

Equivalent Causal Models

Sander Beckers

Munich Center for Mathematical Philosophy – Ludwig Maximilian University, Munich
srekcebrednas@gmail.com

Abstract

The aim of this paper is to offer the first systematic exploration and definition of *equivalent* causal models in the context where both models are not made up of the same variables. The idea is that two models are equivalent when they agree on all “essential” causal information that can be expressed using their common variables. I do so by focussing on the two main features of causal models, namely their structural relations and their functional relations. In particular, I define several relations of causal ancestry and several relations of causal sufficiency, and require that the most general of these relations are preserved across equivalent models.

1 Introduction

Causal models are becoming evermore important and their domains of application are growing rapidly, both within science (most notably AI) and in philosophy. As with all forms of modeling, there usually exist many causal models that can be used to describe a given system of interest. Since different models can lead to widely divergent conclusions, it is of paramount importance to have some method to compare models and group together those that are sufficiently similar. To do so we need to be able to establish when two causal models are *equivalent*. The aim of this paper is to offer the first systematic exploration and definition of equivalent causal models.

Of course what makes for a sensible notion of equivalence varies across different contexts. For example, there exists a well-known notion of equivalence that is used in the epistemically limited context where one is learning a causal model from data. However, such Markov-equivalence (as it is called) does not apply in the context where two causal models are fully specified and do not consist of the same variables. Such contexts are widespread in science, but are of particular concern when trying to capture *stable*, model-independent, features of the world. Furthermore, exploring equivalence in such an omniscient context offers a firm conceptual starting point for generalizing to contexts that involve approximation and uncertainty in future work.

Before diving into the details, I spell out the underlying intuition of both the notion to be developed and the method

by which I shall do so. The aim is to formalize a suitable definition of causal model-equivalence that is relative to a set of variables that the models have in common: the idea is that two models are equivalent when they agree on all “essential” causal information that can be expressed using these common variables. Obviously this will require motivating what I take to be (in)essential and why I choose to focus merely on common variables.

My method focusses on the observation that causal models are characterized using two types of relations: structural relations, i.e., having to do with the causal graph, and functional relations, i.e., having to do with how particular values of variables are determined by values of other variables. I develop two notions of equivalence, one for each type. Causal equivalence will then be defined as being both structurally and functionally equivalent.

The details are as follows. The most basic structural relation is that of one variable X being an ancestor of another variable Y . I generalize this relation along three dimensions. First, I take into account the specific values x, x' and y, y' of the variables that are responsible for the ancestry. Second, I consider ancestry for *sets* of variables \vec{X} and \vec{Y} . Third, I distinguish between potential ancestors, which are relative to a causal model, and actual ancestors, which are relative to a causal model and a specific context. Structural equivalence is then defined as having identical potential and actual ancestors.

The most basic functional relation is that of one set of variable values \vec{x} being causally sufficient for another set \vec{y} . I first look into spelling out an appropriate notion of causal sufficiency, resulting in direct sufficiency. I then generalize this relation by defining sufficiency as its transitive closure. Functional equivalence is defined as having identical sufficiency relations.

The structure of this paper is as follows: the next section introduces the structural equations framework. I then offer a motivation and clarification of the problem that this paper is concerned with in Section 3. Section 4 focusses on the causal structure and develops a notion of structural equivalence. In Section 5 I shift focus to functional equivalence, by considering what makes for a good definition of causal sufficiency. Causal equivalence is then defined as the combination of both forms of equivalence. Section 6 compares functional equivalence to Halpern’s notion of conservative

extension. Section 7 discusses other related work and reflects on possible future directions.

2 Structural Equation Models

This section reviews the definition of causal models as understood in the structural modeling tradition started by (Pearl 2000), mostly using the notation and terminology from (Halpern 2016a).

Definition 1 A signature \mathcal{S} is a tuple $(\mathcal{U}, \mathcal{V}, \mathcal{R})$, where \mathcal{U} is a set of exogenous variables, \mathcal{V} is a set of endogenous variables, and \mathcal{R} a function that associates with every variable $Y \in \mathcal{U} \cup \mathcal{V}$ a nonempty set $\mathcal{R}(Y)$ of possible values for Y (i.e., the set of values over which Y ranges). If $\vec{X} = (X_1, \dots, X_n)$, $\mathcal{R}(\vec{X})$ denotes the crossproduct $\mathcal{R}(X_1) \times \dots \times \mathcal{R}(X_n)$. ■

Exogenous variables represent factors whose causal origins are outside the scope of the causal model, such as background conditions and noise. The values of the endogenous variables, on the other hand, are causally determined by other variables within the model (both endogenous and exogenous).

Definition 2 A causal model M is a pair $(\mathcal{S}, \mathcal{F})$, where \mathcal{S} is a signature and \mathcal{F} defines a function that associates with each endogenous variable X a structural equation F_X giving the value of X in terms of the values of other endogenous and exogenous variables. Formally, the equation F_X maps $\mathcal{R}(\mathcal{U} \cup \mathcal{V} - \{X\})$ to $\mathcal{R}(X)$, so F_X determines the value of X , given the values of all the other variables in $\mathcal{U} \cup \mathcal{V}$. ■

We call a setting $\vec{u} \in \mathcal{R}(\mathcal{U})$ of values of exogenous variables a *context*. The value of Y may depend on the values of only a few other variables. Y *depends on* X if there is some context \vec{u} and some setting of the endogenous variables other than Y and X such that varying the value of X in that context results in a variation in the value of Y ; that is, there is a setting \vec{z} of the endogenous variables other than Y and X and values x and x' of X such that $F_Y(x, \vec{z}, \vec{u}) \neq F_Y(x', \vec{z}, \vec{u})$. (Or, in case $X \in \mathcal{U}$, $F_Y(x, \vec{z}, \vec{u} - x) \neq F_Y(x', \vec{z}, \vec{u} - x)$.) We then say that X is a *parent* of Y . It will prove useful later on to have some notation to capture the details of the foregoing situation: we write $(x, x') \rightsquigarrow (y, y')$ and say that (\vec{u}, \vec{z}) is a *witness* to this. (In case $X \in \mathcal{U}$, the witness is of the form $(\vec{u} - x, \vec{z})$. For sake of notational simplicity we leave this nuance implicit.)

