



HAL
open science

A clarification on the links between potential outcomes and do-interventions

Lucas de Lara

► **To cite this version:**

Lucas de Lara. A clarification on the links between potential outcomes and do-interventions. 2024.
hal-04203003v4

HAL Id: hal-04203003

<https://hal.science/hal-04203003v4>

Preprint submitted on 15 Dec 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A clarification on the links between potential outcomes and do-interventions

Lucas De Lara*

Institut Elie Cartan de Lorraine, Université de Lorraine

Abstract

Most of the scientific literature on causal modeling considers the structural framework of Pearl and the potential-outcome framework of Rubin to be formally equivalent, and therefore interchangeably uses do-interventions and the potential-outcome framework to define counterfactual outcomes. In this paper, we agnostically superimpose a structural causal model and a Rubin causal model compatible with the same observations to specify under which mathematical conditions counterfactual outcomes obtained via do-interventions and potential outcomes need to, do not need to, can, or cannot be equal (almost surely or in law). Our comparison builds upon the fact that such causal models do not have to produce the same counterfactuals outcomes, and highlights real-world problems where they generally cannot correspond under classical causal-inference assumptions. Then, we examine common claims and practices from the causality literature in the light of this comparison. In doing so, we aim at clarifying the links between the two causal frameworks, and the interpretation of their respective counterfactuals.

Keywords: structural causal models, Rubin causal models, equivalence of models, counterfactuals

MSC: 62A09, 62D20

1 Introduction

Understanding causation between phenomena rather than mere association is a fundamental scientific challenge. Over the last decades, two mathematical frameworks using a terminology based on random variables have become the gold standards to address this problem.

On the one hand, the *structural account* of Pearl [2009] rests on the knowledge of a *structural causal model* (SCM) which specifies all cause-effect equations between observed random variables (often depicted by a graph). The interest of such equations comes from the possibility of carrying out *do-interventions*: forcing a variable to take a given value while keeping the rest of the mechanism untouched. More concretely, let T and Y be observational variables of the model such that we would like to understand the downstream effect of T onto Y . Replacing the formula generating T by $T = t$ for a given possible value $t \in \mathcal{T}$ and propagating this change through the other equations defines the altered variable $Y_{T=t}$, representing Y had T been equal to t .

On the other hand, the *potential-outcome account* of Rubin [1974] mathematically formalizes causal inference in clinical trials. Letting T denote a *treatment status* (e.g., taking a drug or not) and Y an *outcome* of interest (e.g., recovering or not), a so-called *Rubin causal model* (RCM) postulates the existence of *potential outcomes* $(Y_t)_{t \in \mathcal{T}}$ representing what the outcome would be *were T equal to t* for any $t \in \mathcal{T}$. The *fundamental problem of causal inference* [Holland, 1986] refers to the fact that in practice we cannot observe simultaneously all the potential outcomes, rendering unidentifiable the causal effect of T onto Y . Nevertheless, causal inference can still be achieved thanks to a mix of untestable assumptions and statistical tools: adjusting on a set of available covariates X containing all possible *confounders* between the treatment and the potential outcomes permits to identify the law of counterfactual outcomes.

Each of these causal theories enables one to carry out *counterfactual reasoning*, that is answering contrary-to-fact questions such as “Had they taken the drug, would have they recovered?”: by applying

*E-mail: lucas.de-lara@univ-lorraine.fr

do-interventions on an SCM, one can compute the outcome $Y_{T=t}$ for every possible treatment status $t \in \mathcal{T}$; using an RCM with appropriate hypotheses, one can infer the law of the every potential outcome Y_t . Both approaches involve variables describing counterfactual outcomes, more precisely outcomes *had the variable T taken a certain value*. This naturally raises the question: are these outcome variables equal (almost surely or in law) across frameworks? A plethora of scientific books and survey papers interchangeably use Pearl’s do notation and the potential-outcome subscript notation to write outcomes after interventions, suggesting that the corresponding definitions of counterfactuals are identical and differ only from their perspectives [Imbens, 2020, Neal, 2020, Barocas et al., 2023, Colnet et al., 2024, Makhlouf et al., 2024]. To justify this, they often refer to Pearl, who argued that “the two frameworks can be used interchangeably and symbiotically”.¹ However, influential works on “equivalences” between the two causal frameworks have mostly focused on translating graphical assumptions into conditional-independence restrictions instead of actually proving whether counterfactual outcomes were equal across models, or implicitly addressed specific cases. Notably, [Pearl, 2009, Chapter 7] and [Richardson and Robins, 2013] consider *ex nihilo* the exchangeability of the two associated notations in their unifications of both causal frameworks.

In this paper, we essentially aim at clarifying in which sense using interchangeably two distinct causal models is appropriate. To this end, we compare a potential-outcome model and a structural causal model compatible with a same distribution of observations from an agnostic perspective. We introduce three levels of comparisons, corresponding to different degrees of counterfactual reasoning, and neutrally ask under which conditions two models are (un)distinguishable at these levels. This analysis crucially reminds that the models are *not* mathematically bound to correspond, meaning that using them symbiotically generally rests on a *choice*. Moreover, it classifies real-world scenarios where the models can(not) be exchanged, depending on their respective assumptions. Then, we interpret the counterfactual statements and causal effects respectively induced by $(Y_t)_{t \in \mathcal{T}}$ and $(Y_{T=t})_{t \in \mathcal{T}}$ when the models do not coincide, and explain how such results relate to the so-called *formal equivalence* between causal frameworks accepted by the causal-inference community.

In a similar vein, Ibeling and Icard [2024] recently provided an in-depth theoretical comparison of the two frameworks by adopting a neutral viewpoint, notably proving that a well-behaved RCM can always be *represented* by an SCM. Altogether, our contributions supplement their work by specifying graphical assumptions that must generally satisfy such an SCM, and by giving a real-world interpretation to these assumptions. In particular, analyzing theoretical representability results through the prism of practically relevant problems enables us to point out overlooked paradoxes in the causal-inference literature. On the basis of our results and discussions, we call the community to rigorously justify their exchanges of models across frameworks, as it could lead to misleading conclusions. In doing so, we hope to further clarify the role of each causal modeling in the past, current, and future causal-inference research.

1.1 Motivating example

To motivate this work, we illustrate on a concrete example how potential outcomes and structural counterfactuals can be respectively used to address a specific causal-inference problem.

1.1.1 Problem

For simplicity and concision, this example uses informal definitions of causal models; Sections 2 and 3 introduce formal definitions. All generic notations are specified in Section 2.1.

We consider the following fairness-inspired problem. Let the *treatment status* $T : \Omega \rightarrow \{0, 1\}$ be a binary random variable indicating the gender, $T(\omega) = 0$ standing for women and $T(\omega) = 1$ standing for men; let the *covariate* $X : \Omega \rightarrow \mathbb{R}$ be a random variable quantifying the level of work experience, a higher score encoding a more adapted experience; let the *outcome* $Y : \Omega \rightarrow \mathbb{R}$ be a random variable evaluating a candidate’s application for some position, a better score giving a higher probability of acceptance. Causal analysts are tasked with assessing the fairness of the evaluation process, by answering the question “What is the average outcome variation for individuals with a given profile had their gender changed?”. To do so, analyst \mathcal{M} relies on a structural causal model to compute the do-intervention outputs $(Y_{T=0}, Y_{T=1})$ (called structural counterfactuals), and then estimates $\mathbb{E}[Y_{T=1} - Y_{T=0} \mid X = x]$. For their part, analyst \mathcal{R} introduces potential outcomes (Y_0, Y_1) satisfying causal-inference assumptions, and then estimates $\mathbb{E}[Y_1 - Y_0 \mid X = x]$. The signs and intensities of the above estimands quantify (un)fairness. We ask

¹<http://causality.cs.ucla.edu/blog/index.php/2012/12/03/judea-pearl-on-potential-outcomes/>

whether analysts \mathcal{M} and \mathcal{R} reach the same conclusion, in other words if the two causal estimands are equal.

Analyst \mathcal{M} postulates that (T, X, Y) is ruled by the following collection of structural assignments:

$$\begin{aligned} T &= U_T, \\ X &= \alpha T + U_X, \\ Y &= X + \beta T + U_Y, \end{aligned}$$

where α and β are deterministic parameters quantifying the causal influence of T onto respectively X and Y , and U_X represents the hidden merit or effort of an individual. Typically, a positive parameter α describes the societal inequalities leading women to have a lower level of work experience than men with equal merit U_X . The variables U_T , U_X , and U_Y are exogenous noises such that $U_Y \perp\!\!\!\perp (T, X)$. For every $t \in \{0, 1\}$, the do-intervention denoted by $\text{do}(T = t)$ enables one to compute the downstream effect onto X and Y of fixing T to the value t . Concretely, it recursively defines $X_{T=t} = \alpha t + U_X$ and $Y_{T=t} = X_{T=t} + \beta t + U_Y = (\alpha + \beta)t + U_X + U_Y$. Therefore,

$$\mathbb{E}[Y_{T=1} - Y_{T=0} \mid X = x] = \alpha + \beta.$$

Analyst \mathcal{R} postulates two potential outcomes Y_0 and Y_1 , relating to the factual outcome Y via the following consistency property: $Y = (1 - T) \cdot Y_0 + T \cdot Y_1$. To compute their estimand from observational data, they suppose that $\mathbb{P}(T = 1 \mid X = x) > 0$ and $(Y_0, Y_1) \perp\!\!\!\perp T \mid X$. These conditions are referred to as the fundamental assumptions of causal inference. They imply that $\mathbb{E}[Y_1 - Y_0 \mid X = x] = \mathbb{E}[Y \mid X = x, T = 1] - \mathbb{E}[Y \mid X = x, T = 0]$ (see Lemma 2). We emphasize that, while they leverage different techniques, the two analysts work with the same (T, X, Y) . Therefore, $\mathbb{E}[Y \mid X = x, T = t] = x + \beta t + \mathbb{E}[U_Y]$ due to $U_Y \perp\!\!\!\perp (T, X)$. This leads to

$$\mathbb{E}[Y_1 - Y_0 \mid X = x] = \beta.$$

Consequently, $\mathbb{E}[Y_{T=1} - Y_{T=0} \mid X = x] \neq \mathbb{E}[Y_1 - Y_0 \mid X = x]$ if $\alpha \neq 0$. This means that the two analysts obtain distinct, possibly contradictory results.

1.1.2 Comments

This example shows that potential outcomes $(Y_t)_{t \in \mathcal{T}}$ equipped with the fundamental assumptions of causal inference do not necessarily define the same causal effects as the post-intervention outcomes $(Y_{T=t})_{t \in \mathcal{T}}$ obtained via do-interventions on T . As such, one cannot always interchangeably employ given potential outcomes and structural counterfactuals. This illustration motivates a general comparison of the two causal-inference approaches, which is precisely the goal of this paper. Let us summarize how the example's specifics connect to the broader results presented in this manuscript.

In Section 3, we propose to analyze the similarities and differences of the two approaches through the prism of equivalence relations between causal models. These relations characterize which causal estimands coincide across two models. For instance, in the above example $\mathbb{E}[Y_{T=1} - Y_{T=0} \mid X = x] \neq \mathbb{E}[Y_1 - Y_0 \mid X = x]$ means that the corresponding SCM and RCM were not *single-outcome equivalent* (Definition 6).

In Section 4, we notably investigate under which conditions an RCM satisfying the fundamental assumptions of causal inference and an SCM are equivalent or not. More precisely, we show that they are generally single-outcome equivalent just in case T does not cause X (Theorem 1). Therefore, in typical fairness problems where T is an immutable variable like sex or race, swapping potential outcomes and structural counterfactuals in a mathematical formula commonly yields distinct estimands. This explains the results of the above illustration. Our analysis also provides a general interpretation of the respective counterfactual variables: a structural counterfactual $Y_{T=t}$ of Y changes by do-intervention T into t and X into $X_{T=t}$ while keeping U_Y equal; a potential outcome Y_t of Y under the assumptions of causal inference changes T into t while keeping X and U_Y equal. This difference in meanings justifies why $\mathbb{E}[Y_{T=1} - Y_{T=0} \mid X = x]$ captured the total effect $\alpha + \beta$ of T on Y , whereas $\mathbb{E}[Y_1 - Y_0 \mid X = x]$ captured the direct effect β .

In Section 5, we detail how these results relate to the formal equivalence between causal frameworks. Our discussion delineates two paradigms for defining and using potential-outcome variables. The first is what analyst \mathcal{M} does: defining potential outcomes via do-interventions on the latent SCM, so that

the properties they satisfy stem from the SCM. The second is what analyst \mathcal{R} does: defining potential outcomes with a different causal interpretation than the one given by do-interventions, in their case by directly placing assumptions upon them—regardless of what holds in the latent SCM. Crucially, only analyst \mathcal{M} relies on the formal equivalence. This is why we caution against unjustified exchanges of the potential-outcome subscript notation and the do notation.

1.2 Outline of the paper

The rest of the paper proceeds as follows. Section 2 furnishes the basic knowledge on structural causal models and potential-outcome models. Section 3 formalizes the problem we address by introducing notions of equivalence between causal models. Section 4 clarifies under which conditions a structural causal model and a potential-outcome model compatible with the same observational data need or need not be equivalent. Section 5 discusses the relation between these results and the formal equivalence between causal frameworks. Appendix A addresses a supplementary illustration. Proofs of intermediary results are deferred to Appendix B.

2 Preliminaries

This section provides the necessary background on structural causal models and potential outcomes. It is meant to keep the paper self-contained. Section 2.1 introduces generic mathematical notations; Section 2.2 presents Pearl’s causal framework; Section 2.3 explains Rubin’s causal framework.

2.1 Basic mathematical notations

Throughout, we consider a probability space $(\Omega, \Sigma, \mathbb{P})$ with Ω a sample space, Σ a σ -algebra, and $\mathbb{P} : \Sigma \rightarrow [0, 1]$ a probability measure. This space does not necessarily have a physical interpretation; it abstractly represents the possible underlying states of the world. Crucially, it serves as the common mathematical basis to define and compare random variables.

A *random variable* W (including *random vectors*) is a measurable function from Ω to a Borel subset of an Euclidean space equipped with the Borel σ -algebra. It produces a probability distribution on its output space: we write $\mathcal{L}(W) := \mathbb{P} \circ W^{-1}$ and $\mathbb{E}[W] := \int W(\omega) d\mathbb{P}(\omega)$ for respectively the *law and expectation under* \mathbb{P} of a random variable W .² We emphasize that the laws of univariate random variables can be completely general in this paper; we do not suppose them to be either Lebesgue-absolutely continuous or discrete. For any Borel set F , we use the common probability-textbook notation $\{W \in F\}$ for the set $\{\omega \in \Omega \mid W(\omega) \in F\} \in \Sigma$. Two variables W_1 and W_2 are \mathbb{P} -almost surely equal, denoted by $W_1 \stackrel{a.s.}{=} W_2$, if $\mathbb{P}(W_1 = W_2) = 1$; they are *equal in law under* \mathbb{P} , denoted by $\mathcal{L}(W_1) = \mathcal{L}(W_2)$, if $\mathbb{P}(W_1 \in F) = \mathbb{P}(W_2 \in F)$ for every Borel set F . The notation $W_1 \perp\!\!\!\perp W_2$ means that W_1 and W_2 are *independent under* \mathbb{P} , that is $\mathbb{P}(W_1 \in F_1, W_2 \in F_2) = \mathbb{P}(W_1 \in F_1) \cdot \mathbb{P}(W_2 \in F_2)$ for all Borel sets F_1, F_2 .

We denote by $\mathbb{P}(\cdot \mid W = w)$ the *regular conditional probability measure with respect to* $\{W = w\}$, which exists and is unique for $\mathcal{L}(W)$ -almost every w . Then, whenever they are well-defined, we write $\mathcal{L}(W_2 \mid W_1 = w_1) := \mathbb{P}(\cdot \mid W_1 = w_1) \circ W_2^{-1}$ and $\mathbb{E}[W_2 \mid W_1 = w_1] := \int W_2(\omega) d\mathbb{P}(\omega \mid W_1 = w_1)$ for respectively the *law and expectation of* W_2 *conditional to* $W_1 = w_1$. The expression $W_1 \perp\!\!\!\perp W_2 \mid W_3$ means that W_1 and W_2 are *independent conditional to* W_3 *under* \mathbb{P} , namely that W_1 and W_2 are independent under $\mathbb{P}(\cdot \mid W_3 = w_3)$ for $\mathcal{L}(W_3)$ -almost every w_3 .

Moreover, for any tuple $w := (w_i)_{i \in \mathcal{I}}$ indexed by a finite index set \mathcal{I} and any subset $I \subseteq \mathcal{I}$ we write $w_I := (w_i)_{i \in I}$. Then, for any index set J we abusively write w_J for $w_{J \cap \mathcal{I}}$ and use the notation $w_J = \emptyset$ whenever $J \cap \mathcal{I} = \emptyset$. Similarly, we define the Cartesian product $\mathcal{W}_I := \prod_{i \in I} \mathcal{W}_i$ for any collection of spaces $(\mathcal{W}_i)_{i \in \mathcal{I}}$.

2.2 Pearl’s causal framework

Pearl’s causal modeling mathematically formalizes associations that standard probability calculus cannot describe through the notions of structural causal models and do-interventions [Pearl, 2009]. This section

²To fix ideas, one can consider that $\Omega = [0, 1]$ and that \mathbb{P} is the uniform distribution on Ω so that for any Borel probability distribution P there exists W such that $P = \mathbb{P} \circ W^{-1}$ (as reminded in the proof of Proposition 3). As such, this choice allows one to define random variables or vectors with any laws.

recalls the basics on this topic, borrowing the introduction proposed in [Blom et al., 2020, Bongers et al., 2021].

2.2.1 Structural causal models

A *structural causal model* (SCM) represents the causal relationships between the studied variables. It is the cornerstone of Pearl’s causal framework.

