
Towards Estimating Bounds on the Effect of Policies under Unobserved Confounding

Alexis Bellot* Silvia Chiappa
Google DeepMind
London, UK

Abstract

As many practical fields transition to provide personalized decisions, data is increasingly relevant to support the evaluation of candidate plans and policies (*e.g.*, guidelines for the treatment of disease, government directives, etc.). In the machine learning literature, significant efforts have been put into developing machinery to predict the effectiveness of policies efficiently. The challenge is that, in practice, the effectiveness of a candidate policy is not always identifiable, *i.e.*, not uniquely estimable from the combination of the available data and assumptions about the domain at hand (*e.g.*, encoded in a causal graph). In this paper, we develop graphical characterizations and estimation tools to bound the effect of policies given a causal graph and observational data collected in non-identifiable settings. Specifically, our contributions are two-fold: (1) we derive analytical bounds for general probabilistic and conditional policies that are tighter than existing results, (2) we develop an estimation framework to estimate bounds from finite samples, applicable in higher-dimensional spaces and continuously-valued data. We further show that the resulting estimators have favourable statistical properties such as fast convergence and robustness to model misspecification.

1 Introduction

Understanding how to act upon the world around us requires humans and artificial systems to thoughtfully evaluate the effect of different plans, policies, and interventions one might consider. Data and methods that facilitate this process are increasingly relevant for high-stakes decision-making, notably in medical care with the rise of precision medicine [41], but also in education [25], law enforcement [13], public policy [44], and economics [2]. The interventions that a practitioner might consider, could consist of complex policies where a variable is set to follow a conditional or stochastic relationship depending on other variables in the system. For example, policy-makers might contemplate higher taxes on processed foods and stronger campaigns on its health risks targeted to overweight individuals. A sensible question in this context could be “what is the effect of a policy that reduces the consumption of processed foods by 50% in people with a body mass index above 30?”. Contrary to atomic (also called hard) interventions that force a particular value, this policy suggests a *softer* intervention.

The identifiability and estimation of policies from data is a widely studied problem in the reinforcement learning [37, 39] and causal inference [2, 12, 15, 31, 32, 40] literatures. Despite the generality entailed by many common approaches in these fields, they often rely on an impractical condition: the assumption that the effectiveness of the policy is uniquely computable from the observed data and

*Correspondence to Alexis Bellot: abellot@google.com

assumptions about the data generating process. This can be violated in real-world settings subject to unobserved confounding, and lead to the non-identifiability of the effectiveness of a given policy. This holds irrespective of the number of samples collected so that in reality multiple answers for the effect of a policy may be equally plausible. In their seminal work in the early 1990's, Manski and Robins [27, 33] showed that, nevertheless, the effect of an atomic intervention may always be bounded in a non-trivial interval (*i.e.*, probabilities strictly contained in $[0, 1]$), irrespective of unobserved confounders or the causal structure underlying the variables involved. Causal effects are therefore in general said to be *partially identifiable*, *i.e.*, one can derive bounds that shrink the range of a priori possible values for a given effect and potentially serve as a useful support for decision-making.

Starting from this insight, the problem of partial identification has been gaining growing attention in the literature. To improve upon these bounds, one common restriction on the data generating mechanism is to assume knowledge of a causal graph describing the phenomenon of interest. Several results have derived progressively tighter bounds by exploiting the independencies in observational and interventional distributions implied by the causal graph [4, 5, 14, 46, 47, 48]. For instance, Balke and Pearl [4] (and subsequent refinements, *e.g.*, [14, 48]) defined polynomial optimization programs to derive bounds that are provably optimal. Tighter analytical bounds have also been derived in selected settings, such as for the "instrumental variable" graph [47] and discrete systems with general graphs or equivalence classes [5, 46]. These results, however, are almost exclusively concerned with atomic interventions and small systems of discretely-valued variables. Despite their generality, there is still a gap towards making the partial-identification and estimation of the effectiveness of policies practical. In particular, making scalable inferences from finite samples with continuous outcomes and higher-dimensional covariates is not possible with existing tools.

This paper aims to provide a novel estimation framework to support decision-making in non-identifiable settings, overcoming some of these challenges. We introduce several new results for the partial identification and estimation of the effect of stochastic and conditional policies from a combination of observational data and assumptions on the domain, encoded in a causal graph. Our contributions may be summarized as follows. We introduce several graphical criteria to derive new analytical bounds on the effect of policies with continuous outcomes and covariates, that improve upon the non-parametric bounds of [27, 33] and [47]. Given these analytical bounds, we then construct estimators leveraging the double machine learning [10] toolkit for scalable inference, and demonstrate that estimators exhibit favourable statistical properties such as robustness to noise and fast convergence. These results are applicable to arbitrary stochastic or conditional policies that an investigator might design, and high-dimensional covariate spaces. Recent work has made progress in developing powerful estimators for identifiable causal effects [6, 8, 19, 20, 36, 49]. However, these estimators are not applicable in non-identifiable settings. To our knowledge, the proposed estimation machinery is the first result for bounding the effect of policies given a causal diagram.

Preliminaries. We use capital letters to denote variables (X), small letters for their values (x), bold letters for sets of variables (\mathbf{X}) and their values (\mathbf{x}), and use supp to denote their domains of definition ($x \in \text{supp}_X$). $\mathbb{1}_x(X)$ is the indicator function that equals 1 if the statement $\{X = x\}$ is true, and equal to 0 otherwise.

The framework we use to underpin the estimation of the effect of policies rests on *structural causal models* (SCMs) [31, Def. 7.1.1]. An SCM \mathcal{M} is a tuple $\langle \mathbf{V}, \mathbf{U}, \mathcal{F}, P(\mathbf{U}) \rangle$, where \mathbf{V} is a set of endogenous (observed) variables, \mathbf{U} is a set of exogenous latent variables, and $\mathcal{F} = \{f_V\}_{V \in \mathbf{V}}$ is a set of functions such that f_V determines values of V taking as argument variables $\mathbf{Pa}_V \subseteq \mathbf{V}$ and $\mathbf{U}_V \subseteq \mathbf{U}$, *i.e.* $V \leftarrow f_V(\mathbf{Pa}_V, \mathbf{U}_V)$. Values of \mathbf{U} are drawn from an exogenous distribution $P(\mathbf{u})$. We assume the model to be recursive, *i.e.* that there are no cyclic dependencies among the variables. Graphically, each SCM \mathcal{M} is associated with a causal diagram \mathcal{G} over \mathbf{V} , where $V \rightarrow W$ if V appears as an argument of f_W in \mathcal{M} , and $V \dashrightarrow W$ if $\mathbf{U}_V \cap \mathbf{U}_W \neq \emptyset$, *i.e.* V and W share an unobserved confounder. We will use graph-theoretic family abbreviations to represent graphical relations, *e.g.* *an*, *de*, etc. Moreover, for a causal diagram \mathcal{G} over \mathbf{V} , the \mathbf{X} -lower-manipulation of \mathcal{G} deletes all those edges that are out of variables in \mathbf{X} , and otherwise keeps \mathcal{G} as it is. The resulting graph is denoted as $\mathcal{G}_{\underline{\mathbf{X}}}$. The \mathbf{X} -upper-manipulation of \mathcal{G} deletes all those edges that are into variables in \mathbf{X} , and otherwise keeps \mathcal{G} as it is. The resulting graph is denoted as $\mathcal{G}_{\overline{\mathbf{X}}}$. We use $\perp\!\!\!\perp_d$ to denote d -separation in causal diagrams [31, Def. 1.2.3].

For a sample set $\mathcal{D} := \{\mathbf{v}^{(i)}\}_{i=1,\dots,n} \sim P$, we use $\mathbb{E}_{\mathcal{D}}[f(\mathbf{V})] := (1/n) \sum_{i=1}^n f(\mathbf{v}^{(i)})$. We use $\|f\|_P := \sqrt{\mathbb{E}_P[\{f(\mathbf{V})\}^2]}$. $\hat{f} - f = o_P(r_n)$ denotes that a function \hat{f} is a consistent estimator of f at a rate r_n , $\hat{f} - f = O_P(r_n)$ denotes that it is bounded in probability at rate r_n .

2 Partial Identification of the Effect of Policies

A policy π over a subset $\mathbf{X} \subset \mathbf{V}$ is a sequence of decision rules or plans $\pi := \{\pi_X\}_{X \in \mathbf{X}}$ for determining the assignment of variables \mathbf{X} . In its most general form, every $\pi_X : \text{supp}_X \times \text{supp}_{C_X} \rightarrow [0, 1]$ is a probability mapping from domain of $C_X \subset C \subset \mathbf{V}$ to the domain of X .

Qualitatively different types of interventions may be modelled with π_X . Specifically, a deterministic intervention setting $X := g(\mathbf{c}_X)$ based on the values of C_X can be encoded as $\pi_X(x | \mathbf{c}_X) := \mathbb{1}_{g(\mathbf{c}_X)}(x)$ while a probabilistic intervention may be written $\pi_X(x | \mathbf{c}_X) := P_{\text{induced by } \pi_X}(X = x | \mathbf{c}_X)$. The intervened model \mathcal{M}_π represents a different regime in which the assignment $\{f_X\}_{X \in \mathbf{X}}$ is replaced with the assignment induced by π . For an outcome $Y \in \mathbf{V}$, the interventional distribution $P_{\mathcal{M}_\pi}(Y)$, equivalently $P_\pi(Y)$, is defined as the distribution over Y in \mathcal{M}_π . It will be useful to adopt the notation $\bar{\pi} := \prod_{X \in \mathbf{X}} \pi_X(x | \mathbf{c}_X)$ to denote the product of policy assignments on individual variables subject to intervention in π .




Figure 1: Graphs for Example 1.

Example 1 (Illustration of policies). In the context of a public health program, let X be a measure of an individual’s weekly physical exercise routine, C an individual’s family history of metabolic diseases, and Y hospital admissions related to heart disease. Consider the causal diagram Fig. 1a; currently, an individual’s exercise voluntarily, which depends on unobserved factors that could be associated with risk of heart disease. The government is considering an incentive plan aimed at increasing the frequency of exercise depending individual’s family history of metabolic diseases (with some probability). The proposed intervention can be encoded as $\pi = \{\pi_X\}$ and sets the new assignment \tilde{f}_X such that the variable X follows the pre-specified distribution. Graphically, this policy is represented by Fig. 1b (with the new edge corresponding to the implementation of the policy highlighted in blue) and may be evaluated with the following quantity,

$$\mathbb{E}_{P_\pi}[Y] = \sum_{x,c} \mathbb{E}_{P_\pi}[Y | x, c] \pi_X(x | c) P_\pi(c). \quad (1)$$

The inferential challenge is that this quantity is not uniquely computable from $P(x, y, c)$ and \mathcal{G} , requiring more complex and nuanced notions of policy evaluation.

Formally, we are interested in the evaluation of the effectiveness of a plan or policy π , acting on a (potentially multivariate) discrete action \mathbf{X} based on values of observed covariates $\mathbf{C} \in \mathbb{R}^d$, on an outcome of interest $Y \in [0, 1]^2$.

Definition 1 (Average treatment effect). *The effectiveness of a policy π on Y is $\mathbb{E}_{P_\pi}[Y]$.*

From the investigator’s perspective, only the causal diagram \mathcal{G} of the environment \mathcal{M} is available. No assumptions about the form or shape of $P(\mathbf{U})$ and \mathcal{F} are made, but for the structural knowledge encoded in \mathcal{G} . In general, there might exist multiple SCMs \mathcal{M} that induce the same causal diagram \mathcal{G} and entail $P(\mathbf{V})$ but result in different values of $\mathbb{E}_{P_\pi}[Y]$. The identification of a *set* of possible solutions that contain the true value of $\mathbb{E}_{P_\pi}[Y]$ leads to the notion of *partial identification*.

Definition 2 (Partial Identification). *The effectiveness of a policy π is said to be partially identifiable from \mathcal{G} and $P(\mathbf{v})$ if*

$$\psi^\ell(P) \leq \mathbb{E}_{P_{\mathcal{M}_\pi}}[Y] \leq \psi^u(P), \quad \text{for any } \mathcal{M} \text{ such that } \mathcal{G}_\mathcal{M} = \mathcal{G}, P_\mathcal{M}(\mathbf{v}) = P(\mathbf{v}), \quad (2)$$

where (ψ^ℓ, ψ^u) are functionals of P that are bounded away from 0 and 1, respectively.

The earliest bound on the effect of policies, the so-called *natural bounds*, were developed considering discrete atomic interventions, i.e. of the form $\pi_X = \mathbb{1}_x(X)$, $X \in \{1, \dots, d_X\}$.

²An upper bound of 1 is assumed for simplicity; the following results could be generalized to more general bounded random variables.




Figure 2: Graphs used in Sec. 2 and 3.

Proposition 1 (Natural Bounds). *For any x , $\mathbb{E}_P[Y \mathbb{1}_x(\mathbf{X})] \leq \mathbb{E}_{P_x}[Y] \leq \mathbb{E}_P[(Y - 1) \mathbb{1}_x(\mathbf{X})] + 1$.*

Remarkably, these inequalities hold irrespective of the causal graph of the system \mathcal{G} ; a result that dates back to [27, 33]. Similarly, we could adapt the underlying proof strategy to derive functionals of P that bound the effect of more general probabilistic or conditional policies π , given as follows.

Proposition 2 (Natural Policy Bounds (NPB)). *For any π , $\mathbb{E}_P[Y \pi] \leq \mathbb{E}_{P_\pi}[Y] \leq \mathbb{E}_P[(Y - 1)\pi] + 1$.*

These inequalities also hold irrespective of the causal graph of the system. The natural bounds from [27, 33] are a special case of the natural policy bounds (NPBs) by setting $\pi := \{\pi_X\}_{X \in \mathbf{X}}$ where $\pi_X := \mathbb{1}_x(X)$, and therefore $\bar{\pi} = \mathbb{1}_x(\mathbf{X})$. We could show that the NPBs are tight in some cases. For example, in Example 1, provably, no better bounds for $\mathbb{E}_{P_\pi}[Y]$ than those defined by Prop. 2 could be derived (a result given in Prop. 10 in Appendix B). For other systems that involve variables that are “separated” in \mathcal{G} , however, better bounds may be derived by exploiting the implications of the causal graph and definition of the policy on the induced observational and interventional distributions. Consider the following example as a first illustration of this idea.

Example 2 (Tighter bounds with causal diagram). We are interested in evaluating the effect of a policy $\pi := \{\pi_{X_1}(x_1 | c), \mathbb{1}_{x_2}(X_2)\}$ from $P(x_1, c, x_2, y)$ and \mathcal{G} in Fig. 2a. In particular, this problem involves a stochastic conditional intervention on X_1 given C and a deterministic intervention on X_2 ; both variables having differing dependencies onto the rest of the system. We could show that,

$$\mathbb{E}_{P_\pi}[Y] = \mathbb{E}_{P_{\pi_{X_1}}}[Y] \geq \mathbb{E}_P[Y \pi_{X_1}] = \psi^\ell(P). \quad (3)$$

The inequality gives an expression for the bound in terms of P that exploits the fact that intervening on X_2 does not influence Y , and is tighter than the natural policy lower bound as $\psi^\ell(P) \geq \mathbb{E}_P[Y \pi_{X_1}] P(X_2 = x_2) = \mathbb{E}_P[Y \pi_{X_1} \mathbb{1}_{x_2}(X_2)]$ (= NPB). For the upper bound we could similarly establish that,

$$\mathbb{E}_{P_\pi}[Y] \leq \mathbb{E}_P[(Y - 1)\pi_{X_1}] + 1 = \psi^u(P). \quad (4)$$

which is smaller than the natural policy bound as $\psi^u(P) \leq \mathbb{E}_P[(Y - 1)\pi_{X_1} \mathbb{1}_{x_2}(X_2)] + 1$ (= NPB).

This example, although relatively straightforward, serves to illustrate the potential of causal diagrams (and the constraints they imply on the effect of policies) for defining tighter bounds.

3 Graphical Criteria for Partial Identification

This section aims to consider more general separation statements between variables encoded in a causal diagram and the decomposition they imply to provide a systematic algorithm to bound the effectiveness of policies. We start by introducing the notion of *partial adjustment sets* (Def. 3) that is applicable with multiple intervention variables.

Definition 3 (Partial adjustment set). *Let $\pi := \{\pi_{X_1}, \pi_{X_2}\}$ be a policy on $\{X_1, X_2\}$ with a conditioning set C . A set $W \subseteq V \setminus (X_1 \cup X_2 \cup C \cup Y)$ is said to be a partial adjustment set for π_{X_2} in \mathcal{G} if $(Y \perp\!\!\!\perp_d X_2 \mid W, C, X_1)_{\mathcal{G}_{\pi_{X_1}, X_2}}$.*

Proposition 3. *Let $\pi := \{\pi_{X_1}, \pi_{X_2}\}$ be a policy mapping a set of covariates C to a set of treatment variables $\{X_1, X_2\}$. Let W be a partial adjustment set for π_{X_2} in \mathcal{G} . Then,*

$$\mathbb{E}_P[Y \pi \gamma] \leq \mathbb{E}_{P_\pi}[Y] \leq \mathbb{E}_P[(Y - 1)\pi \gamma] + 1, \quad (5)$$

where $\gamma := \gamma(\mathbf{X}, C, W) = 1/P(X_2 \mid C, W)$.

In words, partial adjustment sets are designed to exploit the unconfounded status of some intervention variables with respect to the outcome, and could be leveraged to derive tighter bounds when available.

Example 3 (Tighter bounds with partial adjustment sets). Consider the problem of evaluating the effect of a policy $\pi := \{\pi_{X_1}(x_1 | c), \pi_{X_2}(x_2 | c)\}$ from $P(x_1, c, x_2, w, y)$ and \mathcal{G} in Fig. 2b. Following Def. 3, we could establish that $\mathbf{W} = \emptyset$ is a valid partial adjustment set for π_{X_2} since we can verify that $(Y \perp\!\!\!\perp_d X_2 | C, X_2)_{\mathcal{G}_{\pi_{X_1}, X_2}}$. Prop. 3 then gives us a valid expression for bounding the effect of the policy,

$$\mathbb{E}_{P_\pi}[Y] \geq \mathbb{E}_P[Y \bar{\pi} / P(X_2 | C)] \quad (\geq \mathbb{E}_P[Y \bar{\pi}], \text{ the NPB}). \quad (6)$$

For the upper bound we similarly find that,

$$\mathbb{E}_{P_\pi}[Y] \leq \mathbb{E}_P[(Y - 1)\bar{\pi} / P(X_2 | C)] + 1 \quad (\leq \mathbb{E}_P[(Y - 1)\bar{\pi}] + 1, \text{ the NPB}). \quad (7)$$

Next, we define a second useful notion, so-called *partial conditional instrumental variables sets*, that can be exploited to evaluate bounds in z -specific distributions $P(\mathbb{1}_z(\mathbf{Z})/P(\mathbf{Z}))$ instead of in P as described in Prop. 4.

Definition 4 (Partial conditional instrumental set). *A set $\mathbf{Z} \subseteq \mathbf{V}$ is said to be a partial instrumental set conditional on \mathbf{R} for a policy π in \mathcal{G} if $(Y \perp\!\!\!\perp_d \mathbf{Z} | \mathbf{R})_{\mathcal{G}_\pi}$.*

For illustration, we give a simple criterion below to show how this subgroup structure could be exploited to derive tighter bounds.

Proposition 4. *Let \mathbf{Z} be an unconditional partial instrumental set with respect to a policy π , i.e. $(Y \perp\!\!\!\perp_d \mathbf{Z})_{\mathcal{G}_\pi}$. Then,*

$$\max_z \mathbb{E}_P[Y \bar{\pi} \mathbb{1}_z(\mathbf{Z}) / P(\mathbf{Z})] \leq \mathbb{E}_{P_\pi}[Y] \leq \min_z \mathbb{E}_P[(Y - 1)\bar{\pi} \mathbb{1}_z(\mathbf{Z}) / P(\mathbf{Z})] + 1. \quad (8)$$

The following example shows that partial adjustment sets and partial conditional instrumental sets may be usefully combined to derive tighter bounds than would be available had each proposition (Props. 3 and 4) been applied in isolation.