In this paper we restrict attention to *strongly recursive* (or *acyclic*) models, that is, models where there is a partial order \leq on variables such that if Y depends on X , then $X < Y$.¹ In a strongly recursive model, given a context \vec{u} , the values of all the endogenous variables are determined (we can just solve for the value of the variables in the order given by $<$). We often write the equation for an endogenous variable as $Y = f(\vec{X})$; this denotes that the value of Y depends only on the values of the variables in \vec{X} , and the connection is given by the function f . For example, we might have $Y = X + 5$.

¹All definitions could easily be generalized to *recursive* models, i.e., models where the partial order may depend on the context.

An *intervention* has the form $\vec{X} \leftarrow \vec{x}$, where \vec{X} is a set of endogenous variables. Intuitively, this means that the values of the variables in \vec{X} are set to the values \vec{x} . The structural equations define what happens in the presence of interventions. Setting the value of some variables \vec{X} to \vec{x} in a causal model $M = (\mathcal{S}, \mathcal{F})$ results in a new causal model, denoted $M_{\vec{X} \leftarrow \vec{x}}$, which is identical to M , except that \mathcal{F} is replaced by $\mathcal{F}^{\vec{X} \leftarrow \vec{x}}$: for each variable $Y \notin \vec{X}$, $F_Y^{\vec{X} \leftarrow \vec{x}} = F_Y$ (i.e., the equation for Y is unchanged), while for each X' in \vec{X} , the equation $F_{X'}$ for X' is replaced by $X' = x'$ (where x' is the value in \vec{x} corresponding to X').

Given a signature $\mathcal{S} = (\mathcal{U}, \mathcal{V}, \mathcal{R})$, an *atomic formula* is a formula of the form $X = x$, for $X \in \mathcal{V}$ and $x \in \mathcal{R}(X)$. A *causal formula* (over \mathcal{S}) is one of the form $[Y_1 \leftarrow y_1, \dots, Y_k \leftarrow y_k] \phi$, where

- ϕ is a Boolean combination of atomic formulas,
- Y_1, \dots, Y_k are distinct variables in \mathcal{V} , and
- $y_i \in \mathcal{R}(Y_i)$ for each $1 \leq i \leq k$.

Such a formula is abbreviated as $[\vec{Y} \leftarrow \vec{y}] \phi$. The special case where $k = 0$ is abbreviated as ϕ . Intuitively, $[Y_1 \leftarrow y_1, \dots, Y_k \leftarrow y_k] \phi$ says that ϕ would hold if Y_i were set to y_i , for $i = 1, \dots, k$.

A causal formula ψ is true or false in a *causal setting*, which is a causal model given a context. As usual, we write $(M, \vec{u}) \models \psi$ if the causal formula ψ is true in the causal setting (M, \vec{u}) . The \models relation is defined inductively. $(M, \vec{u}) \models X = x$ if the variable X has value x in the unique (since we are dealing with recursive models) solution to the equations in M in context \vec{u} (i.e., the unique vector of values that simultaneously satisfies all equations in M with the variables in \mathcal{U} set to \vec{u}). The truth of conjunctions and negations is defined in the standard way. Finally, $(M, \vec{u}) \models [\vec{Y} \leftarrow \vec{y}] \phi$ if $(M_{\vec{Y} \leftarrow \vec{y}}, \vec{u}) \models \phi$.

3 Motivation

There are many different contexts in which one might consider one causal model M_1 to be equivalent to another model M_2 . A crucial dimension along which we can separate these contexts is by looking at the kind of relation they require to hold between the sets of variables of both models.

A first context, that this paper is not directly concerned with, is the one in which both models are required to have identical variables. It is within this context that most causal search algorithms take place. The output of such an algorithm is usually a partially directed graph that represents the Markov Equivalence Class. Markov equivalence means that two models have identical observational probability distributions. Yet two Markov equivalent models only partially agree about the ancestral relations, and need not agree at all about the structural equations.

This paper is concerned with a second context, namely the one in which we are comparing two causal models only with regard to (a subset of) their common variables \mathcal{V} . In this context, the primary interest of the modeler lies in the causal relations between these common variables. We are looking for a notion of equivalence that is strong enough such that

when confronted with a choice between two models M_1 and M_2 that contain (at least) all of these variables, we can state that if M_1 and M_2 are equivalent with regards to \mathcal{V} then they agree on all the essential causal relations that hold between the common variables and thus one is free to choose either model.

This notion of equivalence matters because we want to be able to draw *stable* conclusions about the causal relations between the variables of interest: no matter what further variables we discover to be causally connected to the variables of interest, and no matter what variables that we are not interested in are being marginalized away, as long as the resulting causal model is equivalent to the original one the conclusions about their common variables remain identical.

As a consequence, the goal of this paper is to find a much stronger notion of equivalence than any notion that could be useful in the first context, one that gets much closer to identity (which is of course the strongest possible notion of equivalence). In particular, any difference between two equivalent causal models in this strong sense should be entirely explained by the fact that the two models do not have identical variable sets. A fortiori, when applying this notion of equivalence to models with identical variable sets, it should reduce to identity. For example, say we have M_1 consisting of the equation $E = C$ and M_2 consisting of equations $E = D$ and $D = C$. If we are exclusively interested in the variables E and C , then any sensible notion of equivalence that allows for different variable sets should consider both models to be equivalent with regards to E and C . Of course this does not mean that it is impossible to state a causal relation between C and E over which both models disagree: one could simply count the number of edges it takes to get from C to E , and in that respect the two models clearly are different. Yet I assume that such “numerical” properties are not essential. If such properties were essential, then it seems impossible to come up with any useful notion of equivalence, for presumably one can always zoom in and inject a hitherto ignored intermediate variable along any edge. Likewise, the number of distinct paths going from one variable to another is also deemed to be inessential for the purposes of this paper. I make do with these intuitive considerations regarding numerical properties and refrain from giving a precise definition.

