

Nondeterministic Causal Models

Sander Beckers

Cornell University

SREKCEBREDNAS@GMAIL.COM

Editors: Biwei Huang and Mathias Drton

Abstract

I generalize acyclic deterministic structural causal models to the nondeterministic case and argue that this offers an improved semantics for counterfactuals. The standard, deterministic, semantics developed by Halpern (and based on the initial proposal of Galles & Pearl) assumes that for each assignment of values to parent variables there is a unique assignment to their child variable, and it assumes that the actual world (an assignment of values to all variables of a model) specifies a unique counterfactual world for each intervention. Both assumptions are unrealistic, and therefore I drop both of them in my proposal. I do so by allowing *multi-valued* functions in the structural equations. In addition, I adjust the semantics so that the solutions to the equations that obtained in the actual world are preserved in any counterfactual world. I provide a sound and complete axiomatization of the resulting logic and compare it to the standard one by Halpern and to more recent proposals that are closer to mine. Finally, I extend these models to the probabilistic case and show that they open up the way to identifying counterfactuals even in Causal Bayesian Networks.

Keywords: nondeterminism, counterfactuals, axiomatization

1. Introduction

Deterministic Structural Causal Models – DSCMs from now on – represent the causal relations between a set of endogenous variables by specifying an equation for each endogenous variable that determines the variable’s value as a function of the values of some other variables, both endogenous (\mathcal{V}) and exogenous (\mathcal{U}) (Pearl, 2009). Exogenous variables represent unobserved sources of variation whose existence has to be assumed in order to obtain *deterministic* equations.

DSCMs serve as the mathematical and conceptual foundation for Pearl’s causal modelling approach. As they lack probabilities, probabilistic causal models are built on top of DSCMs in two stages. First, we get probabilistic DSCMs – PDSCMs – by adding a probability distribution over the exogenous variables \mathcal{U} . This induces a joint distribution over $(\mathcal{U}, \mathcal{V})$. Second, assuming that there are no cyclic causal relations, that exogenous variables are independent, and that no two endogenous variables share an exogenous parent (i.e., that the model is *Markovian*), we get *Causal Bayesian Networks* by marginalizing out the exogenous variables and directly considering the marginal distribution over \mathcal{V} and its Markov factorization (Pearl, 2009, Ch.1)(Bareinboim et al., 2022, Th.2).

I here present *Nondeterministic Structural Causal Models* – NSCMs – as a more general and improved foundation for causal models. The reason for doing so is that the heavy reliance on exogenous variables in DSCMs is unnecessarily restrictive. The generalization from DSCMs to NSCMs is given by dropping the assumption that there must exist exogenous variables such that the value of each endogenous variable can be uniquely determined by its causal parents. Note that dropping this assumption is not the same as dropping the assumption that the world is itself fundamentally deterministic. It is one thing to assume that there must exist some properties of the world such that the value of a variable is uniquely determined, it is quite another to assume that

these properties can be neatly compartmentalized into sets of variables whose values determine the other value according to a stable functional relationship. In many situations the latter may not hold, and such situations are currently beyond the scope of SCMs. (And in even more situations, we may want to – or have to due to practical constraints – represent the world *as if* it does not hold.)

Furthermore, the use of exogenous variables also involves a commitment to an overly strong semantics for causal counterfactual statements, for it implies that there is always a *unique* solution to the model for any counterfactual query, given the actual values of all variables. That assumption is even stronger than the previous one. For example, consider a patient who may or may not receive treatment for some condition, and when she does she may or may not recover, depending on whether the treatment is effective. We observe that in the actual world, she does not receive treatment, and she does not recover. Representing this world using DSCMs requires specifying *actual* values of exogenous variables such that if she had been treated, she would have either certainly recovered, or she would certainly not have. Using NSCMs there is no such requirement, and thus we can simply consider both counterfactual worlds as possible. (See Examples 1 and 2 for the formal details.)

The benefit of dropping this counterfactual *uniqueness property* should not be underestimated. The property has been severely criticized in the general philosophical literature on counterfactuals ever since the seminal work of Lewis (1973) first did so. More relevant for the present purposes, it is one of the driving forces behind Dawid (2000)’s influential criticism of the counterfactual semantics offered by both DSCMs and the related *Potential Outcomes* framework (Rubin, 1974), and this criticism has resurfaced in the recent back-and-forth between Mueller and Pearl (2023) on the one hand and Dawid and Senn (2023) and Sarvet and Stensrud (2023) on the other over the correct methodology for personalised treatment decisions, such as in the patient example above.

It is worth quoting Dawid and Senn at length to illustrate how much their animosity towards counterfactuals is based on these two deterministic assumptions:

Indeed, there are serious philosophical objections to regarding potential responses [unique counterfactual outcomes] as having real existence. Only if we take a fully Laplacean view of the universe, in which the future of the universe is entirely determined by its present state and the laws of Physics, does this make any sense at all – and even then, it is difficult to incorporate the whims of an unconstrained external agent who decides whether or not to give treatment, or to account for the effect of external conditions arising after treatment. Even under Laplacean determinism, our ignorance of the information needed to predict the future means that we are unable to make use of it. Whether or not we believe in a deep-down deterministic universe, our predictions of the future can only be based on the limited information we do have at our disposal, and must necessarily be probabilistic. (Dawid and Senn, 2023, p. 5)

By offering a semantics for counterfactuals in nondeterministic causal models, I aim to provide room for both sides of this hotly contested – and literally vital – dispute to meet in the middle. Although I agree with Dawid and Senn that probabilistic predictions are crucial, I agree with Pearl and Halpern (2016) that an understanding of the logic for counterfactuals should precede a probabilistic account of counterfactuals, and hence that is where I start.

Concretely, I proceed as follows. I define NSCMs (Sec. 2), present a formal language and several corresponding semantics for causal formulas that hold in NSCMs (Sec. 3), compare them to the standard one for DSCMs (Sec. 4), offer a sound and complete axiomatization for two of the novel logics and compare them to other recent proposals that allow nondeterminism (Sec. 5), and

offer a preliminary investigation into the probabilistic generalization of NSCMs, illustrating how it allows for computing counterfactuals even in Causal Bayesian Networks (Sec. 6).

2. Nondeterministic Structural Causal Models

I take the definition of deterministic causal models by Halpern (2016) and generalize it to the non-deterministic case by using multi-valued functions.¹ As a first step, we need to define a signature as the variables out of which a causal model is built up.

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 $\mathbf{X} = (X_1, \dots, X_n)$, $\mathcal{R}(\mathbf{X})$ denotes the Cartesian product $\mathcal{R}(X_1) \times \dots \times \mathcal{R}(X_n)$. ■

A causal model expresses the causal relations between the endogenous variables of a signature. In addition to using multi-valued functions, I depart from Halpern by explicitly including the causal graph as an element of the causal model.

Definition 2 A causal model (or a Nondeterministic Structural Causal Model – NSCM) M is a triple $(\mathcal{S}, \mathcal{F}, \mathcal{G})$, where \mathcal{S} is a signature, \mathcal{G} is a directed graph such that there is one node for each variable in \mathcal{S} , and \mathcal{F} defines a function that associates with each endogenous variable X a structural equation F_X giving the possible values of X in terms of the values of some of the other endogenous and exogenous variables. A structural equation F_X takes on the form $X = f_X(\mathbf{Pa}_X)$, where $\mathbf{Pa}_X \subseteq (\mathcal{U} \cup \mathcal{V} - \{X\})$ are the parents of X as they appear in \mathcal{G} , and $f_X : \mathcal{R}(\mathbf{Pa}_X) \rightarrow \mathcal{P}(\mathcal{R}(X))$. (Here $\mathcal{P}(\mathcal{R}(X))$ is the powerset of $\mathcal{R}(X)$: the set that contains as its elements all subsets of $\mathcal{R}(X)$). ■

We here restrict attention to the case in which \mathcal{G} is acyclic (so that \mathcal{G} is a DAG – a Directed Acyclic Graph). If for each f_X the co-domain does not contain the empty set, we say that a causal model is *total*. We here restrict ourselves to total causal models. This amounts to the assumption that for all possible settings \mathbf{pa}_X of the parents, there exists at least one solution x for the child.

There are no equations for exogenous variables \mathcal{U} , as these are taken to represent the background conditions that are simply given. We call $\mathbf{u} \in \mathcal{R}(\mathcal{U})$ a *context*, $\mathbf{v} \in \mathcal{R}(\mathcal{V})$ a *state*, and $(\mathbf{u}, \mathbf{v}) \in \mathcal{R}(\mathcal{U} \cup \mathcal{V})$ is a *world*. In *deterministic* causal models all the functions f_X are standard as opposed to multi-valued, and thus each equation has a unique solution x for each choice of values \mathbf{pa}_X . In nondeterministic models, a *solution* of the equation $X = f_X(\mathbf{Pa}_X)$ is a tuple (x, \mathbf{pa}_X) such that $x \in f_X(\mathbf{pa}_X)$. (The term “equation” is thus somewhat strange, but we stick with it given the common reference to causal models as “structural equation models”. Furthermore, as pointed out to me by Joris Mooij, an equivalent characterization can be given in terms of a literal equation as well.² A solution of M is a world (\mathbf{u}, \mathbf{v}) that is a solution of all equations in \mathcal{F} .

1. Halpern (2000) already suggested this generalization, but never implemented it. He did recently offer an even further generalization in order to allow for an infinite number of variables with infinite ranges, by doing away with equations altogether (Peters and Halpern, 2021; Halpern and Peters, 2022). Recently Barbero (2024) and Wysocki (2023) have likewise offered generalizations to the nondeterministic case. I compare my approach to these ones in Section 5.

2. Concretely, this can be done using the approach of Bongers et al. (2021). Although they do not explicitly consider nondeterministic models, their models do allow “self-cycles”, meaning they allow for X to depend on X . Note that $X = f_X(\mathbf{Pa}_X)$ is equivalent to $X = X + (1 - 1_{f_X}(X, \mathbf{Pa}_X))$, where 1_{f_X} is the indicator function that returns 1 if $X \in f_X(\mathbf{Pa}_X)$ and 0 otherwise. This correspondence opens up the possibility of integrating both approaches.

In acyclic models the solutions of M given a context \mathbf{u} can be determined recursively: determine the solutions of each equation in the partial order given by \mathcal{G} , and pass these on to the next equation. In general, deterministic causal models (DSCMs) need not have solutions for each context, nor do any such solutions have to be unique, but unicity and existence are guaranteed for acyclic deterministic models. When we move to nondeterministic causal models, unicity and existence are no longer guaranteed even for acyclic models, but existence is recovered in total models.

3. The Causal Language

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 *basic formula (over \mathcal{S})* φ is a Boolean combination of atomic formulas. An *intervention* has the form $\mathbf{Y} \leftarrow \mathbf{y}$, where Y_1, \dots, Y_k are distinct variables in \mathcal{V} , and $y_i \in \mathcal{R}(Y_i)$ for each $1 \leq i \leq k$. A *basic causal formula* has the form $[Y_1 \leftarrow y_1, \dots, Y_k \leftarrow y_k]\varphi$, where φ is a basic formula and Y_1, \dots, Y_k are distinct variables in \mathcal{V} . Such a formula is abbreviated as $[\mathbf{Y} \leftarrow \mathbf{y}]\varphi$. The special case where $k = 0$ is abbreviated as $\llbracket \varphi$. Finally, a *causal formula* is a Boolean combination of basic causal formulas. The language $\mathcal{L}(\mathcal{S})$ that we consider consists of all causal formulas.

A causal formula ψ is true or false in a causal model, given a world. We write $(M, \mathbf{u}, \mathbf{v}) \models \psi$ if the causal formula ψ is true in causal model M given world (\mathbf{u}, \mathbf{v}) . We call a model-world pair $(M, \mathbf{u}, \mathbf{v})$ a *causal setting*.

We first define the \models relation for basic formulas. $(M, \mathbf{u}, \mathbf{v}) \models X = x$ if x is the restriction of \mathbf{v} to X . (From now on we abuse notation and write this simply as $x \in \mathbf{v}$, with the understanding that lower-case letters refer to values of the variables with that name. Similarly, we write $\mathbf{x} \subseteq \mathbf{v}$ whenever we want to allow that \mathbf{X} consists of multiple variables.) We extend \models to basic formulas φ in the standard way. Note that the truth of basic formulas is determined solely by the state \mathbf{v} , and thus we often also write $\mathbf{v} \models \varphi$.

In order to define the \models relation for causal formulas, we introduce two operations on a causal model, the *actualized refinement* that is the result of integrating the actual behavior of the equations as observed in a world (\mathbf{u}, \mathbf{v}) into the equations of a model M , and the *intervened* model that is the result of performing an intervention on the equations of a model M .

Definition 3 *Given causal models M' and M over identical signatures \mathcal{S} and with identical graphs \mathcal{G} , we say that M' is a refinement of M if for all $X \in \mathcal{V}$ and all $\mathbf{pa}_X \in \mathcal{R}(\mathbf{Pa}_X)$: $f_X^{M'}(\mathbf{pa}_X) \subseteq f_X^M(\mathbf{pa}_X)$. ■*

Since we are restricting ourselves to total causal models, the only refinement of a deterministic model is itself. (Concretely, if we drop totality, then a refinement allows for $f_X(\mathbf{pa}_X) = \emptyset$, which does not result in a deterministic model.) Hence we say that a deterministic model is *maximally refined*. Concretely, since the way that the equations determine the outcome in a deterministic model is identical across each world (\mathbf{u}, \mathbf{v}) , there is no need to additionally consider how the equations behave in a specific world. This is no longer true for nondeterministic models, as there the actual values that obtained in a world inform us about how the equations behaved for those values (in that specific world). When evaluating formulas in a world, we need to take this information into account, and to do so requires refining the equations so that they incorporate this actual behavior.

