let us compare the Halpern-Pearl definitions of actual causation to our Ramsey Test analysis. We will have a look at conjunctive scenarios, preemption, and symmetric overdetermination, before we summarize the differences. We end the section with a note on causal models.

5.1 Conjunctive Scenarios

Recall the conjunctive scenario of the Arsonists Example, where both arsonists need to drop lit matches for the forest to burn down. In the actual context, both arsonists intend to start a fire. On our analysis, each individual arsonist dropping a lit match, respectively, counts as an actual cause of the forest burning down; moreover, the conjunction of the two arsonists dropping their lit matches is recognized

as an actual cause of the forest fire.¹⁵ Like our analysis, the definitions of Halpern and Pearl (2005, p. 857) and Halpern (2015) both recognize each arsonist individually as an actual cause. The two definitions have a minimality condition, AC3, in common. AC3 forbids an actual cause of some effect to be a strict superset of another actual cause of said effect. Hence, both definitions rule out the conjunction of the two arsonists as an actual cause of the forest fire because one (and indeed both) of its conjuncts (or ‘subsets’) already counts as an actual cause for this effect.

Let us take a closer look at the modified HP definition in Halpern (2015) so as to see why this definition fails to recognize conjunctive causes in conjunctive scenarios.

Definition 7. $\vec{X} = \vec{x}$ is an Actual Cause of ϕ relative to $\langle M, \vec{u} \rangle$

Let $\langle M, \vec{u} \rangle$ be a causal model, $\vec{X} = \vec{x}$ a set of primitive events and ϕ a Boolean combination of primitive events. $\vec{X} = \vec{x}$ is an actual cause of ϕ relative to $\langle M, \vec{u} \rangle$ iff

AC1 $\langle M, \vec{u} \rangle \models \vec{X} = \vec{x} \wedge \phi$, and

AC2 There is a set \vec{W} of variables in \mathcal{V} and a setting \vec{x}' of the variables in \vec{X} such that:

if $\langle M, \vec{u} \rangle \models \vec{W} = \vec{w}$, then $\langle M, \vec{u} \rangle \models [\vec{X} = \vec{x}', \vec{W} = \vec{w}] \neg \phi$.

AC3 \vec{X} is minimal; no subset of \vec{X} satisfies conditions AC1 and AC2.

According to this definition, $ML_1 = 1 \wedge ML_2 = 1$ (or rather $\{ML_1, ML_2\}$) is not an actual cause of $FB = 1$ with respect to the conjunctive scenario. As already indicated, the reason is AC3. There are subsets of $\{ML_1, ML_2\}$, e.g. $\{ML_1\}$, that satisfy AC1 and AC2. Let $\vec{W} = \{ML_2\}$. Consequently,

if $\langle M, \vec{u} \rangle \models ML_2 = 1$, then $\langle M, \vec{u} \rangle \models [ML_1 = 0, ML_2 = 1]FB = 0$,

which is true because $\langle M_{ML_1=0, ML_2=1}, \vec{u} \rangle \models FB = 0$.

The modified definition, so reads Halpern’s (2015) abstract, “gives reasonable answers (that agree with those of the [HP definition]) in the standard problematic examples”. The two definitions agree indeed in the conjunctive scenario. But the conjunction of the two events, which are necessary for the effect to occur, does not qualify as an actual cause. This result does not seem to be ‘reasonable’. Rather, it

¹⁵According to our definition, why is $ML_1 = 1$ an actual cause of $FB = 1$ in the conjunctive scenario? Informally, because the actual context has it that $ML_2 = 1$ and the suspension of judgment does not require to give up this belief.

seems that making an actual cause explicit in addition to another should not invalidate their conjunction as an actual cause. Merely bringing an event of the actual context from the background into consideration should not invalidate a relation of actual causation, given that this actually occurring event is a necessary factor for the effect to occur. We make reasonable claims like ‘the lightning *together* with the preceding drought caused the forest fire’, thereby implying that the lightning alone would not have been sufficient to cause the fire. This is a conjunctive causal claim since it emphasizes that it is two events that only jointly brought about the effect. Hence, it is desirable that our formal analysis captures causal statements of the form ‘ C and C' caused E '.

