

GRAPHICAL MODELS OF ENTANGLED MISSINGNESS

BY RANJANI SRINIVASAN^{1,a}, ROHIT BHATTACHARYA^{2,b}, RAZIEH NABI^{3,c}, ELIZABETH L. OGBURN^{4,d}, AND ILYA SHPITSER^{5,e}

¹*Department of Electrical and Computer Engineering, Johns Hopkins University, ranjani@jhu.edu*

²*Department of Computer Science, Williams College, rb17@williams.edu*

³*Department of Biostatistics and Bioinformatics, Emory University, razieh.nabi@emory.edu*

⁴*Department of Biostatistics, Johns Hopkins University, eogburn@jhsph.edu*

⁵*Department of Computer Science, Johns Hopkins University, ilyas@cs.jhu.edu*

Despite the growing interest in causal and statistical inference for settings with data dependence, few methods currently exist to account for missing data in dependent data settings; most classical missing data methods in statistics and causal inference treat data units as independent and identically distributed (i.i.d.). We develop a graphical modeling based framework for causal inference in the presence of *entangled missingness*, defined as missingness with data dependence. We distinguish three different types of *entanglements* that can occur, supported by real-world examples. We give sound and complete identification results for all three settings. We show that existing missing data models may be extended to cover entanglements arising from (1) *target law dependence* and (2) *missingness process dependence*, while those arising from (3) *missingness interference* require a novel approach. We demonstrate the use of our entangled missingness framework on synthetic data. Finally, we discuss how, subject to a certain reinterpretation of the variables in the model, our model for missingness interference extends missing data methods to novel missing data patterns in i.i.d. settings.

1. Introduction. Missing data is a ubiquitous problem: data may be missing due to survey non-response, dropout or loss-to-followup, imperfect data collection, or other complications. Data records missing *systematically* can substantially bias subsequent analyses if not properly addressed. Most of the literature on missing data treats data units as independent and identically distributed (i.i.d.) [1–16]. This assumption is reasonable in many types of problems where interactions between units are negligible and can be ignored. However, in recent years there has been increasing recognition that data in many settings are subject to dependence across units and, in particular, to *interference*, where variables measured on one unit may have a causal effect on those measured on another unit. The canonical example is infectious disease settings where treating one unit may have a protective effect on others, but more recently attention has turned to the myriad ways in which humans influence one another through their social networks (broadly construed to include familial, work, social, neighborhood, or online relationships) [17–45].

MSC2020 subject classifications: Primary 00X00, 00X00; secondary 00X00.

Keywords and phrases: First keyword, second keyword.

The current literature lags behind in recognizing that data dependence and missing data might occur simultaneously, with the exception of a few pieces of work: [46] investigated multiple imputation techniques for missing data in health data networks; [47] empirically characterized bias in the analysis of network data with systematic missing data; [48] used a likelihood-based modeling approach for health studies with partially observed data; [49] proposed estimation methods for network logistic regression models in the presence of missing data.

To our knowledge, this paper presents the first nonparametric¹ method to deal with settings with *entangled missingness*, i.e. with both missingness and data dependence, and also the first approach to such settings based on causal graphical models. We develop a graphical modeling framework to deal with three different types of entangled missingness termed *target law dependence*, *missingness process dependence*, and *missingness interference*. We derive sound and complete nonparametric identification results for all the three scenarios in which missingness and data dependence may interact. Our results offer full generality for causal estimands in the presence of entangled missingness.

Target law dependence occurs when the full data distribution exhibits either statistical or causal dependence. For example, consider a study of vaccine effectiveness in which one unit's vaccination status may help to protect their friends or family members from contracting an infectious disease. The underlying full data distribution is non-iid and failing to account for this dependence when imputing or otherwise controlling for missing data would generally introduce bias.

Missingness process dependence occurs when the missingness status of one unit may depend on (variables of) other units. For example, an individual's decision to participate in a survey may depend on their friends' or family members' decisions. Similarly, consider a mass public health surveillance effort, where demographic and family data are collected on a large population of individuals from which a smaller group is invited to participate in substudy where more detailed data are collected. An individual's choice to enroll in the substudy may depend both on their own characteristics and on the characteristics of, or choices made by, other individuals in their social network.