Definition 4 *Given a solution (\mathbf{u}, \mathbf{v}) of a model $M = (\mathcal{S}, \mathcal{F}, \mathcal{G})$, we define the actualized refinement $M^{(\mathbf{u}, \mathbf{v})}$ as the refinement of M in which \mathcal{F} is replaced by $\mathcal{F}^{(\mathbf{u}, \mathbf{v})}$, as follows: for each variable*

$X \in \mathcal{V}$, its function f_X is replaced by $f_X^{(\mathbf{pa}_X, x)}$ that behaves identically to f_X for all inputs except for \mathbf{pa}_X . Instead, $f_X^{(\mathbf{pa}_X, x)}(\mathbf{pa}_X) = x$, where $(x, \mathbf{pa}_X) \subseteq (\mathbf{u}, \mathbf{v})$. ■

Setting the value of some variables \mathbf{Y} to \mathbf{y} in a causal model $M = (\mathcal{S}, \mathcal{F}, \mathcal{G})$ results in a new causal model, denoted $M_{\mathbf{Y} \leftarrow \mathbf{y}}$, which is identical to M , except that \mathcal{F} is replaced by $\mathcal{F}^{\mathbf{Y} \leftarrow \mathbf{y}}$: for each variable $X \notin \mathbf{Y}$, $F_X^{\mathbf{Y} \leftarrow \mathbf{y}} = F_X$ (i.e., the equation for X is unchanged), while for each Y' in \mathbf{Y} , the equation $F_{Y'}$ for Y' is replaced by $Y' = y'$ (where $y' \in \mathbf{y}$). Similarly, \mathcal{G} is replaced with $\mathcal{G}^{\mathbf{Y} \leftarrow \mathbf{y}}$.

With these operations in place, we can define the \models relation for basic causal formulas, relative to settings $(M, \mathbf{u}, \mathbf{v})$ such that (\mathbf{u}, \mathbf{v}) is a solution of M . $(M, \mathbf{u}, \mathbf{v}) \models [\mathbf{Y} \leftarrow \mathbf{y}] \varphi$ iff $\mathbf{v}' \models \varphi$ for all states \mathbf{v}' such that $(\mathbf{u}, \mathbf{v}')$ is a solution of $(M^{(\mathbf{u}, \mathbf{v})})_{\mathbf{Y} \leftarrow \mathbf{y}}$. We inductively extend the semantics to causal formulas in the standard way, that is, $(M, \mathbf{u}, \mathbf{v}) \models [\mathbf{Y} \leftarrow \mathbf{y}] \varphi_1 \wedge [\mathbf{Z} \leftarrow \mathbf{z}] \varphi_2$ iff $(M, \mathbf{u}, \mathbf{v}) \models [\mathbf{Y} \leftarrow \mathbf{y}] \varphi_1$ and $(M, \mathbf{u}, \mathbf{v}) \models [\mathbf{Z} \leftarrow \mathbf{z}] \varphi_2$, and similarly for \neg and \vee .

We define $\langle \mathbf{Y} \leftarrow \mathbf{y} \rangle \varphi$ as an abbreviation of $\neg[\mathbf{Y} \leftarrow \mathbf{y}] \neg \varphi$. So $(M, \mathbf{u}, \mathbf{v}) \models \langle \mathbf{Y} \leftarrow \mathbf{y} \rangle \text{true}$ iff there is some world $(\mathbf{u}, \mathbf{v}')$ that is a solution of $(M^{(\mathbf{u}, \mathbf{v})})_{\mathbf{Y} \leftarrow \mathbf{y}}$. We then write $(M, \mathbf{u}, \mathbf{v}) \models \langle \mathbf{Y} \leftarrow \mathbf{y} \rangle \mathcal{V} = \mathbf{v}'$.

We can also evaluate formulas with respect to just a partial causal setting (M, \mathbf{u}) , or even with respect to a model M by itself. For basic causal formulas, we define that $(M, \mathbf{u}) \models [\mathbf{Y} \leftarrow \mathbf{y}] \varphi$ iff $(M, \mathbf{u}, \mathbf{v}) \models [\mathbf{Y} \leftarrow \mathbf{y}] \varphi$ holds for all states \mathbf{v} such that (\mathbf{u}, \mathbf{v}) is a solution of M . In a similar fashion, we define that $M \models [\mathbf{Y} \leftarrow \mathbf{y}] \varphi$ iff $(M, \mathbf{u}, \mathbf{v}) \models [\mathbf{Y} \leftarrow \mathbf{y}] \varphi$ holds for all solutions (\mathbf{u}, \mathbf{v}) of M . We again inductively extend to causal formulas in the standard way: $(M, \mathbf{u}) \models [\mathbf{Y} \leftarrow \mathbf{y}] \varphi_1 \wedge [\mathbf{Z} \leftarrow \mathbf{z}] \varphi_2$ iff $(M, \mathbf{u}) \models [\mathbf{Y} \leftarrow \mathbf{y}] \varphi_1$ and $(M, \mathbf{u}) \models [\mathbf{Z} \leftarrow \mathbf{z}] \varphi_2$, and similarly for \neg and \vee . Likewise for $M \models \psi$.

We have now defined three different semantics for $\mathcal{L}(\mathcal{S})$: the first with respect to full causal settings $(M, \mathbf{u}, \mathbf{v})$, the second with respect to partial causal settings (M, \mathbf{u}) , and the third with respect to M . The first evaluates causal formulas relative to a single world (\mathbf{u}, \mathbf{v}) , and thus we will call the resulting logic the *single world counterfactual logic*, or **swc** logic for short. It formalizes counterfactual statements relative to a specific *actual world*. The second only requires specifying a single context \mathbf{u} , and thus it considers the entire set of worlds (\mathbf{u}, \mathbf{v}) that extend \mathbf{u} together, and thus we call the resulting logic the *single context counterfactual logic*, or **scc** logic for short. Lastly, the third semantics evaluates formulas without restricting the possible worlds any further than the model M already does, and thus we call the resulting logic the *single model counterfactual logic*, or **smc** logic for short. In the remainder of this paper our attention goes to the first two counterfactual logics, but I flag the **smc** logic as a worthwhile subject for future investigation.

4. Comparison to the Standard Semantics

To recap, the standard logic for causal models by Halpern (2000, 2016) is a logic for DSCMs instead of NSCMs, and its semantics is defined identically to mine except that it does not use the actualized refinement. Given that for any DSCM the actualized refinement $M^{(\mathbf{u}, \mathbf{v})}$ will simply be M anyway, the two semantics are in fact entirely equivalent for DSCMs.

Concretely, the standard semantics defines the \models relation for basic causal formulas as: $(M, \mathbf{u}, \mathbf{v}) \models [\mathbf{Y} \leftarrow \mathbf{y}] \varphi$ iff $\mathbf{v}' \models \varphi$ for all states \mathbf{v}' such that $(\mathbf{u}, \mathbf{v}')$ is a solution of $M_{\mathbf{Y} \leftarrow \mathbf{y}}$. Note that, importantly, this semantics does not depend on the actual state \mathbf{v} , and therefore it can be (and usually is) written and interpreted as a semantics for partial causal settings (M, \mathbf{u}) . As a result, the distinction

between the counterfactual logics **swc** and **scc** collapses in the case of DSCMs, and thus there is only a single counterfactual logic.

If the DSCMs are acyclic, then for each u there is a unique v such that (u, v) is a solution of M , and thus the counterfactual logic is with respect to a single actual world, as is the case for our **swc**. Also, for each u , there is a unique v such that (u, v) is a solution of $M_{Y \leftarrow y}$ for any $Y \leftarrow y$, and hence the standard semantics for acyclic DSCMs satisfies the controversial uniqueness property mentioned in Section 1. If the DSCMs are cyclic, then for each context there may be multiple solutions, or none. Thus, in this case the counterfactual logic is with respect to the set of worlds that extend a context u , as is the case for our **scc** (except that for **scc** a solution is guaranteed to exist, given totality). As I discuss in Section 5, these semantics still satisfy a property that is conceptually very similar to the uniqueness property.

Importantly, as the following Theorem shows, when we move from DSCMs to NSCMs, the actualized refinement *only matters for swc*, and thus for both the **scc** and **smc** logics my semantics are simply generalizations of the standard semantics to the nondeterministic case.

Theorem 5 *Given a nondeterministic causal model M , we have that for all $Y \subseteq \mathcal{V}$, for all $y \in \mathcal{R}(Y)$, and for all basic formulas φ :*

- $M \models [Y \leftarrow y]\varphi$ iff $v \models \varphi$ for all solutions (u, v) of $M_{Y \leftarrow y}$.
- For all contexts u : $(M, u) \models [Y \leftarrow y]\varphi$ iff $v \models \varphi$ for all states v such that (u, v) is a solution of $M_{Y \leftarrow y}$.

Proof: All proofs are to be found in the Appendix. ■

The intuition behind Theorem 5 is that the actualized refinement operation updates M 's equations to include their *actual behavior*, and there is no actual behavior unless one specifies an actual world (u, v) , so the operation can be ignored for the **scc** and **smc** logics.

As far as I am aware, the distinction between **swc** and **scc** logics has gone unnoticed so far. In fact, recently Halpern and Peters (2022) introduced generalized structural equation models – GSEMs – and they provide an axiomatization for the more general logic that is the result, but in doing so Halpern has moved entirely from an **swc** type of logic – that was the subject matter of (Halpern, 2000) and (Galles and Pearl, 1998) – towards a **scc** type of logic, for the logic for GSEMs does not allow evaluating formulas with respect to a single world. In the next section I conclude the definition of the **swc** and **scc** logics by offering a sound and complete axiomatization for both, and interpret the results in light of other recently proposed logics for nondeterministic causal reasoning.

5. Axiomatization and Recent Related Work

Throughout this section we hold fixed some finite signature $\mathcal{S} = (\mathcal{U}, \mathcal{F}, \mathcal{R})$, i.e., \mathcal{U} and \mathcal{V} are finite, and $\mathcal{R}(X)$ is finite for all $X \in \mathcal{U} \cup \mathcal{V}$. Let AX be the axiom system for the language $\mathcal{L}(\mathcal{S})$ that consists of the following list of axioms and inference rule MP.

- D0. All instances of propositional tautologies.
- D1. $[Y \leftarrow y](X = x \Rightarrow X \neq x')$ if $x, x' \in \mathcal{R}(X)$, $x \neq x'$ (functionality)
- D2. $[Y \leftarrow y](\bigvee_{x \in \mathcal{R}(X)} X = x)$ (definiteness)

- D3(a). $\langle \mathbf{X} \leftarrow \mathbf{x} \rangle (W = w \wedge \varphi) \Rightarrow \langle \mathbf{X} \leftarrow \mathbf{x}, W \leftarrow w \rangle (\varphi)$ if $W \notin \mathbf{X}$ (weak composition)
- D3(b). $[\mathbf{X} \leftarrow \mathbf{x}] (W = w \wedge \varphi) \Rightarrow [\mathbf{X} \leftarrow \mathbf{x}, W \leftarrow w] (\varphi)$ if $W \notin \mathbf{X}$ (strong composition)
- D4. $[\mathbf{X} \leftarrow \mathbf{x}] (\mathbf{X} = \mathbf{x})$ (effectiveness)
- D5. $(\langle \mathbf{X} \leftarrow \mathbf{x}, Y \leftarrow y \rangle (W = w \wedge \mathbf{Z} = \mathbf{z}) \wedge \langle \mathbf{X} \leftarrow \mathbf{x}, W \leftarrow w \rangle (Y = y \wedge \mathbf{Z} = \mathbf{z})) \Rightarrow \langle \mathbf{X} \leftarrow \mathbf{x} \rangle (W = w \wedge Y = y \wedge \mathbf{Z} = \mathbf{z})$ if $\mathbf{Z} = \mathcal{V} - (\mathbf{X} \cup \{W, Y\})$ (reversibility)
- D6. $(X_0 \rightsquigarrow X_1 \wedge \dots \wedge X_{k-1} \rightsquigarrow X_k) \Rightarrow \neg (X_k \rightsquigarrow X_0)$ (recursiveness)
- D7. $([\mathbf{X} \leftarrow \mathbf{x}] \varphi \wedge [\mathbf{X} \leftarrow \mathbf{x}] (\varphi \Rightarrow \psi)) \Rightarrow [\mathbf{X} \leftarrow \mathbf{x}] \psi$ (distribution)
- D8. $[\mathbf{X} \leftarrow \mathbf{x}] \varphi$ if φ is a propositional tautology (generalization)
- D9. $\langle \mathbf{Y} \leftarrow \mathbf{y} \rangle \text{true} \wedge (\langle \mathbf{Y} \leftarrow \mathbf{y} \rangle \varphi \Rightarrow [\mathbf{Y} \leftarrow \mathbf{y}] \varphi)$ if $\mathbf{Y} = \mathcal{V}$ or, for some $X \in \mathcal{V}$, $\mathbf{Y} = \mathcal{V} - \{X\}$ (unique outcomes for \mathcal{V} and $\mathcal{V} - \{X\}$)
- D10(a). $\langle \mathbf{Y} \leftarrow \mathbf{y} \rangle \text{true}$ (at least one outcome)
- D10(b). $\langle \mathbf{Y} \leftarrow \mathbf{y} \rangle \varphi \Rightarrow [\mathbf{Y} \leftarrow \mathbf{y}] \varphi$ (at most one outcome)
- D10(c). $\langle \rangle \varphi \Rightarrow \Box \varphi$ (at most one actual outcome)
- MP. From φ and $\varphi \Rightarrow \psi$, infer ψ (modus ponens)

Here, $Y \rightsquigarrow Z$ means that $Y \neq Z$ and $\forall \mathbf{X} \subseteq \mathcal{V}, \mathbf{x} \in \mathcal{R}(\mathbf{X}), y \neq y' \in \mathcal{R}(Y), z \in \mathcal{R}(Z) (\langle \mathbf{X} \leftarrow \mathbf{x}, Y \leftarrow y \rangle (Z = z) \wedge [\mathbf{X} \leftarrow \mathbf{x}, Y \leftarrow y'] (Z \neq z))$.