To further support the point that conjunctions in a conjunctive scenario should count as actual causes, consider the following example:

Example 3. Joe’s Death

Hannah puts poison into Joe’s coffee. However, the poison does not bring about any ill effects for Joe. Later, Joe takes his medication prescribed for a harmless infection he is suffering from. The medication interacts with the poison in his system. As a consequence, Joe dies.

What caused Joe’s death? Most naturally, we are inclined to say Hannah’s poisoning the coffee together with Joe’s taking his prescribed medication. Having no further information as regards the scenario, this conjunction is a sufficient actual cause for Joe’s death, while the conjuncts individually are not sufficient. Not only in a court of law, we are often interested in an actual cause of some event E that is sufficient for E according to some presupposed background, which excludes too far-fetched possibilities. In such cases, we tend to refer to a sufficient actual cause of E (relative to the background) simply as ‘the cause of E ’.¹⁶ Hence, when answering the above question, it is natural to explicitly cite both conjuncts, that is Hannah’s poisoning and Joe’s intake of prescribed medicine. These conjuncts do not ‘act independently’ with respect to Joe’s death, and are thus, individually, not sufficient for Joe’s death (relative to some appropriate model of the scenario). Rather, they are together ‘the’ cause of Joe’s death.¹⁷

¹⁶In the preemption example in Section 4.2, observe how natural it is to say “Suzy’s throw is *the* actual cause of the bottle’s shattering” (p. 19, our emphasis).

¹⁷Similarly, Hiddleston (2005, p. 43, footnote 9) argues that striking a match together with the presence of oxygen causes the match to light. In his theory of causal powers, he says that “the generative power” in the conjunctive match-oxygen example “attaches to the conjunction in the first instance. But often we can get away with treating some elements of a conjunctive cause as background enabling conditions.” (ibid.) If there is a relation of actual causation according to our analysis, we may

There are even concepts that, by definition, refer to an effect of a conjunctive cause. A case in point is the concept of a spring tide. Such a tide is caused by the joint gravitational forces of the moon and the sun. Hence, we would like to say that the gravitational forces of the moon, together with such forces of the sun, caused the spring tide. Again, this type of causal claim suggests allowing for conjunctive causes in the formal analysis of actual causation. While our analysis does not yet distinguish between

- (i) C and C' is an actual cause of E , and
- (ii) C together with C' is an actual cause of E ,

such a distinction could easily be introduced. The difference in meaning between the two causal claims is that (i) is indifferent as regards the distinction between disjunctive and conjunctive scenarios, while (ii) indicates that we are in a conjunctive scenario. Note that (ii) implies (i). Hence, if an analysis of causation were to account for (ii), (i) must be verified in a conjunctive causal scenario.

Let us now briefly sketch how our analysis could be extended to a coherent account of ‘ C together with C' is an actual cause of E ’. The decisive difference between the disjunctive and conjunctive scenario is that $C \wedge \neg C' \gg E$ holds true in the former but not in the latter. By this additional test condition, our analysis can distinguish between ‘ C together with C' is an actual cause of E ’ and ‘ C and C' is an actual cause of E ’. By contrast, in the context of the definitions by Halpern and Pearl (2005) and Halpern (2015), it is not feasible to draw the distinction between (i) and (ii) in a coherent way. The problem is that, in a conjunctive causal scenario, (ii) is supposed to hold true, (i) is not verified, and (ii) implies (i).¹⁸

Moreover, we think we should be allowed to make the following inference, at least in disjunctive and conjunctive scenarios:

$$\frac{C \text{ is an actual cause of } E, \quad C' \text{ is an actual cause of } E}{C \text{ and } C' \text{ is an actual cause of } E}$$

always consider a single conjunct to be an actual cause by relegating some other actual factors to the background. However, we may also make the other actual factors explicit by adding them as further conjuncts to the actual cause.

