arXiv:2003.12659v1 [stat.ML] 27 Mar 2020

Semiparametric Inference For Causal Effects In Graphical Models With Hidden Variables

Rohit Bhattacharya† RBHATTACHARYA@JHU.EDU
Department of Computer Science
Johns Hopkins University
Baltimore, MD 21218, USA

Razieh Nabi† RNABI@JHU.EDU
Department of Computer Science
Johns Hopkins University
Baltimore, MD 21218, USA

Ilya Shpitser ILYAS@CS.JHU.EDU
Department of Computer Science
Johns Hopkins University
Baltimore, MD 21218, USA

† **Equal contribution**

Abstract

The last decade witnessed the development of algorithms that completely solve the identifiability problem for causal effects in hidden variable causal models associated with directed acyclic graphs. However, much of this machinery remains underutilized in practice owing to the complexity of estimating identifying functionals yielded by these algorithms. In this paper, we provide simple graphical criteria and semiparametric estimators that bridge the gap between identification and estimation for causal effects involving a single treatment and a single outcome. First, we provide influence function based doubly robust estimators that cover a significant subset of hidden variable causal models where the effect is identifiable. We further characterize an important subset of this class for which we demonstrate how to derive the estimator with the lowest asymptotic variance, i.e., one that achieves the *semiparametric efficiency bound*. Finally, we provide semiparametric estimators for *any* single treatment causal effect parameter identified via the aforementioned algorithms. The resulting estimators resemble influence function based estimators that are sequentially reweighted, and exhibit a partial double robustness property, provided the parts of the likelihood corresponding to a set of weight models are correctly specified. Our methods are easy to implement and we demonstrate their utility through simulations.

Keywords: causal inference; unmeasured confounding, semiparametric inference; doubly robust estimation; efficient influence function.

1. Introduction

Causal inference is concerned with the use of observed data to reason about cause effect relationships encoded by counterfactual parameters, such as the average causal effect. Since counterfactual quantities are not directly observed in the data, they must be expressed as functionals of the observed data distribution using assumptions encoded in a *causal*

model. The ease of conveying such assumptions pictorially, via a directed acyclic graph (DAG) (Pearl, 2009; Spirtes et al., 2000), prompted further study of the identifiability of counterfactual quantities in causal models that factorize according to a DAG, when some variables may be hidden or unobserved (Tian and Pearl, 2002a). This led to the development of a complete characterization of the identifiability of the average causal effect (ACE) of a given treatment on a given outcome in all hidden variable causal models associated with a directed acyclic graph (Shpitser and Pearl, 2006; Huang and Valtorta, 2006).

Despite the sophistication of causal identification theory, estimators based on simple covariate adjustment remain the most common strategy for evaluating the ACE from data. Estimates obtained in this way are often biased due to the presence of unmeasured confounding and/or model misspecification. A popular approach for addressing the latter issue has been to use semiparametric estimators developed using the theory of *influence functions* (Bang and Robins, 2005; Tsiatis, 2007). The most popular of these estimators is known as the *augmented inverse probability weighted (AIPW)* estimator and is *doubly robust* in that it gives the analyst two chances to obtain a valid estimate for the ACE – either by specifying the correct model for the treatment assignment given observed covariates that render the treatment assignment *ignorable*, or by specifying the correct model for the dependence of the outcome on the treatment and these covariates. Recent work by Henckel et al. (2019) and Rotnitzky and Smucler (2019) yields methods for constructing statistically efficient versions of AIPW that take advantage of Markov restrictions implied on the obseved data by a fully observed causal model associated with a DAG.

If a causal model contains hidden variables, a.k.a. unmeasured confounders, causal inference becomes considerably more complicated. In the present work, we provide semiparametric estimators for the average causal effect of a single treatment variable on a single outcome variable in increasingly general scenarios, culminating in semiparametric estimators for *any* hidden variable causal model of a DAG in which this effect is identifiable. Weight-based estimators for a subclass of models considered in this paper, were studied in (Jung et al., 2020), and to the best of our knowledge, the *front-door* model (Pearl, 1995) is the only graphical model with unmeasured confounders for which an influence function based estimator has been derived (Fulcher et al., 2017). Other related work includes numerical procedures for approximating the influence function proposed by (Frangakis et al., 2015; Carone et al., 2019). However, such methods are either restricted to settings where simple covariate adjustment is valid, or involve numerical approximations of the function itself which may be computationally prohibitive. Our contributions can be summarized as follows.

In Section 3, we introduce a simple extension of AIPW that we term *generalized AIPW (gAIPW)*, to settings involving unmeasured variables. We then propose two novel IPW estimators, *primal IPW* and *dual IPW*, for settings in which covariate adjustment is not possible due to the presence of unmeasured confounders. We show that these estimators use variationally independent components of the joint likelihood on the observed margin of the hidden variable DAG and demonstrate that this leads to an influence function based semiparametric estimator that can be viewed as augmentation of the primal form. We call this *augmented primal IPW (APIPW)*. We then study the robustness properties of APIPW and show that it can be reformulated in a way that is doubly robust in the models involved in the primal and dual IPW estimators.

In Sections 4 and 5, we study equality restrictions on the tangent space implied by a hidden variable DAG model. Such restrictions are important as they play a role in deriving the most efficient influence function based estimator for a given parameter. In the special case where the model is nonparametric saturated, no restrictions are imposed on the tangent space, and the influence function is unique (and thus efficient). We provide Algorithm 1 as a *sound* and *complete* procedure for checking whether a hidden variable causal model that factorizes as a DAG imposes equality restrictions on the observed data tangent space, provided the hidden variables in the model are unrestricted. We then define a class of hidden variable causal models for which the restrictions on the tangent space resemble those of a DAG model with no hidden variables. For this class of models, we derive the space of all influence functions and consequently identify all *regular and asymptotically linear* estimators of the causal effects that are consistent and asymptotically normal. We further show how to derive the most efficient gAIPW or APIPW estimators, as appropriate, within this class of causal models. That is, we provide forms for the influence function based estimator that attains the *semiparametric efficiency bound*.

In Section 6, we describe semiparametric estimators for increasingly general classes of functionals representing identifiable causal effects of a single treatment on a single outcome, culminating with an estimator for *any* such functional. We first operationalize primal and dual IPW via the *primal* and *dual fixing* operators so that they may be applied recursively to a hidden variable causal model. We show how repeated applications of these operators to other variables may lead to a model where they can be applied to the treatment variable itself. The resulting estimators resemble influence function based estimators that are sequentially reweighted, and exhibit a partial double robustness property, i.e., it is doubly robust provided the parts of the likelihood corresponding to a set of weight models are correctly specified. We then propose the *nested IPW* estimator that generalizes IPW to all hidden variable causal models where the target parameter is identified. We propose a *sound* and *complete* algorithm (Algorithm 2) that derives the corresponding nested IPW estimator when possible. Finally, we derive the *augmented nested IPW* estimator that resembles gAIPW and is partially doubly robust to specification of a treatment assignment and outcome model, albeit after correct specification of certain pieces of the observed data likelihood that serve as inverse weights in our estimator.

2. Causal Inference with Graphical Models

The cause-effect relationship between a single treatment T and an outcome Y is typically established through the use of potential outcomes, a.k.a. counterfactuals. For example, the potential outcomes $Y(1)$ and $Y(0)$ may be used to represent a hypothetical randomized controlled trial where units are randomly assigned to the treatment arm (corresponding to $T = 1$), or the control arm (corresponding to $T = 0$). A comparison of the distribution of these counterfactual random variables is then typically conducted on the mean difference scale, known as the *average causal effect* (ACE), i.e., $\mathbb{E}[Y(1)] - \mathbb{E}[Y(0)]$. More generally, we may define the potential outcome $Y(t)$ corresponding to the potential outcome had treatment T been *assigned* to some value t. This allows for the contrast of arbitrary treatment assignments t and t' as $\mathbb{E}[Y(t)] - \mathbb{E}[Y(t')]$. Thus, throughout the paper, we set our target

of inference to be the mean of the counterfactual random variable $Y(t)$. That is,

$$\psi(t) \equiv \mathbb{E}[Y(t)]. \qquad \textit{(target parameter)} \tag{1}$$

2.1 Statistical and Causal Models of a Directed Acyclic Graph

It is well known that the target parameter $\psi(t)$ cannot be expressed as a function of the observed data or, in other words, is not *identified* if no assumptions are made about the data generating process (Pearl, 2009). The use of graphs for causal inference has become increasingly popular as they provide a succinct pictorial representation of substantive non-parametric assumptions made by the data analyst (Greenland et al., 1999; Williams et al., 2018; Hünermund and Bareinboim, 2019). A directed acyclic graph (DAG) $\mathcal{G}(V)$ is defined as a set of nodes V connected through directed edges such that there are no directed cycles. When the vertex set is clear from the given context, we will often abbreviate $\mathcal{G}(V)$ as simply $\mathcal{G}$. DAGs have been used to define both statistical and causal models. *Statistical models* of a DAG $\mathcal{G}(V)$ are sets of distributions that factorize as,

$$p(V) = \prod_{V_i \in V} p(V_i \mid \text{pa}_{\mathcal{G}}(V_i)), \qquad \textit{(DAG factorization)} \tag{2}$$

where $\text{pa}_{\mathcal{G}}(V_i)$ are the parents of V_i in $\mathcal{G}$ (Pearl, 1988).

The absence of edges between variables in $\mathcal{G}$, relative to a complete DAG entails conditional independence facts in $p(V)$. These can be read off directly from the DAG $\mathcal{G}$ by the well-known d-separation criterion (Pearl, 2009). That is, for disjoint sets $X, Y,$ and Z, the following *global Markov property* holds: $(X \perp\!\!\!\perp_{\text{d-sep}} Y \mid Z)_{\mathcal{G}} \implies (X \perp\!\!\!\perp Y \mid Z)_{p(V)}$. When the context is clear, we will simply use $X \perp\!\!\!\perp Y \mid Z$ to denote the conditional independence between X and Y given Z.

Causal models are sets of distributions defined over counterfactual random variables. Causal models of a DAG $\mathcal{G}(V)$ may be defined over counterfactual random variables $V_i(pa_i)$ for each $V_i \in V$ where pa_i is a set of values for $\text{pa}_{\mathcal{G}}(V_i)$. These counterfactuals can alternatively be viewed as being determined by a system of *structural equations* $f_i(pa_i, \epsilon_i)$ that map values pa_i, as well as values of an exogenous noise term ϵ_i to values of V_i (Pearl, 2009; Malinsky et al., 2019). Other counterfactuals may be defined from above via recursive substitution. Specifically for any set $A \subseteq V$, and a variable V_i, we have:

$$V_i(a) \equiv V_i(a \cap \text{pa}_{\mathcal{G}}(V_i), \{V_j(a) : V_j \in \text{pa}_{\mathcal{G}}(V_i) \setminus A\}).$$

For any set $A \subset V$, the distribution of the potential outcome $p(\{V \setminus A\}(a))$ or $p(V(a))$ for short, is identified in a causal model of a DAG $\mathcal{G}$ by the g-formula functional (Robins, 1986):

$$p(V(a)) = \prod_{V_i \in V \setminus A} p(V_i \mid a \cap \text{pa}_{\mathcal{G}}(V_i), \text{pa}_{\mathcal{G}}(V_i) \setminus A). \qquad \textit{(g-formula)} \tag{3}$$

Note that if A is the empty set, we obtain the DAG factorization for $\mathcal{G}$, meaning that a causal model of a DAG $\mathcal{G}$ implies the statistical model of a DAG $\mathcal{G}$.

2.2 Statistical Inference for the Adjustment Functional

Having briefly discussed causal models of a DAG, we now provide a short overview of estimation theory surrounding the target $\psi(t)$ when it is derived from such a model. In all causal models of a DAG $\mathcal{G}$ that are typically used, the target parameter $\psi(t)$ is identified via the *back-door adjustment* formula as follows, where $C = \text{pa}_{\mathcal{G}}(T)$,

$$\psi(t) = \mathbb{E}\big[\mathbb{E}[Y \mid T = t, C]\big]. \qquad \textit{(adjustment functional)} \tag{4}$$

Once the target parameter is identified, causal inference reduces to statistical inference, specifically to an estimation problem of the identifying functional – in our case $\psi(t)$. There is a rich literature on estimation methods for this functional (Robins et al., 1994; Hahn, 1998; Robins, 2000; van der Laan and Rose, 2011; Kennedy et al., 2017).

If a parametric likelihood can be correctly specified for the statistical DAG model of the observed data distribution, then an efficient estimator for $\psi(t)$ may be derived using the plug-in principle. In the commonly assumed case where the DAG corresponding to the observed data distribution is complete, the plug-in estimator for $\psi(t)$ reduces to $\mathbb{P}_n[\mu_t(C; \widehat{\eta}_1)]$, where $\mathbb{P}_n[.] := \frac{1}{n}\sum_{i=1}^{n}(.)$, $\mu_t(C; \eta_1)$ is the correctly specified parametric form for $\mathbb{E}[Y \mid T = t, C]$, and $\widehat{\eta}_1$ are the maximum likelihood values of η_1.

Since assuming a correctly specified parametric observed data likelihood, or even a correctly specified outcome regression $\mu_t(C; \eta)$ is unrealistic in practice, a variety of other estimators have been developed that place *semiparametric* restrictions on the observed data distribution. One such estimator, based on *inverse probability weighting (IPW)*, seeks to compensate for a biased treatment assignment by reweighing observed outcomes of units assigned $T = t$ by the inverse of the normalized treatment assignment probability $p(T = t \mid C)$. If this probability has a known parametric form $\pi_t(C; \eta_2) \equiv p(T = t \mid C)$, the IPW estimator takes the form $\mathbb{P}_n[\frac{\mathbb{I}(T=t)}{\pi_t(C;\widehat{\eta}_2)} \times Y]$, where $\mathbb{I}(.)$ is the indicator function, and $\widehat{\eta}_2$ are the maximum likelihood estimates of η_2. While the IPW estimator is inefficient, it is simple to implement, and is often used in cases where the treatment assignment model $\pi_t(C; \eta_2)$ is known by design, as is often the case in controlled trials.

The plug-in and IPW estimators of $\psi(t)$ are both $\sqrt{n}$*-consistent and asymptotically normal* if the models they rely on, $\mu_t(C; \eta_1)$ and $\pi_t(C; \eta_2)$ respectively, are parametric and correctly specified. Otherwise, these estimators are no longer consistent. If flexible models are used for $\mu_t(C)$ and $\pi_t(C)$ instead, the resulting estimators may remain consistent, but converge to the true value of $\psi(t)$ at unacceptably slow rates; see (Chernozhukov et al., 2018) for examples.

A principled alternative is to consider *regular and asymptotically linear* (RAL) estimators (Robins et al., 1994; Tsiatis, 2007). An estimator $\widehat{\psi}_n$ of a scalar parameter ψ based on n i.i.d copies $Z_1, \ldots, Z_n$ drawn from $p(Z)$ is a RAL estimator if there exists a measurable random function $U_\psi(Z)$ with mean zero and finite variance such that

$$\sqrt{n} \times (\widehat{\psi}_n - \psi) = \frac{1}{\sqrt{n}} \times \sum_{i=1}^{n} U_\psi(Z_i) + o_p(1), \qquad \textit{(RAL estimator of } \psi\textit{)} \tag{5}$$

where $o_p(1)$ is a term that converges in probability to zero as n goes to infinity. The random variable $U_\psi(Z)$ is called the *influence function (IF)* of the estimator $\widehat{\psi}_n$.

A RAL estimator $\widehat{\psi}_n$ is *consistent and asymptotically normal* (CAN), with asymptotic variance equal to the variance of its influence function U_ψ: $\sqrt{n}\times(\widehat{\psi}_n-\psi) \xrightarrow{d} \mathcal{N}\big(0,\ \mathrm{var}(U_\psi)\big)$; we use U_ψ as a shorthand for $U_\psi(Z)$. IFs in semiparametric models are derived as normalized elements of the orthogonal complement of the tangent space of the model. In a nonparametric saturated model (one with an unrestricted tangent space), the IF is unique. In a semiparametric model, there are many IFs, with a unique one that attains the semiparametric efficiency bound obtained by projecting onto the tangent space of the model; see Appendix D and (Bickel et al., 1993; Tsiatis, 2007) for more details.

In the nonparametric saturated model, corresponding to the complete DAG, the unique influence function for $\psi(t)$ is given by $U_{\psi_t} = \frac{\mathbb{I}(T=t)}{\pi_t(C)} \times \{Y - \mu_t(C)\} + \mu_t(C) - \psi(t)$, yielding the *AIPW* estimator: $\mathbb{P}_n\Big[\frac{\mathbb{I}(T=t)}{\pi_t(C;\widehat{\eta}_2)} \times \{Y - \mu_t(C;\widehat{\eta}_1)\} + \mu_t(C;\widehat{\eta}_1)\Big]$. Given the standard factorization of the complete DAG as $p(Y \mid A, C) \times p(A \mid C) \times p(C)$, the propensity score model $\pi_t(C)$ and the outcome regression model $\mu_t(C)$ are variationally independent. Further, the bias of this estimator is a product of the biases of its nuisance functions $\pi_t(C)$ and $\mu_t(C)$. As a result, the AIPW estimator exhibits the *double robustness property*, where it remains consistent if *either* of the two nuisance models $\pi_t(C)$ or $\mu_t(C)$ is specified correctly, even if the other is arbitrarily misspecified.

In a semiparametric model of a DAG, which is defined by conditional independence restrictions on the tangent space implied by the DAG factorization, the above influence function can be projected onto the tangent space of the model to improve efficiency, see (Rotnitzky and Smucler, 2019) for details.

2.3 Causal Graphical Models with Hidden Variables

While estimation theory for fully observed causal models represented by DAGs is very well developed, causal models most relevant to practical applications are sure to contain variables that are unmeasured or hidden to the data analyst. In such cases, the observed data distribution $p(V)$ may be considered to be a margin of a distribution $p(V \cup H)$ associated with a DAG $\mathcal{G}(V \cup H)$ where vertices in V correspond to observed variables and vertices in H correspond to unmeasured or hidden variables. Two complications arise from the presence of hidden variables. First, the target parameter $\psi(t)$ may not always be identified as a function of the observed data and second, parameterizations of latent variable models are generally not fully identifiable and may contain singularities (Drton, 2009).

A natural alternative to the latent variable model is one that places no restrictions on $p(V)$ aside from those implied by the Markov restrictions given by the factorization of $p(V \cup H)$ with respect to $\mathcal{G}(V \cup H)$. It was shown in (Evans, 2018) that all *equality constraints* implied by such a factorization are captured by a nested factorization of $p(V)$ with respect to an *acyclic directed mixed graph (ADMG)* $\mathcal{G}(V)$ derived from $\mathcal{G}(V \cup H)$ via the *latent projection* operation described by Verma and Pearl (1990). Such an ADMG is a smooth supermodel of infinitely many hidden variable DAGs that share the same identification theory for $\psi(t)$ and imply the same equality constraints on the margin $p(V)$ (Richardson et al., 2017; Evans and Richardson, 2019). Thus, our use of ADMGs for identification and estimation of the target $\psi(t)$ is without loss of generality.

The latent projection of a hidden variable DAG $\mathcal{G}(V \cup H)$ onto observed variables V is an ADMG $\mathcal{G}(V)$ with directed ($\rightarrow$) and bidirected ($\leftrightarrow$) edges constructed as follows. The

edge $A \to B$ exists in $\mathcal{G}(V)$ if there exists a directed path from A to B in $\mathcal{G}(V \cup H)$ with all intermediate vertices in H. An edge $A \leftrightarrow B$ exists in $\mathcal{G}(V)$ if there exists a collider-free path (i.e., there are no consecutive edges of the form $\to \circ \leftarrow$) from A to B in $\mathcal{G}(V \cup H)$ with all intermediate vertices in H, such that the first edge on the path is an incoming edge into A and the final edge is an incoming edge into B. An example of the latent projection is provided in Appendix B. Conditional independence facts pertaining to $p(V)$ can then be read off from the ADMG $\mathcal{G}(V)$ by a simple analogue of the d-separation criterion, known as m-separation, that generalizes the notion of a collider to include mixed edges of the form $\to \circ \leftrightarrow, \leftrightarrow \circ \leftarrow$, and $\leftrightarrow \circ \leftrightarrow$, (Richardson, 2003).

The bidirected connected components of an ADMG $\mathcal{G}(V)$ partition its vertices into distinct subsets known as *districts*. A set $S \subseteq V$ is a district in $\mathcal{G}(V)$ if it forms a maximal connected component via only bidirected edges. We use $\mathrm{dis}_{\mathcal{G}}(V_i)$ to denote the district of V_i in $\mathcal{G}$, which includes V_i itself, and $\mathcal{D}(\mathcal{G})$ to denote the set of all districts in $\mathcal{G}$. A distribution $p(V)$ is said to *district factorize* or *Tian factorize* with respect to an ADMG $\mathcal{G}(V)$ if

$$p(V) = \prod_{D \in \mathcal{D}(\mathcal{G})} q_D(D \mid \mathrm{pa}_{\mathcal{G}}(D)), \qquad \textit{(Tian ADMG factorization)} \tag{6}$$

where the parents of a set of vertices D is defined as the set of parents of D not already in D, i.e., $\mathrm{pa}_{\mathcal{G}}(D) \equiv \bigcup_{D_i \in D} \mathrm{pa}_{\mathcal{G}}(D_i) \setminus D$. We follow the same convention for children of a set S, denoted $\mathrm{ch}_{\mathcal{G}}(S)$. For other standard genealogical relations defined for a single vertex V_i, such as ancestors $\mathrm{an}_{\mathcal{G}}(V_i) \equiv \{V_j \in V \mid \exists\, V_j \to \cdots \to V_i \text{ in } \mathcal{G}\}$ and descendants $\mathrm{de}_{\mathcal{G}}(V_i) \equiv \{V_j \in V \mid \exists\, V_i \to \cdots \to V_j \text{ in } \mathcal{G}\}$, both of which include V_i itself by convention, the extension to a set S uses the disjunctive definition which also includes the set itself. For example, $\mathrm{an}_{\mathcal{G}}(S) = \bigcup_{S_i \in S} \mathrm{an}_{\mathcal{G}}(S_i)$. A list of notation and definitions used in this paper can be found in Appendix A.

The use of q in lieu of p in Eq. 6, emphasizes the fact that these factors are not necessarily ordinary conditional distributions. Each factor $q_D(D \mid \mathrm{pa}_{\mathcal{G}}(D))$ may in fact be treated as a post-intervention distribution where all variables outside of D are intervened on and held fixed to some constant value (Tian and Pearl, 2002a). Thus, we use $q_S(\cdot \mid \cdot)$ to denote probability distributions where only variables in S are random and all others are fixed. Such densities are often referred to as *kernels* and are similar to conditional densities in the sense that they provide a mapping from values of elements past the conditioning bar to normalized densities over variables prior to the conditioning bar (Lauritzen, 1996). Conditioning and marginalization in kernels are defined in the usual way.

Tian and Pearl (2002a) also showed that each kernel $q_D(D \mid \mathrm{pa}_{\mathcal{G}}(D))$ in Eq. 6 is a function of $p(V)$ as follows. Define the *Markov blanket* of a vertex V_i as the district of V_i and the parents of its district, excluding V_i itself, i.e., $\mathrm{mb}_{\mathcal{G}}(V_i) = \mathrm{dis}_{\mathcal{G}}(V_i) \cup \mathrm{pa}_{\mathcal{G}}(\mathrm{dis}_{\mathcal{G}}(V_i)) \setminus V_i$. Consider a valid topological order τ on all k vertices in V, that is a sequence $(V_1, \ldots, V_k)$ such that no vertex appearing later in the sequence is an ancestor of vertices earlier in the sequence. Let $\{\preceq_\tau V_i\}$ denote the set of vertices that precede V_i in this sequence, including V_i itself. Then for each $D \in \mathcal{D}(\mathcal{G})$,

$$q_D(D \mid \mathrm{pa}_{\mathcal{G}}(D)) = \prod_{D_i \in D} p(D_i \mid \mathrm{mp}_{\mathcal{G}}(D_i)), \qquad \textit{(Identification of Tian factors)} \tag{7}$$

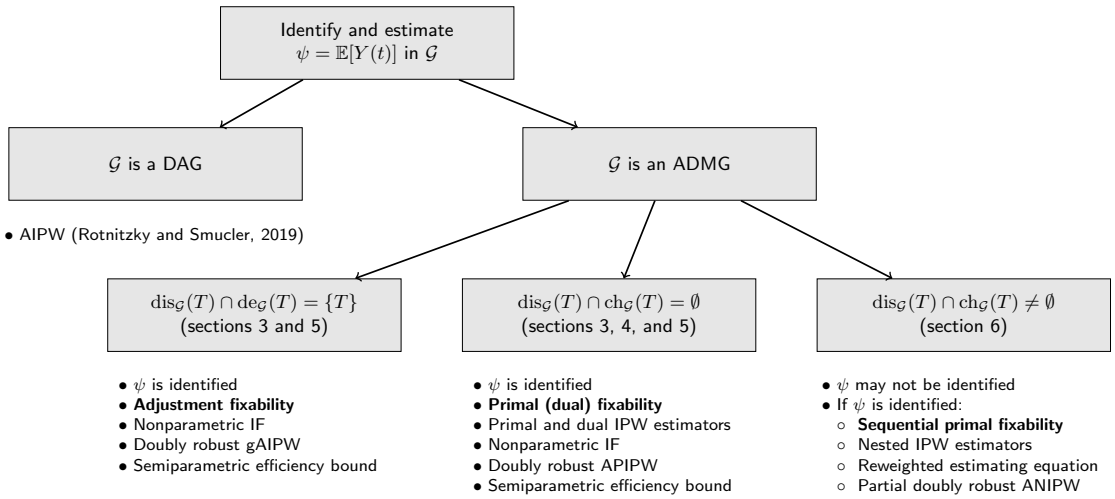

Figure 1: Roadmap

where $\mathrm{mp}_{\mathcal{G}}(V_i)$, the *Markov pillow* of V_i, is defined as its Markov blanket in a subgraph restricted to V_i and its predecessors according to the topological ordering. More formally, $\mathrm{mp}_{\mathcal{G}} \equiv \mathrm{mb}_{\mathcal{G}_S}(V_i)$ where $S = \{\preceq_\tau V_i\}$, and $\mathcal{G}_S$ is the subgraph of $\mathcal{G}$ that is restricted to vertices in S and the edges between these vertices. This leads to a factorization of the observed law as a product of simple conditional factors according to the topological order,

$$p(V) = \prod_{V_i \in V} p(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i)). \qquad \textit{(Topological ADMG factorization)} \tag{8}$$

The above factorization (and the Tian factorization from which it is derived), does not always capture every conditional independence constraint in $p(V)$ implied by the Markov property of the underlying hidden variable DAG $\mathcal{G}(V \cup H)$. However, it is particularly simple to work with, and under some conditions, which we derive in Section 4, is capable of capturing all such constraints. A more general set of factorizations that is guaranteed to capture all ordinary conditional independence constraints, as well as a nested factorization that captures all equality constraints was described in (Richardson, 2003) and (Richardson et al., 2017) respectively. Details on the nested factorization are provided in Section 6, where we use it in order to derive estimators for any identifiable $\psi(t)$.

Fig. 1 shows the roadmap for the rest of our paper. All proofs can be found in Appendix E. For the remainder of the paper, we will assume without loss of generality that the outcome Y is a descendant of T, else the target $\psi(t)$ is identified trivially as $\mathbb{E}[Y]$ which does not pose an interesting inference problem.

3. Doubly Robust Estimation of Causal Effects

In this section, we consider two classes of ADMGs where $\psi(t)$ is nonparametrically identified: one where $\mathrm{dis}_{\mathcal{G}}(T) \cap \mathrm{de}_{\mathcal{G}}(T) = \{T\}$, and one where $\mathrm{dis}_{\mathcal{G}}(T) \cap \mathrm{ch}_{\mathcal{G}}(T) = \{\}$. We call the

first criterion *adjustment* fixability and the second criterion *primal* (or *dual*) fixability. We evaluate our target $\psi(t)$ in a semiparametric model associated with these ADMGs by deriving an influence function based estimator with desriable statistical properties as discussed below.

3.1 Generalized Augmented IPW Estimators

Consider the class of ADMGs where T has no bidirected path to any of its descendants i.e., $\mathrm{dis}_{\mathcal{G}}(T) \cap \mathrm{de}_{\mathcal{G}}(T) = \{T\}$. We term this criterion *adjustment fixability* or *a-fixability* due to its resemblance to the fixability criterion put forward in (Richardson et al., 2017) and its close relation to the validity of covariate adjustment for identification of the target. Fix any valid topological order τ on V such that T appears last among the members of its district. Such an ordering is possible due to a-fixability requiring that no member of the district of T is also a descendant of T. Under such an ordering $\mathrm{mp}_{\mathcal{G}}(T) = \mathrm{mb}_{\mathcal{G}}(T)$ and forms a valid adjustment set for the target $\psi(t)$ as formalized in the following lemma.

Lemma 1 (Identifying functional when T is a-fixable)
Given a distribution $p(V)$ that district factorizes with respect to an ADMG $\mathcal{G}(V)$ in which T is adjustment fixable, $\psi(t)$ is identified as $\psi(t) = \mathbb{E}[\mathbb{E}[Y \mid T = t, \mathrm{mp}_{\mathcal{G}}(T)]]$.

The identifiability of the target in this manner, immediately yields a nonparametric influence function that has the same form as AIPW for DAGs, except the conditioning set is now extended to include members of the district of T and parents of this district. Hence, we refer to the resulting influence function and its estimator as *generalized AIPW* or *gAIPW* for short.

Theorem 2 (Nonparametric influence function of gAIPW)
Under the same conditions as Lemma 1, the nonparametric influence function for the target parameter $\psi(t)$ is as follows.

$$U_{\psi_t} = \frac{\mathbb{I}(T = t)}{p(T \mid \mathrm{mp}_{\mathcal{G}}(T))} \times \Big(Y - \mathbb{E}[Y \mid T = t, \mathrm{mp}_{\mathcal{G}}(T)]\Big) + \mathbb{E}[Y \mid T = t, \mathrm{mp}_{\mathcal{G}}(T)] - \psi(t). \tag{9}$$

It is easy to see that an estimator derived from gAIPW should exhibit similar properties of double robustness as ordinary AIPW. In particular, gAIPW estimators for $\psi(t)$ are consistent if either the propensity score model $p(T \mid \mathrm{mp}_{\mathcal{G}}(T))$ or the outcome regression model $\mathbb{E}[Y \mid T = t, \mathrm{mp}_{\mathcal{G}}(T)]$ is correctly specified.

Lemma 3 (Double robustness of gAIPW)
The estimator obtained by solving the estimating equation $\mathbb{E}[U_{\psi_t}] = 0$, where U_{ψ_t} is given in Theorem 2, is consistent if either $p(T \mid \mathrm{mp}_{\mathcal{G}}(T))$ or $\mathbb{E}[Y \mid T = t, \mathrm{mp}_{\mathcal{G}}(T)]$ is correctly specified.

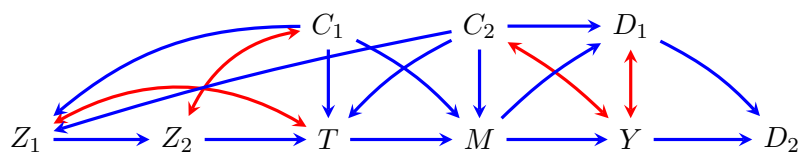

Figure 2: An acyclic directed mixed graph where T is adjustment fixable.

3.1.1 Example: a-fixability and gAIPW Estimator

As a concrete example, consider the ADMG in Fig. 2. The variable T is a-fixable since $\text{dis}_{\mathcal{G}}(T) = \{Z_1, T\}$ and $\text{de}_{\mathcal{G}}(T) = \{T, M, Y, D_1, D_2\}$, and hence $\text{dis}_{\mathcal{G}}(T) \cap \text{de}_{\mathcal{G}}(T) = \{T\}$. Consequently, $\psi(t)$ is identified via Lemma 1 as $\mathbb{E}\big[\mathbb{E}[Y \mid T = t, \text{mp}_{\mathcal{G}}(T)]\big]$, where $\text{mp}_{\mathcal{G}}(T) = \{C, Z\}, C = \{C_1, C_2\}, Z = \{Z_1, Z_2\}$. The corresponding gAIPW estimator is obtained by solving $\mathbb{E}[U_{\psi_t}] = 0$, where U_{ψ_t} is the nonparametric IF given in Theorem 2. This yields the following doubly robust gAIPW estimator of $\psi(t)$,

$$\text{(Example: Fig. 2)} \qquad \psi_{\text{gAIPW}} = \mathbb{P}_n\Big[\, \frac{\mathbb{I}(T = t)}{\widehat{\pi}_t(Z, C)} \times \big(Y - \widehat{\mu}_t(Z, C)\big) + \widehat{\mu}_t(Z, C) \,\Big]$$

where $\widehat{\pi}_t(Z, C)$ and $\widehat{\mu}_t(Z, C)$ are estimates of the propensity score $p(T = t \mid Z, C)$, and the outcome regression $\mathbb{E}[Y \mid T = t, Z, C]$, respectively.

While the estimator proposed above is doubly robust in the propensity score and outcome regression models, it is not the most *efficient*, i.e., one with the lowest asymptotic variance. This is because there are several conditional independence constraints (e.g., $C_1 \perp\!\!\!\perp C_2$) implied by the ADMG $\mathcal{G}(V)$ that are not exploited in the nonparametric gAIPW estimator. We note that a-fixability of T is a *sufficient* but not *necessary* condition under which covariate adjustment is a valid strategy for identification of $\psi(t)$. For a more general criterion, we refer the reader to (Shpitser et al., 2010; Perković et al., 2015). However, requiring that T is a-fixable proves useful for deriving a closed form expression of the efficient influence function as we will see in Section 5 when we talk about the semiparametric efficiency bound and derive the most efficient estimator for $\psi(t)$ in a class of ADMGs where T is a-fixable.

3.2 Primal and Dual IPW Estimators

We now shift our focus to a broader class of ADMGs where adjustment is not a valid strategy in estimating the target parameter $\psi(t)$. Consider the ADMGs shown in Fig. 3. Since there exists a bidirected edge from T to L and L is a descendant of T, T is not a-fixable in either of these ADMGs. Further, it is easy to see that no valid adjustment set exists to identify the causal effect of T on Y. However, such an effect is indeed identified in both graphs. The defining characteristic of both these ADMGs that permits identification of the target $\psi(t)$, is that the district of T does not intersect with the children of T.

Thus, in this section, we consider the class of ADMGs where T is not a-fixable (and more generally, there exists no valid adjustment set), however, T has no bidirected path to any of its children, i.e., $\text{dis}_{\mathcal{G}}(T) \cap \text{ch}_{\mathcal{G}}(T) = \emptyset$. We term this criterion *primal fixability* or *p-fixability*, due to its close relation to the primal fixing operator introduced in Section 6.

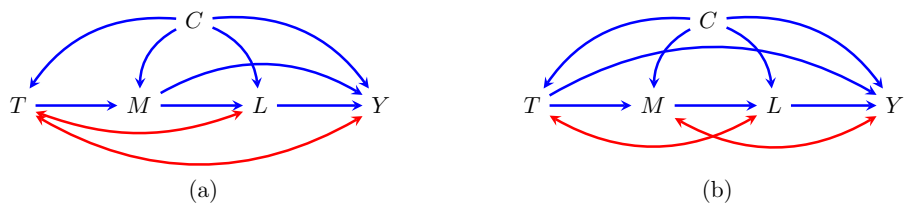

Figure 3: Examples of acyclic directed mixed graphs where T is primal fixable.

This encompasses many popular models in the literature including those that satisfy the front-door criterion (Pearl, 1995). Further, primal fixability was proved to be necessary and sufficient for identification of the effect of T on all other variables $V \setminus T$ by Tian and Pearl (2002a). This makes primal fixability an appealing starting point for the development of doubly robust semiparametric estimators that go beyond gAIPW.

In observed data distributions $p(V)$ that district factorize according to an ADMG $\mathcal{G}(V)$ where T is primal fixable, the identifying functional for the target is as follows,

$$\psi(t) = \sum_{V \setminus T} Y \times \prod_{V_i \in V \setminus D_T} p(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i)) \sum_{T} \prod_{V_i \in D_T} p(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i)) \Big|_{T=t}, \tag{10}$$

where D_T denotes the district of T (Tian and Pearl, 2002a). We provide special notation for the district of T as D_T due to its frequent occurrence in subsequent results.

In the following lemmas, we describe two IPW estimators of the target parameter $\psi(t)$, that are consistent if a particular set of models are specified correctly. We show that the sets of models used for these estimators are variationally independent in much the same way as the propensity score and outcome regression models used in the AIPW estimator are variationally independent.

Since these IPW estimators offer different perspectives in estimating the same parameter, we draw inspiration from the optimization literature in naming the probabilistic operations that lead to them as the primal and dual fixing operators (wrt primal and dual formulations in optimization (Dantzig et al., 1956; Boyd and Vandenberghe, 2004)), and the estimators themselves as primal and dual IPW.

For the remainder of this paper, we assume a fixed valid topological ordering τ where the treatment T appears later than all of its non-descendants i.e., $T \succ_\tau V \setminus \mathrm{de}_{\mathcal{G}}(T)$ and the outcome Y appears earlier than all of its non-descendants i.e., $Y \prec_\tau V \setminus \mathrm{de}_{\mathcal{G}}(Y)$. This allows for easier exposition by fixing the definition of pre-treatment covariates as being any variable that appears earlier than T under the ordering τ. In introducing primal IPW below, we use $\{\succeq T\}$ (dropping subscript τ for readability) to mean the set of vertices (including T) that succeed T under the topological order τ.