Definition 1 (Structural causal model). Let \mathcal{I} and \mathcal{J} be two disjoint finite index sets, and write $\mathcal{V} := \prod_{i \in \mathcal{I}} \mathcal{V}_i \subseteq \mathbb{R}^{|\mathcal{I}|}$, $\mathcal{U} := \prod_{j \in \mathcal{J}} \mathcal{U}_j \subseteq \mathbb{R}^{|\mathcal{J}|}$ for two Borel product spaces. A *structural causal model* \mathcal{M} is a couple $\langle U, g \rangle$ where:

1. $U := (U_j)_{j \in \mathcal{J}}$ is a collection of mutually independent random variables called the *random noises*, such that U_j is from Ω to \mathcal{U}_j for every $j \in \mathcal{J}$;
2. $g := (g_i)_{i \in \mathcal{I}}$ is a collection of measurable \mathbb{R} -valued functions, where for every $i \in \mathcal{I}$ there exist two subsets of indices $\text{Endo}(i) \subseteq \mathcal{I}$ and $\text{Exo}(i) \subseteq \mathcal{J}$, respectively called the *endogenous* and *exogenous parents* of i , such that g_i is from $\mathcal{V}_{\text{Endo}(i)} \times \mathcal{U}_{\text{Exo}(i)}$ to \mathcal{V}_i .³

A random vector $V : \Omega \rightarrow \mathcal{V}$ is a solution of \mathcal{M} if for every $i \in \mathcal{I}$,

$$V_i \stackrel{a.s.}{=} g_i(V_{\text{Endo}(i)}, U_{\text{Exo}(i)}). \quad (1)$$

The equations defined by (1) and characterized by g and U are called the *structural equations*.⁴

Such a model explains how some *endogenous* variables V , representing observed data, are generated from *exogenous* variables U , describing background factors. The structural equations quantify the causal dependencies between all these variables and are frequently illustrated by the directed graph $\mathcal{G}_{\mathcal{M}}$ with nodes $\mathcal{I} \cup \mathcal{J}$, and such that a directed edge points from node k to node l if and only if $k \in \text{Endo}(l) \cup \text{Exo}(l)$ (we say in this case that k is a parent of l). For convenience, we make the common assumption that the studied models are *acyclic*, which means that their associated graphs do not contain any cycles.

Assumption 1 (Acyclicity). \mathcal{M} is such that $\mathcal{G}_{\mathcal{M}}$ is a directed *acyclic* graph.

Not only acyclicity simplifies the interpretation of causal dependencies, but it entails *unique solvability* of the SCM: according to [Bongers et al., 2021, Proposition 3.4], Equation (1) admits a unique solution up to \mathbb{P} -negligible sets. We will abusively refer to such a solution as *the* solution of the SCM. Also, on the basis of this well-posedness of the solution V , we will often replace the indices $i \in \mathcal{I}$ and $j \in \mathcal{J}$ by the associated random variables V_i and U_j in the relevant notations. For example we may write $\text{Endo}(V_i)$ instead of $\text{Endo}(i)$.

The purpose of causal structures is to capture the assumption that variables are not independently manipulable. As we detail next, they enable to understand the downstream effect of fixing some variables to certain values onto nonintervened variables.

2.2.2 The do-intervention

A *perfect do-intervention* is an operation forcing a set of endogenous variables to take predefined values while keeping all the rest of the causal mechanism equal.

Definition 2 (Perfect do-intervention). Let $\mathcal{M} = \langle U, g \rangle$ be an SCM, $I \subseteq \mathcal{I}$ a subset of endogenous variables, and $\tilde{v}_I \in \mathcal{V}_I$ a value. The action $\text{do}(I, \tilde{v}_I)$ defines the modified model $\mathcal{M}_{\text{do}(I, \tilde{v}_I)} = \langle U, \tilde{g} \rangle$ where \tilde{g} is given by: for any $(v, u) \in \mathcal{V} \times \mathcal{U}$ and $i \in \mathcal{I}$,

$$\tilde{g}_i(v_{\text{Endo}(i)}, u_{\text{Exo}(i)}) := \begin{cases} \tilde{v}_i & \text{if } i \in I, \\ g_i(v_{\text{Endo}(i)}, u_{\text{Exo}(i)}) & \text{if } i \in \mathcal{I} \setminus I. \end{cases}$$

³This definition tolerates that distinct endogenous variables share the same exogenous parents, that is $\text{Exo}(i) \cap \text{Exo}(i') \neq \emptyset$ for some $i \neq i'$. Therefore, the variables in $(U_{\text{Exo}(i)})_{i \in \mathcal{I}}$ are not necessarily mutually independent.

⁴In the paper, we will often informally define an SCM by specifying the structural equations rather than g .

A do-intervention preserves acyclicity, and therefore unique solvability. As a consequence, if V is the solution of an acyclic \mathcal{M} , one can define (up to \mathbb{P} -negligible sets) its post-intervention counterpart $V_{\text{do}(I, v_I)}$ solution to $\mathcal{M}_{\text{do}(I, v_I)}$. It describes an alternative world where every V_i for $i \in I$ is set to value v_i while U is preserved. In the rest of the paper, we simply write $\text{do}(V_I = v_I)$ for the operation $\text{do}(I, v_I)$, and use the subscript $V_I = v_I$ to indicate results of this operation. For instance, we write $\mathcal{M}_{V_I = v_I}$ for $\mathcal{M}_{\text{do}(I, v_I)}$ and $V_{V_I = v_I}$ for $V_{\text{do}(I, v_I)}$. Crucially, intervening does not amount to conditioning in general, that is $\mathcal{L}(V \mid V_I = v_I) \neq \mathcal{L}(V_{V_I = v_I})$.

The next proposition provides a general expression of the solution before and after intervention, and will play a key role throughout this paper. For any $I \subseteq \mathcal{I}$ we write $\text{Endo}(I) = \cup_{i \in I} \text{Endo}(i)$, $\text{Exo}(I) = \cup_{i \in I} \text{Exo}(i)$, and $I^c = \mathcal{I} \setminus I$.

Lemma 1 (do-intervention on variables). *Let $\mathcal{M} = \langle U, g \rangle$ be an SCM satisfying acyclicity (Assumption 1) with solution V , and consider $I \subseteq \mathcal{I}$. There exists a deterministic measurable function f_{I^c} such that*

$$V_{I^c} \stackrel{\text{a.s.}}{=} f_{I^c}(V_{\text{Endo}(I^c) \setminus I^c}, U_{\text{Exo}(I^c)}).$$

Moreover, for any intervention $\text{do}(V_I = v_I)$ the solution \tilde{V} of $\mathcal{M}_{V_I = v_I}$ satisfies

$$\begin{aligned} \tilde{V}_{I^c} &\stackrel{\text{a.s.}}{=} f_{I^c}(v_{\text{Endo}(I^c) \setminus I^c}, U_{\text{Exo}(I^c)}), \\ \tilde{V}_I &\stackrel{\text{a.s.}}{=} v_I. \end{aligned}$$

Importantly, this is the same deterministic function f_{I^c} and the same random noises $U_{\text{Exo}(I^c)}$ that generate V_{I^c} and its post-intervention counterpart \tilde{V}_{I^c} , the only change being the assignment $V_I = v_I$. Slightly abusing notations, we will often artificially extend the input variables of f_{I^c} to write $V_{I^c} \stackrel{\text{a.s.}}{=} f_{I^c}(V_I, U_{\text{Exo}(I^c)})$ and $\tilde{V}_{I^c} \stackrel{\text{a.s.}}{=} f_{I^c}(v_I, U_{\text{Exo}(I^c)})$ using the fact that $\text{Endo}(I^c) \setminus I^c \subseteq I$. Lemma 1 can be seen as a “vectorization” of the structural equations.

2.2.3 Counterfactual inference with structural causal models

Counterfactual inference aims at predicting outcomes had a certain event occurred given some factual observations. Typically, it addresses what-if questions such as “Had they been a woman, would have they gotten the position?”. Perfect interventions combined with conditioning provides a natural probabilistic framework to address counterfactual queries. Let for instance $V := (T, X, Y)$ be the solution to an acyclical SCM $\mathcal{M} := \langle U, g \rangle$. Pearl answers the question “had T been equal to t , what would have been the law of Y given the factual context $X = x$?” by using the so-called *three-step procedure* [Pearl, 2009]:

1. **(Abduction)** compute $\mathcal{L}(U \mid X = x)$, the posterior distribution of U -values compatible with the context $\{X = x\}$;
2. **(Action)** carry out a do-intervention on \mathcal{M} to obtain the intervened causal mechanism $g_{T=t}$ of $\mathcal{M}_{T=t}$;
3. **(Prediction)** pass the posterior distribution $\mathcal{L}(U \mid X = x)$ through $g_{T=t}$ to generate the distribution $\mathcal{L}(Y_{T=t} \mid X = x)$ of counterfactual outcomes.

More generally, an SCM enables one to sample from probability distributions of counterfactual outcomes for any choices of context, variables to alter by do-intervention, and outcomes of interest. Note that, while a fully specified SCM permits to compute any counterfactual distributions via the above algorithm, some distributions can also be identified from only the observational distribution $\mathcal{L}(V)$ and the causal graph $\mathcal{G}_{\mathcal{M}}$ (that is without completely knowing g or the law of U) under specific graphical assumptions (see for example the counterfactual interpretation of the *backdoor criterion* [Pearl et al., 2016, Theorem 4.3.1]). Before turning to the potential-outcome framework, let us precise our notations and definitions related to the “do”.

Remark 1 (do-interventions versus the do-operator). Some readers may be familiar with the *do-operator* and the associated notations [Pearl, 2009, Equation 1.37], which do not exactly correspond to the do-intervention operation from Definition 2. This operator serves to distinguish between probability distributions of endogenous variables after distinct do-interventions, similarly to how conditioning distinguishes between probability distributions given distinct realizations of random variables. The respective

notations of the do-operator and the conditioning operation are similar (but the do-operator respects different rules). Typically, one writes $\mathbb{E}[V \mid \text{do}(V_I = v_I)]$ for the expectation of the solution of $\mathcal{M}_{V_I=v_I}$, and $\mathbb{E}[Y \mid \text{do}(T = t), X = x]$ for the conditional expectation in $\mathcal{M}_{T=t}$ of the outcome given the covariates. These quantities can be expressed with the subscript-variable notations we introduced as follows: $\mathbb{E}[V \mid \text{do}(V_I = v_I)] = \mathbb{E}[V_{V_I=v_I}]$ and $\mathbb{E}[Y \mid \text{do}(T = t), X = x] = \mathbb{E}[Y_{T=t} \mid X_{T=t} = x]$. We refer to [Pearl et al., 2016, Chapter 4] for more details.

The do-operator along with the rules of do-calculus enable one to reframe post-intervention probabilities in terms of standard (conditional) probabilities [Pearl, 2009, Theorem 3.4.1], making them useful for conducting *interventional* inference. However, as explained by Pearl [Pearl et al., 2016, Section 4.1], they do not fully allow *counterfactual* inference. Notably, the do-operator distinguishes between probability distributions entailed by distinct interventions but not between random variables entailed by distinct interventions, and thereby fails to represent cross-world dependencies. For example, a quantity like $\mathbb{E}[Y_{T=t} \mid X = x]$, which involves variables defined across $\mathcal{M}_{T=t}$ and \mathcal{M} , cannot always be captured by the do-operator. In particular, $\mathbb{E}[Y_{T=t} \mid X = x] \neq \mathbb{E}[Y \mid \text{do}(T = t), X = x]$ in general [Pearl et al., 2016, Equation 4.8].

All in all, the do-operator does not provide a sufficient vocabulary for the problem we address: determining under which conditions one can substitute in causal estimands the counterfactual *variables* $(Y_t)_{t \in \mathcal{T}}$ of Rubin’s framework with the post-intervention *variables* $(Y_{T=t})_{t \in \mathcal{T}}$ of Pearl’s framework. This is why we do not use notations as $\mathbb{E}[\cdot \mid \text{do}(X = x)]$ in this paper. We see a do-intervention as a transformation of a random variable (instead of a distribution), and indicate post-intervention variables by a subscript as commonly done in the structural-counterfactual literature [Pearl et al., 2016, Chapter 4]. Moreover, we only employ classical operators from probability theory, like $\mathbb{E}[\cdot \mid X = x]$ (as defined in Section 2.1).

2.3 Rubin’s causal framework

The potential-outcome framework, also known as *Neyman-Rubin causal modeling* [Neyman, 1923, Rubin, 1974], was designed to understand the causal effect of a treatment onto an outcome of interest, for instance when one aims at assessing the contribution of a drug to recovering from some disease in clinical trials. In this section, we introduce this framework in the specific case of a binary treatment.

2.3.1 Potential outcomes

Let $T : \Omega \rightarrow \{0, 1\}$ represent a binary *treatment status*, typically such that $T(\omega) = 0$ indicates the absence of treatment and $T(\omega) = 1$ indicates a treatment. More generally, it can encode any distinction between some groups (*e.g.*, men and women). Assuming *no interference between units*,⁵ this framework postulates two *potential outcomes* $Y_0 : \Omega \rightarrow \mathbb{R}$ and $Y_1 : \Omega \rightarrow \mathbb{R}$, one for each treatment status. These potential outcomes as well as the treatment may depend on some covariates $X : \Omega \rightarrow \mathbb{R}^d$ (such as the patient’s weight, height, or historical data in clinical trials). Critically, we cannot observe simultaneously $Y_0(\omega)$ and $Y_1(\omega)$ for a same ω : a problem referred as the *fundamental problem of causal inference* [Holland, 1986]. We only have access to the realized *outcome variable* $Y : \Omega \rightarrow \mathbb{R}$ which is supposed to be *consistent* with (Y_0, Y_1) , that is satisfying $Y = (1 - T) \cdot Y_0 + T \cdot Y_1$. Concretely, if $T(\omega) = 1$ for some $\omega \in \Omega$, then $Y(\omega) = Y_1(\omega)$, and $Y_0(\omega)$ becomes unidentifiable by mere observations. We refer to the random vector (T, X, Y_0, Y_1) as the *Rubin causal model* (RCM), which is an augmented version of (T, X, Y) due to consistency.

Understanding the causal relationship between the treatment and the outcome in this framework ideally consists in answering counterfactual questions such as “What would have been the value of $Y(\omega)$ had $T(\omega)$ been equal to 1 instead of 0 for a specific ω (such that $X(\omega) = x$)?”. This cannot be answered since the value of either $Y_0(\omega)$ or $Y_1(\omega)$ will always be missing. Instead, in practice, one estimates and compares under some assumptions features of $\mathcal{L}(Y_1)$ and $\mathcal{L}(Y_0)$, or $\mathcal{L}(Y_1 \mid X = x)$ and $\mathcal{L}(Y_0 \mid X = x)$. People commonly focus on computing the *average treatment effect* $\mathbb{E}[Y_1 - Y_0]$ or the *conditional average treatment effect* (CATE) $\mathbb{E}[Y_1 - Y_0 \mid X = x]$. The main challenge lies in the fact that *association is not causation* in general. In particular, the quantity $\mathbb{E}[Y \mid T = t]$ does not necessarily coincide with the quantity $\mathbb{E}[Y_t]$ for $t \in \{0, 1\}$. Typically, if some medical treatment is more likely to be taken by weaker

⁵Paraphrasing Rubin Rubin [2010], a *unit* refers to a study object (like a person). This assumption excludes cases where the treatment of one unit may affect the outcome of another.

patients, we may observe a lower rate of recovery among the treated group compared to the nontreated group due to the health condition even though the medicine does increase recovery all other things being kept equal: we would observe $\mathbb{E}[Y | T = 1] < \mathbb{E}[Y | T = 0]$ while $\mathbb{E}[Y_1] > \mathbb{E}[Y_0]$ (a phenomenon that can be seen as a consequence of *Simpson’s paradox* [Blyth, 1972]). In this case, the health condition is called a *confounder*: a variable associated with both the distribution of the treatment and the outcome. However, causal inference from observational data is still possible, as explained next.

2.3.2 Counterfactual inference with fundamental assumptions

Estimating a feature of $\mathcal{L}((T, X, Y_0, Y_1))$, like a treatment effect, requires expressing it in terms of features of $\mathcal{L}((T, X, Y))$ which generates the empirical observations. Such *identifications* can be achieved under two fundamental assumptions. The first one goes by many names through the literature: *conditional ignorability*, *conditional exchangeability*, *conditional exogeneity*, and *conditional unconfoundedness* (among others). Originally formulated by Rosenbaum and Rubin [1983], it states that the potential outcomes are independent of the treatment conditional to the covariates, that is $(Y_0, Y_1) \perp\!\!\!\perp T | X$. Said differently, it ensures that all confounders between the treatment and the potential outcomes are included in the covariates. Note that this assumption is untestable, as it would require to observe simultaneously the two potential outcomes. The second key hypothesis is *positivity*, which ensures that all units can be exposed to both treatment statuses, that is $0 < \mathbb{P}(T = 1 | X = x) < 1$ for $\mathcal{L}(X)$ -almost every $x \in \mathbb{R}^d$. It readily follows from positivity that the probability distribution $\mathcal{L}(Y | X = x, T = t)$ is well defined for $\mathcal{L}(X)$ -almost every $x \in \mathbb{R}^d$ and every $t \in \{0, 1\}$, and from conditional ignorability that it coincides with $\mathcal{L}(Y_t | X = x)$, meaning that association-based outcomes have a causal interpretation. Several statistical methods coexist to estimate the (conditional) average causal effect, all building upon this implication (see for instance [Imbens, 2004, Yao et al., 2021]). We do not detail them for concision and clarity since it is not the topic of this paper. We only point out that, similarly to SCMs, the potential-outcome framework enables one to infer distributions of counterfactual outcomes.

3 Problem setup

This section precises the problem we address: analyzing the mathematical similarities and differences between two models respectively derived from the two causal frameworks. Section 3.1 introduces a potential-outcome model and a structural causal model compatible with a same dataset, and formalizes the assumptions we may place upon them. Section 3.2 defines notions of equivalence between the two models, corresponding to different levels of comparison.

3.1 Causal models and assumptions

The next definition formalizes the notion of observational data in a causal-inference experiment. It generalizes Section 2 by considering a nonbinary treatment and a multivariate outcome.

Definition 3 (Observational vector). Let $N, d, p \geq 1$ be integers, and $\mathcal{T} := \{0, 1, \dots, N\}$. An *observational vector* \mathcal{O} is a random vector (T, X, Y) where $T : \Omega \rightarrow \mathcal{T}$ is the *treatment status*, $X : \Omega \rightarrow \mathbb{R}^d$ the *covariates*, $Y : \Omega \rightarrow \mathbb{R}^p$ the *outcome*, and such that $0 < \mathbb{P}(T = t)$ for all $t \in \mathcal{T}$.

In order to compare the two causal frameworks, we consider in what follows a superimposed construction where an observational vector $\mathcal{O} := (T, X, Y)$ is concurrently governed by an RCM and an SCM. Figure 1 illustrates this construction. We emphasize that we adopt an agnostic approach where there is no presumed relation between the two causal models except their compatibility with \mathcal{O} . This is meant to highlight when equality (possibly in law) between potential and structural counterfactual outcomes is a mathematical necessity or the result of specific assumptions.

3.1.1 Potential-outcome model

On the one hand, we assume that T is the treatment status, X some covariates, and Y the outcome of interest in a potential-outcome model in the sense of the definition below.