Missingness interference occurs when there exist multiple potential outcomes of a variable depending on whether outcomes of other units in the network are measured or not (corresponding to participation of these other units in the study.) Consider an evaluation of an event or program, e.g. a course evaluation, using a survey filled out by the attendees. Answers to survey questions in this case would certainly be unobserved for any individual who did not attend or who dropped out before the end of the course. However, because classes have social and collaborative components, the underlying values of responses of a particular individual, *had they attended*, will potentially differ depending on the attendance of that individual's peers. A similar story could be told about student performance assessments, where there is evidence that classroom composition affects the performance of individual students [50]. These settings have not previously been addressed in any literature on missing data or interference. We propose modeling multiple versions of underlying satisfaction or achievement variables in a way that is structurally similar to counterfactual variables or potential outcomes in causal inference.

The paper is organized as follows. We present a review of relevant background on graphical causal models, interference, and current graphical models for missing data in Sec. 2. We

¹Here, we use the term nonparametric to mean that we place no restrictions on the functional forms of relationships between variables.

introduce new graphical models representing entangled missingness in Sec. 3. Our new identification results for entangled missingness are in Sec. 4. In Sec. 5 we show that our model for missingness interference extends missing data methods to novel missing data patterns in i.i.d. settings. Sec. 6 illustrates the utility of our framework in simulations. We conclude the paper in Sec. 7 by outlining remaining open challenges for non-parametric identification of arbitrary estimands in settings with entangled missingness.

2. Graphical Causal Models, Interference, and Missing Data. Causal inference is traditionally concerned with *counterfactual* or *potential* outcomes like $Y(a)$, defined as the outcome that would have been observed if, possibly contrary to fact, the unit had received treatment or exposure value $A = a$. For simplicity we will assume binary treatment throughout, but our results generalize to multi-level and, with some additional subtleties, continuous A . Causal effects are typically defined as contrasts of counterfactual outcomes, e.g. $\beta := \mathbb{E}[Y(1) - Y(0)]$ is a comparison of the expected value of counterfactual outcomes in a world in which every unit receives treatment $A = 1$ compared with the expected value of counterfactual outcome in a world in which every unit receives treatment $A = 0$. Since potential outcomes are not necessarily observed in available data realizations, assumptions are needed to link the counterfactual distributions and the *observed data distribution*, from which available samples are actually drawn.

A standard set of assumptions used to express β as a functional of observed data are: (i) *consistency*, stating that the observed outcome realization is equal to the counterfactual outcome realization had treatment been set to the actually observed value, or $Y(a) = Y$ if $A = a$, (ii) *positivity*, stating that all treatment assignments, possibly conditioned on a set of baseline covariates \mathbf{C} , have positive support, or if $p(\mathbf{C} = \mathbf{c}) > 0$ then $p(a | \mathbf{C} = \mathbf{c}) > 0$ for all a in the support of A , and (iii) *conditional ignorability*, stating that potential outcomes are independent of the treatment assignment, possibly conditioned on a set of baseline covariates \mathbf{C} , or $Y(a) \perp\!\!\!\perp A | \mathbf{C}$ for all a . Under these assumptions, the parameter β of the full data distribution $p(Y(1), Y(0), A, \mathbf{C})$ is identified from the observed data distribution $p(Y, A, \mathbf{C})$ via the *adjustment functional*: $\mathbb{E}[\mathbb{E}[Y | A = 1, \mathbf{C}] - \mathbb{E}[Y | A = 0, \mathbf{C}]]$.

Conditional independence assumptions like conditional ignorability can be encoded by graphical models; this is the approach that we will take in this paper to present identification results in the presence of entangled missingness. We briefly introduce graphical models notation, terminology and concepts that will be used in this work. We have attempted to contain all required concepts within this section, but for a more thorough treatment of graphical models, see [51–53].

2.1. Graphical Causal Models. A directed acyclic graph (DAG) \mathcal{G} with vertex set \mathbf{V} representing variables \mathbf{V} , consists of only directed edges (\rightarrow), with no directed cycles. The statistical model of a DAG $\mathcal{G}(\mathbf{V})$ is the set of distributions $p(\mathbf{V})$ that factorize as $p(\mathbf{V}) = \prod_{V \in \mathbf{V}} p(V | \text{pa}_{\mathcal{G}}(V))$ where $\text{pa}_{\mathcal{G}}(V)$ denotes parents of V in $\mathcal{G}(\mathbf{V})$; U is a parent of V if $U \rightarrow V$, with $U, V \in \mathbf{V}$. All conditional independence restrictions in $p(\mathbf{V})$ encoded in the DAG $\mathcal{G}(\mathbf{V})$ can be read by examining paths in $\mathcal{G}(\mathbf{V})$ via d-separation rules [54].