Lastly a brief word about a third context that one could consider for the relation between the sets of variables of two models. A natural suggestion is to also compare two models such that we have a mapping $\tau : \mathcal{R}_2(\mathcal{V}_2) \rightarrow \mathcal{R}_1(\mathcal{V}_1)$ between the values of the variables of both models. Indeed, that was the context investigated in previous work (Beckers and Halpern 2019; Beckers, Eberhardt, and Halpern 2019). Yet there the particular notion of equivalence that we used was left entirely implicit, and thus other – more interesting – notions were never considered. So in that respect the current paper is more general. (This will be discussed in Section 6.)

Note also that the third and second context overlap: if M_2 is an *extension* of M_1 , meaning that $\mathcal{V}_1 \subseteq \mathcal{V}_2$, then one can consider the special case where τ is the identity-mapping for the common variables and marginalization for the oth-

ers, i.e., $\tau(\bar{v}_2) = \bar{v}_1$ for all $\bar{v}_2 \in \mathcal{R}_2(\mathcal{V}_2)$ where \bar{v}_1 is the restriction of \bar{v}_2 to $\mathcal{R}_1(\mathcal{V}_1)$. Therefore an interesting line of future work is to generalize the notion of equivalence here developed to the third context by considering equivalence for any mapping τ .

4 Structural Equivalence

This section focuses on the *causal structure* of the causal model. Several useful concepts are introduced in this regard, which are then combined to suggest a definition of structural equivalence.

Section 2 introduced the concept of one variable being a parent of another. It is customary to further generalize this genealogical terminology by taking *ancestor* to be the transitive closure of the parent relation, i.e., X is an ancestor of Y if there are variables V_1, \dots, V_n , such that X is a parent of V_1 , V_1 is a parent of V_2, \dots , and V_n is a parent of Y . By treating variables as nodes and drawing directed edges from parents to their children, each strongly recursive causal model induces a unique DAG (Directed Acyclic Graph) that captures the causal structure of the model. We say that X is an *ancestor* of Y if X is an ancestor of Y in this DAG.

This graph by itself offers only a rough characterization of all the causal information that the model contains: it tells you which variable can depend on some other variable, but it does not provide details of this dependence. In particular, it does not say which values of some variable depend on what values of some other variable. Such details can be added by refining the standard genealogical relations so that they also take into account the *specific values* which instantiate the parent/ancestral relations.

(Throughout the rest of the paper we assume that $M_1 = ((\mathcal{U}, \mathcal{V}_1, \mathcal{R}_1), \mathcal{F})$, $M_2 = ((\mathcal{U}, \mathcal{V}_2, \mathcal{R}_2), \mathcal{F}_2)$, and $\mathcal{V} \subseteq (\mathcal{V}_1 \cap \mathcal{V}_2)$. We further assume that X and Y are distinct members of \mathcal{V} unless stated otherwise, x, x' are distinct values in $\mathcal{R}(X)$, and likewise for y, y' . Similarly, \bar{X} and \bar{Y} are distinct – possibly overlapping – subsets of \mathcal{V} , and \bar{x}, \bar{x}' are non-overlapping sets of values in $\mathcal{R}(X)$. For brevity, we often leave the variables implicit, writing \bar{x} instead of $\bar{X} = \bar{x}$, for example. All definitions assume that $\bar{X} \subseteq \mathcal{V}$, but this is merely for notational simplicity. The case where $\bar{X} \cap \mathcal{U} \neq \emptyset$ is left implicit.)

Definition 3 We say that x rather than x' is a potential parent of y rather than y' if there exists a witness (\bar{u}, \bar{z}) for $(x, x') \rightsquigarrow (y, y')$.²

The ancestral generalization is straightforwardly defined as the transitive closure. We use the following notation for potential ancestry: $(x, x') \rightsquigarrow_p (y, y')$, where $p = (X, V_1, \dots, V_n, Y)$ is called a directed path from X to Y .

Should potential ancestral relations be preserved among equivalent causal models? As such this relation contains crucial causal information, namely that there’s a *stepwise connection* between changing some value of X and changing a

²Note that potential parenthood is a symmetric relation for the values of the variables: $(x, x') \rightsquigarrow (y, y')$ iff $(x', x) \rightsquigarrow (y', y)$. This is no longer true for actual parenthood, introduced later on.

value of Y . But the potential ancestral relation also contains more specific information that is numeric in nature: it states that this connection is given by a single path. As mentioned before, I consider it a mistake to focus on numerical properties. Indeed, consider the following example.

Example 1 Assume the range of E is $\{0, 1, 2\}$ and all other variables are binary. (Throughout all examples, all variables are binary unless otherwise stated.) The original model M_1 consists simply of the equation $E = 2 * C$. The extended model M_2 consists of $A = C$, $B = C$, and $E = A + B$. Intuitively, I believe that these models should come out as equivalent regarding C and E . Yet although C is an ancestor of E in both models, $(1_C, 0_C) \rightsquigarrow_p (2_E, 0_E)$ for some path p only in M_1 . (The subscripts indicate to which variables the values correspond.)

Still, the essential property that there exists a stepwise connection between $(1_C, 0_C)$ and $(2_E, 0_E)$ is preserved in M_2 , it is just that this connection now consists of simultaneously considering two paths rather than one: $(1_C, 0_C) \rightsquigarrow (1_A, 1_B; 0_A, 0_B) \rightsquigarrow (2_E, 0_E)$. In other words, we can generalize the potential ancestral relation to allow for steps between sets of variable values rather than steps between single variable values. Graphical relations are not naturally equipped to deal with sets of values, but functional relations are. So in order to make this precise we first express the single-path potential ancestral relations using functional relations, i.e., relations expressed in terms of which formulas are true under certain interventions.