Halpern and Peters (2022) define \rightsquigarrow differently. However, my definition is easily seen to be equivalent to theirs in the presence of D10(b). Note that none of the axioms mentions basic formulas φ . For both the standard semantics and for mine this is without loss of generality, because φ is easily seen to be equivalent to $\Box \varphi$. (Alternatively, I could have extended the causal language and add the axiom $\varphi \Leftrightarrow \Box \varphi$, which is the choice made by Barbero (2024).) Lastly, note that in the presence of D10(b), AX contains redundant axioms: D3(b) is a consequence of D3(a) and D10(b), D9 is a consequence of D2 and D10(a,b), D10(c) is a consequence of D10(b), and – as shown by Halpern and Peters (2022) – D5 is a consequence of D2, D3, D6, D7, D8, and D10(a,b).

Halpern and Peters (2022) show that AX without D3(b) and D10(c) is a sound and complete axiomatization with respect to acyclic DSCMs that have signature \mathcal{S} , and thus the same holds for AX . Furthermore, none of the D10 axioms holds for cyclic DSCMs. In fact, Halpern and Peters (2022) show that AX without D3(b), D10(a,b,c), and D6, is a sound and complete axiomatization for cyclic DSCMs.

D10(b) implies the *uniqueness* property mentioned earlier. To see why, consider some solution (\mathbf{u}, \mathbf{v}) (or context \mathbf{u}) relative to which we are evaluating formulas, and take φ to be $\mathcal{V} = \mathbf{v}'$ for some $\mathbf{v}' \in \mathcal{V}$. D10(b) then commits us to the claim that if the world $(\mathbf{u}, \mathbf{v}')$ is possible under the counterfactual supposition that \mathbf{Y} were \mathbf{y} , then it is the only world that is possible under that supposition. Although cyclic DSCMs do not validate D10(b), they do validate D9, and that is conceptually not at all different: it simply restricts the above claim to the special case in which the counterfactual supposition contains all but one endogenous variable. As Example 1 illustrates, neither axioms are sound for either the **swc** logic or the **scc** logic, and thus they free causal models from their commitment to this controversial property. (For sake of completeness, note that neither axiom is sound for the **smc** logic either, nor is it for the **smc** logic over DSCMs only.)

Example 1 Consider again our patient, and assume for simplicity that their condition is certain to kill them if untreated but they may recover if treated, and we observe that they are not treated. We can model this with just two binary endogenous variables X, Y , no exogenous variables, and a graph $X \rightarrow Y$. The equation for Y is $Y \in \{0, 1\}$ if $X = 1$ and $Y = 0$ if $X = 0$. The equation for X is irrelevant. Note that D10(b) and D9 – combined with D10(a) – demand that one counterfactual outcome is certain, i.e., that $(M, X = 0, Y = 0) \models [X \leftarrow 1]Y = 1 \vee [X \leftarrow 1]Y = 0$. This consequence of uniqueness is called the principle of conditional excluded middle and its unacceptability is the main source of Lewis (1973)’s classical criticism of the uniqueness property. Yet here we have that $(M, X = 0, Y = 0) \models \langle X \leftarrow 1 \rangle Y = 1 \wedge \langle X \leftarrow 1 \rangle Y = 0$, expressing that if the patient were treated they might have recovered but also if they were treated they might not have recovered. As for the **scc**, note that similarly $M \models \langle X \leftarrow 1 \rangle Y = 1$ and yet $M \not\models [X \leftarrow 1]Y = 1$.

We let AX_{non}^{scc} denote the axiom system D0-D8 and D10(a), and AX_{non}^{swc} denotes AX_{non}^{scc} plus D10(c). The following result shows that the **swc** logic is stronger than the **scc** logic, that both of them are weaker than the counterfactual logic for acyclic DSCMS, and that both of them are incomparable to the counterfactual logic for cyclic DSCMS.

Theorem 6 AX_{non}^{swc} (resp. AX_{non}^{scc}) is a sound and complete axiomatization for the language $\mathcal{L}(\mathcal{S})$ with respect to the **swc** logic (resp. the **scc** logic) over acyclic NSCMs that have signature \mathcal{S} .

D10(c) is what distinguishes the two logics, as it does not hold for the **scc** logic. Simply consider the model M consisting of a single binary endogenous variable X with an equation such that $X = \{0, 1\}$. Then we have that both $M \models \langle \rangle X = 1$ and $M \models \langle \rangle X = 0$ and thus also $M \not\models \Box X = 1$. That D10(c) is not sound for the **scc** logic is what prevents the first problematic assumption of acyclic DSCMS mentioned in Section 1, namely that the values u for all exogenous variables uniquely determine an actual world (u, v) . (As with D10(b), to see this take φ to be of the form $\mathcal{V} = v'$.) In the context of the **swc** logic, D10(c) simply expresses the sensible property that if the world were as it actually is, then the actual world is the only world possible. To see this, note that for any (M, u, v) , the only way for $\langle \rangle \mathcal{V} = v'$ to hold is when $v' = v$.

5.1. Related Work

Theorem 6 shows that both of the **swc** and **scc** logics are stronger than the one introduced in (Halpern and Peters, 2022) for so-called GSEMs, *Generalized Structural Equation Models*, even when restricting to the case that is most similar to ours (the case in which all interventions are considered, the signature is finite, and the graph is acyclic). Briefly, GSEMs do away with structural equations altogether, and simply define there to be some set of states v corresponding to each context-intervention pair $(u, Y \leftarrow y)$. As a result, this means that the connection between formulas that hold under some intervention $X \leftarrow x$ and those that should hold under an extension $X \leftarrow x, W \leftarrow w$ is lost, and thus the axioms D3 (both of them) and D5 are no longer sound, in addition to losing D9 and D10(a,b,c). (See their Theorem 5.3.)

The closest logic to mine is that of Barbero (2024), as it also generalizes structural causal models by allowing for multi-valued functions in the equations. (His work formed part of the motivation for the current paper.) Except for the fact that he also considers the cyclic case, there are two main differences with my proposal.

Firstly, his logic builds on *team semantics*, meaning that instead of evaluating formulas with respect to either a single solution (u, v) or a single context u of a model M , he does so with respect

to a *team*, where a team is any subset of solutions of a model M . As a result, the logic is a mixture of all three of my logics: consider a singleton team and causal formulas are effectively single world counterfactual formulas, consider a team that consists of all solutions extending some context \mathbf{u} , and causal formulas are effectively single context counterfactual formulas, consider a team that consists of all solutions, and causal formulas are single model counterfactual formulas, consider a team that belongs to none of these categories and the causal formula is some form of hybrid.

Second, although his semantics also aims to incorporate the actual behavior of the equations into all counterfactual worlds, his semantics is less demanding than my actualized refinements are. It merely requires that the values of all non-descendants of the variables in some intervention $\mathbf{X} \leftarrow \mathbf{x}$ remain identical to the actual ones in counterfactual worlds.³ It is easy to show that this property also holds for my semantics, i.e., for any variable Y that is not a descendant of any variable in \mathbf{X} , we have for any solution (\mathbf{u}, \mathbf{v}) of M that if $(M, \mathbf{u}, \mathbf{v}) \models Y = y$ then $(M, \mathbf{u}, \mathbf{v}) \models [\mathbf{X} \leftarrow \mathbf{x}]Y = y$.

As a result of the second difference, weak composition – D3(a) – comes apart from strong composition – D3(b) – for Barbero. Just as with the **swc** and **scc** logics, D3(a) is sound and D10(b) is not. However, D3(b) is no longer sound for his logic. I believe this is an undesirable property. Consider again the model from Example 1, and consider the world $(M, X = 1, Y = 1)$. Then according to both Barbero and the **swc** and **scc** logics it holds that $(M, X = 1, Y = 1) \models \Box Y = 1$, yet for Barbero’s logic we do not have that $(M, X = 1, Y = 1) \models [X \leftarrow 1]Y = 1$, the reason being that for his logic $(M, X = 1, Y = 1) \models \langle X \leftarrow 1 \rangle Y = 0$. As Barbero points out, this results in his logic not validating *conjunction conditionalization*. Yet in the causal inference literature this property is considered fundamental and goes by the name *consistency* (Robins, 1987; Pearl, 2009). (See (Walters and Williams, 2013) for a thorough defense of this property in the philosophical literature.) Therefore I view it as an argument in favor of my approach that it *does* satisfy it:⁴

Proposition 7 *If $(M, \mathbf{u}, \mathbf{v}) \models \mathbf{X} = \mathbf{x} \wedge \varphi$ then $(M, \mathbf{u}, \mathbf{v}) \models [\mathbf{X} \leftarrow \mathbf{x}]\varphi$.*

As to the first difference, it is not clear to me what the notion of a team is supposed to capture. The most straightforward interpretation would be that it expresses the set of “possible worlds” in some sense or other. However, that interpretation conflicts with how teams behave. Again considering Example 1, we add that the team consists exclusively of the solution $(X = 1, Y = 1)$. If we then consider the intervention $X \leftarrow 1$, the team suddenly changes to also including the solution $(X = 1, Y = 0)$. I fail to understand why interventions that do not change anything should be able to result in changing what worlds we consider possible. Another consequence of his use of teams is that his logic invalidates D10(c), contrary to the **swc** logic.

Wysocki (2023) also recently proposed to generalize structural causal models using multi-valued functions. His logic, however, only considers formulas with respect to an entire causal model M , and is thus what I have called a single model counterfactual logic. Therefore his work complements the current work, as we here focus on the other two logics instead. There is one important difference between his logic and my **smc** logic though, namely in how we view the idea of a causal parent.

He follows Halpern’s standard approach in that he *defines* the parent relation between variables rather than reading it of an additional graph \mathcal{G} that I (and Barbero) have added to the definition of a

3. At least, that is the intended semantics. When moving from the cyclic to the acyclic case, Barbero drops the requirement on non-descendants. Personal communication with the author confirms that this was unintended.

4. For what it’s worth, I also believe that Barbero’s defense is confused. He discusses exactly this example ($X = 1$ is Alice flipping a coin and $Y = 1$ is its landing heads) and he is explicit that the formula is a counterfactual statement, yet then he says that it expresses what would happen if Alice flips the coin *again*, which is of course not a counterfactual statement about Alice’s present coin flip but a prediction about Alice’s future coin flip.

causal model. Concretely, for DSCMs Halpern (2016) takes the equation for a variable Y to consist of a function $f_Y : \mathcal{R}(\mathcal{U} \cup \mathcal{V} - \{Y\}) \rightarrow \mathcal{R}(Y)$, and then defines Y depends on X to mean that there exist $x \neq x' \in \mathcal{R}(X)$ and $z \in \mathcal{R}(\mathcal{U} \cup \mathcal{V} - \{X, Y\})$ such that $f_Y(z, x) \neq f_Y(z, x')$. Wysocki uses the same definition, except that he generalizes it to the multi-valued case. Both Halpern and Wysocki then take “ X is a parent of Y ” to be synonymous with “ Y depends on X ”. Clearly, if Y depends on X then it will also be a parent in my framework, but the reverse need not hold. In fact, following Barbero, we can distinguish between the graph \mathcal{G}_M that comes with a model and the graph \mathcal{G}_D that is built up out of all the dependence relations, noting that \mathcal{G}_D is a subgraph of \mathcal{G}_M .⁵ Although nothing prevents one from restricting my framework to models M in which $\mathcal{G}_M = \mathcal{G}_D$, I believe it is a mistake to enforce this stricter definition of a parent in general, because it prohibits a natural extension to probabilistic causal models. Viewed probabilistically, if conditional on all other variables, X can change the *distribution* of Y , then X is a parent of Y . Changing the distribution does not require than any value y changes from being possible ($p > 0$) to impossible ($p = 0$) or vice versa, and yet that is what is required for Y to depend on X (in the non-probabilistic sense defined above). Furthermore, one might want to allow for X being a parent of Y even if it *does not* meet this probabilistic requirement, because of a failure of *faithfulness* (Spirtes et al., 2001). This is consistent with my more permissive notion of a parent.

6. Probabilistic Causal Models

As mentioned in the introduction, DSCMs give rise to probabilistic DSCMs and – if Markovian – these in turn give rise to Causal Bayesian Networks – CBNs. I here offer an initial study of the nondeterministic counterpart to this construction, starting with a generalization of probabilistic causal models. (Interestingly, Dawid and Musio (2022) already briefly consider this generalization, calling them *Stochastic Causal Models*. Unfortunately they quickly dismiss them (see Sec. 12.4) by conflating the problem of offering a semantics for counterfactuals with issues involving actual causation. I discuss the latter within the context of nondeterministic models in (Beckers, 2025).)

Definition 8 A probabilistic causal model (or a PNSCM) M is a 4-tuple $(\mathcal{S}, \mathcal{F}, \mathcal{G}, \mathcal{P}_U)$, where \mathcal{S} is a signature, \mathcal{G} is an acyclic directed graph such that there is one node for each variable in \mathcal{S} , and \mathcal{F} defines a function that associates with each endogenous variable X a family of conditional probability distributions $P_X(X|\mathbf{Pa}_X)$ over $\mathcal{R}(X)$, giving the probability of the values of X in terms of the parents \mathbf{Pa}_X of X as they appear in \mathcal{G} . \mathcal{P}_U is a probability distribution over $\mathcal{R}(U)$. ■

A PNSCM is *Markovian* if all exogenous variables are mutually independent and do not share a child. We assume that Markovian PNSCMs satisfy the *Causal Markov Condition*, which means that the joint distribution P_M over $\mathcal{R}(U \times \mathcal{V})$ is given as: $P_M(U, \mathcal{V}) = \prod_{X \in \mathcal{V}} P_X(X|\mathbf{Pa}_X) \prod_{U \in \mathcal{U}} P_U(U)$. A *solution* of M is a world (\mathbf{u}, \mathbf{v}) such that $P_M(\mathbf{u}, \mathbf{v}) > 0$.

The Markov condition continues to hold when we marginalize out the exogenous variables to consider $P_M(\mathcal{V})$, for we have that $P_M(\mathcal{V}) = \prod_{X \in \mathcal{V}} P_X^\mathcal{V}(X|\mathbf{EPa}_X)$, where \mathbf{EPa}_X are X ’s endogenous parents and $P_X^\mathcal{V}$ is the result of marginalizing P_X over X ’s exogenous parents \mathbf{XPa}_X : $P_X^\mathcal{V}(X|\mathbf{EPa}_X) = \sum_{\{\mathbf{xpa}_X \in \mathcal{R}(\mathbf{XPa}_X)\}} P_X(X|\mathbf{EPa}_X, \mathbf{xpa}_X) P_U(\mathbf{xpa}_X)$. (See the Appendix for details.) Therefore we can proceed entirely analogously to the deterministic case and construct

5. As I discuss in the Appendix, axiom D6 corresponds to the property that \mathcal{G}_D is acyclic, which is a consequence of the acyclicity of \mathcal{G}_M . It is unclear whether it is possible to also express the acyclicity of \mathcal{G}_M in the causal language.