Lemma 4 (Primal IPW)
Given a distribution $p(V)$ that district factorizes with respect to an ADMG $\mathcal{G}(V)$ where T is primal fixable, $\psi(t) = \psi(t)_{primal} \equiv \mathbb{E}[\beta(t)_{primal}]$ where

$$\begin{aligned}\beta(t)_{primal} &\equiv \frac{\mathbb{I}(T=t)}{q_{D_T}(T \mid \mathrm{mb}_{\mathcal{G}}(T))} \times Y \\ &= \mathbb{I}(T=t) \times \frac{\sum_T \prod_{V_i \in D_T \cap \{\succeq T\}} p(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i))}{\prod_{V_i \in D_T \cap \{\succeq T\}} p(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i))} \times Y. \end{aligned} \tag{11}$$

Corollary 5 (Primal IPW estimator)
This representation of $\psi(t)$ immediately yields a primal IPW estimator that is consistent if the set of models in $\{p(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i)), \forall V_i \in D_T \cap \{\succeq T\}\}$ is correctly specified.

The kernel $q_{D_T}(T \mid \mathrm{mb}_{\mathcal{G}}(T))$ in Lemma 4 may be viewed as a *nested* propensity score derived from the post-intervention distribution $q_{D_T}(D_T \mid \mathrm{pa}_{\mathcal{G}}(D_T))$ where all variables outside of D_T are intervened on and held fixed to some constant value. Recall that the kernel $q_{D_T}(D_T \mid \mathrm{pa}_{\mathcal{G}}(D_T))$ is identified as $\prod_{V_i \in D_T} p(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i))$ as in Eq. 7. Consequently, $q_{D_T}(T \mid \mathrm{mb}_{\mathcal{G}}(T))$ is identified by the definition of conditioning on all elements in D_T outside of T in the kernel $q_{D_T}(D_T \mid \mathrm{pa}_{\mathcal{G}}(D_T))$ as,

$$\begin{aligned} q_{D_T}(T \mid \mathrm{mb}_{\mathcal{G}}(T)) = q_{D_T}(T \mid D_T \cup \mathrm{pa}_{\mathcal{G}}(D_T) \setminus T) &= \frac{q_{D_T}(D_T \mid \mathrm{pa}_{\mathcal{G}}(D_T))}{q_{D_T}(D_T \setminus T \mid \mathrm{pa}_{\mathcal{G}}(D_T))} \\ &= \frac{q_{D_T}(D_T \mid \mathrm{pa}_{\mathcal{G}}(D_T))}{\sum_T q_{D_T}(D_T \mid \mathrm{pa}_{\mathcal{G}}(D_T))} = \frac{\prod_{V_i \in D_T} p(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i))}{\sum_T \prod_{V_i \in D_T} p(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i))}. \end{aligned}$$

The final expression simplifies further by noticing that all vertices appearing prior to T under the topological order τ, do not contain T in their Markov pillows. Consequently, $p(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i))$ is not a function of T if $V_i \prec T$. Thus, these terms may be pulled out of the summation in the denominator, and cancel with the corresponding term in the numerator. This gives us the resulting primal IPW formulation in Eq. 11.

We now introduce the dual formulation. Define the *inverse Markov pillow* of a vertex V_i to be all other vertices V_j outside of the district of V_i, such that V_i is a member of the Markov pillow of V_j. More formally, $\mathrm{mp}_{\mathcal{G}}^{-1}(V_i) = \{V_j \in V \mid V_j \notin \mathrm{dis}_{\mathcal{G}}(V_i), V_i \in \mathrm{mp}_{\mathcal{G}}(V_j)\}$.

Lemma 6 (Dual IPW)
Given a distribution $p(V)$ that district factorizes with respect to an ADMG $\mathcal{G}(V)$ where T is primal fixable, $\psi(t) = \psi(t)_{dual} \equiv \mathbb{E}[\beta(t)_{dual}]$ where

$$\beta(t)_{dual} = \frac{\prod_{V_i \in \mathrm{mp}_{\mathcal{G}}^{-1}(T)} p(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i)) \mid_{T=t}}{\prod_{V_i \in \mathrm{mp}_{\mathcal{G}}^{-1}(T)} p(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i))} \times Y. \tag{12}$$

Corollary 7 (Dual IPW estimator)
This representation of $\psi(t)$ immediately yields a dual IPW estimator that is consistent if the set of models in $\{p(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i)), \forall V_i \in \mathrm{mp}_{\mathcal{G}}^{-1}(T)\}$ is correctly specified.

It is easy to see that in the regular conditionally ignorable model, the primal and dual IPW estimators correspond to the standard IPW and outcome regression plug-in estimators respectively. More generally, primal IPW can be viewed as a generalization of the g-formula to kernel factorizations that arise in ADMGs. The ordinary g-formula for a DAG model involves truncation of the DAG factorization, namely dropping a simple conditional factor of the treatment given its parents, i.e., $p(V(t)) = \{p(V)/p(T = t \mid \text{pa}_{\mathcal{G}}(T))\}|_{T=t}$. On the other hand, the primal formulation, or the *nested* g-formula, can be viewed as truncation of the Tian factorization in Eq. 10, where the nested conditional factor for the treatment given its Markov blanket is dropped from the observed joint distribution, i.e., $p(V(t)) = \{p(V)/q_{D_T}(T \mid \text{mb}_{\mathcal{G}}(T))\}|_{T=t}$. Intuition for the dual IPW can be gained by viewing it as a probabilistic formalization of the node splitting operation in single world intervention graphs (SWIGs) described in (Richardson and Robins, 2013).

For the remainder of the paper, we occasionally assume dependence on t to be implicit and for simplicity of notation, write $\psi(t)_{\text{primal}}$ as simply ψ_{primal} and $\beta_{\text{primal}}(t)$ as β_{primal} for example. We now present a powerful result which guarantees that the primal and dual IPW estimators use variationally independent components of the observed distribution $p(V)$.

Lemma 8 (Variational independence of primal and dual IPW estimators) *Given a distribution $p(V)$ that district factorizes with respect to an ADMG $\mathcal{G}(V)$ where T is primal fixable, the IPW estimators ψ_{primal} and ψ_{dual} proposed in Lemmas 4 and 6 respectively, use variationally independent components of the observed distribution $p(V)$.*

In order to provide concrete intuition of the primal and dual IPW estimators, we now discuss its application to the ADMGs shown in Fig. 3. We revisit primal and dual fixability later in Section 6 where we define the primal operator in a recursive manner to obtain semiparametric estimating equations for our target parameter $\psi(t)$ in ADMGs where T is not immediately primal fixable.

3.2.1 Examples: primal and dual ipw estimators

Consider the ADMG in Fig. 3(a). The a-fixability criterion dictates that T is not fixable due to the presence of bidirected paths from T to its descendants, namely L and Y. However, T is primal fixable as there is no bidirected path from T to any of its children, namely M. The inverse Markov pillow of T in Fig. 3(a) is just M. Per Lemmas 4 and 6, the primal and dual IPW estimators for the target parameter $\psi(t)$ in Fig. 3(a) are given by,

$$\text{(Fig. 3a)} \qquad \psi_{\text{primal}} = \mathbb{E}\Big[\, \mathbb{I}(T = t) \times \frac{\sum_T \; p(T \mid C) \times p(L \mid T, M, C) \times p(Y \mid T, M, L, C)}{p(T \mid C) \times p(L \mid T, M, C) \times p(Y \mid T, M, L, C)} \times Y\Big],$$
$$\psi_{\text{dual}} = \mathbb{E}\Big[\frac{p(M \mid T = t, C)}{p(M \mid T, C)} \times Y\Big].$$

In order to estimate $\psi(t)$ using finite samples, we proceed as follows. In case of the primal IPW, we fit conditional densities $p(T \mid C), p(L \mid T, M, C)$, and $p(Y \mid T, M, L, C)$, either parametrically (using logistic regression for instance), or using flexible models like generalized additive models (GAMs) or nonparametric kernel regression methods. The target parameter is then obtained by empirically evaluating the outer expectation using the fitted

models. Note that we can also avoid modeling the conditional density of Y, as the outcome regression $\mathbb{E}[Y \mid T, M, L, C]$ suffices to estimate $\psi(t)$, i.e., ψ_{primal} can be expressed equivalently as

$$\mathbb{E}\Big[\, \mathbb{I}(T=t) \times \frac{\sum_T \; p(T \mid C) \times p(L \mid T, M, C)}{p(T \mid C) \times p(L \mid T, M, C)} \times \mathbb{E}[Y \mid T, M, L, C]\Big].$$

A simple procedure to estimate the dual IPW involves modeling the conditional density $p(M \mid T, C)$. However, a more sophisticated procedure may take advantage of modeling the density ratio directly as suggested by Sugiyama et al. (2010).

We now turn our attention to the ADMG in Fig. 3(b). The inverse Markov pillow of T in Fig. 3(b) is $\{M, Y\}$. The corresponding primal and dual IPW estimators are given by,

$$\text{(Fig. 3b)} \qquad \psi_{\text{primal}} = \mathbb{E}\Big[\mathbb{I}(T=t) \times \frac{\sum_T \; p(T \mid C) \times p(L \mid T, M, C)}{p(T \mid C) \times p(L \mid T, M, C)} \times Y\Big].$$

$$\psi_{\text{dual}} = \mathbb{E}\Big[\frac{p(M \mid T=t, C)}{p(M \mid T, C)} \times \frac{p(Y \mid T=t, M, L, C)}{p(Y \mid T, M, L, C)} \times Y\Big].$$

Similar strategies can be used to estimate $\psi(t)$ as in the previous example. Also, note that the conditional density of Y in ψ_{dual} can be replaced by the outcome regression $\mathbb{E}[Y \mid T = t, M, L, C]$, i.e., ψ_{dual} can be expressed equivalently as

$$E\Big[\frac{p(M \mid T=t, C)}{p(M \mid T, C)} \times \mathbb{E}[Y \mid T=t, M, L, C]\Big].$$

3.3 Augmented Primal IPW

In the previous section we have shown the existence of two estimators for the target $\psi(t)$ that use variationally independent portions of the likelihood when T is p-fixable. The question naturally arises if it is possible to combine these estimators to yield a single estimator that exhibits double robustness in the sets of models used in each one. We now show that the nonparametric influence function in this setting yields such an estimator.

Assume $p(V)$ factorizes with respect to an ADMG $\mathcal{G}(V)$ where T is primal fixable and for simplicity of exposition, assume that Y has no descendants in $\mathcal{G}$. The latter assumption is not necessary and our results extend trivially to the setting where this is not true; we use it only to avoid notational complexity and we prove in later sections that the efficient IF is not a function of descendants of Y. Recall from the previous section, that we use a fixed topological order τ where T is preceded by all its non-descendants and Y is succeeded by all its non-descendants. The set of nodes V can then be partitioned into three disjoint sets: $V = \{\mathbb{C}, \mathbb{L}, \mathbb{M}\}$, where

$$\begin{aligned}
\mathbb{C} &= \{C_i \in V \mid C_i \prec T\}, \\
\mathbb{L} &= \{L_i \in V \mid L_i \in D_T, L_i \succeq T\}, \\
\mathbb{M} &= \{M_i \in V \mid M_i \notin \mathbb{C} \cup \mathbb{L}\}.
\end{aligned} \tag{13}$$

Rearranging some of the terms in Eq. 10, $\psi(t)$ is identified as the following function of the observed data in terms of the sets defined above.

$$\psi(t) = \sum_{V \setminus T} Y \times \prod_{M_i \in \mathbb{M}} p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))\Big|_{T=t} \times \sum_{T} \prod_{L_i \in \mathbb{L}} p(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i)) \times p(\mathbb{C}). \tag{14}$$

We derive the corresponding influence function in the theorem below using the *pathwise derivative* (see Appendix D for details.) As estimators such as primal IPW that use the indicator function are inefficient in general, we view this influence function as augmenting the primal with terms (including dual IPW which does not use an indicator) that increase its efficiency. For readability, we use the form $\prod_{L_i \prec M_i}$ as shorthand for $\prod_{L_i \in \mathbb{L} | L_i \prec M_i}$.

Theorem 9 (Nonparametric influence function of augmented primal IPW)
Given a distribution $p(V)$ that district factorizes with respect to an ADMG $\mathcal{G}(V)$ where T is primal fixable, the nonparametric influence function for the target parameter $\psi(t)$ is as follows.

$$\begin{aligned}
U_{\psi_t} = \sum_{M_i \in \mathbb{M}} \Bigg\{ & \frac{\mathbb{I}(T=t)}{\prod_{L_i \prec M_i} p(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i))} \times \Bigg(\sum_{T \cup \{\succ M_i\}} Y \times \prod_{\substack{V_i \in \mathbb{L} \cup \\ \{\succ M_i\}}} p(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i))|_{T=t \text{ if } V_i \in \mathbb{M}} \\
& - \sum_{T \cup \{\succeq M_i\}} Y \times \prod_{\substack{V_i \in \mathbb{L} \cup \\ \{\succeq M_i\}}} p(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i)) \mid_{T=t \text{ if } V_i \in \mathbb{M}} \Bigg) \Bigg\} \\
+ \sum_{L_i \in \mathbb{L} \setminus T} \Bigg\{ & \frac{\prod_{M_i \prec L_i} p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))|_{T=t}}{\prod_{M_i \prec L_i} p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))} \times \Bigg(\sum_{\{\succ L_i\}} Y \times \prod_{V_i \succ L_i} p(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i))|_{T=t \text{ if } V_i \in \mathbb{M}} \\
& - \sum_{\{\succeq L_i\}} Y \times \prod_{V_i \succeq L_i} p(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i)) \mid_{T=t \text{ if } V_i \in \mathbb{M}} \Bigg) \Bigg\} \\
+ \sum_{V \setminus \{T, \mathbb{C}\}} & Y \times \prod_{M_i \in \mathbb{M}} p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))|_{T=t} \times \prod_{L_i \in \mathbb{L} \setminus T} p(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i)) - \psi(t),
\end{aligned} \tag{15}$$

where $\mathbb{C}, \mathbb{L}, \mathbb{M}$ are defined in display (13).

In the following lemma, we show that the influence function U_{ψ_t} in Theorem 9 uses information in the models for $M_i \in \mathbb{M}$ and $L_i \in \mathbb{L}$ in order to yield an estimator that is doubly robust in these sets.

Lemma 10 (Double robustness of augmented primal IPW)
The estimator obtained by solving the estimating equation $\mathbb{E}[U_{\psi_t}] = 0$, where U_{ψ_t} is given in Theorem 9, is consistent if all models in either $\{p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i)),\ \forall M_i \in \mathbb{M}\}$ or $\{p(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i)),\ \forall L_i \in \mathbb{L}\}$ are correctly specified.

According to Lemma 10, the estimator derived from the nonparametric IF is a *doubly robust* estimator. This allows us to perform consistent inferences for the target parameter $\psi(t)$ even in settings where a large part of the model likelihood is arbitrarily misspecified, provided that conditional models for variables in either $\mathbb{M}$ or $\mathbb{L}$ are specified correctly. In addition, double robustness implies that the bias of the estimator has a product form which

allows parametric ($\sqrt{n}$) convergence rates for $\psi(t)$ to be obtained even if flexible machine learning models with slower than parametric convergence rates are used to fit nuisance models. See (Chernozhukov et al., 2018) for details.

3.3.1 Cancellation of Terms in the IF

Given a post treatment variable V_i and its conditional density $p(V_i \mid \text{mp}_{\mathcal{G}}(V_i))$ in the identified functional for $\psi(t)$ in Eq. 14, there is a corresponding term in the influence function U_{ψ_t} in Theorem 9 of the form

$$f_1(\prec V_i) \times \Big(f_2(\preceq V_i) - \sum_{V_i} f_2(\preceq V_i) \times p\big(V_i \mid \text{mp}_{\mathcal{G}}(V_i)\big)\Big), \tag{16}$$

where $f_1(\prec V_i)$ denotes a function of variables that precedes V_i in the topological order. Similarly, $f_2(\preceq V_i)$ is a function of $\prec V_i$ and V_i itself.

Sometimes, these terms in the influence function U_{ψ_t} may cancel each other out. For instance, assume there are two consecutive variables $V_i, V_{i+1} \in \mathbb{L}$ (or $\in \mathbb{M}$) such that $\text{mp}_{\mathcal{G}}(V_{i+1}) \setminus V_i \subseteq \text{mp}_{\mathcal{G}}(V_i)$. The corresponding terms in the influence function share some common terms: First, the two share the same weight terms, i.e., $f_1(\prec V_{i+1}) = f_1(\prec V_i)$, and second $f_2(\preceq V_i) = \sum_{V_{i+1}} f_2(\preceq V_{i+1}) \times p\big(V_{i+1} \mid \text{mp}_{\mathcal{G}}(V_{i+1})\big)$. Therefore, by simply factoring out this weight term, we note that V_i and V_{i+1} can be viewed as contributing a single term to the influence function of the form shown in Eq. 16, and that is

$$f_1(\prec V_{i+1}) \times \Big(f_2(\preceq V_{i+1}) - \sum_{V_i, V_{i+1}} f_2(\preceq V_{i+1}) \times p\big(V_{i+1}, V_i \mid \text{mp}_{\mathcal{G}}(V_i)\big)\Big).$$

This cancellation occurs regardless of whether $V_i \in \mathbb{L}$ or $V_i \in \mathbb{M}$. However it is important that both V_i and V_{i+1} be in the same set (since they need a common f_1 term to be factored out.) Such cancellations may be applied recursively to consecutive variables in $\mathbb{L}$ or $\mathbb{M}$.

Another possible cancellation of terms may occur in the weights that correspond to "dual weights" in Eq. 15. The factors in the numerator and the denominator are exactly the same except for the fact that the numerator is evaluated at $T = t$. However, if there exists $M_i \in \mathbb{M}$ such that T is not in its Markov pillow, i.e., $M_i \perp\!\!\!\perp T \mid \text{mp}_{\mathcal{G}}(M_i)$, then $p(M_i \mid \text{mp}_{\mathcal{G}}(M_i))|_{T=t} / p(M_i \mid \text{mp}_{\mathcal{G}}(M_i)) = 1$. Note that such cancellations only involves variables in $\mathbb{M}$. In fact, the set of conditional densities that stay in these weight terms correspond to the variables in $\text{mp}_{\mathcal{G}}^{-1}(T)$, the inverse Markov pillow of T. Therefore, we have the following general simplification to the influence function U_{ψ_t} in Theorem 9,

$$\frac{\prod\limits_{M_i \prec L_i} p(M_i \mid \text{mp}_{\mathcal{G}}(M_i))\Big|_{T=t}}{\prod\limits_{M_i \prec L_i} p(M_i \mid \text{mp}_{\mathcal{G}}(M_i))} = \frac{\prod\limits_{M_i \in \text{mp}_{\mathcal{G}}^{-1}(T) \cap \{\prec L_i\}} p(M_i \mid \text{mp}_{\mathcal{G}}(M_i))\Big|_{T=t}}{\prod\limits_{M_i \in \text{mp}_{\mathcal{G}}^{-1}(T) \cap \{\prec L_i\}} p(M_i \mid \text{mp}_{\mathcal{G}}(M_i))}.$$

More intuition on the nonparametric IF is provided in Appendix D. An implication of the two aforementioned forms of cancellation is that the robustness statement in Lemma 10 is somewhat conservative. In other words, it may not be necessary to model all the conditional

terms mentioned in the doubly robust statement of Lemma 10. It is sometimes possible to prune vertices from the ADMG and still achieve a doubly robust estimator that requires fitting less models as demonstrated via an example in the next subsection.

The observations above prompt an alternative representation of the influence function mentioned in Theorem 9 that is expressed solely in terms of the primal and dual IPW statements in Lemmas 4 and 6. This alternative formulation of the influence function has practical implications for estimating the target parameter from finite samples and, more importantly, will lead to an elegant formulation of the efficient influence function that we derive in Section 5. The reformulation is as follows.

Theorem 11 (Reformulation of the IF for augmented primal IPW)
Under the same conditions stated in Theorem 9, the nonparametric influence function for augmented primal IPW can be re-expressed as follows.

$$\begin{aligned} U_{\psi_t} = &\sum_{M_i \in \mathbb{M}} \mathbb{E}[\beta_{primal} \mid \{\preceq M_i\}] - \mathbb{E}[\beta_{primal} \mid \{\prec M_i\}] \\ &+ \sum_{L_i \in \mathbb{L}} \mathbb{E}[\beta_{dual} \mid \{\preceq L_i\}] - \mathbb{E}[\beta_{dual} \mid \{\prec L_i\}] \\ &+ \mathbb{E}[\beta_{primal/dual} \mid \mathbb{C}] - \psi(t), \end{aligned}$$

where β_{primal} and β_{dual} are obtained as in Lemmas 4 and 6 respectively, and $\beta_{primal/dual}$ means that we may use either β_{primal} or β_{dual}.

According to the above lemma, the portion of the IF that relates to elements in $\mathbb{C}$, may be recovered using either β_{primal} or β_{dual}. That is, $\mathbb{E}[\beta_{\text{primal}} \mid \mathbb{C}] = \mathbb{E}[\beta_{\text{dual}} \mid \mathbb{C}]$ and $\mathbb{E}[\beta_{\text{primal}}] = \mathbb{E}[\beta_{\text{dual}}] = \psi(t)$. Reformulation of the IF offers the advantage of restricting the modeling of conditional densities to only those involved in β_{primal} and β_{dual}. The analyst may then rely on flexible regression methods in order to model each $\mathbb{E}[\cdot \mid \cdot]$ above.

3.3.2 Examples: Augmented Primal IPW

We now revisit the ADMGs in Fig. 3 and derive the corresponding nonparametric influence functions. Consider the ADMG in Fig. 3(a). The sets in display (13) are as follows, $\mathbb{C} = \{C\}, \mathbb{L} = \{T, L, Y\}$, and $\mathbb{M} = \{M\}$. Applying Theorem 9 to this graph, yields the influence function:

(Fig. 3a)
$$\begin{aligned} U_{\psi_t} = &\frac{\mathbb{I}(T=t)}{p(T \mid C)} \times \Big(\sum_{T,L} p(T \mid C) \times p(L \mid T, M, C) \times \mathbb{E}[Y \mid T, M, L, C] \\ &\qquad - \sum_{T,L,M} p(M \mid T=t, C) \times p(T \mid C) \times p(L \mid T, M, C) \times \mathbb{E}[Y \mid T, M, L, C] \Big) \\ &+ \frac{p(M \mid T=t, C)}{p(M \mid T, C)} \times \Big(Y - \mathbb{E}[Y \mid T, M, L, C] \Big) \\ &+ \frac{p(M \mid T=t, C)}{p(M \mid T, C)} \times \Big(\mathbb{E}[Y \mid T, M, L, C] - \sum_{L} p(L \mid T, M, C) \times \mathbb{E}[Y \mid T, M, L, C] \Big) \\ &+ \sum_{M,L} p(M \mid T=t, C) \times p(L \mid T, M, C) \times \mathbb{E}[Y \mid T, M, L, C] - \psi(t). \end{aligned}$$

Note that in the above influence function, the term $\frac{p(M|T=t,C)}{p(M|T,C)} \times \mathbb{E}[Y \mid T, M, L, C]$ appears twice with opposite signs. This is an example of the kind of cancellation mentioned in the previous section, where Y and L are consecutive elements in the set $\mathbb{L}$ that share essentially the same Markov pillow, i.e., $\text{mp}_{\mathcal{G}}(Y) \setminus L = \text{mp}_{\mathcal{G}}(L)$. In fact, this observation allows us to simplify the influence function even further by deriving the influence function in the ADMG $\mathcal{G}(V \setminus L)$ where L is treated as latent; projecting out L in this example, corresponds to removing all the edges into and out of L. This ADMG is simply the front-door graph with baseline confounding. Given Theorem 9, the IF is as follows.

$$\text{(Fig. 3a)} \quad U_{\psi_t} = \frac{\mathbb{I}(T=t)}{p(T \mid C)} \times \Big(\sum_{T} p(T \mid C) \times \mathbb{E}[Y \mid T, M, C] - \sum_{T,M} p(M \mid T=t, C) \times p(T \mid C) \times \mathbb{E}[Y \mid T, M, C] \Big) + \frac{p(M \mid T=t, C)}{p(M \mid T, C)} \times \Big(Y - \mathbb{E}[Y \mid T, M, C] \Big) + \sum_{M} p(M \mid T=t, C) \times \mathbb{E}[Y \mid T, M, C] - \psi(t).$$

While a general algorithm may be devised to perform such graphical simplifications, we leave this to future work as such a procedure must also ensure that the efficiency of the final estimator is not affected by these simplifications. We now derive the nonparametric IF for the ADMG in Fig. 3(b) where no simplification is possible. The sets in display (13) are $\mathbb{C} = \{C\}, \mathbb{L} = \{T, L\}$, and $\mathbb{M} = \{M, Y\}$. The nonparametric IF per Theorem 9 is,

$$\begin{aligned}
\text{(Fig. 3b)} \qquad U_{\psi_t} = {} & \frac{\mathbb{I}(T=t)}{p(T \mid C) \times p(L \mid T, M, C)} \times \Big(\sum_{T} p(T \mid C) \times p(L \mid T, M, C) \times Y \\
& - \sum_{T} p(T \mid C) \times p(L \mid T, M, C) \times \mathbb{E}[Y \mid T=t, M, L, C] \Big) \\
& + \frac{\mathbb{I}(T=t)}{p(T \mid C)} \times \Big(\sum_{T,L} p(T \mid C) \times p(L \mid T, M, C) \times \mathbb{E}[Y \mid T=t, M, L, C] \\
& - \sum_{T,L} p(T \mid C) \times p(M \mid T=t, C) \times p(L \mid T, M, C) \times \mathbb{E}[Y \mid T=t, M, L, C] \Big) \\
& + \frac{p(M \mid T=t, C)}{p(M \mid T, C)} \times \Big(\mathbb{E}[Y \mid T=t, M, L, C] - \sum_{L} p(L \mid T) \times \mathbb{E}[Y \mid T=t, M, L, C] \Big) \\
& + \sum_{M,L} p(M \mid T=t, C) \times p(L \mid T, M, C) \times \mathbb{E}[Y \mid T=t, M, L, C] - \psi(t). \qquad (17)
\end{aligned}$$

We briefly describe estimation strategies for estimators resulting from the nonparametric IF in Theorem 9 using the influence function in Eq. 17 as an example. An estimator for the target $\psi(t)$ is obtained by solving the estimating equation $\mathbb{E}[U_{\psi_t}] = 0$. In the resulting estimator, conditional densities for $p(T \mid C), p(M \mid T, C), p(L \mid T, M, C)$ and the outcome regression $\mathbb{E}[Y \mid T, M, L, C]$ can be fit either parametrically or using flexible models as described in Section 3.2.1. The outer expectation is then evaluated empirically using the fitted models in order to yield the target parameter. Per Lemma 10, the estimator for $\psi(t)$ is consistent as long as one of the sets $\{p(T \mid C), p(L \mid T, M, C)\}$ or $\{p(M \mid T, C), \mathbb{E}[Y \mid T, M, L, C]\}$ is correctly specified while allowing for arbitrary misspecification of the other.

Another estimation strategy that is computationally simpler stems from the usage of Theorem 11 to the ADMG in Fig. 3(b). The resulting estimator for the target is,

$$\text{(Fig. 3b)} \quad \psi_{\text{reform}} = \mathbb{E}\big[\, \mathbb{E}[\beta_{\text{primal}} \mid Y, T, M, L, C] - \mathbb{E}[\beta_{\text{primal}} \mid T, M, L, C] + \mathbb{E}[\beta_{\text{primal}} \mid M, T, C] - \mathbb{E}[\beta_{\text{primal}} \mid T, C] + \mathbb{E}[\beta_{\text{dual}} \mid L, T, M, C] - \mathbb{E}[\beta_{\text{dual}} \mid T, M, C] + \mathbb{E}[\beta_{\text{dual}} \mid T, C] \,\big]. \tag{18}$$

The above can be estimated from finite samples by first obtaining estimates for β_{primal} and β_{dual} for each sample as in Section 3.2.1 and then fitting flexible regressions for each $\mathbb{E}[\cdot \mid \cdot]$ shown in Eq. 18 using these estimates. The outer expectation is then evaluated empirically using these fitted models, yielding an estimate for the target parameter $\psi(t)$.

4. Efficient IFs in ADMG Models with Unrestricted Tangent Spaces

Influence functions provide a geometric view of the behavior of RAL estimators. Assume a statistical model $\mathcal{M} = \{p_\eta(V) : \eta \in \Gamma\}$ where Γ is the parameter space and η is the parameter indexing a specific model. Consider a Hilbert space[1] H of all mean-zero scalar functions, equipped with an inner product defined as $\mathbb{E}[h_1 \times h_2], h_1, h_2 \in \mathrm{H}$. The *tangent space* in the statistical model $\mathcal{M}$, denoted by Λ, is defined to be the mean-square closure of parametric submodel tangent spaces, where a parametric submodel tangent space is the set of elements $\Lambda_{\eta_\kappa} = \{\alpha \times S_{\eta_\kappa}(V)\}, \alpha$ is a constant and S_{η_κ} is the score for the parameter ψ_{η_κ} for some parametric submodel, $\mathcal{M}_{\text{sub}} = \{P_{\eta_\kappa} : \kappa \in [0, 1)$ where $P_{\eta_{\kappa=0}} = P_{\eta_0}\}$ of the model $\mathcal{M}$. In mathematical form, this is denoted as $\Lambda \equiv \overline{[\Lambda_{\eta_\kappa}]}$.

The tangent space Λ is a closed linear subspace of the Hilbert space H ($\Lambda \subseteq \mathrm{H}$). The orthogonal complement of the tangent space, denoted by $\Lambda^\perp$, is defined as $\Lambda^\perp = \{h \in \mathrm{H} \mid \mathbb{E}[h \times h'] = 0, \forall h' \in \Lambda\}$. Note that $\mathrm{H} = \Lambda \oplus \Lambda^\perp$, where $\oplus$ denotes the direct sum, and $\Lambda \cap \Lambda^\perp = \{0\}$. The vector space $\Lambda^\perp$ is of particular importance because we can construct the class of all influence functions, denoted by $\mathcal{U}$, as $\mathcal{U} = \{U_\psi + \Lambda^\perp\}$. In other words, upon knowing a single influence function U_ψ and $\Lambda^\perp$, we can obtain the class of all possible RAL estimators that admit the CAN property.

Out of all the influence functions in $\mathcal{U}$, there exists a unique one which lies in the tangent space Λ and yields the most efficient RAL estimator by recovering the *semiparametric efficiency bound*. This efficient influence function can be obtained by projecting any influence function, call it U^*_ψ, onto the tangent space Λ. This operation is denoted by $U^{\text{eff}}_\psi = \pi[U^*_\psi \mid \Lambda]$, where U^{eff}_ψ denotes the efficient influence function. For a more detailed description of the concepts outlined here, see Appendix D and (Tsiatis, 2007; Bickel et al., 1993).

4.1 Nonparametric Saturated Models

As we start to derive estimators that achieve the semiparametric efficiency bound for the class of causal graphical models described in Section 3, we begin with the simplest case of *nonparametric saturated* models where the statistical model is completely unrestricted. If

1. The Hilbert space of all mean-zero scalar functions is the L^2 space. For a precise definition of Hilbert spaces see Luenberger (1997).

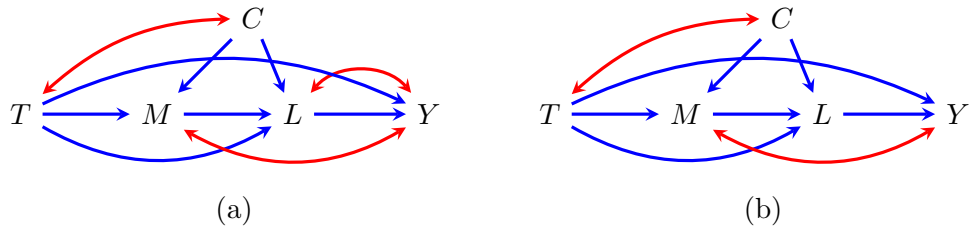

Figure 4: (a) Example of an ADMG whose underlying nested Markov model is NPS even though there is a missing edge between C and Y (b) Removing the bidirected edge between L and Y introduces a Verma constraint $C \perp\!\!\!\perp Y$ in $p(V)/p(L \mid T, M, C)$.

the tangent space contains the entire Hilbert space, i.e., $\Lambda = \mathrm{H}$, then the statistical model $\mathcal{M}$ is called a nonparametric saturated model (NPS).

Given a fixed topological order τ, the distribution of an NPS model can be expressed as the product of conditional densities, $p(V) = \prod_{i=1}^{k} p(V_i \mid \prec_\tau V_i)$, where k denotes the total number of variables. A well-known result states that the tangent space of an NPS model can be partitioned into a direct sum of orthogonal subspaces (Tsiatis, 2007),

$$\Lambda = \mathrm{H} = \Lambda_1 \oplus \ldots \oplus \Lambda_k, \qquad \textit{(Tangent space of an NPS model)} \tag{19}$$

where $\Lambda_i = \big\{\alpha_i(V_1, \ldots, V_i) \in \mathrm{H} \text{ s.t. } \mathbb{E}\big[\alpha_i \mid V_1, \ldots, V_{i-1}\big] = 0\big\}, i = 1, \ldots, k$. Alternatively, the linear space Λ_i can be defined as $\Lambda_i = \big\{\alpha_i(V_1, \ldots, V_i) - \mathbb{E}[\alpha_i \mid V_1, \ldots, V_{i-1}], \ \forall \alpha_i \in \mathrm{H}\big\}$. In addition, any element $h \in \mathrm{H}$ can be decomposed into orthogonal elements $h = h_1 + \cdots + h_k$, where h_i is the projection of h onto Λ_i, i.e., $h_i(.) = \pi[h(.) \mid \Lambda_i]$.

In an NPS model, there exists a single unique influence function and therefore, there is no need to explore the space of all influence functions in order to find the most efficient one. As a result, the estimator that we obtain by solving $\mathbb{E}[U_\psi] = 0$, where U_ψ is given by Theorems 2 or 9 as appropriate, is not only doubly robust but also the most efficient estimator. We formalize this in the following lemma.

Lemma 12 (Efficiency of IFs in Theorems 2 and 9)
Given a nonparametric saturated model that factorizes according to an ADMG $\mathcal{G}(V)$ where T is a-fixable (p-fixable), the influence function U_{ψ_t} provided in Theorem 2 (Theorem 9) is the efficient influence function for the target parameter $\psi(t)$.

An easy way to confirm the nonparametric saturation status of the model implied by an ADMG $\mathcal{G}(V)$ is to simply check that all vertices are pairwise connected by a directed or bidirected edge. However, the absence of edges between two vertices in an ADMG do not necessarily correspond to a conditional independence or even generalized conditional independence (Verma) constraint (Tian and Pearl, 2002b); see Fig. 4(a) for example where the missing edge between C and Y does not correspond to any constraint. The missing edge between C and Y in Fig. 4(b) on the other hand, does correspond to the Verma constraint $C \perp\!\!\!\perp Y$ in the distribution $p(V)/p(L \mid T, M, C)$; see (Verma and Pearl, 1990; Shpitser et al., 2014) for more details.

Algorithm 1 Check Nonparametric Saturation ($\mathcal{G}$)

Define $\mathrm{pa}^{\mathrm{d}}_{\mathcal{G}}(S) \equiv \bigcup_{S_i \in S} \mathrm{pa}_{\mathcal{G}}(S_i)$
for all V_i, V_j pairs in $\mathcal{G}$ such that $i \neq j$ **do**
    if not $\{$ $V_i \in \mathrm{pa}^{\mathrm{d}}_{\mathcal{G}}(\langle V_j \rangle_{\mathcal{G}})$ **or** $V_j \in \mathrm{pa}^{\mathrm{d}}_{\mathcal{G}}(\langle V_i \rangle_{\mathcal{G}})$ **or**
        $\langle V_i, V_j \rangle_{\mathcal{G}}$ is bidirected connected in $\mathcal{G}$ $\}$ **then**
        return Not NPS
return NPS

We now propose a procedure to check if the model implied by an ADMG $\mathcal{G}(V)$ with missing edges is nonparametric saturated by checking if it is equivalent to or can be rephrased as the model implied by another ADMG where there are no missing edges.

4.2 Algorithm to Detect Nonparametric Saturation

We fist introduce the necessary graphical preliminaries in order to describe our procedure. *Conditional ADMGs* (CADMGs) $\mathcal{G}(V, W)$ are acyclic directed mixed graphs whose vertices can be partitioned into random variables V and *fixed* variables W with the restriction that only outgoing edges may be adjacent to variables in W (Richardson et al., 2017). Such graphs are particularly useful in representing post-intervention distributions where variables in W have been intervened on or *fixed*. The usual definitions of genealogic sets such as parents, descendants, and ancestors as well as other special sets such as the Markov blanket and Markov pillow of random variables $V_i \in V$ in a CADMG $\mathcal{G}(V, W)$, extend naturally by allowing for the inclusion of fixed variables into these sets when appropriate.

As implied by the name and as alluded to earlier, the criterion of a-fixability is closely related to the fixing operator defined by Richardson et al. (2017), which can be formalized as follows. A vertex $V_i \in V$ is said to be fixable in a CADMG $\mathcal{G}(V, W)$ if $\mathrm{dis}_{\mathcal{G}}(V_i) \cap \mathrm{de}_{\mathcal{G}}(V_i) = \{V_i\}$. The graphical operation of fixing V_i denoted by $\phi_{V_i}(\mathcal{G})$, yields a new CADMG $\mathcal{G}(V \setminus V_i, W \cup V_i)$ where bidirected and directed edges into V_i are removed and V_i is fixed to a particular value v_i. The notion of fixability can then be extended to a set of vertices S by requiring that there exists an ordering $(S_1, \ldots, S_p)$ such that S_1 is fixable in $\mathcal{G}$, S_2 is fixable in $\phi_{S_1}(\mathcal{G})$ and so on. The *reachable closure* of a set of vertices S, denoted by $\langle S \rangle_{\mathcal{G}}$ is the unique minimal superset of S such that $V \setminus \langle S \rangle_{\mathcal{G}}$ is fixable (Shpitser et al., 2018). If the reachable closure of a set evaluates to the set itself, then the set is said to be *reachable*. That is, a set $S \subseteq V$ is reachable if $\langle S \rangle_{\mathcal{G}} = S$.