Causal models of a DAG combine a generative model of $p(\mathbf{V})$ with the theory of interventions to yield distributions over counterfactual random variables. These models are often defined using *structural equations* (which are a mathematical formulation of an invariant causal mechanism), as follows. Each variable V in a causal model is determined from values of its parents $\text{pa}_{\mathcal{G}}(V)$ and an exogenous noise variable ϵ_V via a structural equation $f_V(\text{pa}_{\mathcal{G}}(V), \epsilon_V)$. Given an intervention on $V \in \mathbf{V}$ that sets V to v , the structural equation $f_V(\text{pa}_{\mathcal{G}}(V), \epsilon_V)$ is

replaced by one that outputs the constant value v . Structural equations and noise variables together produce the observed data distribution $p(\mathbf{V})$. If any subset $\mathbf{A} \subseteq \mathbf{V}$ are set to values \mathbf{a} , $p(\mathbf{Y}(\mathbf{a}))$ denotes the joint distribution over potential outcomes $\mathbf{Y} = \mathbf{V} \setminus \mathbf{A}$ to this intervention, written as $p(\mathbf{Y} \mid \text{do}(\mathbf{a}))$ in [51]. For the purposes of this paper, we will assume the structural equations and noise terms are such that the resulting $p(\mathbf{V})$ is a positive distribution.

A popular causal model called the *structural causal model*, or the *non-parametric structural equation model with independent errors (NPSEM-IE)* [55] assumes that the joint distribution of all exogenous terms are marginally independent: $p(\epsilon) = \prod_{V \in \mathbf{V}} p(\epsilon_V)$. The NPSEM-IE implies the DAG factorization of $p(\mathbf{V})$ with respect to $\mathcal{G}(\mathbf{V})$, and a modified DAG factorization known as the *g-formula* that identifies every joint intervention distribution as the following functional of the observed data distribution $p(\mathbf{V})$:

$$(1) \quad p(\mathbf{Y}(\mathbf{a})) = \prod_{Y \in \mathbf{Y}} p(Y \mid \text{pa}_{\mathcal{G}}(Y)) \Big|_{\mathbf{A}=\mathbf{a}},$$

for every $\mathbf{A} \subset \mathbf{V}$ and $\mathbf{Y} = \mathbf{V} \setminus \mathbf{A}$.

The conditionally ignorable model described earlier can be represented via the DAG shown in Fig. 1(a). With $\mathbf{V} = \{Y, A, C\}$, $\mathbf{A} = \{A\}$ and $\mathbf{V} \setminus \mathbf{A} = \{Y, C\}$ in (1), $p(Y(a), C(a)) = p(Y(a), C) = p(Y \mid A = a, C) \times p(C)$. As a consequence, the counterfactual mean $\mathbb{E}[Y(a)]$ is identified via the adjustment formula as $\mathbb{E}[Y(a)] = \sum_C \mathbb{E}[Y \mid A = a, C] \times p(C)$.² Equivalently, the interventional target may be obtained using an *inverse probability weighting (IPW)* functional $\mathbb{E}[\{\mathbb{I}(A = a)/p(A \mid C)\} \times Y]$, which involves using propensity score $p(A \mid C)$ weights [56]. The weights approach will be particularly useful when we consider identification in missing data models, where the missingness mechanism will be akin to the propensity scores in this functional (see Sec. 2.3).

Causal models with hidden variables may be associated with a DAG $\mathcal{G}(\mathbf{V} \cup \mathbf{H})$, where \mathbf{V} are observed and \mathbf{H} are unobserved variables. In such models not every causal effect is identified under the standard assumptions, and identification theory may be expressed on a latent projection graph $\mathcal{G}(\mathbf{V})$ [57], containing only vertices corresponding to observed variables. Latent projections are a type of mixed graph called an acyclic directed mixed graph (ADMG) containing directed edges, representing direct causation, and bidirected edges (\leftrightarrow) representing unobserved confounding. As an example, the latent projection of a DAG in Fig. 1(a) if C was an unobserved variable, is given in Fig. 1(b). Conditional independence statements from a distribution $p(\mathbf{V})$ associated with an ADMG $\mathcal{G}(\mathbf{V})$ may be read off by a generalization of d-separation called *m-separation* [58].