Causal Bayesian Networks by taking a Markovian probabilistic causal model and marginalizing over the exogenous variables \mathcal{U} (Pearl, 2009, Th. 1.4.1)(Bareinboim et al., 2022, Th.2).

Definition 9 Given a Markovian PNSCM $M = (\mathcal{S}, \mathcal{F}, \mathcal{G})$, the Causal Bayesian Network C_M induced by M is the tuple $(\mathcal{S}_V, P_M(\mathcal{V}), \mathcal{G}_V)$, where \mathcal{S}_V and \mathcal{G}_V are obtained from \mathcal{S} and \mathcal{G} by removing the exogenous variables and all of their edges. ■

Per construction, a CBN and its underlying PNSCM have identical *observational distributions* $P_M(\mathcal{V})$, which represent the first layer of Pearl’s Causal Hierarchy (Bareinboim et al., 2022). We arrive at the second layer, that of *interventional distributions* $P_{M_{X \leftarrow x}}(\mathcal{V})$ (usually written as $P(\mathcal{V}|do(x))$) in exactly the same manner as in the deterministic case, namely by considering the distribution corresponding to the intervened model $M_{X \leftarrow x}$:

$$P_{M_{X \leftarrow x}}(\mathcal{V} = v) = \begin{cases} \prod_{\{Y \in \mathcal{V} \setminus X\}} P_Y^{\mathcal{V}}(Y | \mathbf{E} \mathbf{P} \mathbf{a}_Y) & \text{if } x \subseteq v \\ 0 & \text{otherwise} \end{cases}$$

Just as in the deterministic case, this distribution only depends on knowledge of $P_M(\mathcal{V})$ and the graph \mathcal{G}_V , and therefore it can be viewed interchangeably as an interventional distribution of M or as an interventional distribution of C_M . So far nothing new. However, the nondeterministic and deterministic case start diverging once we move up to the third and most expressive *counterfactual* layer of Pearl’s Causal Hierarchy. To get there, we first need to define the probabilistic generalization of the actualized refinement (Def. 4).

Definition 10 Given a solution (u, v) of a probabilistic model $M = (\mathcal{S}, \mathcal{F}, \mathcal{G}, \mathcal{P}_U)$, we define the actualized refinement $M^{(u, v)}$ as the model in which \mathcal{F} is replaced by $\mathcal{F}^{(u, v)}$, as follows: for each variable $X \in \mathcal{V}$ and $(x, \mathbf{pa}_X) \subseteq (u, v)$, the family of distributions P_X is replaced with the distributions $P_X^{(\mathbf{pa}_X, x)}(X | \mathbf{pa}_X')$ that are identical to $P_X(X | \mathbf{pa}_X')$ for all $\mathbf{pa}_X' \in \mathcal{R}(\mathbf{pa}_X)$ except for \mathbf{pa}_X . Instead, $P_X^{(\mathbf{pa}_X, x)}(x' | \mathbf{pa}_X) = 1$ if $x' = x$ and 0 otherwise. ■

Just as for the non-probabilistic case, counterfactual statements are suppositions relative to an intervention $X \leftarrow x$ that is performed in some actual world (u, v) , and as before, we formalize them by combining the actual refinement and the intervention operators. The difference compared to the logical setting is that we move from statements regarding either *all* solutions of the model $(M^{(u, v)})_{X \leftarrow x}$ (using $[X \leftarrow x]$) or *some* solutions (using $\langle X \leftarrow x \rangle$) to probabilistic statements about $(M^{(u, v)})_{X \leftarrow x}$, the joint distribution of which takes on the form:

$$P_{(M^{(u, v)})_{X \leftarrow x}}(\mathcal{U} = u^*, \mathcal{V} = v^*) = \begin{cases} \prod_{\{Y \in \mathcal{V} \setminus X\}} P_Y^{(\mathbf{pa}_Y, y)}(y^* | \mathbf{pa}_Y^*) P_U(u^*) & \text{if } x^* = x \\ 0 & \text{otherwise.} \end{cases}$$

Here $(\mathbf{pa}_Y^*, y^*, x^*) \subseteq (u^*, v^*)$ and $(\mathbf{pa}_Y, y, x) \subseteq (u, v)$. If we only consider a specific actual context $\mathcal{U} = u$ and apply Definition 10, this results in:

$$P_{(M^{(u, v)})_{X \leftarrow x}}(\mathcal{V} = v^* | \mathcal{U} = u) = \begin{cases} 0 & \text{if } x^* \neq x \\ 0 & \text{if } \emptyset \neq \{Y \in \mathcal{V} \setminus X | \mathbf{epa}_Y = \mathbf{epa}_Y^* \text{ and } y^* \neq y\} \\ 1 & \text{if } \emptyset = \{Y \in \mathcal{V} \setminus X | \mathbf{epa}_Y \neq \mathbf{epa}_Y^*\} \\ \prod_{\{Y \in \mathcal{V} \setminus X | \mathbf{epa}_Y \neq \mathbf{epa}_Y^*\}} P_Y(y^* | \mathbf{epa}_Y^*, x \mathbf{pa}_Y) & \text{otherwise.} \end{cases}$$

We write $P_M(\mathbf{y}_x|\mathbf{u}, \mathbf{v})$ as shorthand for $P_{(M(\mathbf{u}, \mathbf{v}))_{\mathbf{X} \leftarrow \mathbf{x}}}(\mathbf{Y} = \mathbf{y}|\mathbf{U} = \mathbf{u})$ for $\mathbf{Y}, \mathbf{X} \subseteq \mathcal{V}$, which is the probabilistic counterpart of evaluating a basic causal formula in our **swc** logic (Sec. 3): $(M, \mathbf{u}, \mathbf{v}) \models [\mathbf{X} \leftarrow \mathbf{x}]\mathbf{Y} = \mathbf{y}$.

We now have all the tools in hand to construct the third layer, consisting of *probabilities of counterfactuals*, which are expressions of the form $P_M(\mathbf{y}_x, \dots, \mathbf{z}_w)$. What makes them counterfactual as opposed to merely interventional is that the various interventions \mathbf{x} , \mathbf{w} may conflict with each other, in which case their respective outcomes \mathbf{y} and \mathbf{z} occur across worlds that are counterfactual to each other. (Note that these expressions include the basic case in which we intervene after having made an *actual* observation, as those can be expressed as $P_M(\mathbf{y}_x|\mathbf{z}) = P_M(\mathbf{y}_x, \mathbf{z})/P_M(\mathbf{z})$.) Our definition proceeds analogously to the standard deterministic case (Bareinboim et al., 2022).

Definition 11 *A probabilistic causal model M induces a family of joint distributions over counterfactual events of the form $\mathbf{y}_x, \dots, \mathbf{z}_w$ for any $\mathbf{Y}, \mathbf{X}, \dots, \mathbf{Z}, \mathbf{W} \subseteq \mathcal{V}$, as follows:*

$$P_M(\mathbf{y}_x, \dots, \mathbf{z}_w) = \sum_{\{\mathbf{u}, \mathbf{v}\}} P_M(\mathbf{y}_x, \dots, \mathbf{z}_w|\mathbf{u}, \mathbf{v})P_M(\mathbf{u}, \mathbf{v}) = \sum_{\{\mathbf{u}, \mathbf{v}\}} P_M(\mathbf{y}_x|\mathbf{u}, \mathbf{v}) \dots P_M(\mathbf{z}_w|\mathbf{u}, \mathbf{v})P_M(\mathbf{u}, \mathbf{v}).$$

(In the Appendix I show that this probability distribution is indeed well-defined.)■

Intuitively, this definition generalizes the idea that conditional on some actual world (\mathbf{u}, \mathbf{v}) , different counterfactual queries \mathbf{y}_x and \mathbf{z}_w are independent of each other.

Coming back to our patient example, we can now easily compute the probability of necessity and sufficiency, which Mueller and Pearl (2023) here interpret as the *probability of benefit* (the probability of harm is similar).⁶ As we will see, in our – admittedly very simplistic – example this probability is point-identifiable from just observational data. Since the lack of identifiability is one of Dawid and Senn (2023)’s main concerns with the former’s approach, this illustrates how non-deterministic models can provide room for compromise in the debate over personalized medicine.

Example 2 *Imagine again our patient from before, except that we now also consider it possible that the patient recovers even when untreated. So we get: $P_Y(y|x) = p$ and $P_Y(y|x') = q$ and $P_X(x) = r$ for some $0 < p, q, r < 1$, where $Y = y/Y = y'$ represents recovery/death and $X = x/X = x'$ represents treatment/no treatment. Computing the PNS gives: $P_M(y_x, y'_{x'}) = P_Y(y|x)P_Y(y'|x') = p(1 - q)$. (See the Appendix for the intermediate steps.)*

The following result confirms that the probabilistic setting generalizes the logical one.

Theorem 12 *Say we have a PNSEM M and a NSEM M^* over the same signature and graph which are consistent with each other, meaning that for all $X \in \mathcal{V}$ and all $x \in \mathcal{R}(X)$, $\mathbf{pa}_X \in \mathcal{R}(\mathbf{Pa}_X)$: $P_X(x|\mathbf{pa}_X) > 0$ iff $x \in f_X(\mathbf{pa}_X)$. Then for all worlds (\mathbf{u}, \mathbf{v}) and all $\mathbf{y}_x, \dots, \mathbf{z}_w$ for some $\mathbf{Y}, \mathbf{X}, \dots, \mathbf{Z}, \mathbf{W} \subseteq \mathcal{V}$:*

$$P_M(\mathbf{y}_x, \dots, \mathbf{z}_w|\mathbf{u}, \mathbf{v}) = 1 \text{ iff } (M^*, \mathbf{u}, \mathbf{v}) \models [\mathbf{X} \leftarrow \mathbf{x}]\mathbf{Y} = \mathbf{y} \wedge \dots \wedge [\mathbf{W} \leftarrow \mathbf{w}]\mathbf{Z} = \mathbf{z}.$$

Furthermore, the following results confirm that Definition 11 also generalizes the standard definition of probabilities of counterfactuals (Bareinboim et al., 2022, Defns. 5 and 7). (Note that a standard, deterministic, probabilistic causal model, i.e., a PDSCM, is simply a PNSCM for which all $P_X(X|\mathbf{Pa}_X)$ only take value in $\{0, 1\}$, meaning that P_X reduces to a function f_X .)

6. To be clear, I do not endorse their analysis of harm. Instead I have developed my own analysis of harm and benefit in (Beckers et al., 2023a, 2024). What matters here is that both of our accounts invoke probabilities of counterfactuals.

Theorem 13 *Given a PDSCM M and $Y, X, \dots, Z, W \subseteq \mathcal{V}$, it holds that*

$$P_M(\mathbf{y}_x, \dots, \mathbf{z}_w) = \sum_{\{u \mid (M, u) \models [X \leftarrow x] Y = \mathbf{y} \wedge \dots \wedge [W \leftarrow w] Z = \mathbf{z}\}} P_M(u).$$

Corollary 14 *Given a PDSCM M and $Y, X \subseteq \mathcal{V}$, it holds that $P_M(\mathbf{y}_x) = P_{M_{X \leftarrow x}}(\mathbf{y})$.*

6.1. Counterfactuals in Causal Bayesian Networks

Just as in the deterministic case, probabilities of counterfactuals cannot be evaluated directly within CBNs because they depend on knowledge of the full graph \mathcal{G} and the full distribution $P_M(\mathcal{U}, \mathcal{V})$. However, the nondeterministic case *does* open up an entirely novel perspective on procedures for partially identifying such probabilities from partial knowledge about a causal model – such as knowledge of the CBN it induces, for example – since all of the standard procedures rely on the assumption that the ground truth causal model has to be a PDSCM (Balke and Pearl, 1994; Tian and Pearl, 2000; Zhang et al., 2022). I defer a full discussion of this change in perspective to future work, and here restrict attention to just one important observation.

In the deterministic case, a conditional probability $P_X(X | \mathbf{EPa}_X)$ expressing the relation between a variable and its endogenous parents in a CBN is assumed to reduce entirely to an unknown probability $P(\mathbf{XPa}_X)$ over some unknown exogenous variables \mathbf{XPa}_X and an unknown function $f_X(\mathbf{EPa}_X, \mathbf{XPa}_X)$. So unless all the values of $P_X(X | \mathbf{EPa}_X)$ happen to belong to $\{0, 1\}$, this reductive viewpoint forces us to conclude that there *must* be some missing variables, i.e., $\mathbf{XPa}_X \neq \emptyset$. In the nondeterministic case we do not have to make this reductive assumption (although of course nothing prevents us from making it if there is good reason to). Thus one option that is always available for any CBN is to take the probabilities P_X at face value and simply assume that the CBN already contains all relevant variables, so that the ground truth PNSCM is simply identical to the CBN, i.e., to assume that $\mathcal{U} = \emptyset$. Under this option probabilities of counterfactuals (Def 11) can be computed directly from the CBN. Of course the assumption that $\mathcal{U} = \emptyset$ is a very strong one and thus it requires careful justification. In future work I intend to investigate much weaker assumptions that nonetheless offer informative results on partial identification.

Interestingly, since the writing of this article, Galhotra and Halpern (2024) have likewise offered a novel semantics for probabilities of counterfactuals that can be applied directly to CBNs, and therefore a comparison with their approach is in order as well. For now I tentatively conjecture that the semantics we develop are equivalent, possibly up to some minor technical details in formulation.