Example 4 (Tighter bounds with partial adjustment and instrumental sets). For this example, consider the evaluation of an atomic intervention $\pi := \{\mathbb{1}_{x_1}(X_1), \mathbb{1}_{x_2}(X_2)\}$ in the causal diagram \mathcal{G} given in Fig. 2c. Note that $\bar{\pi} = \mathbb{1}_x(\mathbf{X})$. Following Def. 3, $\{\emptyset\}$ is a valid partial adjustment set for π_{X_1} since we can verify that $(Y \perp\!\!\!\perp_d X_2 | X_1)_{\mathcal{G}_{\pi_{X_1}, X_2}}$. Further, we could verify that $\{Z\}$ is a partial instrumental set conditional on $\{X_2\}$ with respect to π_{X_1} since $(Y \perp\!\!\!\perp_d Z | X_2)_{\mathcal{G}_{\pi_{X_1}}}$. These two separation statements in (manipulated versions of) \mathcal{G} could be leveraged to derive a tighter bound than previously considered:

$$\begin{aligned} \max_z \mathbb{E}_P[Y \bar{\pi} \mathbb{1}_z(Z) / P(X_2, Z)] &\leq \mathbb{E}_{P_\pi}[Y] \\ &\leq \min_z \mathbb{E}_P[(Y - 1)\bar{\pi} \mathbb{1}_z(Z) / P(X_2, Z)] + 1. \end{aligned}$$

To derive bounds in a more systematic fashion, combining the notions developed so far, we present Alg. 1 that recursively seeks to find partial adjustment and partial instrumental sets in an efficient and automatic manner.

Intuitively, Alg. 1 seeks to recursively simplify the query. First by omitting the intervened variables that have no effect on the outcome; second by finding the set of intervened variables for which a partial adjustment set could be used to tighten the bound, and third by finding the set of variables that act as valid partial instrumental sets to evaluate bounds on the most favorable conditional distributions rather than on the joint distribution.

Algorithm 1 Bounds for the effect of policies

Input: Graph \mathcal{G} , policy π , outcome Y .

Output: Bounds for $\mathbb{E}_{P_\pi}[Y]$.

- 1: Let $\mathbf{C} = \cup_{X \in \mathbf{X}} \mathbf{C}_X$, $\mathbf{K} = \mathbf{V} \setminus (\mathbf{X} \cup \mathbf{C} \cup Y)$.
 - /* 1. Omit redundant intervention variables. */
 - 2: Let $\mathbf{R} = \mathbf{X} \cap \text{An}(Y)$ in \mathcal{G}_π .
 - 3: Let $\pi_{\mathbf{R}} := \{\pi_X(\mathbf{c}_X)\}_{X \in \mathbf{R}}$.
 - /* 2. Find partial adjustment sets \mathbf{W} . */
 - 4: Initialize $\mathbf{W} = \emptyset$, $\mathbf{S} = \mathbf{R}$.
 - 5: **for** $R \in \mathbf{R}$ **do**
 - 6: Let $\mathbf{S} = \mathbf{S} \setminus R$.
 - 7: Let $\mathbf{W}_R = \text{An}(\mathbf{C} \cup Y) \cap \mathbf{K}$ in $\mathcal{G}_{\bar{\mathbf{S}}, R}$.
 - 8: **if** $(Y \perp\!\!\!\perp_d R | \mathbf{C}, \mathbf{S}, \mathbf{W}, \mathbf{W}_R)$ in $\mathcal{G}_{\bar{\mathbf{S}}, R}$ **then**
 - 9: $\mathbf{W} = \mathbf{W} \cup \mathbf{W}_R$.
 - 10: **end if**
 - 11: **end for**
 - /* 3. Find partial instrumental sets \mathbf{Z} . */
 - 12: Let $\mathbf{T} = \{R : \mathbf{W}_R \in \mathbf{W}\}$, $\mathbf{U} = \mathbf{R} \setminus \mathbf{T}$.
 - 13: Initialize $\mathbf{Z} = \emptyset$.
 - 14: **for** $Z \in \mathbf{K}$ **do**
 - 15: **if** $(Y \perp\!\!\!\perp_d Z | \mathbf{W} \setminus Z, \mathbf{Z}, \mathbf{T}, \mathbf{C})$ in $\mathcal{G}_{\bar{\mathbf{U}}}$ **then**
 - 16: $\mathbf{Z} = \mathbf{Z} \cup Z$.
 - 17: **end if**
 - 18: **end for**
 - /* 4. Return bounds. */
 - 19: Let $\gamma := \bar{\pi}_{\mathbf{R}} \mathbb{1}_z(\mathbf{Z}) / P(\mathbf{T}, \mathbf{Z} | \mathbf{W} \setminus \mathbf{Z}, \mathbf{C})$.
 - 20: Let $\psi_z^\ell := \mathbb{E}_P[Y \gamma]$.
 - 21: Let $\psi_z^u := \mathbb{E}_P[(Y - 1)\gamma] + 1$.
 - 22: Return bounds: $(\max_z \psi_z^\ell, \min_z \psi_z^u)$.
-

Proposition 5. *Alg. 1 is sound.*

In words, Prop. 5 says that for a given causal diagram, policy π , and joint distribution $P(v)$, the average treatment effect $\mathbb{E}_{P_\pi}[Y]$ is contained in the bounds produced by Alg. 1.

Proposition 6. *Let k be the number of variables and m be the number of edges in \mathcal{G} . The run time of Alg. 1 is $\mathcal{O}(k(k^2 + m))$.*

Example 5 (Steps of Alg. 1). For this example we consider evaluating a policy $\pi := \{\pi_{X_1}(\cdot | C), \pi_{X_2}(\cdot | C), \pi_{X_3}(\cdot | C)\}$ from observational data $P(x_1, c, x_2, x_3, a, w, y)$ compatible with \mathcal{G} in Fig. 2d by explicitly following Alg. 1. On line 1, we define $\mathbf{C} = \{C\}$, $\mathbf{K} = \{A, W\}$. We start by omitting potentially irrelevant intervention variables by evaluating $\mathbf{R} = \{X_1, X_2\}$ and noticing that $\mathbb{E}_{P_\pi}[Y] = \mathbb{E}_{P_{\pi_{\mathbf{R}}}}[Y]$. To find partial adjustment sets in line 4, we evaluate the for loop iterating over \mathbf{R} . For $R = X_1$, we find that $\mathbf{S} = X_2$, $\mathbf{W}_{X_1} = \{A, W\}$ but the if condition fails as Y is not d -separated from X_1 given $\{X_2, C, A, W\}$. We turn onto $R = X_2$, where $\mathbf{S} = X_1$, $\mathbf{W}_{X_2} = \{A, W\}$. The if condition is triggered as $(Y \perp_d X_2 | X_1, C, A, W)$ in $\mathcal{G}_{\pi_{X_1}, X_2}$ and therefore we update $\mathbf{W} = \mathbf{W}_{X_2} = \{A, W\}$ (giving a partial adjustment set). We continue with line 12, set $\mathbf{T} = \{X_2\}$, $\mathbf{U} = \{X_1\}$, and iterate over $\mathbf{K} = \{A, W\}$ in the for loop. For $Z = A$, we have that $(Y \perp_d A | W, X_2)$ in $\mathcal{G}_{\pi_{X_1}}$, the if condition is triggered and therefore we update $\mathbf{Z} = \{A\}$. For $Z = W$, we find that $(Y \perp_d W | A, X_2)$ in $\mathcal{G}_{\pi_{X_1}}$ and therefore terminate the for loop. Finally putting together the pieces, we evaluate the bounds to be,

$$\max_z \mathbb{E}_P[Y\gamma] \leq \mathbb{E}_{P_\pi}[Y] \leq \min_z \mathbb{E}_P[(Y-1)\gamma] + 1, \quad \gamma := \bar{\pi}_{\mathbf{R}} \mathbb{1}_a(A)/P(X_2, A | W, C) \quad (9)$$

Remark (Multiple bound expressions). Alg. 1 is designed to make use of partial adjustment and conditional instrumental variable sets but it will return a single bounding expression. In general multiple different expressions could be derived for a given causal diagram and data distribution. For example, given \mathcal{G} in Fig. 3 and a policy π on $\{X_1, X_2\}$, $\mathbb{E}_P[\bar{\pi}Y/P(X_2 | W_1)]$ and $\mathbb{E}_P[\bar{\pi}Y/P(X_2 | W_2)]$ give valid, but numerically different lower bounds for $\mathbb{E}_{P_\pi}[Y]$. The superiority of one bound over another might depend on the associations in the underlying distribution $P(v)$. In this case, we could construct two different SCMs compatible with \mathcal{G} that entail different orders in the numerical values for the two bounds, such that neither is always optimal. In general, we conjecture that finding an optimal analytical bound (using the proposed procedure) is undecidable from the graph alone.




Figure 3: \mathcal{G}

The observation in this remark motivates to develop a procedure to automatically enumerate all partial adjustment sets to facilitate the search for tighter bounds. The following proposition adapts the ListSep algorithm by [43] to enumerate all partial adjustment sets.

Proposition 7 (Enumerating partial adjustment sets). *Let $\pi := \{\pi_{X_1}, \pi_{X_2}\}$ be a policy on $\{X_1, X_2\}$ with a conditioning set \mathbf{C} . All partial adjustment sets may be enumerated in time $\mathcal{O}(k(k + m))$ where k are the number of variables and m be the number of edges in \mathcal{G} .*

Remark (Related methods for policy evaluation). In parallel to the graphical approach for encoding structural assumptions about the domain at hand, a number of works have adopted sensitivity assumptions that quantify the degree of unobserved confounding through various data statistics, such as odds ratios, propensity scores, etc. A rich literature on sensitivity assumptions exists, including Tan’s sensitivity model [38] and Rosenbaum’s sensitivity model [34]. Under these models, [17, 21, 28, 45], among others, present methods that achieve validity and rate properties for the resulting estimators. These approaches start with estimators that optimize the average or conditional treatment effect bounds subject to constraints implied by the sensitivity model. Several of these works leverage Neyman orthogonality techniques to obtain rate guarantees and doubly-robustness properties. For instance, [17, 21, 28] study the estimation of bounds under Tan’s sensitivity model which quantifies the degree of unobserved confounding through odds ratios, and propose estimators with various validity, sharpness, and favourable convergence rate guarantees. In contrast, [45] study the estimation of bounds under Rosenbaum’s sensitivity model also deriving estimators with sharpness and fast convergence guarantees. Some works under these assumptions have also considered policy evaluation (as opposed to the evaluation of atomic interventions) under various sensitivity models in the context of Reinforcement Learning, *e.g.*, [7, 22, 23]. We interpret this line of work as complementary (applicable under different assumptions, *i.e.*, sensitivity models rather than causal diagrams) to the techniques proposed in this paper.

4 Estimation of the Effect of Policies

This section aims to develop an estimation framework for the effect of a policy $\mathbb{E}_{P_\pi}[Y]$ given finite samples from P that partially identifies its value. The key observation of this section is that multiple characterizations for estimation could be derived for a given bound returned by lines 20-21 in Alg. 1.

In a first instance, parameterized by the probability ratio given in Alg. 1 defined by $\gamma = (\gamma_1, \gamma_2)$,

$$T^{\text{PW},\ell} := \mathbb{E}_P[\gamma_2 Y] \quad (= \psi_z^\ell), \quad \gamma_2 := \bar{\pi}_U \gamma_1, \quad \gamma_1 := \bar{\pi}_T \mathbb{1}_z(\mathbf{Z})/P(\mathbf{T}, \mathbf{Z} \mid \mathbf{W} \setminus \mathbf{Z}, \mathbf{C}). \quad (10)$$

In a second instance, parameterized by a collection of regression parameters $\boldsymbol{\mu} = (\mu_0, \mu_1, \tilde{\mu}_1, \mu_2)$ where

$$\mu_2 := \mu_2(\mathbf{R}, \mathbf{C}, \mathbf{W}, \mathbf{Z}) = \mathbb{E}_P[Y \mid \mathbf{R}, \mathbf{C}, \mathbf{W}, \mathbf{Z}], \quad (11)$$

$$\tilde{\mu}_1 := \tilde{\mu}_1(\mathbf{R}, \mathbf{C}, \mathbf{W}, \mathbf{Z}) = \bar{\pi}_U(\mathbf{U} \mid \mathbf{C})\mu_2(\mathbf{R}, \mathbf{C}, \mathbf{W}, \mathbf{Z}), \quad (12)$$

$$\mu_1 := \mu_1(\mathbf{T}, \mathbf{C}, \mathbf{W}, \mathbf{Z}) = \mathbb{E}_P[\tilde{\mu}_1(\mathbf{R}, \mathbf{C}, \mathbf{W}, \mathbf{Z}) \mid \mathbf{T}, \mathbf{C}, \mathbf{W}, \mathbf{Z}], \quad (13)$$

$$\mu_0 := \mu_0(\mathbf{C}, \mathbf{W}, \mathbf{Z}) = \sum_{\mathbf{t}} \mu_1(\mathbf{t}, \mathbf{C}, \mathbf{W}, \mathbf{Z}) \bar{\pi}_T(\mathbf{t} \mid \mathbf{C}). \quad (14)$$

$T^{\text{REG},\ell} := \mathbb{E}_P[\mu_0(\mathbf{C}, \mathbf{W}, \mathbf{z})]$ could be shown to equal ψ_z^ℓ .

Both formulations define equivalent but different estimation targets for the lower bound³ and may be combined leveraging the double machine learning (DML) toolkit for more efficient and robust inferences [10].

The following procedure is the main contribution of this section. It defines a DML estimator for bounding the effectiveness $\mathbb{E}_{P_\pi}[Y]$ of a policy π .

Definition 5 (DML Estimator). *Given π and \mathcal{G} , let $\{\mathbf{R}, \mathbf{T}, \mathbf{W}, \mathbf{C}, \mathbf{Z}, Y\}$ be defined as in Alg. 1. Consider a finite sample of data $\mathcal{D} \sim P$, randomly split into K folds. The k 'th partition of the sample is denoted $\mathcal{D}^{(k)}$ and $\mathcal{D}^{(-k)} := \mathcal{D} \setminus \mathcal{D}^{(k)}$. For each k , learn approximate nuisances $(\hat{\gamma}_k, \hat{\boldsymbol{\mu}}_k)$ with $\mathcal{D}^{(-k)}$. Then, define*

$$\left(\max_{\mathbf{z}} \hat{T}^{\text{DML},\ell}, \min_{\mathbf{z}} \hat{T}^{\text{DML},u} \right) \quad (15)$$

to be an estimate for the bounds on $\mathbb{E}_{P_\pi}[Y]$ where,

$$\hat{T}^{\text{DML},\ell} := \frac{1}{K} \sum_{k=1}^K \mathbb{E}_{\mathcal{D}^{(k)}}[\hat{\gamma}_{2,k}\{Y - \hat{\mu}_{2,k}\}] + \mathbb{E}_{\mathcal{D}^{(k)}}[\hat{\gamma}_{1,k}\{\hat{\mu}_{1,k} - \hat{\mu}_{1,k}\}] + \mathbb{E}_{\mathcal{D}^{(k)}}[\hat{\mu}_{0,k}]$$

$$\hat{T}^{\text{DML},u} := 1 + \frac{1}{K} \sum_{k=1}^K \mathbb{E}_{\mathcal{D}^{(k)}}[\hat{\gamma}_{2,k}\{(Y - 1) - \hat{\mu}_{2,k}\}] + \mathbb{E}_{\mathcal{D}^{(k)}}[\hat{\gamma}_{1,k}\{\hat{\mu}_{1,k} - \hat{\mu}_{1,k}\}] + \mathbb{E}_{\mathcal{D}^{(k)}}[\hat{\mu}_{0,k}].$$

To analyse the error of the DML estimator, we consider the case that nuisances can be estimated consistently. Thus requirement is relatively mild in practice as accurate probability estimation employing off-the-shelf classification and regression methods is feasible in general. Its error with respect to the true bounds is given by the following proposition.

Proposition 8 (Error rates). *Suppose the nuisance estimates $(\hat{\boldsymbol{\mu}}, \hat{\boldsymbol{\gamma}})$ are L_2 -consistent and bounded. Then, the error of the DML estimator $\hat{T}^{\text{DML}} \in \{\hat{T}^{\text{DML},\ell}, \hat{T}^{\text{DML},u}\}$ in Def. 5 is given as follows*

$$\hat{T}^{\text{DML}} - T^{\text{DML}} = \frac{1}{K} \sum_{k=1}^K R_k + O_P\left(\|\hat{\gamma}_{2,k} - \gamma_2\| \|\hat{\mu}_{2,k} - \mu_2\|\right) + O_P\left(\|\hat{\gamma}_{1,k} - \gamma_1\| \|\hat{\mu}_{1,k} - \mu_1\|\right)$$

where R_k is a random variable that converges to zero at a rate $\mathcal{O}_P(1/\sqrt{n})$.

In words, the DML estimator exhibits a robustness property since the error of \hat{T}^{DML} is bounded in probability at $n^{-1/2}$ rate whenever the nuisances converge at a rate $n^{-1/4}$. Note that the term

³For the upperbound, a similar construction could be derived by defining $T^{\text{PW},u} := \mathbb{E}_P[\gamma_2(Y - 1)] + 1$, and $T^{\text{REG},u} := \mathbb{E}_P[\mu_0(\mathbf{C}, \mathbf{W}, \mathbf{z})] + 1$ with Y replaced by $Y - 1$ in the definition of μ_2 . Both estimators could be shown to be unbiased, i.e. equal to ψ_u^ℓ defined in Alg. 1.

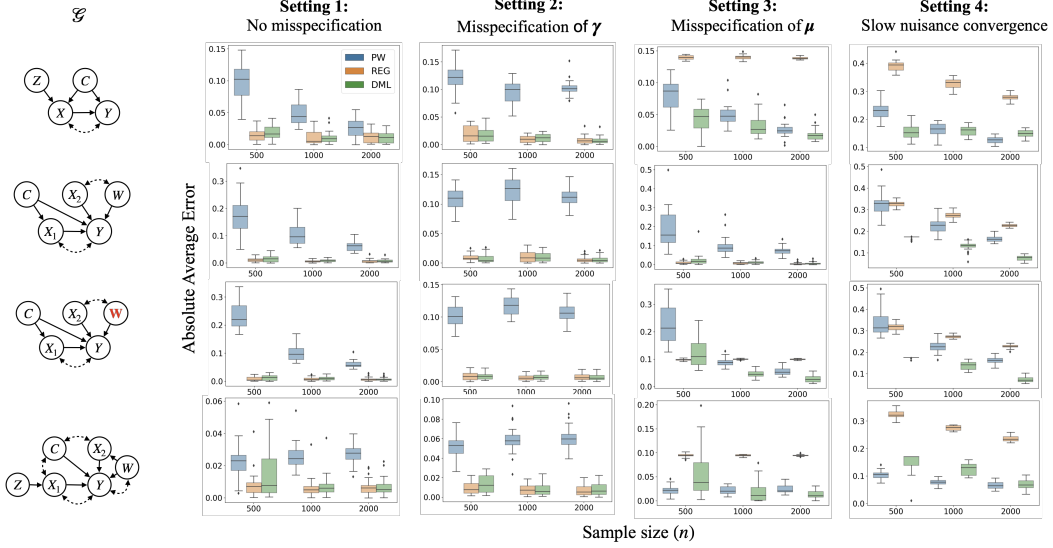


Figure 4: Experimental results on bounding the effectiveness of policies with the proposed estimators. Different rows highlight evaluations on different data generating mechanisms: the first row tests estimation with a partial instrumental set, the second row tests estimation with a partial adjustment set, the third row tests estimation with a high-dimensional partial adjustment set ($W \in \mathbb{R}^{100}$), and the fourth row tests estimation with a combination of partial adjustment and instrumental sets.

$\|\hat{\mu}_{1,k} - \hat{\mu}_{1,k}\|$ quantifies the error in approximating the conditional expectation in Eq. (13) for a given $\hat{\mu}_1$ estimated at that stage, and that the estimation error of $\hat{\gamma}_2$ and $\hat{\gamma}_2$ are equivalent since they are deterministic transformations of each other. The following proposition is a corollary that shows that the DML estimator is unbiased under misspecification.

Proposition 9 (Bias under misspecification). *Suppose either $\hat{\gamma}_1 = \gamma_1$ or $\hat{\mu}_2 = \mu_2$ and that either $\hat{\gamma}_1 = \gamma_1$ or $\hat{\mu}_1 = \hat{\mu}_1$. Then, $\hat{T}^{DML} \in \{\hat{T}^{DML,\ell}, \hat{T}^{DML,u}\}$ is an unbiased estimator of the corresponding bound defined in Alg. 1.*

5 Experiments

This section evaluates the quality of policy effect estimation from finite samples. Our goal is to illustrate the computation of bounds and provide empirical evidence of the fast convergence and robustness to misspecification of estimators.

Finite sample bounds are estimated with gradient boosting classification and regression models (for conditional expectations) or by taking sample averages (for unconditional expectations). We truncated estimates of probability mass functions in the interval $[0.01, 0.99]$ to ensure positivity. We assess the quality of an estimator \hat{T} by computing the absolute average error (AAE) with respect to (a proxy for) the true bounds (estimated in practice with larger sample sizes), *i.e.*, $AAE = |\hat{T}^u - \min_z \psi_z^u| + |\hat{T}^\ell - \max_z \psi_z^\ell|$. Throughout, we report various statistics: 25th, 50th, 75th percentile, etc., across evaluation runs with ten different random seeds. Further details of the simulations are provided in Appendix C.