Identified causal effects in hidden variable models are given by the ID algorithm, which may be viewed as a generalization of the g-formula in the sense that causal effects identified by the ID algorithm are modified factorizations of ADMGs, just as the g-formula is the modified factorization of DAGs. For general identification results for average causal effects in causal models represented by DAGs and ADMGs, see [59–63].

2.2. Interference. Historically, causal inference approaches assumed all units to be causally and statistically independent. This is, however, frequently violated and we say that *interference* is present when one unit may causally affect other units [17–45]. Sometimes interference

²It is common practice in missing data and some causal inference literature to identify expectations as the target of interest, but we will discuss identification of distributions in this paper to better connect with causal graphical models identification literature.




Figure 1: (a) Example of a DAG representation in a conditionally ignorable model; (b) A latent projection of the DAG in (a) if C is an unobserved variable; (c) Example of a DAG illustration for interference in a dyad.

is more narrowly defined as the causal effect of one unit’s treatment on another’s outcome; our use of the term is more general and may apply to covariates, missingness indicators, treatments, or outcomes.

In this paper, in order to isolate issues of nonparametric identification from issues of statistical dependence, we will assume *partial interference*: interference may occur within collections or *blocks* of units of fixed finite size, but the blocks are mutually independent. We will additionally assume that the blocks are identically distributed. This simplifies our presentation considerably, but in principle all of our identification results can be extended to full interference or to blocks with heterogeneous structure for different realizations. We say that a full data estimand (such as the full data distribution) is identified if it can be expressed as a functional of the observed data distribution; this is agnostic to partial versus full interference. A separate but related issue concerns statistical estimability, which typically requires identification plus a large number of independent (or weakly dependent) observations, and non-i.i.d. blocks could pose a challenge for estimability. Although the purpose of this paper is identification rather than estimation, we note that in the setting of i.i.d. blocks of fixed size, estimators may be derived using standard techniques for i.i.d. data, simply by treating each *block of units* as a *unit* for the purposes of estimation. Traditional estimation results will hold immediately under an asymptotic regime that lets the number of blocks go to infinity while the block composition remains fixed and i.i.d. This suggests that we may choose to view each block either as a collection of units or as a single multivariate observation; we switch between these two views throughout.

Let there be b blocks, each block consisting of $m = n/b$ units, where n is the total number of units. Let $\mathbf{A} \equiv (A_1, \dots, A_m)$ be the vector of treatment assignments for units $i = 1, \dots, m$ in a block, and the m -dimensional vector \mathbf{a} be a realization in the support of \mathbf{A} . Similarly, let $\mathbf{Y} \equiv (Y_1, \dots, Y_m)$ and $\mathbf{C} \equiv (C_1, \dots, C_m)$ be the vector of outcomes and covariates respectively. Since units within the same block may interact, we must index a counterfactual outcome of unit i by interventions not only performed on unit i , but also on other units that belong to the same block. In particular, we define $Y_i(\mathbf{a})$ to be the counterfactual outcome of unit i , where the treatment vector \mathbf{A} is intervened on and set to \mathbf{a} . Compare this to the counterfactual $Y(a)$ in (1) for a setting with no interference, where the outcome of a unit is also a response to intervention setting A to a , but where A represents only the unit’s own treatment.

Counterfactuals $Y_i(\mathbf{a})$ that are indexed by treatments of every unit in the network may be used to define causal effects for individual units, or an average of effects for all units in the network. Parameters of interest in interference problems and related estimation strategies are described in detail in [64–67].

Ogburn and VanderWeele [64] extended causal models of DAGs from the i.i.d. setting to interference problems. Fig. 1(c) is a typical example of how DAGs are used to represent

interference, in this case in a block of 2 units (a dyad). In this DAG, the tuple (C_1, A_1, Y_1) corresponds to variables of unit 1 and the tuple (C_2, A_2, Y_2) corresponds to variables of unit 2. The presence of edges between these two tuples encodes the causal influence of one unit's variables on another unit's variables, i.e. the presence of interference.