The concept of the reachable closure, plays a key role in constructing the *maximal arid projection* of an ADMG as described in (Shpitser et al., 2018). Such projections yield another ADMG that implies the same set of equality restrictions as the original, albeit one in which the absence of edges facilitates easier study of the constraints in the model. The algorithm that we now describe closely relates to this projection in that it declares a model to be NPS when the input ADMG's maximal arid projection is a complete graph, and not NPS otherwise. For more details, see proof of Theorem 13 and (Shpitser et al., 2018).

We provide our procedure for checking if a model is nonparametrically saturated in Algorithm 1. We show that our algorithm is *sound* and *complete* for this purpose in the following theorem. Further, an informal complexity analysis of Algorithm 1 shows that it is computationally tractable as it runs in polynomial time with respect to the number

of vertices and edges in the graph $\mathcal{G}$. The complexity of the outer loop is $\mathcal{O}(|V|^2)$ as it requires the selection of all possible pairs of random vertices. Further, naive implementations for computing reachable closures of sets are $\mathcal{O}(|V|^2 + |V| \times |E|)$ as it involves repeated applications of depth first search (popular algorithms for which are linear in complexity $\mathcal{O}(|V| + |E|)$ (Tarjan, 1972)) in order to determine the fixability of a set of vertices.

Theorem 13 (Soundness and completeness of Algorithm 1)
Algorithm 1 is sound and complete for deciding the nonparametric saturation status of the model implied by an ADMG $\mathcal{G}(V)$ by determining the absence of equality constraints.

4.2.1 Example: Nonparametric Saturation

As an example of the application of Algorithm 1, we return to the ADMGs in Fig. 4. As all pairs of vertices besides C and Y are adjacent in these ADMGs, the negation of the condition in lines 3 and 4 trivially evaluates to False for these pairs. However, in the case of Fig. 4(a), the algorithm finds that C is indeed a parent of the reachable closure of Y i.e., $C \in \text{pa}_{\mathcal{G}}(\langle Y \rangle_{\mathcal{G}})$ and hence, completes execution by confirming that the model is NPS. In the case of Fig. 4(b), this statement is no longer true, and indeed neither is $Y \in \text{pa}_{\mathcal{G}}(\langle C \rangle_{\mathcal{G}})$ nor is $\langle C, Y \rangle_{\mathcal{G}}$ bidirected connected. Thus, with all these conditions evaluating to False, the negation is True, resulting in the algorithm correctly identifying Fig. 4(b) as not NPS.

5. Semiparametric Efficiency Bound in a Special Class of ADMGs

As seen in the previous section, there exists a class of NPS models represented via ADMGs that imposes no restriction on any distribution in the model. On the other hand, constraints in a semiparametric model shrink the tangent space Λ, and thus expand its orthogonal complement $\Lambda^{\perp}$. As $\Lambda^{\perp}$ expands, we will have more than one influence function (note that the class of all influence functions is $\{U_{\psi} + \Lambda^{\perp}\}$.) In this section, we are interested in constructing the set of all RAL estimators along with the most efficient one for our target parameter $\psi(t) = \mathbb{E}[Y(t)]$.

Constraints in a semiparametric model can take on many forms. We restrict our attention exclusively to constraints that are encoded by a given ADMG. An ADMG encodes two types of restrictions: regular conditional statements $A \perp\!\!\!\perp B \mid C$, and more general equality constraints such as conditional independence statements in post-intervention distributions known as *Verma constraints* (Verma and Pearl, 1990). In this paper, we focus on ordinary conditional independence statements.

Assume a class of ADMGs, where given a topological order τ, all constraints implied by the ADMG $\mathcal{G}(V)$ can be written as ordinary conditional independence statements of the form,

$$V_i \perp\!\!\!\perp \{\prec_{\tau} V_i\} \setminus \text{mp}_{\mathcal{G}}(V_i) \mid \text{mp}_{\mathcal{G}}(V_i) \tag{20}$$

Such a property immediately implies that the topological factorization of the observed data distribution $p(V)$ shown in Eq. 8 captures all constraints implied by the ADMG $\mathcal{G}(V)$. A sound criterion for identifying ADMGs that satisfy this property is to check that an edge between two vertices V_i and V_j in $\mathcal{G}$ is absent only if $V_i \notin \text{mb}_{\mathcal{G}}(V_j)$ *and* $V_j \notin \text{mb}_{\mathcal{G}}(V_i)$. We

call this class of ADMGs *mb-shielded ADMGs*, as pairs of vertices are always adjacent if either one is in the Markov blanket of the other.

Lemma 14 (mb-shielded ADMGs)
Consider a distribution $p(V)$ that district factorizes with respect to an ADMG $\mathcal{G}(V)$ where an edge between two vertices is absent only if $V_i \notin \mathrm{mb}_{\mathcal{G}}(V_j)$ and $V_j \notin \mathrm{mb}_{\mathcal{G}}(V_i)$. Then, given any valid topological order on V, all conditional independence constraints in $p(V)$ are implied by the set of restrictions: $V_i \perp\!\!\!\perp \{\prec V_i\} \setminus \mathrm{mp}_{\mathcal{G}}(V_i) \mid \mathrm{mp}_{\mathcal{G}}(V_i), \forall V_i \in V$.

In what follows, we first construct the class of all RAL estimators in mb-shielded ADMGs by constructing the tangent space and its orthogonal complement. We then obtain the unique element in the tangent space that corresponds to the most efficient IF. This leads to construction of an estimator that attains the semiparametric efficiency bound, and hence is the estimator with the lowest variance in the class of all RAL estimators.

5.1 The Class of All RAL Estimators in Models of mb-shielded ADMGs

As we discussed earlier, the class of all IFs is $\mathcal{U} = \{U_{\psi_t} + \Lambda^{\perp}\}$. Consequently, the class of alll RAL estimators is characterized via the orthogonal complement of the tangent space. The said space is defined as $\Lambda^{\perp} = \{h(V) - \pi[h(V) \mid \Lambda], h \in \mathrm{H}\}$. This definition provides a guideline for deriving $\Lambda^{\perp}$: first derive Λ and then project elements of the Hilbert space onto Λ. The residual then belongs to $\Lambda^{\perp}$. We start by constructing the tangent space of mb-shielded ADMGs in the following theorem. We denote this tangent space by Λ^* to distinguish it from the tangent space of NPS models Λ, given in Eq. 19.

Theorem 15 (Tangent space Λ^* of mb-shielded ADMGs)
Given an mb-shielded ADMG $\mathcal{G}(V)$ with k number of vertices, the tangent space Λ^* is given by $\Lambda^* = \oplus_{i=1}^{k} \Lambda_i^*$, where

$$\begin{aligned}\Lambda_i^* &= \big\{\alpha_i(V_i, \mathrm{mp}_{\mathcal{G}}(V_i)) \in H \text{ s.t. } \mathbb{E}[\alpha_i \mid \mathrm{mp}_{\mathcal{G}}(V_i)] = 0\big\} \\ &= \big\{\alpha_i(V_i, \mathrm{mp}_{\mathcal{G}}(V_i)) - \mathbb{E}[\alpha_i \mid \mathrm{mp}_{\mathcal{G}}(V_i)],\ \forall \alpha_i(V_i, \mathrm{mp}_{\mathcal{G}}(V_i)) \in H\big\}.\end{aligned}$$

and $\Lambda_i^*, i = 1, \ldots, k$ are mutually orthogonal spaces. In addition, the projection of an element $h(V) \in H$, onto Λ_i^*, denoted by h_i, is given as follows,

$$h_i \equiv \Pi[h(V) \mid \Lambda_i^*] = \mathbb{E}\big[h(V) \mid V_i, \mathrm{mp}_{\mathcal{G}}(V_i)\big] - \mathbb{E}\big[h(V) \mid \mathrm{mp}_{\mathcal{G}}(V_i)\big].$$

The projection in Theorem 15 enables construction of the orthogonal complement of the tangent space denoted by $\Lambda^{*\perp}$.

Theorem 16 (Orthogonal complement $\Lambda^{*\perp}$ in mb-shielded ADMGs)
Given an mb-shielded ADMG $\mathcal{G}(V)$ with k number of vertices, the orthogonal complement of the tangent space $\Lambda^{*\perp}$ is given as

$$\Lambda^{*\perp} = \Big\{ \sum_{V_i \in V} \alpha_i(V_1, \ldots, V_i) - \mathbb{E}\Big[\alpha_i(V_1, \ldots, V_i) \mid V_i, \mathrm{mp}_{\mathcal{G}}(V_i)\Big]\Big\},$$

where $\alpha_i(V_1, \ldots, V_i)$ is any function of V_1 through V_i in H, such that $\mathbb{E}[\alpha_i \mid V_1, \ldots, V_{i-1}] = 0$. In other words, $\alpha_i(V_1, \ldots, V_i) \in \Lambda_i$.

According to Theorem 16, the class of all RAL estimators in mb-shielded ADMGs is $\{U_{\psi_t} + \Lambda^{*\perp}\}$, where U_{ψ_t} is given in Theorem 2 if T is a-fixable and Theorem 9 if T is p-fixable. The efficient influence function can then be found by projecting U_{ψ_t} onto the tangent space Λ^*. We now derive the most efficient RAL estimator for two different classes of mb-shielded ADMGs: one where T is a-fixable and one where T is p-fixable.

5.2 Efficient IF in mb-shielded ADMGs Where T is Adjustment Fixable

The efficient influence function for any mb-shielded ADMG $\mathcal{G}(V)$ where T is a-fixable such as the one shown in Fig 2, is given by projection of the gAIPW influence function given in Theorem 2. We formalize this in the following theorem. The conditional independences below rely on a slight abuse of notation where $A \perp\!\!\!\perp B \mid C$ when $B \cap C \neq \emptyset$ is taken to mean $A \perp\!\!\!\perp B \setminus \{B \cap C\} \mid C$.

Theorem 17 (Efficient gAIPW for mb-shielded ADMGs)
Given a distribution $p(V)$ that district factorizes with respect to an mb-shielded ADMG $\mathcal{G}(V)$ where T is adjustment fixable, the efficient influence function for the target parameter $\psi(t)$ is given as follows,

$$U^{eff}_{\psi_t} = \sum_{V_i \in V^*} \mathbb{E}\Big[\frac{\mathbb{I}(T=t)}{p(T \mid \mathrm{mp}_{\mathcal{G}}(T))} \times Y \;\Big|\; V_i, \mathrm{mp}_{\mathcal{G}}(V_i)\Big] - \mathbb{E}\Big[\frac{\mathbb{I}(T=t)}{p(T \mid \mathrm{mp}_{\mathcal{G}}(T))} \times Y \;\Big|\; \mathrm{mp}_{\mathcal{G}}(V_i)\Big],$$

where $V^* = V \setminus (T \cup Z \cup D)$ and

$$Z = \{Z_i \in V \mid Z_i \perp\!\!\!\perp Y \mid \mathrm{mp}_{\mathcal{G}}(Z_i) \text{ in } \mathcal{G}_{V \setminus T} \text{ and } Z_i \not\perp\!\!\!\perp T \mid \mathrm{mp}_{\mathcal{G}}(Z_i)\},$$
$$D = \{D_i \in V \mid D_i \perp\!\!\!\perp T, \mathrm{mp}_{\mathcal{G}}(T), Y \mid \mathrm{mp}_{\mathcal{G}}(D_i)\}.$$

Several interesting facts follow from the form of the efficient influence function shown in Theorem 17. First, the efficient influence function can be obtained by simply projecting the IPW portion of the gAIPW influence function given in Theorem 2. Second, the set $V \setminus V^*$ enumerates several vertices that do not affect the efficiency of estimating the target parameter $\psi(t)$. These include vertices Z_i that meet the criteria for a *conditional instrumental variable* (conditional on their Markov pillow) as defined in (Pearl, 2009; van der Zander et al., 2015). Further, no efficiency is lost by disregarding other vertices D_i that include descendants of Y, and irrelevant non-descendants of Y. It is also easy to see that Theorem 17 generalizes the efficient influence function put forward for distributions Markov relative to a DAG in (Rotnitzky and Smucler, 2019) to settings that account for unmeasured or hidden variables, and is exactly equivalent when the input to Theorem 17 is a DAG.

5.2.1 Example: Efficient gAIPW

Returning to Fig. 2, fix the topological order $\tau = \{C_1, C_2, Z_1, Z_2, T, M, Y, D_1, D_2\}$. One can check that the vertices labeled Z_1 and Z_2 meet the criteria for conditional instruments and

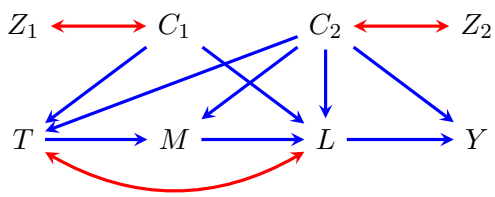

Figure 5: An mb-shielded ADMG that is not NPW and where T is primal fixable.

the vertices D_1 and D_2 meet the criteria of $D_i \perp\!\!\!\perp T, \mathrm{mp}_{\mathcal{G}}(T), Y \mid \mathrm{mp}_{\mathcal{G}}(D_i)$ and thus do not appear in the terms of the efficient influence function given in Theorem 17. Note that when T is a-fixable, it is always true that $\beta_{\text{primal}} = \frac{\mathbb{I}(T=t)}{p(T|\mathrm{mp}_{\mathcal{G}}(T))} \times Y$. Consequently, by applying Theorem 17, we obtain the following efficient estimator for the target $\psi(t)$.

$$\text{(Fig. 2)} \qquad \psi_{\text{eff}} = \mathbb{E}\Big[\, \mathbb{E}[\beta_{\text{primal}} \mid Y, M, C_2 - \mathbb{E}[\beta_{\text{primal}} \mid M, C_2] + \mathbb{E}[\beta_{\text{primal}} \mid M, T, C_1, C_2] - \mathbb{E}[\beta_{\text{primal}} \mid T, C_1, C_2] + \mathbb{E}[\beta_{\text{primal}} \mid C_2] + \mathbb{E}[\beta_{\text{primal}} \mid C_1] - \mathbb{E}[\beta_{\text{primal}}] \,\Big], \tag{21}$$

where $\beta_{\text{primal}} = \frac{\mathbb{I}(T=t)}{p(T|Z_1,Z_2,C_1,C_2)} \times Y$. The above can be estimated by following a similar strategy discussed for the functional in Eq. 18.

5.3 Efficient IF in mb-shielded ADMGs Where T is Primal Fixable

Consider the ADMG shown in Fig. 5. Such an ADMG may reflect additional background knowledge or conditional independences known to the analyst. For example, in Fig. 5, $C_1 \perp\!\!\!\perp C_2$ and $M \perp\!\!\!\perp C_1, Z_1, Z_2 \mid T, C_2$. As this model is no longer NPS, the IF obtained via Theorem 9 is not the most efficient IF. However, it is easy to see that this ADMG is mb-shielded and therefore the efficient IF is given by projection of U_{ψ_t} in Theorem 9 onto the tangent space in Theorem 15. In the following theorem, we provide the general form of the efficient IF in an arbitrary mb-shielded ADMG where T is p-fixable.

Theorem 18 (Efficient augmented primal IPW for mb-shielded ADMGs) *Given a distribution $p(V)$ that district factorizes with respect to an mb-shielded ADMG $\mathcal{G}(V)$ where T is primal fixable, the efficient influence function for the target parameter $\psi(t)$ is given as follows,*

$$\begin{aligned} U^{eff}_{\psi_t} = &\sum_{M_i \in \mathbb{M}} \mathbb{E}[\beta_{primal} \mid M_i, \mathrm{mp}_{\mathcal{G}}(M_i)] - \mathbb{E}[\beta_{primal} \mid \mathrm{mp}_{\mathcal{G}}(M_i)] \\ &+ \sum_{L_i \in \mathbb{L}} \mathbb{E}[\beta_{dual} \mid L_i, \mathrm{mp}_{\mathcal{G}}(L_i)] - \mathbb{E}[\beta_{dual} \mid \mathrm{mp}_{\mathcal{G}}(L_i)] \\ &+ \sum_{C_i \in \mathbb{C}} \mathbb{E}[\beta_{primal/dual} \mid C_i, \mathrm{mp}_{\mathcal{G}}(C_i)] - \mathbb{E}[\beta_{primal/dual} \mid \mathrm{mp}_{\mathcal{G}}(C_i)] \end{aligned} \tag{22}$$

where $\mathbb{C}, \mathbb{L}, \mathbb{M}$ are defined in display (13), and β_{primal} and β_{dual} are obtained as in Lemmas 4 and 6 respectively. $\beta_{primal/dual}$ means that we can either use β_{primal} or β_{dual} for the terms in $\mathbb{C}$.

Hence, the primal and dual IPWs comprise the fundamental elements of the efficient influence function in the setting where T is primal fixable. Further, when T is adjustment fixable as in the previous section, $\mathbb{L}$ comprises of just T as all elements in the district of T are pre-treatment covariates. Since the enumeration over vertices excludes T in the efficient influence function of gAIPW, the corresponding term in Eq. 22 should also be ignored. Thus, the form of the efficient influence function in Eq. 22 directly yields the efficient influence function of gAIPW shown in Theorem 17 as a special case where ψ_{dual} is not used.

5.3.1 Example: Efficient APIPW

Applying Theorem 18 to Fig. 5 gives us the following efficient estimator. Fix a valid topological order $(C_1, C_2, Z_1, Z_2, T, M, L, Y)$. Then

$$\text{(Fig. 5)} \qquad \beta_{\text{primal}} = \frac{\sum_T p(T \mid C_1, C_2) \times p(L \mid T, M, C_1, C_2)}{p(T \mid C_1, C_2) \times p(L \mid T, M, C_1, C_2)} \times Y$$
$$\beta_{\text{dual}} = \frac{p(M \mid T = t, C_2)}{p(M \mid T, C_2)} \times Y. \tag{23}$$

Define the sets $\mathbb{M} = \{M, Y\}$, $\mathbb{L} = \{T, L\}$, and $\mathbb{C} = \{C_1, C_2\}$. Note that we have dropped terms involving the vertices Z_1 and Z_2 as it is easy to check that $\mathbb{E}[\beta_{\text{dual}} \mid Z_i, \text{mp}_{\mathcal{G}}(Z_i)] = \mathbb{E}[\beta_{\text{dual}} \mid \text{mp}_{\mathcal{G}}(Z_i)]$, resulting in a cancellation of these terms. Then

$$\begin{aligned}\text{(Fig. 5)} \qquad \psi_{\text{eff}} = \mathbb{E}\big[&\mathbb{E}[\beta_{\text{primal}} \mid Y, L, C_2] - \mathbb{E}[\beta_{\text{primal}} \mid L, C_2] \\ &+ \mathbb{E}[\beta_{\text{primal}} \mid M, T, C_2] - \mathbb{E}[\beta_{\text{primal}} \mid T, C_2] \\ &+ \mathbb{E}[\beta_{\text{dual}} \mid L, M, T, C_1, C_2] - \mathbb{E}[\beta_{\text{dual}} \mid M, T, C_1, C_2] \\ &+ \mathbb{E}[\beta_{\text{dual}} \mid T, C_1, C_2] - \mathbb{E}[\beta_{\text{dual}} \mid C_1, C_2] \\ &+ \mathbb{E}[\beta_{\text{dual}} \mid C_2] + \mathbb{E}[\beta_{\text{dual}} \mid C_1] - \mathbb{E}[\beta_{\text{dual}}]\big]\end{aligned} \tag{24}$$

The estimation strategy for the above functional is very similar to the one used for Eq. 18.

6. Estimation of Arbitrary Identified Target Parameters

Thus far we have discussed targets of inference $\psi(t)$ that are identified by the adjustment functional Eq. 4, due to a-fixability of T, or by a more complex functional given by Eq. 10, due to primal fixability of T. In arbitrary hidden variable causal models, $\psi(t)$ may be identified even if T is neither a-fixable, nor p-fixable. All such identified $\psi(t)$ are given by an expectation taken with respect to a density given by a truncated factorization of the *nested Markov model*. In this section, we introduce this model, and review the general form of all $\psi(t)$ identified in hidden variable causal models. Finally, for those identifiable $\psi(t)$ where T is neither a-fixable, nor p-fixable, we give a consistent semiparametric estimator, and show that this estimator exhibits a partial double robustness property.

6.1 Identification via Truncated Nested Markov Factorization

The definition of the nested Markov model of an ADMG $\mathcal{G}(V)$ relies on the notion of fixability and fixing that we already used in Section 4 in the description of Algorithm 1. Recall, given a CADMG $\mathcal{G}(V, W)$, a vertex V_i is said to be (adjustment) fixable if $\text{dis}_{\mathcal{G}}(V_i) \cap \text{de}_{\mathcal{G}}(V_i) = \{V_i\}$. Also recall, the graphical operation of fixing V_i denoted by $\phi_{V_i}(\mathcal{G})$, which yields a new CADMG $\mathcal{G}(V \setminus V_i, W \cup V_i)$ where bidirected and directed edges into V_i are removed and V_i is fixed to a particular value v_i. Given a kernel $q_V(V \mid W)$ associated with the CADMG $\mathcal{G}(V, W)$, the corresponding *probabilistic* operation of fixing, denoted by $\phi_{V_i}(q_V; \mathcal{G})$, yields a new kernel

$$\phi_{V_i}(q_V; \mathcal{G}) \equiv q_{V \setminus V_i}(V \setminus V_i \mid W \cup V_i) \equiv \frac{q_V(V \mid W)}{q_V(V_i \mid \text{mb}_{\mathcal{G}}(V_i), W)}. \quad \textit{(Probabilistic fixing operator)} \tag{25}$$

Recall, a set $S \subseteq V$ is said to be fixable in $\mathcal{G}(V, W)$ if there exists a sequence $\sigma_S \equiv (S_1, S_2, \ldots, S_p)$ of elements in S, such that S_1 is fixable in $\mathcal{G}(V, W)$, S_2 is fixable in $\phi_{S_1}(\mathcal{G}(V, W))$, and so on. It is known that any two fixable sequences on S yield the same CADMG, which we will denote by $\phi_S(\mathcal{G}(V, W))$. Fix a CADMG $\mathcal{G}(V, W)$ and a corresponding kernel $q(V \mid W)$. Given a valid fixing sequence σ_S on $S \subseteq V$ valid in $\mathcal{G}(V, W)$, define $\phi_{\sigma_S}(q_V; \mathcal{G})$ inductively to be $q(V \mid W)$ when S is empty, and $\phi_{\sigma_S \setminus S_1}(\phi_{S_1}(q_V; \mathcal{G}); \phi_{S_1}(\mathcal{G}))$ otherwise, where $\sigma_S \setminus S_1$ corresponds to the remainder of the sequence after S_1.

The nested Markov factorization of an ADMG $\mathcal{G}$ relies on the notion of *intrinsic* sets. A set $S \subseteq V$ is said to be intrinsic in $\mathcal{G}$ if $V \setminus S$ is fixable, and $\phi_{V \setminus S}(\mathcal{G})$ contains a single district. The set of intrinsic sets of $\mathcal{G}$ is denoted by $\mathcal{I}(\mathcal{G})$. A distribution $p(V)$ is then said to obey the nested Markov factorization relative to an ADMG $\mathcal{G}(V)$ if for every fixable set S and every valid fixing sequence σ_S,

$$\phi_{\sigma_S}(p(V); \mathcal{G}) = \prod_{D \in \mathcal{D}(\phi_S(\mathcal{G}))} q_D(D \mid \text{pa}_{\mathcal{G}}(D)), \qquad \textit{(Nested Markov factorization)}$$

where all kernels appearing in the product above can be constructed from the set of kernels corresponding to intrinsic sets in $\mathcal{G}$, i.e., $\{q_S(S \mid \text{pa}_{\mathcal{G}}(S)) \mid S \in \mathcal{I}(\mathcal{G})\}$.

It has been shown that if $p(V)$ obeys such a factorization, then for any fixable set S, applying any two distinct valid sequences σ_S^1, σ_S^2 to $p(V)$ and $\mathcal{G}(V)$ also yields the same kernel, which we define as $\phi_S(p(V); \mathcal{G}(V))$. Moreover, for every $D \in \mathcal{I}(\mathcal{G})$, $q_D(D \mid \text{pa}_{\mathcal{G}}(D)) = \phi_{V \setminus D}(p(V); \mathcal{G}(V))$. The nested Markov factorization above defines the *nested Markov model*, with associated Markov properties, described in (Richardson et al., 2017).

An important result from Richardson et al. (2017) states that for any distribution $p(V \cup H)$ that is Markov relative to a hidden variable DAG $\mathcal{G}(V \cup H)$, the observed margin $p(V)$ is nested Markov relative to the latent projection ADMG $\mathcal{G}(V)$ derived from $\mathcal{G}(V \cup H)$. Moreover, identification of the target parameter $\psi(t)$ in a hidden variable causal model associated with $\mathcal{G}(V \cup H)$ may be rephrased, without loss of generality, using $\mathcal{G}(V)$. Specifically, for $Y^* \equiv \text{an}_{\mathcal{G}_{V \setminus T}}(Y)$,

$$\psi(t) = \sum_{Y^*} Y \times \prod_{D \in \mathcal{D}(\mathcal{G}_{Y^*})} \phi_{V \setminus D}(p(V); \mathcal{G}(V)) \Big|_{T=t}, \qquad \textit{(Truncated nested Markov factorization)} \tag{26}$$

provided every $D \in \mathcal{D}(\mathcal{G}_{Y^*})$ is intrinsic. Otherwise $\psi(t)$ is not identifiable.

In special cases described in (Shpitser et al., 2018; Evans and Richardson, 2019), a parametric likelihood can be specified for the nested Markov model, which leads naturally to estimation of $\psi(t)$ in Eq. 26 by the plug-in principle. However, in applications, assuming a full parametric likelihood is unrealistic. In the following section, we describe a semiparametric estimator which does not require a full parametric likelihood. Specifically, we will define the primal and dual fixing operators that can be recursively applied in order to arrive at a CADMG and corresponding kernel where we obtain an estimator for the target which exhibits a partial double robustness property.

6.2 Primal and Dual Fixing Operators

We now introduce the primal and dual fixing operators, the forms of which should look quite familiar to readers as they represent generalizations of the primal and dual formulations introduced in Section 3. In particular, the primal and dual operators can be applied to any CADMG $\mathcal{G}(V, W)$ and kernel $q(V \mid W)$ that is nested Markov with respect to $\mathcal{G}(V, W)$. This allows us to reason about and obtain semiparametric estimators in post-intervention (counterfactual) distributions where some variables have already been intervened on or fixed.

We first recap the criterion of primal fixability and then define the graphical operations of primal fixing (p-fixing) and dual fixing (d-fixing) a vertex V_i in a CADMG $\mathcal{G}(V, W)$. A vertex V_i is said to be p-fixable in $\mathcal{G}(V, W)$ if $\mathrm{dis}_{\mathcal{G}}(V_i) \cap \mathrm{ch}_{\mathcal{G}}(V_i) = \emptyset$. The graphical operation of primal fixing V_i in $\mathcal{G}(V, W)$, denoted by $\Phi_{V_i}(\mathcal{G})$, operates in exactly the same way as ordinary fixing, i.e., $\Phi_{V_i}(\mathcal{G})$ yields a new CADMG $\mathcal{G}(V \setminus V_i, W \cup V_i)$ where all incoming edges into V_i are removed and V_i is fixed to some value v_i. We set the criterion of dual fixability and graphical operation of d-fixing a vertex V_i, denoted by $\Delta_{V_i}(\mathcal{G})$, to be exactly the same as p-fixability and p-fixing.

We now describe the probabilistic *nested* or *primal fixing operator* as a generalization of the fixing operator, and the criterion put forward by Tian and Pearl (2002a) in the context of the identifiability of a post-intervention distribution of a single treatment on all other variables in the graph.

Lemma 19 (Primal fixing operator)
Given a distribution $q_V(V \mid W)$ that district factorizes with respect to a CADMG $\mathcal{G}(V, W)$, T is said to be primal fixable if it has no bidirected path to any of its children, i.e., $\mathrm{dis}_{\mathcal{G}}(T) \cap \mathrm{ch}_{\mathcal{G}}(T) = \emptyset$. *Let D_T be the district of T. The probabilistic operation of primal fixing denoted by $\Phi_T(q_V; \mathcal{G})$ is given by*

$$\begin{aligned}\Phi_T(q_V; \mathcal{G}) \equiv q_{V \setminus T}(V \setminus T \mid W \cup T) &\equiv \frac{q_V(V \mid W)}{q_{D_T}(T \mid \mathrm{mb}_{\mathcal{G}}(T), W)} \\ &= q_V(V \mid W) \times \frac{\prod_{V_i \in D_T} q_V(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i), W)}{\sum_T \prod_{V_i \in D_T} q_V(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i), W)}. \end{aligned} \tag{27}$$

The resulting kernel district factorizes with respect to a CADMG $\mathcal{G}(V \setminus T, W \cup T)$ where all incoming edges into T are removed. We denote this graphical operation of primal fixing variable T in $\mathcal{G}$ as $\Phi_T(\mathcal{G})$.

The key difference between the ordinary fixing operator introduced at the beginning of this section and the primal fixing operator lies in the denominators of Eqs. 25 and 27. The ordinary fixing operator uses $q_V(T \mid \mathrm{mb}_{\mathcal{G}}(T), W)$ which can be interpreted as an ordinary propensity score of the treatment given its Markov blanket derived directly from the distribution $q_V(V \mid W)$. The primal fixing operator instead uses $q_{D_T}(T \mid \mathrm{mb}_{\mathcal{G}}(T), W)$ which, as mentioned in Section 3, can be viewed as a *nested* propensity score derived from the post-intervention distribution $q_{D_T}(D_T \mid \mathrm{pa}_{\mathcal{G}}(D_T), W)$. The identification of $q_{D_T}(T \mid \mathrm{mb}_{\mathcal{G}}(T), W)$ from the kernel $q_{D_T}(D_T \mid \mathrm{pa}_{\mathcal{G}}(T), W)$ is the same as before, yielding Eq. 27.

Remark 20 *The primal fixability criterion accepts a CADMG and the corresponding distribution with variables that have possibly already been fixed or intervened on and determines the identifiability of a given variable T on all other variables on the graph. Unlike the original criterion proposed by Tian and Pearl (2002a), this allows for recursive applications of the operator, starting from any observed data distribution $q_V(V \mid W)$ that is nested Markov with respect to a CADMG $\mathcal{G}(V, W)$, where W may also be an empty set. Note that the proposed primal fixabilty criterion $\mathrm{dis}_{\mathcal{G}}(T) \cap \mathrm{ch}_{\mathcal{G}}(T) = \emptyset$ is also strictly more general than the ordinary fixability criterion proposed by Richardson et al. (2017).*

We now describe the dual fixing operator that serves the same purpose as the primal operator in that it is used to arrive at a CADMG and corresponding distribution obtained by intervening on a single variable T but does so through different means. Specifically, similar to the dual IPW described in Section 3, it uses kernels related to the inverse Markov pillow of T.

Lemma 21 (Dual fixing operator)
Given a distribution $q_V(V \mid W)$ that district factorizes with respect to a CADMG $\mathcal{G}(V, W)$, T is said to be dual fixable if it has no bidirected path to any of its children, i.e., $\mathrm{dis}_{\mathcal{G}}(T) \cap \mathrm{ch}_{\mathcal{G}}(T) = \emptyset$. The probabilistic operation of dual fixing, denoted by $\Delta_T(q_V; \mathcal{G})$, is given by

$$\Delta_T(q_V; \mathcal{G}) \equiv \sum_T q_V(V \mid W) \times \frac{\prod_{V_i \in \mathrm{mp}_{\mathcal{G}}^{-1}(T)} q_V(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i), W) \mid_{T=t}}{\prod_{V_i \in \mathrm{mp}_{\mathcal{G}}^{-1}(T)} q_V(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i), W)}. \tag{28}$$

The resulting kernel district factorizes with respect to a CADMG $\mathcal{G}(V \setminus T, W \cup T)$ where all incoming edges into T are removed. We denote this graphical operation of dual fixing variable T in $\mathcal{G}$ as $\Delta_T(\mathcal{G})$.

The results of the primal and dual fixing operators are equivalent in the following sense.

Lemma 22 (Equivalence of the primal and dual operators)
Given a CADMG $\mathcal{G}(V, W)$ and corresponding kernel $q_V(V \mid W)$ where T is primal (dual) fixable, $\Phi_T(\mathcal{G}) = \Delta_T(\mathcal{G})$, and $\Phi_T(q_V; \mathcal{G}) = \Delta_T(q_V; \mathcal{G})$ for any fixed value $T = t$.

This immediately yields the following corollary, connecting the primal and dual operators to the task of causal inference.

Corollary 23 (Identification via primal and dual fixing)
Given a causal model associated with a hidden variable DAG $\mathcal{G}(V \cup H)$, that induces the

observed marginal distribution $p(V)$ nested Markov with respect to the latent projection $\mathcal{G}(V)$, and T is primal (or dual) fixable, $\psi(t)$ is identifiable as,

$$\psi(t) = \sum_{V \setminus T} Y \times \Phi_T(p(V); \mathcal{G}(V))\big|_{T=t} = \sum_{V \setminus T} Y \times \Delta_T(p(V); \mathcal{G}(V)).$$

As is the case for the ordinary fixing operator ϕ, the primal and dual fixing operations may be applied to sequences of vertices $\sigma_S = (S_1, \ldots, S_p)$, provided this sequence is *p-fixable* in $\mathcal{G}(V, W)$. That is, S_1 is p-fixable in $\mathcal{G}$, S_2 is p-fixable in $\Phi_{S_1}(\mathcal{G})$, and so on. Given a sequence $\sigma_S = (S_1, \ldots, S_p)$ p-fixable in $\mathcal{G}(V, W)$, define $\Phi_{\sigma_S}(\mathcal{G}(V, W))$ as $\mathcal{G}$ if S is empty, and $\Phi_{\sigma_S \setminus S_1}(\Phi_{S_1}(\mathcal{G}))$ otherwise. Similarly, define $\Phi_{\sigma_S}(q_V; \mathcal{G})$ to be $q_V(V \mid W)$ if S is empty and $\Phi_{\sigma_S \setminus S_1}(\Phi_{S_1}(q_V; \mathcal{G}); \Phi_{S_1}(\mathcal{G}))$ otherwise. Since the graphical operator of p-fixing is equivalent to ordinary fixing, it follows that two valid p-fixing sequences yield the same CADMG. The following lemma formalizes that any two valid p-fixing sequences also yield the same kernel.

Lemma 24 (Commutativity of p-fixing)
If σ_S^1 and σ_S^2 are both sequences p-fixable in the ADMG $\mathcal{G}(V)$, then for any $p(V \cup H)$ Markov relative to a DAG $\mathcal{G}(V \cup H)$ that yields the latent projection $\mathcal{G}(V)$, $\Phi_{\sigma_S^1}(p(V); \mathcal{G}(V)) = \Phi_{\sigma_S^2}(p(V); \mathcal{G}(V))$.

Due to this lemma, for marginals $p(V)$ induced by causal models associated with a hidden variable DAG $\mathcal{G}(V \cup H)$, if S is p-fixable in the latent projection $\mathcal{G}(V)$, we define $\Phi_S(p(V); \mathcal{G}(V))$ to be the result of applying the fixing operator Φ to any p-fixable sequence on S and $p(V), \mathcal{G}(V)$.

6.3 Sequential Fixing Identification When T Is Not Primal Fixable

In this section, we show that identification of the target $\psi(t)$ may be possible even when the treatment T is not p-fixable in $\mathcal{G}(V)$. In particular, we focus on the class of ADMGs where a sequence of p-fixing operations applied to vertices Z that cause Y only through T yields a new CADMG $\mathcal{G}(V \setminus Z, Z)$ where T itself becomes p-fixable.

Consider the ADMG shown in Fig. 6(a). Clearly T is not p-fixable as it does not meet the condition that $\mathrm{dis}_{\mathcal{G}}(T) \cap \mathrm{ch}_{\mathcal{G}}(T) = \emptyset$. However, Z_1 is p-fixable, and yields a CADMG $\Phi_{Z_1}(\mathcal{G})$ where T is p-fixable as shown in Fig. 6(d). Further, Z_1 is an ancestor of Y, but only via a directed path through T. Hence, while the CADMG in Fig. 6(d) corresponds to the post-intervention distribution $p(C, Z_2(z_1), T(z_1), Y(z_1))$, fixing T in this CADMG yields the post-intervention distribution $p(C, Z_2(z_1), Y(t))$ from which $p(Y(t))$ can be easily obtained as $Y(t, z_1) = Y(t)$ (Malinsky et al., 2019). A similar argument can be made to show that p-fixing according to the sequence (Z_2, Z_1), resulting in the ADMGs shown in Figs. 6(b, c), also gives us the desired post-intervention distribution as $p(C, Y(t, z_1, z_2)) = p(C, Y(t))$.

The general procedure to obtain identification by a sequence of p-fixing operations is given by the following lemma. For the remainder of the paper we assume, without loss of generality, that $V = \mathrm{an}_{\mathcal{G}(V)}(Y)$ by considering the marginal distribution $p(\mathrm{an}_{\mathcal{G}}(Y))$ and graph $\mathcal{G}_{\mathrm{an}_{\mathcal{G}}(Y)}$. This greatly simplifies the notation in the proceeding sections.