5.1 Synthetic Simulations

The synthetic simulations consider 4 data generating mechanisms constructed according to the graphs in Fig. 4. They highlight the use of partial instrumental sets, partial adjustment sets, high-dimensional partial adjustment sets, and combinations of partial instrumental and adjustment sets of varying dimensionality. The task is to estimate bounds on the effectiveness of a policy $\pi := \{\pi_{X_1}, \pi_{X_2}\}$ defined as follows.

$$\pi_{X_1} := \pi_{X_1}(X_1 = 1) = 0.5, \quad \pi_{X_2} := \pi_{X_2}(X_2 = 1 | c) = 1/(1 - \exp\{-c\}). \quad (16)$$

That is the policy assigns $X_1 \in \{0, 1\}$ randomly with probability 0.5, and assigns $X_2 \in \{0, 1\}$ as a function of C with the probability of $X_2 = 1$ increasing with the value of C .

To estimate bounds on $\mathbb{E}_{P_\pi}[Y]$, we consider the proposed estimators, labelled: T^{PW} , estimated with the nuisances in Eq. (10), T^{REG} , estimated with the nuisances in Eq. (11), and T^{DML} estimated with the procedure in Def. 5. Our evaluations test performance across 4 different settings designed to highlight various properties.

- **Setting 1:** All nuisances estimated correctly. This setting aims to show that all estimators converge to the bound of interest.
- **Setting 2:** Nuisances $\hat{\gamma}$ are sampled from a uniform distribution to induce misspecification in the estimation of γ .
- **Setting 3:** Nuisances $\hat{\mu}$ are sampled from a uniform distribution to induce misspecification in the estimation of μ . Settings 2 and 3 aim to demonstrate the doubly-robustness property of the DML estimator.
- **Setting 4:** Noise ϵ is introduced in the estimation of all nuisances (γ, μ) to emphasize error due to finite sample variation. Specifically, noise $\epsilon \sim \text{Normal}(n^{-\alpha}, n^{-\alpha})$, $\alpha = 1/4$, that induces a slower rate of convergence as a function of sample size, inspired by [24, 20]. This setting aims to show that the fast convergence behavior of the DML estimator compared to competing estimators.

The results are given in Fig. 4⁴. We observe that across all data generating mechanisms, estimators improve with the size of the dataset and converge under no misspecification (Setting 1) to the underlying bounds. It is interesting to note also the differing accuracy of estimators in the small sample regime. T^{PW} , based the estimation of a ratio of probabilities, can be unstable with low sample sizes if the ratio denominator is estimated to be close to zero while T^{REG} , based on a sequence of regression tasks, tends to be better behaved. T^{DML} in contrast is constructed as a combination of elements of T^{PW} and T^{REG} . In particular, note in Def. 5 the use of nuisances μ and γ . As a result, T^{DML} has quite a different performance profile. Settings 2 and 3 in Fig. 4 show that the DML estimator T^{DML} is robust to misspecification in either the nuisances μ or γ that highlights the robustness property. Further, when decaying noise is introduced in the estimation of nuisances (Setting 4), the DML estimator outperforms in general with a faster convergence rate. We also observe that performance remains close to optimal with high-dimensional variables \mathbf{W} , demonstrating that the DML estimator provides a practical toolkit for bounding in practice.

5.1.1 Width of Bounds According to Different Graphical Criteria

This section evaluates the width of the bounds returned by exploiting the different graphical criteria provided in Sec. 3. The simulations are based on the data generating mechanism described by causal diagram illustrated in the fourth row of Fig. 4. We consider evaluating the policy in Eq. (16) and compute the bounds obtained by applying Prop. 2 (most conservative), Prop. 3 (using the partial adjustment set W only), Prop. 4 (using the partial instrumental set Z only), and finally Alg. 1 (that combines all propositions and is the proposed approach). Fig. 5 gives the results over 10 seeds of the data and across multiple data sizes, highlighting the gain achieved by exploiting the causal structure using the proposed approaches.




Figure 5: Width of bounds.

Remark (Actual width of bounds in practice). The majority of our empirical evaluations are spent on evaluating the accuracy of different methods at estimating bounds, without addressing whether the returned bounds are actually informative. In practice, the graph structure can play an important role in tightening bounds but the actual width of the bounds in a particular problem are primarily driven by the distribution of data. Here is an example to make this more concrete.

For a given policy π , Prop. 2 defines tight bounds (under some circumstances) on $\mathbb{E}_{P_\pi}[Y]$. The width of this bound is given by $1 - \mathbb{E}_P[\bar{\pi}]$ that is ultimately driven by $P(x)$. This term may therefore evaluate to anything between 0 and 1 depending on $P(x)$ and π .

⁴For the first row, the policy evaluated is: $\pi_X := \pi_X(X = 1 | c) = 1/(1 - \exp\{-c\})$.




Figure 6: Health campaign evaluations.

For more complex causal structures, the bounds proposed in Props. 3 and 4 (and Alg. 1) reduce the width of the interval above. Loosely written, from $1 - \mathbb{E}_P[\bar{\pi}]$ to $1 - \mathbb{E}_P[\bar{\pi}] \times \alpha$ for some α that is a function of the joint distribution P and the structure of the graph. But again, the actual width of the interval is ultimately determined by the values of P and π that may be large or small depending on the value probabilities involved.

5.2 Evaluating Health Campaigns

This section illustrates the evaluation of lifestyle recommendations for the mitigation of obesity in individuals from Colombia, Peru and Mexico [30]. The data was collected from anonymous users using a web platform, and includes reported obesity levels measured according to BMI (Y), age (A), smoking status (S), frequency of consumption of high caloric food (H), whether individuals monitored their calorie intake (M), family history being overweight (F), exercise frequency (E), and time using technology devices (T). Obesity is a multi-factored medical condition for which several causes have been acknowledged in the literature; causal diagrams relating the variables above have been curated in several related studies [1, 9]. We considered these findings to construct the causal diagram in Fig. 6a that we assumed for this example.

We aim to study the effect of a health campaign designed to lower the intake of high caloric food (H) and increase the frequency of exercise (E) on obesity levels (Y). For instance, we could hypothesize that the campaign leads to an increase in the observed proportion of individuals rarely consuming of high caloric food (H) from 0.12 to 0.5 and that of individuals doing exercise (E) regularly from 0.05 to 0.5. These statements could be formulated as a stochastic policy $\pi^\alpha := \{\pi_H^\alpha, \pi_E^\alpha\}$ acting on H and E , with new assignments,

$$\pi_H^\alpha := \pi_H^\alpha(H = \text{rarely}) = \alpha, \quad \pi_E^\alpha := \pi_E^\alpha(E = \text{regularly}) = \alpha. \quad (17)$$

We consider the evaluation of expected BMI levels $\mathbb{E}_{P_{\pi^\alpha}}[Y]$ that range from 12 to 50 in the population, with a mean of 29.3. First note that this or other policies acting on (H, E) are not identifiable due to the bi-directed edge $\{H \leftrightarrow Y\}$, but may nevertheless be bounded using Alg. 1. We find that

$$\max_t \mathbb{E}_P[Y\gamma] \leq \mathbb{E}_{P_{\pi^\alpha}}[Y] \leq \min_t \mathbb{E}_P[(Y-1)\gamma] + 1, \quad \gamma := \bar{\pi}_{\pi^\alpha} \mathbb{1}_t(T)/P(E, T | A, S, F). \quad (18)$$

To illustrate the inference of bounds with the DML estimator, we consider evaluating policies with $\alpha = 0.2, 0.4, 0.6, 0.8$. Fig. 6b gives the results. The end-points of the intervals denote estimated lower and upper bounds. We see that policies that promote a healthier lifestyle (larger values of α) are expected to reduce obesity levels on average but substantial uncertainty is still expected.

6 Conclusions

The evaluation of policies is arguably one of the critical ingredients enabling more personalized decision-making systems. When the effect of policies is not identifiable, bounds can provide an effective support for making informed decisions. In this paper we developed partial identification and estimation tools for bounding the effect of a (stochastic or conditional) policy given data and assumptions encoded in a causal graph. We introduced several graphical characterizations that induce tighter bounds, and developed an estimation framework that exhibit robustness to noise and fast convergence. The results of this paper were illustrated through synthetic simulations and a real-world health campaign example for the reduction of obesity levels.

Acknowledgements

We thank the anonymous reviewers for helpful comments.

References

- [1] Steven Allender, Brynle Owen, Jill Kuhlberg, Janette Lowe, Phoebe Nagorcka-Smith, Jill Whelan, and Colin Bell. A community based systems diagram of obesity causes. *PloS one*, 10(7):e0129683, 2015.
- [2] Susan Athey and Guido W Imbens. The state of applied econometrics: Causality and policy evaluation. *Journal of Economic perspectives*, 31(2):3–32, 2017.
- [3] Vahid Balazadeh Meresht, Vasilis Syrgkanis, and Rahul G Krishnan. Partial identification of treatment effects with implicit generative models. *Advances in Neural Information Processing Systems*, 35:22816–22829, 2022.
- [4] Alexander Balke and Judea Pearl. Bounds on treatment effects from studies with imperfect compliance. *Journal of the American Statistical Association*, 92(439):1171–1176, 1997.
- [5] Alexis Bellot. Towards bounding causal effects under Markov equivalence. In *The 40th Conference on Uncertainty in Artificial Intelligence*. PMLR, 2024.
- [6] Alexis Bellot, Anish Dhir, and Giulia Prando. Generalization bounds and algorithms for estimating conditional average treatment effect of dosage. *arXiv preprint arXiv:2205.14692*, 2022.
- [7] Andrew Bennett and Nathan Kallus. Policy evaluation with latent confounders via optimal balance. *Advances in neural information processing systems*, 32, 2019.
- [8] Rohit Bhattacharya, Razieh Nabi, and Ilya Shpitser. Semiparametric inference for causal effects in graphical models with hidden variables. *The Journal of Machine Learning Research*, 23(1):13325–13400, 2022.
- [9] Bryony Butland, Susan Jebb, Peter Kopelman, Klim McPherson, Sandy Thomas, Jane Mardell, Vivienne Parry, et al. *Tackling obesities: future choices-project report*, volume 10. Citeseer, 2007.
- [10] Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. Double/debiased machine learning for treatment and structural parameters, 2018.
- [11] David Maxwell Chickering and Judea Pearl. A clinician’s tool for analyzing non-compliance. In *Proceedings of the National Conference on Artificial Intelligence*, pages 1269–1276, 1996.
- [12] Juan Correa and Elias Bareinboim. A calculus for stochastic interventions: Causal effect identification and surrogate experiments. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 10093–10100, 2020.
- [13] Elizabeth K Drake, Steve Aos, and Marna G Miller. Evidence-based public policy options to reduce crime and criminal justice costs: Implications in washington state. *Victims and offenders*, 4(2):170–196, 2009.
- [14] Noam Finkelstein and Ilya Shpitser. Deriving bounds and inequality constraints using logical relations among counterfactuals. In *Conference on Uncertainty in Artificial Intelligence*, pages 1348–1357. PMLR, 2020.
- [15] Limor Gultchin, Virginia Aglietti, Alexis Bellot, and Silvia Chiappa. Functional causal Bayesian optimization. In *Uncertainty in Artificial Intelligence*, pages 756–765. PMLR, 2023.
- [16] Yaowei Hu, Yongkai Wu, Lu Zhang, and Xintao Wu. A generative adversarial framework for bounding confounded causal effects. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 12104–12112, 2021.

- [17] Andrew Jesson, Sören Mindermann, Yarin Gal, and Uri Shalit. Quantifying ignorance in individual-level causal-effect estimates under hidden confounding. In *International Conference on Machine Learning*, pages 4829–4838. PMLR, 2021.
- [18] Shalmali Joshi, Junzhe Zhang, and Elias Bareinboim. Towards safe policy learning under partial identifiability: A causal approach. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 13004–13012, 2024.
- [19] Yonghan Jung, Jin Tian, and Elias Bareinboim. Estimating identifiable causal effects on Markov equivalence class through double machine learning. In *International Conference on Machine Learning*, pages 5168–5179. PMLR, 2021.
- [20] Yonghan Jung, Jin Tian, and Elias Bareinboim. Estimating joint treatment effects by combining multiple experiments. In *International Conference on Machine Learning*, pages 15451–15527. PMLR, 2023.
- [21] Nathan Kallus, Xiaojie Mao, and Angela Zhou. Interval estimation of individual-level causal effects under unobserved confounding. In *The 22nd international conference on artificial intelligence and statistics*, pages 2281–2290. PMLR, 2019.
- [22] Nathan Kallus and Angela Zhou. Confounding-robust policy improvement. *Advances in neural information processing systems*, 31, 2018.
- [23] Nathan Kallus and Angela Zhou. Confounding-robust policy evaluation in infinite-horizon reinforcement learning. *Advances in neural information processing systems*, 33:22293–22304, 2020.
- [24] Edward H Kennedy. Towards optimal doubly robust estimation of heterogeneous causal effects. *arXiv preprint arXiv:2004.14497*, 2020.
- [25] Travis Mandel, Yun-En Liu, Sergey Levine, Emma Brunskill, and Zoran Popovic. Offline policy evaluation across representations with applications to educational games. In *AAMAS*, volume 1077, 2014.
- [26] Henry B Mann and Abraham Wald. On stochastic limit and order relationships. *The Annals of Mathematical Statistics*, 14(3):217–226, 1943.
- [27] Charles F Manski. Nonparametric bounds on treatment effects. *The American Economic Review*, 80(2):319–323, 1990.
- [28] Miruna Oprescu, Jacob Dorn, Marah Ghoummaid, Andrew Jesson, Nathan Kallus, and Uri Shalit. B-learner: Quasi-oracle bounds on heterogeneous causal effects under hidden confounding. *arXiv preprint arXiv:2304.10577*, 2023.
- [29] Kirtan Padh, Jakob Zeitler, David Watson, Matt Kusner, Ricardo Silva, and Niki Kilbertus. Stochastic causal programming for bounding treatment effects. *arXiv preprint arXiv:2202.10806*, 2022.
- [30] Fabio Mendoza Palechor and Alexis de la Hoz Manotas. Dataset for estimation of obesity levels based on eating habits and physical condition in individuals from colombia, peru and mexico. *Data in brief*, 25:104344, 2019.
- [31] Judea Pearl. *Causality*. Cambridge university press, 2009.
- [32] Judea Pearl and James M Robins. Probabilistic evaluation of sequential plans from causal models with hidden variables. In *UAI*, volume 95, pages 444–453. Citeseer, 1995.
- [33] James M Robins. The analysis of randomized and non-randomized aids treatment trials using a new approach to causal inference in longitudinal studies. *Health service research methodology: a focus on AIDS*, pages 113–159, 1989.
- [34] Paul R Rosenbaum, P Briskman Rosenbaum, and Briskman. *Design of observational studies*, volume 10. Springer, 2010.

- [35] Andrea Rotnitzky, James Robins, and Lucia Babino. On the multiply robust estimation of the mean of the g-functional. *arXiv preprint arXiv:1705.08582*, 2017.
- [36] Uri Shalit, Fredrik D Johansson, and David Sontag. Estimating individual treatment effect: generalization bounds and algorithms. In *International conference on machine learning*, pages 3076–3085. PMLR, 2017.
- [37] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- [38] Zhiqiang Tan. A distributional approach for causal inference using propensity scores. *Journal of the American Statistical Association*, 101(476):1619–1637, 2006.
- [39] Philip Thomas and Emma Brunskill. Data-efficient off-policy policy evaluation for reinforcement learning. In *International Conference on Machine Learning*, pages 2139–2148. PMLR, 2016.
- [40] Jin Tian. Identifying dynamic sequential plans. *arXiv preprint arXiv:1206.3292*, 2012.
- [41] Anastasios A Tsiatis, Marie Davidian, Shannon T Holloway, and Eric B Laber. *Dynamic treatment regimes: Statistical methods for precision medicine*. CRC press, 2019.
- [42] Mark J van der Laan and Susan Gruber. Targeted minimum loss based estimation of causal effects of multiple time point interventions. *The international journal of biostatistics*, 8(1), 2012.
- [43] Benito van der Zander. *Algorithmics of Identifying Causal Effects in Graphical Models*. PhD thesis, Universität zu Lübeck, 2020.
- [44] Evert Vedung. *Public policy and program evaluation*. Routledge, 2017.
- [45] Steve Yadlowsky, Hongseok Namkoong, Sanjay Basu, John Duchi, and Lu Tian. Bounds on the conditional and average treatment effect with unobserved confounding factors. *arXiv preprint arXiv:1808.09521*, 2018.
- [46] Junzhe Zhang. Designing optimal dynamic treatment regimes: A causal reinforcement learning approach. In *International Conference on Machine Learning*, pages 11012–11022. PMLR, 2020.
- [47] Junzhe Zhang and Elias Bareinboim. Bounding causal effects on continuous outcome. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 12207–12215, 2021.
- [48] Junzhe Zhang, Jin Tian, and Elias Bareinboim. Partial counterfactual identification from observational and experimental data. *arXiv preprint arXiv:2110.05690*, 2021.
- [49] Yao Zhang, Alexis Bellot, and Mihaela Schaar. Learning overlapping representations for the estimation of individualized treatment effects. In *International Conference on Artificial Intelligence and Statistics*, pages 1005–1014. PMLR, 2020.

Appendix

Table of Contents

A Preliminaries, Related Work, and Impact Statement	15
A.1 Preliminaries	15
A.2 Related Work	15
A.3 Broader Impact Statement	16
B Proofs	18
B.1 Proofs of statements in Sec. 2	18
B.2 Proofs of statements in Sec. 4	25
C Details on experiments	30
C.1 Simulations	30
C.2 Health Campaign Evaluation	31
D NeurIPS Paper Checklist	33

A Preliminaries, Related Work, and Impact Statement

A.1 Preliminaries

For the derivation of results, we use the notion of counterfactuals defined as follows.

For an SCM \mathcal{M} , arbitrary subsets of endogenous variables \mathbf{X}, \mathbf{Y} , the potential outcome of \mathbf{Y} to intervention $do(\mathbf{x})$, denoted by $\mathbf{Y}_x(\mathbf{u})$, is the solution for \mathbf{Y} with $\mathbf{U} = \mathbf{u}$ in the sub-model \mathcal{M}_x . It can be read as the counterfactual sentence “the value that \mathbf{Y} would have obtained in situation $\mathbf{U} = \mathbf{u}$, had \mathbf{X} been \mathbf{x} ?”. Statistically, averaging \mathbf{u} over the distribution $P(\mathbf{u})$ leads to the counterfactual variables \mathbf{Y}_x . The distribution of the variable \mathbf{Y}_x is denoted $P(\mathbf{Y}_x)$. With this formalism, we can express all the quantities in the main body of this paper in the language of counterfactuals. For instance, the distribution of \mathbf{Y} in sub-model \mathcal{M}_x can alternatively be written $P(\mathbf{Y}_x) \equiv P_x(\mathbf{Y})$. Moreover, $\mathbb{E}_P[f(\mathbf{Y}_x) \mid \mathbf{z}]$, also written $\mathbb{E}_{P(\cdot|\mathbf{z})}[f(\mathbf{Y}_x)]$, denotes the conditional expectation of $f(\mathbf{Y})$ over $P(\mathbf{y}_x \mid \mathbf{z})$. Similarly, the distribution over \mathbf{Y} in the sub-model \mathcal{M}_π can be written $P_{\mathcal{M}}(\mathbf{Y}_\pi)$, or $P(\mathbf{Y}_\pi) \equiv P_\pi(\mathbf{Y})$ for short. See [31, Chapter 7] for further context on counterfactuals.

Definition 6 (The Axioms of Counterfactuals, Chapter 7.3.1 [31]). *For any three sets of endogenous variables $\mathbf{X}, \mathbf{Y}, \mathbf{W}$ in a causal model and \mathbf{x}, \mathbf{w} in the domains of \mathbf{X} and \mathbf{W} , the following holds:*

- *Composition: $\mathbf{W}_x = \mathbf{w}$ implies that $\mathbf{Y}_{x,\mathbf{w}} = \mathbf{Y}_x$.*
- *Effectiveness: $\mathbf{X}_{\mathbf{w},\mathbf{x}} = \mathbf{x}$.*
- *Reversibility: $\mathbf{Y}_{x,\mathbf{w}} = \mathbf{y}$ and $\mathbf{W}_{x,\mathbf{y}} = \mathbf{w}$ imply that $\mathbf{Y}_x = \mathbf{y}$.*

Theorem 1 (Soundness and Completeness of the Axioms Theorems 7.3.3, 7.3.6 [31]). *The Axioms of counterfactuals are sound and complete for all causal models.*

The following rules to manipulate experimental distributions produced by policies extend the do-calculus and will be used for the proof of several theoretical statements [12].