The principles of graphical causal models generalize from the i.i.d. setting to settings with interference, the only difference being that the graph now represents an entire block of units rather than a single unit representing an i.i.d. realization. As noted above, it is just a matter of switching our lens between two views: that of a collection of units or of a single unit's multivariate observation. Further, we are able to extend graph-based identification algorithms from i.i.d settings to the partial interference setting.

2.3. Graphical Representation of Missing Data Models. A missing data model encodes assumptions about how missingness arises and how it relates to underlying variables of interest. The model is a set of distributions defined over a set of random variables $\{\mathbf{O}, \mathbf{Z}^{(1)}, \mathbf{R}, \mathbf{Z}\}$, where \mathbf{O} denotes the set of variables that are always observed, $\mathbf{Z}^{(1)}$ denotes the set of variables that are potentially missing, \mathbf{R} denotes the set of missingness indicators of the variables in $\mathbf{Z}^{(1)}$, and \mathbf{Z} denotes the set of *observed proxies* of the variables in $\mathbf{Z}^{(1)}$. Missingness indicator $R_Z \in \mathbf{R}$, corresponding to $Z^{(1)} \in \mathbf{Z}^{(1)}$, takes the value 1 when the variable Z is observed, and 0 when it is not. The observed proxy Z is deterministically defined as follows:

$$(2) \quad Z = \begin{cases} Z^{(1)} & \text{if } R_Z = 1, \\ ? & \text{if } R_Z = 0. \end{cases}$$

Using the terminology of causal models, each missingness indicator $R \in \mathbf{R}$ may be viewed as a treatment variable that can be intervened on, and each $Z \in \mathbf{Z}$ as an observed outcome. Thus, $Z^{(1)}$ is a counterfactual random variable had we, possibly contrary to fact, intervened and set the corresponding missingness indicator R_Z to 1. We can thus view (2) as missing data equivalent of the consistency assumption in causal inference.

We call the distribution $p(\mathbf{O}, \mathbf{Z}^{(1)}, \mathbf{R})$ the *full law*, the distribution $p(\mathbf{O}, \mathbf{Z}^{(1)})$ the *target law*, the distribution $p(\mathbf{O}, \mathbf{Z}, \mathbf{R})$ the *observed data law*, and the conditional distribution $p(\mathbf{R} | \mathbf{Z}^{(1)}, \mathbf{O})$ the *missingness process*, or the *missing data propensity score*.

Typically we are interested in a functional of the target law, often of the form $h(\mathbf{O}, \mathbf{Z}^{(1)})$, and the goal of missing data methods is to nonparametrically identify such functionals in terms of the observed data law. This requires assumptions, and it is common to assume the missingness consistency assumption (2) and a missing data version of the positivity assumption:

$$(3) \quad p(\mathbf{R} = 1 | \mathbf{Z}^{(1)}, \mathbf{O}) > 0.$$

As is the case with causal inference, consistency and positivity do not suffice for nonparametric identification of target law parameters, but the addition of certain independence assumptions may suffice for nonparametric identification of the full law and therefore of any functionals thereof. The current literature considers missing data models where such assumptions can be encoded on a DAG $\mathcal{G}(\mathbf{V})$, where vertices \mathbf{V} correspond to random variables in $\mathbf{O} \cup \mathbf{Z}^{(1)} \cup \mathbf{R} \cup \mathbf{Z}$, and certain additional restrictions are placed on $\mathcal{G}(\mathbf{V})$: (i) each $Z \in \mathbf{Z}$ has only two parents, R_Z and $Z^{(1)}$, and (ii) variables in \mathbf{R} cannot point to variables in $\mathbf{O} \cup \mathbf{Z}^{(1)}$. Restriction (i) is imposed by definition of Z via (2). In order to distinguish the deterministic relations implied by (2) from probabilistic relations, we draw edges pointing into $Z \in \mathbf{Z}$ in gray. Restriction (ii) is imposed to ensure that, while changes in R_Z cause changes in the observed proxies $Z \in \mathbf{Z}$,




Figure 2: A missing data DAG example corresponding to (a) a MCAR model, (b) a MAR model, and (c) a MNAR model. (d) A missing data MAR model with $Z \equiv \{A, Y\}$.

they do not result in changes to the underlying full data $Z^{(1)} \in \mathbf{Z}^{(1)}$. Our proposed framework, outlined in the following sections, relaxes the edge restrictions in (ii); see Sec. 5 for more details.