7. Conclusion

I here developed a nondeterministic generalization of causal models and offered an axiomatization for two of the resulting counterfactual logics, arguing that their semantics are superior to other proposals. Crucially, they inherit all the benefits of Pearl’s framework whilst dropping the controversial uniqueness property. As many applications require probabilistic causal models, I also initiated the probabilistic extension of this framework, but I consider it far from concluded. Among other things, I anticipate that it offers a fruitful middle ground between the two camps disputing the role of counterfactuals in personalised decision-making. Lastly, in a companion paper (Beckers, 2025) I take advantage of the additional expressivity of nondeterministic causal models to offer a novel and improved account of *actual causation*. In doing so I also develop a generalization of the notion of *counterfactual dependence* and reflect on the role of both notions for the logic of causal discovery.

Acknowledgments

I thank Fausto Barbero, Joe Halpern, Joris Mooij, and Dzelila Siljak, for helpful discussions on some aspects of this work, and I thank the reviewers for their insightful reviews. I am particularly grateful for the pushback I received from the reviewers on the probabilistic setting, because this resulted in a thorough revision of that section. This work was funded by ARO grant: W911NF-22-1-0061.

References

- Alexander Balke and Judea Pearl. Counterfactual probabilities: Computational methods, bounds and applications. In *Proceedings of the 10th Conference on Uncertainty in Artificial Intelligence (UAI 1994)*, volume 10, pages 46–54, 1994.
- Fausto Barbero. On the logic of interventionist counterfactuals under indeterministic causal laws. In *International Symposium on Foundations of Information and Knowledge Systems (FoIKS 2024)*, pages 203–221, 2024.
- Elias Bareinboim, Juan D. Correa, Duligur Ibeling, and Thomas Icard. On pearl’s hierarchy and the foundations of causal inference. In *Probabilistic and Causal Inference: The Works of Judea Pearl*, pages 507–556. Association for Computing Machinery, 2022.
- Sander Beckers. Actual causation and nondeterministic causal models. In *Proceedings of the 4th Conference on Causal Learning and Reasoning (CLear 2025)*. PMLR, 2025.
- Sander Beckers, Hana Chockler, and Joseph Y. Halpern. Quantifying harm. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI-23*, pages 363–371, 2023a.
- Sander Beckers, Joseph Y. Halpern, and Christopher Hitchcock. Causal models with constraints. In *Proceedings of the 2nd Conference on Causal Learning and Reasoning (CLear 2023)*, 2023b.
- Sander Beckers, Hana Chockler, and Joseph Y. Halpern. A causal analysis of harm. *Minds and Machines*, 34(3):34, 2024.
- Stephan Bongers, Patrick Forré, Jonas Peters, and Joris M. Mooij. Foundations of structural causal models with cycles and latent variables. *The Annals of Statistics*, 49(5):2885–2915, 2021.
- A. Philip Dawid. Causal inference without counterfactuals (with discussion). *Journal of the American Statistical Association*, 95:465–480, 2000.
- A. Philip Dawid and Monica Musio. Effects of causes and causes of effects. *Annual Review of Statistics and Its Application*, 9:261–287, 2022.
- A. Philip Dawid and Stephen Senn. Personalised decision-making without counterfactuals. *arXiv preprint*, <https://arxiv.org/abs/2301.11976>, 2023.
- Sainyam Galhotra and Joseph Y. Halpern. Intervention and conditioning in causal bayesian networks. In *Advances in Neural Information Processing Systems 37 (NeurIPS 2024)*, volume 37, pages 89019–89041, 2024.

- David Galles and Judea Pearl. An axiomatic characterization of causal counterfactuals. *Foundations of Science*, 3(1):151–182, 1998.
- Joseph Y. Halpern. Axiomatizing causal reasoning. *Journal of Artificial Intelligence Research*, 12: 317–337, 2000.
- Joseph Y. Halpern. *Actual Causality*. MIT Press, 2016.
- Joseph Y. Halpern and Spencer Peters. Reasoning about causal models with infinitely many variables. *Proceedings of the 36th AAAI Conference on Artificial Intelligence*, 36(5668-5675), 2022.
- David Lewis. *Counterfactuals*. Blackwell Publishers and Harvard University Press, 1973.
- Scott Mueller and Judea Pearl. Personalized decision making – a conceptual introduction. *Journal of Causal Inference*, 11(1), 2023.
- Judea Pearl. *Causality: Models, Reasoning, and Inference; 2nd edition*. Cambridge University Press, 2009.
- Spencer Peters and Joseph Y. Halpern. Causal modeling with infinitely many variables. *arXiv preprint*, <https://arxiv.org/abs/2112.09171>, 2021.
- James M. Robins. Addendum to “a new approach to causal inference in mortality studies with a sustained exposure period—application to control of the healthy worker survivor effect”. *Computers & Mathematics With Applications*, 14:923–945, 1987.
- Donald B. Rubin. Estimating causal effects of treatment in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66:688–701, 1974.
- Aaron L Sarvet and Mats J Stensrud. Perspective on ‘harm’ in personalized medicine. *American Journal of Epidemiology*, 2023.
- Peter Spirtes, Clark Glymour, and Richard Scheines. *Causation, Prediction, and Search*. MIT Press, 2nd edition edition, 2001.
- Jin Tian and Judea Pearl. Probabilities of causation: Bounds and identification. *Annals of Mathematics and Artificial Intelligence*, 28(1):287–313, 2000.
- Lee Walters and J. Robert G. Williams. An argument for conjunction conditionalization. *The Review of Symbolic Logic*, 6(4):573–588, 2013.
- Tomasz Wysocki. The underdeterministic framework. *British journal for the philosophy of science*, forthcoming, 2023.
- Junzhe Zhang, Jin Tian, and Elias Bareinboim. Partial counterfactual identification from observational and experimental data. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162, pages 26548–26558, 2022.

Appendix A.

A.1. Proofs of Theorems

Theorem 5 *Given a nondeterministic causal model M , we have that for all $Y \subseteq \mathcal{V}$, for all $y \in \mathcal{R}(Y)$, and for all basic formulas φ :*

- $M \models [Y \leftarrow y]\varphi$ iff $v \models \varphi$ for all solutions (u, v) of $M_{Y \leftarrow y}$.
- For all contexts u : $(M, u) \models [Y \leftarrow y]\varphi$ iff $v \models \varphi$ for all states v such that (u, v) is a solution of $M_{Y \leftarrow y}$.

Proof: We first show that the two claims are equivalent. Per definition $M \models [Y \leftarrow y]\varphi$ iff $(M, u, v) \models [Y \leftarrow y]\varphi$ for all solutions (u, v) of M , which is equivalent to the statement that for all contexts u : $(M, u, v) \models [Y \leftarrow y]\varphi$ for all states v such that (u, v) is a solution of M . In turn, the latter is per definition equivalent to the statement that for all contexts u : $(M, u) \models [Y \leftarrow y]\varphi$. By the second claim, the latter is equivalent to the statement that for all contexts u : $v \models \varphi$ for all states v such that (u, v) is a solution of $M_{Y \leftarrow y}$. Finally, this is equivalent to the right side of the first claim.

We now prove the first claim. Recall that, given a solution (u, v) of M , $(M, u, v) \models [Y \leftarrow y]\varphi$ iff $v' \models \varphi$ for all states v' such that (u, v') is a solution of $(M^{(u,v)})_{Y \leftarrow y}$. Therefore we need to show that the following claims are equivalent:

- $v' \models \varphi$ for all states v' for which there exists some solution (u, v) of M such that (u, v') is a solution of $(M^{(u,v)})_{Y \leftarrow y}$.
- $v' \models \varphi$ for all solutions (u, v') of $M_{Y \leftarrow y}$.

This equivalence follows directly from the statement that for all states v' and all contexts u , the following claims are equivalent:

- (u, v') is a solution of $(M^{(u,v)})_{Y \leftarrow y}$ for some solution (u, v) of M .
- (u, v') is a solution of $M_{Y \leftarrow y}$.

The implication from the first claim to the second is a direct consequence of the definition of an actualized refinement. Therefore we proceed with proving the reverse implication.

Assume (u, v') is a solution of $M_{Y \leftarrow y}$. We need to show that there exists some v such that (u, v') is a solution of $(M^{(u,v)})_{Y \leftarrow y}$ and (u, v) is a solution of M . We proceed by induction over the partial order given by \mathcal{G}_M , by considering the claim: given an ancestrally closed set $W \subseteq \mathcal{V}$, there exists some w such that (u, w') appears in a solution of $(M^{(u,w)})_{Y \leftarrow y}$ and the set of solutions extending (u, w) of M is non-empty. Here $w' \subseteq v'$, and $M^{(u,w)}$ instantiates the obvious generalization of Definition 4 to ancestrally closed partial settings.

Given that we only consider total models, there exists at least one solution (to both M and to $M_{Y \leftarrow y}$) for each context u , and thus the claim holds for the base case with $W = \emptyset$.

We now proceed with the inductive case. Assume that the set of solutions of M extending (u, w) is non-empty and (u, w') appears in some solution of $(M^{(u,w)})_{Y \leftarrow y}$.

Say V is the next variable in $\mathcal{V} - W$ according to the partial order of \mathcal{G}_M , and $v' \in v'$. We need to prove that there exists some $v \in \mathcal{R}(V)$ so that the set of solutions of M extending (u, w, v) is non-empty and such that (u, w', v') appears in some solution of $(M^{(u,w,v)})_{Y \leftarrow y}$.

First we consider the case where $V \in \mathbf{Y}$. Since the set of solutions extending (\mathbf{u}, \mathbf{w}) is non-empty, there exists some v such that the set of solutions extending $(\mathbf{u}, \mathbf{w}, v)$ is non-empty as well. The equation for V in $(M^{(\mathbf{u}, \mathbf{w}, v)})_{\mathbf{Y} \leftarrow \mathbf{y}}$ will be the constant equation $V = v^*$, where $v^* \in \mathbf{y}$. Furthermore, since $(\mathbf{u}, \mathbf{v}')$ is a solution of $M_{\mathbf{Y} \leftarrow \mathbf{y}}$, $v^* = v'$. It follows that $(\mathbf{u}, \mathbf{w}', v')$ appears in a solution of $(M^{(\mathbf{u}, \mathbf{w}, v)})_{\mathbf{Y} \leftarrow \mathbf{y}}$, as had to be shown.

Second we consider the case where $V \notin \mathbf{Y}$, which we separate into two sub-cases. Here $\mathbf{pa}_{\mathbf{V}'} \subseteq \mathbf{v}'$. Note that, as \mathbf{W} is ancestrally closed, $\mathbf{Pa}_{\mathbf{V}} \subseteq \mathbf{W}$.

Consider the case where $\mathbf{pa}_{\mathbf{V}} \neq \mathbf{pa}_{\mathbf{V}'}$, where $\mathbf{pa}_{\mathbf{V}} \subseteq \mathbf{w}$. As before, we can consider any v such that the set of solutions extending $(\mathbf{u}, \mathbf{w}, v)$ is non-empty. Per definition of an actualized refinement, we have that $f_V(\mathbf{pa}_{\mathbf{V}'}) = f_V^{(v, \mathbf{pa}_{\mathbf{V}})}(\mathbf{pa}_{\mathbf{V}'})$. Furthermore, since $(\mathbf{u}, \mathbf{v}')$ is a solution of $M_{\mathbf{Y} \leftarrow \mathbf{y}}$, we have that $\mathbf{v}' \in f_V(\mathbf{pa}_{\mathbf{V}'})$, and thus also $\mathbf{v}' \in f_V^{(v, \mathbf{pa}_{\mathbf{V}})}(\mathbf{pa}_{\mathbf{V}'})$. It follows that $(\mathbf{u}, \mathbf{w}', \mathbf{v}')$ appears in a solution of $(M^{(\mathbf{u}, \mathbf{w}, v)})_{\mathbf{Y} \leftarrow \mathbf{y}}$.

Lastly, consider the case where $\mathbf{pa}_{\mathbf{V}} = \mathbf{pa}_{\mathbf{V}'}$. As each $v \in f_V(\mathbf{pa}_{\mathbf{V}})$ is a solution of $V = f_V(\mathbf{Pa}_{\mathbf{X}})$, each solution of M that starts with (\mathbf{u}, \mathbf{w}) can be extended to a solution $(\mathbf{u}, \mathbf{w}, v)$ for each $v \in f_V(\mathbf{pa}_{\mathbf{V}})$. Therefore the set of solutions extending $(\mathbf{u}, \mathbf{w}, v)$ is non-empty for each $v \in f_V(\mathbf{pa}_{\mathbf{V}})$. Choosing $v = v'$, we get that there exists at least one solution that extends $(\mathbf{u}, \mathbf{w}, v)$, and $(\mathbf{u}, \mathbf{w}', \mathbf{v}')$ appears in a solution of $(M^{(\mathbf{u}, \mathbf{w}, v)})_{\mathbf{Y} \leftarrow \mathbf{y}}$. This concludes the proof. ■

Theorem 6 AX_{non}^{swc} (resp. AX_{non}^{scc}) is a sound and complete axiomatization for the language $\mathcal{L}(\mathcal{S})$ with respect to the **swc** logic (resp. the **scc** logic) over acyclic NSCMs that have signature \mathcal{S} .

Proof: We start with completeness. First we consider AX_{non}^{swc} and the **swc** logic. It suffices to show that for any formula $\varphi \in \mathcal{L}(\mathcal{S})$ that is consistent with $AX_{non}^{swc}(\mathcal{S})$, there is an acyclic NSCM M such that $(M, \mathbf{u}, \mathbf{v}) \models \varphi$ for some causal setting $(M, \mathbf{u}, \mathbf{v})$. The proof follows the same technique as used by Halpern (2000) and Halpern and Peters (2022).

Suppose that $\varphi \in \mathcal{L}(\mathcal{S})$ is consistent with $AX_{non}^{swc}(\mathcal{S})$ (i.e., we cannot prove $\neg\varphi$ in $AX_{non}^{swc}(\mathcal{S})$). Then φ can be extended to a maximal consistent set C of formulas, meaning that $\varphi \in C$, every finite subset C' of C is consistent with $AX_{non}^{swc}(\mathcal{S})$, and no strict superset C^* of C has the property that every finite subset of C^* is consistent with $AX_{non}^{swc}(\mathcal{S})$. Standard arguments show that, for every formula $\psi \in \mathcal{L}(\mathcal{S})$, either ψ or $\neg\psi$ must be in C . Moreover, every instance of the axioms in $AX_{non}^{swc}(\mathcal{S})$ must be in C .