Fact 1 $(x, x') \rightsquigarrow (y, y')$ with witness (\bar{u}, \bar{z}) in M iff $(M, \bar{u}) \models [\bar{Z} \leftarrow \bar{z}, X \leftarrow x]Y = y \wedge [\bar{Z} \leftarrow \bar{z}, X \leftarrow x']Y = y'$.

As was pointed out in Section 2, in strongly recursive models the parents of a variable suffice to determine its value. Therefore we can always restrict the witness to only include the context and a variable's endogenous parents. In fact, sometimes we can restrict it even further. For example, if the equation for Y is $Y = (A \wedge B) \vee C$, then $(0_A, 0_C)$ suffices to determine that $Y = 0$ and $(0_A, 1_C)$ suffices to determine that $Y = 1$. This allows us to conclude that $(0_C, 1_C) \rightsquigarrow (0_Y, 1_Y)$ with just using (0_A) as a witness.

Generalizing the notion of a witness so that it need not consist of values for all other variables allows for comparing potential parenthood across models with different signatures. We make this precise by defining what it means for a set of variable values to be *directly sufficient* for another set of variable values. Informally, a set is sufficient precisely when the outcome is independent of the values of variables not in the set.

Definition 4 We say that \bar{x} is directly sufficient for \bar{y} if for all settings \bar{z} of the other endogenous variables and all contexts \bar{u} we have that $(M, \bar{u}) \models [\bar{X} \leftarrow \bar{x}, \bar{Z} \leftarrow \bar{z}]\bar{Y} = \bar{y}$.

(If $\bar{X}_{\mathcal{U}} = \bar{X} \cap \mathcal{U}$ is not empty, the quantification is restricted to contexts \bar{u} such that their restriction to $\bar{X}_{\mathcal{U}}$ agrees with $\bar{x}_{\mathcal{U}}$.) ■

Now that we have identified the crucial functional relation out of which the potential parent relation is built up, we can generalize this parent relation to sets of variable values.

Definition 5 We say that \bar{x} rather than \bar{x}' are potential joint parents of \bar{y} rather than \bar{y}' if there is a set $\bar{Z} \subseteq \mathcal{V} \cup \mathcal{U}$, and a setting $\bar{z} \in \mathcal{R}(\bar{Z})$ such that

- (\bar{z}, \bar{x}) is directly sufficient for \bar{y} ,
- (\bar{z}, \bar{x}') is directly sufficient for \bar{y}' , and
- \bar{X} is minimal (i.e., at least one of the previous statements is false when considering any strict subset of \bar{X}).

We notate this as $(\bar{x}, \bar{x}') \rightsquigarrow (\bar{y}, \bar{y}')$ and say that \bar{z} is a witness. ■

It is trivial to see that if restricted to singletons, Definition 5 reduces to Definition 3. We demand \bar{X} to be minimal in order to avoid calling redundant elements parents. (For example, if we have as structural equation $Y = A$, then we do not want to call $(1_A, 1_B)$ joint parents of $Y = 1$: $B = 1$ is entirely redundant here.) We demand that \bar{X} and \bar{Y} are non-identical in order to avoid self-parenthood. But it would be too much to demand that their intersection is empty. To see why, we consider a slight variant on Example 1.

Example 2 All variables are as before. The original model M_1 consists simply of the equation $E = 2 * C$. The extended model M_2 consists of $A = C$ and $E = A + C$. As before, I believe that these models should come out as equivalent regarding C and E . But in order to establish that $(1_C, 0_C) \rightsquigarrow_p (2_E, 0_E)$ for some path p in M_2 , we need to break up the single step into $(1_C, 0_C) \rightsquigarrow (1_C, 1_A; 0_C, 0_A) \rightsquigarrow (2_E, 0_E)$.

Now that we have generalized the potential parenthood relation along the dimension of sets of variables (which is already built into functional relations), we can take the transitive closure to further generalize this relation along the dimension that is characteristic of the ancestral relation (but is lacking from many functional relations).

Definition 6 We say that \bar{x} rather than \bar{x}' are potential joint ancestors of \bar{y} rather than \bar{y}' if there exist $\bar{V}_1, \dots, \bar{V}_n$ and settings $\bar{v}_i, \bar{v}'_i \in \mathcal{R}(\bar{V}_i)$ for each $i \in \{1, \dots, n-1\}$ such that $(\bar{v}_i, \bar{v}'_i) \rightsquigarrow (\bar{v}_{i+1}, \bar{v}'_{i+1})$, $(\bar{x}, \bar{x}') \rightsquigarrow (\bar{v}_1, \bar{v}'_1)$ and $(\bar{v}_n, \bar{v}'_n) \rightsquigarrow (\bar{y}, \bar{y}')$. We call the tuple $n = (\bar{X}, \bar{V}_1, \dots, \bar{V}_n, \bar{Y})$ a network from \bar{X} to \bar{Y} , and say that \bar{x} rather than \bar{x}' are potential joint ancestors of \bar{y} rather than \bar{y}' along n . We notate this as $(\bar{x}, \bar{x}') \rightsquigarrow_n (\bar{y}, \bar{y}')$. ■

Preservation of potential joint ancestors deals with Examples 1 and 2, but it does not suffice for our notion of equivalence, because it entirely ignores that much causal information is context-dependent. The above defined ancestral relations are all *general properties of the causal model M* , in the sense that they do not make any reference to an actual context \bar{u} . Yet causal models also capture causal information regarding specific contexts. Indeed, it is such information that is at stake when speaking of actual causation, an important concept that has benefitted greatly from the advent of causal models. A good notion of structural equivalence should take this into account as well. Therefore we also define a context-specific counterpart of the potential ancestral relations.