Theorem 2 (Inference Rules σ -calculus [12]). *Let \mathcal{G} be a causal diagram compatible with an SCM \mathcal{M} , with endogenous variables \mathbf{V} . For any disjoint subsets $\mathbf{X}, \mathbf{Y}, \mathbf{Z} \subseteq \mathbf{V}$, two disjoint subsets $\mathbf{T}, \mathbf{W} \subseteq \mathbf{V} \setminus (\mathbf{Z} \cup \mathbf{Y})$ (i.e., possibly including \mathbf{X}), the following rules are valid for any intervention strategies $\pi_{\mathbf{X}}, \pi_{\mathbf{Z}}$, and $\pi'_{\mathbf{Z}}$ such that $\mathcal{G}_{\pi_{\mathbf{X}}\pi_{\mathbf{Z}}}, \mathcal{G}_{\pi_{\mathbf{X}}\pi'_{\mathbf{Z}}}$ have no cycles:*

- *Rule 1 (Insertion/Deletion of observations):*

$$P_{\pi_{\mathbf{X}}}(\mathbf{y} \mid \mathbf{w}, \mathbf{t}) = P_{\pi_{\mathbf{X}}}(\mathbf{y} \mid \mathbf{w}) \quad \text{if } (\mathbf{T} \perp\!\!\!\perp_d \mathbf{Y} \mid \mathbf{W}) \text{ in } \mathcal{G}_{\pi_{\mathbf{X}}}.$$

- *Rule 2 (Change of regimes under observation):*

$$P_{\pi_{\mathbf{X}},\pi_{\mathbf{Z}}}(\mathbf{y} \mid \mathbf{z}, \mathbf{w}) = P_{\pi_{\mathbf{X}},\pi'_{\mathbf{Z}}}(\mathbf{y} \mid \mathbf{z}, \mathbf{w}) \quad \text{if } (\mathbf{Y} \perp\!\!\!\perp_d \mathbf{Z} \mid \mathbf{W}) \text{ in } \mathcal{G}_{\pi_{\mathbf{X}},\pi_{\mathbf{Z}}}, \underline{\mathbf{z}} \text{ and } \mathcal{G}_{\pi_{\mathbf{X}},\pi'_{\mathbf{Z}}}, \underline{\mathbf{z}}$$

- *Rule 3 (Change of regimes without observation):*

$$P_{\pi_{\mathbf{X}},\pi_{\mathbf{Z}}}(\mathbf{y} \mid \mathbf{w}) = P_{\pi_{\mathbf{X}},\pi'_{\mathbf{Z}}}(\mathbf{y} \mid \mathbf{w}) \quad \text{if } (\mathbf{Y} \perp\!\!\!\perp_d \mathbf{Z} \mid \mathbf{W}) \text{ in } \mathcal{G}_{\pi_{\mathbf{X}},\pi_{\mathbf{Z}}}, \overline{\mathbf{Z}(\mathbf{W})} \text{ and } \mathcal{G}_{\pi_{\mathbf{X}},\pi'_{\mathbf{Z}}}, \overline{\mathbf{Z}(\mathbf{W})}$$

where $\mathbf{Z}(\mathbf{W})$ is the set of elements in \mathbf{Z} that are not ancestors of \mathbf{W} in $\mathcal{G}_{\pi_{\mathbf{X}}}$.

A.2 Related Work

Active experimentation by physically manipulating reality is generally not feasible for many consequential applications. This motivates the study of causal effect identification and estimation. The former studies the problem of inferring whether a unique expression or bound of the desired causal effect could be derived from the available data and assumptions. The latter studies the problem of providing efficient estimators from finite samples to compute causal effects or bounds on causal effects in practice. In the following, we review these two lines of work to better contextualize our contribution.

Partial identification. The natural bounds were derived to demonstrate that useful inference about causal effects could be drawn without making identifying assumptions beyond the observed data [27]. Relatedly, an analysis of bounds of causal effects was also provided for studies with imperfect

compliance under a set of instrumental variable assumptions that are not sufficient however to identify the causal effect of interest [33]. These bounds may be written in closed form and have recently been extended as a more general strategy for bounding causal effects given assumptions encoded in causal diagrams in discrete systems [5, 46]. Recent work has also considered bounds in closed form with access to both observational and interventional distributions [18]. These techniques were developed alongside a second line of research that employs a polynomial optimization program to compute causal bounds given a causal diagram [4]. They proposed a family of canonical models parameterized according to the causal diagram, reducing the bounding problem to a series of equivalent linear programs. [11] further used Bayesian techniques to investigate the sharpness of these bounds with regard to the observational sample size. Recently, [14, 48] describe a polynomial programming approach to solve the partial identification for general causal graphs. More recent proposals consider parameterizations in the space of linear combinations of a set of fixed basis functions [29] and neural networks [3, 16].

In parallel, a number of works have adopted sensitivity assumptions (as an alternative or in combination with a causal diagram) that quantify the degree of unobserved confounding through various data statistics, such as odds ratios, propensity scores, etc. A rich literature on sensitivity assumptions exists, including Tan’s sensitivity model [38] and Rosenbaum’s sensitivity model [34]. Under these models, [17, 21, 28, 45], among others, present methods that achieve validity and rate properties for the resulting estimators. These approaches start with estimators that optimize the average or conditional treatment effect bounds subject to constraints implied by the sensitivity model. Several of these works leverage Neyman orthogonality techniques to obtain rate guarantees and doubly-robustness properties. For instance, [17, 21, 28] study the estimation of bounds under Tan’s sensitivity model which quantifies the degree of unobserved confounding through odds ratios, and propose estimators with various validity, sharpness, and favourable convergence rate guarantees. In contrast, [45] study the estimation of bounds under Rosenbaum’s sensitivity model also deriving estimators with sharpness and fast convergence guarantees.

While some works have considered policy evaluation under various sensitivity models in the context of Reinforcement Learning, e.g., [7, 22, 23], most works target estimation of bounds on the effect of atomic interventions. We interpret this line of work as complementary (applicable under different assumptions, i.e., sensitivity models rather than causal diagrams) to the techniques proposed in this paper.

Causal effect estimation. With a causal diagram as input, causal effect estimation has traditionally focused on a subset of identifiable scenarios, relying on assumptions such as backdoor criterion or the availability of adjustment sets. Identification expressions in those cases are given by (sequential) covariate adjustments that have lead to statistically appealing estimators from observational data [32]. Notable examples are doubly robust estimators [10, 35, 42]. Recently these techniques have been extended to settings identifiable from multiple experimental distributions [8, 19, 20]. To our knowledge, no estimation framework specific to the estimation of analytical expressions derived for bounds on the effect of policies or atomic interventions has been developed.

A.3 Broader Impact Statement

Our work investigates the conditions under which policies may be bounded from observational data. Issues of policy evaluation are central to fields of application involved in decision-making as well as AI and ML. We believe that a tool to bound the effect of stochastic and conditional policies in systems with unobserved confounding is an important addition to the scientific toolbox. Reasoning instead without acknowledging for the potential of unobserved confounding may lead researchers to operate on a more heuristical basis. For instance, the hypothetical effect of public policies might be misrepresented if informed by the evaluation of an idealized atomic intervention that is difficult to implement in practice. And, once implemented, might have unintended consequences. With this background, we believe that research on the partial identification and estimation of policies based on stochastic or conditional interventions can help scientists and individuals make more informed decisions.

In this work, we start from the assumption that a causal diagram that is consistent with the underlying data generating system of interest is available. In general, this requires domain knowledge. While some of the bounds provided do not require full knowledge of the causal graph, whenever a d -separation is assumed its truth value should be justified by prior knowledge or experiment. It is

important also to make the distinction between the task of partial identification, that is inferring an expression for bounds on causal effects, and that of estimation, that is providing efficient estimators from finite samples to compute bounds in practice. This set of results concerns both of these tasks. The first objective of our procedure is to provide an expression for lower and upper bounds, irrespective of the accuracy with which one is able to approximate $P(\mathbf{V})$ from finite samples. The second objective is introduce efficient estimators for bounds on the effect of policies using finite samples from $P(\mathbf{V})$. In higher-dimensional systems, the computational complexity of estimating the conditional distributions that define lower and upper bounds on causal effects is a substantial challenge. Consequently, practitioners must exercise caution when deploying the proposed method in small sample scenarios where estimators may be inaccurate. Moreover, we have stated our convergence guarantees in the infinite sample limit, without quantifying the finite-sample estimation uncertainty. Finite-sample properties could be explored similarly to [24] given a particular choice of function class to extend our results with high-probability bounds. Finally, we emphasize that simulations on real and synthetic data are provided for illustration purposes only. These results do not recommend or advocate for the implementation of a particular policy, and should be considered in practice in combination with other aspects of the decision-making process.

B Proofs

This section presents proofs for theoretical statements in the main body of this paper.

B.1 Proofs of statements in Sec. 2

Prop. 1 restated. For any \mathbf{x} , $\mathbb{E}_P[Y \mathbb{1}_{\mathbf{x}}(\mathbf{X})] \leq \mathbb{E}_{P_{\mathbf{x}}}[Y] \leq \mathbb{E}_P[(Y - 1) \mathbb{1}_{\mathbf{x}}(\mathbf{X})] + 1$.

Proof. Consider the derivation of the lower bound.

$$\begin{aligned} \mathbb{E}_P[Y_{\mathbf{x}}] &= \sum_{\mathbf{x}'} \mathbb{E}_P[Y_{\mathbf{x}} | \mathbf{x}'] P(\mathbf{x}') \\ &\geq \mathbb{E}_P[Y_{\mathbf{x}} | \mathbf{x}] P(\mathbf{x}) \\ &\stackrel{(1)}{=} \mathbb{E}_P[Y | \mathbf{x}] P(\mathbf{x}) \\ &= \sum_y y P(y | \mathbf{x}) P(\mathbf{x}) \\ &= \sum_{y, \mathbf{x}} y \mathbb{1}_{\mathbf{x}}(\mathbf{x}) P(y, \mathbf{x}) \\ &= \mathbb{E}_P[Y \mathbb{1}_{\mathbf{x}}(\mathbf{X})] \end{aligned}$$

(1) follows by the composition axiom of counterfactuals.

Consider the derivation of the upper bound. It holds that,

$$\begin{aligned} \mathbb{E}_P[Y_{\mathbf{x}}] &= \sum_{\mathbf{x}'} \mathbb{E}_P[Y_{\mathbf{x}} | \mathbf{x}'] P(\mathbf{x}') \\ &= \mathbb{E}_P[Y_{\mathbf{x}} | \mathbf{x}] P(\mathbf{x}) + \sum_{\mathbf{x}' \neq \mathbf{x}} \mathbb{E}_P[Y_{\mathbf{x}} | \mathbf{x}'] P(\mathbf{x}') \\ &\stackrel{(1)}{\leq} \mathbb{E}_P[Y | \mathbf{x}] P(\mathbf{x}) + \sum_{\mathbf{x}' \neq \mathbf{x}} P(\mathbf{x}') \\ &= \mathbb{E}_P[Y | \mathbf{x}] P(\mathbf{x}) + 1 - P(\mathbf{x}) \\ &= \mathbb{E}_P[Y \mathbb{1}_{\mathbf{x}}(\mathbf{X})] + 1 - \mathbb{E}_P[\mathbb{1}_{\mathbf{x}}(\mathbf{X})] \\ &= \mathbb{E}_P[(Y - 1) \mathbb{1}_{\mathbf{x}}(\mathbf{X})] + 1. \end{aligned}$$

(1) follows by the boundedness of Y , here assumed bounded by 1 for simplifying the derivation. \square

Proposition 10. *The natural bounds are tight in general.*

Proof. We show this statement by providing a pair of SCMs that agree on the input observational distribution but evaluate to the lower and upper bounds, respectively, specified by the natural bounds. Let $\mathbb{M}(\mathcal{G})$ denote the space of SCMs that induce the causal diagram \mathcal{G} in Fig. 1a. We introduce a pair of SCMs compatible with the causal diagram \mathcal{G} that evaluate to lower and upper bounds respectively. Let X be binary, $Y \in [0, 1]$, and $C \in \mathbb{R}$, and consider $\mathcal{M}_1, \mathcal{M}_2 \in \mathbb{M}(\mathcal{G})$ defined as follows,

$$\mathcal{M}_1 := \begin{cases} x := f_X(u_1) \\ c := f_C(u_2) \\ y := \begin{cases} f_Y(x, c, u_1, u_2) & \text{if } x = f_X(u_1), \\ 0 & \text{otherwise.} \end{cases} \end{cases}$$

and,

$$\mathcal{M}_2 := \begin{cases} x := f_X(u_1) \\ c := f_C(u_2) \\ y := \begin{cases} f_Y(x, c, u_1, u_2) & \text{if } x = f_X(u_1), \\ 1 & \text{otherwise.} \end{cases} \end{cases}$$

For a $P_{\mathcal{M}_1}(\mathbf{u}) = P_{\mathcal{M}_2}(\mathbf{u})$, both SCMs agree on observational distributions $P_{\mathcal{M}_1}(x, y, \mathbf{c}) = P_{\mathcal{M}_2}(x, y, \mathbf{c})$. However the following derivations show that the interventional expectation $\mathbb{E}_{P_{\mathcal{M}_1}}[Y \mid do(x = 1)]$ differs across models: for \mathcal{M}_1 equal to the analytical lower bound, and for \mathcal{M}_2 equal to the analytical upper bound demonstrating that (in this case) the bound is tight. In particular,

$$\begin{aligned}
& \mathbb{E}_{P_{\mathcal{M}_1}}[Y \mid do(x = 1)] \\
&= \mathbb{E}_{P_{\mathcal{M}_1}}[Y \mid x = 1, u : x = f_X(u)]P(u : x = f_X(u)) \\
&+ \mathbb{E}_{P_{\mathcal{M}_1}}[Y \mid x = 1, u : x \neq f_X(u)]P(u : x \neq f_X(u)) \\
&= \mathbb{E}_{P_{\mathcal{M}_1}}[Y \mid x = 1]P(x = 1) \\
&= \mathbb{E}_P[Y \mid x = 1]P(x = 1) \\
&= \mathbb{E}_P[Y \mathbb{1}_1(X)] \\
& \mathbb{E}_{P_{\mathcal{M}_2}}[Y \mid do(x = 1)] \\
&= \mathbb{E}_{P_{\mathcal{M}_2}}[Y \mid x = 1, u : x = f_X(u)]P(u : x = f_X(u)) \\
&+ \mathbb{E}_{P_{\mathcal{M}_2}}[Y \mid x = 1, u : x \neq f_X(u)]P(u : x \neq f_X(u)) \\
&= \mathbb{E}_{P_{\mathcal{M}_2}}[Y \mid x = 1]P(x = 1) + \mathbb{E}_{P_{\mathcal{M}_2}}[Y \mid x = 1, u : x = f_X(u)]P(u : x = f_X(u)) \\
&= \mathbb{E}_{P_{\mathcal{M}_2}}[Y \mid x = 1]P(x = 1) + 1 - P(x = 1) \\
&= \mathbb{E}_P[(Y - 1)\mathbb{1}_1(X)] + 1.
\end{aligned}$$

□

Prop. 2 restated. For any π , $\mathbb{E}_P[Y\bar{\pi}] \leq \mathbb{E}_{P_\pi}[Y] \leq \mathbb{E}_P[(Y - 1)\bar{\pi}] + 1$.

Proof. Consider the derivation of the lower bound. Let π denote a policy from an arbitrary set of covariates \mathbf{C} to \mathbf{X} . By marginalizing over \mathbf{X} , \mathbf{C} we find that

$$\mathbb{E}_{P_\pi}[Y] = \mathbb{E}_P[Y_\pi] = \int \mathbb{E}_P[Y_\pi \mid \mathbf{x}, \mathbf{c}] \prod_{X \in \mathbf{X}} \pi_X(x \mid \mathbf{c}_X) P(\mathbf{c}) d\mathbf{c} d\mathbf{x}.$$

\mathbf{C}_X denotes the subset of \mathbf{C} that is used to inform the intervention on X . Moreover,

$$\begin{aligned}
\mathbb{E}_P[Y_\pi \mid \mathbf{x}, \mathbf{c}] &\stackrel{(1)}{=} \mathbb{E}_P[Y_{\mathbf{x}} \mid \mathbf{c}] \\
&= \sum_{\mathbf{x}'} \mathbb{E}_P[Y_{\mathbf{x}} \mid \mathbf{x}', \mathbf{c}] P(\mathbf{x}' \mid \mathbf{c}) \\
&\geq \mathbb{E}_P[Y_{\mathbf{x}} \mid \mathbf{x}, \mathbf{c}] P(\mathbf{x} \mid \mathbf{c}) \\
&\stackrel{(2)}{=} \mathbb{E}_P[Y \mid \mathbf{x}, \mathbf{c}] P(\mathbf{x} \mid \mathbf{c}).
\end{aligned}$$

(1) holds by Rule 2 in Thm. 2 by swapping the policy π by the do intervention $do(\mathbf{X} = \mathbf{x})$: given that the policy acts on \mathbf{X} taking as inputs \mathbf{C} it is always true that $Y \perp\!\!\!\perp_d \mathbf{X} \mid \mathbf{C}$ in $\mathcal{G}_{\pi \underline{\mathbf{X}}}$ and $\mathcal{G}_{\underline{\mathbf{X}} \overline{\mathbf{X}}}$ in any graph \mathcal{G} . The following equality follows by marginalizing over the domain of \mathbf{X} , which is assumed discrete. (2) follows by the composition axiom of counterfactuals. By combining these two expressions we get

$$\begin{aligned}
\mathbb{E}_P[Y_\pi] &\geq \int \prod_{X \in \mathbf{X}} \pi(x \mid \mathbf{c}_X) P(\mathbf{c}) \mathbb{E}_P[Y \mid \mathbf{x}, \mathbf{c}] P(\mathbf{x} \mid \mathbf{c}) d\mathbf{c} d\mathbf{x} \\
&= \int \bar{\pi} y P(y \mid \mathbf{x}, \mathbf{c}) P(\mathbf{x} \mid \mathbf{c}) P(\mathbf{c}) dy d\mathbf{c} d\mathbf{x} \\
&= \mathbb{E}_P[Y\bar{\pi}].
\end{aligned}$$

Consider the derivation of the upper bound. It holds that,

$$\begin{aligned}
\mathbb{E}_P[Y_\pi | \mathbf{x}, \mathbf{c}] &= \mathbb{E}_P[Y_{\mathbf{x}} | \mathbf{c}] \\
&= \sum_{\mathbf{x}'} \mathbb{E}_P[Y_{\mathbf{x}} | \mathbf{x}', \mathbf{c}] P(\mathbf{x}' | \mathbf{c}) \\
&= \mathbb{E}_P[Y_{\mathbf{x}} | \mathbf{x}, \mathbf{c}] P(\mathbf{x} | \mathbf{c}) + \sum_{\mathbf{x}' \neq \mathbf{x}} \mathbb{E}_P[Y_{\mathbf{x}} | \mathbf{x}', \mathbf{c}] P(\mathbf{x}' | \mathbf{c}) \\
&\stackrel{(1)}{\leq} \mathbb{E}_P[Y | \mathbf{x}, \mathbf{c}] P(\mathbf{x} | \mathbf{c}) + \sum_{\mathbf{x}' \neq \mathbf{x}} P(\mathbf{x}' | \mathbf{c}) \\
&= \mathbb{E}_P[Y | \mathbf{x}, \mathbf{c}] P(\mathbf{x} | \mathbf{c}) + 1 - P(\mathbf{x} | \mathbf{c}).
\end{aligned}$$

(1) follows by the boundedness of Y , here assumed bounded by 1 for simplifying the derivation. As above by combining this inequality with the decomposition of the policy effect we get

$$\begin{aligned}
\mathbb{E}_{P_\pi}[Y] &\leq \int \prod_{X \in \mathbf{X}} \pi(x | \mathbf{c}_X) P(\mathbf{c}) \{ \mathbb{E}_P[Y | \mathbf{x}, \mathbf{c}] P(\mathbf{x} | \mathbf{c}) + 1 - P(\mathbf{x} | \mathbf{c}) \} d\mathbf{c} d\mathbf{x} \\
&= \mathbb{E}_P[Y \bar{\pi}] + 1 - \mathbb{E}_P[\bar{\pi}] \\
&= \mathbb{E}_P[(Y - 1) \bar{\pi}] + 1.
\end{aligned}$$

□

Prop. 3 restated. Let $\pi := \{\pi_{\mathbf{X}_1}, \pi_{\mathbf{X}_2}\}$ be a policy mapping a set of covariates \mathbf{C} to a set of treatment variables $\{\mathbf{X}_1, \mathbf{X}_2\}$. Let \mathbf{W} be a partial adjustment set for $\pi_{\mathbf{X}_2}$ in \mathcal{G} . Then,

$$\mathbb{E}_P[Y \bar{\pi} \gamma] \leq \mathbb{E}_{P_\pi}[Y] \leq \mathbb{E}_P[(Y - 1) \bar{\pi} \gamma] + 1,$$

where $\gamma := \gamma(\mathbf{X}, \mathbf{C}, \mathbf{W}) = 1/P(\mathbf{X}_2 | \mathbf{C}, \mathbf{W})$.