Definition 4 (Rubin causal model). Let $\mathcal{O} := (T, X, Y)$ be an observational vector. A *Rubin causal model (or potential-outcome model) compatible with \mathcal{O}* is a random vector $\mathcal{R} := (T, X, (Y_t)_{t \in \mathcal{T}})$ such that the tuple $(Y_t)_{t \in \mathcal{T}} : \Omega \rightarrow (\mathbb{R}^p)^{N+1}$ meets the *consistency rule*:

$$Y \stackrel{a.s.}{=} \sum_{t \in \mathcal{T}} \mathbf{1}_{\{T=t\}} Y_t.$$

For any $t \in \mathcal{T}$, Y_t is referred to as the *potential outcome* had the treatment been equal to t . We denote by $\mathfrak{R}_{\mathcal{O}}$ the class of Rubin causal models compatible with \mathcal{O} .

A generic RCM $\mathcal{R} := (T, X, (Y_t)_{t \in \mathcal{T}}) \in \mathfrak{R}_{\mathcal{O}}$ is basically an augmented version of \mathcal{O} by $N + 1$ random vectors $(Y_t)_{t \in \mathcal{T}}$ satisfying the consistency rule. The variable Y is not explicit in \mathcal{R} as it can be recovered from T and $(Y_t)_{t \in \mathcal{T}}$ via consistency. In this setting, the first fundamental assumption for causal inference can be written as follows.

Assumption 2 (Positivity). \mathcal{O} is such that for all $t \in \mathcal{T}$ and $\mathcal{L}(X)$ -almost every x , $0 < \mathbb{P}(T = t \mid X = x) < 1$.

Remark that, strictly speaking, positivity in an hypothesis on the observational vector rather than the RCM. But people typically suppose it in the context of potential outcomes. We distinguish two formulations for the second fundamental assumption, namely conditional ignorability.

Assumption 3 (Cross-outcome conditional ignorability). \mathcal{R} is such that $(Y_t)_{t \in \mathcal{T}} \perp\!\!\!\perp T \mid X$.

Assumption 4 (Single-outcome conditional ignorability). \mathcal{R} is such that for all $t \in \mathcal{T}$, $Y_t \perp\!\!\!\perp T \mid X$.

The stronger version (Assumption 3) is the original one, but most causal-inference methods only require the weaker version (Assumption 4). In our main results, we will clearly specify which form of conditional ignorability is required. Crucially, Assumptions 2 and 4 permit to fully identify the law of (T, X, Y_t) by the law of (T, X, Y) for any $t \in \mathcal{T}$, as recalled below.

Lemma 2 (Single-outcome identification of potential outcomes). *Let $\mathcal{O} := (T, X, Y)$ be an observational vector satisfying positivity (Assumption 2). For any $\mathcal{R} := (T, X, (Y_t)_{t \in \mathcal{T}}) \in \mathfrak{R}_{\mathcal{O}}$, if \mathcal{R} meets single-outcome conditional ignorability (Assumption 4) then for every $t \in \mathcal{T}$ and $\mathcal{L}(X)$ -almost any $x \in \mathbb{R}^d$:*

$$\mathcal{L}(Y_t \mid X = x) = \mathcal{L}(Y \mid X = x, T = t).$$

Thereby, for any Borel set $F \subseteq \mathcal{T} \times \mathbb{R}^d \times \mathbb{R}^p$,

$$\mathbb{P}((T, X, Y_t) \in F) = \int \mathbb{P}(Y \in F(t', x) \mid X = x, T = t) d\mathbb{P}(X = x, T = t'),$$

where $F(t', x) := \{y \in \mathbb{R}^p \mid (t', x, y) \in F\}$ for every $(t', x) \in \mathcal{T} \times \mathbb{R}^d$.

3.1.2 Structural causal model

On the other hand, we assume that the variables in \mathcal{O} are generated by a latent SCM \mathcal{M} , as specified below.

Definition 5 (Compatible structural causal model). Let \mathcal{O} be an observational vector. An SCM $\mathcal{M} := \langle U, g \rangle$ is *compatible with \mathcal{O}* if it admits a unique solution V equal to \mathcal{O} . We denote by $\mathfrak{M}_{\mathcal{O}}$ the class of SCMs satisfying acyclicity (Assumption 1) and compatible with \mathcal{O} .

The latent model \mathcal{M} defines for every $t \in \mathcal{T}$ the post-intervention outcome $Y_{T=t}$ under $\text{do}(T = t)$ using Definition 2. We refer to $Y_{T=t}$ as the *structural counterfactual outcome* had the treatment been equal to t . Similarly to potential outcomes, structural counterfactuals satisfy the consistency rule, and therefore induce an RCM.

Lemma 3 (Consistency rule for structural counterfactuals). *Let $\mathcal{O} := (T, X, Y)$ be an observational vector. For any $\mathcal{M} \in \mathfrak{M}_{\mathcal{O}}$, the tuple $(Y_{T=t})_{t \in \mathcal{T}}$ obtained via do-interventions in \mathcal{M} meets the consistency rule,*

$$Y \stackrel{a.s.}{=} \sum_{t \in \mathcal{T}} \mathbf{1}_{\{T=t\}} Y_{T=t}.$$

Therefore, $\mathcal{R}_{\mathcal{M}} := (T, X, (Y_{T=t})_{t \in \mathcal{T}})$ belongs to $\mathfrak{R}_{\mathcal{O}}$. We call $\mathcal{R}_{\mathcal{M}}$ the entailed RCM of \mathcal{M} .




Figure 1: Superimposed construction of an RCM \mathcal{R} and an SCM \mathcal{M} compatible with a same observational vector \mathcal{O} .

The notion of entailed RCM bridges SCMs to RCMs, thereby enables one to compare these two classes of models which differ fundamentally in their constructions of counterfactual outcomes (see Remark 2). Contrasting $\mathcal{R}_{\mathcal{M}} := (T, X, (Y_{T=t})_{t \in \mathcal{T}})$ against $\mathcal{R} := (T, X, (Y_t)_{t \in \mathcal{T}})$ almost surely or in law is precisely the goal of this paper, as formalized next in Section 3.2.

Remark 2 (Primitives versus derivatives). As noted by Pearl Pearl [2010], the potential outcomes $(Y_t)_{t \in \mathcal{T}}$ are “undefined *primitives*” of the RCM, not related to any formal of measurable quantities, while the post-intervention outcomes $(Y_{T=t})_{t \in \mathcal{T}}$ are “*derivatives*” of the SCM by application of do-interventions. Said differently, the firsts are inputs *defining* the causal model, whereas the seconds are post-intervention outputs *defined by* the causal model. Figure 1 illustrates this aspect. As such, consistency is a *theorem* for structural counterfactuals and a constitutive *assumption* for potential outcomes. This holds more generally for any conceivable properties on counterfactual outcomes: in \mathcal{R} they usually are assumptions; in $\mathcal{R}_{\mathcal{M}}$ they are derived from \mathcal{M} .

3.2 Notions of equivalence between causal models

We aim at studying the mathematical similarities and differences between potential outcomes and structural counterfactuals from a theoretically neutral perspective. We consider three levels of comparison that will guide our analysis throughout the paper.

Definition 6 (Equivalences between causal models). Let \mathcal{O} be an observational vector, and $\mathcal{R}_1 := (T, X, (Y_t^{(1)})_{t \in \mathcal{T}})$, $\mathcal{R}_2 := (T, X, (Y_t^{(2)})_{t \in \mathcal{T}})$ be two models in $\mathfrak{R}_{\mathcal{O}}$. We say that \mathcal{R}_1 and \mathcal{R}_2 are:

- (i) *almost-surely equivalent* if $Y_t^{(1)} \stackrel{a.s.}{=} Y_t^{(2)}$ for any $t \in \mathcal{T}$;
- (ii) *cross-outcome equivalent* if $\mathcal{L}\left((T, X, (Y_t^{(1)})_{t \in \mathcal{T}})\right) = \mathcal{L}\left((T, X, (Y_t^{(2)})_{t \in \mathcal{T}})\right)$;
- (iii) *single-outcome equivalent* if $\mathcal{L}\left((T, X, Y_t^{(1)})\right) = \mathcal{L}\left((T, X, Y_t^{(2)})\right)$ for every $t \in \mathcal{T}$.

By overloading terminology, for $\mathcal{M} \in \mathfrak{M}_{\mathcal{O}}$ and $\mathcal{R} \in \mathfrak{R}_{\mathcal{O}}$, we analogously say that \mathcal{M} is equivalent to \mathcal{R} in one of the above senses if $\mathcal{R}_{\mathcal{M}}$ is equivalent to \mathcal{R} in this sense.

These types of equivalence focus on (entailed) RCMs in contrast to the notions proposed in [Bongers et al., 2021, Beckers, 2021] which are specifically tailored to SCMs. Remark that (i) \implies (ii) \implies (iii). Mathematically, each relation characterizes a class of RCM-based estimands that are invariant under the swap of two equivalent models \mathcal{R}_1 and \mathcal{R}_2 , which amounts to exchanging their potential outcomes. If the models are almost-surely equivalent, then *all* estimands are invariant under the swap. If the models are cross-outcome equivalent, then any estimands *based on their cross-outcome distributions* are invariant under the swap: for instance, $\mathbb{E}\left[\|Y_1^{(1)} - Y_0^{(1)}\|^2 \mid X = x\right] = \mathbb{E}\left[\|Y_1^{(2)} - Y_0^{(2)}\|^2 \mid X = x\right]$. If the models are single-outcome equivalent, then any estimands *based on their single-outcome distributions* are invariant under the swap: for example, $\mathbb{E}\left[Y_1^{(1)} - Y_0^{(1)} \mid X = x\right] = \mathbb{E}\left[Y_1^{(2)} - Y_0^{(2)} \mid X = x\right]$. When equivalence

does not hold, some causal effects of the corresponding relation may vary under the swap. Less formally, the motivation for the three ladders in Definition 6 comes from different levels at which people reason counterfactually. Let us detail them with an RCM and an SCM.

The almost-sure level focuses on counterfactual questions at the scale of $\omega \in \Omega$. It asks (for example) “What would have been the value of $Y(\omega)$ had $T(\omega)$ been equal to 1 instead of 0?”. The answer is deterministic in both the potential-outcome approach and the structural approach, given by $Y_1(\omega)$ in the RCM and by $Y_{T=1}(\omega)$ in the SCM. Should the models be almost-surely equivalent, the answer would be identical for almost-every $\omega \in \Omega$. We emphasize that this level is methodologically inessential: in practice, one does not have access to the random variables of the models themselves but to some realizations of their laws. Nevertheless, comparing causal models on this almost-sure baseline is theoretically important to completely understand their mathematical differences.

We now turn to the cross-outcome level. Because counterfactual questions at the scale of ω cannot be answered, people rather address surrogate queries like “What would have been the law of Y had T been equal to 1 instead of 0 given that $X = x$?”. In both causal approaches, the answer is generally not deterministic (that is uniquely determined). Possible answers along with their probabilities are respectively described by $\mathcal{L}(Y_1 | X = x)$ (which can be estimated using the techniques mentioned in Section 2.3.2) in the potential-outcome framework and by $\mathcal{L}(Y_{T=1} | X = x)$ (which can be inferred via the three-step procedure from Section 2.2.3) in the SCM. More generally, most answers to counterfactual questions in the structural framework and in the potential-outcome framework are respectively characterized by the joint probability distributions $\mathcal{L}((T, X, (Y_{T=t})_{t \in \mathcal{T}}))$ and $\mathcal{L}((T, X, (Y_t)_{t \in \mathcal{T}}))$. Therefore, should the models be cross-outcome equivalent, they would yield the same conclusions when reasoning counterfactually at this level.

The single-outcome level resembles the cross-outcome level in the sense that it also focuses on distributions of outcomes; it differs by being mathematically weaker. It is motivated by the fact that researchers and practitioners predominantly ask counterfactual questions involving single counterfactual outcomes instead of joint counterfactual outcomes, as in the above paragraph. As such, knowing only the marginal laws $\mathcal{L}((T, X, Y_t))$ and $\mathcal{L}((T, X, Y_{T=t}))$ for every $t \in \mathcal{T}$ suffices to compute most of the practically relevant counterfactual estimands. This is why it is arguably the most critical level in practice. For illustration, consider average treatment effects like $\mathbb{E}[Y_1 - Y_0]$ or $\mathbb{E}[Y_1 - Y_0 | X = x]$ (as in Section 1.1.1), distributional treatment effects like $D(\mathcal{L}(Y_1), \mathcal{L}(Y_0))$ or $D(\mathcal{L}(Y_1 | X = x), \mathcal{L}(Y_0 | X = x))$ where D is a discrepancy between probability measures [Muandet et al., 2021, Park et al., 2021], or the *standard* counterfactual-fairness condition on Y : $\mathcal{L}(Y_{T=t'} | X = x, T = t) = \mathcal{L}(Y_{T=t} | X = x, T = t)$ for every $t \in \mathcal{T}$ [Kusner et al., 2017]. They all concern the single-outcome level.⁶ Therefore, should the models be single-outcome equivalent, they would yield the same results in these state-of-the-art methodologies.

Remark 3 (At which level do the fundamental assumptions matter?). According to Lemma 2, positivity (Assumption 2) and single-outcome conditional ignorability (Assumption 4) uniquely determines the single-outcome level of RCMs in $\mathfrak{R}_{\mathcal{O}}$ by $\mathcal{L}(\mathcal{O})$. As such, these assumptions allow the estimation of most practically-relevant causal effects. However, we point out that even cross-outcome conditional ignorability (Assumption 3) does not uniquely determines the cross-outcome level. In Remark 4 from Section 4.2, we study two RCMs $\mathcal{R}_1 := (T, X, (Y_t^{(1)})_{t \in \mathcal{T}})$ and $\mathcal{R}_2 := (T, X, (Y_t^{(2)})_{t \in \mathcal{T}})$ satisfying positivity and cross-outcome conditional ignorability such that $\mathcal{L}((T, X, (Y_t^{(1)})_{t \in \mathcal{T}})) \neq \mathcal{L}((T, X, (Y_t^{(2)})_{t \in \mathcal{T}}))$. All in all, the scope of the fundamental assumptions of causal inference is limited to the single-outcome level of counterfactual reasoning. Identifying a cross-outcome effect such as $\mathbb{E}[Y_1 | Y = 0, T = 0]$ requires stronger assumptions.

To summarize, when two causal models compatible with a same observational vector meet an equivalence condition from Definition 6, they can be used interchangeably at a certain level of counterfactual reasoning. This mathematically signifies that a certain class of causal effects is invariant under exchanges of their counterfactual outcomes (*e.g.*, the CATE for single-outcome equivalent models). In the next section, we study whether $\mathcal{M} \in \mathfrak{M}_{\mathcal{O}}$ and $\mathcal{R} \in \mathfrak{R}_{\mathcal{O}}$ need to, do not need to, can, or cannot be equivalent under different degrees of assumptions.

⁶There also exists a notion of *strong* counterfactual fairness [Kusner et al., 2017, Equation 4]. It demands $\mathbb{P}(Y_{T=t} = Y_{T=t'} | X = x, T = t) = 1$ for every $t, t' \in \mathcal{T}$, and is thereby a proper *cross-outcome* condition.

4 Main results

In this section, we compare the generic RCM and SCM introduced in Section 3.1 according to the three levels presented in Section 3.2. More precisely, Section 4.1 firstly reminds that equivalence does not necessarily hold whatever the level, then Section 4.2 identifies and studies cases where equivalence does (not) hold at the single-outcome level. Finally, Section 4.3 proposes conceptual and practical interpretations of these mathematical results.

4.1 Equivalence does not generally hold

We start by a crucial reminder justifying why it is relevant to compare the models at the aforementioned levels: an RCM and an SCM compatible with a same observational vector are not necessarily equivalent.

4.1.1 General case

The proposition below formalizes this claim in the most general scenario.

Proposition 1 (Equivalence is not necessary). *Let \mathcal{O} be an observational vector. For any $\mathcal{M} \in \mathfrak{M}_{\mathcal{O}}$, there exists $\mathcal{R} \in \mathfrak{R}_{\mathcal{O}}$ such that \mathcal{M} and \mathcal{R} are equivalent in none of the senses from Definition 6.*

This result rests on a simple remark: the consistency rule does not suffice to characterize potential outcomes \mathbb{P} -almost surely.⁷ More precisely, while necessarily $Y_t(\omega) = Y(\omega) = Y_{T=t}(\omega)$ on $\{T = t\}$ for $t \in \mathcal{T}$ (up to \mathbb{P} -negligible subsets) due to Lemma 3, there is no constraint on $Y_t(\omega)$ over $\Omega \setminus \{T = t\}$; it could take any value over it without violating the consistency rule. In contrast, $Y_{T=t}$ is defined (almost) everywhere through the altered SCM $\mathcal{M}_{T=t}$. The proof of Proposition 1 exploits this specification issue of potential outcomes: for any $t \in \mathcal{T}$, Y_t can be any function on $\{T \neq t\}$, thereby can be chosen distinct to $Y_{T=t}$.

Proof of Proposition 1 Let $(Y_{T=t})_{t \in \mathcal{T}}$ be the structural counterfactuals of \mathcal{M} , and define potential outcomes $(Y_t)_{t \in \mathcal{T}}$ as follows. For any $t \in \mathcal{T}$,

$$Y_t := \mathbf{1}_{\{T=t\}}Y_{T=t} + \mathbf{1}_{\{T \neq t\}}(Y_{T=t} + y),$$

where $y \in \mathbb{R}^p$ is not the null vector. The tuple $(Y_t)_{t \in \mathcal{T}}$ satisfies the consistency rule according to Lemma 3. Moreover, for any $t \in \mathcal{T}$, Y_t is not almost-surely equal nor equal in law to $Y_{T=t}$ due to $\mathbb{P}(T \neq t) > 0$ and $Y_{T=t} + y \neq Y_{T=t}$. ■

It may seem counter-intuitive to design oneself the potential outcomes (as done in the proof and in upcoming examples), since people frequently consider them as externally imposed. We emphasize that Proposition 1 is theoretically neutral, and only serves to remind that equality between counterfactual outcomes across causal models does not *necessarily* hold.

4.1.2 Causal-inference setting

Proposition 1 focuses on the most general setting where the potential outcomes meet only consistency. Said differently, it simply shows that consistency alone is not enough to guarantee equality (almost-surely or in law) between counterfactual outcomes. Nevertheless, people suppose most often that the potential outcomes also satisfy conditional ignorability in order to apply causal-inference techniques. This raises the question whether such an additional hypothesis could render the causal models equivalent. The next theorem ensures that equivalence does not always hold even under the fundamental assumptions of causal inference.