Definition 7 We say that x rather than x' is an actual parent of y rather than y' w.r.t. (M, \bar{u}) if $(x, x') \rightsquigarrow (y, y')$ with a

witness (\vec{u}, \vec{z}) such that $(M, \vec{u}) \models \vec{Z} = \vec{z} \wedge X = x$. We notate this as $(x, x') \rightsquigarrow^{\vec{u}} (y, y')$. ■

Both the generalization of actual parenthood to sets of variable values and the analogous ancestral counterpart are defined straightforwardly. I only explicitly give the combination of both generalizations, which is the context-relative counterpart of Definition 6.

Definition 8 We say that \vec{x} rather than \vec{x}' are actual joint ancestors of \vec{y} rather than \vec{y}' w.r.t. (M, \vec{u}) if $(\vec{x}, \vec{x}') \rightsquigarrow_n (\vec{y}, \vec{y}')$ for some n such that the exogenous parts of all the witnesses are contained in \vec{u} and all the endogenous parts \vec{z} are such that $(M, \vec{u}) \models \vec{Z} = \vec{z} \wedge \vec{X} = \vec{x}$. We notate this as $(\vec{x}, \vec{x}') \rightsquigarrow_n^{\vec{u}} (\vec{y}, \vec{y}')$. ■

The following example illustrates why preservation of actual joint ancestors should be combined with the preservation of potential joint ancestors to get structural equivalence.

Example 3 Model M_1 contains $E = B \vee D$, whereas model M_2 contains $E = (A \wedge \neg C) \vee B \vee D$. Both models contain $A = C$, $B = C$, and $D = \neg C$. Note that both models have identical signatures. It is easy to verify that for each context both models will have the same actual joint ancestors. Yet they disagree on the potential ancestors, for $(1_A, 0_A) \rightsquigarrow (1_E, 0_E)$ in M_2 (with witness $(0_C, 0_B, 0_D)$) and not so in M_1 . The two models are indeed different, for they disagree on whether A is an ancestor of E , and thus do not even have the same DAG.

We are finally ready to define what it means for the structural causal information to be preserved across causal models. (As a reminder, an equivalence relation is defined as being reflexive, symmetric, and transitive. It is easy to see that Definition 9 satisfies those properties.)

Definition 9 M_1 and M_2 are structurally equivalent relative to \mathcal{V} if for all $\vec{X} \subseteq \mathcal{U} \cup \mathcal{V}$, $\vec{Y} \subseteq \mathcal{V}$, and \vec{x}, \vec{y} the following two conditions hold:

$(\vec{x}, \vec{x}') \rightsquigarrow_n (\vec{y}, \vec{y}')$ for some n in M_1 iff $(\vec{x}, \vec{x}') \rightsquigarrow_m (\vec{y}, \vec{y}')$ for some m in M_2 .

For all $\vec{u} \in \mathcal{R}(\mathcal{U})$ it holds that

$(\vec{x}, \vec{x}') \rightsquigarrow_n^{\vec{u}} (\vec{y}, \vec{y}')$ for some n in M_1 iff $(\vec{x}, \vec{x}') \rightsquigarrow_m^{\vec{u}} (\vec{y}, \vec{y}')$ for some m in M_2 . ■

As mentioned, an interesting special case occurs when we are comparing two models with regards to *all* of their common variables, i.e., $\mathcal{V} = \mathcal{V}_1 \cap \mathcal{V}_2$. In that case we simply say that M_1 and M_2 are *structurally equivalent*. (Note that contrary to Definition 9, this is not a transitive relation, so here we are using the term “equivalence” loosely.)

The following variation on Example 3 illustrates that structural equivalence does not suffice to ensure that equivalence of models with identical signatures reduces to identity.

Example 4 In M_1 we have $E = (A \wedge \neg C \wedge F) \vee B \vee D \vee (G \wedge F)$, whereas in M_2 we have $E = (A \wedge \neg C \wedge G) \vee B \vee D \vee (G \wedge F)$. Both models contain $A = C$, $B = C$, $F = C$, $G = C$, and $D = \neg C$. The reader can confirm that in addition to agreeing on all actual joint ancestral relations, as was the case with Example 3, they also agree on all potential joint ancestors, and are thus structurally equivalent.

(Note that they sometimes disagree about the witnesses for the potential relations: $(1_F, 0_F) \rightsquigarrow (1_E, 0_E)$ with witness $(1_A, 0_C, 0_B, 0_D, 0_G)$ only in M_1 , and likewise with F and G reversed for M_2 .) Intuitively the models are not equivalent, because they disagree on whether it is F or G that can contribute to $E = 1$ together with A and C .

The next section develops several notions of equivalence which do give the desired outcome for cases with identical signature, and argues for the superiority of one of them.

5 Functional Equivalence

Now we switch attention to the basic functional relations that are captured by causal models, meaning those relations capturing how the values of variables are determined by the values of other variables. Recall that if \vec{x} is directly sufficient for \vec{y} (Def. 4), this means precisely that \vec{x} fully determines the variables \vec{Y} taking on the values \vec{y} . So the functional relations between a set of variables are specified completely by stating all relations of direct sufficiency that hold between their values. A naive suggestion might therefore be to define functional equivalence as having identical relations of direct sufficiency for all sets only containing variables that are common to both models. However, that would give an overly strong notion of equivalence, for it means that the only way of extending a model is by adding children to childless variables. The reason is that the moment one adds a new parent A to a set of existing parents \vec{X} of a variable Y then, by definition of being a parent, at least one setting \vec{x} will cease being directly sufficient for some y . To illustrate, even the trivial models we considered as a sanity check in Section 3 would not be considered equivalent. Recall that M_1 is made up of $E = D$ and $D = C$, and M_2 consists simply of $E = C$. Given that $C = c$ is directly sufficient for $E = e$ only in the latter, these models would not come out as equivalent for the variables $\{E, C\}$.