Proof. Consider the derivation of the lower bound. By marginalizing over \mathbf{X}, \mathbf{C} we find that

$$\mathbb{E}_{P_\pi}[Y] = \mathbb{E}_P[Y_\pi] = \int \mathbb{E}_P[Y_\pi | \mathbf{x}, \mathbf{c}] \prod_{X \in \mathbf{X}} \pi_X(x | \mathbf{c}_X) P(\mathbf{c}) d\mathbf{c} d\mathbf{x}.$$

\mathbf{C}_X denotes the subset of \mathbf{C} that is used to inform the intervention on X . Denote $\bar{\pi} := \prod_{X \in \mathbf{X}} \pi_X(x | \mathbf{c}_X)$. Let $\mathbf{X}_1, \mathbf{X}_2$ be a partition of \mathbf{X} such that $\mathbf{W} \subseteq \mathbf{V} \setminus (\mathbf{X} \cup \mathbf{C} \cup Y)$ is partial adjustment set for $\pi_{\mathbf{X}_2}$. Then,

$$\begin{aligned}
\mathbb{E}_P[Y_\pi | \mathbf{x}, \mathbf{c}] &= \mathbb{E}_P[Y_{\pi_{\mathbf{X}_1}, \pi_{\mathbf{X}_2}} | \mathbf{x}, \mathbf{c}] \\
&= \int \mathbb{E}_P[Y_{\mathbf{x}_1, \mathbf{x}_2} | \mathbf{c}, \mathbf{w}] P(\mathbf{w} | \mathbf{c}) d\mathbf{w} \\
&\stackrel{(1)}{=} \int \mathbb{E}_P[Y_{\mathbf{x}_1} | \mathbf{x}_2, \mathbf{w}, \mathbf{c}] P(\mathbf{w} | \mathbf{c}) d\mathbf{w} \\
&= \int \sum_{\mathbf{x}'_1} \{ \mathbb{E}_P[Y_{\mathbf{x}_1} | \mathbf{x}'_1, \mathbf{x}_2, \mathbf{w}, \mathbf{c}] P(\mathbf{x}'_1 | \mathbf{x}_2, \mathbf{w}, \mathbf{c}) \} P(\mathbf{w} | \mathbf{c}) d\mathbf{w} \\
&\geq \int \mathbb{E}_P[Y_{\mathbf{x}_1} | \mathbf{x}_1, \mathbf{x}_2, \mathbf{w}, \mathbf{c}] P(\mathbf{x}_1 | \mathbf{x}_2, \mathbf{w}, \mathbf{c}) P(\mathbf{w} | \mathbf{c}) d\mathbf{w} \\
&= \int \mathbb{E}_P[Y | \mathbf{x}_1, \mathbf{x}_2, \mathbf{w}, \mathbf{c}] P(\mathbf{x}_1 | \mathbf{x}_2, \mathbf{w}, \mathbf{c}) P(\mathbf{w} | \mathbf{c}) d\mathbf{w}.
\end{aligned}$$

(1) follows by the definition of \mathbf{X}_2 as a partial adjustment set: it holds that $(Y \perp\!\!\!\perp_d \mathbf{X}_2 | \mathbf{W}, \mathbf{C}, \mathbf{X}_1)$ in $\mathcal{G}_{\overline{\mathbf{X}_1} \mathbf{X}_2}$ which induces the equality $\mathbb{E}_P[Y_{\mathbf{x}_1, \mathbf{x}_2} | \mathbf{c}, \mathbf{w}] = \mathbb{E}_P[Y_{\mathbf{x}_1} | \mathbf{x}_2, \mathbf{w}, \mathbf{c}]$. The last equality follows by the composition axiom of counterfactuals. Combining this expression with the decomposition of the policy effect we get,

$$\begin{aligned}
\mathbb{E}_P[Y_\pi] &\geq \int \prod_{X \in \mathbf{X}} \pi(x | \mathbf{c}_X) P(\mathbf{c}) \mathbb{E}_P[Y | \mathbf{x}_1, \mathbf{x}_2, \mathbf{w}, \mathbf{c}] P(\mathbf{x}_1 | \mathbf{x}_2, \mathbf{w}, \mathbf{c}) P(\mathbf{w} | \mathbf{c}) d\mathbf{w} d\mathbf{c} d\mathbf{x} \\
&= \int \bar{\pi} y P(y | \mathbf{x}_1, \mathbf{x}_2, \mathbf{w}, \mathbf{c}) P(\mathbf{x}_1 | \mathbf{x}_2, \mathbf{w}, \mathbf{c}) P(\mathbf{w} | \mathbf{c}) P(\mathbf{c}) dy d\mathbf{c} d\mathbf{x} \\
&= \mathbb{E}_P[Y \bar{\pi} / P(\mathbf{X}_2 | \mathbf{W}, \mathbf{C})].
\end{aligned}$$

For the upper bound, consider the following derivation,

$$\begin{aligned}
& \mathbb{E}_P[Y_\pi \mid \mathbf{x}, \mathbf{c}] \\
&= \int \mathbb{E}_P[Y_{\mathbf{x}_1} \mid \mathbf{x}_2, \mathbf{w}, \mathbf{c}] P(\mathbf{w} \mid \mathbf{c}) d\mathbf{w} \\
&= \int \sum_{\mathbf{x}'_1} \{ \mathbb{E}_P[Y_{\mathbf{x}_1} \mid \mathbf{x}'_1, \mathbf{x}_2, \mathbf{w}, \mathbf{c}] P(\mathbf{x}'_1 \mid \mathbf{x}_2, \mathbf{w}, \mathbf{c}) \} P(\mathbf{w} \mid \mathbf{c}) d\mathbf{w} \\
&= \int \left\{ \mathbb{E}_P[Y_{\mathbf{x}_1} \mid \mathbf{x}_1, \mathbf{x}_2, \mathbf{w}, \mathbf{c}] P(\mathbf{x}_1 \mid \mathbf{x}_2, \mathbf{w}, \mathbf{c}) P(\mathbf{w} \mid \mathbf{c}) \right. \\
&\quad \left. + \sum_{\mathbf{x}'_1 \neq \mathbf{x}_1} \mathbb{E}_P[Y_{\mathbf{x}_1} \mid \mathbf{x}'_1, \mathbf{x}_2, \mathbf{w}, \mathbf{c}] P(\mathbf{x}'_1 \mid \mathbf{x}_2, \mathbf{w}, \mathbf{c}) P(\mathbf{w} \mid \mathbf{c}) \right\} d\mathbf{w} \\
&\leq \int \left\{ \mathbb{E}_P[Y \mid \mathbf{x}_1, \mathbf{x}_2, \mathbf{w}, \mathbf{c}] P(\mathbf{x}_1 \mid \mathbf{x}_2, \mathbf{w}, \mathbf{c}) P(\mathbf{w} \mid \mathbf{c}) + \sum_{\mathbf{x}'_1 \neq \mathbf{x}_1} P(\mathbf{x}'_1 \mid \mathbf{x}_2, \mathbf{w}, \mathbf{c}) P(\mathbf{w} \mid \mathbf{c}) \right\} d\mathbf{w} \\
&= \int \left\{ \mathbb{E}_P[Y \mid \mathbf{x}_1, \mathbf{x}_2, \mathbf{w}, \mathbf{c}] P(\mathbf{x}_1 \mid \mathbf{x}_2, \mathbf{w}, \mathbf{c}) P(\mathbf{w} \mid \mathbf{c}) + (1 - P(\mathbf{x}_1 \mid \mathbf{x}_2, \mathbf{w}, \mathbf{c})) P(\mathbf{w} \mid \mathbf{c}) \right\} d\mathbf{w} \\
&= \int \mathbb{E}_P[Y \mid \mathbf{x}_1, \mathbf{x}_2, \mathbf{w}, \mathbf{c}] P(\mathbf{x}_1 \mid \mathbf{x}_2, \mathbf{w}, \mathbf{c}) P(\mathbf{w} \mid \mathbf{c}) d\mathbf{w} + 1 - \int P(\mathbf{x}_1 \mid \mathbf{x}_2, \mathbf{w}, \mathbf{c}) P(\mathbf{w} \mid \mathbf{c}) d\mathbf{w}.
\end{aligned}$$

The inequality follows from the boundedness of Y and the rest of the arguments are analogous to the lower bound derivation. Combining this expression with the decomposition of the policy effect implies that,

$$\mathbb{E}_{P_\pi}[Y] \leq \mathbb{E}_P[(Y - 1)\bar{\pi}/P(\mathbf{X}_2 \mid \mathbf{W}, \mathbf{C})] + 1.$$

□

Prop. 4 restated. Let \mathbf{Z} be an unconditional partial instrumental set with respect to a policy π , i.e. $(Y \perp_d \mathbf{Z})_{\mathcal{G}_\pi}$. Then,

$$\begin{aligned}
\mathbb{E}_{P_\pi}[Y] &\geq \max_{\mathbf{z}} \mathbb{E}_P[Y \bar{\pi} \mathbb{1}_{\mathbf{z}}(\mathbf{Z})/P(\mathbf{Z})] \\
\mathbb{E}_{P_\pi}[Y] &\leq \min_{\mathbf{z}} \mathbb{E}_P[(Y - 1)\bar{\pi} \mathbb{1}_{\mathbf{z}}(\mathbf{Z})/P(\mathbf{Z})] + 1.
\end{aligned}$$

Proof. Consider the derivation of the lower bound. Let π denote a policy from an arbitrary set of covariates \mathbf{C} to \mathbf{X} . Given that $(Y \perp_d \mathbf{Z})_{\mathcal{G}_\pi}$, by marginalizing over \mathbf{X}, \mathbf{C} we find that

$$\mathbb{E}_{P_\pi}[Y] = \mathbb{E}_P[Y_\pi] = \mathbb{E}_P[Y_\pi \mid \mathbf{z}] = \int \mathbb{E}_P[Y_\pi \mid \mathbf{x}, \mathbf{c}, \mathbf{z}] \prod_{X \in \mathbf{X}} \pi_X(x \mid \mathbf{c}_X) P(\mathbf{c} \mid \mathbf{z}) d\mathbf{c} d\mathbf{x}.$$

\mathbf{C}_X denotes the subset of \mathbf{C} that is used to inform the intervention on X . Moreover,

$$\begin{aligned}
\mathbb{E}_P[Y_\pi \mid \mathbf{x}, \mathbf{c}, \mathbf{z}] &\stackrel{(1)}{=} \mathbb{E}_P[Y_{\mathbf{x}} \mid \mathbf{c}, \mathbf{z}] \\
&= \sum_{\mathbf{x}'} \mathbb{E}_P[Y_{\mathbf{x}} \mid \mathbf{x}', \mathbf{c}, \mathbf{z}] P(\mathbf{x}' \mid \mathbf{c}, \mathbf{z}) \\
&\geq \mathbb{E}_P[Y_{\mathbf{x}} \mid \mathbf{x}, \mathbf{c}, \mathbf{z}] P(\mathbf{x} \mid \mathbf{c}, \mathbf{z}) \\
&\stackrel{(2)}{=} \mathbb{E}_P[Y \mid \mathbf{x}, \mathbf{c}, \mathbf{z}] P(\mathbf{x} \mid \mathbf{c}, \mathbf{z}).
\end{aligned}$$

(1) holds by Rule 2 in Thm. 2 by swapping the policy π by the do intervention $do(\mathbf{X} = \mathbf{x})$: given that the policy acts on \mathbf{X} taking as inputs \mathbf{C} it is always true that $Y \perp_d \mathbf{X} \mid \mathbf{C}$ in $\mathcal{G}_{\pi \mathbf{X}}$ and $\mathcal{G}_{\mathbf{X} \bar{\mathbf{X}}}$ in any graph \mathcal{G} . The following equality follows by marginalizing over the domain of \mathbf{X} , which is assumed discrete. (2) follows by the composition axiom of counterfactuals. By combining these two

expressions we get

$$\begin{aligned}
\mathbb{E}_P[Y_\pi] &\geq \int \prod_{X \in \mathbf{X}} \pi(x | \mathbf{c}_X) P(\mathbf{c} | \mathbf{z}) \mathbb{E}_P[Y | \mathbf{x}, \mathbf{c}, \mathbf{z}] P(\mathbf{x} | \mathbf{c}, \mathbf{z}) d\mathbf{c} d\mathbf{x} \\
&= \int \bar{\pi} y P(y | \mathbf{x}, \mathbf{c}, \mathbf{z}) P(\mathbf{x} | \mathbf{c}, \mathbf{z}) P(\mathbf{c} | \mathbf{z}) dy d\mathbf{c} d\mathbf{x} \\
&= \mathbb{E}_P[Y \bar{\pi} | \mathbf{z}] \\
&= \mathbb{E}_P[Y \bar{\pi} \mathbb{1}_z(\mathbf{Z}) / P(\mathbf{Z})].
\end{aligned}$$

And since the l.h.s. does not depend on \mathbf{Z} we can further tighten the bound by writing,

$$\mathbb{E}_P[Y_\pi] \geq \max_{\mathbf{z}} \mathbb{E}_P[Y \bar{\pi} \mathbb{1}_z(\mathbf{Z}) / P(\mathbf{Z})].$$

Consider the derivation of the upper bound. It holds that,

$$\begin{aligned}
\mathbb{E}_P[Y_\pi | \mathbf{x}, \mathbf{c}, \mathbf{z}] &= \mathbb{E}_P[Y_{\mathbf{x}} | \mathbf{c}, \mathbf{z}] \\
&= \sum_{\mathbf{x}'} \mathbb{E}_P[Y_{\mathbf{x}} | \mathbf{x}', \mathbf{c}, \mathbf{z}] P(\mathbf{x}' | \mathbf{c}, \mathbf{z}) \\
&= \mathbb{E}_P[Y_{\mathbf{x}} | \mathbf{x}, \mathbf{c}, \mathbf{z}] P(\mathbf{x} | \mathbf{c}, \mathbf{z}) + \sum_{\mathbf{x}' \neq \mathbf{x}} \mathbb{E}_P[Y_{\mathbf{x}} | \mathbf{x}', \mathbf{c}, \mathbf{z}] P(\mathbf{x}' | \mathbf{c}, \mathbf{z}) \\
&\stackrel{(1)}{\leq} \mathbb{E}_P[Y | \mathbf{x}, \mathbf{c}, \mathbf{z}] P(\mathbf{x} | \mathbf{c}, \mathbf{z}) + \sum_{\mathbf{x}' \neq \mathbf{x}} P(\mathbf{x}' | \mathbf{c}, \mathbf{z}) \\
&= \mathbb{E}_P[Y | \mathbf{x}, \mathbf{c}, \mathbf{z}] P(\mathbf{x} | \mathbf{c}, \mathbf{z}) + 1 - P(\mathbf{x} | \mathbf{c}, \mathbf{z}).
\end{aligned}$$

(1) follows by the boundedness of Y , here assumed bounded by 1 for simplifying the derivation. As above by combining this inequality with the decomposition of the policy effect we get

$$\begin{aligned}
\mathbb{E}_{P_\pi}[Y] &\leq \int \prod_{X \in \mathbf{X}} \pi(x | \mathbf{c}_X) P(\mathbf{c} | \mathbf{z}) \{ \mathbb{E}_P[Y | \mathbf{x}, \mathbf{c}, \mathbf{z}] P(\mathbf{x} | \mathbf{c}, \mathbf{z}) + 1 - P(\mathbf{x} | \mathbf{c}, \mathbf{z}) \} d\mathbf{c} d\mathbf{x} \\
&= \mathbb{E}_P[Y \bar{\pi} | \mathbf{z}] + 1 - \mathbb{E}_P[\bar{\pi} | \mathbf{z}] \\
&= \mathbb{E}_P[(Y - 1) \bar{\pi} | \mathbf{z}] + 1 \\
&= \mathbb{E}_P[(Y - 1) \bar{\pi} \mathbb{1}_z(\mathbf{Z}) / P(\mathbf{Z})] + 1.
\end{aligned}$$

And since the l.h.s. does not depend on \mathbf{Z} we can further tighten the bound by writing,

$$\mathbb{E}_P[Y_\pi] \geq \min_{\mathbf{z}} \mathbb{E}_P[(Y - 1) \bar{\pi} \mathbb{1}_z(\mathbf{Z}) / P(\mathbf{Z})] + 1.$$

□

Prop. 5 restated. *Alg. 1 is sound.*

Proof. For the soundness of Alg. 1, we will consider each operation in turn and show that recovered adjustment sets and conditional instrumental sets lead to a valid bound.

1. *Omit redundant intervention variables.* For a policy $\pi := \{\pi_X(\mathbf{c}_X)\}_{X \in \mathbf{X}}$ denote $\pi_{\mathbf{R}} := \{\pi_X(\mathbf{c}_X)\}_{X \in \mathbf{R}}$, $\mathbf{R} \subseteq \mathbf{X}$. Let $\mathbf{R} = \mathbf{X} \cap \text{An}(Y)$ in \mathcal{G}_π . By Rule 3 of the σ -calculus we have that,

$$\mathbb{E}_{P_\pi}[Y] = \mathbb{E}_{P_{\pi_{\mathbf{R}}}}[Y],$$

since $Y \perp\!\!\!\perp_d \mathbf{X} \setminus \mathbf{R}$ in $\mathcal{G}_{\pi, \overline{\mathbf{X} \setminus \mathbf{R}}}$ and $\mathcal{G}_{\pi_{\mathbf{R}}, \overline{\mathbf{X} \setminus \mathbf{R}}}$ as there are no directed paths from $\mathbf{X} \setminus \mathbf{R}$ to Y by definition of \mathbf{R} . The first two lines therefore reduce the number of intervention variables. We proceed to bound the equivalent $\mathbb{E}_{P_{\pi_{\mathbf{R}}}}[Y]$.

2. *Find partial adjustment sets \mathbf{W} .* Line 4-11 consider finding partial adjustment sets by traversing the set of treatment variables $R \in \mathbf{R}$. Recall that

$$\mathbb{E}_{P_{\pi_{\mathbf{R}}}}[Y] = \mathbb{E}_P[Y_{\pi_{\mathbf{R}}}] = \int \mathbb{E}_P[Y_{\pi_{\mathbf{R}}} | \mathbf{r}, \mathbf{c}] \prod_{X \in \mathbf{R}} \pi_X(x | \mathbf{c}_X) P(\mathbf{c}) d\mathbf{c} d\mathbf{r}.$$

Denote $\bar{\pi}_R := \prod_{X \in \mathbf{R}} \pi_X(x \mid \mathbf{c}_X)$. In words lines 4-11 iteratively considers the existence of separators sets between $R \in \mathbf{R}$ and Y to reduce the scope of the policy on Y , i.e. Y_{π_R} . By [43, Lemma 3.18.], if there exists a separator \mathbf{W} between two sets of variables \mathbf{X} and \mathbf{Y} such that $S \subseteq \mathbf{W} \subseteq \mathbf{K}$ in a graph \mathcal{G} , i.e. $\mathbf{X} \perp_d \mathbf{Y} \mid \mathbf{W}$ in \mathcal{G} , then $\mathbf{W} = An(\mathbf{X} \cup \mathbf{Y} \cup S) \cap \mathbf{K}$ is also a valid separator. The definition of \mathbf{W}_R in line 7 aims to leverage this result to check for the existence of separators in our case.

Consider a $R \in \mathbf{R}$, let $S = \mathbf{R} \setminus R$, $\mathbf{K} = \mathbf{V} \setminus (\mathbf{X} \cup \mathbf{C} \cup \mathbf{Y})$. In particular, if there exist a set $\{\mathbf{C}, \mathbf{R}\} \subset \mathbf{W} \subset \mathbf{K}$ such that $(Y \perp_d R \mid \mathbf{W})$ in $\mathcal{G}_{\pi_S, \bar{R}}$ then $\mathbf{W}_R = An(\mathbf{R} \cup \mathbf{Y} \cup \mathbf{C}) \cap \mathbf{K} = An(\mathbf{Y} \cup \mathbf{C}) \cap \mathbf{K}$ in \mathcal{G}_{π_S} satisfies $(Y \perp_d R \mid \mathbf{W}_R)$ in $\mathcal{G}_{\pi_S, \bar{R}}$, i.e. \mathbf{W}_R is a separator if one exists. Note that \mathbf{W}_R does not include any member of $\{\mathbf{C}, \mathbf{R}, \mathbf{Y}\}$. Assume this separation holds and that the if statement in line 8 is triggered. Then,

$$\begin{aligned} \mathbb{E}_P[Y_{\pi_R} \mid \mathbf{r}, \mathbf{c}] &= \mathbb{E}_P[Y_r \mid \mathbf{c}] \\ &= \mathbb{E}_P[Y_{s,r} \mid \mathbf{c}] \\ &= \int \mathbb{E}_P[Y_{s,r} \mid \mathbf{c}, \mathbf{w}_R] P(\mathbf{w}_R \mid \mathbf{c}) d\mathbf{w}_R \\ &= \int \mathbb{E}_P[Y_s \mid r, \mathbf{c}, \mathbf{w}_R] P(\mathbf{w}_R \mid \mathbf{c}) d\mathbf{w}_R. \end{aligned}$$

The last equality holds by assumption $(Y \perp_d R \mid \mathbf{W}_R)$ in $\mathcal{G}_{\bar{S}, \bar{R}}$.