Proposition 2 (Equivalence is not necessary under the fundamental assumptions of causal inference). *Let \mathcal{O} be an observational vector satisfying positivity (Assumption 2). There exist $\mathcal{M} \in \mathfrak{M}_{\mathcal{O}}$ and $\mathcal{R} \in \mathfrak{R}_{\mathcal{O}}$ satisfying conditional ignorability (Assumption 3 or 4) such that \mathcal{M} and \mathcal{R} are equivalent in none of the senses from Definition 6.*

⁷This formally means that the set $\mathfrak{R}_{\mathcal{O}}$ is not reduced to one class of almost-surely equivalent RCMs.

This results comes from the fact that conditional ignorability fully determines the single-outcome distributions of an RCM (Lemma 2) and does not hold by design in an entailed RCM.

Proof of Proposition 2 We address the specific case where $\mathcal{T} := \{0, 1\}$ and both X and Y are \mathbb{R} -valued; one can readily generalize the proof. Consider the following SCM:

$$\begin{aligned} T &\stackrel{a.s.}{=} U_T, \\ X &\stackrel{a.s.}{=} T + U_X \\ Y &\stackrel{a.s.}{=} T + X + U_Y, \end{aligned}$$

where U_T follows a Bernoulli distribution with parameter $1/2$, while U_X and U_Y both follow the centered Gaussian distribution with unit variance, such that U_T, U_X and U_Y are mutually independent.

By do-interventions, $X_{T=t} \stackrel{a.s.}{=} t + U_X$ and $Y_{T=t} \stackrel{a.s.}{=} t + X_{T=t} + U_Y$ for any $t \in \mathcal{T}$. Therefore, by substituting the expression of $X_{T=t}$ in the expression of $Y_{T=t}$ we obtain $Y_{T=0} \stackrel{a.s.}{=} U_X + U_Y$ and $Y_{T=1} \stackrel{a.s.}{=} 2 + U_X + U_Y$. Then, we define potential outcomes as:

$$Y_0 := (1 - T) \cdot (X + U_Y) + T \cdot (X - U_Y), \quad (2)$$

$$Y_1 := (1 - T) \cdot (1 + X - U_Y) + T \cdot (1 + X + U_Y). \quad (3)$$

A key property of this construction is that $\mathcal{L}(U_Y) = \mathcal{L}(-U_Y)$ by symmetry of U_Y 's distribution, while $\mathbb{P}(U_Y = -U_Y) = 0$. After simplification,

$$Y_0 = (1 - T) \cdot (T + U_X + U_Y) + T \cdot (T + U_X - U_Y) = T + (1 - 2T) \cdot U_Y + U_X, \quad (4)$$

$$Y_1 = (1 - T) \cdot (1 + T + U_X - U_Y) + T \cdot (1 + T + U_X + U_Y) = (1 + T) - (1 - 2T) \cdot U_Y + U_X. \quad (5)$$

Let us check the required assumptions using Equations (2) and (3). Firstly, the pair (Y_0, Y_1) clearly satisfies $Y \stackrel{a.s.}{=} (1 - T) \cdot Y_0 + T \cdot Y_1$. Secondly, for $\mathcal{L}(X)$ -almost every $x \in \mathbb{R}$ and any $t \in \mathcal{T}$, $\mathbb{P}(T = t \mid X = x) = \frac{\varphi(x-t)}{\varphi(x)+\varphi(x-1)} > 0$ where φ is the density function of the centered Gaussian distribution with unit variance. Therefore, positivity holds. Thirdly, $\mathcal{L}((Y_0, Y_1) \mid X = x, T = 0) = \mathcal{L}((x + U_Y, 1 + x - U_Y) \mid X = x, T = 0)$ by definition. Next, Lemma 4 ensures that $U_Y \perp\!\!\!\perp (T, X)$, thereby $\mathcal{L}((Y_0, Y_1) \mid X = x, T = 0) = \mathcal{L}((x + U_Y, 1 + x - U_Y)) = \mathcal{L}((x - U_Y, 1 + x + U_Y))$ since $\mathcal{L}(U_Y) = \mathcal{L}(-U_Y)$. Then, notice that

$$\mathcal{L}((Y_0, Y_1) \mid X = x, T = 1) = \mathcal{L}((x - U_Y, 1 + x + U_Y)) = \mathcal{L}((Y_0, Y_1) \mid X = x, T = 0)$$

using once again $U_Y \perp\!\!\!\perp (T, X)$. Therefore, cross-outcome conditional ignorability holds. To conclude, verify from Equations (4) and (5) that $\mathbb{E}[Y_0] = 1/2$ while $\mathbb{E}[Y_{T=0}] = 0$, and $\mathbb{E}[Y_1] = 3/2$ while $\mathbb{E}[Y_{T=1}] = 2$. Therefore, $\mathcal{L}(Y_t) \neq \mathcal{L}(Y_{T=t})$ for any $t \in \mathcal{T}$.

This proves the result with an \mathcal{R} that satisfies cross-outcome conditional ignorability (Assumption 3). We now show that the same holds when only single-outcome conditional ignorability (Assumption 4) is true. To do so, we slightly modify the definition of Y_0 in the above example while keeping Y_1 identical. Set $Y_0 := X + U_Y = T + U_X + U_Y$, so that $\mathbb{E}[Y_0] = 1/2$ again. Hence, we still have that $\mathcal{L}(Y_t) \neq \mathcal{L}(Y_{T=t})$ for any $t \in \mathcal{T}$. Moreover, note that the consistency rule still holds. It remains to verify conditional ignorability.

Let us check that single-outcome conditional ignorability *does* hold. For $t \in \mathcal{T}$,

$$\mathcal{L}(Y_0 \mid X = x, T = t) = \mathcal{L}(x + U_Y)$$

$$\mathcal{L}(Y_1 \mid X = x, T = t) = (1 - t) \cdot \mathcal{L}(1 + x - U_Y) + t \cdot \mathcal{L}(1 + x + U_Y) = \mathcal{L}(1 + x + U_Y),$$

due to $U_Y \perp\!\!\!\perp (T, X)$ and $\mathcal{L}(U_Y) = \mathcal{L}(-U_Y)$. From the above expressions and the fact that $U_Y \perp\!\!\!\perp T$, it follows that $Y_t \perp\!\!\!\perp T \mid X$ for $t \in \mathcal{T}$. We turn to proving that cross-outcome conditional ignorability *does not* hold. By a similar calculus to before,

$$\mathcal{L}((Y_0, Y_1) \mid X = x, T = 0) = \mathcal{L}((x + U_Y, 1 + x - U_Y))$$

$$\mathcal{L}((Y_0, Y_1) \mid X = x, T = 1) = \mathcal{L}((x + U_Y, 1 + x + U_Y)).$$

Therefore, $\mathcal{L}((Y_0, Y_1) \mid X = x, T = 0) \neq \mathcal{L}((Y_0, Y_1) \mid X = x, T = 1)$. This concludes the proof. ■

To summarize Section 4.1: whether it be at the variable or distributional level, with or without the fundamental assumptions of causal inference, potential outcomes and structural counterfactuals are not necessarily equal. Therefore, whatever the level of counterfactual reasoning, using $(T, X, (Y_t)_{t \in \mathcal{T}})$ and $(T, X, (Y_{T=t})_{t \in \mathcal{T}})$ interchangeably (as commonly done in the scientific literature) must be either the manifestation of an arbitrary choice—a selection among all the pairs of equivalent causal models—or the mathematical consequence of further assumptions. Notably, Proposition 2 demands strong hypotheses on the potential outcomes but let the latent SCM almost unrestrained. This motivates a sharper analysis of the conditions (in particular on the SCM) that could imply equivalence or on the contrary make it impossible. This is precisely what the following subsection proposes.

4.2 Comparison of causal models in the causal-inference setting

In what follows, we aim at providing from a purely mathematical perspective a better understanding of what renders (or not) an SCM and an RCM equivalent.

4.2.1 Structural assumptions

We start by introducing additional assumptions on the underlying SCM that will simplify the comparison with an RCM. Let us denote by $U_T := U_{\text{Exo}(T)}$, $U_X := U_{\text{Exo}(X)}$ and $U_Y := U_{\text{Exo}(Y)}$ the exogenous parents of respectively T , X , and Y in the latent SCM $\mathcal{M} := \langle g, U \rangle$. These notations have already been used in the SCMs from the introductory example of Section 1.1.1 and from the proof of Proposition 2 (in the specific case where T , X , and Y were all univariate). Then, we consider the following graphical conditions on \mathcal{M} :

Assumption 5 (Outcome). \mathcal{M} is such that $Y_{\text{Endo}(T)} = Y_{\text{Endo}(X)} = \emptyset$.

Assumption 6 (Independent random noises). \mathcal{M} is such that $\text{Exo}(Y) \cap (\text{Exo}(T) \cup \text{Exo}(X)) = \emptyset$, which means that $U_Y \perp\!\!\!\perp (U_T, U_X)$.

Assumption 5 structurally defines the variable Y as the *outcome*; it changes in response to X and T but not the contrary. Through Lemma 1, it permits to write

$$\begin{cases} T & \stackrel{\text{a.s.}}{=} f_T(X, U_T), \\ X & \stackrel{\text{a.s.}}{=} f_X(T, U_X), \\ Y & \stackrel{\text{a.s.}}{=} f_Y(T, X, U_Y), \end{cases} \quad \text{and} \quad \begin{cases} T_{T=t} & \stackrel{\text{a.s.}}{=} t, \\ X_{T=t} & \stackrel{\text{a.s.}}{=} f_X(t, U_X), \\ Y_{T=t} & \stackrel{\text{a.s.}}{=} f_Y(t, X_{T=t}, U_Y), \end{cases}$$

for any $t \in \mathcal{T}$, where f_T, f_X and f_Y are measurable functions derived from g . The artificial cycle in these formulas (*i.e.*, X and T are both functions of each other) merely serves to consider all configurations of causal links between T and X (see Figure 2); strictly, \mathcal{M} satisfies Assumption 1.

Assumption 6 is a random-noise condition that notably holds in the widely-used class of SCMs referred to as Markovian, where all the noises in $(U_{\text{Exo}(i)})_{i \in \mathcal{I}}$ are mutually independent. As clarified by the following lemma, this item guarantees that all potential confounders between T and Y —except T itself—are included in X .

Lemma 4 (No hidden confounder). *Let $\mathcal{O} := (T, X, Y)$ be an observational vector. For any $\mathcal{M} \in \mathfrak{M}_{\mathcal{O}}$ satisfying Assumptions 5 and 6, if U_T, U_X and U_Y denote the exogenous parents of respectively T, X and Y in \mathcal{M} , then $U_Y \perp\!\!\!\perp (T, X)$.*

Altogether, Assumptions 5 and 6 mean that the randomness of the outcome $Y \stackrel{\text{a.s.}}{=} f_Y(T, X, U_Y)$ can be causally interpreted by three sources: the direct effect of the treatment status T , the direct effect of the covariates X , and any other possible effects U_Y independent to T and X .

4.2.2 Single-outcome comparison

Supposing that the underlying SCM satisfies the above assumptions, we derive theoretical results on the single-outcome level of equivalence between causal models. To contrast at this level an RCM with the SCM, we need a common basis to compare laws across them. The theorem below characterizes the law of (T, X, Y_t) for any $t \in \mathcal{T}$ through the latent SCM \mathcal{M} under the fundamental assumptions of causal inference, thereby enabling us to compare $\mathcal{L}((T, X, Y_t))$ with $\mathcal{L}((T, X, Y_{T=t}))$.

Theorem 1 (Single-outcome structural expression of potential outcomes). *Let $\mathcal{O} := (T, X, Y)$ be an observational vector satisfying positivity (Assumption 2), and $\mathcal{M} \in \mathfrak{M}_{\mathcal{O}}$ meet Assumptions 5 plus 6. According to Lemma 4, $U_Y \perp\!\!\!\perp (T, X)$ where U_Y denotes the exogenous parents of Y in \mathcal{M} . According to Lemma 1, there exists a measurable function f_Y such that $Y \stackrel{a.s.}{=} f_Y(T, X, U_Y)$ and $Y_{T=t} \stackrel{a.s.}{=} f_Y(t, X_{T=t}, U_Y)$ for any $t \in \mathcal{T}$. Then, for any $\mathcal{R} := (T, X, (Y_t)_{t \in \mathcal{T}}) \in \mathfrak{R}_{\mathcal{O}}$ satisfying conditional ignorability (Assumption 4),*

$$\mathcal{L}((T, X, Y_t)) = \mathcal{L}((T, X, f_Y(t, X, U_Y))),$$

for any $t \in \mathcal{T}$.

The proof combines the RCM single-outcome identification from Lemma 2 with the SCM expression of Y from Lemma 1 to connect the two causal models at the single-outcome level.

Proof of Theorem 1 Let $t \in \mathcal{T}$. Let $F \subseteq \mathcal{T} \times \mathbb{R}^d \times \mathbb{R}^p$ be a Borel set, and define the sets $F(t', x) := \{y \in \mathbb{R}^p \mid (t', x, y) \in F\}$ for every $(t', x) \in \mathcal{T} \times \mathbb{R}^d$. According to Lemma 2,

$$\mathbb{P}((T, X, Y_t) \in F) = \int \mathbb{P}(Y \in F(t', x) \mid X = x, T = t) d\mathbb{P}(X = x, T = t'). \quad (6)$$

Additionally, according to Lemma 1 the pre-intervention and post-intervention outcomes can be written as $Y \stackrel{a.s.}{=} f_Y(T, X, U_Y)$ and $Y_{T=t} \stackrel{a.s.}{=} f_Y(t, X_{T=t}, U_Y)$. Therefore, for any $t' \in \mathcal{T}$ and $\mathcal{L}(X)$ -almost any $x \in \mathbb{R}^d$,

$$\mathcal{L}(Y \mid X = x, T = t) = \mathcal{L}(f_Y(T, X, U_Y) \mid X = x, T = t) = \mathcal{L}(f_Y(t, x, U_Y) \mid X = x, T = t).$$

Also, it follows from Assumptions 5 and 6 through Lemma 4 that $U_Y \perp\!\!\!\perp (T, X)$, hence that $\mathcal{L}(f_Y(t, x, U_Y) \mid X = x, T = t) = \mathcal{L}(f_Y(t, x, U_Y) \mid X = x, T = t')$. By continuing the above computation we therefore have for any $t' \in \mathcal{T}$:

$$\mathcal{L}(Y \mid X = x, T = t) = \mathcal{L}(f_Y(t, x, U_Y) \mid X = x, T = t'). \quad (7)$$

Combining (6) with (7) leads to

$$\mathbb{P}((T, X, Y_t) \in F) = \int \mathbb{P}(f_Y(t, x, U_Y) \in F(t', x) \mid X = x, T = t') d\mathbb{P}(X = x, T = t').$$

Finally,

$$\begin{aligned} \mathbb{P}((T, X, Y_t) \in F) &= \int \mathbb{P}((T, X, f_Y(t, X, U_Y)) \in F \mid X = x, T = t') d\mathbb{P}(X = x, T = t') \\ &= \mathbb{P}((T, X, f_Y(t, X, U_Y)) \in F). \end{aligned}$$

This means that $\mathcal{L}((T, X, Y_t)) = \mathcal{L}((T, X, f_Y(t, X, U_Y)))$, which concludes the proof. \blacksquare

Remark 4 (Single-outcome expression vs. cross-outcome expression). Theorem 1 expresses for any $t \in \mathcal{T}$ the marginal distribution $\mathcal{L}((T, X, Y_t))$ with SCM-based quantities under single-outcome conditional ignorability (Assumption 4), but does not provide a similar formula for the whole joint distribution $\mathcal{L}((T, X, (Y_t)_{t \in \mathcal{T}}))$. There is an explanation: $\mathcal{L}((T, X, (Y_t)_{t \in \mathcal{T}})) \neq \mathcal{L}((T, X, (f_Y(t, X, U_Y))_{t \in \mathcal{T}}))$ in general—even under cross-outcome conditional ignorability (Assumption 3). To prove this point, we consider the same setting as the proof of Proposition 2, in which Assumptions 2, 3, 5 and 6 hold.

Using the SCM notations, $f_Y(0, X, U_Y) = X + U_Y$ and $f_Y(1, X, U_Y) = 1 + X + U_Y$. It then follows from $\mathcal{L}(-U_Y) = \mathcal{L}(U_Y)$ and Equations (2) and (3) that $\mathcal{L}(Y_0) = \mathcal{L}(f_Y(0, X, U_Y))$ and $\mathcal{L}(Y_1) = \mathcal{L}(f_Y(1, X, U_Y))$, as anticipated by Theorem 1. However, $\mathcal{L}((Y_0, Y_1)) \neq \mathcal{L}((f_Y(0, X, U_Y), f_Y(1, X, U_Y)))$. It suffices to remark that $\mathcal{L}(f_Y(1, X, U_Y) - f_Y(0, X, U_Y)) = \delta_1$ (the Dirac measure at 1) whereas $\mathcal{L}(Y_1 - Y_0) = \mathcal{L}(1 + 2(2T - 1) \cdot U_Y) \neq \delta_1$ using Equations (4) and (5).

This remark is consistent with the point made in Remark 3 that we shall prove now. Continuing the same example, we define the potential outcomes $(Y_0^{(1)}, Y_1^{(1)}) := (Y_0, Y_1)$ and $(Y_0^{(2)}, Y_1^{(2)}) := (f_Y(0, X, U_Y), f_Y(1, X, U_Y))$ along with their associated RCMs $\mathcal{R}^{(1)}$ and $\mathcal{R}^{(2)}$ in $\mathfrak{R}_{\mathcal{O}}$. Note that the pair

$(Y_0^{(2)}, Y_1^{(2)})$ clearly satisfies the consistency rule. Moreover,

$$\begin{aligned}\mathcal{L}\left((Y_0^{(2)}, Y_1^{(2)}) \mid X = x, T = 0\right) &= \mathcal{L}\left((x + U_Y, 1 + x + U_Y) \mid X = x, T = 0\right) \\ &= \mathcal{L}\left((x + U_Y, 1 + x + U_Y) \mid X = x, T = 1\right) \\ &= \mathcal{L}\left((Y_0^{(2)}, Y_1^{(2)}) \mid X = x, T = 1\right),\end{aligned}$$

due to $U_Y \perp\!\!\!\perp (T, X)$. Therefore, $\mathcal{R}^{(1)}$ and $\mathcal{R}^{(2)}$ both meet positivity and cross-outcome conditional ignorability, but are not cross-outcome equivalent.