The problem is that, as the name suggests, *direct* sufficiency focusses on parent-children relations, and is thus a numerical property. Instead, we generalize the notion of direct sufficiency so that it also captures sufficiency for descendants that are not children. Recall that we generalized ancestral relations by incorporating a feature of functional relations, namely that it also applies to sets of variable values and not just to single variables. Here we make the symmetrical counterpart of that move: we generalize the functional relation of direct sufficiency by incorporating a feature of ancestral relations, namely its transitivity.

Definition 10 We say that \vec{x} is sufficient for \vec{y} if there exist $\vec{w}_0, \dots, \vec{w}_n$ such that for each i , \vec{w}_i is directly sufficient for \vec{w}_{i+1} , where $\vec{x} = \vec{w}_0$ and $\vec{y} = \vec{w}_n$. ■

This gives rise to our notion of functional equivalence. As with structural equivalence, we say that M_1 and M_2 are *functionally equivalent* if they are functionally equivalent relative to $\mathcal{V} = \mathcal{V}_1 \cap \mathcal{V}_2$.

Definition 11 M_1 and M_2 are functionally equivalent relative to \mathcal{V} if for all $\vec{X} \subseteq \mathcal{V} \cup \mathcal{U}$, $\vec{Y} \subseteq \mathcal{V}$ and \vec{x}, \vec{y} , it holds that \vec{x} is sufficient for \vec{y} in M_1 iff \vec{x} is sufficient for \vec{y} in M_2 . ■

It is easy to see that functional equivalence deals with Example 4. For example, $(1_A, 0_C, 1_F, 0_B, 0_D, 0_G)$ is sufficient for $E = 1$ in M_1 but not in M_2 . Nonetheless, the following example illustrates why functional equivalence by itself does not suffice to guarantee structural equivalence.

Example 5 M_1 consists simply of $Y = D$ and $X = D$. M_2 on the other hand consists of $Y = (V \wedge X) \vee D$, $V = D$, and $X = D$. It is easy to verify that the models are functionally equivalent. Yet both models disagree about whether X is an ancestor of Y , and thus are not structurally equivalent.

We therefore define causal equivalence as the combination of structural and functional equivalence.

Definition 12 M_1 and M_2 are causally equivalent relative to \mathcal{V} if they are both structurally and functionally equivalent relative to \mathcal{V} . M_1 and M_2 are causally equivalent if they are both structurally and functionally equivalent. ■

6 Weak Equivalence

We have come across two sensible notions of causal sufficiency already (Defs. 4 and 10). The following definition also offers a very natural notion of causal sufficiency.

Definition 13 We say that \bar{x} is weakly sufficient for \bar{y} w.r.t. (M, \bar{u}) if $(M, \bar{u}) \models [\bar{X} \leftarrow \bar{x}] \bar{Y} = \bar{y}$.

(If $\bar{X}_{\mathcal{U}} = \bar{X} \cap \mathcal{U}$ is not empty, then the restriction of the context \bar{u} to $\bar{X}_{\mathcal{U}}$ has to agree with $\bar{x}_{\mathcal{U}}$.) ■

Direct sufficiency required \bar{x} to guarantee \bar{y} regardless of the context and of the values of the other variables. Weak sufficiency on the other hand merely requires that \bar{x} helps in getting \bar{y} (even if only by not preventing it), but it does not guarantee it: it relies on the other variables to play nice and take on their usual values. Note that direct sufficiency implies sufficiency, and sufficiency implies weak sufficiency. Therefore one might wonder whether defining equivalence using weak sufficiency could offer a plausible alternative to functional equivalence.

(Halpern 2016b) considers precisely this definition, be it that he limits attention to the special case where M_2 is an extension of M_1 , i.e., a case where $\mathcal{V}_1 \subseteq \mathcal{V}_2$. Informally, his definition states that both models agree on all formulas consisting solely of variables that they have in common.

Definition 14 M_2 is a conservative extension of M_1 if for all $\bar{X}, \bar{Y} \subseteq \mathcal{V}_1$, all \bar{x}, \bar{y} , and all \bar{u} , it holds that \bar{x} is weakly sufficient for \bar{y} w.r.t. (M_1, \bar{u}) iff \bar{x} is weakly sufficient for \bar{y} w.r.t. (M_2, \bar{u}) . ■

Halpern calls the extended model a “conservative” extension precisely because he takes it to be intuitive that the larger model preserves all relevant causal information of the smaller model (as far as their mutual variables go of course). I define weak equivalence as the generalization of Definition 14 to all sets of common variables \mathcal{V} .

Definition 15 M_1 and M_2 are weakly equivalent relative to \mathcal{V} if for all $\bar{X} \subseteq \mathcal{U} \cup \mathcal{V}$, $\bar{Y} \subseteq \mathcal{V}$, all \bar{x}, \bar{y} , and all \bar{u} , it holds that \bar{x} is weakly sufficient for \bar{y} w.r.t. (M_1, \bar{u}) iff \bar{x} is weakly sufficient for \bar{y} w.r.t. (M_2, \bar{u}) . ■

As before, we say that M_1 and M_2 are *weakly equivalent* if they are weakly equivalent relative to all of their common variables. Just as functional equivalence, weak equivalence captures important features that are ignored by structural equivalence. In particular, it meets our demand that it reduces to identity when applied to models with identical variable sets (and thus also deals with Example 4). Yet functional equivalence is stronger than weak equivalence.

Theorem 1 Given M_1 and M_2 , it holds that if M_1 and M_2 are functionally equivalent relative to \mathcal{V} then they are weakly equivalent relative to \mathcal{V} .

Proof: Assume M_1 and M_2 are functionally equivalent relative to \mathcal{V} . Say \bar{x} is weakly sufficient for \bar{y} w.r.t. (M_1, \bar{u}) , where $\bar{X} \subseteq \mathcal{U} \cup \mathcal{V}$ and $\bar{Y} \subseteq \mathcal{V}$. (The reasoning for M_2 is identical.) We first show that (\bar{u}, \bar{x}) is sufficient for \bar{y} in M_1 .