Now considering the term $\mathbb{E}_P[Y_s \mid r, \mathbf{c}, \mathbf{w}_R]$, for the second pass through the for loop, we can see that if the d -separation statement in the if condition is fulfilled we can further reduce the scope of the intervention on S . In particular, for $R' \in \mathbf{R} \setminus R$, assuming that the if statement is triggered for $\mathbf{W}_{R'}$ we have that,

$$\begin{aligned} \mathbb{E}_P[Y_s \mid r, \mathbf{c}, \mathbf{w}_R] &= \int \mathbb{E}_P[Y_s \mid r, \mathbf{c}, \mathbf{w}_R, \mathbf{w}_{R'}] P(\mathbf{w}_{R'} \mid \mathbf{w}_R, \mathbf{c}) d\mathbf{w}_{R'} \\ &= \int \mathbb{E}_P[Y_{s \setminus r'} \mid r', r, \mathbf{c}, \mathbf{w}_R, \mathbf{w}_{R'}] P(\mathbf{w}_{R'} \mid \mathbf{w}_R, \mathbf{c}) d\mathbf{w}_{R'}. \end{aligned}$$

Upon reaching the end of the for loop we have for \mathbf{W} defined in line 9, $\mathbf{T} = \{R : \mathbf{W}_R \in \mathbf{W}\}$, $\mathbf{U} = \mathbf{R} \setminus \mathbf{T}$ that,

$$\mathbb{E}_P[Y_\pi \mid \mathbf{r}, \mathbf{c}] = \int \mathbb{E}_P[Y_u \mid \mathbf{t}, \mathbf{c}, \mathbf{w}] P(\mathbf{w} \mid \mathbf{c}) d\mathbf{w}.$$

3. *Find partial instrumental sets \mathbf{Z} .* Starting line 12, we look for partial conditional instrumental sets. Consider first the case that a potential instruments $Z \in \mathbf{W}$ is evaluated in the for loop. On line 15, assume that the if statement is triggered, i.e. $(Y \perp_d Z \mid \mathbf{W} \setminus Z, \mathbf{T})$ in $\mathcal{G}_{\bar{U}}$. It then holds that,

$$\begin{aligned} \mathbb{E}_P[Y_\pi \mid \mathbf{r}, \mathbf{c}] &:= \int \mathbb{E}_P[Y_u \mid \mathbf{w}, \mathbf{t}, \mathbf{c}] P(\mathbf{w} \mid \mathbf{c}) d\mathbf{w} \\ &= \int \mathbb{E}_P[Y_u \mid \mathbf{w} \setminus z, z, \mathbf{t}, \mathbf{c}] \int_z P(\mathbf{w} \mid \mathbf{c}) d\mathbf{w} \\ &= \int \mathbb{E}_P[Y_u \mid \mathbf{w} \setminus z, z, \mathbf{t}, \mathbf{c}] P(\mathbf{w} \setminus z \mid \mathbf{c}) d\mathbf{w} \setminus z. \end{aligned}$$

In turn, for the case that some $Z' \notin \mathbf{W}$ is evaluated in the for loop, if $(Y \perp_d Z' \mid \mathbf{W}, \mathbf{T})$ in $\mathcal{G}_{\bar{U}}$,

$$\begin{aligned} \mathbb{E}_P[Y_\pi \mid \mathbf{r}, \mathbf{c}] &:= \int \mathbb{E}_P[Y_u \mid \mathbf{w}, \mathbf{t}, \mathbf{c}] P(\mathbf{w} \mid \mathbf{c}) d\mathbf{w} \\ &= \int \mathbb{E}_P[Y_u \mid \mathbf{w}, \mathbf{t}, z', \mathbf{c}] P(\mathbf{w} \mid \mathbf{c}) d\mathbf{w} \end{aligned}$$

For the \mathbf{Z} obtained in line 15, we derive that,

$$\mathbb{E}_P[Y_\pi \mid \mathbf{r}, \mathbf{c}] = \int \mathbb{E}_P[Y_u \mid \mathbf{w} \setminus z, z, \mathbf{t}, \mathbf{c}] P(\mathbf{w} \setminus z \mid \mathbf{c}) d\mathbf{w} \setminus z.$$

4. *Return bounds.* We now proceed to use the expression above to lower and upper bound the effect of interest. In particular, for the lower bound we have that,

$$\begin{aligned}
\mathbb{E}_P[Y_\pi | \mathbf{x}, \mathbf{c}] &= \int \mathbb{E}_P[Y_{\mathbf{u}} | \mathbf{w} \setminus \mathbf{z}, \mathbf{z}, \mathbf{t}, \mathbf{c}] P(\mathbf{w} \setminus \mathbf{z} | \mathbf{c}) d\mathbf{w} \setminus \mathbf{z} \\
&= \int \sum_{\mathbf{u}'} \mathbb{E}_P[Y_{\mathbf{u}} | \mathbf{u}', \mathbf{w} \setminus \mathbf{z}, \mathbf{z}, \mathbf{t}, \mathbf{c}] P(\mathbf{u}' | \mathbf{w} \setminus \mathbf{z}, \mathbf{z}, \mathbf{t}, \mathbf{c}) P(\mathbf{w} \setminus \mathbf{z} | \mathbf{c}) d\mathbf{w} \setminus \mathbf{z} \\
&\geq \int \mathbb{E}_P[Y_{\mathbf{u}} | \mathbf{u}, \mathbf{w} \setminus \mathbf{z}, \mathbf{z}, \mathbf{t}, \mathbf{c}] P(\mathbf{u} | \mathbf{w} \setminus \mathbf{z}, \mathbf{z}, \mathbf{t}, \mathbf{c}) P(\mathbf{w} \setminus \mathbf{z} | \mathbf{c}) d\mathbf{w} \setminus \mathbf{z} \\
&= \int \mathbb{E}_P[Y | \mathbf{u}, \mathbf{w} \setminus \mathbf{z}, \mathbf{z}, \mathbf{t}, \mathbf{c}] P(\mathbf{u} | \mathbf{w} \setminus \mathbf{z}, \mathbf{z}, \mathbf{t}, \mathbf{c}) P(\mathbf{w} \setminus \mathbf{z} | \mathbf{c}) d\mathbf{w} \setminus \mathbf{z}.
\end{aligned}$$

By combining the inequality above with the decomposition of the policy effect we get

$$\begin{aligned}
&\mathbb{E}_{P_\pi}[Y] \\
&\geq \int \prod_{X \in \mathbf{R}} \pi(x | \mathbf{c}_X) P(\mathbf{c}) \mathbb{E}_P[Y | \mathbf{u}, \mathbf{w} \setminus \mathbf{z}, \mathbf{z}, \mathbf{t}, \mathbf{c}] P(\mathbf{u} | \mathbf{w} \setminus \mathbf{z}, \mathbf{z}, \mathbf{t}, \mathbf{c}) P(\mathbf{w} \setminus \mathbf{z} | \mathbf{c}) d\mathbf{c} d\mathbf{r} d\mathbf{w} \setminus \mathbf{z} \\
&= \int \bar{\pi}_{\mathbf{R}} y P(y | \mathbf{r}, \mathbf{w} \setminus \mathbf{z}, \mathbf{z}, \mathbf{c}) P(\mathbf{u} | \mathbf{w} \setminus \mathbf{z}, \mathbf{z}, \mathbf{t}, \mathbf{c}) P(\mathbf{w} \setminus \mathbf{z} | \mathbf{c}) P(\mathbf{c}) d\mathbf{c} d\mathbf{r} d\mathbf{w} \setminus \mathbf{z} dy \\
&= \int \bar{\pi}_{\mathbf{R}} y P(y | \mathbf{r}, \mathbf{w} \setminus \mathbf{z}, \mathbf{z}, \mathbf{c}) P(\mathbf{u} | \mathbf{w} \setminus \mathbf{z}, \mathbf{z}, \mathbf{t}, \mathbf{c}) \frac{P(\mathbf{t}, \mathbf{z} | \mathbf{w} \setminus \mathbf{z}, \mathbf{c})}{P(\mathbf{t}, \mathbf{z} | \mathbf{w} \setminus \mathbf{z}, \mathbf{c})} P(\mathbf{w} \setminus \mathbf{z} | \mathbf{c}) P(\mathbf{c}) \mathbb{1}_z(\mathbf{z}) d\mathbf{c} d\mathbf{r} d\mathbf{w} \setminus \mathbf{z} dy d\mathbf{z} \\
&= \int \bar{\pi}_{\mathbf{R}} y P(y, \mathbf{r}, \mathbf{w} \setminus \mathbf{z}, \mathbf{z}, \mathbf{c}) \frac{\mathbb{1}_z(\mathbf{z})}{P(\mathbf{t}, \mathbf{z} | \mathbf{w} \setminus \mathbf{z}, \mathbf{c})} d\mathbf{c} d\mathbf{r} d\mathbf{w} \setminus \mathbf{z} dy d\mathbf{z} \\
&= \mathbb{E}_P[Y \gamma]
\end{aligned}$$

where $\gamma := \bar{\pi}_{\mathbf{R}} \mathbb{1}_z(\mathbf{Z}) / P(\mathbf{T}, \mathbf{Z} | \mathbf{W} \setminus \mathbf{Z}, \mathbf{C})$. Note that we have used the definition $\mathbf{R} = \mathbf{T} \cup \mathbf{U}$.

In turn, for the upper bound we have that,

$$\begin{aligned}
\mathbb{E}_P[Y_\pi | \mathbf{x}, \mathbf{c}] &= \int \mathbb{E}_P[Y_{\mathbf{u}} | \mathbf{w} \setminus \mathbf{z}, \mathbf{z}, \mathbf{t}, \mathbf{c}] P(\mathbf{w} \setminus \mathbf{z} | \mathbf{c}) d\mathbf{w} \setminus \mathbf{z} \\
&= \int \sum_{\mathbf{u}'} \mathbb{E}_P[Y_{\mathbf{u}} | \mathbf{u}', \mathbf{w} \setminus \mathbf{z}, \mathbf{z}, \mathbf{t}, \mathbf{c}] P(\mathbf{u}' | \mathbf{w} \setminus \mathbf{z}, \mathbf{z}, \mathbf{t}, \mathbf{c}) P(\mathbf{w} \setminus \mathbf{z} | \mathbf{c}) d\mathbf{w} \setminus \mathbf{z} \\
&= \int \mathbb{E}_P[Y_{\mathbf{u}} | \mathbf{u}, \mathbf{w} \setminus \mathbf{z}, \mathbf{z}, \mathbf{t}, \mathbf{c}] P(\mathbf{u} | \mathbf{w} \setminus \mathbf{z}, \mathbf{z}, \mathbf{t}, \mathbf{c}) P(\mathbf{w} \setminus \mathbf{z} | \mathbf{c}) d\mathbf{w} \setminus \mathbf{z} \\
&+ \int \sum_{\mathbf{u}' \neq \mathbf{u}} \mathbb{E}_P[Y_{\mathbf{u}} | \mathbf{u}', \mathbf{w} \setminus \mathbf{z}, \mathbf{z}, \mathbf{t}, \mathbf{c}] P(\mathbf{u}' | \mathbf{w} \setminus \mathbf{z}, \mathbf{z}, \mathbf{t}, \mathbf{c}) P(\mathbf{w} \setminus \mathbf{z} | \mathbf{c}) d\mathbf{w} \setminus \mathbf{z} \\
&\leq \int \mathbb{E}_P[Y_{\mathbf{u}} | \mathbf{u}, \mathbf{w} \setminus \mathbf{z}, \mathbf{z}, \mathbf{t}, \mathbf{c}] P(\mathbf{u} | \mathbf{w} \setminus \mathbf{z}, \mathbf{z}, \mathbf{t}, \mathbf{c}) P(\mathbf{w} \setminus \mathbf{z} | \mathbf{c}) d\mathbf{w} \setminus \mathbf{z} \\
&+ \int \sum_{\mathbf{u}' \neq \mathbf{u}} P(\mathbf{u}' | \mathbf{w} \setminus \mathbf{z}, \mathbf{z}, \mathbf{t}, \mathbf{c}) P(\mathbf{w} \setminus \mathbf{z} | \mathbf{c}) d\mathbf{w} \setminus \mathbf{z} \\
&= \int \mathbb{E}_P[Y | \mathbf{u}, \mathbf{w} \setminus \mathbf{z}, \mathbf{z}, \mathbf{t}, \mathbf{c}] P(\mathbf{u} | \mathbf{w} \setminus \mathbf{z}, \mathbf{z}, \mathbf{t}, \mathbf{c}) P(\mathbf{w} \setminus \mathbf{z} | \mathbf{c}) d\mathbf{w} \setminus \mathbf{z} \\
&+ \int (1 - P(\mathbf{u} | \mathbf{w} \setminus \mathbf{z}, \mathbf{z}, \mathbf{t}, \mathbf{c})) P(\mathbf{w} \setminus \mathbf{z} | \mathbf{c}) d\mathbf{w} \setminus \mathbf{z} \\
&= 1 + \int \mathbb{E}_P[Y - 1 | \mathbf{u}, \mathbf{w} \setminus \mathbf{z}, \mathbf{z}, \mathbf{t}, \mathbf{c}] P(\mathbf{u} | \mathbf{w} \setminus \mathbf{z}, \mathbf{z}, \mathbf{t}, \mathbf{c}) P(\mathbf{w} \setminus \mathbf{z} | \mathbf{c}) d\mathbf{w} \setminus \mathbf{z}
\end{aligned}$$

By combining the inequality above with the decomposition of the policy effect we get

$$\begin{aligned}
& \mathbb{E}_{P_\pi}[Y] \\
& \leq 1 + \int \prod_{X \in \mathbf{R}} \pi(x | \mathbf{c}_X) P(\mathbf{c}) \mathbb{E}_P[Y - 1 | \mathbf{u}, \mathbf{w} \setminus \mathbf{z}, \mathbf{z}, \mathbf{t}, \mathbf{c}] P(\mathbf{u} | \mathbf{w} \setminus \mathbf{z}, \mathbf{z}, \mathbf{t}, \mathbf{c}) P(\mathbf{w} \setminus \mathbf{z} | \mathbf{c}) d\mathbf{c} d\mathbf{r} d\mathbf{w} \setminus \mathbf{z} \\
& = 1 + \int \bar{\pi}_{\mathbf{R}}(y - 1) P(y | \mathbf{r}, \mathbf{w} \setminus \mathbf{z}, \mathbf{z}, \mathbf{c}) P(\mathbf{u} | \mathbf{w} \setminus \mathbf{z}, \mathbf{z}, \mathbf{t}, \mathbf{c}) P(\mathbf{w} \setminus \mathbf{z} | \mathbf{c}) P(\mathbf{c}) d\mathbf{c} d\mathbf{r} d\mathbf{w} \setminus \mathbf{z} d\mathbf{y} \\
& = 1 + \int \bar{\pi}_{\mathbf{R}}(y - 1) P(y, \mathbf{r}, \mathbf{w} \setminus \mathbf{z}, \mathbf{z}, \mathbf{c}) \frac{\mathbb{1}_{\mathbf{z}}(\mathbf{z})}{P(\mathbf{t}, \mathbf{z} | \mathbf{w} \setminus \mathbf{z}, \mathbf{c})} d\mathbf{c} d\mathbf{r} d\mathbf{w} \setminus \mathbf{z} d\mathbf{y} d\mathbf{z} \\
& = 1 + \mathbb{E}_P[(Y - 1)\gamma]
\end{aligned}$$

where $\gamma := \bar{\pi}_{\mathbf{R}} \mathbb{1}_{\mathbf{z}}(\mathbf{Z}) / P(\mathbf{T}, \mathbf{Z} | \mathbf{W} \setminus \mathbf{Z}, \mathbf{C})$. Note that we have used the definition $\mathbf{R} = \mathbf{T} \cup \mathbf{U}$. \square

Prop. 6 restated. *Let k be the number of variables and m be the number of edges in \mathcal{G} . The run time of Alg. 1 is $\mathcal{O}(k(k^2 + m))$.*

Proof. Let k be the number of variables and m be the number of edges in \mathcal{G} . Operations in Alg. 1, such as computing ancestors (e.g. lines 2,7), could be done in $\mathcal{O}(k^2)$ time, e.g. with a Breadth-First Search algorithm. Checking for d -separation is commonly done with the Bayes-Ball algorithm that can be implemented with a reachability search method, e.g. [43, Sec. 3.2.1], and could be done in $\mathcal{O}(k + m)$ time by [43, Prop. 3.17]. The sets \mathbf{R} and \mathbf{Z} have size at most k , so the for loops in lines 5 and 14 are executed at most k times each. Combining these we get that Alg. 1 requires $\mathcal{O}(k(k^2 + m))$ time to return the bounds. \square

Prop. 7 restated. *Let $\pi := \{\pi_{\mathbf{X}_1}, \pi_{\mathbf{X}_2}\}$ be a policy on $\{\mathbf{X}_1, \mathbf{X}_2\}$ with a conditioning set \mathbf{C} . All partial adjustment sets may be enumerated in time $\mathcal{O}(k(k + m))$ where k are the number of variables and m be the number of edges in \mathcal{G} .*

Proof. This proposition is adapts Proposition 3.20 in [43] to find partial adjustment sets using the ListSep algorithm. ListSep performs backtracking to enumerate all separator sets \mathbf{Z} between \mathbf{X} and \mathbf{Y} such that $\mathbf{I} \subseteq \mathbf{Z} \subseteq \mathbf{R}$, aborting branches that will not find a valid separator. In particular, it calls the TestSep and FindSep algorithms recursively.

The TestSep algorithm takes as input a graph, two sets of nodes to be separated, and a candidate separator set. By setting the graph to be $\mathcal{G}_{\pi_{\mathbf{X}_1}, \mathbf{X}_2}$, the two sets to be separated to be \mathbf{Y} and \mathbf{X}_2 , and the separator set \mathbf{Z} , TestSep will provably return whether $\mathbf{W} = \mathbf{Z} \setminus (\mathbf{C}, \mathbf{X}_1)$ is a partial adjustment set.

The FindSep algorithm uses the observation that if there exists a separator \mathbf{Z} between \mathbf{X} and \mathbf{Y} such that $\mathbf{I} \subseteq \mathbf{Z} \subseteq \mathbf{R}$ then $\mathbf{Z} := An(\mathbf{X} \cup \mathbf{Y} \cup \mathbf{I}) \cap \mathbf{R}$ is a separator. For each possible partition $(\mathbf{X}_1, \mathbf{X}_2)$ of \mathbf{X} , we can therefore find a separator by testing (using TestSep) whether $An(\mathbf{X} \cup \mathbf{Y} \cup \mathbf{C}) \cap (\mathbf{V} \setminus (\mathbf{X} \cup \mathbf{Y}))$ in $\mathcal{G}_{\pi_{\mathbf{X}_1}, \mathbf{X}_2}$ is a separator. If it is then we have found a partial adjustment set.

To find all partial adjustment sets, we then proceed as follows. For each possible partition $(\mathbf{X}_1, \mathbf{X}_2)$ of \mathbf{X} , apply the ListSep algorithm with graph $\mathcal{G}_{\pi_{\mathbf{X}_1}, \mathbf{X}_2}$, sets to be separated \mathbf{Y} and \mathbf{X}_2 , and possible separator sets \mathbf{Z} constrained as $\mathbf{C}, \mathbf{X}_1 \subseteq \mathbf{Z} \subseteq \mathbf{V} \setminus (\mathbf{X} \cup \mathbf{Y})$. Then, for every separator set \mathbf{Z} found, output partial adjustment sets $\mathbf{W} = \mathbf{Z} \setminus (\mathbf{C}, \mathbf{X}_1)$.

This procedure finds all possible partial adjustment sets in time $\mathcal{O}(k(k + m))$ where k are the number of variables and m be the number of edges in \mathcal{G} by Proposition 3.20 [43]. \square

B.2 Proofs of statements in Sec. 4

We make use of two auxiliary lemmas from the literature.