Lemma 25 (Identification via a sequence of p-fixing)
Fix a causal model associated with a hidden variable DAG $\mathcal{G}(V \cup H)$, that induces the

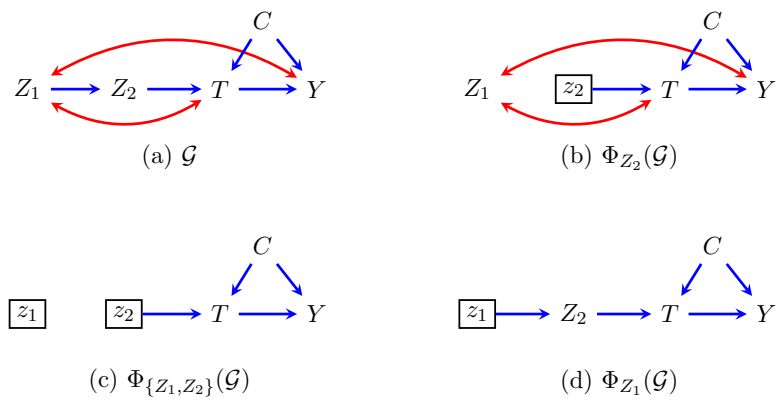

Figure 6: (a) An ADMG where T is not fixable. (b, c) A valid sequence of fixing that yields a CADMG where T is p-fixable and $p(Y(t))$ can be obtained as $p(Y(t)) = p(Y(t, z_1, z_2))$. (d) A graph obtained after Z_1 is p-fixed, where $p(Y(t))$ can be obtained by p-fixing T, as $p(Y(t)) = p(Y(t, z_1))$.

observed marginal distribution $p(V)$. Given $Y^* \equiv \text{an}_{\mathcal{G}_{V \setminus T}}(Y)$, $\psi(t)$ is identified if there exists a subset $Z \subseteq V \setminus (Y^* \cup T)$ that is p-fixable in $\mathcal{G}(V)$ such that T is p-fixable in $\Phi_Z(\mathcal{G}(V))$. Moreover, if $\psi(t)$ is identified, we have

$$\psi(t) = \sum_{V \setminus \{Z \cup T\}} Y \times \Phi_{Z \cup T}(p(V); \mathcal{G}(V)) \Big|_{T=t}. \tag{29}$$

The above result directly yields an inverse weighted estimator as described in Theorem 27.

Let $\widetilde{V} \equiv V \setminus Z$. The kernel $\Phi_Z(p(V); \mathcal{G}(V))$ evaluated at a particular set of values z yields a functional of $p(V)$ corresponding to the counterfactual distribution $p(\{V_i(z) : V_i \in \widetilde{V}\}) \equiv q_{\widetilde{V}}(\widetilde{V} \mid Z = z)$. In this distribution, $\psi(t)$ is identified by an analogue of Eq. 10, with $q_{\widetilde{V}}(\widetilde{V} \mid Z = z)$ substituted for $p(V)$:

$$\psi(t) = \sum_{V \setminus T} Y \times \prod_{V_i \in V \setminus D_T} q_{\widetilde{V}}(V_i \mid \text{mp}_{\mathcal{G}}(V_i), Z = z) \times \sum_{T} \prod_{V_i \in D_T} q_{\widetilde{V}}(V_i \mid \text{mp}_{\mathcal{G}}(V_i), Z = z) \Big|_{T=t}. \tag{30}$$

Note that $\psi(t)$ is not a function of Z since every element in Z is only an ancestor of Y through T, so $Y(t, z) = Y(t)$. As a result, if we view $q_{\widetilde{V}}(\widetilde{V} | Z = z)$ (for any values z) as the observed data distribution over $\widetilde{V}$, the following analogue of Theorem 9 is immediate.

Lemma 26 (Nonparametric IF in models of a CADMG)
Fix the observed data distribution $q_{\widetilde{V}}(\widetilde{V} \mid Z = z)$ over $\widetilde{V}$ (for some values z of Z), where the target of inference $\psi(t)$ is given by Eq. 30. Then an influence function for $\psi(t)$, denoted by $U(\psi(t); q_{\widetilde{V}})$, has the same form as the influence function in Theorem 9 after substitution of all conditional densities $p(\cdot \mid \cdot)$ in Eq. 15 with $q_{\widetilde{V}}(\cdot \mid \cdot, Z = z)$.

The influence function in Lemma 26 maintains the double robustness property of the IF in Theorem 9, except with respect to nuisance functions derived from $q_{\widetilde{V}}$, rather than p. Specifically, define $\mathbb{L}_Z$ and $\mathbb{M}_Z$ via display (13), applied to $\Phi_Z(\mathcal{G}(V))$, rather than the original graph $\mathcal{G}(V)$. Then the estimator is consistent if all models in either the set $\{q_{\widetilde{V}}(M_i \mid \mathrm{mp}_{\Phi_Z(\mathcal{G})}(M_i), Z = z),\ \forall M_i \in \mathbb{M}_Z\}$ or $\{q_{\widetilde{V}}(L_i \mid \mathrm{mp}_{\Phi_Z(\mathcal{G})}(L_i), Z = z),\ \forall L_i \in \mathbb{L}_Z\}$ are correctly specified.

Lemma 26 provides an estimating equation for $\psi(t)$ as $\mathbb{E}_{q_{\widetilde{V}}}[U(\psi(t); q_{\widetilde{V}})] = 0$, where the expectation is taken with respect to $q_{\widetilde{V}}$. In reality, data on $q_{\widetilde{V}}$ is not directly available. However, since $q_{\widetilde{V}}$ is a known functional of $p(V)$ given by sequentially applying the p-fixing operator, a natural estimator for $\psi(t)$ given realizations of $p(V)$ would use inverse probability weighting to create a reweighted distribution that resembles $q_{\widetilde{V}}$. These weights are given by nuisance functionals $\pi(Z_1), \pi(Z_2), \ldots, \pi(Z_p)$ of $p(V)$ implied by the primal fixing operator. We formalize this in the following theorem.

Theorem 27 (Reweighted estimating equations)
Let $(Z_1, \ldots, Z_p)$ be a p-fixable sequence on Z in the identifying functional for $\psi(t)$, given in Eq. 29. Let π_{Z_i} be defined inductively as follows. π_{Z_1} is defined as Eq. 27 with $q_V(V \mid W) = p(V)$, and π_{Z_i} is defined as Eq. 27 with $q_V(V \mid W) = \Phi_{\{Z_1,\ldots,Z_{i-1}\}}(p(V); \mathcal{G}(V))$. Consider the estimating equation

$$\mathbb{E}\left[\ \frac{p^*(Z)}{\prod_{Z_i \in Z} \pi_{Z_i}} \times U(\psi(t); \Phi_Z(p(V); \mathcal{G}(V))\ \right] = 0, \tag{31}$$

where $p^*(Z)$ is any normalized density of Z and $U(\psi(t); \Phi_Z(p(V); \mathcal{G}(V))$ is obtained via Lemma 26. This esitmating equation yields a consistent estimator for $\psi(t)$, provided all models in $\{\pi_{Z_i} : Z_i \in Z\}$ and models in either the set $\{q_{\widetilde{V}}(M_i \mid \mathrm{mp}_{\Phi_Z(\mathcal{G})}(M_i), Z = z),\ \forall M_i \in \mathbb{M}_Z\}$ or $\{q_{\widetilde{V}}(L_i \mid \mathrm{mp}_{\Phi_Z(\mathcal{G})}(L_i), Z = z),\ \forall L_i \in \mathbb{L}_Z\}$ are correctly specified.

Remark 28 *When T is primal fixable in $\mathcal{G}$, Theorem 27 reduces to Lemma 10, yielding a full double robustness property. When T is a-fixable in $\Phi_Z(\mathcal{G})$, it is easy to see that Lemma 26 can also be adapted to yield a gAIPW style estimator for $\psi(t)$ where the nuisances are models for the propensity score $q_{\widetilde{V}}(T \mid \mathrm{mp}_{\Phi_Z(\mathcal{G})}(T), Z = z)$ and the outcome regression $\mathbb{E}_{q_{\widetilde{V}}}[Y \mid T = t, \mathrm{mp}_{\Phi_Z(\mathcal{G})} T, Z = z]$.*

6.3.1 Example: Reweighted Estimating Equations

We now work through a usage of Theorem 27 for the ADMG $\mathcal{G}(V)$ in Fig. 6(a), where we aim to p-fix Z_1 prior to p-fixing T in the graph $\Phi_{Z_1}(\mathcal{G}(V))$ shown in Fig. 6(d). By Eq. 27,

$$\pi_{Z_1} = \frac{p(Y \mid T, Z, C) \times p(T \mid Z, C) \times p(Z_1)}{\sum_{Z_1}\ p(Y \mid T, Z, C) \times p(T \mid Z, C) \times p(Z_1)},$$

where $Z = \{Z_1, Z_2\}$. The resulting estimating equation is given by,

$$\mathbb{E}\Big[p^*(Z_1) \times \frac{\sum_{Z_1}\ p(Y \mid T, Z, C) \times p(T \mid Z, C) \times p(Z_1)}{p(Y \mid T, Z, C) \times p(T \mid Z, C) \times p(Z_1)} \times U\big(\psi(t); \Phi_{Z_1}(p(V); \mathcal{G}(V))\big)\Big] = 0,$$

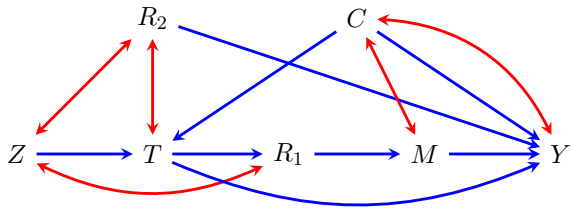

Figure 7: An ADMG where sequential p-fixing is insufficient but $\psi(t)$ is still identified via the truncated nested Markov factorization.

where $U(\psi(t); \Phi_{Z_1}(p(V); \mathcal{G}(V)))$ is given in Theorem 9, with nuisance models specified with respect to $p(V)/\pi_{Z_1}$, rather than with respect to $p(V)$. In the current example, the nuisances in this distribution are simply the propensity score model for the treatment given covariates C and the outcome regression model for Y given the treatment and C, as seen from the corresponding ADMG in Fig. 6(d).

Alternatively, one could also consider the fixing sequence (Z_2, Z_1) which also yields a CADMG $\Phi_{Z_1,Z_2}(\mathcal{G})$ where T is p-fixable as in Fig. 6(c). In this case, π_{Z_2} is simply $p(Z_2 \mid Z_1)$. In the CADMG $\Phi_{Z_2}(\mathcal{G})$, the variable Z_1 is childless. Therefore, p-fixing Z_1 in the corresponding distribution corresponds to marginalization of Z_1 (Richardson et al., 2017). Thus, any p-fixings that occur in the sequence that correspond to marginalization in this manner, do not require the specification of an additional p-fixing weight. Hence, the estimating equation in this case is simply,

$$\mathbb{E}\left[\frac{p^*(Z_1, Z_2)}{p(Z_2 \mid Z_1)} \times U\big(\psi(t); \Phi_{Z_1,Z_2}(p(V); \mathcal{G}(V))\big)\right] = 0. \tag{32}$$

We now briefly describe an implementation that would allow us to estimate $\psi(t)$ using the above estimating equation. We first fit a model for the conditional density $p(Z_2 \mid Z_1)$. We then use this model to predict inverse weights for each sample. We then fit the nuisance models $q_{CTY}(T \mid C)$ and $\mathbb{E}_{q_{CTY}}[Y \mid T = t, C]$ using logistic/linear regression or any flexible method that can be fit by utilizing the inverse weights in its objective function. If additional weights π_{Z_i} were needed, these models would be fit recursively using a product of the inverse weights from the previous stages. The final nuisance models in U_{ψ_t} would then use $1/\prod_{Z_i \in Z} \pi_{Z_i}$ as weights. Finally, the estimate for $\psi(t)$ is obtained by plugging in predictions from these models according to the functional form of the estimator derived from U_{ψ_t} in Theorem 26, which in the case of Fig. 6(c) reduces to the form of the gAIPW estimator and empirically evaluating the expectation.

6.4 Augmented Nested IPW Estimators For Any Identifiable $\psi(t)$

We now describe our most general algorithm — one that yields an IPW estimator for *any* $\psi(t)$ that is identifiable from the observed margin $p(V)$ corresponding to an ADMG $\mathcal{G}(V)$. Consider the ADMG shown in Fig. 7. T is neither p-fixable nor can we directly

Algorithm 2 Nested IPW $(\mathcal{G}, p(V))$

1: Let $Y^* = \text{an}_{\mathcal{G}_{V \setminus T}}(Y)$ and $D_T = \text{dis}_{\mathcal{G}}(T)$ and $\mathcal{D}^* \leftarrow \{D \in \mathcal{D}(\mathcal{G}_{Y^*}) \mid D \cap D_T \neq \emptyset\}$
2: **if** $\exists D \in \mathcal{D}^*$ such that D is not intrinsic in $\mathcal{G}$ **then**
3: **return** Fail
4: Define non-descendants of a vertex V_i as $\text{nd}_{\mathcal{G}}(V_i) \equiv V \setminus \text{de}_{\mathcal{G}}(V_i)$
5: Fix a topological order τ such that $V_i \succ_\tau \text{nd}_{\mathcal{G}}(V_i) \setminus Y^*,\ \forall V_i \in Y^*$
6: Define $q_D(D \mid \text{pa}_{\mathcal{G}}(D)) \equiv \phi_{V \setminus D}(p(V); \mathcal{G}(V))$
7: $p^\dagger(V) \leftarrow p(V) \times \prod_{D \in \mathcal{D}^*} \left(q_D(D \mid \text{pa}_{\mathcal{G}}(D)) \times \prod_{D_i \in D} \frac{1}{p(D_i \mid \text{mp}_{\mathcal{G}}(D_i))} \right)$
8: **return** $\psi(t)_{\text{nested}} \equiv \mathbb{E}_{p^\dagger}\left[\frac{\mathbb{I}(T=t)}{p(T \mid \text{mp}_{\mathcal{G}}(T))} \times Y \right]$

apply Theorem 27 as there does not exist a valid p-fixing sequence for $V \setminus Y^* = \{Z, T\}$. However, $\psi(t)$ is in fact still identified via the truncated nested Markov factorization as follows. $Y^* \equiv \{Y, M, C, R_1, R_2\}$ and $\mathcal{D}(\mathcal{G}_{Y^*}) = \{\{Y, M, C\}, \{R_1\}, \{R_2\}\}$. Then

$$\psi(t) = \sum_{Y^*} Y \times \phi_{V \setminus \{Y,M,C\}}(p(V); \mathcal{G}) \times \phi_{V \setminus R_1}(p(V); \mathcal{G}) \times \phi_{V \setminus R_2}(p(V); \mathcal{G}) \Big|_{T=t}.$$

In the following theorem, we provide the corresponding IPW estimator for all targets $\psi(t)$ that are identifiable from the observed margin of a hidden variable causal DAG $\mathcal{G}(V \cup H)$. As these estimators are derived from the nested Markov factorization of the latent projection ADMG $\mathcal{G}(V)$, we coin the term nested IPW in referring to them. We show that Algorithm 2 which we use to derive such estimators is *sound* and *complete*. That is, when Algorithm 2 returns a nested IPW functional, $\psi(t)_{\text{nested}} = \psi(t)$ and when the algorithm fails to return a functional, $\psi(t)$ is not identifiable within the given model.

Theorem 29 (Soundness and completeness of nested IPW)
Let $p(V)$ and $\mathcal{G}(V)$ be the observed marginal distribution and ADMG induced by a hidden variable causal model associated with DAG $\mathcal{G}(V \cup H)$. Then if $\psi(t)$ is identifiable in the model, $\psi(t) = \psi(t)_{nested}$. If $\psi(t)$ is not identifiable in the model, Algorithm 2 returns 'fail'.

Based on the above nested IPW estimator, we now derive a natural augmented nested IPW estimator that exhibits a partial double robustness property similar to gAIPW in Theorem 2. That is, provided the rebalancing weights in line 7 of Algorithm 2 are correctly specified, the estimator that we derive is consistent if either a propensity score model $p^\dagger(T \mid \text{mp}_{\mathcal{G}}(T))$ or outcome regression model $\mathbb{E}_{p^\dagger}[Y \mid T = t, \text{mp}_{\mathcal{G}}(T)]$ is correctly specified. Such a result, to our knowledge, is the first of its kind, and represents a powerful unification of the nested Markov model of an ADMG and semiparametric statistics.

Theorem 30 (Augmented nested IPW for any identifiable $\psi(t)$)
Under the same conditions as Lemma 29, define $q_D(D \mid \text{pa}_{\mathcal{G}}(D)) \equiv \phi_{V \setminus D}(p(V); \mathcal{G}(V))$, and let ρ_D as the rebalancing weights for all $D \in \mathcal{D}^*$ used in line 7 of Algorithm 2. That is,

$$\rho_D = \frac{q_D(D \mid \text{pa}_{\mathcal{G}}(D))}{\prod_{D_i \in D}\ p(D_i \mid \text{mp}_{\mathcal{G}}(D_i))}.$$

Then $\mathbb{E}\big[\prod_{D\in\mathcal{D}^*}\rho_D \times U^{\dagger}_{\psi_t}\big] = 0$ *where* $U^{\dagger}_{\psi_t}$ *takes the form of gAIPW in Theorem 2, yields a consistent estimator for* $\psi(t)$ *provided all models in* $\{\rho_D : D \in \mathcal{D}^*\}$ *and either the propensity score model* $p^{\dagger}(T \mid \mathrm{mp}_{\mathcal{G}}(T))$ *or the outcome regression* $\mathbb{E}_{p^{\dagger}}[Y \mid T = t, \mathrm{mp}_{\mathcal{G}}(T)]$ *are correctly specified.*

6.4.1 Example: ANIPW

We now return to the ADMG shown in Fig. 7 and discuss the application of Theorem 30, in order to obtain an estimator for $\psi(t)$. Recall, $Y^* \equiv \{Y, M, C, R_1, R_2\}$ and $\mathcal{D}(\mathcal{G}_{Y^*}) = \{\{Y, M, C\}, \{R_1\}, \{R_2\}\}$. Note that $\mathcal{D}^*$ simply focuses on the districts related to $\mathcal{G}_{Y^*}$ that do not overlap with D_T. Therefore, $\mathcal{D}^*$ in line 1 of the algorithm is $\{\{R_1\}, \{R_2\}\}$. Since both of these districts are intrinsic in $\mathcal{G}$, Algorithm 2 does not fail. Fix the topological order $(Z, C, T, R_1, R_2, M, Y)$ according to line 5. Then,

$$\text{(Fig. 7)} \qquad \begin{aligned} p^{\dagger}(V) &= p(V) \times \frac{q_{R_1}(R_1 \mid T)}{p(R_1 \mid T, Z)} \times \frac{q_{R_2}(R_2)}{p(R_2 \mid T, Z, C, R_1)} \\ &= p(V) \times \frac{\sum_{Z,C} p(Z, C) \times p(R_1 \mid T, Z, C)}{p(R_1 \mid T, Z)} \times \frac{p(R_2)}{p(R_2 \mid T, Z, C, R_1)} \end{aligned} \tag{33}$$

Fitting models for each ρ_D shown above then serves to produce inverse weights for fitting the final propensity score model $p^{\dagger}(T \mid Z, C)$ and outcome regression $\mathbb{E}_{p^{\dagger}}[Y \mid Z, C]$ which gives us an estimate for $\psi(t)$ using the functional form of gAIPW. This estimate is doubly robust in the final nuisances of the propensity score and outcome regression, after requiring correct specification of the initial models required to produce $p^{\dagger}(V)$.

7. Experiments

In this section, we briefly describe experiments and results for each major contribution in our paper. For each experiment, we generate data according to hidden variable DAGs that give rise to the latent projection ADMGs used in the motivating examples throughout the paper. Specifically, for each bidirected edge in the latent projection ADMG, we allow for the presence of two unmeasured confounders that are parents of both end points of the bidirected edge; one that is sampled from a uniform distribution and one that is sampled from a Bernoulli distribution. For example in Fig. 2, for the bidirected edge $Z_1 \leftrightarrow T$, the underlying hidden variable DAG contains variables H_1 and H_2 which are parents of both Z_1 and T. We provide an example of such a data generating process in Appendix C.

We use generalized additive models to fit all of our nuisances. For each experiment on double robustness, boxplots for the average causal effect under different settings of model misspecification are calculated from 100 trials with a sample size of $n = 1000$. Our form of model misspecification is simple, involving simply dropping some of the appropriate conditioning variables in the nuisance models to demonstrate robustness to *arbitrary* model misspecification. Tables for demonstrating the reduction in variance using the efficient influence function are calculated from 100 trials with increasing sample size.

Simulation 1. Double robustness and efficiency of gAIPW
In this simulation, we generated data according to a hidden variable DAG $\mathcal{G}(V \cup H)$ that gives rise to the latent projection ADMG $\mathcal{G}(V)$ shown in Fig. 2. As expected, gAIPW

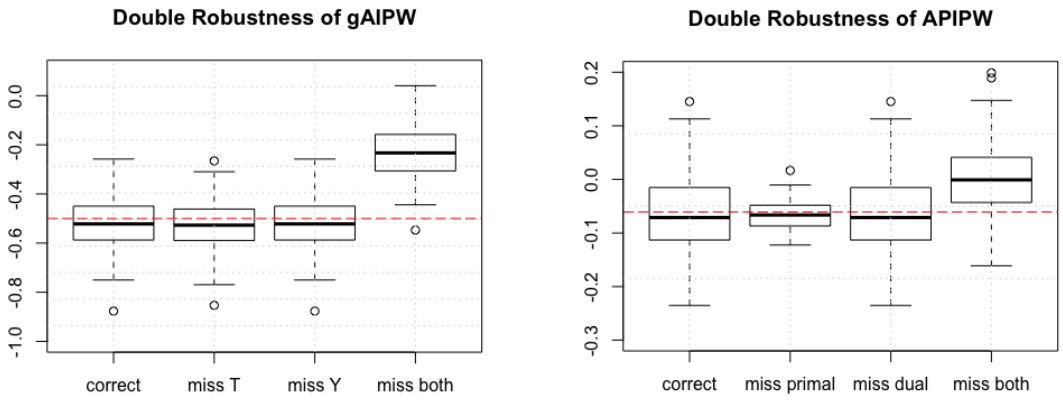

n	Nonparametric IF	Efficient IF
500	0.127, 0.406	0.127, 0.298
1000	0.091, 0.212	0.091, 0.133
2000	0.071, 0.097	0.071, 0.069
5000	0.044, 0.041	0.044, 0.026

n	Nonparametric IF	Efficient IF
500	0.079, 0.354	0.079, 0.207
1000	0.059, 0.166	0.059, 0.103
2000	0.040, 0.088	0.040, 0.054
5000	0.027, 0.035	0.027, 0.021

Figure 8: Experiments for gAIPW (first column) and APIPW (second column) using the ADMGs in Figs. 2 and 5, respectively. The panels on the first row demonstrate the double robustness property. The red dashed lines indicate the true values of the ACE. Tables on the second row compare the bias and variance between the nonparametric IF and the efficient one as a function of sample size. The format of each cell is (bias, variance).

remains unbiased (see top left panel of Fig. 8) when at least one of the nuisance models corresponding to the propensity score or outcome regression is specified correctly and is biased if they are both misspecified. The table on the bottom left panel of Fig. 8 summarizes the experiment comparing variances of the nonparametric IF and efficient IF based estimators. As expected, though both estimators are unbiased, the one based on the efficient IF offers lower variance at all sample sizes. The absolute reduction in variance is greater and most beneficial at smaller sample sizes, but the relative reduction in variance remains the same across all sample sizes. This highlights the benefit of using the efficient IF in many practical applications where sample sizes may be small.

Simulation 2. Double robustness and efficiency of APIPW
In this simulation, we generated data according to a hidden variable DAG $\mathcal{G}(V \cup H)$ that gives rise to the latent projection ADMG $\mathcal{G}(V)$ shown in Fig. 5. The experiments on double robustness of APIPW are shown in the second column of Fig. 8. Estimates are unbiased when at least one set of models corresponding to those used in β_{primal} or β_{dual} are correctly

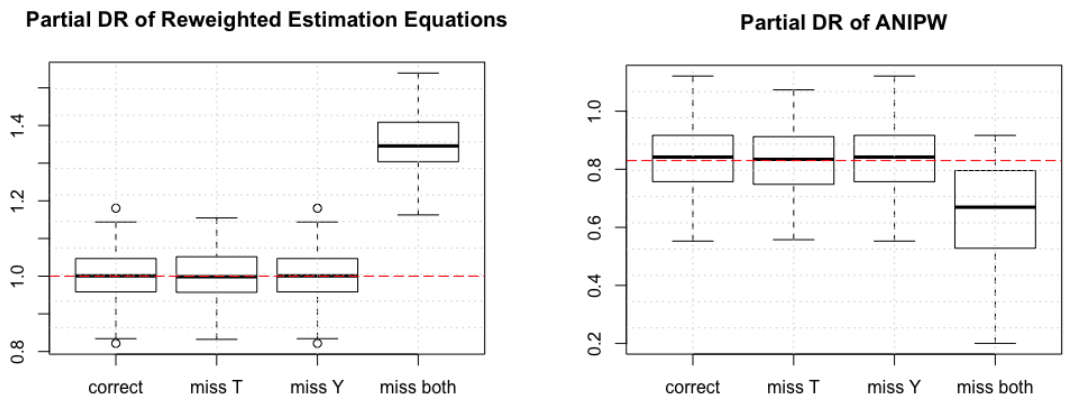

Figure 9: Demonstrating partial double robustness for sequential reweighting (left) and ANIPW (right) using the ADMGs in Figs. 6(a) and 7, respectively. The red dashed lines indicate the true values of the ACE.

specified. The nuisances involved in each set can be seen in Eq. 23. The table on the bottom right panel in Fig. 8 summarizes the experiment in comparing the variances of the nonparametric and efficient IF based estimators. Once again, the efficient IF based estimator yields lower variance at all samples sizes.

Simulation 3. Partial double robustness of sequentially reweighted estimating equations
In this simulation, we generated data according to a hidden variable DAG $\mathcal{G}(V \cup H)$ that gives rise to the latent projection ADMG $\mathcal{G}(V)$ shown in Fig. 6(a). We opt to use the simpler estimating equation presented in Eq. 32 that requires us to correctly specify a nuisance model for $p(Z_2 \mid Z_1)$. Upon correct specification of this nuisance model, the boxplots on the left in Fig. 9 demonstrate partial double robustness for the remaining nuisance models — a propensity score model $q_{CTY}(T \mid C)$ and outcome regression model $\mathbb{E}_{q_{CTY}}[Y \mid T = t, C]$.

Simulation 4. Partial double robustness of ANIPW
In this simulation, we generated data according to a hidden variable DAG $\mathcal{G}(V \cup H)$ that gives rise to the latent projection ADMG $\mathcal{G}(V)$ shown in Fig. 7. After estimating the nuisances required to yield the distribution $p^{\dagger}(V)$ as in Eq. 32, boxplots on the right in Fig. 9 demonstrate partial double robustness of ANIPW in the remaining nuisances of the propensity score model and outcome regression model in $p^{\dagger}(V)$.

8. Conclusion

In this paper, we bridged the gap between identification and estimation theory for the causal effect of a single treatment on a single outcome in hidden variable causal models associated with directed acyclic graphs (DAGs). We provided a simple graphical criterion,

a-fixability, that permits the use of an influence function based estimator that generalizes the augmented inverse probability weighted (AIPW) estimator to settings with hidden variables. We then provided another graphical criterion, p-fixability, which when satisfied allows for the derivation of two novel IPW estimators – primal and dual IPW. We further derived the nonparametric influence function under p-fixability of the treatment that yields the augmented primal IPW estimator and showed that it is doubly robust in the models used in primal and dual IPW estimators.

We considered restrictions on the tangent space implied by the latent projection acyclic directed mixed graph (ADMG) of the hidden variable causal model. We provided an algorithm that is sound and complete for the purposes of checking the nonparametric saturation status of a hidden variable causal model as long these hidden variables are unrestricted. Further, through the use of mb-shielded ADMGs, we provided a graphical criterion that defines a class of hidden variable causal models whose score restrictions resemble those of a DAG with no hidden variables. For the class of causal models that can be expressed as an mb-shielded ADMG, we then derived the form of the efficient influence function under a-fixability and p-fixability, that takes advantage of the Markov restrictions implied on the observed data. These results are completely generic and may be used to derive the efficient version of *any* nonparametric influence function in the model with these restrictions.

Next, we defined the primal and dual fixing operators that operationalize primal and dual IPW so that they can be recursively applied to simplify problems where the treatment is not directly p-fixable. This resulted in estimators that resemble influence function based estimators that are sequentially reweighted and partially doubly robust in the final nuisance models used after reweighting.

Finally, we developed a semiparametric estimator for *any* identifiable causal effect involving a single treatment and a single outcome. This *nested AIPW estimator* takes the form of a reweighted AIPW estimator with a set of weights given by terms in a modified nested Markov factorization, and enjoys a partial double robustness property, provided the models for the set of weights are correctly specified.

A natural extension of the present work is deriving influence function based estimators for any identifiable causal effect (including those that involve multiple treatment variables), and finding their most efficient versions by projecting onto the tangent space defined by equality restrictions, such as conditional independences and Verma constraints, implied by the causal model.

Appendix A. Glossary of Terms and Notations

Symbol	Definition	Symbol	Definition
T	Treatment	$\mathcal{G}(V)$	Graph $\mathcal{G}$ with vertices V
$Y, Y(t)$	Outcome, potential outcome	$\mathrm{pa}_{\mathcal{G}}(V_i)$	Parents of V_i in $\mathcal{G}$
V	Observed variables	$\mathrm{ch}_{\mathcal{G}}(V_i)$	Children of V_i in $\mathcal{G}$
H	Unmeasured variables	$\mathrm{de}_{\mathcal{G}}(V_i)$	Descendants of V_i in $\mathcal{G}$
W	Fixed variables	$\mathrm{an}_{\mathcal{G}}(V_i)$	Ancestors of V_i in $\mathcal{G}$
$\psi(t)$	Target parameter $\mathbb{E}[Y(t)]$	$\mathrm{mb}_{\mathcal{G}}(V_i)$	Markov blanket of V_i in $\mathcal{G}$
U_{ψ_t}	Influence function for $\psi(t)$	$\mathrm{mp}_{\mathcal{G}}(V_i)$	Markov pillow of V_i in $\mathcal{G}$
H	Hilbert space	$\mathrm{mp}_{\mathcal{G}}^{-1}(V_i)$	Inverse of Markov pillow of V_i
Λ	Tangent space	$\mathcal{D}(\mathcal{G})$	Districts in $\mathcal{G}$
$\Lambda^{\perp}$	Orthogonal complement	$\mathrm{dis}_{\mathcal{G}}(V_i)$	District of V_i in $\mathcal{G}$
$\mathcal{U}$	Class of all influence functions	D_T	District of T
U_{ψ}^{eff}	Efficient influence function	τ	A valid topological order
$\pi[h \mid \Lambda]$	Projection of h onto Λ	$\mathcal{G}_S$	Subgraph of $\mathcal{G}$ on vertices S
$\mathbb{C}$	Pre-treatment variables	$V_i \prec V_j$	V_i precedes V_j
$\mathbb{L}$	Post-treatment variables in D_T	$\{\prec V_i\}$	Vertices preceding V_i
$\mathbb{M}$	Variables not in $\mathbb{C} \cup \mathbb{L}$	$\phi_{V_i}(\mathcal{G})$	Fixing V_i in $\mathcal{G}$
$\mathbb{P}_n$	Empirical distribution	$\phi_{V_i}(q_V; \mathcal{G})$	Fixing V_i in $q_V(\cdot \mid \cdot)$
$S(V)$	Score of $p(V)$	$\langle S \rangle_{\mathcal{G}}$	Reachable closure of S in $\mathcal{G}$
Y^*	$\mathrm{an}_{\mathcal{G}_{V \setminus T}}(Y)$	$\mathcal{G}(V, W)$	A CADMG with fixed W
$\Phi_{V_i}(\mathcal{G})$	Primal fixing V_i in $\mathcal{G}$	$\Phi_{V_i}(q_V; \mathcal{G})$	Primal fixing V_i in q_V
$\Delta_{V_i}(\mathcal{G})$	Dual fixing V_i in $\mathcal{G}$	$\Delta_{V_i}(q_V; \mathcal{G})$	Dual fixing V_i in q_V

Appendix B. Example of Latent Projection

(a) $\mathcal{G}(V, H)$ (b) $\mathcal{G}(V)$

Appendix C. Details on Simulated Data

We first generate ten hidden variables that are used across all four of our simulations. H_i are sampled from a Binomial distribution, with corresponding probability of p_{u_i} for

$i = 1, 3, 5, 7, 9$ and H_j are sampled from a Uniform distribution with corresponding lower bound a_{u_j} and upper bound b_{u_j} for $j = 2, 4, 6, 8, 10$. For the observed variables in each simulation, continuous variables are sampled from normal distributions and binary variables are sampled from Bernoulli distributions. In Fig. 2, we assume Z_1, T, D_2 are binary random variables and the rest are continuous. In Fig. 5, we assume Z_1, C_1, T, M, L are binary. In Fig. 6(a), we assume Z_2, T are binary. In Fig. 7, we assume R_1, R_2, T are binary [2]. As a concrete example, we illustrate the data generating process for Fig. 2 used in our first simulation. We use $H_{i,j}$ as a shorthand for $\{H_i, H_j\}$, $C = \{C_1, C_2\}$, and $Z = \{Z_1, Z_2\}$.

$$
\begin{aligned}
\text{(Fig. 2)} \qquad C_1 \mid H_{3,4} &\sim \alpha_0^{c_1} + \alpha_1^{c_1} H_3 + \alpha_2^{c_1} H_4 + \mathcal{N}(0,1) \\
C_2 &\sim \alpha_0^{c_2} + \alpha_1^{c_2} H_{5,6} + \mathcal{N}(0,1) \\
p(Z_1 = 1 \mid C, H_{3,4}) &\sim \text{expit}(\alpha_0^{z_1} + \alpha_1^{z_1} C_1 + \alpha_2^{z_1} C_2 + \alpha_3^{z_1} H_1 + \alpha_4^{z_1} H_2) \\
Z_2 \mid Z_1, H_{1,2} &\sim \alpha_0^{z_2} + \alpha_1^{z_2} Z_1 + \alpha_2^{z_2} H_3 + \alpha_3^{z_2} H_4 + \mathcal{N}(0,1) \\
p(T = 1 \mid C, Z_2, H_{3,4}) &\sim \text{expit}(\alpha_0^{t} + \alpha_1^{t} C_1 + \alpha_2^{t} C_2 + \alpha_3^{t} Z_2 + \alpha_4^{t} H_3 + \alpha_5^{t} H_4) \\
M \mid C_1, C_2, T &\sim \alpha_0^{m} + \alpha_1^{m} C_1 + \alpha_2^{m} C_2 + \mathcal{N}(0,1) \\
Y \mid M, H_{5,6,7,8} &\sim \alpha_0^{y} + \alpha_1^{y} M + \alpha_2^{y} H_5 + \alpha_3^{y} H_6 + \alpha_4^{y} H_7 + \alpha_5^{y} H_8 + \mathcal{N}(0,1) \\
D_1 \mid C_2, M, H_{5,6,7,8} &\sim \alpha_0^{d_1} + \alpha_1^{d_1} C_2 + \alpha_2^{d_1} M + \alpha_3^{d_1} H_5 + \alpha_4^{d_1} H_6 + \alpha_5^{d_1} H_7 + \alpha_6^{d_1} H_8 + \mathcal{N}(0,1) \\
p(D_2 = 1 \mid D_1, Y) &\sim \text{expit}(\alpha_0^{d_2} + \alpha_1^{d_2} D_1 + \alpha_3^{d_2} Y)
\end{aligned}
$$

Appendix D. A Brief Overview of Semiparametric Estimation Theory

Assume a statistical model $\mathcal{M} = \{p_\eta(Z) : \eta \in \Gamma\}$ where Γ is the parameter space and η is the parameter indexing a specific model. We are often interested in a function $\psi : \eta \in \Gamma \mapsto \psi(\eta) \in \mathbb{R}$; i.e., a parameter that maps the distribution P_η to a scalar number in $\mathbb{R}$, such as an identified average causal effect. (For brevity, we sometimes use ψ instead of $\psi(\eta)$, which should be obvious from context.) Truth is denoted by P_{η_0} and ψ_0. An estimator $\widehat{\psi}_n$ of a scalar[3] parameter ψ based on n i.i.d copies $Z_1, \ldots, Z_n$ drawn from $p_\eta(Z)$, is *asymptotically linear* if there exists a measurable random function $U_\psi(Z)$ with mean zero and finite variance such that

$$\sqrt{n} \times (\widehat{\psi}_n - \psi) = \frac{1}{\sqrt{n}} \times \sum_{i=1}^{n} U_\psi(Z_i) + o_p(1), \tag{34}$$

where $o_p(1)$ is a term that converges in probability to zero as n goes to infinity. The random variable $U_\psi(Z)$ is called the *influence function* of the estimator $\widehat{\psi}_n$. The term influence function comes from the robustness literature (Hampel, 1974).

Before mentioning the asymptotic properties of an asymptotically linear estimator, it is worth noting that in asymptotic theory, we can sometimes construct *super efficient* estimators, e.g. Hodges estimator, that have undesirable local properties associated with them.

2. The code is attached to the submission as part of the supplementary materials.
3. Here, our focus is on estimation of $\psi = \mathbb{E}[Y(t)]$ which is a scalar parameter. For an extension to a vector valued functional in $\mathbb{R}^q, q > 1$, refer to (Tsiatis, 2007; Bickel et al., 1993).

Therefore, the analysis is oftentimes restricted to *regular*[4] and asymptotically linear (RAL) estimators to avoid such complications. Although most reasonable estimators are RAL, regular estimators do exist that are not asymptotically linear. However, as a consequence of Hájek (1970) representation theorem, the most efficient regular estimator is asymptotically linear; hence, it is reasonable to restrict attention to RAL estimators. According to Newey (1990), the influence function of a RAL estimator is the same as the influence function of its estimand. Further, there is a bijective correspondence between RAL estimators and influence functions.