We now define an acyclic NSCM M with signature \mathcal{S} and the requisite causal setting $(M, \mathbf{u}, \mathbf{v})$ as follows. Fix some context $\mathbf{u} \in \mathcal{R}(\mathcal{U})$. The following Lemma from Halpern and Peters (2022) is useful, where AX_{bas} is the axiom system consisting of D0, D7, D8, and MP.

Lemma 15

- (a) $AX_{bas} \vdash [\mathbf{Y} \leftarrow \mathbf{y}]\varphi_1 \wedge [\mathbf{Y} \leftarrow \mathbf{y}]\varphi_2 \Leftrightarrow [\mathbf{Y} \leftarrow \mathbf{y}](\varphi_1 \wedge \varphi_2)$
- (b) $AX_{bas} \vdash \langle \mathbf{Y} \leftarrow \mathbf{y} \rangle(\varphi_1 \vee \varphi_2) \Leftrightarrow \langle \mathbf{Y} \leftarrow \mathbf{y} \rangle\varphi_1 \vee \langle \mathbf{Y} \leftarrow \mathbf{y} \rangle\varphi_2$

Applying Lemma 15 and the axioms D1-2, D7, and D10(a) to the empty intervention, it follows by standard modal reasoning that there exists some $\mathbf{v} \in \mathcal{R}(\mathcal{V})$ so that $\langle \rangle \mathbf{V} = \mathbf{v} \in C$. We will construct M such that $(M, \mathbf{u}, \mathbf{v}) \models \varphi$.

For all $X \in \mathcal{V}$ we define $\mathbf{Pa}_{\mathbf{X}}$ as the set of variables $Y \in \mathcal{V}$ such that $Y \rightsquigarrow X \in C$. Given D6, we have hereby defined an acyclic graph \mathcal{G} . We define F_X for $x \in \mathcal{R}(X)$ and $\mathbf{pa}_{\mathbf{X}} \in \mathcal{R}(\mathbf{Pa}_{\mathbf{X}})$ by taking $x \in f_X(\mathbf{pa}_{\mathbf{X}})$ iff $\langle \mathbf{Pa}_{\mathbf{X}} \leftarrow \mathbf{pa}_{\mathbf{X}} \rangle X = x \in C$. As before, it follows from Lemma 15

and the axioms D1-2, D7, and D10(a), that for each choice \mathbf{pa}_X there will be at least one value x such that $x \in f_X(\mathbf{pa}_X)$, and thus our NSCM is total. Therefore M is an acyclic and total NSCM, as required.

Furthermore, $\langle V = v \in C$ combined with D10(c) gives $\llbracket V = v \in C$. Applying D3(b) and Lemma 15, for each X we get that $[\mathbf{Pa}_X \leftarrow \mathbf{pa}_X]X = x \in C$, where $(x, \mathbf{pa}_X) \subseteq v$. Therefore the function f_X is deterministic for the actual parent values, (i.e, $x = f_X(\mathbf{pa}_X)$), and thus taking $f_X^{(\mathbf{pa}_X, x)}(\mathbf{pa}_X) = f_X(\mathbf{pa}_X)$ results in a deterministic actualized refinement for the actual parent values, as required for an NSCM. It remains to be shown that $(M, \mathbf{u}, v) \models \varphi$.

Following exactly the same reasoning as in the proof of Theorem 5.2 in Halpern and Peters (2022) and in Beckers et al. (2023b), it follows that this reduces to showing for all formulas of the form $\langle Y \leftarrow y \rangle X = x$ with $X = \mathcal{V} - Y$ that $\langle Y \leftarrow y \rangle X = x \in C$ iff $(M, \mathbf{u}, v) \models \langle Y \leftarrow y \rangle X = x$.

Suppose that $\langle Y \leftarrow y \rangle X = x \in C$. Let $v' = (y, x)$. It suffices to show that for each $X \in \mathbf{X}$, $\langle \mathbf{Pa}_X \leftarrow \mathbf{pa}_x \rangle X = x \in C$, where $(x, \mathbf{pa}_x) \subseteq v'$. Consider some X and the requisite values \mathbf{pa}_X and x . By D3(a), we have that $\langle \mathbf{Pa}_X \leftarrow \mathbf{pa}_X, Z \leftarrow z \rangle X = x \in C$, where $Z = \mathcal{V} - (\mathbf{Pa}_X \cup \{X\})$, and $z \subseteq v'$.

Let us consider some $z'' \in \mathcal{R}(Z)$ such that $\langle \mathbf{Pa}_X \leftarrow \mathbf{pa}_X, X \leftarrow x \rangle (Z = z'') \in C$. (As before, such $z'' \in \mathcal{R}(Z)$ must exist.) If also $\langle \mathbf{Pa}_X \leftarrow \mathbf{pa}_X, Z \leftarrow z'' \rangle X = x \in C$, then by D5 and D7 we get that $\langle \mathbf{Pa}_X \leftarrow \mathbf{pa}_X \rangle X = x \in C$, as required.

Remains to consider the case where $[\mathbf{Pa}_X \leftarrow \mathbf{pa}_X, Z \leftarrow z'']X \neq x \in C$. It follows that $z \neq z''$. Per construction of \mathbf{Pa}_X , we have for each $Z \in \mathbf{Z}$ that $\neg(Z \rightsquigarrow X) \in C$. We show by induction that this results in a contradiction.

For the base case, take $A_0 = \emptyset$, and $a_0 \subseteq z$. Let $W_0 = Z - A_0$, and $w_0 \subseteq z, w''_0 \subseteq z''$. We have that $[\mathbf{Pa}_X \leftarrow \mathbf{pa}_X, W_0 \leftarrow w''_0, A_0 \leftarrow a_0]X \neq x \in C$.

The inductive case consists of considering $A_{k+1} = A_k \cup \{Z\}$ for some $Z \in W_k$. We let $W_{k+1} = Z - A_{k+1}$, $a_{k+1} \subseteq z$, and $w_{k+1} \subseteq z, w''_{k+1} \subseteq z''$. By the induction hypothesis, we have that $[\mathbf{Pa}_X \leftarrow \mathbf{pa}_X, W_k \leftarrow w''_k, A_k \leftarrow a_k]X \neq x \in C$, which can be rewritten as $[\mathbf{Pa}_X \leftarrow \mathbf{pa}_X, W_{k+1} \leftarrow w''_{k+1}, A_k \leftarrow a_k, Z \leftarrow z'']X \neq x \in C$. If $\langle \mathbf{Pa}_X \leftarrow \mathbf{pa}_X, W_{k+1} \leftarrow w''_{k+1}, A_k \leftarrow a_k, Z \leftarrow z \rangle X = x \in C$, it follows that $Z \rightsquigarrow X \in C$. Therefore, $[\mathbf{Pa}_X \leftarrow \mathbf{pa}_X, W_{k+1} \leftarrow w''_{k+1}, A_{k+1} \leftarrow a_{k+1}]X \neq x \in C$. Given that $|Z|$ is finite, for some $k \in \mathbb{N}$ this results in a contradiction.

For the other way, suppose that $(M, \mathbf{u}, v) \models \langle Y \leftarrow y \rangle X = x$. Let $v' = (y, x)$. Per construction of M , we know that for each $X \in \mathbf{X}$, $\langle \mathbf{Pa}_X \leftarrow \mathbf{pa}_x \rangle X = x \in C$, where $(x, \mathbf{pa}_x) \subseteq v'$.

For each $X \in \mathbf{X}$, we have that for any $Z \subseteq \mathcal{V} - (\mathbf{Pa}_X \cup \{X\})$: $\neg(Z_i \rightsquigarrow X) \in C$ for all $Z_i \in Z$. Therefore, for any z we have that $\langle Z \leftarrow z, \mathbf{Pa}_X \leftarrow \mathbf{pa}_x \rangle X = x \in C$. Letting $\mathbf{X} = \{X_1, \dots, X_k\}$, we have in particular that for each $i \in \{1, \dots, k\}$: $\langle Y \leftarrow y, \mathbf{X}^{-i} \leftarrow x^{-i} \rangle X_i = x_i \in C$, where $\mathbf{X}^{-i} := (X_1, \dots, X_{i-1}, X_{i+1}, \dots, X_k)$.

Taking $\langle Y \leftarrow y, \mathbf{X}^{-1} \leftarrow x^{-1} \rangle X_1 = x_1 \in C$ and $\langle Y \leftarrow y, \mathbf{X}^{-2} \leftarrow x^{-2} \rangle X_2 = x_2 \in C$, we can apply D5 to derive that $\langle Y \leftarrow y, \mathbf{X}^{-1,2} \leftarrow x^{-1,2} \rangle (X_1 = x_1 \wedge X_2 = x_2) \in C$. By the same reasoning, we get that $\langle Y \leftarrow y, \mathbf{X}^{-2,3} \leftarrow x^{-2,3} \rangle (X_2 = x_2 \wedge X_3 = x_3) \in C$. Again applying D5 to the last two statements, we get that $\langle Y \leftarrow y, \mathbf{X}^{-1,2,3} \leftarrow x^{-1,2,3} \rangle (X_1 = x_1 \wedge X_2 = x_2 \wedge X_3 = x_3) \in C$. By straightforward induction, we get that $\langle Y \leftarrow y \rangle X \leftarrow x \in C$, which is what had to be shown. This concludes the proof of completeness.

For AX_{non}^{scc} and the **scc** logic, the proof proceeds identically except for three differences. The first difference is that we have no need for some $v \in \mathcal{R}(\mathcal{V})$ so that $\langle \rangle V = v \in C$, as we only need to construct M and \mathbf{u} such that $(M, \mathbf{u}) \models \varphi$. The second (related) difference is that in this case there is no need to consider some actual values \mathbf{pa}_X for each $X \in \mathbf{X}$ and verify that $f_X(\mathbf{pa}_X)$ is deterministic, for the actualized refinement is not relevant to the semantics of \models for (M, \mathbf{u}) , as shown by Theorem 5. This explains why D10(c) is not part of AX_{non}^{scc} . The third difference is that we here need to show $\langle Y \leftarrow \mathbf{y} \rangle X = \mathbf{x} \in C$ iff $(M, \mathbf{u}) \models \langle Y \leftarrow \mathbf{y} \rangle X = \mathbf{x}$ (as opposed to having v included on the RHS). By Theorem 5, this is equivalent to showing that $\langle Y \leftarrow \mathbf{y} \rangle X = \mathbf{x} \in C$ iff $(\mathbf{u}, \mathbf{y}, \mathbf{x})$ is a solution of $M_{Y \leftarrow \mathbf{y}}$. The remainder of the proof remains identical.

Now we prove soundness. We leave it as a simple exercise to the reader that D0-D1-D2-D4-D7-D8-D10(a) are sound for both of our logics, and that D10(c) is sound for our **swc** logic.

D6:

As we explain in the discussion of Wysocki (2023)’s work in Section 5, contrary to our NSCMs, DSCMs do not come with a graph. Rather, a graph \mathcal{G}_D is induced by invoking “ Y depends on X ”: there is an edge from X to Y iff there exist settings $\mathbf{z} \in (\mathcal{U} \cup \mathcal{V} - \{X, Y\})$, and $x, x' \in \mathcal{R}(X)$, such that $f_Y(\mathbf{z}, x) \neq f_Y(\mathbf{z}, x')$. As mentioned in Halpern and Peters (2022), D6 expresses the acyclicity of the induced graph \mathcal{G}_D . Of course NSCMs can be used to invoke a graph \mathcal{G}_D in exactly the same manner: there is an edge from X to Y iff there exist settings $\mathbf{z} \in (\mathcal{U} \cup \mathcal{V} - \{X, Y\})$, and $x, x' \in \mathcal{R}(X)$, such that $f_Y(\mathbf{z}, x) \neq f_Y(\mathbf{z}, x')$. NSCMs also come with an explicit graph \mathcal{G}_M , and this need not be identical to \mathcal{G}_D . But the graph \mathcal{G}_D is easily seen to be a subgraph of \mathcal{G}_M , since X has to be an argument of f_Y for Y to depend on X , and per definition this means that there is an edge from X to Y in \mathcal{G}_M . Therefore the acyclicity of \mathcal{G}_M implies the acyclicity of \mathcal{G}_D , and thus D6 is sound for both of our logics. Concretely: if $X \rightsquigarrow Y$ for some solution (\mathbf{u}, \mathbf{v}) , then X has to be an ancestor of Y in \mathcal{G}_D . So the falsity of D6 would imply that \mathcal{G}_D is cyclic.

D3(a):

We start with soundness for \models relative to $(M, \mathbf{u}, \mathbf{v})$. Assume $W \notin \mathbf{X}$, and (\mathbf{u}, \mathbf{v}) is a solution of M , and $(M, \mathbf{u}, \mathbf{v}) \models \langle \mathbf{X} \leftarrow \mathbf{x} \rangle (W = w \wedge \varphi)$. This means that there exists \mathbf{v}' so that $(\mathbf{u}, \mathbf{v}')$ is a solution of $(M^{(\mathbf{u}, \mathbf{v})})_{\mathbf{X} \leftarrow \mathbf{x}}$ and $\mathbf{v}' \models (W = w \wedge \varphi)$. By D4, it directly follows that $(\mathbf{u}, \mathbf{v}')$ is also a solution of $(M^{(\mathbf{u}, \mathbf{v})})_{\mathbf{X} \leftarrow \mathbf{x}, W \leftarrow w}$, and we also have that $\mathbf{v}' \models \varphi$. This means precisely that $(M, \mathbf{u}, \mathbf{v}) \models \langle \mathbf{X} \leftarrow \mathbf{x}, W \leftarrow w \rangle \varphi$.