As such, Theorem 1 enables one to compare potential outcomes and structural counterfactuals at the single-outcome level. It critically implies under the fundamental assumptions of causal inference that for any $t \in \mathcal{T}$:

$$\mathcal{L}(Y_t) = \mathcal{L}(f_Y(t, X, U_Y)) \text{ and } \mathcal{L}(Y_{T=t}) = \mathcal{L}(f_Y(t, X_{T=t}, U_Y)). \quad (8)$$

This result sheds light on Proposition 2: Y_t and $Y_{T=t}$ are not necessarily equal in law since $\mathcal{L}(X) \neq \mathcal{L}(X_{T=t})$ in general. Moreover, it furnishes sufficient conditions on the latent SCM for single-outcome equivalence to hold. Observe that the probability distributions from Equation (8) are equal if: (1) X is not altered by do-interventions on T , or (2) Y is not impacted by X . The following assumption captures scenario (1), which is the most relevant in causal inference.

Assumption 7 (No posttreatment covariate). \mathcal{M} is such that, $T_{\text{Endo}(X)} = \emptyset$.

Assumption 7 together with Assumption 5 mean that X contains no descendant of T in \mathcal{M} . It implies that $X_{T=t} \stackrel{\text{a.s.}}{=} X$ for any $t \in \mathcal{T}$, which guarantees single-outcome equivalence under the assumptions of Theorem 1.

Corollary 1 (Single-outcome equivalence under no posttreatment covariate). *Let \mathcal{O} be an observational vector satisfying positivity (Assumption 2), $\mathcal{M} \in \mathfrak{M}_{\mathcal{O}}$ meet Assumptions 5 plus 6, and $\mathcal{R} \in \mathfrak{R}_{\mathcal{O}}$ satisfy single-outcome conditional ignorability (Assumption 4). If Assumption 7 holds, then \mathcal{M} and \mathcal{R} are single-outcome equivalent but not necessarily cross-outcome equivalent, even under cross-outcome conditional ignorability (Assumption 3).*

Proof of Corollary 1 Let $t \in \mathcal{T}$. Under the considered assumptions, we have according to Theorem 1: $\mathcal{L}((T, X, Y_t)) = \mathcal{L}((T, X, f_Y(t, X, U_Y)))$. Moreover, recall that $Y_{T=t} \stackrel{\text{a.s.}}{=} f_Y(t, X_{T=t}, U_Y)$. If X contains no descendant of T in \mathcal{M} , then $X_{T=t} \stackrel{\text{a.s.}}{=} X$, leading to $\mathcal{L}((T, X, Y_t)) = \mathcal{L}((T, X, Y_{T=t}))$.

We now turn to proving that cross-outcome equivalence does not necessarily hold, even under cross-outcome conditional ignorability. To this end, we consider the following SCM:

$$\begin{aligned}T &\stackrel{\text{a.s.}}{=} U_T, \\ X &\stackrel{\text{a.s.}}{=} U_X, \\ Y &\stackrel{\text{a.s.}}{=} T + X + U_Y,\end{aligned}$$

where U_T follows a Bernoulli distribution with parameter $1/2$, U_X is any \mathbb{R} -valued random variable, and U_Y follows a centered Gaussian distribution with unit variance, such that U_T, U_X and U_Y are mutually independent. This model satisfies Assumptions 5, 6 and 7 by construction. By do-interventions, $Y_{T=0} \stackrel{\text{a.s.}}{=} U_X + U_Y$ and $Y_{T=1} = 1 + U_X + U_Y$. Next, define the potential outcomes as follows:

$$\begin{aligned}Y_0 &:= (1 - T) \cdot (X + U_Y) + T \cdot (X - U_Y) = U_X + (1 - 2T) \cdot U_Y, \\ Y_1 &:= (1 - T) \cdot (1 + X - U_Y) + T \cdot (1 + X + U_Y) = 1 + U_X - (1 - 2T) \cdot U_Y.\end{aligned}$$

Let us verify the required assumptions. Firstly, the potential outcomes clearly satisfy consistency, that is $Y = (1 - T) \cdot Y_0 + T \cdot Y_1$. Secondly, since $T \perp\!\!\!\perp X$, we have $\mathbb{P}(T = 1 \mid X) = \mathbb{P}(T = 1) = 1/2$ which entails positivity. Thirdly, $\mathcal{L}((Y_0, Y_1) \mid X = x, T = 1) = \mathcal{L}((Y_0, Y_1) \mid T = 1)$ because $U_Y \perp\!\!\!\perp (T, X) \stackrel{\text{a.s.}}{=} (U_T, U_X)$. Additionally, $\mathcal{L}((Y_0, Y_1) \mid T = 1) = \mathcal{L}((U_X - U_Y, 1 + U_X + U_Y)) = \mathcal{L}((U_X + U_Y, 1 + U_X - U_Y))$ since $\mathcal{L}(U_Y) = \mathcal{L}(-U_Y)$, and $\mathcal{L}((U_X + U_Y, 1 + U_X - U_Y)) = \mathcal{L}((Y_0, Y_1) \mid T = 0) = \mathcal{L}((Y_0, Y_1) \mid X = x, T = 0)$ using again $U_Y \perp\!\!\!\perp (T, X)$. Wrapping up: $\mathcal{L}((Y_0, Y_1) \mid X = x, T = 1) = \mathcal{L}((Y_0, Y_1) \mid X = x, T = 0)$, meaning that cross-outcome conditional ignorability holds according to the first part of the proof.

Therefore, the causal models are single-outcome equivalent. To conclude that they are not cross-outcome equivalent, note that $\mathcal{L}(Y_1 - Y_0) = \mathcal{L}(1 + 2(2T - 1) \cdot U_Y)$ whereas $\mathcal{L}(Y_{T=1} - Y_{T=0}) = \delta_1$. ■

This result specifies a large class of SCMs that are single-outcome equivalent to an RCM meeting positivity and conditional ignorability: those in which Y is a conditionally unconfounded outcome (Assumptions 5 plus 6) and T does not impact X (Assumption 7). If one of these two sets of hypotheses does not hold, then single-outcome equivalence is not guaranteed. We emphasize that Assumption 7 is the most critical; Assumptions 5 plus 6 mainly serve to simplify the analysis (see Remark 6). Notably, if Assumptions 5 and 6 hold but not Assumption 7, then $\mathbb{P}(X_{T=t} \neq X) > 0$ for some $t \in \mathcal{T}$, and it follows from Equation (8) that the invariance $\mathcal{L}(Y_t) = \mathcal{L}(Y_{T=t})$ for any $t \in \mathcal{T}$ basically occurs in the useless scenarios where f_Y does not really change in response to its X -input. Therefore, excluding these pathological cases, Assumption 7 through Theorem 1 fully classifies single-outcome equivalence between an SCM satisfying Assumptions 5 plus 6 and an RCM under the standard causal-inference regime.⁸ We now turn to contextualizing this abstract comparison within real-world problems.

4.3 Interpretation

Section 4.2 explained under which mathematical conditions potential outcomes $(Y_t)_{t \in \mathcal{T}}$ in the causal-inference setting and structural counterfactuals $(Y_{T=t})_{t \in \mathcal{T}}$ define—or not—the same single-outcome causal estimands. In this section, we discuss these formal results from practical and conceptual perspectives, and exemplify the meaning of reasoning counterfactually with nonequivalent models.

4.3.1 A practical perspective on cases of (non)equivalence

Up until now, we adopted an analytical viewpoint where the models were abstract mathematical objects. However, an SCM is not a mere detached tool; it is meant to genuinely capture the world’s functioning. As such, practitioners do not decide the graphical relationships between endogenous variables themselves, they are imposed by Nature. In contrast, an RCM introduces hypothetical counterfactual variables whose interpretation may depend on the assumptions placed upon them. This critically signifies that single-outcome equivalence between the *true* latent SCM and a potential-outcome model satisfying conditional ignorability does not generically hold in many real-world problems, due to the dichotomy ruled by Assumption 7.

To clarify this point, consider a problem where one relies on an RCM $\mathcal{R} \in \mathfrak{R}_{\mathcal{O}}$ equipped with positivity and conditional ignorability to apply standard causal-inference techniques. They wonder whether this model is equivalent to the underlying SCM $\mathcal{M} \in \mathfrak{M}_{\mathcal{O}}$, which is supposed to meet Assumptions 5 and 6. Answering this question amounts to translating the pivotal causal-ordering assumption between T and X (Assumption 7) and its negation into common language.

Assumption 7 basically signifies that the covariates are not altered by the treatment, as illustrated in Figure 2a. Notably, this configuration encompasses various typical causal-inference scenarios: in clinical trials, the covariates X (sometimes called *pretreatment* variables) may influence the treatment allocation T but never the contrary. In these common situations, both the RCM and the SCM produce the same single-outcome counterfactuals due to Corollary 1. In particular, $\mathbb{E}[Y_1 - Y_0 \mid X = x] = \mathbb{E}[Y_{T=1} - Y_{T=0} \mid X = x]$. We crucially remind that equivalence does not necessarily hold at the cross-outcome level (as also stated in Corollary 1).

Importantly, the negation of Assumption 7—which states that T is a parent of several (or even all) covariates in \mathcal{M} —is also mathematically possible. Figures 2b and 2c illustrate the possible causal graphs. In these situations, Theorem 1 does not guarantee single-outcome equivalence. Consequently, confusion between the two causal approaches can lead to misleading results: according to Equation (8), the RCM considers counterfactual outcomes at fixed X , whereas the SCM alters the covariates into $X_{T=t}$. These cases are empirically relevant, since people also rely on causal inference outside the scope of clinical trials, in settings where the treatment impacts the covariates. For example, T drives X but not the contrary (as in Figure 2c) in emblematic causal problems such as the Berkeley’s admission paradox where T represents the sex and X the course choice [Bickel et al., 1975]. This is more generally true

⁸A similar “almost-true” equivalence between a graphical condition like Assumption 7 and its structural implications was already discussed in [Pearl, 2009, Section 7.4], where the so-called graphical and counterfactual criteria of exogeneity coincide after excluding incidental cases.

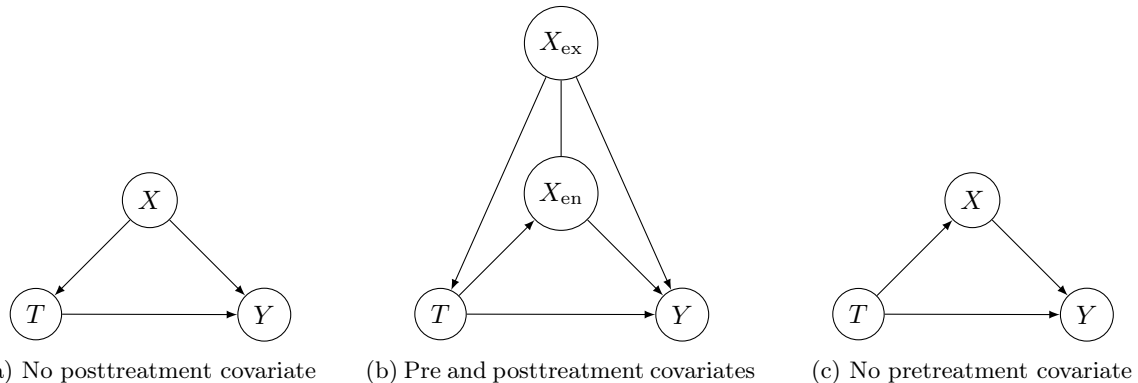


Figure 2: Three possible configurations of the treatment in \mathcal{M} under Assumption 5. $X_{ex} := X_{\text{Endo}(T)}$ denotes the parents of T in X while X_{en} are the remaining covariates. Exogenous variables are not represented but Assumption 6 holds. A single node can represent several variables. In (a), T does not impact X (Assumption 7); in (b), T may impact some X -variables and some X -variables may impact T ; in (c), X does not impact T .

in the causal-fairness literature, where the variable to alter typically encodes an intrinsic feature like the sex, the race, or the age of individuals (see for instance [Kusner et al., 2017, Chiappa and Isaac, 2019, Plecko and Meinshausen, 2020, Nilforoshan et al., 2022, Barocas et al., 2023] for machine-learning-related research). Other scenarios, notably motivated by epidemiological research, include both pre and posttreatment covariates (as in Figure 2b) [VanderWeele et al., 2014]. In these cases, the treatment often represents an unstable characteristic of individuals like a habit (smoking) or a fluctuating biological factor (obesity). Section 4.3.3 and Appendix A provide detailed illustrations of real-world-inspired problems where Assumption 7 does not hold.

Remark 5 (No causation without manipulation). The “no causation without manipulation” principle of Holland [1986] states that the causal influence of immutable variables like a person’s race is ill-defined since we cannot conceptualize an assignment mechanism for such a nonmanipulable status (in contrast, for example, to allocating a medical treatment). There is a long-standing debate among causality scholars on this topic. Some agree that this principle should be a general rule [Winship and Morgan, 1999, Freedman, 2004, Berk, 2004, Greiner and Rubin, 2011], other argue that addressing nonmanipulable causes is scientifically relevant and proposed interpretations of this practice [VanderWeele and Robinson, 2014, Glymour and Spiegelman, 2017, Pearl, 2018, 2019]. This question relates to the current discussion around Assumption 7, illustrated by Figure 2, at two levels.

The first level is independent of the employed causal models. Manipulability is an experimental concept that does not have a definite formal translation. In contrast, Assumption 7 is a formal condition on the causal ordering between endogenous variables. As such manipulability and causal ordering do not overlap, but they may intersect when approaching problems from a real-world standpoint, since manipulable and immutable causes usually conform to different graphical configurations in practice. As aforementioned, a *manipulable* medical treatment commonly corresponds to Figure 2a, an *immutable* attribute like race commonly corresponds to Figure 2c, an *unstable* characteristic like being a smoker commonly corresponds to Figure 2b.⁹ As such, some people consider that the “no causation without manipulation” maxim implies that the covariates must be fixed, pretreatment variables only [Imbens, 2020, Section 4.1]. According to someone who follows this implication, it does not make sense to tackle problems corresponding to Figures 2b and 2c, regardless of whether they work with an RCM or an SCM.

The second level opposes RCMs and SCMs. From a purely mathematical perspective, both causal approaches accept immutable causes: an RCM poses no formal constraints on its treatment variable, and any endogenous variable of an SCM—whatever its graphical position—is eligible to do-interventions. However, RCM analysts have traditionally precluded pretreatment covariates from their models [Imbens and Rubin, 2015], while SCM analysts have frequently applied do-interventions on variables with no endogenous par-

⁹It is usual but not necessary: an arguably nonmanipulable feature like height can be positioned as T in Figure 2a for being a responder of sex; an arguably manipulable feature like chocolate consumption can be placed as T in Figure 2a for having no determinant among the other studied endogenous variables.

ents [Kusner et al., 2017, Chiappa, 2019, Plecko and Meinshausen, 2020]. According to someone who deems that “no posttreatment covariate” specifically constrains RCMs, it does not make sense to compare (as we did) an RCM with the latent SCM in the settings of Figures 2b and 2c.

We emphasize that our work does not focus on philosophical arguments: it neutrally analyzes the literature. In particular, we do not contribute to the general debate on manipulability and do not argue in favor or against including implications of Holland’s principle specifically in Rubin’s framework. From this angle, we justify studying situations with posttreatment covariates and even comparing RCMs and SCMs in such cases for theoretical and empirical reasons. Firstly, as aforementioned, there is nothing that mathematically forbids to study these cases, be it with an RCM or an RCM. Therefore, addressing them is necessary to completely compare the two causal approaches. Secondly, many researchers have actually relied on (or referred to) the potential-outcome framework to understand the influence of immutable variables like sex, race, or biological factors [Li et al., 2017, Glymour and Spiegelman, 2017, Khademi et al., 2019, Khademi and Honavar, 2020, Qureshi et al., 2020], typically placed like the treatment in Figure 2c.¹⁰ Moreover, some references have presented side by side Rubin’s approach and Pearl’s approach specifically in such contexts, notably in fairness problems [Barocas et al., 2023, Makhlof et al., 2024]. If we consider this corpus of the causal-inference literature to be admissible, then for the sake of empirical relevance we must also compare RCMs and SCMs in cases possibly disobeying “no causation without manipulation” or Assumption 7.

All in all, under the fundamental assumptions of causal inference, equivalence of single-outcome counterfactuals across the two causal models depends on the relationships between the treatment and the covariates, as described in Figure 2. What typically distinguishes the different configurations is the nature of the so-called treatment. If the treatment can be assigned *a posteriori* (as in Figure 2a), then the two notions of counterfactuals coincide. Otherwise (as in Figures 2b and 2c) structural counterfactuals and potential-outcome counterfactuals are generically not equal. The concrete example from Section 1.1.1, where the SCM fits Figure 2c and the potential outcomes of the RCM produce different causal estimands than structural counterfactuals, epitomizes this point.

4.3.2 A conceptual perspective on cases of (non)equivalence

We now advance from the practical standpoint to a conceptual perspective. The fact that \mathcal{R} and $\mathcal{R}_{\mathcal{M}}$ are not necessarily single-outcome equivalent under the assumptions of Theorem 1 can be interpreted as follows: $(Y_t)_{t \in \mathcal{T}}$ meeting conditional ignorability and $(Y_{T=t})_{t \in \mathcal{T}}$ obtained by do-interventions of T do not carry the same counterfactual semantic. Said differently, their “had T been equal to t ” do not mean the same thing. This can similarly be understood as: $(Y_t)_{t \in \mathcal{T}}$ results from formally distinct interventions than *do*-interventions on T . Beyond simply observing that these two approaches may produce distinct outcomes, understanding their respective significations is crucial, since practitioners must be able to explain and justify the estimands they aim to compute. We address this point by reminding some basics of counterfactual reasoning in a broad sense.

Counterfactual reasoning can be defined as thinking about outcomes in hypothetical worlds where some circumstances change from what factually happened while others are kept equal. Crucially, there is not a single way of reasoning counterfactually.¹¹ Each way notably depends on what is kept equal across worlds. Theorem 1 outlines the difference in this regard between the potential-outcome framework under the fundamental assumptions of causal inference and the structural account of counterfactuals. Recall that under the required assumptions, for every $t \in \mathcal{T}$

$$\begin{aligned} Y &\stackrel{a.s.}{=} f_Y(T, X, U_Y) \stackrel{a.s.}{=} f_Y(T, f_X(T, U_X), U_Y), \\ Y_{T=t} &\stackrel{a.s.}{=} f_Y(t, X_{T=t}, U_Y) \stackrel{a.s.}{=} f_Y(t, f_X(t, U_X), U_Y), \\ \mathcal{L}((T, X, Y_t)) &= \mathcal{L}((T, X, f_Y(t, X, U_Y))). \end{aligned}$$

Compared to the factual outcome Y , structural counterfactuals differ in T but share the same background factors (U_X, U_Y) in an almost-sure sense, while potential outcomes differ in T but share the same

¹⁰Other articles, for instance [Bertrand and Mullainathan, 2004, Ridgeway, 2006, Gaebler et al., 2022], explicitly focus on sex and race as *perceived* by a decider. Such a perception could depend on the covariates and not be immutable.