Lemma 1 (Continuous Mapping Theorem, [26]). *Let $\{X_n\}_{n \in \mathbb{N}}$, X be random elements defined on a metric space S . Consider a continuous function $g : S \rightarrow S'$ (where S' is another metric space). Then,*

$$X_n \rightarrow_p X \Rightarrow g(X_n) \rightarrow_p g(X).$$

Lemma 2 (Lemma 2, [24]). Let $f_\eta := f(\mathbf{V}; \eta)$ denote a finite and continuous functional and η denote its nuisances. For n independent samples of P , $D := \{\mathbf{v}^{(i)}\}_{i=1, \dots, n} \sim P$, let $\hat{T} = \mathbb{E}_D[f_{\hat{\eta}}]$ and $T := \mathbb{E}_P[f_\eta]$ for some η . Let $\mathbb{E}_{D-P}[f_\eta] := \mathbb{E}_D[f_\eta] - \mathbb{E}_P[f_\eta]$. Then, the following decomposition holds:

$$\mathbb{E}_D[f_{\hat{\eta}}] - \mathbb{E}_P[f_\eta] = \underbrace{\mathbb{E}_{D-P}[f_\eta]}_{=A} + \underbrace{\mathbb{E}_{D-P}[f_{\hat{\eta}} - f_\eta]}_{=B} + \mathbb{E}_P[f_{\hat{\eta}} - f_\eta]. \quad (19)$$

Suppose further that samples used for estimating η are independent and separate; and the nuisances are consistent. Then, $R = A + B$ is a random variable that converges to zero at a rate $\mathcal{O}_P(1/\sqrt{n})$.

Prop. 8 restated. Suppose the nuisance estimates $(\hat{\boldsymbol{\mu}}, \hat{\boldsymbol{\gamma}})$ are L_2 -consistent and bounded. Then, the error of the DML estimator $\hat{T}^{\text{DML}} \in \{\hat{T}^{\text{DML}, \ell}, \hat{T}^{\text{DML}, u}\}$ in Def. 5 is given as follows

$$\hat{T}^{\text{DML}} - T^{\text{DML}} = \frac{1}{K} \sum_{k=1}^K R_k + \mathcal{O}_P\left(\|\hat{\boldsymbol{\gamma}}_{2,k} - \boldsymbol{\gamma}_2\| \|\hat{\boldsymbol{\mu}}_{2,k} - \boldsymbol{\mu}_2\|\right) + \mathcal{O}_P\left(\|\hat{\boldsymbol{\gamma}}_{1,k} - \boldsymbol{\gamma}_1\| \|\hat{\boldsymbol{\mu}}_{1,k} - \hat{\boldsymbol{\mu}}_{1,k}\|\right)$$

where R_k is a random variable that converges to zero at a rate $\mathcal{O}_P(1/\sqrt{n})$.

Proof. In this proof we consider the estimation of the lower bound without loss of generality. Recall the definition of nuisance functions and estimators:

$$\boldsymbol{\gamma}_2 := \bar{\pi}_U \boldsymbol{\gamma}_1, \quad \boldsymbol{\gamma}_1 := \bar{\pi}_T \mathbb{1}_z(\mathbf{Z})/P(\mathbf{T}, \mathbf{Z} \mid \mathbf{W} \setminus \mathbf{Z}, \mathbf{C}),$$

with $T^{\text{PW}, \ell} := \mathbb{E}_P[\boldsymbol{\gamma}_2 Y] = \boldsymbol{\psi}_z^\ell$. And $\boldsymbol{\mu} = (\mu_0, \mu_1, \tilde{\mu}_2, \mu_2)$ defined by

$$\begin{aligned} \mu_2 &:= \mu_2(\mathbf{R}, \mathbf{C}, \mathbf{W}, \mathbf{Z}) = \mathbb{E}_P[Y \mid \mathbf{R}, \mathbf{C}, \mathbf{W}, \mathbf{Z}], \\ \tilde{\mu}_1 &:= \tilde{\mu}_2(\mathbf{R}, \mathbf{C}, \mathbf{W}, \mathbf{Z}) = \bar{\pi}_U(\mathbf{U} \mid \mathbf{C}) \mu_2(\mathbf{R}, \mathbf{C}, \mathbf{W}, \mathbf{Z}), \\ \mu_1 &:= \mu_1(\mathbf{T}, \mathbf{C}, \mathbf{W}, \mathbf{Z}) = \mathbb{E}_P[\tilde{\mu}_2(\mathbf{R}, \mathbf{C}, \mathbf{W}, \mathbf{Z}) \mid \mathbf{T}, \mathbf{C}, \mathbf{W}, \mathbf{Z}], \\ \mu_0 &:= \mu_0(\mathbf{C}, \mathbf{W}, \mathbf{Z}) = \sum_{\mathbf{t}} \mu_1(\mathbf{t}, \mathbf{C}, \mathbf{W}, \mathbf{Z}) \bar{\pi}_T(\mathbf{t} \mid \mathbf{C}). \end{aligned}$$

We first verify that the population level value of the regression estimator coincides with the lower bound $\boldsymbol{\psi}_z^\ell$. This can be seen with the following derivation,

$$\begin{aligned} T^{\text{REG}, \ell} &:= \mathbb{E}_P[\mu_0(\mathbf{W} \setminus \mathbf{Z}, \mathbf{C}, \mathbf{z})] \\ &= \sum_{\mathbf{w} \setminus \mathbf{z}, \mathbf{c}} \mu_0(\mathbf{w}, \mathbf{c}, \mathbf{z}) P(\mathbf{w} \setminus \mathbf{z}, \mathbf{c}) \\ &= \sum_{\mathbf{w} \setminus \mathbf{z}, \mathbf{c}} \left(\sum_{\mathbf{t}} \mu_1(\mathbf{t}, \mathbf{w}, \mathbf{c}, \mathbf{z}) \bar{\pi}_T \right) P(\mathbf{w} \setminus \mathbf{z}, \mathbf{c}) \\ &= \sum_{\mathbf{t}, \mathbf{w} \setminus \mathbf{z}, \mathbf{c}} \left(\sum_{\mathbf{u}} \mu_2(\mathbf{t}, \mathbf{u}, \mathbf{c}, \mathbf{w} \setminus \mathbf{z}, \mathbf{z}) \bar{\pi}_U P(\mathbf{u} \mid \mathbf{t}, \mathbf{z}, \mathbf{w} \setminus \mathbf{z}, \mathbf{c}) \right) \frac{\bar{\pi}_T}{P(\mathbf{t}, \mathbf{z} \mid \mathbf{w} \setminus \mathbf{z}, \mathbf{c})} P(\mathbf{t}, \mathbf{z}, \mathbf{w} \setminus \mathbf{z}, \mathbf{c}) \\ &= \sum_{\mathbf{t}, \mathbf{u}, \mathbf{w} \setminus \mathbf{z}, \mathbf{c}} \left(\sum_{\mathbf{y}} y P(\mathbf{y} \mid \mathbf{t}, \mathbf{u}, \mathbf{z}, \mathbf{w} \setminus \mathbf{z}, \mathbf{c}) \right) \frac{\bar{\pi}_R}{P(\mathbf{t}, \mathbf{z} \mid \mathbf{w} \setminus \mathbf{z}, \mathbf{c})} P(\mathbf{t}, \mathbf{u}, \mathbf{z}, \mathbf{w}, \mathbf{c}) \\ &= \sum_{\mathbf{y}, \mathbf{t}, \mathbf{u}, \mathbf{z}', \mathbf{w} \setminus \mathbf{z}', \mathbf{c}} y \frac{\bar{\pi}_R \mathbb{1}_z(\mathbf{z}')}{P(\mathbf{t}, \mathbf{z}' \mid \mathbf{w} \setminus \mathbf{z}', \mathbf{c})} P(\mathbf{y}, \mathbf{t}, \mathbf{u}, \mathbf{z}', \mathbf{w} \setminus \mathbf{z}', \mathbf{c}) \\ &= \mathbb{E}_P[Y \bar{\pi}_R \mathbb{1}_z(\mathbf{Z})/P(\mathbf{T}, \mathbf{Z} \mid \mathbf{W} \setminus \mathbf{Z}, \mathbf{C})] \\ &= \boldsymbol{\psi}_z^\ell \end{aligned}$$

The DML estimator is similarly unbiased as,

$$\begin{aligned} T^{\text{DML}, \ell} &:= \mathbb{E}_P[\boldsymbol{\gamma}_2 \{Y - \mu_2\}] + \mathbb{E}_P[\boldsymbol{\gamma}_1 \{\tilde{\mu}_1 - \mu_1\}] + \mathbb{E}_P[\mu_0] \\ &= \mathbb{E}_P[\boldsymbol{\gamma}_2 \{\mathbb{E}_P[Y \mid \mathbf{R}, \mathbf{W} \setminus \mathbf{Z}, \mathbf{C}, \mathbf{Z}] - \mu_2\}] + \mathbb{E}_P[\boldsymbol{\gamma}_1 \{\tilde{\mu}_1 - \mu_1\}] + \mathbb{E}_P[\mu_0] \\ &= \mathbb{E}_P[\boldsymbol{\gamma}_2 \{\mu_2 - \mu_2\}] + \mathbb{E}_P[\boldsymbol{\gamma}_1 \{\mathbb{E}_P[\tilde{\mu}_1 \mid \mathbf{T}, \mathbf{C}, \mathbf{W} \setminus \mathbf{Z}, \mathbf{Z}] - \mu_1\}] + \mathbb{E}_P[\mu_0] \\ &= \mathbb{E}_P[\boldsymbol{\gamma}_1 \{\mu_1 - \mu_1\}] + \mathbb{E}_P[\mu_0] \\ &= \mathbb{E}_P[\mu_0] \\ &= \boldsymbol{\psi}_z^\ell. \end{aligned}$$

Consider the estimated value of $T^{\text{DML},\ell}$ following the procedure in Def. 5,

$$\hat{T}^{\text{DML},\ell} := \frac{1}{K} \sum_{k=1}^K \hat{T}_k^{\text{DML},\ell}, \quad \hat{T}_k^{\text{DML},\ell} := \mathbb{E}_{\mathcal{D}^{(k)}}[\hat{\gamma}_2\{Y - \hat{\mu}_2\}] + \mathbb{E}_{\mathcal{D}^{(k)}}[\hat{\gamma}_1\{\hat{\mu}_1 - \hat{\mu}_1\}] + \mathbb{E}_{\mathcal{D}^{(k)}}[\hat{\mu}_0],$$

We could then write,

$$\hat{T}_k^{\text{DML},\ell} - T^{\text{DML},\ell} = A + B + C$$

where,

$$\begin{aligned} A &= \mathbb{E}_{\mathcal{D}^{(k)}-P}[\gamma_2\{Y - \mu_2\} + \gamma_1\{\tilde{\mu}_1 - \mu_1\} + \mu_0] \\ B &= \mathbb{E}_{\mathcal{D}^{(k)}-P}\left[\left(\gamma_2\{Y - \mu_2\}\right) + \gamma_1\{\tilde{\mu}_1 - \mu_1\} + \mu_0\right] - \left(\hat{\gamma}_2\{Y - \hat{\mu}_2\} + \hat{\gamma}_1\{\hat{\mu}_1 - \hat{\mu}_1\} + \hat{\mu}_0\right) \\ C &= \mathbb{E}_P\left[\left(\gamma\{Y - \mu_2\}\right) + \gamma\{\tilde{\mu}_1 - \mu_1\} + \mu_0\right] - \left(\hat{\gamma}_2\{Y - \hat{\mu}_2\} + \hat{\gamma}_1\{\hat{\mu}_1 - \hat{\mu}_1\} + \hat{\mu}_0\right) \end{aligned}$$

By the construction in Def. 5, the samples used for estimating $(\gamma_1, \gamma_2, \mu_0, \mu_1, \tilde{\mu}_1, \mu_2)$ and evaluating the outer expectation are independent and separate. Under the assumption that nuisance parameters $(\hat{\gamma}_1, \hat{\gamma}_2, \hat{\mu}_0, \hat{\mu}_1, \hat{\mu}_1, \hat{\mu}_2)$ are consistent, $R = A + B$ converges to zero at a rate $\mathcal{O}_P(1/\sqrt{|\mathcal{D}^{(k)}|})$ by Lem. 2.

Before manipulating C and deriving its large sample behaviour, consider the following intermediate results:

$$\begin{aligned} \mathbb{E}_P[\mu_0] &= \sum_{\mathbf{c}, \mathbf{w}} \mu_0(\mathbf{c}, \mathbf{w}, \mathbf{z})P(\mathbf{c}, \mathbf{w}) \\ &= \sum_{\mathbf{c}, \mathbf{w}} \sum_{\mathbf{t}} \bar{\pi}_T \mu_1(\mathbf{t}, \mathbf{c}, \mathbf{w}, \mathbf{z})P(\mathbf{c}, \mathbf{w}) \\ &= \sum_{\mathbf{c}, \mathbf{w}, \mathbf{z}', \mathbf{t}} \bar{\pi}_T \mathbb{1}_{\mathbf{z}}(\mathbf{z}') \mu_1(\mathbf{t}, \mathbf{c}, \mathbf{w}, \mathbf{z}')P(\mathbf{c}, \mathbf{w}) \\ &= \sum_{\mathbf{c}, \mathbf{w}, \mathbf{z}', \mathbf{t}} \frac{\bar{\pi}_T \mathbb{1}_{\mathbf{z}}(\mathbf{z}')}{P(\mathbf{t}, \mathbf{z}' | \mathbf{c}, \mathbf{w})} \mu_1(\mathbf{t}, \mathbf{c}, \mathbf{w}, \mathbf{z}')P(\mathbf{t}, \mathbf{c}, \mathbf{w}, \mathbf{z}') \\ &= \mathbb{E}_P[\gamma_1 \mu_1], \end{aligned}$$

and similarly $\mathbb{E}_P[\hat{\mu}_0] = \mathbb{E}_P[\gamma_1 \hat{\mu}_1]$. Note further that,

$$\begin{aligned} \mathbb{E}_P[\gamma_2 \mu_2] &= \sum_{\mathbf{c}, \mathbf{w}, \mathbf{t}, \mathbf{u}, \mathbf{z}'} \frac{\bar{\pi}_T \bar{\pi}_U \mathbb{1}_{\mathbf{z}}(\mathbf{z}')}{P(\mathbf{t}, \mathbf{z}' | \mathbf{c}, \mathbf{w})} \mu_2(\mathbf{c}, \mathbf{w}, \mathbf{t}, \mathbf{u}, \mathbf{z}')P(\mathbf{c}, \mathbf{w}, \mathbf{t}, \mathbf{u}, \mathbf{z}') \\ &= \sum_{\mathbf{c}, \mathbf{w}, \mathbf{t}, \mathbf{u}, \mathbf{z}'} \gamma_1 \bar{\pi}_U \mu_2(\mathbf{c}, \mathbf{w}, \mathbf{t}, \mathbf{u}, \mathbf{z}')P(\mathbf{c}, \mathbf{w}, \mathbf{t}, \mathbf{u}, \mathbf{z}') \\ &= \sum_{\mathbf{c}, \mathbf{w}, \mathbf{t}, \mathbf{u}, \mathbf{z}'} \gamma_1 \tilde{\mu}_1(\mathbf{c}, \mathbf{w}, \mathbf{t}, \mathbf{u}, \mathbf{z}')P(\mathbf{c}, \mathbf{w}, \mathbf{t}, \mathbf{z}') \\ &= \mathbb{E}_P[\gamma_1 \tilde{\mu}_1], \end{aligned}$$

and similarly $\mathbb{E}_P[\gamma_2 \hat{\mu}_2] = \mathbb{E}_P[\gamma_1 \hat{\mu}_1]$.

Now consider $C = C_1 + C_2 + C_3$ where,

$$\begin{aligned} C_1 &:= \mathbb{E}_P[\hat{\gamma}_2\{Y - \hat{\mu}_2\} - \gamma_2\{Y - \mu_2\}] = \mathbb{E}_P[\hat{\gamma}_2\{\mu_2 - \hat{\mu}_2\}] \\ C_2 &:= \mathbb{E}_P[\hat{\gamma}_1\{\hat{\mu}_1 - \hat{\mu}_1\} - \gamma_1\{\tilde{\mu}_1 - \mu_1\}] \\ C_3 &:= \mathbb{E}_P[\hat{\mu}_0 - \mu_0] = \mathbb{E}_P[\gamma_1 \hat{\mu}_1 - \gamma_1 \mu_1] \end{aligned}$$

where the last equality follows from the result above. Therefore

$$\begin{aligned}
C &= \mathbb{E}_P \left[\hat{\gamma}_2 \{\mu_2 - \hat{\mu}_2\} + \hat{\gamma}_1 \{\hat{\mu}_1 - \hat{\mu}_1\} - \gamma_1 \{\tilde{\mu}_1 - \mu_1\} + \gamma_1 \{\hat{\mu}_1 - \mu_1\} \right] \\
&= \mathbb{E}_P \left[\{\hat{\gamma}_2 - \gamma_2\} \{\mu_2 - \hat{\mu}_2\} + \gamma_2 \{\mu_2 - \hat{\mu}_2\} + \hat{\gamma}_1 \{\hat{\mu}_1 - \hat{\mu}_1\} - \gamma_1 \{\tilde{\mu}_1 - \mu_1\} + \gamma_1 \{\hat{\mu}_1 - \mu_1\} \right] \\
&= \mathbb{E}_P \left[\{\hat{\gamma}_2 - \gamma_2\} \{\mu_2 - \hat{\mu}_2\} \right] + \mathbb{E}_P \left[\gamma_2 \{\mu_2 - \hat{\mu}_2\} + \hat{\gamma}_1 \{\hat{\mu}_1 - \hat{\mu}_1\} - \gamma_1 \{\tilde{\mu}_1 - \mu_1\} \right] \\
&\stackrel{(1)}{=} \mathbb{E}_P \left[\{\hat{\gamma}_2 - \gamma_2\} \{\mu_2 - \hat{\mu}_2\} \right] + \mathbb{E}_P \left[\gamma_1 \{\tilde{\mu}_1 - \hat{\mu}_1\} + \hat{\gamma}_1 \{\hat{\mu}_1 - \hat{\mu}_1\} - \gamma_1 \{\tilde{\mu}_1 - \hat{\mu}_1\} \right] \\
&= \mathbb{E}_P \left[\{\hat{\gamma}_2 - \gamma_2\} \{\mu_2 - \hat{\mu}_2\} \right] + \mathbb{E}_P \left[-\gamma_1 \{\hat{\mu}_1 - \hat{\mu}_1\} + \hat{\gamma}_1 \{\hat{\mu}_1 - \hat{\mu}_1\} \right] \\
&= O_P \left(\|\hat{\gamma}_2 - \gamma_2\| \|\mu_2 - \hat{\mu}_2\| \right) + O_P \left(\|\hat{\gamma}_1 - \gamma_1\| \|\hat{\mu}_1 - \hat{\mu}_1\| \right)
\end{aligned}$$

(1) holds by the equalities $\mathbb{E}_P[\gamma_2 \mu_2] = \mathbb{E}_P[\gamma_1 \tilde{\mu}_1]$, $\mathbb{E}_P[\gamma_2 \hat{\mu}_2] = \mathbb{E}_P[\gamma_1 \hat{\mu}_1]$.

Finally, this implies that

$$\hat{T}^{\text{DML},\ell} - T^{\text{DML},\ell} = R + O_P \left(\|\hat{\gamma}_2 - \gamma_2\| \|\mu_2 - \hat{\mu}_2\| \right) + O_P \left(\|\hat{\gamma}_1 - \gamma_1\| \|\hat{\mu}_1 - \hat{\mu}_1\| \right) \quad (20)$$

where R is a random variable that converges to zero at a rate $O_P(1/\sqrt{n})$.