Soundness for \models relative to (M, \mathbf{u}) is a consequence of the soundness for \models relative to $(M, \mathbf{u}, \mathbf{v})$. For concreteness, we here write out the intermediate steps. Assume that $(M, \mathbf{u}) \models \langle \mathbf{X} \leftarrow \mathbf{x} \rangle (W = w \wedge \varphi)$. Per definition of $\langle \rangle$, this is equivalent to: $(M, \mathbf{u}) \models \neg[\mathbf{X} \leftarrow \mathbf{x}](W \neq w \vee \neg\varphi)$. In turn, this is equivalent to it not being the case that for all states \mathbf{v} so that (\mathbf{u}, \mathbf{v}) is a solution of M , we have $(M, \mathbf{u}, \mathbf{v}) \models [\mathbf{X} \leftarrow \mathbf{x}](W \neq w \vee \neg\varphi)$. This is equivalent to there existing some state \mathbf{v}' such that $(\mathbf{u}, \mathbf{v}')$ is a solution of M and $(M, \mathbf{u}, \mathbf{v}') \models \langle \mathbf{X} \leftarrow \mathbf{x} \rangle (W = w \wedge \varphi)$. By D3(a) for \models relative to causal settings, we get that for some solution $(\mathbf{u}, \mathbf{v}'')$ of M , namely $(\mathbf{u}, \mathbf{v}')$, $(M, \mathbf{u}, \mathbf{v}'') \models \langle \mathbf{X} \leftarrow \mathbf{x}, W \leftarrow w \rangle \varphi$. Applying all of the above equivalences in the other direction, this is seen to be equivalent to $(M, \mathbf{u}) \models \langle \mathbf{X} \leftarrow \mathbf{x}, W \leftarrow w \rangle \varphi$, which is what had to be shown.

D3(b):

Soundness for \models relative to (M, \mathbf{u}) is a direct consequence of the soundness for \models relative to $(M, \mathbf{u}, \mathbf{v})$, so we proceed with the latter.

Assume $W \notin \mathbf{X}$, and (\mathbf{u}, \mathbf{v}) is a solution of M , and $(M, \mathbf{u}, \mathbf{v}) \models [\mathbf{X} \leftarrow \mathbf{x}](W = w \wedge \varphi)$. It suffices to show that for any state \mathbf{v}' , if $(\mathbf{u}, \mathbf{v}')$ is a solution of $(M^{(\mathbf{u}, \mathbf{v})})_{\mathbf{X} \leftarrow \mathbf{x}, W \leftarrow w}$, then it is a solution of $(M^{(\mathbf{u}, \mathbf{v})})_{\mathbf{X} \leftarrow \mathbf{x}}$.

We proceed by a reductio. Assume that $(\mathbf{u}, \mathbf{v}')$ is a solution only of the former. Note that the two models have identical equations for all variables except W , and that it has to be that $w \in \mathbf{v}'$. Therefore it must be that (w, \mathbf{pa}_W) is not a solution of W 's equation in $(M^{(\mathbf{u}, \mathbf{v})})_{\mathbf{X} \leftarrow \mathbf{x}}$, where $\mathbf{pa}_W \subseteq \mathbf{v}'$. Given that, per the first assumption, $W = w$ for all solutions of $(M^{(\mathbf{u}, \mathbf{v})})_{\mathbf{X} \leftarrow \mathbf{x}}$, it must be that \mathbf{pa}_W does not appear in any such solution. However, given that both models are acyclic and have identical equations for all non-descendants of W (including \mathbf{pa}_W), they have exactly the same partial solutions (\mathbf{u}, \mathbf{a}) , where $\mathbf{A} \subseteq \mathcal{V}$ consists of all non-descendants of W . Thus it cannot be that \mathbf{pa}_W only appears in a solution to one of them.

D5:

We start with soundness for \models relative to $(M, \mathbf{u}, \mathbf{v})$. Assume (\mathbf{u}, \mathbf{v}) is a solution of M , and $(M, \mathbf{u}, \mathbf{v}) \models \langle \mathbf{X} \leftarrow \mathbf{x}, Y \leftarrow y \rangle (W = w \wedge \mathbf{Z} = \mathbf{z}) \wedge \langle \mathbf{X} \leftarrow \mathbf{x}, W \leftarrow w \rangle (Y = y \wedge \mathbf{Z} = \mathbf{z})$, where $\mathbf{Z} = \mathcal{V} - (\mathbf{X} \cup \{W, Y\})$. We need to show that $(M, \mathbf{u}, \mathbf{v}) \models \langle \mathbf{X} \leftarrow \mathbf{x} \rangle (W = w \wedge Y = y \wedge \mathbf{Z} = \mathbf{z})$.

Per assumption, there exists a solution $(\mathbf{u}, \mathbf{v}_1)$ of $(M^{(\mathbf{u}, \mathbf{v})})_{\mathbf{X} \leftarrow \mathbf{x}, Y \leftarrow y}$ such that its restriction to (W, \mathbf{Z}) is (w, \mathbf{z}) , and there exists a solution $(\mathbf{u}, \mathbf{v}_2)$ of $(M^{(\mathbf{u}, \mathbf{v})})_{\mathbf{X} \leftarrow \mathbf{x}, W \leftarrow w}$ such that its restriction to (Y, \mathbf{Z}) is (y, \mathbf{z}) . Clearly also both solutions have that $\mathbf{X} = \mathbf{x}$, and it must be that $y \in \mathbf{v}_1$, and also $w \in \mathbf{v}_2$.

Since $\mathcal{V} = \mathbf{Z} \cup \mathbf{X} \cup \{W, Y\}$, the tuple $\mathbf{v}_3 = (\mathbf{x}, \mathbf{z}, y, w)$ is a state. Therefore it follows that $\mathbf{v}_1 = \mathbf{v}_2 = \mathbf{v}_3$. It now suffices to show that $(\mathbf{u}, \mathbf{v}_3)$ is a solution of $(M^{(\mathbf{u}, \mathbf{v})})_{\mathbf{X} \leftarrow \mathbf{x}}$. As the equations for all variables in $\mathbf{Z} \cup \mathbf{X}$ are identical across the three models, $(\mathbf{u}, \mathbf{v}_3)$ contains solutions to all these equations in $(M^{(\mathbf{u}, \mathbf{v})})_{\mathbf{X} \leftarrow \mathbf{x}}$. As the equation for Y is identical across $(M^{(\mathbf{u}, \mathbf{v})})_{\mathbf{X} \leftarrow \mathbf{x}}$ and $(M^{(\mathbf{u}, \mathbf{v})})_{\mathbf{X} \leftarrow \mathbf{x}, W \leftarrow w}$, $(\mathbf{u}, \mathbf{v}_3)$ contains a solution for Y in $(M^{(\mathbf{u}, \mathbf{v})})_{\mathbf{X} \leftarrow \mathbf{x}}$. The same holds for W and the other model, concluding the proof.

Now we prove soundness for \models relative to (M, \mathbf{u}) . Assume that $(M, \mathbf{u}) \models \langle \mathbf{X} \leftarrow \mathbf{x}, Y \leftarrow y \rangle (W = w \wedge \mathbf{Z} = \mathbf{z}) \wedge \langle \mathbf{X} \leftarrow \mathbf{x}, W \leftarrow w \rangle (Y = y \wedge \mathbf{Z} = \mathbf{z})$, where $\mathbf{Z} = \mathcal{V} - (\mathbf{X} \cup \{W, Y\})$. We need to show that $(M, \mathbf{u}) \models \langle \mathbf{X} \leftarrow \mathbf{x} \rangle (W = w \wedge Y = y \wedge \mathbf{Z} = \mathbf{z})$. By Theorem 5, our assumption is equivalent to the statement that $(\mathbf{u}, \mathbf{x}, \mathbf{z}, y, w)$ is a solution of both $M_{\mathbf{X} \leftarrow \mathbf{x}, Y \leftarrow y}$ and $M_{\mathbf{X} \leftarrow \mathbf{x}, W \leftarrow w}$, and what we need to show is that $(\mathbf{u}, \mathbf{x}, \mathbf{z}, y, w)$ is a solution of $M_{\mathbf{X} \leftarrow \mathbf{x}}$. This follows by applying the reasoning from the previous paragraph to the three models at hand. ■

Proposition 7 *If $(M, \mathbf{u}, \mathbf{v}) \models \mathbf{X} = \mathbf{x} \wedge \varphi$ then $(M, \mathbf{u}, \mathbf{v}) \models [\mathbf{X} \leftarrow \mathbf{x}] \varphi$.*

Proof: We show that for any φ' : $(M, \mathbf{u}, \mathbf{v}) \models \varphi'$ iff $(M, \mathbf{u}, \mathbf{v}) \models \Box \varphi'$. The result then follows by repeated application of D3(b), starting with the antecedent $\Box(\mathbf{X} = \mathbf{x} \wedge \varphi)$.

$(M, \mathbf{u}, \mathbf{v}) \models \varphi'$ simply means that $\mathbf{v} \models \varphi'$. So it suffices to show that (\mathbf{u}, \mathbf{v}) is the only solution of $M^{(\mathbf{u}, \mathbf{v})}$ that extends the context \mathbf{u} . We proceed by induction by extending \mathbf{u} one variable at the time, where we respect the partial order determined by \mathcal{G} . For the base case, given that exogenous variables have no equations, we trivially have that \mathbf{u} is part of any solution extending \mathbf{u} . For the inductive case, assume that we have solved the equations for some set of variables $\mathbf{W} \subseteq \mathcal{V}$ and using $\mathcal{U} = \mathbf{u}$, and that this resulted in a unique solution (\mathbf{u}, \mathbf{w}) such that \mathbf{w} agrees with \mathbf{v} on \mathbf{W} . Consider some $X \in \mathcal{V}$ such that $\mathbf{Pa}_X \subseteq \mathbf{W} \cup \mathcal{U}$ (which has to exist giving acyclicity), and let \mathbf{pa}_X be the restriction of (\mathbf{u}, \mathbf{w}) to \mathbf{Pa}_X . Let $x \in \mathbf{v}$. Per definition of an actualized refinement, $f_X^{(\mathbf{pa}_X, x)}(\mathbf{pa}_X) = x$, and thus $(\mathbf{u}, \mathbf{w}, x)$ is the unique solution to the equations for the variables $\mathbf{W} \cup \{X\}$, which is what had to be shown. ■

Theorem 12 Say we have a PNSEM M and a NSEM M^* over the same signature and graph which are consistent with each other, meaning that for all $X \in \mathcal{V}$ and all $x \in \mathcal{R}(X)$, $\mathbf{pa}_X \in \mathcal{R}(\mathbf{Pa}_X)$: $P_X(x|\mathbf{pa}_X) > 0$ iff $x \in f_X(\mathbf{pa}_X)$. Then for all worlds (\mathbf{u}, \mathbf{v}) and all $\mathbf{y}_x, \dots, \mathbf{z}_w$ for some $\mathbf{Y}, \mathbf{X}, \dots, \mathbf{Z}, \mathbf{W} \subseteq \mathcal{V}$:

$$P_M(\mathbf{y}_x, \dots, \mathbf{z}_w | \mathbf{u}, \mathbf{v}) = 1 \text{ iff } (M^*, \mathbf{u}, \mathbf{v}) \models [\mathbf{X} \leftarrow \mathbf{x}] \mathbf{Y} = \mathbf{y} \wedge \dots \wedge [\mathbf{W} \leftarrow \mathbf{w}] \mathbf{Z} = \mathbf{z}.$$

Proof: Note that we have (1):

$$P_M(\mathbf{y}_x, \dots, \mathbf{z}_w | \mathbf{u}, \mathbf{v}) = 1 \text{ iff } P_M(\mathbf{y}_x | \mathbf{u}, \mathbf{v}) \dots P_M(\mathbf{z}_w | \mathbf{u}, \mathbf{v}) = 1$$

$$\text{iff } P_M(\mathbf{y}_x | \mathbf{u}, \mathbf{v}) = 1 \text{ and } \dots \text{ and } P_M(\mathbf{z}_w | \mathbf{u}, \mathbf{v}) = 1.$$

And also (2):

$$(M^*, \mathbf{u}, \mathbf{v}) \models [\mathbf{X} \leftarrow \mathbf{x}] \mathbf{Y} = \mathbf{y} \wedge \dots \wedge [\mathbf{W} \leftarrow \mathbf{w}] \mathbf{Z} = \mathbf{z}$$

$$\text{iff } (M^*, \mathbf{u}, \mathbf{v}) \models [\mathbf{X} \leftarrow \mathbf{x}] \mathbf{Y} = \mathbf{y} \text{ and } \dots \text{ and } (M^*, \mathbf{u}, \mathbf{v}) \models [\mathbf{W} \leftarrow \mathbf{w}] \mathbf{Z} = \mathbf{z}.$$

Therefore it suffices to show that $P_M(\mathbf{y}_x | \mathbf{u}, \mathbf{v}) = 1$ iff $(M^*, \mathbf{u}, \mathbf{v}) \models [\mathbf{X} \leftarrow \mathbf{x}] \mathbf{Y} = \mathbf{y}$, which is a direct consequence of Lemma 16 ■

Lemma 16 Given a PNSEM M and a NSEM M^* that are consistent with each other, for all pair of worlds (\mathbf{u}, \mathbf{v}) , $(\mathbf{u}, \mathbf{v}^*)$, and all \mathbf{y}_x for some $\mathbf{Y}, \mathbf{X} \subseteq \mathcal{V}$:

$$(\mathbf{u}, \mathbf{v}^*) \text{ is a solution of } (M^{(\mathbf{u}, \mathbf{v})})_{\mathbf{X} \leftarrow \mathbf{x}} \text{ iff } (\mathbf{u}, \mathbf{v}^*) \text{ is a solution of } (M^{*(\mathbf{u}, \mathbf{v})})_{\mathbf{X} \leftarrow \mathbf{x}}.$$

Proof: $(\mathbf{u}, \mathbf{v}^*)$ is a solution of $(M^{(\mathbf{u}, \mathbf{v})})_{\mathbf{X} \leftarrow \mathbf{x}}$ iff

- $\mathbf{x} \subseteq \mathbf{v}^*$ and
- for all $Y \in \mathcal{V} \setminus \mathbf{X}$: $y^* \in f_Y^{(\mathbf{pa}_Y, \mathbf{y})}(\mathbf{pa}_Y^*)$, where $(\mathbf{y}, \mathbf{pa}_Y) \subseteq (\mathbf{u}, \mathbf{v})$ and $(\mathbf{y}^*, \mathbf{pa}_Y^*) \subseteq (\mathbf{u}, \mathbf{v}^*)$.

iff

- $\mathbf{x} \subseteq \mathbf{v}^*$ and
- for all $Y \in \mathcal{V} \setminus \mathbf{X}$ such that $\mathbf{pa}_Y = \mathbf{pa}_Y^*$: $y^* = y$
- for all $Y \in \mathcal{V} \setminus \mathbf{X}$ such that $\mathbf{pa}_Y \neq \mathbf{pa}_Y^*$: $y^* \in f_Y(\mathbf{pa}_Y^*)$

iff

- $\mathbf{x} \subseteq \mathbf{v}^*$ and
- for all $Y \in \mathcal{V} \setminus \mathbf{X}$ such that $\mathbf{pa}_Y = \mathbf{pa}_Y^*$: $y^* = y$
- for all $Y \in \mathcal{V} \setminus \mathbf{X}$ such that $\mathbf{pa}_Y \neq \mathbf{pa}_Y^*$: $y^* \in f_Y(\mathbf{pa}_Y^*)$

iff

- $\mathbf{x} \subseteq \mathbf{v}^*$ and
- for all $Y \in \mathcal{V} \setminus \mathbf{X}$ such that $\mathbf{epa}_Y = \mathbf{epa}_Y^*$: $y^* = y$

- for all $Y \in \mathcal{V} \setminus \mathbf{X}$ such that $\text{epa}_Y \neq \text{epa}_Y^*$: $P_Y(y^* | \text{epa}_Y^*, \text{xpa}_Y) > 0$

iff

$$P_{(M^{(u,v)})_{\mathbf{X} \leftarrow \mathbf{x}}}(\mathbf{u}, \mathbf{v}^*) > 0$$

iff

$(\mathbf{u}, \mathbf{v}^*)$ is a solution of $(M^{(u,v)})_{\mathbf{X} \leftarrow \mathbf{x}}$.