By a simple consequence of the central limit theorem and Slutsky's theorem, it is straightforward to show that the RAL estimator $\widehat{\psi}_n$ is *consistent and asymptotically normal* (CAN), with asymptotic variance equal to the variance of its influence function U_ψ,

$$\sqrt{n} \times (\widehat{\psi}_n - \psi) \xrightarrow{d} N\big(0,\ \mathrm{var}(U_\psi)\big). \tag{35}$$

The first step in dealing with a semiparametric model, is to consider a simpler finite-dimensional parametric submodel that is contained within the semiparametric model and it contains the truth. Consider a (regular) parametric submodel $\mathcal{M}_{\text{sub}} = \{P_{\eta_\kappa} : \kappa \in [0,1)$ where $P_{\eta_{\kappa=0}} = P_{\eta_0}\}$ of the model $\mathcal{M}$. Given P_{η_0}, define the corresponding score to be $S_{\eta_0}(Z) = \dfrac{d}{d\kappa} \log p_{\eta_\kappa}(Z)\Big|_{\kappa=0}$. It is known that

$$\frac{d}{d\kappa}\psi(\eta_\kappa)\Big|_{\kappa=0} = \mathbb{E}\Big[U_\psi(Z) \times S_{\eta_0}(Z)\Big], \tag{36}$$

where $\psi(\eta_\kappa)$ is the target parameter in the parametric submodel, $U_\psi(Z)$ is the corresponding influence function evaluated at law P_{η_0}, $S_{\eta_0}(Z)$ is the score of the law P_{η_0}, and the expectation is taken with respect to P_{η_0}. Equation 36 provides an easy way to derive an influence function for the parameter ψ. In the next subsection, we use this equation to derive an influence function for our target $\psi = \mathbb{E}[Y(t)]$ and discuss its properties.

Influence functions provide a geometric view of the behavior of RAL estimators. Consider a Hilbert space[5] H of all mean-zero scalar functions, equipped with an inner product defined as $\mathbb{E}[h_1 \times h_2], h_1, h_2 \in \mathrm{H}$. The *tangent space* in the model $\mathcal{M}$, denoted by Λ, is defined to be the mean-square closure of parametric submodel tangent spaces, where a parametric submodel tangent space is the set of elements $\Lambda_{\eta_\kappa} = \{\alpha S_{\eta_\kappa}(Z)\}, \alpha$ is a constant and S_{η_κ} is the score for the parameter ψ_{η_κ} for some parametric submodel. In mathematical form, $\Lambda = \overline{[\Lambda_{\eta_\kappa}]}$.

The tangent space Λ is a closed linear subspace of the Hilbert space H ($\Lambda \subseteq \mathrm{H}$). The orthogonal complement of the tangent space, denoted by $\Lambda^\perp$, is defined as $\Lambda^\perp = \{h \in \mathrm{H} \mid \mathbb{E}[h \times h'] = 0, \forall h' \in \Lambda\}$. Note that $\mathrm{H} = \Lambda \oplus \Lambda^\perp$, where $\oplus$ is the direct sum, and $\Lambda \cap \Lambda^\perp = \{0\}$. Given an arbitrary element $h \in \Lambda^\perp$, it holds that for any submodel $\mathcal{M}_{\text{sub}}$, with score S_{η_0} corresponding to P_{η_0}, $\mathbb{E}[h \times S_{\eta_0}] = 0$. Consequently, using Eq. 36, $h + U_\psi(Z)$ is also an influence function. The vector space $\Lambda^\perp$ is then of particular importance because we can

4. Given a collection of probability laws $\mathcal{M}$, an estimator $\widehat{\psi}$ of $\psi(P)$ is said to be regular in $\mathcal{M}$ at P if its convergence to $\psi(P)$ is locally uniform (van der Vaart, 2000).
5. The Hilbert space of all mean-zero scalar functions is the L^2 space. For a precise definition of Hilbert spaces see Luenberger (1997).

now construct the class of all influence functions, denoted by $\mathcal{U}$, as $\mathcal{U} = U_\psi(Z) + \Lambda^\perp$. Upon knowing a single IF $U_\psi(Z)$ and the tangent space orthogonal complement $\Lambda^\perp$, we can obtain the class of all possible RAL estimators that admit the CAN property.

Out of all the influence functions in $\mathcal{U}$ there exists a unique one which lies in the tangent space Λ, and which yields the most efficient RAL estimator by recovering the *semiparametric efficiency bound*. This efficient influence function can be obtained by projecting any influence function, call it U^*_ψ, onto the tangent space Λ. This operation is denoted by $U^{\text{eff}_\psi} = \pi[U^*_\psi \mid \Lambda]$, where U^{eff}_ψ denotes the efficient IF.

On the other hand, if the tangent space contains the entire Hilbert space, i.e., $\Lambda = \text{H}$, then the statistical model $\mathcal{M}$ is called a *nonparametric* model. In a nonparametric model, we only have one influence function since $\Lambda^\perp = \{0\}$. This unique influence function can be obtained via Eq. 36 and corresponds to the efficient influence function U^{eff}_ψ (the unique element in the tangent space Λ) in the nonparametric model $\mathcal{M}$. For a detailed description of the concepts outlined here, please refer to Tsiatis (2007); Bickel et al. (1993).

D.1 Intuitions Regarding the Nonparametric IF in Primal Fixability

Given a post treatment variable V_i and its conditional density $p(V_i \mid \text{mp}_{\mathcal{G}}(V_i))$ in the identified functional of $\psi(t)$ in Eq. 14, there is a corresponding term in the influence function U_{ψ_t} in Theorem 9 of the form

$$f_1(\prec V_i) \times \Big(f_2(\preceq V_i) - \sum_{V_i} f_2(\preceq V_i) \times p\big(V_i \mid \text{mp}_{\mathcal{G}}(V_i)\big)\Big), \tag{37}$$

where $f_1(\prec V_i)$ denotes a function of variables that come before V_i in the topological order, a.k.a history/past of V_i. Similarly, $f_2(\preceq V_i)$ is a function of past of V_i and including V_i itself. $f_1(\prec V_i)$ is defined as follows,

$$f_1(\prec V_i) = \begin{cases} \dfrac{\mathbb{I}(T=t)}{\prod_{L_i \prec V_i} \; p(L_i \mid \text{mp}_{\mathcal{G}}(L_i))}, & \text{if } V_i \in \mathbb{M} \\[2ex] \dfrac{\prod_{M_i \prec V_i} p(M_i \mid \text{mp}_{\mathcal{G}}(M_i))|_{T=t}}{\prod_{M_i \prec V_i} p(M_i \mid \text{mp}_{\mathcal{G}}(M_i))}, & \text{if } V_i \in \mathbb{L} \end{cases} \tag{38}$$

Interestingly, these weights resemble the ones in ψ_{primal} and ψ_{dual} that we introduced in Lemmas 4 and 6, if the target were the counterfactual mean $\mathbb{E}[V_i(t)]$. That is,

$$\psi_{v_i,\text{primal}} = \mathbb{E}\Big[\frac{\mathbb{I}(T=t)}{\prod_{L_i \prec V_i} \; p(L_i \mid \text{mp}_{\mathcal{G}}(L_i))} \times \sum_T \prod_{L_i \prec V_i} p(L_i \mid \text{mp}_{\mathcal{G}}(L_i)) \times V_i \Big]$$

$$\psi_{v_i,\text{dual}} = \mathbb{E}\Big[\frac{\prod_{M_i \prec V_i} p(M_i \mid \text{mp}_{\mathcal{G}}(M_i))|_{T=t}}{\prod_{M_i \prec V_i} p(M_i \mid \text{mp}_{\mathcal{G}}(M_i))} \times V_i \Big]$$

However, to calculate the effect of T on Y, we do not need to worry about the effect of T on intermediate variables V_i. In Lemma 10, we show that the influence function U_{ψ_t} in Theorem 9 cleverly uses the information in these intermediate primal and dual estimators and yields a doubly robust estimator for ψ_t.

Appendix E. Proofs

Throughout this section, we use the following sets that were defined in Section 3.3. We partition the set of nodes V into three disjoint sets: $V = \{\mathbb{C}, \mathbb{L}, \mathbb{M}\}$, where

$$\begin{aligned}\mathbb{C} &= \{C_i \in V \mid C_i \prec T\},\\ \mathbb{L} &= \{L_i \in V \mid L_i \in D_T, L_i \succeq T\},\\ \mathbb{M} &= \{M_i \in V \mid M_i \notin \mathbb{C} \cup \mathbb{L}\}.\end{aligned}$$

Lemma 1 (Identifiying functional when T is a-fixable)

Proof By the pre-condition that $\text{dis}_{\mathcal{G}}(T) \cap \text{de}_{\mathcal{G}}(T) = \{T\}$, there exists a valid topological ordering on vertices V such that T appears last among the members of its district. Under such an ordering $\text{mp}_{\mathcal{G}}(T) = \text{mb}_{\mathcal{G}}(T)$. That is, under a-fixability $\text{mp}_{\mathcal{G}}(T) = \text{dis}_{\mathcal{G}}(T) \cup \text{pa}_{\mathcal{G}}(\text{dis}_{\mathcal{G}}(T)) \setminus T$. Note that $\text{mp}_{\mathcal{G}}(T)$ does not contain any descendants of T so using it as an adjustment set does not block any directed (causal) paths from T to Y.

We now show that conditioning on the set $\text{mp}_{\mathcal{G}}(T) = \text{mb}_{\mathcal{G}}(T)$ blocks all back-door paths from T to Y.

- Back-door paths of the form $T \leftarrow P \cdots Y$: Since $P \in \text{pa}_{\mathcal{G}}(T)$ and $\text{pa}_{\mathcal{G}}(T) \subseteq \text{mp}_{\mathcal{G}}(T)$, we condition on all parents of T resulting in these paths being blocked.
- Back-door paths of the form $T \leftrightarrow S \rightarrow \cdots \rightarrow Y$: Since $S \in \text{dis}_{\mathcal{G}}(T)$ and $\text{dis}_{\mathcal{G}}(T) \subseteq \text{mp}_{\mathcal{G}}(T)$, we block these paths by conditioning on all such S.
- Finally collider paths of the form $T \leftrightarrow V_1 \leftrightarrow \cdots \leftrightarrow V_p \leftarrow X \cdots Y$ and $T \leftrightarrow V_1 \leftrightarrow \cdots \leftrightarrow V_p \leftrightarrow X \cdots Y$: These paths are normally blocked due to the presence of the colliders but conditioning on $V_i \in \text{dis}_{\mathcal{G}}(T)$ opens up these colliders until finally, the path is blocked by conditioning on some vertex X that is either in $\text{pa}_{\mathcal{G}}(\text{dis}_{\mathcal{G}}(T))$ or in $\text{dis}_{\mathcal{G}}(T)$, both of which are contained in the Markov pillow of T.

As $\text{mp}_{\mathcal{G}}(T)$ does not block any causal paths from T to Y and blocks all back-door paths, it forms a valid adjustment set for the effect of T on Y (Pearl, 2009). Hence, the target $\psi(t)$ is identified as $\psi(t) = \mathbb{E}[\mathbb{E}[Y \mid T = t, \text{mp}_{\mathcal{G}}(T)]]$. ■

Theorem 2 (Nonparametric influence function of gAIPW)

Proof The proof here is very similar to how one would prove regular AIPW and is provided for the sake of completeness. The parameter of interest is,

$$\psi_\kappa(t) = \sum_{Y, \text{mp}_{\mathcal{G}}(T)} Y \times p_\kappa(Y \mid T = t, \text{mp}_{\mathcal{G}}(T)) \times p_\kappa(\text{mp}_{\mathcal{G}}(T)).$$

Using Eq. 36, we have $\frac{d}{d\kappa}\psi_\kappa(t)\Big|_{\kappa=0} = \mathbb{E}[U_{\psi_t} \times S_{\eta_0}(V)]$. Therefore,

$$\frac{d}{d\kappa}\psi_\kappa(t)\Big|_{\kappa=0} = \sum_{Y,\mathrm{mp}_{\mathcal{G}}(T)} Y \times \frac{d}{d\kappa}(p_\kappa(Y \mid T=t, \mathrm{mp}_{\mathcal{G}}(T))) \times p_\kappa(\mathrm{mp}_{\mathcal{G}}(T)) + \sum_{Y,\mathrm{mp}_{\mathcal{G}}(T)} Y \times p_\kappa(Y \mid T=t, \mathrm{mp}_{\mathcal{G}}(T)) \times \frac{d}{d\kappa}(p_\kappa(\mathrm{mp}_{\mathcal{G}}(T)))$$

$$\begin{aligned}
\text{First Term}: &\sum_{Y,\mathrm{mp}_{\mathcal{G}}(T)} Y \times p(Y \mid T=t, \mathrm{mp}_{\mathcal{G}}(T)) \times S(Y \mid T=t, \mathrm{mp}_{\mathcal{G}}(T)) \times p(\mathrm{mp}_{\mathcal{G}}(T)) \\
&= \sum_{Y,\mathrm{mp}_{\mathcal{G}}(T),T} \frac{\mathbb{I}(T=t)}{p(T \mid \mathrm{mp}_{\mathcal{G}}(T))} \times Y \times S(Y \mid T, \mathrm{mp}_{\mathcal{G}}(T)) \times p(Y, T, \mathrm{mp}_{\mathcal{G}}(T)) \\
&= \mathbb{E}\Big[\frac{\mathbb{I}(T=t)}{p(T \mid \mathrm{mp}_{\mathcal{G}}(T))} \times Y \times S(Y \mid T, \mathrm{mp}_{\mathcal{G}}(T))\Big] \\
&= \mathbb{E}\Big[\frac{\mathbb{I}(T=t)}{p(T \mid \mathrm{mp}_{\mathcal{G}}(T))} \times \Big(Y - \mathbb{E}[Y \mid T, \mathrm{mp}_{\mathcal{G}}(T)]\Big) \times S(Y \mid T, \mathrm{mp}_{\mathcal{G}}(T))\Big] \\
&= \mathbb{E}\Big[\frac{\mathbb{I}(T=t)}{p(T \mid \mathrm{mp}_{\mathcal{G}}(T))} \times \Big(Y - \mathbb{E}[Y \mid T, \mathrm{mp}_{\mathcal{G}}(T)]\Big) \times S(Y, T, \mathrm{mp}_{\mathcal{G}}(T))\Big].
\end{aligned}$$

The third and fourth equality hold since

$$\begin{aligned}
&\mathbb{E}\Big[\frac{\mathbb{I}(T=t)}{p(T \mid \mathrm{mp}_{\mathcal{G}}(T))} \times \mathbb{E}[Y \mid T, \mathrm{mp}_{\mathcal{G}}(T)] \times S(Y \mid T, \mathrm{mp}_{\mathcal{G}}(T))\Big] = 0, \text{ and} \\
&\mathbb{E}\Big[\frac{\mathbb{I}(T=t)}{p(T \mid \mathrm{mp}_{\mathcal{G}}(T))} \times \big(Y - \mathbb{E}[Y \mid T, \mathrm{mp}_{\mathcal{G}}(T)]\big) \times S(T, \mathrm{mp}_{\mathcal{G}}(T))\Big] = 0,
\end{aligned}$$

respectively.

$$\begin{aligned}
\text{Second Term}: &\sum_{\mathrm{mp}_{\mathcal{G}}(T)} \mathbb{E}[Y \mid T=t, \mathrm{mp}_{\mathcal{G}}(T)] \times S(\mathrm{mp}_{\mathcal{G}}(T)) \times p(\mathrm{mp}_{\mathcal{G}}(T)) \\
&= \mathbb{E}\big[\mathbb{E}[Y \mid T=t, \mathrm{mp}_{\mathcal{G}}(T)] \times S(\mathrm{mp}_{\mathcal{G}}(T))\big] \\
&= \mathbb{E}\Big[\Big(\mathbb{E}[Y \mid T=t, \mathrm{mp}_{\mathcal{G}}(T)] - \psi(t)\Big) \times S(\mathrm{mp}_{\mathcal{G}}(T))\Big] \\
&= \mathbb{E}\Big[\Big(\mathbb{E}[Y \mid T=t, \mathrm{mp}_{\mathcal{G}}(T)] - \psi(t)\Big) \times S(Y, T, \mathrm{mp}_{\mathcal{G}}(T))\Big].
\end{aligned}$$

The second and third equality hold since $\mathbb{E}[S(\mathrm{mp}_{\mathcal{G}}(T))] = 0$ and $\mathbb{E}\big[\big(\mathbb{E}[Y \mid T=t, \mathrm{mp}_{\mathcal{G}}(T)] - \psi(t)\big) \times S(Y, T \mid \mathrm{mp}_{\mathcal{G}}(T))\big] = 0$, respectively. Combining the first and second terms together yields the nonparametric IF for $\psi(t)$, which we termed gAIPW,

$$U_{\psi_t} = \frac{\mathbb{I}(T=t)}{p(T \mid \mathrm{mp}_{\mathcal{G}}(T))} \times \Big(Y - \mathbb{E}[Y \mid T, \mathrm{mp}_{\mathcal{G}}(T)]\Big) + \mathbb{E}[Y \mid T=t, \mathrm{mp}_{\mathcal{G}}(T)] - \psi(t).$$

■

Lemma 3 (Double robustness of gAIPW)

Proof We need to show under correct specification of either $\mathbb{E}[Y \mid T, \mathrm{mp}_{\mathcal{G}}(T)]$ or $p(T \mid \mathrm{mp}_{\mathcal{G}}(T))$, the estimating equation that we obtain as a result of Theorem 2 remains zero. We break this down into two scenarios.

Scenario 1. Assume $p(T \mid \mathrm{mp}_{\mathcal{G}}(T))$ is correctly specified, and let $\mathbb{E}^*[Y \mid T, \mathrm{mp}_{\mathcal{G}}(T)]$ denote the misspecified model for $\mathbb{E}[Y \mid T, \mathrm{mp}_{\mathcal{G}}(T)]$. Therefore,

$$
\begin{aligned}
&\mathbb{E}\Big[\frac{\mathbb{I}(T=t)}{p(T \mid \mathrm{mp}_{\mathcal{G}}(T))} \times \Big(Y - \mathbb{E}^*[Y \mid T, \mathrm{mp}_{\mathcal{G}}(T)]\Big) + \mathbb{E}^*[Y \mid T=t, \mathrm{mp}_{\mathcal{G}}(T)] - \psi(t)\Big] \\
&= \mathbb{E}\Big[\frac{\mathbb{I}(T=t)}{p(T \mid \mathrm{mp}_{\mathcal{G}}(T))} \times Y\Big] - \psi(t) + \mathbb{E}\Big[\Big(1 - \frac{\mathbb{I}(T=t)}{p(T \mid \mathrm{mp}_{\mathcal{G}}(T))}\Big) \times \mathbb{E}^*[Y \mid T=t, \mathrm{mp}_{\mathcal{G}}(T)]\Big] \\
&= 0 + \mathbb{E}\Big[\mathbb{E}\Big[\big(1 - \frac{\mathbb{I}(T=t)}{p(T \mid \mathrm{mp}_{\mathcal{G}}(T))}\big) \times \mathbb{E}^*[Y \mid T=t, \mathrm{mp}_{\mathcal{G}}(T)] \;\Big|\; \mathrm{mp}_{\mathcal{G}}(T)\Big]\Big] \\
&= \mathbb{E}\Big[\mathbb{E}^*[Y \mid T, \mathrm{mp}_{\mathcal{G}}(T)] | \, \mathrm{mp}_{\mathcal{G}}(T)] \times \Big(1 - \frac{\mathbb{E}\big[\mathbb{I}(T=t) \mid \mathrm{mp}_{\mathcal{G}}(T)\big]}{p(T=t \mid \mathrm{mp}_{\mathcal{G}}(T))}\Big)\Big] \\
&= 0.
\end{aligned}
$$

The last equality holds since $\mathbb{E}\big[\mathbb{I}(T=t) \mid \mathrm{mp}_{\mathcal{G}}(T)\big] = p(T=t \mid \mathrm{mp}_{\mathcal{G}}(T))$, therefore the above expectation evaluates to 0, despite the incorrectly specified model.

Scenario 2. Assume $\mathbb{E}[Y \mid T, \mathrm{mp}_{\mathcal{G}}(T)]$ is correctly specified, and let $p^*(T \mid \mathrm{mp}_{\mathcal{G}}(T))$ denote the miss-specified model for $p(T \mid \mathrm{mp}_{\mathcal{G}}(T))$. Therefore,

$$
\begin{aligned}
&\mathbb{E}\Big[\frac{\mathbb{I}(T=t)}{p^*(T \mid \mathrm{mp}_{\mathcal{G}}(T))} \times \Big(Y - \mathbb{E}[Y \mid T, \mathrm{mp}_{\mathcal{G}}(T)]\Big) + \mathbb{E}[Y \mid T=t, \mathrm{mp}_{\mathcal{G}}(T)] - \psi(t)\Big] \\
&= \mathbb{E}\Big[\frac{\mathbb{I}(T=t)}{p^*(T \mid \mathrm{mp}_{\mathcal{G}}(T))} \times \Big(Y - \mathbb{E}[Y \mid T, \mathrm{mp}_{\mathcal{G}}(T)]\Big)\Big] + \mathbb{E}\Big[\mathbb{E}[Y \mid T=t, \mathrm{mp}_{\mathcal{G}}(T)]\Big] - \psi(t) \\
&= \mathbb{E}\Big[\mathbb{E}\Big[\frac{\mathbb{I}(T=t)}{p^*(T \mid \mathrm{mp}_{\mathcal{G}}(T))} \times \Big(Y - \mathbb{E}[Y \mid T, \mathrm{mp}_{\mathcal{G}}(T)]\Big) \;\Big|\; T, \mathrm{mp}_{\mathcal{G}}(T)\Big] + 0 \\
&= \mathbb{E}\Big[\frac{\mathbb{I}(T=t)}{p^*(T \mid \mathrm{mp}_{\mathcal{G}}(T))} \times \Big(\mathbb{E}[Y \mid T, \mathrm{mp}_{\mathcal{G}}(T)] - \mathbb{E}[Y \mid T, \mathrm{mp}_{\mathcal{G}}(T)]\Big)\Big] \\
&= 0.
\end{aligned}
$$

We have shown that the estimating equation for $\psi(t)$ evaluates to 0 with respect to the observed data distribution, as long as either $\mathbb{E}[Y \mid T, \mathrm{mp}_{\mathcal{G}}(T)]$ or $p(T \mid \mathrm{mp}_{\mathcal{G}}(T))$ is specified correctly. This establishes double robustness of gAIPW. ■

Lemma 4 (Primal IPW)

Proof Our goal is to demonstrate that the primal IPW formulation is equivalent to the identifying functional of the target parameter $\psi(t)$ shown in Eq. 10 and restated below.

$$
\psi(t) = \sum_{V \setminus T} \prod_{V_i \in V \setminus D_T} p(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i))\Big|_{T=t} \times \sum_{T} \prod_{D_i \in D_T} p(D_i \mid \mathrm{mp}_{\mathcal{G}}(D_i)) \times Y.
$$

The primal IPW formulation for the target $\psi(t)$ is,

$$\mathbb{E}[\beta_{\text{primal}}(t)] \equiv \mathbb{E}\left[\frac{\mathbb{I}(T=t)}{q_{D_T}(T \mid \text{mb}_{\mathcal{G}}(T))} \times Y\right]$$

where $q_{D_T}(D_T \mid \text{pa}_{\mathcal{G}}(D_T)) = \prod_{V_i \in D_T} p(V_i \mid \text{mp}_{\mathcal{G}}(V_i))$, and

$$\begin{aligned}
q_{D_T}(T \mid \text{mb}_{\mathcal{G}}(T)) = q_{D_T}(T \mid D_T \cup \text{pa}_{\mathcal{G}}(D_T) \setminus T) &= \frac{q_{D_T}(D_T \mid \text{pa}_{\mathcal{G}}(D_T))}{q_{D_T}(D_T \setminus T \mid \text{pa}_{\mathcal{G}}(D_T))} \\
&= \frac{q_{D_T}(D_T \mid \text{pa}_{\mathcal{G}}(D_T))}{\sum_T q_{D_T}(D_T \mid \text{pa}_{\mathcal{G}}(D_T))} = \frac{\prod_{V_i \in D_T} p(V_i \mid \text{mp}_{\mathcal{G}}(V_i))}{\sum_T \prod_{V_i \in D_T} p(V_i \mid \text{mp}_{\mathcal{G}}(V_i))} \\
&= \frac{\prod_{V_i \in \mathbb{L}} p(V_i \mid \text{mp}_{\mathcal{G}}(V_i))}{\sum_T \prod_{V_i \in \mathbb{L}} p(V_i \mid \text{mp}_{\mathcal{G}}(V_i))}.
\end{aligned}$$

The last equality holds because the conditional densities of $V_i \in \mathbb{C}$, does not depend on T, and they cancel out from the numerator and denominator. Therefore, product in the ratio is over the variables in $D_T \cap \{\succeq T\}$ which we have denoted by $\mathbb{L}$. Therefore,

$$\begin{aligned}
\mathbb{E}[\beta_{\text{primal}}(t)] &= \mathbb{E}\left[\mathbb{I}(T=t) \times \frac{\sum_T \prod_{D_i \in \mathbb{L}} p(D_i \mid \text{mp}_{\mathcal{G}}(D_i))}{\prod_{D_i \in \mathbb{L}} p(D_i \mid \text{mp}_{\mathcal{G}}(D_i))} \times Y\right] \\
&= \sum_V \prod_{V_i \in V} p(V_i \mid \text{mp}_{\mathcal{G}}(V_i)) \times \mathbb{I}(T=t) \times \frac{\sum_T \prod_{D_i \in \mathbb{L}} p(D_i \mid \text{mp}_{\mathcal{G}}(D_i))}{\prod_{D_i \in \mathbb{L}} p(D_i \mid \text{mp}_{\mathcal{G}}(D_i))} \times Y \\
&= \sum_V \mathbb{I}(T=t) \times \prod_{V_i \in V \setminus \mathbb{L}} p(V_i \mid \text{mp}_{\mathcal{G}}(V_i)) \\
&\qquad \times \prod_{D_i \in \mathbb{L}} p(D_i \mid \text{mp}_{\mathcal{G}}(D_i)) \times \frac{\sum_T \prod_{D_i \in \mathbb{L}} p(D_i \mid \text{mp}_{\mathcal{G}}(D_i))}{\prod_{D_i \in \mathbb{L}} p(D_i \mid \text{mp}_{\mathcal{G}}(D_i))} \times Y \\
&= \sum_V \mathbb{I}(T=t) \times \prod_{V_i \in V \setminus \mathbb{L}} p(V_i \mid \text{mp}_{\mathcal{G}}(V_i)) \times \sum_T \prod_{D_i \in \mathbb{L}} p(D_i \mid \text{mp}_{\mathcal{G}}(D_i)) \times Y.
\end{aligned}$$

In the second equality, we evaluated the outer expectation with respect to the joint $p(V)$. In the third equality, we partitioned the joint into factors for the set $\mathbb{L}$ and factors for $V \setminus \mathbb{L}$. In the fourth equality, we canceled out the the factors involved in the denominator of the primal IPW with the corresponding terms in the joint.

We can then move the conditional factors of pre-treatment variables in the district of T past the summation over T as these factors are not functions of T. Finally, we evaluate the indicator function, concluding the proof. That is,

$$\begin{aligned}
\psi_{\text{primal}} &= \sum_V \mathbb{I}(T=t) \times \prod_{V_i \in V \setminus D_T} p(V_i \mid \text{mp}_{\mathcal{G}}(V_i)) \times \sum_T \prod_{D_i \in D_T} p(D_i \mid \text{mp}_{\mathcal{G}}(D_i)) \times Y \\
&= \sum_{V \setminus T} \prod_{V_i \in V \setminus D_T} p(V_i \mid \text{mp}_{\mathcal{G}}(V_i))\Big|_{T=t} \times \sum_T \prod_{D_i \in D_T} p(D_i \mid \text{mp}_{\mathcal{G}}(D_i)) \times Y = \psi(t)
\end{aligned}$$

■

Lemma 6 (Dual IPW)

Proof The proof strategy is similar to the one used for the primal IPW. The dual IPW formulation for the target $\psi(t)$ is,

$$\begin{aligned}
\mathbb{E}[\beta_{\text{dual}}(t)] &= \mathbb{E}\Big[\frac{\prod_{M_i \in \text{mp}_{\mathcal{G}}^{-1}(T)} \; p(M_i \mid \text{mp}_{\mathcal{G}}(M_i)) \mid_{T=t}}{\prod_{M_i \in \text{mp}_{\mathcal{G}}^{-1}(T)} \; p(M_i \mid \text{mp}_{\mathcal{G}}(M_i))} \times Y\Big] \\
&= \sum_{V} \prod_{V_i \in V} p(V_i \mid \text{mp}_{\mathcal{G}}(V_i)) \times \frac{\prod_{M_i \in \text{mp}_{\mathcal{G}}^{-1}(T)} \; p(M_i \mid \text{mp}_{\mathcal{G}}(M_i)) \mid_{T=t}}{\prod_{M_i \in \text{mp}_{\mathcal{G}}^{-1}(T)} \; p(M_i \mid \text{mp}_{\mathcal{G}}(M_i))} \times Y \\
&= \sum_{V} \prod_{V_i \in V \setminus \text{mp}_{\mathcal{G}}^{-1}(T)} p(V_i \mid \text{mp}_{\mathcal{G}}(V_i)) \\
&\qquad \times \prod_{M_i \in \text{mp}_{\mathcal{G}}^{-1}(T)} p(M_i \mid \text{mp}_{\mathcal{G}}(M_i)) \times \frac{\prod_{M_i \in \text{mp}_{\mathcal{G}}^{-1}(T)} \; p(M_i \mid \text{mp}_{\mathcal{G}}(M_i)) \mid_{T=t}}{\prod_{M_i \in \text{mp}_{\mathcal{G}}^{-1}(T)} \; p(M_i \mid \text{mp}_{\mathcal{G}}(M_i))} \times Y \\
&= \sum_{V} \prod_{V_i \in V \setminus \text{mp}_{\mathcal{G}}^{-1}(T)} p(V_i \mid \text{mp}_{\mathcal{G}}(V_i)) \times \prod_{M_i \in \text{mp}_{\mathcal{G}}^{-1}(T)} p(M_i \mid \text{mp}_{\mathcal{G}}(M_i)) \mid_{T=t} \times Y \\
&= \sum_{V \setminus T} \prod_{V_i \in V \setminus \{\text{mp}_{\mathcal{G}}^{-1}(T) \cup D_T\}} p(V_i \mid \text{mp}_{\mathcal{G}}(V_i)) \times \prod_{M_i \in \text{mp}_{\mathcal{G}}^{-1}(T)} p(M_i \mid \text{mp}_{\mathcal{G}}(M_i)) \mid_{T=t} \\
&\qquad \times \sum_{T} \prod_{D_T} p(D_i \mid \text{mp}_{\mathcal{G}}(D_i)) \times Y.
\end{aligned}$$

In the above derivation, we first evaluated the outer expectation with respect to the joint $p(V)$. We then partitioned the joint into factors corresponding to $\text{mp}_{\mathcal{G}}^{-1}(T)$ and $V \setminus \text{mp}_{\mathcal{G}}^{-1}(T)$. The factors involved in the denominator of the dual IPW then canceled out with the corresponding terms in the joint. The last equality holds because by the definition of the inverse Markov pillow, $\text{mp}_{\mathcal{G}}^{-1}(T)$ contains all variables not in the district of T such that T is a member of its Markov pillow. In the above expression, factors corresponding to the inverse Markov pillow of T are evaluated at $T = t$. Consequently, the only factors above that are still functions of T are the ones corresponding to the district of T. This allows us to push the summation over T.

Finally, since the summation over T will prevent factors within the district of T from being evaluated at $T = t$, we can simply apply the evaluation to the entire functional and merge the sets not involved in the district of T above. That is,

$$\psi_{\text{dual}} = \sum_{V \setminus T} \prod_{V_i \in V \setminus D_T} p(V_i \mid \text{mp}_{\mathcal{G}}(V_i)) \times \sum_{T} \prod_{D_i \in D_T} p(D_i \mid \text{mp}_{\mathcal{G}}(D_i)) \times Y \Big|_{T=t} = \psi(t).$$

■

Lemma 8 (Variational independence of primal and dual IPW estimators)

Proof Consider the topological factorization of the observed distribution $p(V)$ for the ADMG as shown in Eq. 8.

$$p(V) = \prod_{V_i \in V} p(V_i \mid \text{mp}_{\mathcal{G}}(V_i)).$$

Note by definition, the inverse Markov pillow of T does not contain elements in the district of T, i.e., $\mathrm{mp}_{\mathcal{G}}^{-1}(T) \cap D_T = \emptyset$. Thus, we can partition V into three disjoint sets as follows:

$$\mathbb{L} = D_T \cap \{\succeq T\}, \qquad \mathbb{M}^* = \mathrm{mp}_{\mathcal{G}}^{-1}(T), \qquad \mathbb{R} = V \setminus (\mathbb{L} \cup \mathbb{M}^*)$$

The set $\mathbb{L}$ is the same as what we defined earlier at the beginning of this proof section. $\mathbb{M}^*$ is a subset of $\mathbb{M}$, and the remainder terms $\mathbb{R} = \mathbb{C} \cup \{\mathbb{M} \setminus \mathbb{M}^*\}$. The topological factorization of the observed joint can then be restated as,

$$p(V) = \prod_{R_i \in \mathbb{R}} p(R_i \mid \mathrm{mp}_{\mathcal{G}}(R_i)) \prod_{M_i \in \mathbb{M}^*} p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i)) \prod_{L_i \in \mathbb{L}} p(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i)).$$

It is then clear from the above factorization that the components of the primal IPW estimator which sit in $\mathbb{L}$, and the components of the dual IPW estimator which sit in $\mathbb{M}$, form congenial and variationally independent pieces of the joint distribution $p(V)$. ■

Theorem 9 (Nonparametric influence function of augmented primal IPW)

Proof The target parameter is identified via the following function of the observed data,

$$\psi_\kappa(t) = \sum_{V \setminus T} Y \times \prod_{M_i \in \mathbb{M}} p_\kappa(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))|_{T=t} \times \sum_{T} \prod_{L_i \in \mathbb{L}} p_\kappa(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i)) \times p_\kappa(\mathbb{C}), \tag{39}$$

and according to Eq. 36, $\frac{d}{d\kappa}\psi_\kappa(t)\big|_{\kappa=0} = \mathbb{E}\big[U_{\psi_t} \times S_{\eta_0}(V)\big]$. Therefore,

$$\begin{aligned}
\frac{d}{d\kappa}\psi_\kappa(t) &= \frac{d}{d\kappa}\Big\{\sum_{V \setminus T} Y \times \prod_{M_i \in \mathbb{M}} p_\kappa(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))|_{T=t} \times \sum_{T} \prod_{L_i \in \mathbb{L} \setminus T} p_\kappa(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i)) \times p_\kappa(T, \mathbb{C})\Big\} \\
&= \sum_{V \setminus T} Y \times \frac{d}{d\kappa}\Big\{\prod_{M_i \in \mathbb{M}} p_\kappa(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))|_{T=t}\Big\} \times \sum_{T} \prod_{L_i \in \mathbb{L} \setminus T} p_\kappa(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i)) \times p_\kappa(T, \mathbb{C}) \quad \text{(1st Term)} \\
&+ \sum_{V \setminus T} Y \times \prod_{M_i \in \mathbb{M}} p_\kappa(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))|_{T=t} \times \sum_{T} \frac{d}{d\kappa}\Big\{\prod_{L_i \in \mathbb{L} \setminus T} p_\kappa(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i))\Big\} \times p_\kappa(T, \mathbb{C}) \quad \text{(2nd Term)} \\
&+ \sum_{V \setminus T} Y \times \prod_{M_i \in \mathbb{M}} p_\kappa(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))|_{T=t} \times \sum_{T} \prod_{L_i \in \mathbb{L} \setminus T} p_\kappa(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i)) \times \frac{d}{d\kappa}\Big\{p_\kappa(T, \mathbb{C})\Big\}. \quad \text{(3rd Term)}
\end{aligned}$$

First Term: The contribution of the first term to the final IF is made of individual contributions of the elements in $\mathbb{M}$. Since the derivation is similar, we only derive it for an element $M_j \in \mathbb{M}$.