¹⁸Our analysis can even be used to identify sufficient actual causes. According to our analysis, let the variable assignment C be an actual cause of the variable assignment E relative to $\langle M, \vec{u} \rangle$. Then the conjunction of C together with all the members of P_σ is a sufficient actual cause of E relative to $\langle M, \vec{u} \rangle$, where P_σ is an element of the remainder set $\langle M, \vec{u} \rangle \perp (C \vee E)$ such that $\langle M_{P_\sigma}, \mathcal{A} \rangle \models [C]E$. We think it is an advantage of an account of actual causation if it can identify sufficient actual causes. The identification provides information about a set of actual factors that are jointly sufficient to bring about the effect, relative to the structural equations.

Why should this inference turn out to be valid in such scenarios? Do the premises say something different from the conclusion at all? In the absence of an analysis of ‘ C together with C' is an actual cause of E ’, there might not be any difference in content between the premises and the conclusions. So, suppose, at the deep semantic level, there is no such difference. Then, the inference in question is valid, just as inferring p from $\neg\neg p$ is valid on the semantics of classical logic. If there is a difference, this difference should not prevent us from judging ‘ C and C' is an actual cause of E ’ to be true in conjunctive and disjunctive scenarios. Note, finally, that the formal language of causal models has the expressive resources for conjunctive events. Hence, it is a meaningful question to ask whether or not the conjunction of C and C' is an actual cause of E . And the answer to this question should clearly be in the affirmative for conjunctive and disjunctive scenarios.¹⁹

5.2 Preemption

Let us consider our example of preemption again. We will find that the modified definition and our analysis use a similar strategy, viz. an intervention of actuality. According to the modified HP definition, Suzy’s throw ($ST = 1$) is an actual cause of the bottle’s shattering ($BS = 1$) with respect to $\langle M, \vec{u} \rangle$. AC1 and AC3 are obviously satisfied. As for AC2, let $\vec{W} = \{BH\}$. Consequently,

$$\text{if } \langle M, \vec{u} \rangle \models BH = 0, \text{ then } \langle M, \vec{u} \rangle \models [ST = 0, BH = 0]BS = 0$$

is satisfied, as $\langle M_{ST=0, BH=0}, \vec{u} \rangle \models BS = 0$. Notice that choosing \vec{W} to be $\{BH\}$ is tantamount to keeping BH fixed at its actual value 0. As the structural equation of BH is set to 0 in $\langle M_{ST=0, BH=0}, \vec{u} \rangle$, even if Billy throws ($BT = 1$), the bottle does not shatter. In our analysis, the actual value assignment $BH = 0$ survives the suspension of judgment, and thus the same intervention of actuality applies. However, in contrast to the modified HP definition, we are not free to simply choose some actual value assignments. The interventions of actuality, so to speak, are determined by the suspension of judgment.

The modified HP definition implies that Billy’s throw ($BT = 1$) is not an actual cause of the bottle’s shattering ($BS = 1$) with respect to $\langle M, \vec{u} \rangle$. The reason is that AC2 is violated: there are no variables that can be kept fixed at their actual values such that setting BT to the value 0 results in BS taking on the value 0. In the actual context \vec{u} , Suzy throws, her rock hits, and so the bottle shatters whether

¹⁹Many thanks to an anonymous referee for suggesting to discuss the above inference scheme.

or not Billy throws. This explanation of why we do judge Suzy's throw to be 'the' cause of the bottle's shattering but not Billy's derives from the same intuition as in our definition: it was her rock that *actually* hit the bottle and not his. If the remainder set is a singleton, then our analysis gives us the pertinent piece of actuality ($BH = 0$).