Given strong recursivity, (\bar{u}, \bar{x}) is directly sufficient for some setting \bar{a}_1 , where $\bar{A}_1 \subseteq \mathcal{V}_1$ are those variables that only depend on the exogenous variables \mathcal{U} and on the variables in \bar{X} . If $\bar{Y} \subseteq \bar{A}_1$, then either $\bar{y} = \bar{a}_1$ or \bar{a}_1 is directly sufficient for \bar{y} , and in both cases the result follows. If not, we can continue this reasoning by observing that $(\bar{u}, \bar{x}, \bar{a}_1)$ is directly sufficient for some setting \bar{a}_2 , where $\bar{A}_2 \subseteq \mathcal{V}_1$ are those variables that only depend on the variables in $\mathcal{U} \cup \bar{X} \cup \bar{A}_1$. Therefore eventually we arrive at some \bar{a}_n so that $\bar{Y} \subseteq \bar{A}_n$ and (\bar{u}, \bar{x}) is sufficient for \bar{a}_n in M_1 .

Because M_1 and M_2 are functionally equivalent relative to \mathcal{V} , we also have that (\bar{u}, \bar{x}) is sufficient for \bar{y} in M_2 , from which it follows that \bar{x} is weakly sufficient for \bar{y} w.r.t. (M_2, \bar{u}) . ■

The following interesting result illustrates that the dichotomy between structural and functional relations is not strict: weak equivalence (and thus also functional equivalence) is not entirely independent from structural equivalence. Concretely, the ancestor relation is preserved when moving from a causal model to a conservative extension. (Example 5 illustrates that the reverse is not true.)

Theorem 2 Given models M_1 and M_2 so that M_2 is a conservative extension of M_1 , for any variables $X, Y \in \mathcal{V}_1$, we have that if X is an ancestor of Y in M_1 then X is an ancestor of Y in M_2 . In other words, the DAG of M_1 is a sub-DAG of M_2 when ignoring all the variables not in \mathcal{V}_1 .

Proof: Assume X is an ancestor of Y in M_1 . Because of the transitivity of the ancestral relation, we may assume that X is a parent of Y in M_1 , i.e., there exist x, x', y, y' and a witness (\bar{u}, \bar{z}) with $\bar{Z} = \mathcal{V}_1 - \{X, Y\}$ so that $(x, x') \rightsquigarrow (y, y')$. By Fact 1 this means that $(M_1, \bar{u}) \models [\bar{Z} \leftarrow \bar{z}, X \leftarrow x] Y = y \wedge [\bar{Z} \leftarrow \bar{z}, X \leftarrow x'] Y = y'$. Because M_2 is a conservative extension of M_1 , we also have $(M_2, \bar{u}) \models [\bar{Z} \leftarrow \bar{z}, X \leftarrow x] Y = y \wedge [\bar{Z} \leftarrow \bar{z}, X \leftarrow x'] Y = y'$.

Say \bar{A} are the parents of Y in M_2 excluding X . Thus any setting of (\bar{A}, X) combined with the context is directly sufficient for some value of Y . If all variables in \bar{A} get the same values under both interventions, i.e., there exists a setting $\bar{a} \in \mathcal{R}(\bar{A})$ such that $(M_2, \bar{u}) \models [\bar{Z} \leftarrow \bar{z}, X \leftarrow x] \bar{A} =$

$\vec{a} \wedge [\vec{Z} \leftarrow \vec{z}, X \leftarrow x'] \vec{A} = \vec{a}$, then (\vec{u}, \vec{a}) is a witness for the fact that X is also a parent of Y in M_2 .

If they do not get the same values then there exists an $A_1 \in \vec{A}$ which gets different values a_1 and a'_1 under the respective interventions. In other words, $(M_2, \vec{u}) \models [\vec{Z} \leftarrow \vec{z}, X \leftarrow x] A_1 = a_1 \wedge [\vec{Z} \leftarrow \vec{z}, X \leftarrow x'] A_1 = a'_1$.

We are thus in the same situation for A_1 , a parent of Y , as we were for Y . Therefore we can repeat the reasoning from above by looking at the parents of A_1 in M_2 excluding X . Given strong recursivity, eventually we have to find a variable A_i such that X is a parent of A_i , A_i is a parent of A_{i-1} , ..., and A_1 is a parent of Y . ■

Despite this result and the intuitive appeal of the idea that equivalence comes down to comparing formulas made up solely of common variables, the following example shows that weak equivalence is too weak to preserve all causal information that I would consider as relevant, and therefore is not suitable to replace functional equivalence.

Example 6 *The model M_2 contains the following equations: $E = B \vee C \vee F$, $B = A \wedge \neg C$, $C = D$, $A = \neg D$, and $F = \neg C$. If we remove variable B we get the model M_1 where the equation for E has become $E = A \vee C \vee F$. The reader can verify that the models are both weakly and structurally equivalent. Yet the role of A in both models is intuitively quite different: if $C = 1$, then $A = 1$ and $C = 1$ fulfill entirely symmetrical roles in bringing about $E = 1$ according to M_1 , which is clearly not the case according to M_2 . Functional equivalence formalizes this intuition: $A = 1$ is sufficient for $E = 1$ in M_1 and not in M_2 .*

7 Other Related Work and Future Directions

(Bongers et al. 2016) offer a definition of “equivalence after marginalization” that generalizes Halpern’s notion of conservative extension to cyclic and non-deterministic models. Their work is highly complementary to mine: contrary to them, I restrict myself to acyclic deterministic models, and contrary to me, they restrict themselves to using just this one type of equivalence (i.e., just as Halpern, they only consider Definition 14).