$$\begin{aligned}
&\sum_{V \setminus T} Y \times \prod_{M_i \in \{\prec M_j\} \cap \mathbb{M}} p_\kappa(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))|_{T=t} \times \frac{d}{d\kappa}\Big\{p_\kappa(M_j \mid \mathrm{mp}_{\mathcal{G}}(M_j))|_{T=t}\Big\} \\
&\qquad \times \prod_{M_i \in \{\succ M_j\} \cap \mathbb{M}} p_\kappa(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))|_{T=t} \times \sum_{T} \prod_{L_i \in \mathbb{L}} p_\kappa(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i)) \times p_\kappa(\mathbb{C}) \\
&\overset{(1)}{=} \sum_{V \setminus \{T, \{\preceq M_j\}\}} \prod_{M_i \in \{\prec M_j\} \cap \mathbb{M}} p_\kappa(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))|_{T=t} \times \frac{d}{d\kappa}\Big\{p_\kappa(M_j \mid \mathrm{mp}_{\mathcal{G}}(M_j))|_{T=t}\Big\} \\
&\qquad \times \sum_{T \cup \{\succ M_j\}} Y \times \prod_{V_i \in \mathbb{L} \cup \{\{\succ M_j\} \cap \mathbb{M}\}} p_\kappa(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i))\Big|_{T=t \text{ if } V_i \in \mathbb{M}} \times p_\kappa(\mathbb{C})
\end{aligned}$$

$$
\begin{aligned}
&\overset{(2)}{=} \sum_{\preceq M_j} \frac{\mathbb{I}(T=t)}{\prod_{L_i \prec M_j} p(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i))} \times S(M_j \mid \mathrm{mp}_{\mathcal{G}}(M_j)) \times \prod_{V_i \in \{\preceq M_j\}} p(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i)) \\
&\qquad \times \sum_{T \cup \{\succ M_j\}} Y \times \prod_{V_i \in \mathbb{L} \cup \{\{\succ M_j\} \cap \mathbb{M}\}} p(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i))\Big|_{T=t \text{ if } V_i \in \mathbb{M}} \\
&\overset{(3)}{=} \mathbb{E}\Bigg[\frac{\mathbb{I}(T=t)}{\prod_{L_i \prec M_j} p(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i))} \times \underbrace{\sum_{T \cup \{\succ M_j\}} Y \times \prod_{V_i \in \mathbb{L} \cup \{\{\succ M_j\} \cap \mathbb{M}\}} p(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i))\Big|_{T=t \text{ if } V_i \in \mathbb{M}}}_{:= f(\preceq M_j)} \\
&\qquad \times S(M_j \mid \mathrm{mp}_{\mathcal{G}}(M_j))\Bigg] \\
&\overset{(4)}{=} \mathbb{E}\Bigg[\frac{\mathbb{I}(T=t)}{\prod_{L_i \prec M_j} p(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i))} \times \Big(f(\preceq M_j) - \sum_{M_j} f(\preceq M_j) \times p(M_j \mid \mathrm{mp}_{\mathcal{G}}(M_j))\Big) \times S(M_j \mid \mathrm{mp}_{\mathcal{G}}(M_j))\Bigg] \\
&\overset{(5)}{=} \mathbb{E}\Bigg[\frac{\mathbb{I}(T=t)}{\prod_{L_i \prec M_j} p(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i))} \times \Big(f(\preceq M_j) - \sum_{M_j} f(\preceq M_j) \times p(M_j \mid \mathrm{mp}_{\mathcal{G}}(M_j))\Big) \times S(V)\Bigg]
\end{aligned}
$$

The first equality follows from the fact that terms corresponding to $M_i \in \{\prec M_j\}$ are not functions of elements in $\{\succ M_j\}$ and of Y. The second equality follows by term grouping, the definition of conditional scores, and term cancellation. The third equality is by definition of joint expectation. The fourth and fifth equalities are implied by the fact that conditional scores have expected value of 0 (given their conditioning set). Therefore, the contribution of $M_j \in \mathbb{M}$ is the following:

$$
\begin{aligned}
\frac{\mathbb{I}(T=t)}{\prod_{L_i \prec M_j} p(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i))} \times \Big(&\sum_{T \cup \{\succ M_j\}} Y \times \prod_{V_i \in \mathbb{L} \cup \{\{\succ M_j\} \cap \mathbb{M}\}} p(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i))\Big|_{T=t \text{ if } V_i \in \mathbb{M}} \\
&- \sum_{T \cup \{\succeq M_j\}} Y \times \prod_{V_i \in \mathbb{L} \cup \{\{\succeq M_j\} \cap \mathbb{M}\}} p(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i))\Big|_{T=t \text{ if } V_i \in \mathbb{M}} \Big).
\end{aligned}
$$

Second Term: The contribution of the second term to the final IF is made of individual contributions of the elements in $\mathbb{L} \setminus T$. Since the derivation is similar, we only derive it for an element $L_j \in \mathbb{L} \setminus T$.

$$
\begin{aligned}
&\sum_{V \setminus T} Y \times \prod_{M_i \in \mathbb{M}} p_\kappa(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))|_{T=t} \times \sum_{T} \Bigg\{ \prod_{L_i \in \{\prec L_j\} \cap \mathbb{L} \setminus T} p_\kappa(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i)) \\
&\qquad \times \frac{d}{d\kappa}\Big\{p_\kappa(L_j \mid \mathrm{mp}_{\mathcal{G}}(L_j))\Big\} \times \prod_{L_i \in \{\succ L_j\} \cap \mathbb{L} \setminus T} p_\kappa(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i))\Bigg\} \times p_\kappa(T, \mathbb{C}) \\
&\overset{(1)}{=} \sum_{V} Y \times \prod_{V_i \in \{\succ L_j\}} p_\kappa(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i))|_{T=t \text{ if } V_i \in \mathbb{M}} \times \frac{d}{d\kappa}\Big\{p_\kappa(L_j \mid \mathrm{mp}_{\mathcal{G}}(L_j))\Big\} \\
&\qquad \times \prod_{V_i \in \{\prec L_j\}} p_\kappa(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i))|_{T=t \text{ if } V_i \in \mathbb{M}} \\
&\overset{(2)}{=} \sum_{\preceq L_j} \underbrace{\sum_{\succ L_j} Y \times \prod_{V_i \in \{\succ L_j\}} p(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i))|_{T=t \text{ if } V_i \in \mathbb{M}}}_{f(\preceq L_j)} \times S(L_j \mid \mathrm{mp}_{\mathcal{G}}(L_j)) \\
&\qquad \times \prod_{V_i \in \{\preceq L_j\}} p(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i))|_{T=t \text{ if } V_i \in \mathbb{M}}
\end{aligned}
$$

$$
\begin{aligned}
&\overset{(3)}{=} \sum_{\preceq L_j} f(\preceq L_j) \times \frac{\prod_{M_i \in \mathbb{M} \cap \{\prec L_j\}} \; p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))\big|_{T=t}}{\prod_{M_i \in \mathbb{M} \cap \{\prec L_j\}} \; p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))} \times S(L_j \mid \mathrm{mp}_{\mathcal{G}}(L_j)) \times \prod_{V_i \in \{\preceq L_j\}} p(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i)) \\
&\overset{(4)}{=} \mathbb{E}\left[\frac{\prod_{M_i \in \mathbb{M}\{\cap \prec L_j\}} \; p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))\big|_{T=t}}{\prod_{M_i \in \mathbb{M} \cap \{\prec L_j\}} \; p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))} \times f(\preceq L_j) \times S(L_j \mid \mathrm{mp}_{\mathcal{G}}(L_j))\right] \\
&\overset{(5)}{=} \mathbb{E}\left[\frac{\prod_{M_i \in \mathbb{M} \cap \{\prec L_j\}} \; p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))\big|_{T=t}}{\prod_{M_i \in \mathbb{M} \cap \{\prec L_j\}} \; p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))} \times \Big(f(\preceq L_j) - \sum_{L_j} f(\preceq L_j) \times p(L_j \mid \mathrm{mp}_{\mathcal{G}}(L_j))\Big) \times S(L_j \mid \mathrm{mp}_{\mathcal{G}}(L_j))\right] \\
&\overset{(6)}{=} \mathbb{E}\left[\frac{\prod_{M_i \in \mathbb{M} \cap \{\prec L_j\}} \; p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))\big|_{T=t}}{\prod_{M_i \in \mathbb{M} \cap \{\prec L_j\}} \; p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))} \times \Big(f(\preceq L_j) - \sum_{L_j} f(\preceq L_j) \times p(L_j \mid \mathrm{mp}_{\mathcal{G}}(L_j))\Big) \times S(V)\right]
\end{aligned}
$$

The first equality follows from the fact that terms corresponding to $M_i \in \mathbb{M}$ are not functions of T, the fact that $\mathbb{C}, \mathbb{M}, \mathbb{L}$ partition V, and term grouping. The second equality is by definition of conditional scores. The third equality is by term cancellation. The fourth is by definition of joint expectations, the fifth and sixth equalities are implied by the fact that conditional scores have expected value of 0 (given their conditioning set). Therefore, the contribution of $L_j \in \mathbb{L} \setminus T$ is the following:

$$
\begin{aligned}
\frac{\prod_{M_i \in \mathbb{M} \cap \{\prec L_j\}} \; p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))\big|_{T=t}}{\prod_{M_i \in \mathbb{M} \cap \{\prec L_j\}} \; p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))} \times \Big(& \sum_{\succ L_j} Y \times \prod_{V_i \in \{\succ L_j\}} p(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i))\big|_{T=t \text{ if } V_i \in \mathbb{M}} \\
& - \sum_{\succeq L_j} Y \times \prod_{V_i \in \{\succeq L_j\}} p(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i))\big|_{T=t \text{ if } V_i \in \mathbb{M}} \Big).
\end{aligned}
$$

Third Term: The contribution of the last term to the final IF is as follows.

$$
\begin{aligned}
&\sum_{V \setminus T} Y \times \prod_{M_i \in \mathbb{M}} p_\kappa(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))|_{T=t} \times \sum_{T} \prod_{L_i \in \mathbb{L} \setminus T} p_\kappa(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i)) \times \frac{d}{d\kappa}\Big\{p_\kappa(T, \mathbb{C})\Big\} \\
&\overset{(1)}{=} \sum_{T, \mathbb{C}} \underbrace{\Bigg\{ \sum_{V \setminus T, \mathbb{C}} Y \times \prod_{M_i \in \mathbb{M}} p_\kappa(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))|_{T=t} \times \prod_{L_i \in \mathbb{L} \setminus T} p_\kappa(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i)) \Bigg\}}_{f(T, \mathbb{C})} \times \frac{d}{d\kappa}\Big\{p_\kappa(T, \mathbb{C})\Big\}. \\
&\overset{(2)}{=} \sum_{T, \mathbb{C}} f(T, \mathbb{C}) \times S(T, \mathbb{C}) \times p(T, \mathbb{C}) = \mathbb{E}\Big[f(T, \mathbb{C}) \times S(T, \mathbb{C})\Big] \\
&\overset{(3)}{=} \mathbb{E}\Big[\Big(f(T, \mathbb{C}) - \sum_{T, \mathbb{C}} f(T, \mathbb{C}) \times p(T, \mathbb{C})\Big) \times S(T, \mathbb{C})\Big] \\
&\overset{(4)}{=} \mathbb{E}\Big[\Big(f(T, \mathbb{C}) - \psi(t)\Big) \times S(V)\Big].
\end{aligned}
$$

The first equality is term grouping, the second is by definition of marginal scores, the third and fourth equalities are implied by the fact that scores have expected value 0. Therefore, the contribution of the last term is the following:

$$
\sum_{V \setminus \{T, \mathbb{C}\}} Y \times \prod_{M_i \in \mathbb{M}} p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))\big|_{T=t} \times \prod_{L_i \in \mathbb{L} \setminus T} p(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i)) - \psi(t).
$$

Putting all these together yields the final influence function. ■

Lemma 10 (Double robustness of augmented primal IPW)

Proof We need to show that under correct specification of conditional densities in either $\{p(M_i \mid \text{mp}_{\mathcal{G}}(M_i)), \forall M_i \in \mathbb{M}\}$ or $\{p(L_i \mid \text{mp}_{\mathcal{G}}(L_i)), \forall L_i \in \mathbb{L}\}$, the influence function in Theorem 9 remains to be mean zero. We break this down into two scenarios.

Scenario 1. Assume models in $\mathbb{L}$ are correctly specified, and let $p^*(M_i \mid \text{mp}_{\mathcal{G}}(M_i))$ denote the misspecified model for $p(M_i \mid \text{mp}_{\mathcal{G}}(M_i)), \forall M_i \in \mathbb{M}$. We note that for any $L_j \in \mathbb{L} \setminus T$, the following line in the IF evaluates to zero in expectation.

$$
\begin{aligned}
&\mathbb{E}\Bigg[\frac{\prod_{M_i \prec L_j} p^*(M_i \mid \text{mp}_{\mathcal{G}}(M_i))|_{T=t}}{\prod_{M_i \prec L_j} p^*(M_i \mid \text{mp}_{\mathcal{G}}(M_i))} \Bigg(\sum_{\succ L_j} Y \times \prod_{L_i \in \mathbb{L} \cap \{\succ L_j\}} p(L_i \mid \text{mp}_{\mathcal{G}}(L_i)) \times \prod_{M_i \in \mathbb{M} \cap \{\succ L_j\}} p^*(M_i \mid \text{mp}_{\mathcal{G}}(M_i))|_{T=t} \\
&\qquad - \sum_{\succeq L_j} Y \times \prod_{L_i \in \mathbb{L} \cap \{\succeq L_j\}} p(L_i \mid \text{mp}_{\mathcal{G}}(L_i)) \times \prod_{M_i \in \mathbb{M} \cap \{\succeq L_j\}} p^*(M_i \mid \text{mp}_{\mathcal{G}}(M_i))|_{T=t} \Bigg) \Bigg] \\
&\overset{(1)}{=} \sum_{\preceq L_j} \frac{\prod_{M_i \prec L_j} p^*(M_i \mid \text{mp}_{\mathcal{G}}(M_i))|_{T=t}}{\prod_{M_i \prec L_j} p^*(M_i \mid \text{mp}_{\mathcal{G}}(M_i))} \times \prod_{V_i \prec L_j} p(V_i \mid \text{mp}_{\mathcal{G}}(V_i)) \times p(L_j \times \text{mp}_{\mathcal{G}}(L_j)) \\
&\qquad \times \Bigg(\sum_{\succ L_j} Y \times \prod_{L_i \in \mathbb{L} \cap \{\succ L_j\}} p(L_i \mid \text{mp}_{\mathcal{G}}(L_i)) \times \prod_{M_i \in \mathbb{M} \cap \{\succ L_j\}} p^*(M_i \mid \text{mp}_{\mathcal{G}}(M_i))|_{T=t} \\
&\qquad\qquad - \sum_{\succeq L_j} Y \times \prod_{L_i \in \mathbb{L} \cap \{\succeq L_j\}} p(L_i \mid \text{mp}_{\mathcal{G}}(L_i)) \times \prod_{M_i \in \mathbb{M} \cap \{\succ L_j\}} p^*(M_i \mid \text{mp}_{\mathcal{G}}(M_i))|_{T=t} \Bigg) \\
&\overset{(2)}{=} \sum_{\prec L_j} \frac{\prod_{M_i \prec L_j} p^*(M_i \mid \text{mp}_{\mathcal{G}}(M_i))|_{T=t}}{\prod_{M_i \prec L_j} p^*(M_i \mid \text{mp}_{\mathcal{G}}(M_i))} \times \prod_{V_i \prec L_j} p(V_i \mid \text{mp}_{\mathcal{G}}(V_i)) \times \sum_{L_j} p(L_j \times \text{mp}_{\mathcal{G}}(L_j)) \\
&\qquad \times \Bigg(\sum_{\succ L_j} Y \times \prod_{L_i \in \mathbb{L} \cap \{\succ L_j\}} p(L_i \mid \text{mp}_{\mathcal{G}}(L_i)) \times \prod_{M_i \in \mathbb{M} \cap \{\succ L_j\}} p^*(M_i \mid \text{mp}_{\mathcal{G}}(M_i))|_{T=t} \\
&\qquad\qquad - \sum_{\succeq L_j} Y \times \prod_{L_i \in \mathbb{L} \cap \{\succeq L_j\}} p(L_i \mid \text{mp}_{\mathcal{G}}(L_i)) \times \prod_{M_i \in \mathbb{M} \cap \{\succ L_j\}} p^*(M_i \mid \text{mp}_{\mathcal{G}}(M_i))|_{T=t} \Bigg) \\
&\overset{(3)}{=} \sum_{\prec L_i} \frac{\prod_{M_i \prec L_j} p^*(M_i \mid \text{mp}_{\mathcal{G}}(M_i))|_{T=t}}{\prod_{M_i \prec L_j} p^*(M_i \mid \text{mp}_{\mathcal{G}}(M_i))} \times \prod_{V_i \prec L_j} p(V_i \mid \text{mp}_{\mathcal{G}}(V_i)) \\
&\qquad \times \Bigg(\sum_{\succeq L_j} Y \times \prod_{L_i \in \mathbb{L} \cap \{\succeq L_j\}} p(L_i \mid \text{mp}_{\mathcal{G}}(L_i)) \times \prod_{M_i \in \mathbb{M} \cap \{\succ L_j\}} p^*(M_i \mid \text{mp}_{\mathcal{G}}(M_i))|_{T=t} \\
&\qquad\qquad - \sum_{\succeq L_j} Y \times \prod_{L_i \in \mathbb{L} \cap \{\succeq L_j\}} p(L_i \mid \text{mp}_{\mathcal{G}}(L_i)) \times \prod_{M_i \in \mathbb{M} \cap \{\succ L_j\}} p^*(M_i \mid \text{mp}_{\mathcal{G}}(M_i))|_{T=t} \Bigg) \\
&\overset{(4)}{=} 0.
\end{aligned}
$$

The first equality is by definition of joint expectation. The second equality is by the fact that terms associated with $\prec L_j$ are not functions of L_j. The third equality is by term grouping.

Moreover, for any $M_j, M_{j-1} \in \mathbb{M}$, the following equality holds,

$$
\begin{aligned}
&\mathbb{E}\Bigg[\frac{\mathbb{I}(T=t)}{\prod_{L_i \prec M_j} p(L_i \mid \text{mp}_{\mathcal{G}}(L_i))} \times \sum_{T \cup \{\succeq M_j\}} Y \times \prod_{L_i \in \mathbb{L}} p(L_i \mid \text{mp}_{\mathcal{G}}(L_i)) \times \prod_{M_i \in \mathbb{M} \cap \{\succeq M_j\}} p^*(M_i \mid \text{mp}_{\mathcal{G}}(M_i)) \mid_{T=t} \Bigg] \\
&= \mathbb{E}\Bigg[\frac{\mathbb{I}(T=t)}{\prod_{L_i \prec M_{j-1}} p(L_i \mid \text{mp}_{\mathcal{G}}(L_i))} \times \sum_{T \cup \{\succ M_{j-1}\}} Y \times \prod_{L_i \in \mathbb{L}} p(L_i \mid \text{mp}_{\mathcal{G}}(L_i)) \times \prod_{M_i \in \mathbb{M} \cap \{\succ M_{j-1}\}} p^*(M_i \mid \text{mp}_{\mathcal{G}}(M_i)) \mid_{T=t} \Bigg],
\end{aligned}
$$

since the left hand side is equal to

$$
\begin{aligned}
&\sum_{\prec M_i} p(\prec M_i) \times \frac{\mathbb{I}(T=t)}{\prod_{L_i \prec M_j} p(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i))} \\
&\qquad\qquad \times \Big[\sum_{T \cup \{\succeq M_j\}} Y \times \prod_{L_i \in \mathbb{L}} p(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i)) \times \prod_{M_i \in \mathbb{M} \cap \{\succeq M_j\}} p^*(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i)) \mid_{T=t} \Big] \\
&\overset{(1)}{=} \sum_{\preceq M_{j-1}} p(\preceq M_{i-1}) \times \Bigg\{ \sum_{M_{j-1} \prec L_k \prec M_j} \frac{\mathbb{I}(T=t)}{\prod_{L_i \prec M_j} p(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i))} \times p(L_k \mid \mathrm{mp}_{\mathcal{G}}(L_k)) \\
&\qquad\qquad \times \Big[\sum_{T \cup \{\succeq M_j\}} Y \times \prod_{L_i \in \mathbb{L}} p(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i)) \times \prod_{M_i \in \mathbb{M} \cap \{\succeq M_j\}} p^*(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i)) \mid_{T=t} \Big] \Bigg\} \\
&\overset{(2)}{=} \sum_{\preceq M_{j-1}} p(\preceq M_{j-1}) \times \frac{\mathbb{I}(T=t)}{\prod_{L_i \prec M_{j-1}} p(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i))} \\
&\qquad \times \sum_{M_{j-1} \prec L_k \prec M_j} \Bigg\{ \sum_{T \cup \{\succeq M_j\}} Y \times \prod_{L_i \in \mathbb{L}} p(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i)) \times \prod_{M_i \in \mathbb{M} \cap \{\succeq M_j\}} p^*(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i)) \mid_{T=t} \Bigg\} \\
&\overset{(3)}{=} \sum_{\preceq M_{j-1}} p(\preceq M_{j-1}) \times \frac{\mathbb{I}(T=t)}{\prod_{L_i \prec M_{j-1}} p(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i))} \\
&\qquad \times \Bigg\{ \sum_{T \cup \{\succ M_{j-1}\}} Y \times \prod_{L_i \in \mathbb{L}} p(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i)) \times \prod_{M_i \in \mathbb{M} \cap \{\succeq M_j\}} p^*(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i)) \mid_{T=t} \Bigg\} \\
&\overset{(4)}{=} \mathbb{E}\Bigg[\frac{\mathbb{I}(T=t)}{\prod_{L_i \prec M_{j-1}} p(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i))} \\
&\qquad \times \Bigg\{ \sum_{T \cup \{\succ M_{j-1}\}} Y \times \prod_{L_i \in \mathbb{L}} p(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i)) \times \prod_{M_i \in \{\mathbb{M} \cap \succ M_{j-1}\}} p^*(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i)) \mid_{T=t} \Bigg\} \Bigg],
\end{aligned}
$$

which is exactly the same as the right hand side. This leaves the IF with only two terms $\psi(t)$ and β_{primal} and according to Lemma 4, $\mathbb{E}[\beta_{\text{primal}}] = \psi(t)$, provided the models in $\mathbb{L}$ are correctly specified, which was assumed. Therefore, $\mathbb{E}[U_{\psi_t}] = 0$.

Scenario 2. Assume models in $\mathbb{M}$ are correctly specified, and let $p^*(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i))$ denote the misspecified model for $p(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i)), \forall L_i \in \mathbb{L}$. We note that for any $M_j \in \mathbb{M}$, the following line in the IF evaluates to zero.

$$
\begin{aligned}
&\mathbb{E}\Big[\frac{\mathbb{I}(T=t)}{\prod_{L_i \prec M_j} p^*(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i))} \Big(\sum_{T \cup \{\succ M_j\}} Y \times \prod_{L_i \in \mathbb{L}} p^*(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i)) \times \prod_{M_i \in \mathbb{M} \cap \{\succ M_j\}} p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i)) |_{T=t} \\
&\qquad\qquad - \sum_{T \cup \{\succeq M_j\}} Y \times \prod_{L_i \in \mathbb{L}} p^*(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i)) \times \prod_{M_i \in \mathbb{M} \cap \{\succeq M_j\}} p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i)) |_{T=t} \Big) \Big] \\
&\overset{(1)}{=} \sum_{\preceq M_j} \frac{\mathbb{I}(T=t)}{\prod_{L_i \prec M_j} p^*(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i))} \times \prod_{V_i \prec M_j} p(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i)) \times p(M_j \mid \mathrm{mp}_{\mathcal{G}}(M_j)) \\
&\qquad \times \Big(\sum_{T \cup \{\succ M_j\}} Y \times \prod_{L_i \in \mathbb{L}} p^*(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i)) \times \prod_{M_i \in \mathbb{M} \cap \{\succ M_j\}} p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i)) |_{T=t} \\
&\qquad\quad - \sum_{T \cup \{\succeq M_j\}} Y \times \prod_{L_i \in \mathbb{L}} p^*(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i)) \times \prod_{M_i \in \mathbb{M} \cap \{\succeq M_j\}} p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i)) |_{T=t} \Big) \\
&\overset{(2)}{=} \sum_{\prec M_j} \frac{\mathbb{I}(T=t)}{\prod_{L_i \prec M_j} p^*(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i))} \times \prod_{V_i \prec M_j} p(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i)) \times \sum_{M_j} p(M_j \mid \mathrm{mp}_{\mathcal{G}}(M_j))
\end{aligned}
$$

$$
\begin{aligned}
&\times \Bigg(\sum_{T \cup \{\succ M_j\}} Y \times \prod_{L_i \in \mathbb{L}} p^*(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i)) \times \prod_{M_i \in \mathbb{M} \cap \{\succ M_j\}} p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))|_{T=t} \\
&\quad - \sum_{T \cup \{\succeq M_j\}} Y \times \prod_{L_i \in \mathbb{L}} p^*(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i)) \times \prod_{M_i \in \mathbb{M} \cap \{\succeq M_j\}} p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))|_{T=t} \Bigg) \\
&\overset{(3)}{=} \sum_{\prec M_j} \frac{\mathbb{I}(T=t)}{\prod_{L_i \prec M_j} p^*(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i))} \times \prod_{V_i \prec M_j} p(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i)) \\
&\qquad \times \Bigg(\sum_{T \cup \{\succeq M_j\}} Y \times \prod_{L_i \in \mathbb{L}} p^*(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i)) \times \prod_{M_i \in \mathbb{M} \cap \{\succeq M_j\}} p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))|_{T=t} \\
&\qquad\quad - \sum_{T \cup \{\succeq M_j\}} Y \times \prod_{L_i \in \mathbb{L}} p^*(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i)) \times \prod_{M_i \in \mathbb{M} \cap \{\succeq M_j\}} p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))|_{T=t} \Bigg) \\
&\overset{(4)}{=} 0.
\end{aligned}
$$

Moreover, for any $L_j, L_{j-1} \in \mathbb{L}$, the following equality holds,

$$
\begin{aligned}
&\mathbb{E}\Bigg[\frac{\prod_{M_i \prec L_j} p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))|_{T=t}}{\prod_{M_i \prec L_j} p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))} \\
&\qquad \times \sum_{\succeq L_j} Y \times \prod_{L_i \in \mathbb{L} \cap \{\succeq L_j\}} p^*(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i)) \times \prod_{M_i \in \mathbb{M} \cap \{\succeq L_j\}} p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i)) \,|_{T=t} \Bigg] \\
&\mathbb{E}\Bigg[\frac{\prod_{M_i \prec L_{j-1}} p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))|_{T=t}}{\prod_{M_i \prec L_{j-1}} p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))} \\
&\qquad \times \sum_{\succ L_{j-1}} Y \times \prod_{L_i \in \mathbb{L} \cap \{\succ L_{j-1}\}} p^*(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i)) \times \prod_{M_i \in \mathbb{M} \cap \{\succ L_{j-1}\}} p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i)) \,|_{T=t} \Bigg],
\end{aligned}
$$

since the left hand side is equal to

$$
\begin{aligned}
&\sum_{\prec L_j} p(\prec L_j) \times \frac{\prod_{M_i \prec L_j} p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))|_{T=t}}{\prod_{M_i \prec L_j} p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))} \\
&\qquad \times \sum_{\succeq L_j} Y \times \prod_{L_i \in \mathbb{L} \cap \{\succeq L_j\}} p^*(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i)) \times \prod_{M_i \in \mathbb{M} \cap \{\succeq L_j\}} p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i)) \,|_{T=t} \\
&\overset{(1)}{=} \sum_{\preceq L_{j-1}} p(\preceq L_{j-1}) \times \Bigg\{ \sum_{L_{j-1} \prec M_k \prec L_j} \frac{\prod_{M_i \prec L_j} p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))|_{T=t}}{\prod_{M_i \prec L_j} p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))} \times p(M_k \mid \mathrm{mp}_{\mathcal{G}}(M_k)) \\
&\qquad \times \sum_{\succeq L_j} Y \times \prod_{L_i \in \mathbb{L} \cap \{\succeq L_j\}} p^*(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i)) \times \prod_{M_i \in \mathbb{M} \cap \{\succeq L_j\}} p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i)) \,|_{T=t} \Bigg\} \\
&\overset{(2)}{=} \sum_{\preceq L_{j-1}} p(\preceq L_{j-1}) \times \frac{\prod_{M_i \prec L_{j-1}} p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))|_{T=t}}{\prod_{M_i \prec L_{j-1}} p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))} \times \Bigg\{ \sum_{L_{j-1} \prec M_k \prec L_j} p(M_k \mid \mathrm{mp}_{\mathcal{G}}(M_k))|_{T=t} \\
&\qquad \times \sum_{\succeq L_j} Y \times \prod_{L_i \in \mathbb{L} \cap \{\succeq L_j\}} p^*(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i)) \times \prod_{M_i \in \mathbb{M} \cap \{\succeq L_j\}} p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i)) \,|_{T=t} \Bigg\} \Bigg\} \\
&\overset{(3)}{=} \sum_{\preceq L_{j-1}} p(\preceq L_{j-1}) \times \frac{\prod_{M_i \prec L_{j-1}} p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))|_{T=t}}{\prod_{M_i \prec L_{j-1}} p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))} \times \\
&\qquad \times \Bigg\{ \sum_{L_{j-1} \prec M_k \prec L_j} \Bigg[\sum_{\succeq L_j} Y \times \prod_{L_i \in \mathbb{L} \cap \{\succeq L_j\}} p^*(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i)) \times \prod_{M_i \in \mathbb{M} \cap \{\succeq L_j\}} p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i)) \,|_{T=t} \Bigg]
\end{aligned}
$$

$$\times\, p(M_k \mid \mathrm{mp}_{\mathcal{G}}(M_k))|_{T=t}\Big\}$$

$$\overset{(4)}{=} \sum_{\preceq L_{j-1}} p(\preceq L_{j-1}) \times \frac{\prod_{M_i \prec L_{j-1}} p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))|_{T=t}}{\prod_{M_i \prec L_{j-1}} p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))}$$

$$\times \sum_{\succ L_{j-1}} Y \times \prod_{L_i \in \mathbb{L} \cap \{\succ L_{j-1}\}} p^*(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i)) \times \prod_{M_i \in \mathbb{M} \cap \{\succ L_{j-1}\}} p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))\,|_{T=t}\Big\}$$

$$\overset{(5)}{=} \mathbb{E}\Bigg[\frac{\prod_{M_i \prec L_{j-1}} p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))|_{T=t}}{\prod_{M_i \prec L_{j-1}} p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))}$$

$$\times \sum_{\succ L_{j-1}} Y \times \prod_{L_i \in \mathbb{L} \cap \{\succ L_{j-1}\}} p^*(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i)) \times \prod_{M_i \in \mathbb{M} \cap \{\succ L_{j-1}\}} p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))\,|_{T=t}\Bigg],$$

which is exactly the same as the right hand side. This leaves the IF with only two terms $\psi(t)$ and β_{dual} and according to Lemma 6, $\mathbb{E}[\beta_{\text{dual}}] = \psi(t)$. Therefore, $\mathbb{E}[U_{\psi_t}] = 0$. ■

Theorem 11 (Reformulation of the IF for augmented primal IPW)

Proof We prove this theorem, by showing what happens to $V_i \in V$, if V_i is in $\mathbb{M}$, or $\mathbb{L}$, or $\mathbb{C}$.

∘ For any $M_j \in \mathbb{M}$, we have,

$$\begin{aligned}
\mathbb{E}\Big[\beta_{\text{primal}} \;\Big|\; \{\preceq M_j\}\Big] &= \mathbb{E}\Bigg[\frac{\mathbb{I}(T=t)}{\prod_{L_i \in \mathbb{L}} p(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i))} \times \sum_{T} \prod_{L_i \in \mathbb{L}} p(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i)) \times Y \;\Bigg|\; \{\preceq M_i\}\Bigg] \\
&= \sum_{V_i \succ M_j} p(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i)) \times \frac{\mathbb{I}(T=t)}{\prod_{L_i \in \mathbb{L}} p(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i))} \times \sum_{T} \prod_{L_i \in \mathbb{L}} p(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i)) \times Y \\
&= \frac{\mathbb{I}(T=t)}{\prod_{L_i \prec M_j} p(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i))} \times \sum_{V_i \succ M_j} \sum_{T} \prod_{V_i \in \mathbb{L} \cup \{\succ M_j\}} p(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i))|_{T=t \text{ if } V_i \in \mathbb{M}} \times Y\Big\} \\
&= \frac{\mathbb{I}(T=t)}{\prod_{L_i \prec M_j} p(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i))} \times \sum_{T \cup \{\succ M_j\}} Y \times \prod_{V_i \in \mathbb{L} \cup \{\succ M_j\}} p(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i))|_{T=t \text{ if } V_i \in \mathbb{M}}\Big\}.
\end{aligned}$$

Similarly,

$$\mathbb{E}\big[\beta_{\text{primal}} \mid \{\prec M_j\}\big] = \frac{\mathbb{I}(T=t)}{\prod_{L_i \prec M_j} p(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i))} \times \sum_{T \cup \{\succeq M_j\}} Y \times \prod_{V_i \in \mathbb{L} \cup \{\succeq M_j\}} p(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i))|_{T=t \text{ if } V_i \in \mathbb{M}}\Big\}.$$

Therefore, $\mathbb{E}\big[\beta_{\text{primal}} \mid \{\preceq M_j\}\big] - \mathbb{E}\big[\beta_{\text{primal}} \mid \prec \{M_j\}\big]$ is equivalent to M_j's corresponding line in the IF.

∘ Now, for any $L_j \in \mathbb{L}$, we have,

$$\begin{aligned}
\mathbb{E}\Big[\beta_{\text{dual}} \;\Big|\; \{\preceq L_j\}\Big] &= \mathbb{E}\Bigg[\frac{\prod_{M_i \in \mathbb{M}} \; p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))|_{T=t}}{\prod_{M_i \in \mathbb{M}} \; p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))} \times Y \;\Bigg|\; \{\preceq L_j\}\Bigg] \\
&= \sum_{V_i \succ L_i} p(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i)) \times \frac{\prod_{M_i \in \mathbb{M}} \; p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))|_{T=t}}{\prod_{M_i \in \mathbb{M}} \; p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))} \times Y \\
&= \frac{\prod_{M_i \prec L_j} \; p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))|_{T=t}}{\prod_{M_i \prec L_j} \; p(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i))} \times \sum_{V_i \succ L_j} Y \times \prod_{V_i \succ L_j} p(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i))|_{T=t \text{ if } V_i \in \mathbb{M}}.
\end{aligned}$$

Similarly,

$$\mathbb{E}\Big[\beta_{\text{dual}} \;\Big|\; \{\prec L_j\}\Big] = \frac{\prod_{M_i \prec L_j} \; p(M_i \mid \text{mp}_{\mathcal{G}}(M_i))|_{T=t}}{\prod_{M_i \prec L_j} \; p(M_i \mid \text{mp}_{\mathcal{G}}(M_i))} \times \sum_{V_i \succeq L_j} Y \times \prod_{V_i \succeq L_j} p(V_i \mid \text{mp}_{\mathcal{G}}(V_i))|_{T=t \text{ if } V_i \in \mathbb{M}}.$$

Therefore, $\mathbb{E}\big[\beta_{\text{dual}} \mid \{\preceq L_j\}\big] - \mathbb{E}\big[\beta_{\text{dual}} \mid \{\prec L_j\}\big]$ is equivalent to L_j's corresponding line in the IF.

∘ For variables in $\mathbb{C}$, we have,

$$\begin{aligned}
\mathbb{E}[\beta_{\text{primal}} \mid \mathbb{C}] &= \mathbb{E}\Big[\frac{\mathbb{I}(T=t)}{\prod_{L_i \in \mathbb{L}} p(L_i \mid \text{mp}_{\mathcal{G}}(L_i))} \times \sum_{T} \prod_{L_i \in \mathbb{L}} p(L_i \mid \text{mp}_{\mathcal{G}}(L_i)) \times Y \;\Big|\; \mathbb{C}\Big] \\
&= \sum_{V \setminus \mathbb{C}} p(V \setminus \mathbb{C}) \times \frac{\mathbb{I}(T=t)}{\prod_{L_i \in \mathbb{L}} p(L_i \mid \text{mp}_{\mathcal{G}}(L_i))} \times \sum_{T} \prod_{L_i \in \mathbb{L}} p(L_i \mid \text{mp}_{\mathcal{G}}(L_i)) \times Y \\
&= \sum_{V \setminus \mathbb{C}} \mathbb{I}(T=t) \times \prod_{M_i \in \mathbb{M}} p(M_i \mid \text{mp}_{\mathcal{G}}(M_i)) \times \sum_{T} \prod_{L_i \in \mathbb{L}} p(L_i \mid \text{mp}_{\mathcal{G}}(L_i)) \times Y \\
&= \sum_{V \setminus \{T, \mathbb{C}\}} \prod_{M_i \in \mathbb{M}} p(M_i \mid \text{mp}_{\mathcal{G}}(M_i))\Big|_{T=t} \times \sum_{T} \prod_{L_i \in \mathbb{L}} p(L_i \mid \text{mp}_{\mathcal{G}}(L_i)) \times Y,
\end{aligned}$$

and based on Lemma 4, $\mathbb{E}[\beta_{\text{primal}}] = \psi(t)$. Therefore, $\mathbb{E}[\beta_{\text{primal}} \mid \mathbb{C}] - \mathbb{E}[\beta_{\text{primal}}]$ corresponds to the last line in the IF. We can also run a similar argument for β_{dual}. According to Lemma 6, $\mathbb{E}[\beta_{\text{dual}}] = \psi(t)$, and

$$\begin{aligned}
\mathbb{E}[\beta_{\text{dual}} \mid \mathbb{C}] &= \mathbb{E}\Big[\frac{\prod_{M_i \in \mathbb{M}} \; p(M_i \mid \text{mp}_{\mathcal{G}}(M_i))|_{T=t}}{\prod_{M_i \in \mathbb{M}} \; p(M_i \mid \text{mp}_{\mathcal{G}}(M_i))} \times Y \;\Big|\; \mathbb{C}\Big] \\
&= \sum_{V \setminus \mathbb{C}} p(V \setminus \mathbb{C}) \times \frac{\prod_{M_i \in \mathbb{M}} \; p(M_i \mid \text{mp}_{\mathcal{G}}(M_i))|_{T=t}}{\prod_{M_i \in \mathbb{M}} \; p(M_i \mid \text{mp}_{\mathcal{G}}(M_i))} \times Y \\
&= \sum_{V \setminus \mathbb{C}} \prod_{M_i \in \mathbb{M}} p(M_i \mid \text{mp}_{\mathcal{G}}(M_i))|_{T=t} \times \prod_{L_i \in \mathbb{L}} p(L_i \mid \text{mp}_{\mathcal{G}}(L_i)) \times Y \\
&= \sum_{V \setminus \{T, \mathbb{C}\}} \prod_{M_i \in \mathbb{M}} p(M_i \mid \text{mp}_{\mathcal{G}}(M_i))\Big|_{T=t} \times \sum_{T} \prod_{L_i \in \mathbb{L}} p(L_i \mid \text{mp}_{\mathcal{G}}(L_i)) \times Y \\
&= \mathbb{E}[\beta_{\text{primal}} \mid \mathbb{C}].
\end{aligned}$$

■

Lemma 12 (Efficiency of IFs in Theorems 2 and 9)

Proof Follows trivially as there is only one influence function when the model is NPS. ■

Theorem 13 (Soundness and completeness of Algorithm 1)

Proof The construction of Algorithm 1 is closely related to the *maximal arid projection* described in (Shpitser et al., 2018). MArGs were proposed as a more general analogue of

maximal ancestral graphs typically used in the context of causal discovery and where the absence of edges may only imply ordinary conditional independence constraints (Richardson and Spirtes, 2002; Zhao et al., 2005; Ali et al., 2009). The absence of an edge between two vertices in a MArG rule out the presence of certain paths between them known as *dense inducing paths* resulting in the so called *maximality* property. We now show that Algorithm 1 declares an input ADMG to be NPS if it is equivalent to a MArG with no missing edges, and not NPS if it is equivalent to one with at least one missing edge. We then use the maximality property to derive the form of the implied equality constraint.