5.3 Symmetric Overdetermination

Let us turn our attention to the disjunctive scenario of the Arsonists Example, where each of the arsonists alone is sufficient for the forest to burn down. We show that, according to the modified HP definition, the conjunction $ML_1 \wedge ML_2 = 1$ qualifies as an actual cause of $FB = 1$. AC1 is satisfied, as $\langle M, \vec{u} \rangle \models ML_1 = 1 \wedge ML_2 = 1 \wedge FB = 1$. As for AC2, let $\vec{W} = \emptyset$. Consequently,

$$\langle M, \vec{u} \rangle \models [ML_1 = 0, ML_2 = 0]FB = 0 \text{ iff } \langle M_{ML_1=0, ML_2=0}, \vec{u} \rangle \models FB = 0.$$

This is satisfied as $F_{FB}^{ML_1=0, ML_2=0} = \max(ML_1, ML_2) = 0$. Interestingly, $\{ML_1, ML_2\}$ is minimal in the sense of AC3 as the two subsets $\{ML_1\}$ and $\{ML_2\}$ do not satisfy AC2. We show violation of AC2 for $\{ML_1\}$. Suppose, for contradiction, (i) $\{ML_2 = 1\}$ is an actual cause of $\{FB = 1\}$. Obviously, (ii) we need to consider two cases for \vec{W} : $\vec{W} = \emptyset$ and $\vec{W} = \{ML_2\}$. By AC2, (i) and (ii) imply that

$$\langle M, \vec{u} \rangle \models [ML_1 = 0]FB = 0,$$

or

$$\langle M, \vec{u} \rangle \models [ML_1 = 0, ML_2 = 1]FB = 0.$$

However, it is easy to see that neither $\langle M_{ML_1=0}, \vec{u} \rangle \models FB = 0$ nor $\langle M_{ML_1=0, ML_2=1}, \vec{u} \rangle \models FB = 0$ as $F_{FB}^{ML_1=0} = F_{FB}^{ML_1=0, ML_2=1} = \max(ML_1, ML_2) = 1$. Contradiction. By the same reasoning pattern, $\{ML_2\}$ does not satisfy AC2 due the symmetry between $\{ML_2\}$ and $\{ML_1\}$ in the Arsonists Example.

We have just seen that the conjunction $ML_1 = 1 \wedge ML_2 = 1$ satisfies the minimality condition AC3 on the modified HP definition. This definition counts none of $ML_1 = 1$ and $ML_2 = 1$ alone as a cause of the forest fire, even though they are individually sufficient to bring about this fire. Hence, the definition by Halpern (2015) fails to solve the problem of overdetermination.

5.4 Summary

The following tables summarize the differences between the considered definitions with respect to the Arsonists Example.

Conjunctive Scenario

Actual Cause of $FB = 1$	HP Definition	Modified HP Def.	RT Analysis
$ML_1 = 1$	Yes	Yes	Yes
$ML_2 = 1$	Yes	Yes	Yes
$ML_1 = 1 \wedge ML_2 = 1$	No	No	Yes

The table compares the HP definition, the modified HP definition, and our Ramsey Test analysis (RT Analysis) in the conjunctive scenario of the Arsonists Example. Remarkably, the HP definitions do not recognize the conjunction $ML_1 = 1 \wedge ML_2 = 1$ as an actual cause of $FB = 1$.