For the upper bound, consider the following definition of nuisance functions and estimators:

$$\gamma_2 := \bar{\pi}_U \gamma_1, \quad \gamma_1 := \bar{\pi}_T \mathbb{1}_z(\mathbf{Z})/P(\mathbf{T}, \mathbf{Z} \mid \mathbf{W} \setminus \mathbf{Z}, \mathbf{C}),$$

with $T^{\text{PW},\ell} := \mathbb{E}_P[\gamma_2(Y-1)] + 1 = \psi_z^\ell$. And $\boldsymbol{\mu} = (\mu_0, \mu_1, \tilde{\mu}_2, \mu_2)$ defined by

$$\begin{aligned}
\mu_2 &:= \mu_2(\mathbf{R}, \mathbf{C}, \mathbf{W}, \mathbf{Z}) = \mathbb{E}_P[Y-1 \mid \mathbf{R}, \mathbf{C}, \mathbf{W}, \mathbf{Z}], \\
\tilde{\mu}_1 &:= \tilde{\mu}_2(\mathbf{R}, \mathbf{C}, \mathbf{W}, \mathbf{Z}) = \bar{\pi}_U(\mathbf{U} \mid \mathbf{C}) \mu_2(\mathbf{R}, \mathbf{C}, \mathbf{W}, \mathbf{Z}), \\
\mu_1 &:= \mu_1(\mathbf{T}, \mathbf{C}, \mathbf{W}, \mathbf{Z}) = \mathbb{E}_P[\tilde{\mu}_2(\mathbf{R}, \mathbf{C}, \mathbf{W}, \mathbf{Z}) \mid \mathbf{T}, \mathbf{C}, \mathbf{W}, \mathbf{Z}], \\
\mu_0 &:= \mu_0(\mathbf{C}, \mathbf{W}, \mathbf{Z}) = \sum_{\mathbf{t}} \mu_1(\mathbf{t}, \mathbf{C}, \mathbf{W}, \mathbf{Z}) \bar{\pi}_T(\mathbf{t} \mid \mathbf{C}).
\end{aligned}$$

We first verify that the population level value of the regression estimator coincides with the lower bound ψ_z^ℓ . This can be seen with the following derivation,

$$\begin{aligned}
T^{\text{REG},u} &:= \mathbb{E}_P[\mu_0(\mathbf{W} \setminus \mathbf{Z}, \mathbf{C}, \mathbf{z})] + 1 \\
&= 1 + \sum_{\mathbf{w} \setminus \mathbf{z}, \mathbf{c}} \mu_0(\mathbf{w}, \mathbf{c}, \mathbf{z}) P(\mathbf{w} \setminus \mathbf{z}, \mathbf{c}) \\
&= 1 + \sum_{\mathbf{w} \setminus \mathbf{z}, \mathbf{c}} \left(\sum_{\mathbf{t}} \mu_1(\mathbf{t}, \mathbf{w}, \mathbf{c}, \mathbf{z}) \bar{\pi}_T \right) P(\mathbf{w} \setminus \mathbf{z}, \mathbf{c}) \\
&= 1 + \sum_{\mathbf{t}, \mathbf{w} \setminus \mathbf{z}, \mathbf{c}} \left(\sum_{\mathbf{u}} \mu_2(\mathbf{t}, \mathbf{u}, \mathbf{c}, \mathbf{w} \setminus \mathbf{z}, \mathbf{z}) \bar{\pi}_U P(\mathbf{u} \mid \mathbf{t}, \mathbf{z}, \mathbf{w} \setminus \mathbf{z}, \mathbf{c}) \right) \frac{\bar{\pi}_T}{P(\mathbf{t}, \mathbf{z} \mid \mathbf{w} \setminus \mathbf{z}, \mathbf{c})} P(\mathbf{t}, \mathbf{z}, \mathbf{w} \setminus \mathbf{z}, \mathbf{c}) \\
&= 1 + \sum_{\mathbf{t}, \mathbf{u}, \mathbf{w} \setminus \mathbf{z}, \mathbf{c}} \left(\sum_{\mathbf{y}} (y-1) P(\mathbf{y} \mid \mathbf{t}, \mathbf{u}, \mathbf{z}, \mathbf{w} \setminus \mathbf{z}, \mathbf{c}) \right) \frac{\bar{\pi}_R}{P(\mathbf{t}, \mathbf{z} \mid \mathbf{w} \setminus \mathbf{z}, \mathbf{c})} P(\mathbf{t}, \mathbf{u}, \mathbf{z}, \mathbf{w}, \mathbf{c}) \\
&= 1 + \sum_{\mathbf{y}, \mathbf{t}, \mathbf{u}, \mathbf{z}', \mathbf{w} \setminus \mathbf{z}', \mathbf{c}} (y-1) \frac{\bar{\pi}_R \mathbb{1}_z(\mathbf{z}')}{P(\mathbf{t}, \mathbf{z}' \mid \mathbf{w} \setminus \mathbf{z}', \mathbf{c})} P(\mathbf{y}, \mathbf{t}, \mathbf{u}, \mathbf{z}', \mathbf{w} \setminus \mathbf{z}', \mathbf{c}) \\
&= 1 + \mathbb{E}_P[(Y-1) \bar{\pi}_R \mathbb{1}_z(\mathbf{Z})/P(\mathbf{T}, \mathbf{Z} \mid \mathbf{W} \setminus \mathbf{Z}, \mathbf{C})] \\
&= \psi_z^u
\end{aligned}$$

The DML estimator is similarly unbiased as,

$$\begin{aligned}
T^{\text{DML},u} &:= \mathbb{E}_P[\gamma_2\{(Y-1) - \mu_2\}] + \mathbb{E}_P[\gamma_1\{\tilde{\mu}_1 - \mu_1\}] + \mathbb{E}_P[\mu_0] + 1 \\
&= \mathbb{E}_P[\gamma_2\{\mathbb{E}_P[Y-1 \mid \mathbf{R}, \mathbf{W} \setminus \mathbf{Z}, \mathbf{C}, \mathbf{Z}] - \mu_2\}] + \mathbb{E}_P[\gamma_1\{\tilde{\mu}_1 - \mu_1\}] + \mathbb{E}_P[\mu_0] + 1 \\
&= \mathbb{E}_P[\gamma_2\{\mu_2 - \mu_2\}] + \mathbb{E}_P[\gamma_1\{\mathbb{E}_P[\tilde{\mu}_1 \mid \mathbf{T}, \mathbf{C}, \mathbf{W} \setminus \mathbf{Z}, \mathbf{Z}] - \mu_1\}] + \mathbb{E}_P[\mu_0] + 1 \\
&= \mathbb{E}_P[\gamma_1\{\mu_1 - \mu_1\}] + \mathbb{E}_P[\mu_0] + 1 \\
&= \mathbb{E}_P[\mu_0] + 1 \\
&= \psi_z^u.
\end{aligned}$$

The arguments of the proof are now analogous to that of the lower bound. We can conclude therefore that,

$$\hat{T}^{\text{DML},u} - T^{\text{DML},u} = R + O_P\left(\|\hat{\gamma}_2 - \gamma_2\| \|\mu_2 - \hat{\mu}_2\|\right) + O_P\left(\|\hat{\gamma}_1 - \gamma_1\| \|\hat{\mu}_1 - \mu_1\|\right). \quad (21)$$

□

Prop. 9 restated. *Suppose either $\hat{\gamma}_1 = \gamma_1$ or $\hat{\mu}_2 = \mu_2$ and that either $\hat{\gamma}_1 = \gamma_1$ or $\hat{\mu}_1 = \mu_1$. Then, $\hat{T}^{\text{DML}} \in \{\hat{T}^{\text{DML},\ell}, \hat{T}^{\text{DML},u}\}$ is an unbiased estimator of the corresponding bound defined in Alg. 1.*

Proof. Consider the estimated value of $T^{\text{DML},\ell}$ following the procedure in Def. 5,

$$\hat{T}^{\text{DML},\ell} := \frac{1}{K} \sum_{k=1}^K \hat{T}_k^{\text{DML},\ell}, \quad \hat{T}_k^{\text{DML},\ell} := \mathbb{E}_{\mathcal{D}^{(k)}}[\hat{\gamma}_2\{Y - \hat{\mu}_2\}] + \mathbb{E}_{\mathcal{D}^{(k)}}[\hat{\gamma}_1\{\hat{\mu}_1 - \hat{\mu}_1\}] + \mathbb{E}_{\mathcal{D}^{(k)}}[\hat{\mu}_0],$$

Under the assumption that,

$$\mathbb{E}_P[\hat{T}^{\text{DML},\ell}] = \mathbb{E}_P[\hat{\gamma}_2\{Y - \hat{\mu}_2\}] + \mathbb{E}_P[\hat{\gamma}_1\{\hat{\mu}_1 - \hat{\mu}_1\}] + \mathbb{E}_P[\hat{\mu}_0]$$

The bias of the estimator is given by,

$$\begin{aligned} & \mathbb{E}_P[\hat{T}^{\text{DML},\ell}] - T^{\text{DML},\ell} \\ &= \mathbb{E}_P\left[\left(\gamma\{Y - \mu_2\} + \gamma\{\tilde{\mu}_1 - \mu_1\} + \mu_0\right) - \left(\hat{\gamma}_2\{Y - \hat{\mu}_2\} + \hat{\gamma}_1\{\hat{\mu}_1 - \hat{\mu}_1\} + \hat{\mu}_0\right)\right] \\ &= \mathbb{E}_P\left[\{\hat{\gamma}_2 - \gamma_2\}\{\mu_2 - \hat{\mu}_2\}\right] + \mathbb{E}_P\left[\{\hat{\gamma}_1 - \gamma_1\}\{\hat{\mu}_1 - \hat{\mu}_1\}\right] \end{aligned}$$

The last equality follows from the derivation of (C) in the proof of Prop. 8.

If either $\hat{\gamma}_1 = \gamma_1$ or $\hat{\mu}_2 = \mu_2$ and that either $\hat{\gamma}_1 = \gamma_1$ or $\hat{\mu}_1 = \mu_1$ then this expression equals 0 and we have that,

$$\mathbb{E}_P[\hat{T}^{\text{DML},\ell}] - T^{\text{DML},\ell} = 0.$$

An analogous result holds for the upper bound. □

C Details on experiments

This section provides details of the data generating mechanisms used for synthetic experiments, and implementations.

As described in the experimental section, for estimating nuisances (γ, μ) we used Gradient Boosting models for classification and regression where appropriate. We implemented the models using Python using the commands `GradientBoostingClassifier()` and `GradientBoostingRegressor()` using default hyperparameters. We systematically bound probability values in the interval $[0.01, 0.99]$ to avoid propagating potentially large approximating errors.

C.1 Simulations

We make use of the following 4 data generating systems, summarized by the graphs in Fig. 4. In this section, we use the notation $\mathbb{1}\{X\}$ that equals 1 if the statement $\{X\}$ is true, and equal to 0 otherwise.

Example 1. The data generating mechanism for the first scenario is given by:

$$\begin{aligned} U_Z, U_{XY}, U_C &\sim \mathcal{N}(0, 1), \\ Z &:= f_Z(U_Z), \\ C &:= f_C(U_C), \\ X &:= f_X(Z, C, U_{XY}), \\ Y &:= f_Y(X, C, U_{XY}), \end{aligned}$$

where,

$$\begin{aligned} f_Z(U_Z) &:= -1 \cdot \mathbb{1}\{U_Z \leq -0.5\} + 0 \cdot \mathbb{1}\{-0.5 < U_Z < 0.5\} + 1 \cdot \mathbb{1}\{U_Z \geq 0.5\}, \\ f_C(U_C) &:= U_C, \\ f_X(C, U_{XY}) &:= \mathbb{1}\{U_{XY} > 1 + \exp\{0.75C + 0.5Z + 0.5\}\}, \\ f_Y(X, C, U_{XY}) &:= (2X - 1) - 2 \cdot (2U_{XY} - 1) \end{aligned}$$

The outcome Y is then re-scaled to lie in the $[0, 1]$ interval. The target for estimation can be derived with Alg. 1 and equals, for the lower and upper bound respectively,

$$\max_{z \in \{-1, 0, 1\}} \mathbb{E}_P[\bar{\pi} \mathbb{1}_z(Z)Y/P(Z | C)], \quad \min_{z \in \{-1, 0, 1\}} \mathbb{E}_P[\bar{\pi} \mathbb{1}_z(Z)(Y - 1)/P(Z | C)] + 1.$$

Example 2. The data generating mechanism for the second scenario is given by:

$$\begin{aligned} U_{X_1Y}, U_{X_2W}, U_{X_2}, U_{X_1}, U_C, U_W &\sim \mathcal{N}(0, 1), \\ C &:= f_C(U_C), \\ W &:= f_W(U_{X_2W}), \\ X_1 &:= f_{X_1}(C, U_{X_1Y}, U_{X_1}), \\ X_2 &:= f_{X_2}(U_{X_2W}, U_{X_2}), \\ Y &:= f_Y(X_1, X_2, C, W, U_{X_1Y}), \end{aligned}$$

where,

$$\begin{aligned} f_C(U_C) &:= U_C, \\ f_W(U_W) &:= U_{X_2W}, \\ f_{X_1}(C, U_{X_1Y}, U_{X_1}) &:= \mathbb{1}\{U_{X_1Y} \cdot (1 + \exp\{0.75C + 0.5\}) > U_{X_1}\}, \\ f_{X_2}(U_{X_2W}, U_{X_2}) &:= \mathbb{1}\{U_{X_2W} + U_{X_2} > 0\}, \\ f_Y(X_1, X_2, C, W, U_{X_1Y}) &:= (2X_1 - 1) \cdot (C + 1) + 2 \sin(2C \cdot (2X_1 - 1)) \\ &\quad - 2(2U_{X_1Y} - 1)(1 + 0.5C) + X_2 + W. \end{aligned}$$

The outcome Y is then re-scaled to lie in the $[0, 1]$ interval. The target for estimation can be derived with Alg. 1 and equals, for the lower and upper bound respectively,

$$\mathbb{E}_P[\bar{\pi}Y/P(X_2 | W, C)], \quad \mathbb{E}_P[\bar{\pi}(Y - 1)/P(X_2 | W, C)] + 1.$$

Example 3. The data generating mechanism for the third scenario is given by:

$$\begin{aligned}
U_{X_1Y}, U_{X_2W}, U_{X_2}, U_{X_1}, U_C &\sim \mathcal{N}(0, 1), \\
\mathbf{W} &:= \{W_1, \dots, W_{100}\} \text{ s.t. } W_i \sim \mathcal{N}(0, 1), i = 1, \dots, 100, \\
C &:= f_C(U_C), \\
X_1 &:= f_{X_1}(C, U_{X_1Y}, U_{X_1}), \\
X_2 &:= f_{X_2}(\mathbf{W}, U_{X_2}), \\
Y &:= f_Y(X_1, X_2, C, \mathbf{W}, U_{X_1Y}),
\end{aligned}$$

where,

$$\begin{aligned}
f_C(U_C) &:= U_C, \\
f_{X_1}(C, U_{X_1Y}, U_{X_1}) &:= \mathbb{1}\{U_{X_1Y} \cdot (1 + \exp\{0.75C + 0.5\}) > U_{X_1}\}, \\
f_{X_2}(\mathbf{W}, U_{X_2}) &:= \mathbb{1}\{W_1 + W_2 + U_{X_2} > 0\}, \\
f_Y(X_1, X_2, C, \mathbf{W}, U_{X_1Y}) &:= (2X_1 - 1) \cdot (C + 1) + 2 \sin(2C \cdot (2X_1 - 1)) \\
&\quad - 2(2U_{X_1Y} - 1)(1 + 0.5C) + X_2 + \sum_{i=1, \dots, 100} W_i/100.
\end{aligned}$$

The outcome Y is then re-scaled to lie in the $[0, 1]$ interval. The target for estimation can be derived with Alg. 1 and equals, for the lower and upper bound respectively,

$$\mathbb{E}_P[\bar{\pi}Y/P(X_2 | \mathbf{W}, C)], \quad \mathbb{E}_P[\bar{\pi}(Y - 1)/P(X_2 | \mathbf{W}, C)] + 1.$$

Example 4. The data generating mechanism for the fourth scenario is given by:

$$\begin{aligned}
U_Z, U_{X_1Y}, U_{X_2}, U_{X_1}, U_C, U_Z &\sim \mathcal{N}(0, 1), \\
Z &:= f_Z(U_Z) \\
\mathbf{W} &:= \{W_1, \dots, W_5\} \text{ s.t. } W_i \sim \mathcal{N}(0, 1), i = 1, \dots, 5, \\
C &:= f_C(U_C), \\
X_1 &:= f_{X_1}(C, Z, U_{X_1Y}, U_{X_1}, Z), \\
X_2 &:= f_{X_2}(\mathbf{W}, U_{X_2}), \\
Y &:= f_Y(X_1, X_2, C, \mathbf{W}, U_{X_1Y}),
\end{aligned}$$

where,

$$\begin{aligned}
f_C(U_C) &:= U_C, \\
f_Z(U_Z) &:= \mathbb{1}\{U_Z > 0\}, \\
f_{X_1}(U_{X_1Y}, U_{X_1}, Z, U_C) &:= \mathbb{1}\{U_{X_1Y} \cdot (1 + \exp\{0.75U_C + 0.5Z + 0.5\}) > U_{X_1}\}, \\
f_{X_2}(\mathbf{W}, U_{X_2}) &:= \mathbb{1}\{W_1 + W_2 + U_{X_2} > 0\}, \\
f_Y(X_1, X_2, C, \mathbf{W}, U_{X_1Y}) &:= (2X_1 - 1) \cdot (C + 1) + 2 \sin(2C \cdot (2X_1 - 1)) \\
&\quad - 2(2U_{X_1Y} - 1)(1 + 0.5C) + X_2 + W_1
\end{aligned}$$

The outcome Y is then re-scaled to lie in the $[0, 1]$ interval. The target for estimation can be derived with Alg. 1 and equals, for the lower and upper bound respectively,

$$\max_{z \in \{0,1\}} \mathbb{E}_P[\bar{\pi} \mathbb{1}_z(Z)Y/P(Z, X_2 | \mathbf{W}, C)], \quad \min_{z \in \{0,1\}} \mathbb{E}_P[\bar{\pi} \mathbb{1}_z(Z)(Y - 1)/P(Z, X_2 | \mathbf{W}, C)] + 1.$$

C.2 Health Campaign Evaluation

The data was curated from anonymous from Colombia, Peru and Mexico, using a web platform [30]. The exact questions considered in the survey can be found in [30]. The authors made the data available under a Creative Commons license⁵ and is currently hosted by Kaggle as a c.s.v file, accessible through the following link: kaggle.com/code/mpwolke/obesity-levels-life-style/.

⁵creativecommons.org/licenses/by/4.0/




Figure 7: Obesity diagrams.

We considered a number of pre-processing steps. In particular, we created a BMI variable from weight and height data that were a posteriori removed from the dataset, imputed (with median or mode among observed instances) missing values, scaled continuously-valued covariates by subtracting the mean and dividing by the standard deviation, and scaled the outcome BMI to lie in the interval $[0, 1]$. Several covariates were ignored, including data on the consumption of water, consumption of vegetables, consumption of food between meals, consumption of alcohol, and number of main meals. These may be interpreted as unobserved variables, potentially confounding other relationships in the data as specified by the assumed causal diagram. The data from a total of 2111 individuals was recorded in this study.

The bounds on the effect of a policy $\pi^\alpha := \{\pi_H^\alpha, \pi_E^\alpha\}$ acting on H and E , with new assignments defined as $\{P_{\text{new}}(H = \text{rarely}) = \alpha, P_{\text{new}}(E = \text{regularly}) = \alpha\}$ are approximated from the expression in Prop. 4. Following this proposition we find that Age (A), Smoking (S), calories consumption Monitoring (M), and Family history with overweight (F), that is $\mathbf{W} = \{A, S, M, F\}$, form a partial adjustment set for π_E^α given that $(Y \perp\!\!\!\perp_d E \mid \mathbf{W}, H)_{\mathcal{G}_{\pi_{HE}}}$. We can verify also that $\mathbf{Z} = \{T\}$ is a partial instrumental set for π since $(Y \perp\!\!\!\perp_d T)_{\mathcal{G}_\pi}$. For illustration, the corresponding mutilated diagrams are shown in Fig. 7. Finally, we partition the data equally to obtain two sets of samples for training first and second stage classifiers and regressors, respectively, obtaining a first estimate of bounds, before then switching the role of the two data samples and averaging over the resulting estimates.

D NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: The claims made match the theoretical and experimental results presented in the paper. The impact statement in the Appendix reflects how much the results can be expected to generalize to other settings.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We include several remarks in the paper to qualify our contribution, e.g. the remark in Sec. 3, and include an Impact Statement in the Appendix to more thoroughly describe the limitations of our analysis, assumptions, and applicability in real-world settings.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: All theoretical statements are quoted in full in the paper. We have attempted to provide an example to illustrate the significance of each theoretical statement and highlight its implications. The formal proof of all statements is given in the Appendix.

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We provide the data generating mechanisms, details of the target of estimation and information as to what python functions and libraries can be used to fit the proposed estimators. We do not, however, disclose an open source implementation of the proposed methods at this moment.

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in

some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes] and [No]

Justification: The data is publicly available and we have provided full details as to where to access the data and how to run the synthetic data generation pipeline. The code will not be open sourced at this moment but we believe to have provided sufficient details to reproduce our results.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We provided details where applicable. In our case, data splits, hyper-parameters, optimizers, etc., are not significant for the implementation the method.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: We report our results as box plots that include several statistics, e.g. median, 25th, 75th, percentiles, etc., to summarize our results for different seeds of the data and algorithm.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [NA]

Justification: Experiments were run on a single CPU in under a few minutes of wall time.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: We have read the guidelines and we do not think that our work presents any notable concerns.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: We include a broader impact statement in the Appendix.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: No risk of misuse.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [NA]

Justification: The paper does not use existing asset

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.

- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: The paper does not release new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.

- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.