Given *any* ADMG $\mathcal{G}(V)$, there exists a nested Markov equivalent MArG $\mathcal{G}^a(V)$ that implies the same set of conditional and generalized independence constraints and can be obtained via the maximal arid projection as follows (Shpitser et al., 2018). Recall the definition of $\mathrm{pa}^{\mathrm{d}}_{\mathcal{G}}(S)$ as $\bigcup_{S_i \in S} \mathrm{pa}_{\mathcal{G}}(S_i)$.

- For $V_i \in V$, the edge $V_i \to V_j$ exists in $\mathcal{G}^a(V)$ if $V_i \in \mathrm{pa}^{\mathrm{d}}_{\mathcal{G}}(\langle V_j \rangle_{\mathcal{G}})$.
- For $V_i, V_j \in V$, the edge $V_i \leftrightarrow V_j$ exists in $\mathcal{G}^a(V)$ if neither $V_i \in \mathrm{pa}^{\mathrm{d}}_{\mathcal{G}}(\langle V_j \rangle_{\mathcal{G}})$ nor $V_j \in \mathrm{pa}^{\mathrm{d}}_{\mathcal{G}}(\langle V_i \rangle_{\mathcal{G}})$ but $\langle V_i, V_j \rangle$ is a bidirected connected set.

Soundness

We prove soundness by showing that Algorithm 1 declares the model to be nonparametric saturated (NPS) only when the input ADMG $\mathcal{G}(V)$ is nested Markov equivalent to a MArG $\mathcal{G}^a(V)$ where all vertices in V are pairwise adjacent. If all vertices are pairwise adjacent, this immediately rules out the possibility of equality constraints.

For each pair of vertices (V_i, V_j) either of the first two conditions in line 3 of Algorithm 1 evaluates to True precisely when the MArG projection operator adds a directed edge between V_i and V_j. Further, the third condition in line 4 evaluates to True when the MArG projection adds a bidirected edge between V_i and V_j. Thus, as long as the MArG projection operator continues to require the presence of an edge between each pair (V_i, V_j) the negation of all the conditions makes it so that line 5 of the algorithm is never executed. Once all pairs have been checked, the model is declared to be nonparametrically saturated in line 6.

Completeness

We prove completeness by showing that Algorithm 1 declares the model to be not NPS only when the input ADMG is nested Markov equivalent to a MArG $\mathcal{G}^a(V)$ that has a pair of vertices (V_i, V_j) that are not connected by a directed or bidirected edge. We then explicate the equality constraint implied by this missing edge.

It is clear from previous arguments in the proof of soundness that the negation of the conditions in line 2 evaluates to True only when the MArG projection operator fails to add an edge between a pair of vertices (V_i, V_j). As soon as this occurs, it is also clear that the resulting MArG $\mathcal{G}^a(V)$ obtained by executing the full projection will still have a missing edge between V_i and V_j. We now show that this missing edge corresponds to an equality constraint involving V_i and V_j.

A path $(V_i, X_1, \ldots, X_p, V_j)$ is said to be inducing if every non-endpoint node X_i is both a collider on this path as well as an ancestor of at least one of the vertices V_i or V_j. Such paths are important because it has been shown that the absence of an inducing path between two

non-adjacent vertices V_i and V_j implies the existence of a set Z such that V_i and V_j are m-separated given Z (Verma and Pearl, 1990). That is, when V_i and V_j are not connected by an inducing path in $\mathcal{G}^a(V)$, there exists a set Z such that $V_i \perp\!\!\!\perp V_j \mid Z$ and this is an equality constraint that rules out nonparametric saturation of $\mathcal{G}$.

Consider the case when there does exist an inducing path between V_i and V_j. By definition of the maximality property of MArGs, there exists a valid fixing sequence for some $S \subset V$ such that this path is no longer inducing in $\phi_S(\mathcal{G}^a(V))$. We now discuss all possible cases of inducing paths between V_i and V_j and the corresponding equality constraint obtained after fixing some subset of vertices in $\mathcal{G}^a(V)$. Note it is sufficient for us to focus on the subgraph $\mathcal{G}^{ant} \equiv \mathcal{G}^a_{\text{an}_{\mathcal{G}}(V_i \cup V_j)}$ (Richardson, 2003). This subgraph also preserves the inducing path as all ancestors of V_i and V_j are included.

Consider the case when the inducing path consists of only bidirected edges i.e., $V_i \leftrightarrow X_1 \leftrightarrow, \ldots, \leftrightarrow X_p \leftrightarrow V_j$. Note that none of the vertices X_i in this path are fixable in $\mathcal{G}^{ant}$ as by definition of an inducing path, X_i is either an ancestor of V_i or of V_j. Thus, $\text{dis}_{\mathcal{G}^{ant}}(X_i) \cap \text{de}_{\mathcal{G}^{ant}}(V_i) \neq \{X_i\}$. However, the construction of the MArG $\mathcal{G}^a$ guarantees that V_i and V_j are not bidirected connected in $\langle V_i, V_j \rangle_{\mathcal{G}}$ and consequently not bidirected connected in the ancestral subgraph $\langle V_i, V_j \rangle_{\mathcal{G}^{ant}}$. In order for this to be true, at least one vertex X_i must become fixable after a sequence of fixing on some vertices S that are descendants of X and ancestors of V_i and V_j (excluding X, V_i, and V_j). In the graph $\phi_S(\mathcal{G}^{ant})$, X_i is fixable precisely because it is no longer an ancestor of either V_i or V_j. Therefore, the path $V_i \leftrightarrow X_1 \leftrightarrow, \ldots, \leftrightarrow X_p \leftrightarrow V_j$ is no longer inducing in $\phi_S(\mathcal{G}^{ant})$. Thus, there exists a set Z such that V_i and V_j can be m-separated in $\phi_S(\mathcal{G}^{ant})$, and the corresponding equality constraint is $V_i \perp\!\!\!\perp V_j \mid Z$ in $\phi_S(p(\text{an}_{\mathcal{G}}(V_i \cup V_j)); \mathcal{G}^{ant})$.

Consider the case when the inducing path is of the form $V_i \rightarrow X_1 \leftrightarrow, \ldots, \leftrightarrow X_p \leftrightarrow V_j$. As the graph $\mathcal{G}^{ant}$ is an ancestral subgraph of $\mathcal{G}^a(V)$, we can apply the Tian factorization to $\mathcal{G}^{ant}$. Define $X \equiv \{X_1, \ldots, X_p\}$, and let V^{ant} denote all vertices in $\mathcal{G}^{ant}$ and D_X denote the district in $\mathcal{G}^{ant}$ that contains $\{X, V_j\}$ or $\{V_i, X, V_j\}$ if V_i is also in the same district. Then, $q_{D_X}(D_X \mid \text{pa}_{\mathcal{G}^{ant}}(D_X))$ is identified and district factorizes with respect to the CADMG $\phi_{V \setminus D_X}(\mathcal{G}^{ant})$ (Richardson et al., 2017). In such a CADMG, the only possible directed paths from any vertex X_i to V_i or V_j are through vertices in D_X as these are the only random vertices that remain in $\phi_{V \setminus D_X}(\mathcal{G}^{ant})$. First consider the case when V_i is not in D_X. Then V_i is fixed in $\phi_{V \setminus D_X}(\mathcal{G}^{ant})$ and has no ancestors so the path $V_i \rightarrow X_1 \leftrightarrow, \ldots, X_p \leftrightarrow V_j$ remains inducing only if all vertices $X_i \in X$ have a directed path to V_j. If such a path exists for every $X_i \in X$ then no X_i is fixable in $\phi_{V \setminus D_X}(\mathcal{G}^{ant})$. Further, no $D_i \in D_X \setminus V_j$ is fixable either as they are all within the same district and have directed paths to V_j. Thus, the reachable closure of V_j in $\mathcal{G}^{ant}$ and as a consequence in $\mathcal{G}^a$, contains X_1. Since V_i is a parent of X_1, the MArG projection should have yielded an edge $V_i \rightarrow V_j$ which is a contradiction. Similarly, if V_i is in D_X, and all $X_i \in X$ have directed paths to either V_i or V_j, then $\langle V_i, V_j \rangle_{\mathcal{G}^a}$ would remain a bidirected connected set and the MArG projection would have yielded an edge $V_i \leftrightarrow V_j$ which is also a contradiction. Therefore, in either case, there exists at least one $X_i \in X$ such that X_i is neither an ancestor of V_i nor V_j in $\phi_{V \setminus D_X}(\mathcal{G}^{ant})$. Thus, the path $V_i \rightarrow X_1 \leftrightarrow, \ldots, \leftrightarrow X_p \leftrightarrow V_j$ is no longer inducing and we have the equality constraint, given some set $Z \subset V^{ant}$ that $V_i \perp\!\!\!\perp V_j \mid Z$ in $\phi_{V \setminus D_X}(\mathcal{G}^{ant})$.

Thus, it must be true that the input ADMG $\mathcal{G}(V)$ implies at least one equality constraint, specifically between the variables V_i and V_j, as it is nested Markov equivalent to a

MArG with a missing edge between these two vertices, and we have provided a form for the implied equality constraint. Hence, whenever line 5 is executed in Algorithm 1, the model is truly *not* nonparametrically saturated. ■

Lemma 14 (mb-shielded ADMGs)

Proof We first show that under these conditions, if two vertices A and B are not adjacent, they cannot be connected by an *inducing path*. It is important to rule out paths of this kind as Verma and Pearl (1990) have shown, that there is no ordinary conditional independence constraint between A and B if and only if there is an inducing path between them.

A path $(A, V_1, \ldots, V_p, B)$ is said to be inducing if every non-endpoint node V_i is both a collider on this path as well as an ancestor of at least one of the vertices A or B. Since the first condition requires every non-endpoint node to be a collider, we list all such paths between A and B and show that they either cannot exist, or cannot be inducing.

- $A \to V_1 \leftrightarrow B$
- $A \leftrightarrow V_1 \leftrightarrow B$
- $A \leftrightarrow V_1 \leftarrow B$
- $A \leftrightarrow V_1 \leftrightarrow, \cdots, \leftrightarrow V_p \leftrightarrow B$
- $A \to V_1 \leftrightarrow, \cdots, \leftrightarrow V_p \leftrightarrow B$
- $A \leftrightarrow V_1 \leftrightarrow, \cdots, \leftrightarrow V_p \leftarrow B$
- $A \to V_1 \leftrightarrow, \cdots, \leftrightarrow V_p \leftarrow B$

All paths except the last one cannot exist in $\mathcal{G}$ as it would contradict the antecedent that a missing edge is permitted only if $A \notin \mathrm{mb}_{\mathcal{G}}(B)$ and $B \notin \mathrm{mb}_{\mathcal{G}}(A)$. The last path is permitted but cannot be an inducing path as it would contradict that $\mathcal{G}$ is acyclic. For this path to be inducing, it must be that V_1 is an ancestor of A or B and same for V_p. V_1 cannot be an ancestor of A as that would lead to a cycle so it may only be an ancestor of B. Similarly, V_p cannot be an ancestor of B so it may only be an ancestor of A. If this is the case, we still get the directed cycle $V_1 \to, \cdots, \to B, \to V_p \to, \cdots, \to A \to V_1$.

We have now shown that when two vertices A and B are not adjacent, there is no inducing path between them. Thus, there must be an ordinary conditional independence constraint. We now show that all such constraints are captured by the topological factorization shown in Eq. 8. Fix any valid topological order τ for the vertices in V. By the ordinary local Markov property, $V_i \perp\!\!\!\perp \{\prec V_i\} \setminus \mathrm{mp}_{\mathcal{G}}(V_i) \mid \mathrm{mp}_{\mathcal{G}}(V_i)$ (Richardson, 2003). Since τ is a *total ordering*, it is either the case that $A \prec B$ or $B \prec A$. Assume, without loss of generality that $A \prec B$. Then, A cannot be in the Markov pillow of B, as we require that A is not in the Markov blanket of B. Therefore, $A \perp\!\!\!\perp B \mid \mathrm{mp}_{\mathcal{G}}(B)$, and more generally, all conditional independence constraints in $\mathcal{G}$ can be expressed in this form. Thus, the topological factorization $p(V) = \prod_{V_i \in V} p(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i))$ captures all constraints in $\mathcal{G}$ that corresponds to the ordinary Markov model of an ADMG (Richardson, 2003; Evans and Richardson, 2012). ■

Theorem 15 (Tangent space Λ^* of mb-shielded ADMGs)

Proof The proof here is similar to the proofs of Theorems 4.4 and 4.5 in Tsiatis (2007). Given $p(V)$ that factorizes with respect to an mb-shielded ADMG $\mathcal{G}(V)$, we can write down the following factorization using the ordinary local Markov property $V_i \perp\!\!\!\perp \{\prec V_i\} \backslash \mathrm{mp}_{\mathcal{G}}(V_i) \mid \mathrm{mp}_{\mathcal{G}}(V_i) : p(V) = \prod_{V_i \in V} p(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i))$. Under no restriction, the conditional density $p(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i)), \forall V_i \in V$, is any positive function such that $\int p(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i)) d\nu(V_i) = 1$, for all values of $\mathrm{mp}_{\mathcal{G}}(V_i))$, where $d\nu(V_i)$ is the dominating measure. Note that $p(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i))$s are variationally independent.

The tangent space Λ^*, corresponding to the model of the mb-shielded ADMG $\mathcal{G}(V)$, is defined as the mean square closure of all parametric submodel tangent spaces. Assume there are k variables in V. The parametric submodel is defined as $\mathcal{M}_{\text{sub}} = \{\prod_{V_i \in V} p(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i); \gamma_i)\}$, where $\gamma_i, i = 1, \ldots, k$ are parameters that are variationally independent and $p(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i); \gamma_{0i})$ denotes the true conditional density of $p(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i))$. The parametric submodel tangent space is defined as the space spanned by the joint score $S_\gamma(V_1, \ldots, V_k)$ given as follows,

$$S_\gamma(V_1, \ldots, V_k) = \frac{\partial}{\partial \gamma} \log p(V) = S_{\gamma_1}(V_1) + \cdots + S_{\gamma_k}(V_k, \mathrm{mp}_{\mathcal{G}}(V_k)).$$

Therefore, the parametric submodel tangent space is $\Lambda^*_\gamma = a_1 \times S_{\gamma_1}(V_1) + \cdots + a_k \times S_{\gamma_k}(V_k, \mathrm{mp}_{\mathcal{G}}(V_k))$, where a_i's are constants. Due to variational independence of γ_is, $\Lambda^*_\gamma = \Lambda^*_{\gamma_1} \oplus \cdots \oplus \Lambda^*_{\gamma_k}$, where $\Lambda^*_{\gamma_i} = \{a_i \times S_{\gamma_i}(V_i, \mathrm{mp}_{\mathcal{G}}(V_i))\}$. The tangent space Λ^* is then the mean-square closure of all parametric submodel tangent spaces, i.e., $\Lambda^* = \Lambda^*_1 \oplus \cdots \oplus \Lambda^*_k$, where Λ^*_i is the mean-square closure of the parametric submodel tangent space $\Lambda^*_{\gamma_i}$ which corresponds to the term $p(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i))$.

By the ordinary local Markov property, i.e., $V_i \perp\!\!\!\perp \{\prec V_i\} \setminus \mathrm{mp}_{\mathcal{G}}(V_i) \mid \mathrm{mp}_{\mathcal{G}}(V_i)$, and properties of score functions for parametric models of conditional densities, the score function $S_i(.)$ must be a function of only $\{V_i, \mathrm{mp}_{\mathcal{G}}(V_i)\}$ and must have conditional expectation $\mathbb{E}[S_i(V_i, \mathrm{mp}_{\mathcal{G}}(V_i)) \mid \mathrm{mp}_{\mathcal{G}}(V_i)] = 0$. Consequently, any element spanned by $S_i(V_i, \mathrm{mp}_{\mathcal{G}}(V_i))$ must belong to Λ^*_i; hence $\Lambda^*_i = \{\alpha(V_i, \mathrm{mp}_{\mathcal{G}}(V_i)) \mid \mathbb{E}[\alpha \mid \mathrm{mp}_{\mathcal{G}}(V_i)] = 0\}$. Further, in order to show that Λ^*_i's are orthogonal, we need to show that $\mathbb{E}[h_i \times h_j] = 0$, where $h_i \in \Lambda^*_i$ and $h_j \in \Lambda^*_j$,

$$\begin{aligned}
\mathbb{E}\big[h_i \times h_j\big] &= \mathbb{E}\Big[h_i \times \mathbb{E}\big[h_j \mid V_i, \mathrm{mp}_{\mathcal{G}}(V_i)\big]\Big] \\
&= \mathbb{E}\Bigg[h_i \times \mathbb{E}\Big[\mathbb{E}\big[h_j \mid V_i, \mathrm{mp}_{\mathcal{G}}(V_i), \mathrm{mp}_{\mathcal{G}}(V_j)\big] \;\Big|\; V_i, \mathrm{mp}_{\mathcal{G}}(V_i)\Big]\Bigg] \\
&= \mathbb{E}\Bigg[h_i \times \mathbb{E}\Big[\mathbb{E}\big[h_j \mid \mathrm{mp}_{\mathcal{G}}(V_j)\big] \mid V_i, \mathrm{mp}_{\mathcal{G}}(V_i)\Big]\Bigg] = 0.
\end{aligned}$$

The projection h_i is in Λ^*_i. Therefore, we only need to show that $h - h_i$ is orthogonal to all elements in Λ^*_i. Consider an arbitrary element $\ell \in \Lambda^*_i$,

$$\begin{aligned}
\mathbb{E}[(h - h_i) \times \ell] &= \mathbb{E}\Big[\ell \times \big(\mathbb{E}[h \mid V_i, \mathrm{mp}_{\mathcal{G}}(V_i)] - h_i\big)\Big] = \mathbb{E}[\ell \times \mathbb{E}[h \mid \mathrm{mp}_{\mathcal{G}}(V_i)]] \\
&= \mathbb{E}\Big[\mathbb{E}\big[\ell \times \mathbb{E}[h \mid \mathrm{mp}_{\mathcal{G}}(V_i)] \mid V_i, \mathrm{mp}_{\mathcal{G}}(V_i)\big]\Big] = \mathbb{E}\Big[\mathbb{E}[h \mid \mathrm{mp}_{\mathcal{G}}(V_i)] \times \mathbb{E}\big[\ell \mid V_i, \mathrm{mp}_{\mathcal{G}}(V_i)\big]\Big] = 0.
\end{aligned}$$

■

Theorem 16 (Orthogonal complement $\Lambda^{*\perp}$ in mb-shielded ADMGs)

Proof $\Lambda^{*\perp} = \{h - \pi[h \mid \Lambda^*], \forall h \in \mathrm{H}\}$. For a given $h \in \mathrm{H}$, we have $h = h_1 + \cdots + h_k$, where $h_i \in \Lambda_i$ as Λ_i is defined in Section 4. According to Section 4, h_i is any function of $V_1, \ldots, V_i$, such that $\mathbb{E}[h_i \mid V_1, \ldots, V_{i-1}] = 0$. Therefore,

$$\begin{aligned} h - \pi[h \mid \Lambda^*] &= \big(h_1 + \cdots + h_k\big) - \pi[h_1 + \cdots + h_k \mid \Lambda_1^* \oplus \ldots \Lambda_k^*] \\ &= \sum_{i=1}^{k} h_i - \pi[h_i \mid \Lambda_1^* \oplus \ldots \Lambda_k^*] \\ &= \sum_{i=1}^{k} h_i - \pi[h_i \mid \Lambda_i^*] \\ &= \sum_{i=1}^{k} h_i - \mathbb{E}[h_i \mid V_i, \mathrm{mp}_{\mathcal{G}}(V_i)] + \mathbb{E}[h_i \mid \mathrm{mp}_{\mathcal{G}}(V_i)]. \end{aligned}$$

The third equality holds since Λ_i is orthogonal to Λ_j^*, for $i, j = 1, \ldots, k$, such that $i \neq j$. Note that $h_i \equiv \mathbb{E}[h(V) \mid V_1, \ldots, V_i] - \mathbb{E}[h(V) \mid V_1, \ldots, V_{i-1}]$. Since $h(V)$ is an arbitrary element of the Hilbert space, without loss of generality, we can replace $\mathbb{E}[h(V) \mid V_1, \ldots, V_i]$ with $\alpha_i(V_1, \ldots, V_i)$. Therefore, $h_i = \alpha_i - \mathbb{E}[\alpha_i \mid V_1, \ldots, V_{i-1}]$. Substituting h_i in the above equation yields the following.

$$\begin{aligned} &h_i - \mathbb{E}[h_i \mid V_i, \mathrm{mp}_{\mathcal{G}}(V_i)] + \mathbb{E}[h_i \mid \mathrm{mp}_{\mathcal{G}}(V_i)] \\ &\quad = \alpha_i - \mathbb{E}[\alpha_i \mid V_1, \ldots, V_{i-1}] \\ &\qquad - \mathbb{E}[\alpha_i \mid V_i, \mathrm{mp}_{\mathcal{G}}(V_i)] + \mathbb{E}\Big[\mathbb{E}[\alpha_i \mid V_1, \ldots, V_{i-1}] \Big| \, V_i, \mathrm{mp}_{\mathcal{G}}(V_i)\Big] \\ &\qquad + \mathbb{E}[\alpha_i \mid \mathrm{mp}_{\mathcal{G}}(V_i)] - \mathbb{E}\Big[\mathbb{E}[\alpha_i \mid V_1, \ldots, V_{i-1}] \Big| \, \mathrm{mp}_{\mathcal{G}}(V_i)\Big] \\ &\quad = \alpha_i - \mathbb{E}[\alpha_i \mid V_1, \ldots, V_{i-1}] - \mathbb{E}[\alpha_i \mid V_i, \mathrm{mp}_{\mathcal{G}}(V_i)] + \mathbb{E}[\alpha_i \mid \mathrm{mp}_{\mathcal{G}}(V_i)] \\ &\quad = \Big\{\alpha_i - \mathbb{E}[\alpha_i \mid V_1, \ldots, V_{i-1}]\Big\} - \Big\{\mathbb{E}\Big[\alpha_i - \mathbb{E}[\alpha_i \mid V_i, \ldots, V_{i-1}] \;\Big|\; V_i, \mathrm{mp}_{\mathcal{G}}(V_i)\Big]\Big\}. \end{aligned}$$

Consequently, the orthogonal complement of the tangent space is the following,

$$\Lambda^{*\perp} = \Bigg\{ \sum_{V_i \in V} \alpha_i(V_1, \ldots, V_i) - \mathbb{E}[\alpha_i \mid V_i, \mathrm{mp}_{\mathcal{G}}(V_i)] \Bigg\},$$

where α_i is any function of $V_1, \ldots, V_i$ such that $\mathbb{E}[\alpha_i \mid V_1, \ldots, V_{i-1}] = 0$, i.e., $\alpha_i \in \Lambda_i$. ■

Theorem 17 (Efficient gAIPW for mb-shielded ADMGs)

Proof Let U_{ψ_t} be the IF given in Theorem 2, that is,

$$U_{\psi_t} = \frac{\mathbb{I}(T = t)}{p(T \mid \mathrm{mp}_{\mathcal{G}}(T))} \times \big(Y - \mathbb{E}[Y \mid T, \mathrm{mp}_{\mathcal{G}}(T)]\big) + \mathbb{E}[Y \mid T = t, \mathrm{mp}_{\mathcal{G}}(T)] - \psi(t).$$

We know that $U^{\text{eff}}_{\psi_t} = \pi[U_{\psi_t} \mid \Lambda^*]$. We first show that $\pi[U_{\psi_t} \mid \Lambda^*] = \pi[U_{\psi_t} \mid \Lambda^* \setminus \Lambda^*_T]$ by showing that $\pi[U_{\psi_t} \mid \Lambda^*_T] = 0$, as follows,

$$\begin{aligned}
\pi[U_{\psi_t} \mid \Lambda^*_T] &= \mathbb{E}[U_{\psi_t} \mid T, \text{mp}_{\mathcal{G}}(T)] - \mathbb{E}[U_{\psi_t} \mid \text{mp}_{\mathcal{G}}(T)] \\
&= \mathbb{E}[U_{\psi_t} \mid T, \text{mp}_{\mathcal{G}}(T)] - \mathbb{E}\Big[\mathbb{E}[U_{\psi_t} \mid T, \text{mp}_{\mathcal{G}}(T)] \;\Big|\; \text{mp}_{\mathcal{G}}(T)\Big] \\
&= \Big(\mathbb{E}[Y \mid T = t, \text{mp}_{\mathcal{G}}(T)] - \psi(t)\Big) - \mathbb{E}\big[\mathbb{E}[Y \mid T = t, \text{mp}_{\mathcal{G}}(T)] - \psi(t) \;\big|\; \text{mp}_{\mathcal{G}}(T)\big] \\
&= 0.
\end{aligned}$$

Therefore,

$$\begin{aligned}
U^{\text{eff}}_{\psi_t} &= \pi[U_{\psi_t} \mid \Lambda^* \setminus \Lambda^*_T] \\
&= \pi\Big[\frac{\mathbb{I}(T = t)}{p(T \mid \text{mp}_{\mathcal{G}}(T))} \times \big(Y - \mathbb{E}[Y \mid T, \text{mp}_{\mathcal{G}}(T)]\big) + \mathbb{E}[Y \mid T = t, \text{mp}_{\mathcal{G}}(T)] - \psi(t) \;\Big|\; \Lambda^* \setminus \Lambda^*_T\Big] \\
&= \pi\Big[\frac{\mathbb{I}(T = t)}{p(T \mid \text{mp}_{\mathcal{G}}(T))} \times Y \;\Big|\; \Lambda^* \setminus \Lambda^*_T\Big] + \pi\Big[\big(1 - \frac{\mathbb{I}(T = t)}{p(T \mid \text{mp}_{\mathcal{G}}(T))}\big) \times \mathbb{E}[Y \mid T = t, \text{mp}_{\mathcal{G}}(T)] \;\Big|\; \Lambda^* \setminus \Lambda^*_T\Big] \\
&= \pi\Big[\frac{\mathbb{I}(T = t)}{p(T \mid \text{mp}_{\mathcal{G}}(T))} \times Y \;\Big|\; \Lambda^* \setminus \Lambda^*_T\Big]
\end{aligned}$$

The last equality holds since $\big(1 - \frac{\mathbb{I}(T=t)}{p(T|\text{mp}_{\mathcal{G}}(T))}\big) \times \mathbb{E}[Y \mid T = t, \text{mp}_{\mathcal{G}}(T)]$ is a function of $\{T, \text{mp}_{\mathcal{G}}(T)\}$ and is mean zero given $\text{mp}_{\mathcal{G}}(T)$. Therefore, it belongs to Λ^*_T and hence is orthogonal to $\Lambda^* \setminus \Lambda^*_T$. We thus have that:

$$U^{\text{eff}}_{\psi_t} = \pi\Big[\frac{\mathbb{I}(T = t)}{p(T \mid \text{mp}_{\mathcal{G}}(T))} \times Y \;\Big|\; \Lambda^* \setminus \Lambda^*_T\Big]$$

Further, we show that $U^{\text{eff}}_{\psi_t}$ is orthogonal to $\{\oplus_{D_i \in D} \Lambda^*_{D_i}\}$ and $\{\oplus_{Z_i \in Z} \Lambda^*_{Z_i}\}$, where D and Z are defined as follows.

$$\begin{aligned}
&D = \{D_i \in V \mid D_i \perp\!\!\!\perp T, \text{mp}_{\mathcal{G}}(T), Y \mid \text{mp}_{\mathcal{G}}(D_i)\}, \\
&Z = \{Z_i \in V \mid Z_i \perp\!\!\!\perp Y \mid \text{mp}_{\mathcal{G}}(Z_i) \text{ in } \mathcal{G}_{V \setminus T} \text{ and } Z_i \not\perp\!\!\!\perp T \mid \text{mp}_{\mathcal{G}}(Z_i)\}.
\end{aligned}$$

For $D_i \in D$, we need to show that $\pi\Big[\frac{\mathbb{I}(T=t)}{p(T|\text{mp}_{\mathcal{G}}(T))} \times Y \;\Big|\; \Lambda^*_{D_i}\Big] = 0$, and that is obvious since $\mathbb{E}\Big[\frac{\mathbb{I}(T=t)}{p(T|\text{mp}_{\mathcal{G}}(T))} \times Y \;\Big|\; D_i, \text{mp}_{\mathcal{G}}(D_i)\Big]$ is not a function of D_i. Therefore,

$$\mathbb{E}\Big[\frac{\mathbb{I}(T = t)}{p(T \mid \text{mp}_{\mathcal{G}}(T))} \times Y \;\Big|\; D_i, \text{mp}_{\mathcal{G}}(D_i)\Big] = \mathbb{E}\Big[\frac{\mathbb{I}(T = t)}{p(T \mid \text{mp}_{\mathcal{G}}(T))} \times Y \;\Big|\; \text{mp}_{\mathcal{G}}(D_i)\Big].$$

For $Z_i \in Z$, we need to show that $\pi\Big[\frac{\mathbb{I}(T=t)}{p(T|\text{mp}_{\mathcal{G}}(T))} \times Y \;\Big|\; \Lambda^*_{Z_i}\Big] = 0$. In other words, we need to show that,

$$\mathbb{E}\Big[\frac{\mathbb{I}(T = t)}{p(T \mid \text{mp}_{\mathcal{G}}(T))} \times Y \;\Big|\; Z_i, \text{mp}_{\mathcal{G}}(Z_i)\Big] = \mathbb{E}\Big[\frac{\mathbb{I}(T = t)}{p(T \mid \text{mp}_{\mathcal{G}}(T))} \times Y \;\Big|\, \text{mp}_{\mathcal{G}}(Z_i)\Big].$$

Note that an mb-shielded ADMG is Markov equivalent to a DAG $\mathcal{G}^d$, which can be constructed as follows. Under the topological order τ fixed on the original ADMG $\mathcal{G}, V_i \to V_j$

exists in $\mathcal{G}^d$ if V_i and V_j are adjacent in $\mathcal{G}$ and $V_i \prec_\tau V_j$. $\mathcal{G}^d$ is a DAG because we only allow for directed edges and there is no directed cycle as we follow a valid topological order in $\mathcal{G}$. Further, $\text{mp}_{\mathcal{G}}(V_i) = \text{pa}_{\mathcal{G}^d}(V_i), \forall V_i \in V$. Therefore, the identifying functional for the target parameter is the same in both $\mathcal{G}$ and $\mathcal{G}^d$, that is $\mathbb{E}[\mathbb{E}[Y \mid T = t, \text{mp}_{\mathcal{G}}(T)] = \mathbb{E}[\mathbb{E}[Y \mid T = t, \text{pa}_{\mathcal{G}^d}(T)]]$. We know for the instrument variables in Z, there always exists a set $F \in V$ that d-separates $Z_i \in Z$ from Y given F, T (van der Zander et al., 2015; Rotnitzky and Smucler, 2019). Showing that the above equation is true then simply follows from the argument outlined in proposition 3 in (Rotnitzky and Smucler, 2019).
Finally, given Λ^* in Theorem 15, the efficient IF is as follows. Let $V^* = V \setminus (T \cup Z \cup D)$,

$$\begin{aligned} U^{\text{eff}}_{\psi_t} &= \pi\Big[\frac{\mathbb{I}(T=t)}{p(T \mid \text{mp}_{\mathcal{G}}(T))} \times Y \;\Big|\; \Lambda^* \setminus \Lambda^*_T\Big] \\ &= \sum_{V_i \in V^*} \mathbb{E}\Big[\frac{\mathbb{I}(T=t)}{p(T \mid \text{mp}_{\mathcal{G}}(T))} \times Y \;\Big|\; V_i, \text{mp}_{\mathcal{G}}(V_i)\Big] - \mathbb{E}\Big[\frac{\mathbb{I}(T=t)}{p(T \mid \text{mp}_{\mathcal{G}}(T))} \times Y \;\Big|\; \text{mp}_{\mathcal{G}}(V_i)\Big]. \end{aligned}$$

■

Theorem 18 (Efficient augmented primal IPW for mb-shielded ADMGs)

Proof Consider the reformulated IF in Theorem 11. In order to get the efficient IF, we project the reformulated IF onto the tangent space Λ^* given by Theorem 15. We first note that we can rewrite the term $\sum_{\mathbb{C}} \mathbb{E}[\beta_{\text{primal/dual}} \mid \mathbb{C}] - \psi(t)$ in the reformulated IF as $\sum_{C_i \in \mathbb{C}} \mathbb{E}[\beta_{\text{primal/dual}} \mid \{\preceq C_i\}] - \mathbb{E}[\beta_{\text{primal/dual}} \mid \{\prec C_i\}]$, where $\beta_{\text{primal/dual}}$ means that we can use either β_{primal} or β_{dual} for the $\mathbb{C}$ term. We have,

$$\begin{aligned} \pi[U^{\text{reform}}_{\psi_t} \mid \Lambda^*] = &\sum_{M_i \in \mathbb{M}} \pi\Big[\mathbb{E}[\beta_{\text{primal}} \mid \{\preceq M_i\}] - \mathbb{E}[\beta_{\text{primal}} \mid \{\prec M_i\}]\Big| \Lambda^*\Big] \\ &+ \sum_{L_i \in \mathbb{L}} \pi\Big[\mathbb{E}[\beta_{\text{dual}} \mid \{\preceq L_i\}] - \mathbb{E}[\beta_{\text{dual}} \mid \{\prec L_i\}] \;\Big|\; \Lambda^*\Big] \\ &+ \sum_{C_i \in \mathbb{C}} \pi\Big[\mathbb{E}[\beta_{\text{primal/dual}} \mid \preceq C_i] - \mathbb{E}[\beta_{\text{primal/dual}} \mid \prec C_i] \;\Big|\; \Lambda^*\Big]. \end{aligned}$$

Let β be either β_{primal} or β_{dual} or $\beta_{\text{primal/dual}}$. Note that $\Big\{\mathbb{E}\big[\beta \mid \{\preceq V_i\}\big] - \mathbb{E}\big[\beta_{\text{primal}} \mid \{\prec V_i\}\big]\Big\}$ lives in Λ_{V_i}, and $\Lambda_{V_i} \perp\!\!\!\perp \Lambda^* \setminus \Lambda^*_{V_i}$. Therefore, their projection onto $\Lambda^* \setminus \Lambda^*_{V_i}$ is zero. We have,

$$\begin{aligned} &\pi\Big[\mathbb{E}\big[\beta \mid \{\preceq V_i\}\big] - \mathbb{E}\big[\beta \mid \{\prec V_i\}\big]\Big| \Lambda^*_{V_i}\Big] \\ &\quad= \mathbb{E}\Big[\mathbb{E}[\beta \mid \{\preceq V_i\}] - \mathbb{E}[\beta \mid \{\prec V_i\}] \;\Big|\; V_i, \text{mp}_{\mathcal{G}}(V_i)\Big] - \mathbb{E}\Big[\mathbb{E}[\beta \mid \{\preceq V_i\}] - \mathbb{E}[\beta \mid \{\prec V_i\}] \;\Big|\; \text{mp}_{\mathcal{G}}(V_i)\Big] \\ &\quad= \mathbb{E}\big[\beta \mid V_i, \text{mp}_{\mathcal{G}}(V_i)\big] - \mathbb{E}\Big[\mathbb{E}\big[\beta \mid \prec V_i\big] \;\Big|\; V_i, \text{mp}_{\mathcal{G}}(V_i)\Big] - \mathbb{E}\big[\beta \mid \text{mp}_{\mathcal{G}}(V_i)\big] + \mathbb{E}\big[\beta \mid \text{mp}_{\mathcal{G}}(V_i)\big] \\ &\quad= \mathbb{E}\big[\beta \mid V_i, \text{mp}_{\mathcal{G}}(V_i)\big] - \mathbb{E}\Big[\mathbb{E}\big[\beta \mid \prec V_i\big] \;\Big|\; V_i, \text{mp}_{\mathcal{G}}(V_i)\Big] \\ &\quad= \mathbb{E}\big[\beta \mid V_i, \text{mp}_{\mathcal{G}}(V_i)\big] - \mathbb{E}\big[\beta \mid \text{mp}_{\mathcal{G}}(V_i)\big]. \end{aligned}$$

Therefore, the efficient IF is as follows.

$$\pi[U^{\text{reform}}_{\psi_t} \mid \Lambda^*] = \sum_{M_i \in \mathbb{M}} \mathbb{E}\big[\,\beta_{\text{primal}} \mid M_i, \text{mp}_{\mathcal{G}}(M_i)\big] - \mathbb{E}\big[\,\beta_{\text{primal}} \mid \text{mp}_{\mathcal{G}}(M_i)\big]$$
$$+ \sum_{L_i \in \mathbb{L}} \mathbb{E}\big[\,\beta_{\text{dual}} \mid L_i, \text{mp}_{\mathcal{G}}(L_i)\big] - \mathbb{E}\big[\,\beta_{\text{dual}} \mid \text{mp}_{\mathcal{G}}(L_i)\big]$$
$$+ \sum_{C_i \in \mathbb{C}} \mathbb{E}\big[\beta_{\text{primal/dual}} \mid C_i, \text{mp}_{\mathcal{G}}(C_i)\big] - \mathbb{E}\big[\beta_{\text{primal/dual}} \mid \text{mp}_{\mathcal{G}}(C_i)\big].$$