Disjunctive Scenario

Actual Cause of $FB = 1$	HP Definition	Modified HP Def.	RT Analysis
$ML_1 = 1$	Yes	No	Yes
$ML_2 = 1$	Yes	No	Yes
$ML_1 = 1 \wedge ML_2 = 1$	No	Yes	Yes

The table compares the HP definition, the modified HP definition, and our Ramsey Test analysis in the disjunctive scenario of the Arsonists Example. Contrary to the claim in Halpern’s (2015) abstract, the modified definition does not completely agree with Halpern and Pearl’s (2005) definition. This is troublesome because overdetermination is clearly a ‘standard problematic example’. Furthermore, the modified definition recognizes the conjunction $ML_1 = 1 \wedge ML_2 = 1$ as an actual cause of $FB = 1$ in the disjunctive scenario. Halpern’s (2015) attempt to explain this conundrum is barely convincing. He resorts to the claim that if $\vec{X} = \vec{x}$ is an actual cause of ϕ , then each “conjunct in $\vec{X} = \vec{x}$ is called *part* of a cause of ϕ ” (his emphasis). In the disjunctive scenario, however, each of $ML_1 = 1$ and $ML_2 = 1$ is individually judged to be an actual cause of $FB = 1$, not merely part of one cause. This is why overdetermination is considered a major problem for Lewis’s (1973) counterfactual account of causation in the first place.

5.5 A Note on Causal Models

We adopted the framework of causal models and so inherited their merits and drawbacks. The problem of spurious causation does not arise due to the recursivity of the structural equations. Moreover, the recursivity lifts the need to obtain the directedness of causation, e.g. by assuming a temporal order.

On the other hand, models of structural equations suffer from a drawback: they encode *some* antecedently given causal structure. As observed by Hitchcock (2007, pp. 503-4), each causal model represents a “*causal structure* of the situation in question”. In a causal model, structural equations relate values of variables that are meant to refer to token events, for example, which impact Billy’s throwing a rock at a certain place and time would have on the glass bottle under certain (non-actual) circumstances. Thereby, structural equations presuppose some information about causal relations, be it in form of sets of counterfactuals (e.g. Hitchcock (2001) and Woodward (2003)), or primitive causal mechanisms construed as law-like relationships that support a counterfactual interpretation (e.g. Pearl (2009) and Halpern and Pearl (2005)).²⁰

Finally, our Ramsey Test analysis of actual causation is characterized by an epistemic interpretation of causal models. That is, we take the structural equations to express beliefs about elementary causal dependences. While for Halpern and Hitchcock (2015, p. 384) the structural equations may be seen as “describing objective features of the world”, they agree that “judgments of actual causation are subjective”. Note, furthermore, that their ‘reasoning about causality’ is suspiciously close to our ‘judgments on causal relations’. “[U]ltimately”, so Halpern and Pearl (2005, p. 878), “the choice of model is a subjective one”. This is in line with the present Ramsey Test definition being relative to an epistemic state. Hence, nothing seems to preclude an interpretation of the model relativity as agent relativity. Nothing essential, however, hinges on our epistemic interpretation of causal models. Although suggested by the Ramsey Test, an epistemic interpretation is not mandatory for our analysis of causation. A more realist reading can be obtained via the notion of an epistemically perfect agent whose beliefs never fail to be true.

²⁰This being said, the proponents of models of structural equations do not find the presupposition of some causal structure all too problematic. For causal models do not directly represent relations of actual causation. You cannot simply read off actual causal relations from a causal model. Rather, the hope is that appropriate causal models encode information about a causal structure in terms of which actual causes can be identified. However, the causal structure is a primitive that cannot be further analyzed in terms of actual causation. For a more detailed justification, see Halpern and Pearl (2005), Halpern and Hitchcock (2010), and Fenton-Glynn (2017).

6 Conclusion

We have proposed a Ramsey Test definition of actual causation in the framework of causal models. For this definition cases of overdetermination, conjunctive scenarios, and preemption are no problem. With respect to these standard examples, we have shown that the definition surpasses the Halpern-Pearl definitions. Underlying this claim are criteria that Halpern and Pearl (2005, p. 846) suggest for evaluating approaches of actual causation:

The best ways to judge the adequacy of an approach are the intuitive appeal of the definitions and how well it deals with examples; we believe that this article shows that our approach fares well on both counts.