¹¹This nonuniqueness can be appreciated through the seminal work of Lewis on counterfactual conditionals [Lewis, 1973a,b, 1979]. According to Lewis, the verification of counterfactual statements depends on the arbitrary choice of a similarity relation.

input variables (X, U_Y) in a distributional sense. Said differently, up to measure-theoretic technicalities, potential-outcome counterfactuals are *ceteris paribus* counterfactuals (*i.e.*, all other things being kept equal) with respect to all input variables—notably the covariates—whereas structural counterfactuals are *ceteris paribus* counterfactuals with respect to all background factors—but *mutatis mutandis* (*i.e.*, after changing what must be changed) with respect to the covariates. According to Corollary 1, these two conceptions of counterfactuals *happen* to coincide in settings like Figure 2a where there is no posttreatment covariate. But essentially, they are different.

We emphasize that both definitions of counterfactuals are perfectly legitimate and bear causal interpretations. However, they convey distinct meanings and thereby correspond to different causal effects. Therefore, *they should not be employed for the same purpose*. Let us illustrate this aspect on a concrete case.

4.3.3 Illustration: an immutable treatment and two different kinds of counterfactuals

We conclude Section 4 by exemplifying the mathematical, practical, and conceptual considerations of the (non)equivalence results from Section 4.2.2. To do so, we compare RCMs and SCMs under the assumptions of Theorem 1 in concrete examples that do not satisfy Assumption 7 so that single-outcome equivalence does not hold. More precisely, in what follows we further study the motivating example from Section 1.1.1 which corresponds to Figure 2c. For concision, we defer to Appendix A the analysis of a scenario adapted to Figure 2b.

We firstly remind the example’s setting while properly defining the considered causal models. The treatment status T indicates the gender; the covariate X quantifies the level of work experience; the outcome Y evaluates a candidate’s application for some position. Suppose that the associated observational vector $\mathcal{O} := (T, X, Y)$ is ruled by the following SCM $\mathcal{M} \in \mathfrak{M}_{\mathcal{O}}$:

$$\begin{aligned} T &\stackrel{a.s.}{=} U_T, \\ X &\stackrel{a.s.}{=} \alpha T + U_X, \\ Y &\stackrel{a.s.}{=} X + \beta T + U_Y, \end{aligned}$$

where α and β are deterministic parameters. This model satisfies Assumption 5 and the negation of Assumption 7 by design. Moreover, we suppose that positivity (Assumption 2) is true (for instance by choosing U_X inducing a Gaussian distribution), and that $U_Y \perp\!\!\!\perp (U_T, U_X)$ so that Assumption 6 holds. Finally, we set two potential outcomes (Y_0, Y_1) meeting the consistency rule, that is $Y \stackrel{a.s.}{=} (1-T) \cdot Y_0 + T \cdot Y_1$, and single-outcome conditional ignorability (Assumption 4). This defines $\mathcal{R} := (T, X, Y_0, Y_1) \in \mathfrak{R}_{\mathcal{O}}$. Note that all the assumptions of Theorem 1 are satisfied, and that the SCM fits Figure 2c.

We recall the estimands that an analyst working with \mathcal{R} and an analyst working with $\mathcal{R}_{\mathcal{M}}$ respectively obtain for the average causal effect of T onto Y conditional to $X = x$. The former analyst gets the following CATE based on potential outcomes:

$$\begin{aligned} \text{CATE}_{\mathcal{R}}(x) &:= \mathbb{E}[Y_1 - Y_0 \mid X = x] \\ &= \mathbb{E}[(X + \beta + U_Y) - (X + U_Y) \mid X = x] \\ &= \beta. \end{aligned}$$

This time, we applied Theorem 1 in the first step of the above calculus to illustrate that it gives the same results as using Lemma 2 plus $U_Y \perp\!\!\!\perp (T, X)$ like in Section 1.1.1. The latter analyst gets the following CATE based on structural counterfactuals:

$$\begin{aligned} \text{CATE}_{\mathcal{R}_{\mathcal{M}}}(x) &:= \mathbb{E}[Y_{T=1} - Y_{T=0} \mid X = x] \\ &= \mathbb{E}[(\alpha + U_X + \beta + U_Y) - (U_X + U_Y) \mid X = x] \\ &= \alpha + \beta. \end{aligned}$$

As noticed in Section 1.1.1, $\text{CATE}_{\mathcal{R}} \neq \text{CATE}_{\mathcal{R}_{\mathcal{M}}}$ if $\alpha \neq 0$. Because these estimands are single-outcome effects, we conclude that \mathcal{R} and \mathcal{M} are not single-outcome equivalent, as expected by the graphical discussion from Section 4.3.1. Similar conclusions hold in settings with pre and posttreatment covariates (Figure 2b), as demonstrated in Appendix A.

Let us interpret in more details the obtained CATEs through the conceptual lens of Section 4.3.2. Observe that $\text{CATE}_{\mathcal{R}}$ measures only the *direct effect* of the treatment in \mathcal{M} : it completely ignores the dependence of Y on T through X , as it involves only β . In contrast, remark that $\text{CATE}_{\mathcal{R}_{\mathcal{M}}}$ measures the *total effect* of the treatment: it takes into account the whole path of influence of T onto Y , involving both α and β . This is due to the *ceteris paribus/mutatis mutandis* significations of the associated counterfactual variables: by holding X fixed, (Y_0, Y_1) blocks information that does not flow *directly* from T ; by letting X change, $(Y_{T=0}, Y_{T=1})$ captures also information that flows *indirectly* from T through X . From a fairness perspective, $\text{CATE}_{\mathcal{R}}$ says that if $\beta = 0$, that is if T is not a *direct* cause of Y , then the application process is fair; whether it is unfair towards men or women when $\beta \neq 0$ depends on the sign of β . In contrast, $\text{CATE}_{\mathcal{R}_{\mathcal{M}}}$ says that if $\beta = -\alpha$, that is if the decision rule Y compensates the discrepancy of work experiences X across genders T , then the application process is fair. Each analysis points out a different notion of fairness: considering $\text{CATE}_{\mathcal{R}_{\mathcal{M}}}$ as a fairness criterion suggests that recruiters should correct societal inequalities by preferring women with potentially lower work experience but higher merit whereas relying on the $\text{CATE}_{\mathcal{R}}$ suggests it is only explicitly including the gender in the decision-rule pipeline that is unfair.¹² Critically, if $\alpha \neq 0$, *practitioners mixing potential outcomes with structural counterfactuals could reach contradictory conclusions on fairness*. This is why having a clear understanding of the semantic carried by the employed counterfactual outcomes is crucial. We conclude the section with additional remarks raised by this example.

Remark 6 (About the structural assumptions). At the end of Section 4.2.2, we argued that the equivalence between an RCM satisfying the fundamental assumptions of causal inference and the latent SCM was mostly governed by Assumption 7 (no posttreatment covariate), while Assumptions 5 and 6 mainly served to derive general results. Let us illustrate this point on the above example.

To emphasize the respective roles of Assumptions 5, 6 and 7, we consider the same setting but we do not assume $U_Y \perp (U_X, U_T)$ anymore. Consequently, Assumption 6 does not necessarily hold and U_Y may depend on (T, X) . This relaxation does not affect $\text{CATE}_{\mathcal{R}_{\mathcal{M}}}$ which remains equal to $\alpha + \beta$. As Theorem 1 cannot be employed, we compute $\text{CATE}_{\mathcal{R}}$ via Lemma 2. This gives:

$$\begin{aligned} \text{CATE}_{\mathcal{R}}(x) &:= \mathbb{E}[Y_1 - Y_0 \mid X = x] \\ &= \mathbb{E}[Y \mid X = x, T = 1] - \mathbb{E}[Y \mid X = x, T = 0] \\ &= \mathbb{E}[x + \beta + U_Y \mid X = x, T = 1] - \mathbb{E}[x + U_Y \mid X = x, T = 0] \\ &= \beta + \mathbb{E}[U_Y \mid X = x, T = 1] - \mathbb{E}[U_Y \mid X = x, T = 0]. \end{aligned}$$

In comparison to before, the difference $\mathbb{E}[U_Y \mid X = x, T = 1] - \mathbb{E}[U_Y \mid X = x, T = 0]$ can be nonzero. Remark that, except in the very specific case where $\mathbb{E}[U_Y \mid X = x, T = 1] - \mathbb{E}[U_Y \mid X = x, T = 0] = \alpha$ for $\mathcal{L}(X)$ -almost every x , we still have $\text{CATE}_{\mathcal{R}} \neq \text{CATE}_{\mathcal{R}_{\mathcal{M}}}$. This means that \mathcal{R} and $\mathcal{R}_{\mathcal{M}}$ are not single-outcome equivalent when even when Assumptions 5 and 6 do not hold.

Remark 7 (Computing direct effects from an SCM). One can still define and compute $\text{CATE}_{\mathcal{R}}$, that is the *direct effect*, using an SCM \mathcal{M} . We propose two different expressions of this quantity with *structural* terms in the setup of the fairness illustration,

$$\begin{aligned} \text{CATE}_{\mathcal{R}}(x) &= \mathbb{E}[(x + \beta + U_Y) - (x + U_Y)], \\ &= \mathbb{E}[Y_{T=1, X=x} - Y_{T=0, X=x}], \end{aligned} \tag{9}$$

$$= \mathbb{E}[Y_{T=1} \mid X_{T=1} = x] - \mathbb{E}[Y_{T=0} \mid X_{T=0} = x]. \tag{10}$$

Leveraging \mathcal{R} instead of \mathcal{M} permits to compute the direct effect without specifying \mathcal{M} . In contrast, one cannot always define and compute the *total effect* $\text{CATE}_{\mathcal{R}_{\mathcal{M}}}$ from \mathcal{R} . More generally, by expressing the single-outcome expression of \mathcal{R} in terms of \mathcal{M} , Theorem 1 enables one to find SCM-based formulas of effects derived from \mathcal{R} under the fundamental assumptions of causal inference.

This remark also serves to precise a key message of the paper. Note that Equation (9) employs interventions on (T, X) , while Equation (10) employs interventions solely on T but involves the post-intervention covariates. Critically, $\text{CATE}_{\mathcal{R}}$ cannot be written with $(T, X, Y_{T=0}, Y_{T=1})$ only. This means that $\text{CATE}_{\mathcal{R}}$ is identifiable in \mathcal{M} but *not* in $\mathcal{R}_{\mathcal{M}}$. Therefore, one should not conclude from our results that the effects derived from an RCM \mathcal{R} in the causal-inference setting are in general incompatible with

¹²Such a differentiation is closely related to the notion of *path-specific counterfactual fairness* [Chiappa, 2019].

the ones derived from the latent SCM \mathcal{M} . What we show is that swapping the respective counterfactual-outcome variables relative to T of these models, which amounts to exchanging \mathcal{R} and $\mathcal{R}_{\mathcal{M}}$, generally defines distinct estimands. In other words, our work highlights a difference between potential outcomes meeting classical assumptions and do-interventions on T in \mathcal{M} , rather than between RCMs and SCMs (see also Remark 10).

To sum-up, each type of counterfactual variables has a different causal signification in this immutable-treatment configuration, and therefore produces different causal estimands. This signifies that the difference between frameworks does not amount to methodological considerations only. Analysts and researchers should also justify the chosen model and its assumptions depending on the kind of causal effects they want to compute. In the final section, we include the conceptual and practical divergences between models pointed out by this illustration into a more comprehensive reflection.

5 Discussion

This section examines common practices in the causal-inference literature in light of the mathematical differences between models explained in Section 4. Firstly, Section 5.1 clarifies the relation between our contribution and the notorious formal equivalence between frameworks, unveiling a fundamental dichotomy in the applications of potential outcomes. Secondly, Section 5.2 leverages this discussion to provide recommendations on the exchange of potential-outcome and do notations.

5.1 On the formal equivalence between frameworks

As mentioned in the introduction, many articles interchangeably use the potential-outcome notation and the do notation, invoking an “equivalence” between causal frameworks. This may seem paradoxical after reading Section 4. In the following, we clarify this aspect. Then, we explain what a conflict between models like Section 4.3.3 signifies to the relationship between SCMs and RCMs, and to how people generally manipulate these models.

5.1.1 Structural representation and graphical translation

Recall that every $\mathcal{M} \in \mathfrak{M}_{\mathcal{O}}$ entails an RCM $\mathcal{R}_{\mathcal{M}} \in \mathfrak{R}_{\mathcal{O}}$ according to Lemma 3. Conversely, any $\mathcal{R} \in \mathfrak{R}_{\mathcal{O}}$ can be *represented* by an SCM.

Proposition 3 (Structural representation). *Let \mathcal{O} be an observational vector. For any $\mathcal{R} \in \mathfrak{R}_{\mathcal{O}}$, there exists $\mathcal{M}_{\mathcal{R}} \in \mathfrak{M}_{\mathcal{O}}$ such that $\mathcal{M}_{\mathcal{R}}$ and \mathcal{R} are almost-surely equivalent. We say that $\mathcal{M}_{\mathcal{R}}$ is a structural representation of \mathcal{R} .*

This result can be seen as a variant of [Ibeling and Icard, 2024, Proposition 1] in our specific setting. The expression *structural representation* is inspired from it.

Interestingly, if one chooses to define potential outcomes as structural counterfactuals from a structural representation, then Rubin’s causal framework and Pearl’s causal framework become two different languages to talk about the same objects. In the RCM, assumptions for causal inference are generally framed as conditional-independence restrictions on counterfactual variables (like Assumptions 3 and 4); in Pearl’s causal framework, assumptions on causal relationships are generally framed in terms of graphical conditions on factual variables (like Assumptions 5, 6 and 7). Both [Pearl, 2009, Chapter 7] and [Richardson and Robins, 2013] focus on unifying these two mathematical languages by providing rules for translating assumptions and theorems from one viewpoint to the other. This ensures what people often refer as the *logical* or *formal equivalence* between frameworks. We crucially emphasize that, in contrast to the notions in Definition 6, this corresponds to an equivalence between formalisms—not between given causal models. This formal equivalence notably allows analysts to work symbiotically with an RCM \mathcal{R} and a structural representation $\mathcal{M}_{\mathcal{R}}$, or with an SCM \mathcal{M} and its entailed RCM $\mathcal{R}_{\mathcal{M}}$.

Remark 8 (Graphical interpretation of conditional ignorability). A classic translation rule is [Pearl et al., 2016, Theorem 4.3.1], which states that if X meets the *backdoor criterion* relative to (T, Y) in \mathcal{M} [Pearl, 2009, Definition 3.3.1], then $Y_{T=t} \perp\!\!\!\perp T \mid X$ for all $t \in \mathcal{T}$: said differently, single-outcome conditional ignorability holds in $\mathcal{R}_{\mathcal{M}}$. This counterfactual interpretation of graphical conditions echoes Corollary 1.

Our result states that if an SCM satisfies Assumptions 5, 6 and 7, then it is single-outcome equivalent to an RCM such that $Y_t \perp\!\!\!\perp T \mid X$ for all $t \in \mathcal{T}$ (Assumption 4). The connection between [Pearl et al., 2016, Theorem 4.3.1] and Corollary 1 follows from the fact that Assumptions 5, 6 and 7 (illustrated by Figure 2a) entail that X meets the backdoor criterion relative to (T, Y) in \mathcal{M} .

Interestingly, the two results differ from their perspectives. Corollary 1 reframes Pearl’s translation rule in terms of equivalence between presumably *separated* SCMs and RCMs. This change of viewpoint helps understanding some unappreciated aspects of the formal equivalence, as discussed below.

5.1.2 Two aspects of the formal equivalence

There is no logical contradiction between the results from Section 4 and this formal equivalence. Nevertheless, our contribution highlights two overlooked features.

Firstly, Proposition 3 or [Ibeling and Icard, 2024, Proposition 1] simply *allow* to represent an RCM by an SCM. Ultimately, defining potential outcomes as structural counterfactuals is a *choice*—it does not rest on any proof. When Pearl writes $Y_t := Y_{T=t}$ in [Pearl, 2009, Equation 3.51], claiming that the operation $\text{do}(T = t)$ on the SCM gives a physical meaning to the vague “had T been t ” of the potential outcome, this is an arbitrary choice. As demonstrated by Propositions 1 and 2, naively looking at the definitions of the Rubin’s causal framework and of Pearl’s causal framework, there is nothing that mathematically constrains potential outcomes to coincide at any level with the structural counterfactuals of a chosen SCM.¹³ What [Pearl, 2009, Chapter 7] and [Richardson and Robins, 2013] really show is that *if* potential outcomes are chosen to be structural counterfactuals, then one can translate the assumptions made on an underlying causal graph into assumptions on potential outcomes—not that they *must* be chosen as such.

Secondly, even though there always exist *theoretical* structural representations of a given RCM \mathcal{R} , nothing guarantees that they correctly capture the real-world causal dependencies. In other words, the true SCM \mathcal{M} may not represent \mathcal{R} . Notably, Section 4.3.1 specifies situations where an RCM \mathcal{R} tailored to causal inference is not single-outcome equivalent to the true SCM \mathcal{M} (as illustrated in Section 4.3.3), which implies that it cannot be represented by \mathcal{M} .

The fact that (1) using the formal equivalence is a choice and (2) applying the formal equivalence on a predefined RCM may produce an SCM conflicting with the true one has important consequences on its practicality that we discuss next.

5.1.3 Two fundamentally distinct paradigms for potential outcomes

An SCM is meant to correctly capture the cause-effect links between the variables of interest. This is why no analyst deliberately works with an SCM that is apparently wrong. In particular, it does not make sense to use an RCM \mathcal{R} in synergy with a structural representation if the true SCM \mathcal{M} itself does not represent \mathcal{R} . Who would axiomatize potential outcomes via do-interventions in a fake SCM? Under this principle, two paradigms for defining and applying a potential-outcome model \mathcal{R} coexist.

On the one hand, Pearl has firmly advocated for long to *always* use the potential-outcome framework in symbiosis with an SCM. This rule induces a first paradigm which amounts to defining the RCM as $\mathcal{R} := \mathcal{R}_{\mathcal{M}}$ where \mathcal{M} is the *true* SCM. This means letting $(Y_t)_{t \in \mathcal{T}} = (Y_{T=t})_{t \in \mathcal{T}}$. In this approach, properties of the potential outcomes (like conditional ignorability) do not come from fundamental assumptions but follow from hypotheses made on \mathcal{M} , generally framed as graphical conditions (recall Remark 8). Note that this forbids the practice from Section 4.3.3 where we employed \mathcal{R} satisfying conditional ignorability whereas Assumption 7 did not hold in \mathcal{M} . According to Pearl, a key interest of this approach comes from the fact that conditional-independence assumptions of RCMs are hard to interpret, while the causal graph and structural equations of an SCM form an intelligible formalization of scientific knowledge from which such assumptions can be derived and justified. We refer to [Pearl, 2009, Section 7.4] for his detailed argumentation. Sections 4.3.2 and 5.1 underline that following this paradigm consists in making a specific choice regarding the hypothetical interventions defining the potential outcomes, thereby regarding their counterfactual semantics: it formally defines their “had T been t ” as the operation $\text{do}(T = t)$ in \mathcal{M} which changes the remaining endogenous variables accordingly.