■

Theorem 13 *Given a PDSCM M and $\mathbf{Y}, \mathbf{X}, \dots, \mathbf{Z}, \mathbf{W} \subseteq \mathcal{V}$, it holds that*

$$P_M(\mathbf{y}_x, \dots, \mathbf{z}_w) = \sum_{\{\mathbf{u} | (M, \mathbf{u}) \models [\mathbf{X} \leftarrow \mathbf{x}] \mathbf{Y} = \mathbf{y} \wedge \dots \wedge [\mathbf{W} \leftarrow \mathbf{w}] \mathbf{Z} = \mathbf{z}\}} P_M(\mathbf{u}).$$

Proof: From the fact that all $P_X(X | \mathbf{Pa}_X)$ only take value in $\{0, 1\}$ it follows directly that also $P_M(\mathbf{y}_x, \dots, \mathbf{z}_w | \mathbf{u}, \mathbf{v}) \in \{0, 1\}$ for any $\mathbf{y}_x, \dots, \mathbf{z}_w$.

Furthermore, note that a PDSCM M can be viewed equivalently as the combination of a DSCM that is consistent with it together with $P_M(\mathcal{U})$. Applying Theorem 12, this gives for any (\mathbf{u}, \mathbf{v}) that:

$$P_M(\mathbf{y}_x, \dots, \mathbf{z}_w | \mathbf{u}, \mathbf{v}) = 1 \text{ iff } (M, \mathbf{u}, \mathbf{v}) \models [\mathbf{X} \leftarrow \mathbf{x}] \mathbf{Y} = \mathbf{y} \wedge \dots \wedge [\mathbf{W} \leftarrow \mathbf{w}] \mathbf{Z} = \mathbf{z}.$$

Also, from the discussion in Section 4 we know that when our semantics is applied to a DSCM it reduces to the standard one, so that we get

$(M, \mathbf{u}, \mathbf{v}) \models [\mathbf{X} \leftarrow \mathbf{x}] \mathbf{Y} = \mathbf{y} \wedge \dots \wedge [\mathbf{W} \leftarrow \mathbf{w}] \mathbf{Z} = \mathbf{z}$ iff (\mathbf{u}, \mathbf{v}) is a solution of M and $(M, \mathbf{u}) \models [\mathbf{X} \leftarrow \mathbf{x}] \mathbf{Y} = \mathbf{y} \wedge \dots \wedge [\mathbf{W} \leftarrow \mathbf{w}] \mathbf{Z} = \mathbf{z}$ iff $(M, \mathbf{u}) \models \mathcal{V} = \mathbf{v}$ and $(M, \mathbf{u}) \models [\mathbf{X} \leftarrow \mathbf{x}] \mathbf{Y} = \mathbf{y} \wedge \dots \wedge [\mathbf{W} \leftarrow \mathbf{w}] \mathbf{Z} = \mathbf{z}$.

Combining all three observations, we get that:

$$\begin{aligned} P_M(\mathbf{y}_x, \dots, \mathbf{z}_w) &= \sum_{\{\mathbf{u}, \mathbf{v}\}} P_M(\mathbf{y}_x, \dots, \mathbf{z}_w | \mathbf{u}, \mathbf{v}) P_M(\mathbf{u}, \mathbf{v}) \\ &= \sum_{\{\mathbf{u}, \mathbf{v} | (M, \mathbf{u}, \mathbf{v}) \models [\mathbf{X} \leftarrow \mathbf{x}] \mathbf{Y} = \mathbf{y} \wedge \dots \wedge [\mathbf{W} \leftarrow \mathbf{w}] \mathbf{Z} = \mathbf{z}\}} P_M(\mathbf{u}, \mathbf{v}) \\ &= \sum_{\{\mathbf{u}, \mathbf{v} | (M, \mathbf{u}) \models \mathcal{V} = \mathbf{v} \text{ and } (M, \mathbf{u}) \models [\mathbf{X} \leftarrow \mathbf{x}] \mathbf{Y} = \mathbf{y} \wedge \dots \wedge [\mathbf{W} \leftarrow \mathbf{w}] \mathbf{Z} = \mathbf{z}\}} P_M(\mathbf{u}) 1_{\{(M, \mathbf{u}) \models \mathcal{V} = \mathbf{v}\}} \\ &= \sum_{\{\mathbf{u} | (M, \mathbf{u}) \models [\mathbf{X} \leftarrow \mathbf{x}] \mathbf{Y} = \mathbf{y} \wedge \dots \wedge [\mathbf{W} \leftarrow \mathbf{w}] \mathbf{Z} = \mathbf{z}\}} P_M(\mathbf{u}). \blacksquare \end{aligned}$$

A.2. Details on the Markov condition

Letting EPa_X and XPa_X respectively denote X 's endogenous and exogenous parents, and letting $\mathcal{V} = \{A, B, \dots, K\}$, we get that $P_M(\mathbf{v}) = \sum_{\mathbf{u} \in \mathcal{R}(\mathcal{U})} P_M(\mathbf{u}, \mathbf{v})$
 $= \sum_{\{(\text{xpa}_A, \text{xpa}_B, \dots, \text{xpa}_K) \in \mathcal{R}(\text{XPa}_A \times \text{XPa}_B \times \dots \times \text{XPa}_K)\}} \prod_{X \in \mathcal{V}} P_X(x | \text{epa}_X, \text{xpa}_X) P_{\mathcal{U}}(\text{xpa}_X)$
 $= \prod_{X \in \mathcal{V}} (\sum_{\{\text{xpa}_X \in \mathcal{R}(\text{XPa}_X)\}} P_X(x | \text{epa}_X, \text{xpa}_X) P_{\mathcal{U}}(\text{xpa}_X))$, where each $(x, \text{epa}_X) \subseteq \mathbf{v}$, and $(\text{xpa}_A, \text{xpa}_B, \dots, \text{xpa}_K) = \mathbf{u}$. (The latter means we are assuming that each $U \in \mathcal{U}$ is a parent of some $X \in \mathcal{V}$, but this is obviously without loss of generality.)

A.3. Details on the well-definedness of Definition 11

Lemma 17 *Given a PNSEM M and $\mathbf{Y}, \mathbf{X}, \dots, \mathbf{Z}, \mathbf{W} \subseteq \mathcal{V}$, then for all values $\mathbf{y}, \mathbf{x}, \dots, \mathbf{z}, \mathbf{w} \in \mathcal{R}(\mathbf{Y}), \mathcal{R}(\mathbf{X}), \dots, \mathcal{R}(\mathbf{Z}), \mathcal{R}(\mathbf{W})$, $P_M(\mathbf{y}_x, \dots, \mathbf{z}_w | \mathbf{u}, \mathbf{v})$ is a well-defined probability distribution.*

Proof: First we show that each factor of the form $P_M(\mathbf{y}_x | \mathbf{u}, \mathbf{v})$ is a well-defined probability distribution. Let $\mathbf{Z} = \mathcal{V} - \mathbf{Y}$. Then $P_M(\mathbf{y}_x | \mathbf{u}, \mathbf{v}) = \sum_{\{\mathbf{z} \in \mathcal{R}(\mathbf{Z})\}} P_{(M^{(u,v)})_{\mathbf{X} \leftarrow \mathbf{x}}}(\mathbf{Y} = \mathbf{y}, \mathbf{Z} = \mathbf{z} | \mathcal{U} = \mathbf{u})$.

Therefore it suffices to show that $P_{(M(u,v))_{X \leftarrow x}}(\mathcal{V} = \mathbf{v}^* | \mathcal{U} = \mathbf{u})$ is a well-defined probability distribution. This follows directly from the fact that it is the joint distribution of a PNSEM, and those are well-defined per assumption of satisfying the Markov condition. Concretely, we have that:

$$P_{(M(u,v))_{X \leftarrow x}}(\mathcal{V} = \mathbf{v}^* | \mathcal{U} = \mathbf{u}) = \prod_{\{Y \in \mathcal{V} \setminus \mathbf{X}\}} P_Y^{(epa_Y, xpa_Y, y)}(y^* | epa_Y^*, xpa_Y) \prod_{\{X \in \mathbf{X}\}} P_X^{alt}(x^*).$$

Here $P_X^{alt}(X)$ is the probability distribution that returns 1 if $X = x$ for the unique value x such that $x \in \mathbf{v}$ and 0 otherwise. Furthermore, note that for each $Y \in \mathcal{V} \setminus \mathbf{X}$, $P_Y^{(epa_Y, xpa_Y, y)}(Y | epa_Y, Xpa_Y)$ defines a family of conditional probability distributions over $\mathcal{R}(Y)$, as required.

Second, note that the product of the probability distributions of independent variables is also a probability distribution. So by *defining* $P_M(\mathbf{y}_x, \dots, \mathbf{z}_w | \mathbf{u}, \mathbf{v})$ as $P_M(\mathbf{y}_x | \mathbf{u}, \mathbf{v}) \dots P_M(\mathbf{z}_w | \mathbf{u}, \mathbf{v})$, we are *choosing* to define the counterfactual variables $\mathbf{y}_x, \dots, \mathbf{z}_w$ as being mutually independent conditional on a world (\mathbf{u}, \mathbf{v}) . (As the joint distribution of these variables had no anterior meaning in the context of PNSCMs, from a purely mathematical point of view we are free to choose to define it in whatever way we want. That the definition is also a good one is of course what I try to argue in the paper.)

Third, we show that $P_M(\mathbf{y}_x, \dots, \mathbf{z}_w, \mathbf{u}, \mathbf{v}) = P_M(\mathbf{y}_x, \dots, \mathbf{z}_w | \mathbf{u}, \mathbf{v}) P_M(\mathbf{u}, \mathbf{v})$, and thus the definition of $P_M(\mathbf{y}_x, \dots, \mathbf{z}_w)$ is simply an application of the law of total probability.

$$\begin{aligned} & P_M(\mathbf{y}_x, \dots, \mathbf{z}_w, \mathbf{u}, \mathbf{v}) \\ &= \sum_{\{\mathbf{u}', \mathbf{v}'\}} P_M(\mathbf{y}_x | \mathbf{u}', \mathbf{v}') \dots P_M(\mathbf{z}_w | \mathbf{u}', \mathbf{v}') P_M(\mathbf{u}, \mathbf{v} | \mathbf{u}', \mathbf{v}') P_M(\mathbf{u}', \mathbf{v}') \\ &= P_M(\mathbf{y}_x | \mathbf{u}, \mathbf{v}) \dots P_M(\mathbf{z}_w | \mathbf{u}, \mathbf{v}) P_M(\mathbf{u}, \mathbf{v}) \\ &= P_M(\mathbf{y}_x, \dots, \mathbf{z}_w | \mathbf{u}, \mathbf{v}) P_M(\mathbf{u}, \mathbf{v}) \blacksquare \end{aligned}$$

A.4. Details on Example 2

Example 2 *Imagine again our patient from before, except that we now also consider it possible that the patient recovers even when untreated. So we get: $P_Y(y|x) = p$ and $P_Y(y|x') = q$ and $P_X(x) = r$ for some $0 < p, q, r < 1$, where $Y = y/Y = y'$ represents recovery/death and $X = x/X = x'$ represents treatment/no treatment. Computing the PNS gives: $P_M(y_x, y'_{x'}) = \sum_{\{x'', y''\}} P_M(y_x, y'_{x'} | x'', y'') P_M(x'', y'') = \sum_{\{x'', y''\}} P_M(y_x | x'', y'') P_M(y'_{x'} | x'', y'') P_M(x'', y'')$
 $= P_M(y_x | x, y) P_M(y'_{x'} | x, y) P_M(x, y) + P_M(y_x | x, y') P_M(y'_{x'} | x, y') P_M(x, y')$
 $+ P_M(y_x | x', y) P_M(y'_{x'} | x', y) P_M(x', y) + P_M(y_x | x', y') P_M(y'_{x'} | x', y') P_M(x', y')$
 $= 1 P_M(y'_{x'}) P_M(x, y) + 0 P_M(y'_{x'}) P_M(x, y') + P_M(y_x) 0 P_M(x', y) + P_M(y_x) 1 P_M(x', y') =$
 $P_M(y'_{x'}) P_M(x, y) + P_M(y_x) P_M(x', y') = (1-q)rp + p(1-r)(1-q) = p(1-q) = P_Y(y|x) P_Y(y'|x')$*