We believe that the present article has shown that our analysis fares even better than the Halpern-Pearl definitions with respect to overdetermination and conjunctive scenarios.

Of course, we need to deal with the other problematic examples of the literature in order to properly compete with the Halpern-Pearl definitions. Fortunately, such a follow-up paper is underway, including a treatment of prevention, double prevention, and switches. There, we shall also compare the agnostic dependence used in our strengthened Ramsey Test with the notion of counterfactual dependence.

References

- Alchourrón, M. A., Gärdenfors, P., and Makinson, D. On the Logic of Theory Change: Partial Meet Contraction Functions and their Associated Revision Functions. *Journal of Symbolic Logic*, 50:510–530, 1985.
- Andreas, H. and Günther, M. On the Ramsey Test Analysis of ‘Because’. *Erkenntnis*, Jun 2018.
- Eiter, T. and Lukasiewicz, T. Complexity Results for Structure-Based Causality. *Artificial Intelligence*, 142(1):53 – 89, 2002.
- Fenton-Glynn, L. A Proposed Probabilistic Extension of the Halpern and Pearl Definition of ‘Actual Cause’. *British Journal for the Philosophy of Science*, 68: 1061–1124, 2017.

- Gärdenfors, P. *Knowledge in Flux*. MIT Press, Cambridge, MA, 1988.
- Hall, N. Structural Equations and Causation. *Philosophical Studies*, 132(1):109–136, 2007.
- Halpern, J. Y. Defaults and Normality in Causal Structures. *CoRR*, abs/0806.2140, 2008.
- Halpern, J. Y. A Modification of the Halpern-Pearl Definition of Causality. *Proc. 24th International Joint Conference on Artificial Intelligence (IJCAI 2015)*, pages 3022–3033, 2015.
- Halpern, J. Y. and Hitchcock, C. Actual Causation and the Art of Modeling. In Dechter, R., Geffner, H., and Halpern, J. Y., editors, *Heuristics, Probability, and Causality: a Tribute to Judea Pearl*, pages 383–406. College Publications, London, 2010.
- Halpern, J. Y. and Hitchcock, C. Graded Causation and Defaults. *British Journal for the Philosophy of Science*, 66(2):413–457, 2015.
- Halpern, J. Y. and Pearl, J. Causes and Explanations: A Structural-Model Approach. Part 1: Causes. In Breese, J. and Koller, D., editors, *Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence*, pages 194–202, San Francisco, 2001. Morgan Kaufman.
- Halpern, J. Y. and Pearl, J. Causes and Explanations: A Structural-Model Approach. Part I: Causes. *British Journal for the Philosophy of Science*, 56(4):843–887, 2005.
- Hiddleston, E. Causal Powers. *The British Journal for the Philosophy of Science*, 56(1):27–59, 2005.
- Hitchcock, C. The Intransitivity of Causation Revealed in Equations and Graphs. *Journal of Philosophy*, 98(6):273–299, 2001.
- Hitchcock, C. Prevention, Preemption, and the Principle of Sufficient Reason. *Philosophical Review*, 116(4):495–532, 2007.
- Lewis, D. Causation. *Journal of Philosophy*, 70(17):556–567, 1973.
- Pearl, J. *Causality: Models, Reasoning and Inference*. Cambridge University Press, New York, NY, USA, 2nd edition, 2009.

- Ramsey, F. P. General Propositions and Causality. *Philosophical Papers*, pages 145–163, 1929/1990.
- Rott, H. Ifs, Though, and Because. *Erkenntnis*, 25(3):345–370, 1986.
- Stalnaker, R. A Theory of Conditionals. In Rescher, N., editor, *Studies in Logical Theory (American Philosophical Quarterly Monograph Series)*, number 2, pages 98–112. Blackwell, Oxford, 1968.
- Woodward, J. *Making Things Happen : A Theory of Causal Explanation*. Oxford University Press, Oxford, 2003.