¹³From a more philosophical angle, [Markus, 2021, Section 2.1] made a similar remark to argue that the two frameworks were *weakly* equivalent rather than *strongly* equivalent.

On the other hand, defining an RCM differently than by do-interventions on the true SCM is mathematically doable, as shown several times throughout Section 4. Adopting this second paradigm does not mean denying the formal equivalence between frameworks; it means not *applying* it. In this approach, potential outcomes $(Y_t)_{t \in \mathcal{T}}$ are typically primitives of the model not necessarily derived from \mathcal{M} (recall Remark 2 and Figure 1): the analyst may directly place assumptions on \mathcal{R} (like conditional ignorability) regardless of what holds in $\mathcal{R}_{\mathcal{M}}$. But they can also be derivatives of \mathcal{M} (or of another type of models) by other actions than do-interventions (see Remark 10). What characterizes the paradigm is essentially not a matter of “primitive versus derivative”, rather the fact that potential outcomes are not axiomatized by do-interventions on T in \mathcal{M} . This construction implies that the structural representations of \mathcal{R} may differ from \mathcal{M} and should thereby be ignored. The chosen (or derived) assumptions define the semantics of the counterfactual statements associated to \mathcal{R} . Recall that Section 4.3.2 notably showed that dressing potential outcomes with conditional ignorability defines their “had T been t ” as switching T into t while keeping all remaining variables unchanged (at the single-outcome level). Other assumptions could give other semantics. A methodological interest of this paradigm notably comes from the possibility to compute the *direct* causal effects of treatments breaking Assumption 7 through statistical methods without knowing \mathcal{M} (as explained in Remark 7). This practice is perhaps controversial and philosophically debatable. Nevertheless, we point out that (even though it is generally implicitly done) *not* defining potential outcomes as the true structural counterfactuals is actually something common in the scientific literature. In Remark 5, we mentioned a series of articles leveraging the potential-outcome framework equipped with the fundamental assumptions of causal inference to compute the effects of immutable characteristics [Li et al., 2017, Glymour and Spiegelman, 2017, Khademi et al., 2019, Khademi and Honavar, 2020, Qureshi et al., 2020, Makhlof et al., 2024]. As explained in Section 4.3, in such cases \mathcal{R} is generally not even single-world equivalent to $\mathcal{R}_{\mathcal{M}}$. This positions these works in the second paradigms. Therefore, if we acknowledge them, then we must accept that unifying potential outcomes and structural counterfactuals is not an obligation.

Remark 9 (About Holland’s principle in the second paradigm). Consider a practitioner who thinks that “no causation without manipulation” applies to RCMs but not to SCMs, and thereby only follows the second paradigm with no posttreatment covariate (Assumption 7). To them, Corollary 1 ensures that whenever the potential-outcome framework can be applied, the employed RCM is at least single-outcome equivalent to the latent SCM under Assumptions 5 and 6. This appear to rule out situations where potential outcomes may not coincide with structural counterfactuals in the second paradigm, such as the one from Sections 1.1.1 and 4.3.3 (and the one from Appendix A). Let us comment on this. Firstly, even in cases like Figure 2a, cross-world or almost-sure equivalence is not guaranteed. This means that substituting \mathcal{R} by $\mathcal{R}_{\mathcal{M}}$ is generally incorrect at stronger counterfactual-reasoning levels. This is why such a practitioner should still be careful. Secondly, as mentioned above, addressing a treatment not satisfying Assumption 7 with an RCM meeting the fundamental assumptions of causal inference is both theoretically possible and empirically done. This is why we discuss the implications of such a choice.

Remark 10 (About interventions and the true SCM in the second paradigm). We emphasize that adopting the second paradigm does not mean abandoning the *true* SCM \mathcal{M} ; it means ignoring the *structural representations* of the RCM. The sole consequence of this approach is that the causal interpretation of $(Y_t)_{t \in \mathcal{T}}$ possibly differs from the one of $(Y_{T=t})_{t \in \mathcal{T}}$ (as exemplified in Section 4.3.2). Said differently, the second paradigm marks a tension between potential outcomes and *do-interventions* on T in \mathcal{M} rather than between potential outcomes and the *model* \mathcal{M} itself. This echoes the conclusion of Remark 7.

Notably, one could define $(Y_t)_{t \in \mathcal{T}}$ through other types of interventions in \mathcal{M} than do-interventions. In particular, Theorem 1 shows that the single-outcome laws of \mathcal{R} under the fundamental assumption of causal inference do not necessarily result from a do-intervention on T in \mathcal{M} , but rather from an action that fixes T to t only in the equations corresponding to Y in \mathcal{M} . We could for instance define the entailed RCM of \mathcal{M} via such an intervention to reconcile the two causal models at this level.

In the motivating example from Section 1.1.1, analyst \mathcal{M} and analyst \mathcal{R} respectively adopt the first and second paradigms to define their counterfactual outcomes. As such, analyst \mathcal{M} exploits the formal equivalence between frameworks—not analyst \mathcal{R} . It feels that Pearl’s rule to always leverage an SCM as the axiomatic characterization of potential outcomes through do-interventions possibly made unclear the existence of these two paradigms. We emphasize that we do not discuss which paradigm people should follow. Instead, we neutrally classify what people *actually* do, and explain the implicit meaning

of these practices. All in all, each approach can be legitimate; *what crucially matters is having a clear understanding of the produced counterfactuals and being transparent about the choices made*. If an analyst aims to compute counterfactual estimands carrying a *mutatis mutandis* signification as defined by do-interventions, then they should explicitly mention that they follow the first paradigm, and specify the true SCM with at least the key graphical relationships. If an analyst aims to compute counterfactual estimands carrying another signification, then they should explicitly mention that they follow the second paradigm, and be clear about the counterfactual interpretation of their assumptions (*e.g.*, *ceteris paribus* via conditional ignorability). The following case must be avoided: an RCM analyst does not explicitly define their model via a partially specified SCM, misinterprets the formal equivalence and hence wrongly believes that their potential outcomes necessarily derive from do-interventions on the true SCM, therefore computes distinct causal estimands than what they expected. In this sense, we conclude this article by defending a more cautious use of notations.

5.2 On the exchange of notations

Following the conclusion of Section 5.1, this subsection discusses what researchers and practitioners should do. More specifically, our suggestions do not concern the methods people use, but how they *present* their approaches and their results. Concretely, we argue that exchanging the do notation and the potential-outcome notation should be done with greater care than what is commonly done in the literature. A notation is the identification of a mathematical object. Therefore, a same notation can be used for two differently-defined objects just in case they are mathematically equal. In the first paradigm, where the potential outcomes are chosen as the true structural counterfactuals, the notational exchange with the do is valid; not in the second paradigm. In what follows, we examine books and articles referring to the formal equivalence between causal frameworks in a confusing way, to underline the importance of specifying and justifying the adopted paradigm as well as using adequate notations.¹⁴

For example, [Colnet et al., 2024, Section 5] commences with a potential-outcome model, and then invoke the formal equivalence to substitute its notations by do notations. Because writing a “do” only makes sense when there is an SCM involved, this exchange implicitly engages a structural representation of their RCM. But recall that nothing guarantees that this structural representation corresponds at the single-outcome level to the true SCM. In their case, everything works fine precisely because only settings with no posttreatment covariate are considered, only single-outcome causal effects are studied, and their RCM satisfies conditional ignorability. However, they never explicitly state Assumption 7 and even less explain its crucial role. While this may seem harmless in practice, this could mislead people to wrongly believe that this equivalence holds in general: it would generally not hold with posttreatment covariates or at stronger counterfactual-reasoning levels. Even from a purely logical viewpoint, working in this favorable scenario should not obviate proper justification to why $\mathcal{L}((T, X, Y_t)) = \mathcal{L}((T, X, Y_{T=t}))$ for every $t \in \mathcal{T}$. We suggest two options to clarify [Colnet et al., 2024, Section 5]. The first is specifying that the first paradigm is adopted to define the RCM, and explaining how its assumptions (conditional ignorability) are derived from the true SCM (backdoor criterion). The second is specifying that the second paradigm is adopted to define the RCM, and justifying why the RCM and the true SCM happen to be single-outcome equivalent (Corollary 1). Replacing the potential-outcome notations of the RCM with do notations only makes sense in the first option.

The case of [Barocas et al., 2023, Chapter 5] and Makhlof et al. [2024] also demands caution. These surveys introduce the two causal frameworks specifically in the context of fairness with identical notations on counterfactual variables and suggest that the appropriate choice of framework is mostly a matter of methodological considerations because of the formal equivalence. Additionally, they present RCM-based inference techniques specifically relying on conditional ignorability. This feels pedagogically dangerous. As repeated, in settings with posttreatment covariates—fairness problems typically—one can generally not interchangeably manipulate the counterfactual outcomes of an RCM under the fundamental assumptions for causal inference and the ones of the *true* SCM. Section 4.3.3 notably exemplified how mixing the notations could lead to contradictory results. In such settings, SCM counterfactuals and RCM counterfactuals correspond to different paradigms. Leveraging both paradigms in a same work requires distinct notations to prevent confusion. This is particularly critical when the treatment may impact the covariates.

¹⁴We do not suggest that these references contain erroneous claims, only that the presentation of specific aspects can be misleading. Beyond that, we strongly recommend reading them.

These examples underline that the formal equivalence in itself is not a correct justification for exchanging do notations and potential-outcome notations in an RCM: the key argument is the paradigm followed to define the RCM. Writing a “do” exclusively makes sense in studies clearly following the first paradigm. Writing a potential-outcome notation is acceptable in both paradigms (since structural counterfactuals are potential outcomes) but distinct notations are needed when relying on both paradigms for a same problem.

To summarize Section 5, we think that the scope and implications of the formal equivalence between causal frameworks can be misleading. In particular, it does not mean that $\mathcal{L}((T, X, (Y_t)_{t \in \mathcal{T}})) = \mathcal{L}((T, X, (Y_{T=t})_{t \in \mathcal{T}}))$ in general; it means that such an exchange can hold *if equivalent assumptions are made across models*. Supposing distinct axioms across models means giving distinct interpretations to their respective counterfactual outcomes, which thereby relate to distinct causal interventions and causal effects. This is why we recommend to present Rubin’s causal models and Pearl’s causal models as distinct mechanisms for reasoning counterfactually that coincide under specific assumptions and choices, rather than merely different perspectives.

6 Conclusion

In this paper, we superimposed Pearl’s causal framework and Rubin’s causal framework without presuppositions to show that structural counterfactual outcomes and potential outcomes do not necessarily coincide at any levels of counterfactual reasoning. To furnish a thorough comparison at the most relevant level, we expressed the laws of potential outcomes in terms of the latent SCM under classical causal-inference assumptions. On the basis of this result, we gave a detailed interpretation of counterfactuals in each causal framework, specifying when they entailed different conclusions. More specifically, counterfactual inference with potential outcomes under conditional ignorability yields *ceteris paribus* counterfactuals with respect to the covariates, whereas counterfactual inference with a do-intervention on an SCM yields *mutatis mutandis* counterfactuals with respect to the covariates. If the cause of interest affects the covariates, these constructions are generally not equal in law. For these reasons, we call the community to not interchangeably use the potential-outcome framework and do-interventions to define counterfactual outcomes, unless the justification is explicitly made.

We emphasize that our contribution is not an argument in favor of using one causal model rather than the other, or against the formal equivalence between frameworks. Instead of taking position or addressing philosophical arguments, it highlights some facts: theoretically, one can perfectly define potential outcomes as distinct to the true structural counterfactuals; empirically, researchers have actually (implicitly) worked with potential outcomes defined as distinct to the true structural counterfactuals. Our work is meant to shed light on the different mathematical choices that analysts can make when working with counterfactual outcomes, and to precise their implications in order to prevent incorrect or ambiguous conclusions in causal studies. In doing this paper, we hope to clarify the similarities and differences between the two major causal approaches.

Acknowledgment

The author thanks Antoine Chambaz and Émilie Devijver for helpful discussions and valuable suggestions, as well as two anonymous reviewers for their useful recommendations. Most of this work has been done when the author was at Institut de Mathématiques de Toulouse (Université Paul Sabatier).

Author’s statement

This work was partially funded by the project CAUSALI-T-AI (ANR-23-PEIA-0007) of the French National Research Agency. The author confirms the sole responsibility for the presented results and manuscript preparation. The author states no conflict of interest.

A Illustration with pre and posttreatment covariates

Section 4.3.3 compared causal effects obtained with respectively potential outcomes under the causal-inference regime and structural counterfactuals in a problem corresponding to Figure 2c. This section makes the same comparison in the context of Figure 2b. It is meant to propose a real-world illustration of Figure 2b, and to prove that nonequivalence does not generally hold in this case as well.

The following example (more precisely its graph) is inspired by Li et al. [2022]. The treatment status $T : \Omega \rightarrow \{0, 1\}$ indicates whether an individual is a smoker or not; a first covariate $X_1 : \Omega \rightarrow \mathbb{R}$ represents the expression of a gene, which is higher when the gene is more active; a second covariate $X_2 : \Omega \rightarrow \mathbb{R}$ quantifies blood pressure; the outcome $Y : \Omega \rightarrow \mathbb{R}$ evaluates a person's health, a higher score meaning a better health. Suppose that these variables are observed in the context of a medical study. Causal analysts are task with assessing the effect of smoking onto health. Assume that the observational vector $\mathcal{O} := (T, X, Y)$ is ruled by the following SCM $\mathcal{M} \in \mathfrak{M}_{\mathcal{O}}$:

$$\begin{aligned} T &\stackrel{a.s.}{=} \mathbf{1}_{\{X_1 + U_T > 0\}}, \\ X_1 &\stackrel{a.s.}{=} U_1, \\ X_2 &\stackrel{a.s.}{=} \alpha T + U_2, \\ Y &\stackrel{a.s.}{=} \gamma X_1 - X_2 + \beta T + U_Y, \end{aligned}$$

where α, β, γ are again deterministic parameters. Typically, $\alpha > 0$ encodes that smoking increases blood pressure, $\beta < 0$ represents the direct negative impact of smoking onto health, and $\gamma < 0$ describes a gene that deteriorates health (and makes people more inclined to smoke). Using the notations of the paper: $X := (X_1, X_2)$ and $U_X := (U_1, U_2)$. As in the fairness illustration, the model satisfies Assumption 5 and the negation of Assumption 7 by design. Moreover, as before we suppose that positivity (Assumption 2) is true and that $U_Y \perp\!\!\!\perp (U_T, U_X)$ so that Assumption 6 holds. Note that the SCM fits Figure 2b; the exact graph is more precisely depicted in Figure 3. Finally, we set two potential outcomes (Y_0, Y_1) meeting the consistency rule, that is $Y \stackrel{a.s.}{=} (1 - T) \cdot Y_0 + T \cdot Y_1$, and single-outcome conditional ignorability (Assumption 4). This defines $\mathcal{R} := (T, X, Y_0, Y_1) \in \mathfrak{R}_{\mathcal{O}}$. All the assumptions of Theorem 1 are satisfied.

Let us compute the estimands that an analyst working with \mathcal{R} and an analyst working with $\mathcal{R}_{\mathcal{M}}$ respectively obtain for the average causal effect of T onto Y conditional to $X = x$:

$$\begin{aligned} \text{CATE}_{\mathcal{R}}(x) &:= \mathbb{E}[Y_1 - Y_0 \mid X = x] \\ &= \mathbb{E}[(\gamma x_1 - x_2 + \beta + U_Y) - (\gamma x_1 - x_2 + U_Y) \mid X = x] \\ &= \beta, \end{aligned}$$

and

$$\begin{aligned} \text{CATE}_{\mathcal{R}_{\mathcal{M}}}(x) &:= \mathbb{E}[Y_{T=1} - Y_{T=0} \mid X = x] \\ &= \mathbb{E}[(\gamma(X_{T=1})_1 - (X_{T=1})_2 + \beta + U_Y) - (\gamma(X_{T=0})_1 - (X_{T=0})_2 + U_Y) \mid X = x] \\ &= \mathbb{E}[\gamma(X_{T=1} - X_{T=0})_1 - (X_{T=1} - X_{T=0})_2 \mid X = x] + \beta \\ &= \mathbb{E}[0 - (X_{T=1} - X_{T=0})_2 \mid X = x] + \beta \\ &= -\mathbb{E}[(\alpha + U_2) - U_2 \mid X = x] + \beta \\ &= -\alpha + \beta. \end{aligned}$$

In this configuration as well $\text{CATE}_{\mathcal{R}}$ recovers the direct effect of T in \mathcal{M} , while $\text{CATE}_{\mathcal{R}_{\mathcal{M}}}$ recovers the total effect. This supports that $\text{CATE}_{\mathcal{R}}$ is a direct effect no matter the graph from Figure 2: it comes from the counterfactual interpretation of fixing everything but the the treatment (Section 4.3.2). Note that compared to the fairness illustration, there is a minus before α in the total effect due to X_2 having a negative impact on Y . Additionally, both effects ignore the noncausal path from T to Y via X_1 quantified by γ . Most importantly, $\text{CATE}_{\mathcal{R}} \neq \text{CATE}_{\mathcal{R}_{\mathcal{M}}}$ if $\alpha \neq 0$. As expected from the discussion in Section 4.3.1, an RCM satisfying the fundamental assumptions of causal inference is generally not single-world equivalent to the latent SCM when both pre and posttreatment covariates coexist.