■

Lemma 19 (Primal fixing operator)

Proof Given a CADMG $\mathcal{G}(V, W)$, the Tian factorization of the distribution $q(V \mid W)$ is:

$$q(V \mid W) = \prod_{D \in \mathcal{D}(\mathcal{G})} q_D(D \mid \text{pa}_{\mathcal{G}}(D), W)$$
$$= q_{D_T}(D_T \mid \text{pa}_{\mathcal{G}}(D_T), W) \times \prod_{D \in \mathcal{D}(\mathcal{G}) \setminus D_T} q_D(D \mid \text{pa}_{\mathcal{G}}(D), W).$$

Therefore, the probabilistic operation of primal fixing a variable T in $q(V \mid W)$ is given by,

$$\Phi_T(q_V; \mathcal{G}) \equiv \frac{q_{D_T}(D_T \mid \text{pa}_{\mathcal{G}}(D_T), W)}{q_{D_T}(T \mid \text{mb}_{\mathcal{G}}(D_T), W)} \times \prod_{D \in \mathcal{D}(\mathcal{G}) \setminus D_T} q_D(D \mid \text{pa}_{\mathcal{G}}(D), W). \tag{40}$$

Lemma 11 in (Richardson et al., 2017) gives the following construction for new independences arising in kernels obtained through fixing. Given disjoint sets $A, B, C, D \in V$, and a kernel $q_V(A \cup B \cup C \mid D, W)$,

$$A \perp\!\!\!\perp B \mid D, W \text{ in } \frac{q_V(A \cup B \cup C \mid D, W)}{q_V(A \mid B, D, W)}.$$

Applying this result to the ratio shown in Eq. 40, we see that the resulting kernel from primal fixing must have the independence $T \perp\!\!\!\perp \text{mb}_{\mathcal{G}}(T) \mid W$. It is easy to check via the extension of the m-separation criterion to CADMGs (see (Richardson et al., 2017) for details) that these independence relations hold in the CADMG $\Phi_{\mathcal{G}}(T)$, where all incoming edges to T have been deleted. ■

Lemma 21 (Dual fixing operator)

Proof Given a CADMG $\mathcal{G}(V, W)$ and corresponding kernel $q(V \mid W)$, the probabilistic operation of dual fixing a variable T is defined as,

$$\Delta_T(q_V; \mathcal{G}) = \sum_T q_V(V \mid W) \times \frac{\prod_{M_i \in \text{mp}_{\mathcal{G}}^{-1}(T)} q_V(M_i \mid \text{mp}_{\mathcal{G}}(M_i), W) \mid_{T=t}}{\prod_{M_i \in \text{mp}_{\mathcal{G}}^{-1}(T)} q_V(M_i \mid \text{mp}_{\mathcal{G}}(M_i), W)}.$$

Fix a reverse topological order $(M_1, \ldots, M_p)$ on elements in $\mathrm{mp}_{\mathcal{G}}^{-1}(T)$. Then consider

$$q^* \equiv \frac{q_V(V \mid W)}{q_V(M_1 \mid \mathrm{mp}_{\mathcal{G}}(M_1), W)}$$

As seen in the proof of the primal fixing operator, we see that $M_1 \perp\!\!\!\perp \mathrm{mp}_{\mathcal{G}}(M_1) \mid W$ in q^* and can be associated with a CADMG where all incoming edges from vertices in $\{\prec M_1\}$ to M_1 are removed. Applying this inductively to all elements M_i in $\mathrm{mp}_{\mathcal{G}}^{-1}(T)$ according to the reverse topological order, we get that

$$\frac{q_V(V \mid W)}{\prod_{M_i \in \mathrm{mp}_{\mathcal{G}}^{-1}(T)} q_V(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i), W)}$$

is associated with a CADMG $\mathcal{G}^*$ where all incoming edges from vertices in $\{\prec M_i\}$ to M_i have been removed for all $M_i \in \mathrm{mp}_{\mathcal{G}}^{-1}(T)$. Note that $T \prec M_i$ and $\mathrm{ch}_{\mathcal{G}}(T) \subset \mathrm{mp}_{\mathcal{G}}^{-1}(T)$ by definition of the inverse Markov pillow. Therefore, T has no children in the resulting CADMG $\mathcal{G}^*$ due to the removal of edges. Now, remultiplying by the same conditionals but with $T = t$ introduces a fixed variable t into the graph and readds the edges that were deleted by the division, with the caveat that the original outgoing edges from T now emerge from the fixed node t. Thus, this results in a kernel that is exactly analogous to the node-splitting operation described in (Richardson and Robins, 2013). As the random variable T is still childless in this CADMG, marginalizing T in the kernel,

$$q' \equiv \frac{q_V(V \mid W)}{\prod_{M_i \in \mathrm{mp}_{\mathcal{G}}^{-1}} q_V(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i), W)} \times \prod_{M_i \in \mathrm{mp}_{\mathcal{G}}^{-1}} q_V(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i), W)|_{T=t},$$

and the corresponding CADMG $\mathcal{G}'$, yields the distribution $\Delta_T(q_V; \mathcal{G})$ corresponding to d-fixing T and the CADMG $\Delta_T(\mathcal{G})$ where all incoming edges into T have been removed and T is fixed to t. ■

Lemma 22 (Equivalence of the primal and dual operators)

Proof Equivalence of the graphical operators $\Phi_T(\mathcal{G})$ and $\Delta_T(\mathcal{G})$ is trivial as they are both defined to be equal to the a-fixing graphical operation $\phi_T(\mathcal{G})$.

We now show that $\Phi_T(q_V; \mathcal{G}) = \Delta_T(q_V; \mathcal{G})$ for any fixed value $T = t$.

$$\Delta_T(q_V; \mathcal{G}) \equiv \sum_T q_V(V \mid W) \times \frac{\prod_{M_i \in \mathrm{mp}_{\mathcal{G}}^{-1}(T)} q_V(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i), W) \mid_{T=t}}{\prod_{M_i \in \mathrm{mp}_{\mathcal{G}}^{-1}(T)} q_V(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i), W)}.$$

We can partition V into $\mathrm{mp}_{\mathcal{G}}^{-1}(T)$ and $V \setminus \mathrm{mp}_{\mathcal{G}}^{-1}(T)$ yielding,

$$\sum_T \prod_{V_i \in V \setminus \mathrm{mp}_{\mathcal{G}}^{-1}(T)} q_V(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i), W) \times$$
$$\prod_{M_i \in \mathrm{mp}_{\mathcal{G}}^{-1}(T)} q_V(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i), W) \times \frac{\prod_{M_i \in \mathrm{mp}_{\mathcal{G}}^{-1}(T)} q_V(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i), W) \mid_{T=t}}{\prod_{M_i \in \mathrm{mp}_{\mathcal{G}}^{-1}(T)} q_V(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i), W)}$$

$$= \sum_T \prod_{V_i \in V \setminus \mathrm{mp}_{\mathcal{G}}^{-1}(T)} q_V(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i), W) \times \prod_{M_i \in \mathrm{mp}_{\mathcal{G}}^{-1}(T)} q_V(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i), W) \mid_{T=t}.$$

Partitioning $V \setminus \mathrm{mp}_{\mathcal{G}}^{-1}(T)$ further into D_T and the remaining vertices we have,

$$\sum_T \prod_{V_i \in V \setminus (\mathrm{mp}_{\mathcal{G}}^{-1}(T) \cup D_T)} q_V(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i), W) \times \prod_{M_i \in \mathrm{mp}_{\mathcal{G}}^{-1}(T)} q_V(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i), W)|_{T=t} \times \prod_{D_i \in D_T} q_V(D_i \mid \mathrm{mp}_{\mathcal{G}}(D_i), W).$$

As only vertices that are in D_T are functions of T, the summation over T can be pushed past the factors over $V \setminus (\mathrm{mp}_{\mathcal{G}}^{-1}(T) \cup D_T)$ and $\mathrm{mp}_{\mathcal{G}}^{-1}(T)$ and these sets can be merged. This gives us,

$$\prod_{V_i \in V \setminus D_T} q_V(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i), W) \sum_T \prod_{D_i \in D_T} q_V(D_i \mid \mathrm{mp}_{\mathcal{G}}(D_i), W) \Big|_{T=t}.$$

Multiplying and dividing by factors involved in D_T we get,

$$q_V(V \mid W) \times \frac{\sum_T \prod_{D_i \in D_T} q_V(D_i \mid \mathrm{mp}_{\mathcal{G}}(D_i), W)}{\prod_{D_i \in D_T} q_V(D_i \mid \mathrm{mp}_{\mathcal{G}}(D_i), W)} \Big|_{T=t} = \Phi_T(q_V; \mathcal{G}) \Big|_{T=t}.$$

■

Corollary 23 (Identification via primal and dual fixing)

Proof As we always assume the presence of an underlying hidden variable DAG model from which we derive the marginal distribution $p(V)$ and the corresponding ADMG $\mathcal{G}(V)$ where T is p-fixable, it is easy to interpret p-fixing and d-fixing as yielding the post-intervention distribution $p(V(t))$. That is, $\Phi_T(p; \mathcal{G})|_{T=t} = \Delta_T(p; \mathcal{G}) = p(V(t))$. $\psi(t)$ is then obtained from $p(V(t))$ by marginalizing over $V \setminus T$. ■

Lemma 24 (Commutativity of p-fixing)

Proof The presence of an underlying hidden variable DAG model that gives rise to the marginal distribution $p(V)$ and ADMG $\mathcal{G}(V)$ makes it so that commutativity immediately follows from results in (Tian and Pearl, 2002b). Alternatively, the presence of the underlying DAG also allows us to interpret execution of two valid p-fixing sequences as steps in an identification procedure that both yield the same post-intervention distribution $p(V(s))$. ■

Lemma 25 (Identification via a sequence of p-fixing)

Proof Consider the p-fixing sequence $(S_1, \ldots, S_p, T)$. The operation $\Phi_{S_1}(p(V); \mathcal{G})$ gives us the post-intervention distribution $p(V(s_1))$. Applying $\Phi_{S_2}(p(V(s_1)))$ gives us $p(V(s_1, s_2))$, and so on. Applying the whole p-fixing sequence gives us $p(V(s_1, \ldots, s_p, t))$. Summing over

all non-fixed vertices and multiplying by Y gives us $\mathbb{E}[Y(s_1,\ldots,s_p,t)] = \sum_{V\setminus(S\cup T)} Y \times p(V(s_1,\ldots,s_p,t)) = \sum_{V\setminus(S\cup T)} Y \times \Phi_{S\cup T}(p(V);\mathcal{G})|_{T=t}$. However, since the vertices S have been chosen such that all directed (causal) paths from S_i to Y in $\mathcal{G}$ must pass through T, it is the case that $Y(s_1,\ldots,s_p,t) = Y(t)$ (Malinsky et al., 2019). Therefore,

$$\mathbb{E}[Y(s_1,\ldots,s_p,t)] = \mathbb{E}[Y(t)] = \sum_{V\setminus\{S\cup T\}} Y \times \Phi_{S\cup T}(p(V);\mathcal{G})\Big|_{T=t} = \psi(t).$$

■

Lemma 26 (Nonparametric IF in models of a CADMG)

Proof This follows from the fact that Eq. 36 holds under any statistical model [6]. Consider the statistical model over variables $\widetilde{V}$ where the density of $p(\widetilde{V})$ is assumed to belong to the class $\{q_{\tilde{V}}(\tilde{V} \mid Z=z;\eta);\eta\in\Gamma\}$. Even though we typically use $p(.)$ to denote densities, the kernel $q_{\tilde{V}}(\tilde{V} \mid Z=z)$ is indeed a valid probability density as well. As mentioned in the main draft, the kernel $q_{\tilde{V}}(\tilde{V}|Z)$ is a mapping from values in Z to normalized densities over $\tilde{V}$, i.e., $\sum_{\tilde{v}\in\tilde{V}} q_{\tilde{V}}(v \mid z) = 1$. Conditioning and marginalization operations in kernels are defined in the usual way. Our parameter of interest in this statistical model is $\psi(t) \equiv \mathbb{E}_{q_{\tilde{V}}}[Y(t)]$, where the expectation is taken with respect to the density $q_{\tilde{V}}(\tilde{V} \mid Z)$. We can partition the set of nodes $\widetilde{V}$ into three disjoint sets: $\widetilde{V} = \{\widetilde{\mathbb{C}},\widetilde{\mathbb{M}},\widetilde{\mathbb{L}}\}$, where

$$\widetilde{\mathbb{C}} = \{C_i \in \widetilde{V} \mid C_i \prec T\}, \quad \widetilde{\mathbb{L}} = \{L_i \in \widetilde{V} \mid L_i \in D_T, L_i \preceq T\}, \quad \widetilde{\mathbb{M}} = \{M_i \in \widetilde{V} \mid M_i \notin \widetilde{\mathbb{C}}\cup\widetilde{\mathbb{L}}\}.$$

The parameter $\psi(t)$ is identified via Eq.30 as follows.

$$\psi(t) = \sum_{V\setminus T} Y \times \prod_{M_i\in\widetilde{\mathbb{M}}} q_{\tilde{V}}(M_i \mid \mathrm{mp}_{\mathcal{G}}(M_i), Z=z) \sum_{T} \prod_{L_i\in\widetilde{\mathbb{L}}} q_{\tilde{V}}(L_i \mid \mathrm{mp}_{\mathcal{G}}(L_i), Z=z) \times q_{\tilde{V}}(\widetilde{\mathbb{C}})\Big|_{T=t}$$

If we replace, $q_{\tilde{V}}(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i), Z=z)$ with $p(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i))$, then the form of the above identifying functional is exactly the same as the one in Eq. 14. According to Eq. 36, the corresponding IF for $\psi(t)$ should take the same form as the IF in Theorem 9, where $p(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i))$ is replaced with $q_{\tilde{V}}(V_i \mid \mathrm{mp}_{\mathcal{G}}(V_i), Z=z)$. ■

Theorem 27 (Reweighted estimating equations)

Proof Let $p(V)$ be the joint density over V that district factorizes with respect to an ADMG $\mathcal{G}$, and $q_{\widetilde{V}}(\widetilde{V} \mid Z)$ be the kernel associated with the CADMG $\mathcal{G}(\widetilde{V},Z)$ that is obtained via fixing variables in Z according to a valid p-fixing sequence (note that $V = \{\widetilde{V},Z\}$.) Define $q(\widetilde{V},Z) = q_{\widetilde{V}}(\widetilde{V} \mid Z) \times p^*(Z)$ to be the joint density associated with the ADMG $\widetilde{\mathcal{G}} \equiv \mathcal{G}(\widetilde{V}\cup Z)$, where $p^*(Z)$ is any normalized function of Z.

6. with sufficient smoothness conditions.

Our target of inference is $\psi(t) \equiv \mathbb{E}[Y(t)]$ where the expectation is taken with respect to the observed data distribution $p(V)$. To find an estimating equation for $\psi(t)$, we use an intermediate target $\psi_q(t) \equiv \mathbb{E}_q[Y(t)]$, which is defined with respect to the ADMG $\widetilde{\mathcal{G}}$. Let $\mathbb{E}_q[\widetilde{U}] = 0$ be the estimating equation for $\psi_q(t)$ where $\widetilde{U}$ is the influence function for $\psi_q(t)$ given by Lemma 26. We know that $\mathbb{E}_q[\widetilde{U}] = 0$. Therefore,

$$\begin{aligned}
\mathbb{E}_q[\widetilde{U}] &= \sum_{\widetilde{V},Z} \widetilde{U} \times q(\widetilde{V}, Z) = \sum_{\widetilde{V},Z} \widetilde{U} \times q_{\widetilde{V}}(\widetilde{V} \mid Z) \times p^*(Z) \\
&= \sum_{\widetilde{V},Z} \widetilde{U} \times q_{\widetilde{V}}(\widetilde{V} \mid Z) \times p^*(Z) \times \frac{\prod_{Z_i \in Z} \pi_{Z_i}}{\prod_{Z_i \in Z} \pi_{Z_i}} \\
&= \sum_{\widetilde{V},Z} \frac{p^*(Z)}{\prod_{Z_i \in Z} \pi_{Z_i}} \times \widetilde{U} \times p(V) = \mathbb{E}\Big[\frac{p^*(Z)}{\prod_{Z_i \in Z} \pi_{Z_i}} \times \widetilde{U}\Big] = 0.
\end{aligned}$$

As we have shown that $\mathbb{E}\Big[\frac{p^*(Z)}{\prod_{Z_i \in Z} \pi_{Z_i}} \times \widetilde{U}\Big] = 0$, yields a consistent and doubly robust estimator for $\psi_q(t)$, conditional on the correct specification of $\pi_{Z_i}, \forall Z_i \in Z$, it exhibits the same properties for $\psi(t)$ as $\psi(t) = \psi_q(t)$. ■

Theorem 29 (Soundness and completeness of nested IPW)

Proof Soundness of the algorithm implies that when our algorithm succeeds, the subsequent identifying functional for $\psi(t)$ is correct. Completeness implies, that when the algorithm fails, the target parameter $\psi(t)$ is not identifiable within the model.

Soundness

We first prove soundness of the algorithm. That is, when Algorithm 2 does not fail, $\psi(t)$ is indeed $\psi(t)_{\text{nested}}$. The algorithm does not fail when all districts $D \in \mathcal{D}^*$ are intrinsic in $\mathcal{G}$. Note that $\mathcal{D}^*$ is a subset of the districts in $\mathcal{G}_{Y^*}$. However, by construction of $\mathcal{D}^*$, the remaining districts in $\mathcal{G}_{Y^*}$ are those that do not have any overlap with D_T. We now show that such districts are always intrinsic in $\mathcal{G}$.

Consider a district $D \in \mathcal{D}(\mathcal{G}_{Y^*})$ such that $D \cap D_T = \emptyset$. The district D forms a subset of a larger district in $\mathcal{G}$, say $D' \in \mathcal{D}(\mathcal{G})$. Due to results in (Tian and Pearl, 2002a), we know that D' is always intrinsic. If $D = D'$ then the result immediately follows. Otherwise, In the CADMG $\phi_{V \setminus D'}(\mathcal{G})$, there exists at least one vertex D_i in D' not in Y^*, that has no children. This is because all directed paths from D_i to vertices in Y^* must go through T and since T is not in D', all incoming edges to T have been deleted. The only other way D_i may not be childless is if there existed a cycle in $\mathcal{G}$, which is a contradiction. Thus, such a vertex D_i is always fixable and furthermore, fixing it corresponds to the marginalization operation $\sum_{D_i} q_{D'}(D' \mid \text{pa}_{\mathcal{G}}(D'))$ (Richardson et al., 2017). Once D_i is fixed, another vertex D_j that is in D' but not in Y^* becomes childless. Applying this argument inductively, we see that all $D_i \in D'$ such that $D_i \notin Y^*$ are fixable through marginalization under a reverse topological order. Hence for districts D in $\mathcal{G}_{Y^*}$ that do not overlap with D_T, the set $D = D' \setminus \{D_i \in D' \mid D_i \notin Y^*\}$ is always intrinsic. Thus, Algorithm 2 succeeds when all districts in $\mathcal{G}_{Y^*}$ are intrinsic.

We now show that under this condition, $\psi(t)_{\text{nested}} \equiv \mathbb{E}_{p^\dagger}\left[\frac{\mathbb{I}(T=t)}{p(T|\text{mp}_{\mathcal{G}}(T))} \times Y\right] = \psi(t)$. By definition, we have

$$\psi(t)_{\text{nested}} = \sum_V p(V) \times \prod_{D^* \in \mathcal{D}^*} \frac{q_{D^*}(D^* \mid \text{pa}_{\mathcal{G}}(D^*))}{\prod_{D_i^* \in D^*} p(D_i^* \mid \text{mp}_{\mathcal{G}}(D_i^*))} \times \frac{\mathbb{I}(T=t)}{p(T \mid \text{mp}_{\mathcal{G}}(T))} \times Y.$$

The districts of $\mathcal{G}$ can be partitioned into three sets. $\mathcal{D}_T$ is the district in $\mathcal{G}$ that contains T (with all elements in $\mathcal{D}^*$, if any, subsets of D_T). $\mathcal{D}'$ is the set of districts in $\mathcal{G}$, excluding D_T, that overlap with Y^*. $\mathcal{D}^z$ is the set of districts in $\mathcal{G}$, excluding D_T, that do not overlap with Y^*. The observed distribution $p(V)$ then Tian factorizes as,

$$p(V) = \prod_{D^z \in \mathcal{D}^z} q_{D^z}(D^z \mid \text{pa}_{\mathcal{G}}(D^z)) \times \prod_{D' \in \mathcal{D}'} q_{D'}(D' \mid \text{pa}_{\mathcal{G}}(D')) \times q_{D_T}(D_T | \text{pa}_{\mathcal{G}}(D_T)).$$

By results in Tian and Pearl (2002a), $q_{D_T}(D_T \mid \text{pa}_{\mathcal{G}}(D_T))$ is identified as $\prod_{D_i \in D_T} p(D_i \mid \text{mp}_{\mathcal{G}}(D_i))$ (for any topological ordering). Since every element in $\mathcal{D}^*$ is a subset of D_T, and since vertices in $D_T \setminus \bigcup_{D^* \in \mathcal{D}^*}$ precede vertices $D_T \cap \bigcup_{D^* \in \mathcal{D}^*} = D_T \cap Y^*$ in the ordering, we have

$$\psi(t)_{\text{nested}} = \sum_V \prod_{D^z \in \mathcal{D}^z} q_{D^z}(D^z \mid \text{pa}_{\mathcal{G}}(D^z)) \times \prod_{D' \in \mathcal{D}'} q_{D'}(D' \mid \text{pa}_{\mathcal{G}}(D')) \times \prod_{D^* \in \mathcal{D}^*} q_{D^*}(D^* \mid \text{pa}_{\mathcal{G}}(D^*))$$
$$\times \sum_{D_T \cap Y^*} q_{D_T}(D_T | \text{pa}_{\mathcal{G}}(D_T)) \times \frac{\mathbb{I}(T=t)}{p(T \mid \text{mp}_{\mathcal{G}}(T))} \times Y.$$

Since T is the last element in the ordering in $D_T \setminus Y^*$, we further have:

$$\psi(t)_{\text{nested}} = \sum_{Y^*} \sum_{V \setminus Y^*} \prod_{D^z \in \mathcal{D}^z} q_{D^z}(D^z \mid \text{pa}_{\mathcal{G}}(D^z)) \times \prod_{D' \in \mathcal{D}'} q_{D'}(D' \mid \text{pa}_{\mathcal{G}}(D')) \times \prod_{D^* \in \mathcal{D}^*} q_{D^*}(D^* \mid \text{pa}_{\mathcal{G}}(D^*))$$
$$\times \sum_{(D_T \cap Y^*) \cup \{T\}} q_{D_T}(D_T | \text{pa}_{\mathcal{G}}(D_T)) \times \mathbb{I}(T=t) \times Y.$$

Consider applying marginalization of elements in $V \setminus Y^*$ to $\psi(t)_{\text{nested}}$ above in the reverse topological ordering on $V \setminus Y^*$. Districts in $\mathcal{G}$ partition V and so, by definition of $\mathcal{D}^*, \mathcal{D}'$ and D_T, elements in $\mathcal{D}^z \cup \{D' \setminus Y^* : D' \in \mathcal{D}'\} \cup \{D_T \setminus (Y^* \cup \{T\})\}$ partition $V \setminus Y^*$. This partition, and the fact that marginalizations are processed in reverse topological order, means that at every stage, the variable to be summed occurs in precisely one place in the expression. This implies that the result of the overall summation of $V \setminus Y^*$ yields:

$$\psi(t)_{\text{nested}} = \sum_{Y^*} \prod_{D' \in \mathcal{D}'} \sum_{D' \setminus Y^*} q_{D'}(D' \mid \text{pa}_{\mathcal{G}}(D')) \times \prod_{D^* \in \mathcal{D}^*} q_{D^*}(D^* \mid \text{pa}_{\mathcal{G}}(D^*)) \times \mathbb{I}(T=t) \times Y$$

By definition, $q_{D^*}(D^* \mid \text{pa}_{\mathcal{G}}(D^*)) \equiv \phi_{V \setminus D^*}(p(V); \mathcal{G}(V))$. Since every D' in $\mathcal{D}'$ is a top level district in $\mathcal{G}$, there exists a valid fixing sequence on $V \setminus D'$. Further, in the CADMG $\phi_{V \setminus D'}(\mathcal{G}(V))$, any element in $D' \setminus Y^*$ cannot be an ancestor of an element in $D' \cap Y^*$ (if a directed path not through T existed from an element V_i in D' to an element in $D' \cap Y^*$,

then V_i must itself be in $D' \cap Y^*$, while a directed path from V_i to $D' \cap Y^*$ through T disappears in $\phi_{V \setminus D'}(\mathcal{G}(V))$ since T is outside D'. Consequently fixing elements $D' \setminus Y^*$ in reverse topological order in $\phi_{V \setminus D'}(\mathcal{G}(V))$ and $\phi_{V \setminus D'}(p(V), \mathcal{G}(V))$ is equivalent to marginalizing those variables. As a result, for every $D' \in \mathcal{D}'$, $\sum_{D' \setminus Y^*} q_{D'}(D' \mid \text{pa}_{\mathcal{G}}(D')) = \phi_{V \setminus (D' \cap Y^*)}(p(V); \mathcal{G}(V))$. Our conclusion follows:

$$\psi(t)_{\text{nested}} = \sum_{Y^*} \prod_{D \in \mathcal{D}(\mathcal{G}_{Y^*})} \phi_{V \setminus D}(p(V); \mathcal{G}) \times Y \Big|_{T=t} = \psi(t).$$

Completeness

Follows trivially as we have shown the failure condition of Algorithm 2 to be equivalent to the failure condition of the identification algorithm in (Richardson et al., 2017) which is known to be sound and complete. ■

Theorem 30 (Augmented nested IPW for any identifiable $\psi(t)$)

Proof The parameter of inference $\psi(t)$ is identified via Algorithm 2 as

$$\psi(t) = \mathbb{E}_{p^\dagger}\Big[\frac{\mathbb{I}(T=t)}{p(T \mid \text{mp}_{\mathcal{G}}(T))} \times Y\Big],$$

where expectation is defined with respect to

$$p^\dagger(V) = p(V) \times \prod_{D \in \mathcal{D}^*} \frac{q_D(D \mid \text{pa}_{\mathcal{G}}(D))}{\prod_{D_i \in D} p(D_i \mid \text{mp}_{\mathcal{G}}(D_i))}.$$

Let $U^\dagger_{\psi_t}$ denote the nonparametric IF corresponding to $\psi(t)$ defined w.r.t $p^\dagger$. We can think of $p^\dagger$ as a kernel. Consequently, according to Lemma 26, $U^\dagger_{\psi_t}$ has the following form

$$U^\dagger_{\psi_t} = \mathbb{E}_{p^\dagger}\Big[\frac{\mathbb{I}(T=t)}{p(T \mid \text{mp}_{\mathcal{G}}(T))} \times \big(Y - \mathbb{E}[Y \mid T=t, \text{mp}_{\mathcal{G}}(T)]\big) + \mathbb{E}[Y \mid T=t, \text{mp}_{\mathcal{G}}(T)] - \psi(t)\Big]$$

and by properties of IFs, it has mean zero: $\mathbb{E}_{p^\dagger}[U^\dagger_{\psi_t}] = 0$. The objective here is to evaluate the expectation with respect to the observed data distribution $p(V)$.

$$\begin{aligned}
\mathbb{E}_{p^\dagger}\big[U^\dagger_{\psi_t}\big] &= \sum_V U^\dagger_{\psi_t} \times p^\dagger(V) \\
&= \sum_V U^\dagger_{\psi_t} \times p(V) \times \prod_{D \in \mathcal{D}^*} \frac{q_D(D \mid \text{pa}_{\mathcal{G}}(D))}{\prod_{D_i \in D} p(D_i \mid \text{mp}_{\mathcal{G}}(D_i))} \\
&= \mathbb{E}\Big[\prod_{D \in \mathcal{D}^*} \frac{q_D(D \mid \text{pa}_{\mathcal{G}}(D))}{\prod_{D_i \in D} p(D_i \mid \text{mp}_{\mathcal{G}}(D_i))} \times U^\dagger_{\psi_t}\Big] \\
&= \mathbb{E}\Big[\prod_{D \in \mathcal{D}^*} \rho_D \times U^\dagger_{\psi_t}\Big] = 0.
\end{aligned}$$

where $\rho_D = \frac{q_D(D|\text{pa}_{\mathcal{G}}(D))}{\prod_{D_i \in D} p(D_i|\text{mp}_{\mathcal{G}}(D_i))}$.

■

References

R. Ayesha Ali, Thomas S. Richardson, and Peter L. Spirtes. Markov equivalence for ancestral graphs. *The Annals of Statistics*, 37(5B):2808–2837, 2009.

Heejung Bang and James M. Robins. Doubly robust estimation in missing data and causal inference models. *Biometrics*, 61(4):962–973, 2005.

Peter J. Bickel, Chris A.J. Klaassen, Yaacov Ritov, and Jon A. Wellner. *Efficient and adaptive estimation for semiparametric models*, volume 4. Johns Hopkins University Press Baltimore, 1993.

Stephen P. Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge University Press, 2004.

Marco Carone, Alexander R. Luedtke, and Mark J. van der Laan. Toward computerized efficient estimation in infinite-dimensional models. *Journal of the American Statistical Association*, 114(527):1174–1190, 2019.

Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney K. Newey, and James M. Robins. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21, 2018.

George B. Dantzig, Lester R. Ford Jr., and Delbert R. Fulkerson. A primal-dual algorithm. *Linear Equalities and Related Systems, Annals of Mathematics Study*, (38):171–181, 1956.

Mathias Drton. Discrete chain graph models. *Bernoulli*, 15(3):736–753, 2009.

Robin J. Evans. Margins of discrete Bayesian networks. *The Annals of Statistics*, 46(6A): 2623–2656, 2018.

Robin J. Evans and Thomas S. Richardson. Maximum likelihood fitting of acyclic directed mixed graphs to binary data. *arXiv preprint arXiv:1203.3479*, 2012.

Robin J. Evans and Thomas S. Richardson. Smooth, identifiable supermodels of discrete DAG models with latent variables. *Bernoulli*, 25(2):848–876, 2019.

Constantine E. Frangakis, Tianchen Qian, Zhenke Wu, and Iván Díaz. Deductive derivation and Turing-computerization of semiparametric efficient estimation. *Biometrics*, 71(4): 867–874, 2015.

Isabel R. Fulcher, Ilya Shpitser, Stella Marealle, and Eric J. Tchetgen Tchetgen. Robust inference on indirect causal effects. *arXiv preprint arXiv:1711.03611*, 2017.

Sander Greenland, Judea Pearl, and James M. Robins. Causal diagrams for epidemiologic research. *Epidemiology*, pages 37–48, 1999.

Jinyong Hahn. On the role of the propensity score in efficient semiparametric estimation of average treatment effects. *Econometrica*, pages 315–331, 1998.

Jaroslav Hájek. A characterization of limiting distributions of regular estimates. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 14(4):323–330, 1970.

Frank R. Hampel. The influence curve and its role in robust estimation. *Journal of the American Statistical Association*, 69(346):383–393, 1974.

Leonard Henckel, Emilija Perković, and Marloes H. Maathuis. Graphical criteria for efficient total effect estimation via adjustment in causal linear models. *arXiv preprint arXiv:1907.02435*, 2019.

Yimin Huang and Marco Valtorta. Pearl's calculus of intervention is complete. In *Proceedings of the 22nd Conference on Uncertainty in Artificial Intelligence*, pages 13–16, 2006.

Paul Hünermund and Elias Bareinboim. Causal inference and data-fusion in econometrics. *arXiv preprint arXiv:1912.09104*, 2019.

Yonghan Jung, Jin Tian, and Elias Bareinboim. Estimating causal effects using weighting-based estimators. In *Proceedings of the 34th Conference on Artificial Intelligence*. AAAI Press, 2020.

Edward H. Kennedy, Zongming Ma, Matthew D. McHugh, and Dylan S. Small. Nonparametric methods for doubly robust estimation of continuous treatment effects. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 79(4):1229–1245, 2017.

Steffen L. Lauritzen. *Graphical models*. Oxford, U.K.: Clarendon, 1996.

David G. Luenberger. *Optimization by vector space methods*. John Wiley & Sons, 1997.

Daniel Malinsky, Ilya Shpitser, and Thomas S. Richardson. A potential outcomes calculus for identifying conditional path-specific effects. In *Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics*, pages 3080–3088, 2019.

Whitney K. Newey. Semiparametric efficiency bounds. *Journal of Applied Econometrics*, 5 (2):99–135, 1990.

Judea Pearl. *Probabilistic reasoning in intelligent systems*. Morgan Kauffman, 1988.

Judea Pearl. Causal diagrams for empirical research. *Biometrika*, 82(4):669–688, 1995.

Judea Pearl. *Causality*. Cambridge University Press, 2009.

Emilija Perković, Johannes Textor, Markus Kalisch, and Marloes H. Maathuis. A complete generalized adjustment criterion. In *Proceedings of the 31st Conference on Uncertainty in Artificial Intelligence*, pages 682–691. AUAI Press, 2015.

Thomas S. Richardson. Markov properties for acyclic directed mixed graphs. *Scandinavian Journal of Statistics*, 30(1):145–157, 2003.

Thomas S. Richardson and James M. Robins. Single world intervention graphs (SWIGs): A unification of the counterfactual and graphical approaches to causality. *Center for the Statistics and the Social Sciences, University of Washington Series. Working Paper*, 128 (30), 2013.

Thomas S. Richardson and Peter L. Spirtes. Ancestral graph Markov models. *The Annals of Statistics*, 30(4):962–1030, 2002.

Thomas S. Richardson, Robin J. Evans, James M. Robins, and Ilya Shpitser. Nested Markov properties for acyclic directed mixed graphs. *arXiv preprint arXiv:1701.06686*, 2017.

James M. Robins. A new approach to causal inference in mortality studies with a sustained exposure period – application to control of the healthy worker survivor effect. *Mathematical Modelling*, 7(9-12):1393–1512, 1986.

James M. Robins. Marginal structural models versus structural nested models as tools for causal inference. In *Statistical Models in Epidemiology, The Environment, and Clinical Trials*, pages 95–133. Springer, 2000.

James M. Robins, Andrea Rotnitzky, and Lue Ping Zhao. Estimation of regression coefficients when some regressors are not always observed. *Journal of the American Statistical Association*, 89(427):846–866, 1994.

Andrea Rotnitzky and Ezequiel Smucler. Efficient adjustment sets for population average treatment effect estimation in non-parametric causal graphical models. *arXiv preprint arXiv:1912.00306*, 2019.

Ilya Shpitser and Judea Pearl. Identification of joint interventional distributions in recursive semi-Markovian causal models. In *Proceedings of the 21st National Conference on Artificial Intelligence*, 2006.

Ilya Shpitser, Tyler J. VanderWeele, and James M. Robins. On the validity of covariate adjustment for estimating causal effects. In *Proceedings of the 26th Conference on Uncertainty in Artificial Intelligence*, pages 527–536. AUAI Press, 2010.

Ilya Shpitser, Robin J. Evans, Thomas S. Richardson, and James M. Robins. Introduction to nested markov models. *Behaviormetrika*, 41(1):3–39, 2014.

Ilya Shpitser, Robin J. Evans, and Thomas S. Richardson. Acyclic linear SEMs obey the nested Markov property. In *Proceedings of the 34th Conference on Uncertainty in Artificial Intelligence*, 2018.

Peter L. Spirtes, Clark N. Glymour, Richard Scheines, David Heckerman, Christopher Meek, Gregory Cooper, and Thomas S. Richardson. *Causation, prediction, and search*. MIT press, 2000.

Masashi Sugiyama, Motoaki Kawanabe, and Pui Ling Chui. Dimensionality reduction for density ratio estimation in high-dimensional spaces. *Neural Networks*, 23(1):44–59, 2010.

Robert Tarjan. Depth-first search and linear graph algorithms. *SIAM Journal on Computing*, 1(2):146–160, 1972.

Jin Tian and Judea Pearl. A general identification condition for causal effects. In *Proceedings of the 18th National Conference on Artificial Intelligence*, pages 567–573. American Association for Artificial Intelligence, 2002a.

Jin Tian and Judea Pearl. On the testable implications of causal models with hidden variables. In *Proceedings of the 18th Conference on Uncertainty in Artificial Intelligence*, pages 519–527, 2002b.

Anastasios Tsiatis. *Semiparametric theory and missing data*. Springer Science & Business Media, 2007.

Mark J. van der Laan and Sherri Rose. *Targeted learning: causal inference for observational and experimental data*. Springer Science & Business Media, 2011.

Aad W. van der Vaart. *Asymptotic statistics*, volume 3. Cambridge University Press, 2000.

Benito van der Zander, Johannes Textor, and Maciej Liśkiewicz. Efficiently finding conditional instruments for causal inference. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence*, 2015.

Thomas Verma and Judea Pearl. Equivalence and synthesis of causal models. In *Proceedings of the 6th Conference on Uncertainty in Artificial Intelligence*, 1990.

Thomas C. Williams, Cathrine C. Bach, Niels B. Matthiesen, Tine B. Henriksen, and Luigi Gagliardi. Directed acyclic graphs: a tool for causal studies in paediatrics. *Pediatric research*, 84(4):487–493, 2018.

Hui Zhao, Zhongguo Zheng, and Baijun Liu. On the Markov equivalence of maximal ancestral graphs. *Science in China Series A: Mathematics*, 48(4):548–562, 2005.