The naive suggestion mentioned earlier of defining functional equivalence using direct sufficiency is present implicitly in (Beckers and Halpern 2019). (In the terminology we used there, it comes down to stating that M_1 is a constructive abstraction of M_2 for the identity-marginalization mapping discussed in Section 3). (Rubenstein et al. 2017) also give such a definition (see their Theorem 9). However, as mentioned in Section 3, the focus of both articles lies with the third rather than the second context. As a result, neither of them discuss the issue of equivalence, nor do they consider any of the other forms of equivalence here introduced.

In future work I intend to generalize my approach to the third context by combining my notions of equivalence with the various notions of abstraction from (Beckers and Halpern 2019). Simply put, the idea would be to require for some causal relations that they hold between \vec{x} and \vec{y} in M_2 iff they hold between $\tau(\vec{x})$ and $\tau(\vec{y})$ in M_1 , where τ is an “abstraction” mapping. (What makes this non-trivial, is the

fact that τ is defined as mapping a setting $\vec{v}_2 \in \mathcal{R}_2(\mathcal{V}_2)$ to some $\vec{v}_1 \in \mathcal{R}_1(\mathcal{V}_1)$, and thus one has to come up with a sensible generalization to random settings \vec{x} and \vec{y} .) As an additional step, I can make use of the work from (Beckers, Eberhardt, and Halpern 2019) to define a notion of approximate equivalence.

Another promising extension concerns the topic of actual causation, i.e., the causal relation that holds between particular events, as opposed to types of events. I intend to use the current approach to analyse existing – and to propose novel – definitions of actual causation by assuming that a good definition ought to be stable across equivalent models: whether \vec{x} causes \vec{y} in some context \vec{u} should not depend on which of two equivalent models one uses. I take the idea of such an analysis from (Gallow forthcoming), whose work offered the impetus for the current project. He offers a very strong sufficient condition for equivalence, but does not develop a notion of equivalence. (The interested reader can easily verify that his sufficient condition indeed suffices to meet the criteria of my definition of causal equivalence.)

The notion of actual joint ancestry (Definition 8) will be central to this extension. It was inspired by a similar – but subtly different – one from (Gallow forthcoming), who presents it as the basis for a stable definition of actual causation.³⁴ Therefore the current paper can be seen as offering a partial formal vindication of his approach (modulo the subtle differences between the definitions). But I believe there is still room for improvement. For example, actual joint ancestry is (by definition) a transitive relation, whereas nowadays there is a broad consensus that actual causation is not transitive. For a more concrete and intuitive objection, consider a model with equation $E = (C \wedge B) \vee (\neg C \wedge A)$, and a context such that $C = 1$, $B = 1$, and $A = 1$. Then $(1_A, 1_C; 0_A, 0_C) \rightsquigarrow^{\vec{u}} (1_E, 0_E)$. However, it seems strange to consider $A = 1$ to be a part of an actual cause of $E = 1$, given that the disjunct it appears in was false.

To sum up, I investigated when two causal models ought to be considered equivalent relative to (a subset of) their common variables. I proposed an answer by looking at the two fundamental features of a causal model, namely its structural and its functional relations. Although the discussion has made clear that this definition satisfies certain intuitive properties, I do not take it to offer the last word. There may very well be further intricacies that the current definition overlooks, which could lead to more refined notions of equivalence.

³For the interested reader: here is an example for which his definition offers a different verdict. The equations are $E = (A \vee B) \wedge C$, and $B = A \vee C$, and we consider a context in which $A = 1$ and $C = 1$. Then $(1_A, 1_C; 0_A, 0_C) \rightsquigarrow^{\vec{u}} (1_A, 1_B; 0_A, 0_B) \rightsquigarrow^{\vec{u}} (1_E, 0_E)$, and thus $(A = 1, C = 1)$ are actual joint ancestors of $E = 1$. Gallow’s definition involves a minimality condition applied to the entire network, as opposed to applied to each step between parents and children. Since also $(1_C, 0_C) \rightsquigarrow^{\vec{u}} (1_E, 0_E)$, $(A = 1, C = 1)$ is not minimal and therefore he does not consider it an actual cause.

⁴His full definition adds to this a focus on the default/deviant distinction between values of variables, which is fairly common in work on actual causation. This distinction and his use of it could easily be added to my analysis, and thus stands orthogonal to it.

Acknowledgments

Many thanks to Dmitri Gallow, Joe Halpern, and AAAI reviewers for comments on an earlier version of this paper. This research was made possible by funding from the Alexander von Humboldt Foundation.

References

- Beckers, S.; Eberhardt, F.; and Halpern, J. Y. 2019. Approximate Causal Abstraction. In *Proc. 35th Conference on Uncertainty in Artificial Intelligence (UAI 2019)*.
- Beckers, S.; and Halpern, J. Y. 2019. Abstracting Causal Models. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence*, volume 33, 2678–2685.
- Bongers, S.; Peters, J.; Schölkopf, B.; and Mooij, J. M. 2016. Structural Causal Models: Cycles, Marginalizations, Exogenous Reparameterizations and Reductions. URL arxiv.org/abs/1611.06221v1.
- Gallow, J. D. forthcoming. A Model-Invariant Theory of Causation. *Philosophical Review* Available on PhilArchive: <https://philarchive.org/archive/GALCPA-8> (Last accessed: 03/12/2021).
- Halpern, J. Y. 2016a. *Actual Causality*. MIT Press.
- Halpern, J. Y. 2016b. Appropriate Causal Models and the Stability of Causation. *Review of Symbolic Logic* 9(1): 76–102.
- Pearl, J. 2000. *Causality: Models, Reasoning, and Inference*. Cambridge University Press. ISBN 0-521-77362-8.
- Rubenstein, P. K.; Weichwald, S.; Bongers, S.; Mooij, J. M.; Janzing, D.; Grosse-Wentrup, M.; and Schölkopf, B. 2017. Causal consistency of structural equation models. In *Proc. 33rd Conference on Uncertainty in Artificial Intelligence (UAI 2017)*.