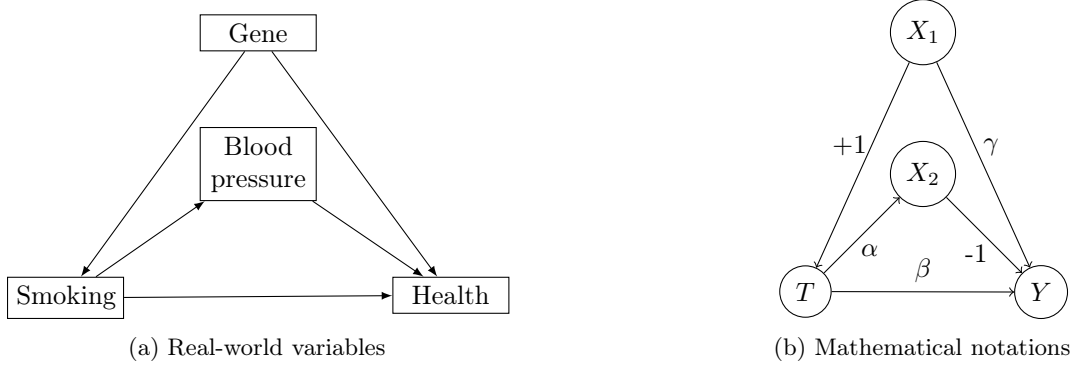


Figure 3: Causal graph of \mathcal{M} in Appendix A. Linear coefficients are indicated on Figure 3b. Exogenous variables are not represented but Assumption 6 holds.

B Remaining proofs

Proof of Lemma 1 Since \mathcal{M} is acyclical, there exists a topological ordering on the indices in \mathcal{I} with respect to the parent-child relation, and therefore on the subset I^c . This means in particular that there exist some indices $j \in I^c$ such that g_j takes only variables in V_I as endogenous inputs. Starting from these indices, and recursively substituting along the topological ordering produces a measurable f_{I^c} such that

$$V_{I^c} \stackrel{a.s.}{=} f_{I^c}(V_{\text{Endo}(I^c) \setminus I^c}, U_{\text{Exo}(I^c)}).$$

Note that $\text{Endo}(I^c) \setminus I^c = \text{Endo}(I^c) \cap I \subseteq I$. Carrying out the same substitution on the intervened model $\mathcal{M}_{V_I=v_I}$ with solution \tilde{V} gives

$$\tilde{V}_{I^c} \stackrel{a.s.}{=} f_{I^c}(v_{\text{Endo}(I^c) \setminus I^c}, U_{\text{Exo}(I^c)}),$$

while by definition $\tilde{V}_I \stackrel{a.s.}{=} v_I$. ■

Proof of Lemma 2 Let $t \in \mathcal{T}$. Positivity ensures that $\mathbb{P}(\cdot | X = x, T = t')$ is well-defined for $\mathcal{L}(X)$ -almost every $x \in \mathbb{R}^d$ and any $t' \in \mathcal{T}$. Thereby, conditional ignorability entails that $\mathcal{L}(Y_t | X = x, T = t) = \mathcal{L}(Y_t | X = x, T = t')$. Moreover, consistency implies that $\mathcal{L}(Y | X = x, T = t) = \mathcal{L}(Y_t | X = x, T = t)$. All in all, for any $t' \in \mathcal{T}$:

$$\mathcal{L}(Y_t | X = x, T = t') = \mathcal{L}(Y | X = x, T = t). \quad (11)$$

Next, we use the above expression involving *conditional* distributions to obtain an expression for the *joint* distribution of (T, X, Y_t) . Let $F \subseteq \mathcal{T} \times \mathbb{R}^d \times \mathbb{R}^p$ be a Borel set. We have,

$$\mathbb{P}((T, X, Y_t) \in F) = \int \mathbb{P}((t', x, Y_t) \in F | X = x, T = t') d\mathbb{P}(X = x, T = t').$$

Then, we define the sets $F(t', x) := \{y \in \mathbb{R}^p | (t', x, y) \in F\}$ for every $(t', x) \in \mathcal{T} \times \mathbb{R}^d$ (which are Borel sets) so that

$$\begin{aligned} \mathbb{P}((T, X, Y_t) \in F) &= \int \mathbb{P}(Y_t \in F(t', x) | X = x, T = t') d\mathbb{P}(X = x, T = t') \\ &= \int \mathbb{P}(Y \in F(t', x) | X = x, T = t) d\mathbb{P}(X = x, T = t'), \end{aligned}$$

where the last equality follows from (11). This concludes the proof. ■

Proof of Lemma 3 Let $t \in \mathcal{T}$. By assumption, the random vector $\mathcal{O} := (T, X, Y)$ is the solution to an acyclical SCM. We write U_X and U_Y the exogenous parents of respectively X and Y . Therefore, by partitioning \mathcal{O} into T and (X, Y) , Lemma 1 guarantees the existence of a measurable function $f_{X,Y}$ such

that

$$\begin{aligned}(X, Y) &\stackrel{a.s.}{=} f_{X,Y}(T, U_X, U_Y), \\ (X_{T=t}, Y_{T=t}) &\stackrel{a.s.}{=} f_{X,Y}(t, U_X, U_Y).\end{aligned}$$

Therefore, selecting the coordinates corresponding to Y furnishes a measurable function \tilde{f}_Y such that

$$\begin{aligned}Y &\stackrel{a.s.}{=} \tilde{f}_Y(T, U_X, U_Y), \\ Y_{T=t} &\stackrel{a.s.}{=} \tilde{f}_Y(t, U_X, U_Y).\end{aligned}$$

These expressions hold on a measurable set $\Omega^* \subseteq \Omega$ such that $\mathbb{P}(\Omega^*) = 1$. To conclude, simply observe that for any $\omega \in \Omega^*$ such that $T(\omega) = t$, we have

$$Y(\omega) = \tilde{f}_Y(t, U_X(\omega), U_Y(\omega)) = Y_{T=t}(\omega).$$

■

Proof of Lemma 4 By assumption, the random vector $\mathcal{O} := (T, X, Y)$ is the solution to an acyclical SCM where U_T, U_X and U_Y denote the exogenous parents of respectively T, X and Y . Since additionally $Y_{\text{Endo}(T)} = Y_{\text{Endo}(X_i)} = \emptyset$ for every $i \in \{1, \dots, d\}$ (Assumption 5), Lemma 1 ensures the existence of a measurable function $f_{T,X}$ such that $(T, X) \stackrel{a.s.}{=} f_{T,X}(U_T, U_X)$. Therefore, if $U_Y \perp\!\!\!\perp (U_T, U_X)$ (Assumption 6), then $U_Y \perp\!\!\!\perp (T, X)$. ■

Proof of Proposition 3 Throughout the proof, we write X_i, Y_i , and $Y_{t,i}$ for the i th component of respectively X, Y , and Y_t for any $t \in \mathcal{T}$.

We start by finding a generative model for the RCM $\mathcal{R} := (T, X, (Y_t)_{t \in \mathcal{T}})$ using a standard result from probability theory: any Borel probability measure on a Polish space can be obtained by *push-forward* of the uniform measure on $[0, 1]$ denoted by Unif . This follows from the existence of a measurable bijection with measurable inverse between any Polish space and \mathbb{R} [Kechris, 2012, Theorem 15.6], and the fact that any Borel probability measure on \mathbb{R} can always be obtained by push-forward of Unif by its generalized-inverse probability distribution function. As a consequence, there always exists a measurable function $\psi : [0, 1] \rightarrow \mathcal{T} \times \mathbb{R}^d \times (\mathbb{R}^p)^{N+1}$ such that $\mathcal{L}((T, X, (Y_t)_{t \in \mathcal{T}})) = \text{Unif} \circ \psi^{-1}$. Thereby, there exists a random variable U such that $\mathcal{L}(U) = \text{Unif}$ and $(T, X, (Y_t)_{t \in \mathcal{T}}) \stackrel{a.s.}{=} \psi(U)$. We divide ψ into $\psi = (\psi_T, (\psi_{X_i})_{i=1}^d, (\psi_{Y_{0,i}})_{i=1}^p, \dots, (\psi_{Y_{N,i}})_{i=1}^p)$ and define $\psi_X := (\psi_{X_i})_{i=1}^d$ and $\psi_{Y_t} := (\psi_{Y_{t,i}})_{i=1}^p$ for any $t \in \mathcal{T}$. This enables us to write $T \stackrel{a.s.}{=} \psi_T(U)$, $X \stackrel{a.s.}{=} \psi_X(U)$, and $Y_t \stackrel{a.s.}{=} \psi_{Y_t}(U)$ for any $t \in \mathcal{T}$.

Then, on the basis of ψ and U , we construct an SCM $\mathcal{M} := \langle U, g \rangle$ that satisfies two properties: (1) \mathcal{M} is acyclical with solution (T, X, Y) , (2) $Y_{T=t} \stackrel{a.s.}{=} Y_t$ for any $t \in \mathcal{T}$. We define $g := (g_T, (g_{X_i})_{i=1}^d, (g_{Y_i})_{i=1}^p)$ where $g_T := \psi_T$, $g_{X_i} = \psi_{X_i}$, and g_{Y_i} is given by $(t, u) \mapsto \psi_{Y_{t,i}}(u)$. The associated structural equations correctly define T and X by definition of ψ_T and ψ_X . They also correctly define Y since the fact that $(Y_t)_{t \in \mathcal{T}}$ satisfies the consistency rule means that $Y \stackrel{a.s.}{=} \sum_{t \in \mathcal{T}} \mathbf{1}_{\{T=t\}} \psi_{Y_t}(U) = \sum_{t \in \mathcal{T}} \mathbf{1}_{\{T=t\}} g_Y(t, U) = g_Y(T, U)$ where $g_Y := (g_{Y_i})_{i=1}^p$. Additionally, g clearly does not represent cyclic relationships between variables. Therefore, $\mathcal{M} := \langle U, g \rangle$ belongs to $\mathfrak{M}_{\mathcal{O}}$. Moreover, do-interventions in this model produces for any $t \in \mathcal{T}$, $Y_{T=t} \stackrel{a.s.}{=} g_Y(t, U) = \psi_{Y_t}(U) \stackrel{a.s.}{=} Y_t$. Therefore, $\mathcal{R}_{\mathcal{M}}$ and \mathcal{R} are almost-surely equivalent. This concludes the proof. ■

References

- S. Barocas, M. Hardt, and A. Narayanan. *Fairness and Machine Learning: Limitations and Opportunities*. MIT Press, 2023.
- S. Beckers. Equivalent causal models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 6202–6209, 2021.
- R. A. Berk. *Regression analysis: A constructive critique*, volume 11. Sage, 2004.

- M. Bertrand and S. Mullainathan. Are emily and greg more employable than lakisha and jamal? a field experiment on labor market discrimination. *American economic review*, 94(4):991–1013, 2004.
- P. J. Bickel, E. A. Hammel, and J. W. O’Connell. Sex bias in graduate admissions: Data from Berkeley: Measuring bias is harder than is usually assumed, and the evidence is sometimes contrary to expectation. *Science*, 187(4175):398–404, 1975.
- T. Blom, S. Bongers, and J. M. Mooij. Beyond structural causal models: Causal constraints models. In R. P. Adams and V. Gogate, editors, *Proceedings of The 35th Uncertainty in Artificial Intelligence Conference*, volume 115 of *Proceedings of Machine Learning Research*, pages 585–594. PMLR, 22–25 Jul 2020. URL <https://proceedings.mlr.press/v115/blom20a.html>.
- C. R. Blyth. On Simpson’s paradox and the sure-thing principle. *Journal of the American Statistical Association*, 67(338):364–366, 1972.
- S. Bongers, P. Forré, J. Peters, and J. M. Mooij. Foundations of structural causal models with cycles and latent variables. *The Annals of Statistics*, 49(5):2885–2915, 2021.
- S. Chiappa. Path-specific counterfactual fairness. In *AAAI Conference on Artificial Intelligence*, volume 33, pages 7801–7808, 2019.
- S. Chiappa and W. S. Isaac. *A Causal Bayesian Networks Viewpoint on Fairness*, pages 3–20. Springer International Publishing, Cham, 2019. ISBN 978-3-030-16744-8. doi: 10.1007/978-3-030-16744-8_1. URL https://doi.org/10.1007/978-3-030-16744-8_1.
- B. Colnet, I. Mayer, G. Chen, A. Dieng, R. Li, G. Varoquaux, J.-P. Vert, J. Josse, and S. Yang. Causal inference methods for combining randomized trials and observational studies: a review. *Statistical science*, 39(1):165–191, 2024.
- D. A. Freedman. Graphical models for causation, and the identification problem. *Evaluation Review*, 28(4):267–293, 2004.
- J. Gaebler, W. Cai, G. Basse, R. Shroff, S. Goel, and J. Hill. A causal framework for observational studies of discrimination. *Statistics and public policy*, 9(1):26–48, 2022.
- M. M. Glymour and D. Spiegelman. Evaluating public health interventions: 5. causal inference in public health research—do sex, race, and biological factors cause health outcomes? *American journal of public health*, 107(1):81–85, 2017.
- D. J. Greiner and D. B. Rubin. Causal effects of perceived immutable characteristics. *Review of Economics and Statistics*, 93(3):775–785, 2011.
- P. W. Holland. Statistics and causal inference. *Journal of the American statistical Association*, 81(396):945–960, 1986.
- D. Ibeling and T. Icard. Comparing causal frameworks: Potential outcomes, structural models, graphs, and abstractions. *Advances in Neural Information Processing Systems*, 36, 2024.
- G. W. Imbens. Nonparametric estimation of average treatment effects under exogeneity: A review. *Review of Economics and statistics*, 86(1):4–29, 2004.
- G. W. Imbens. Potential outcome and directed acyclic graph approaches to causality: Relevance for empirical practice in economics. *Journal of Economic Literature*, 58(4):1129–79, 2020.
- G. W. Imbens and D. B. Rubin. *Causal inference in statistics, social, and biomedical sciences*. Cambridge University Press, 2015.
- A. Kechris. *Classical descriptive set theory*, volume 156. Springer Science & Business Media, 2012.
- A. Khademi and V. Honavar. Algorithmic bias in recidivism prediction: A causal perspective (student abstract). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 13839–13840, 2020.

- A. Khademi, S. Lee, D. Foley, and V. Honavar. Fairness in algorithmic decision making: An excursion through the lens of causality. In *The World Wide Web Conference*, pages 2907–2914, 2019.
- M. J. Kusner, J. Loftus, C. Russell, and R. Silva. Counterfactual fairness. In *Advances in Neural Information Processing Systems*, volume 30, pages 4066–4076. Curran Associates, Inc., 2017.
- D. Lewis. Causation. *Journal of Philosophy*, 70(17):556–567, 1973a.
- D. Lewis. *Counterfactuals*. Blackwell, 1973b.
- D. Lewis. Counterfactual dependence and time’s arrow. *Noûs*, pages 455–476, 1979.
- J. Li, J. Liu, L. Liu, T. D. Le, S. Ma, and Y. Han. Discrimination detection by causal effect estimation. In *2017 IEEE International Conference on Big Data (Big Data)*, pages 1087–1094. IEEE, 2017.
- Q. Li, Z. Wang, S. Liu, G. Li, and G. Xu. Deep treatment-adaptive network for causal inference. *The VLDB Journal*, 31(5):1127–1142, 2022.
- K. Makhoulouf, S. Zhioua, and C. Palamidessi. When causality meets fairness: A survey. *Journal of Logical and Algebraic Methods in Programming*, page 101000, 2024.
- K. A. Markus. Causal effects and counterfactual conditionals: contrasting Rubin, Lewis and Pearl. *Economics & Philosophy*, 37(3):441–461, 2021.
- K. Muandet, M. Kanagawa, S. Saengkyongam, and S. Marukatat. Counterfactual mean embeddings. *The Journal of Machine Learning Research*, 22(1):7322–7392, 2021.
- B. Neal. *Introduction to causal inference from a machine learning perspective*. bradyneal.com, 2020. https://www.bradyneal.com/Introduction_to_Causal_Inference-Dec17_2020-Neal.pdf.
- J. Neyman. Sur les applications de la thar des probabilités aux expériences agaricales: Essay des principe. excerpts reprinted (1990) in english. *Statistical Science*, 5(463-472):4, 1923.
- H. Nilforoshan, J. D. Gaebler, R. Shroff, and S. Goel. Causal conceptions of fairness and their consequences. In *International Conference on Machine Learning*, pages 16848–16887. PMLR, 2022.
- J. Park, U. Shalit, B. Schölkopf, and K. Muandet. Conditional distributional treatment effect with kernel conditional mean embeddings and u-statistic regression. In *International Conference on Machine Learning*, pages 8401–8412. PMLR, 2021.
- J. Pearl. *Causality*. Cambridge university press, 2009.
- J. Pearl. Brief report: On the consistency rule in causal inference: " axiom, definition, assumption, or theorem?". *Epidemiology*, pages 872–875, 2010.
- J. Pearl. Does obesity shorten life? or is it the soda? on non-manipulable causes. *Journal of Causal Inference*, 6(2):20182001, 2018.
- J. Pearl. On the interpretation of do(x). *Journal of Causal Inference*, 7(1):20192002, 2019. doi: doi: 10.1515/jci-2019-2002. URL <https://doi.org/10.1515/jci-2019-2002>.
- J. Pearl, M. Glymour, and N. P. Jewell. *Causal Inference in Statistics: A Primer*. John Wiley & Sons, 2016.
- D. Plecko and N. Meinshausen. Fair data adaptation with quantile preservation. *Journal of Machine Learning Research*, 21:1–44, 2020.
- B. Qureshi, F. Kamiran, A. Karim, S. Ruggieri, and D. Pedreschi. Causal inference for social discrimination reasoning. *Journal of Intelligent Information Systems*, 54:425–437, 2020.
- T. S. Richardson and J. M. Robins. Single world intervention graphs (SWIGs): a unification of the counterfactual and graphical approaches to causality. *Center for the Statistics and the Social Sciences*, 128(30):2013, 2013. Working Paper.

- G. Ridgeway. Assessing the effect of race bias in post-traffic stop outcomes using propensity scores. *Journal of quantitative criminology*, 22:1–29, 2006.
- P. R. Rosenbaum and D. B. Rubin. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1):41–55, 1983.
- D. Rubin. Causal inference. In P. Peterson, E. Baker, and B. McGaw, editors, *International Encyclopedia of Education*, pages 66–71. Elsevier, Oxford, third edition edition, 2010. ISBN 978-0-08-044894-7. doi: <https://doi.org/10.1016/B978-0-08-044894-7.01313-0>. URL <https://www.sciencedirect.com/science/article/pii/B9780080448947013130>.
- D. B. Rubin. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of educational Psychology*, 66(5):688, 1974.
- T. J. VanderWeele and W. R. Robinson. On the causal interpretation of race in regressions adjusting for confounding and mediating variables. *Epidemiology*, 25(4):473–484, 2014.
- T. J. VanderWeele, S. Vansteelandt, and J. M. Robins. Effect decomposition in the presence of an exposure-induced mediator-outcome confounder. *Epidemiology*, 25(2):300–306, 2014.
- C. Winship and S. L. Morgan. The estimation of causal effects from observational data. *Annual review of sociology*, 25(1):659–706, 1999.
- L. Yao, Z. Chu, S. Li, Y. Li, J. Gao, and A. Zhang. A survey on causal inference. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 15(5):1–46, 2021.