A missing data DAG model is a set of distributions defined over variables in $\mathbf{O} \cup \mathbf{Z}^{(1)} \cup \mathbf{R} \cup \mathbf{Z}$ that factorize with respect to a DAG with the above restrictions, as follows:

$$(4) \quad p(\mathbf{O}, \mathbf{Z}^{(1)}, \mathbf{R}, \mathbf{Z}) = \prod_{V \in \mathbf{O} \cup \mathbf{Z}^{(1)} \cup \mathbf{R}} p(V \mid \text{pa}_{\mathcal{G}}(V)) \times \prod_{Z \in \mathbf{Z}} p(Z \mid R_Z, Z^{(1)}).$$

Examples of missing data DAGs are shown in Fig. 2: Fig. 2(a) corresponds to a *missing completely at random* (MCAR) model where $R_Y \perp\!\!\!\perp O, Y^{(1)}$, Fig. 2(b) corresponds to a *missing at random* (MAR) model where $R_Y \perp\!\!\!\perp Y^{(1)} \mid O$, and Fig. 2(c) corresponds to a *missing not at random* (MNAR) model where neither independence holds.

If a missing data model contains hidden variables \mathbf{H} , it may be represented by the latent projection ADMG $\mathcal{G}(\mathbf{O}, \mathbf{Z}, \mathbf{Z}^{(1)}, \mathbf{R})$ of the DAG $\mathcal{G}(\mathbf{H}, \mathbf{O}, \mathbf{Z}, \mathbf{Z}^{(1)}, \mathbf{R})$, just like in the case of hidden variable causal models. In such models, the target of inference is some function of the full data law $p(\mathbf{O}, \mathbf{Z}^{(1)}, \mathbf{R})$ where \mathbf{H} is marginalized out.

There is a substantial literature on non-graphical approaches to dealing with missing data. A simple approach is to use only fully observed rows in complete-case analysis, justifiable only if the underlying mechanism is MCAR [1]. For MAR scenarios, many approaches including expectation maximization [3, 68, 69], multiple imputation [70, 71], inverse probability weighting [72, 73] have been developed. For MNAR problems, often parametric or semiparametric restrictions have been imposed on the underlying data distribution and missingness selection model, such that they yield identification [3, 74–79]. In this work we use graphical models to derive conditions for nonparametric identification.

The full law in a missing data model is identified *if and only if* the missingness mechanism is identified by some functional of the observed data law $g(p(\mathbf{R}, \mathbf{O}, \mathbf{Z}))$, because

$$(5) \quad p(\mathbf{O}, \mathbf{Z}^{(1)}, \mathbf{R}) = p(\mathbf{O}, \mathbf{Z}^{(1)}) \times \underbrace{p(\mathbf{R} \mid \mathbf{O}, \mathbf{Z}^{(1)})}_{g(p(\mathbf{R}, \mathbf{O}, \mathbf{Z}))} = \frac{p(\mathbf{O}, \mathbf{Z}^{(1)}, \mathbf{R} = 1)}{p(\mathbf{R} = 1 \mid \mathbf{O}, \mathbf{Z}^{(1)})} \times \underbrace{p(\mathbf{R} \mid \mathbf{O}, \mathbf{Z}^{(1)})}_{g(p(\mathbf{R}, \mathbf{O}, \mathbf{Z}))}.$$

Note that this functional mirrors the IPW-based estimator we outlined for causal models earlier, with the missingness mechanism serving as the inverse weight.

As a quick example, the full law in the MAR missing data DAG model of Fig. 2(d) is identified because the missingness mechanism $p(R_A, R_Y \mid \mathbf{C}, A^{(1)}, Y^{(1)}) = p(R_A, R_Y \mid \mathbf{C})$ is a function of observed data.

The hierarchy of MCAR, MAR and MNAR mechanisms holds relevance in the context of identification simply due to the fact that if missingness is not at random, the full law may not

be identified. Hence, when we provide illustrative examples and graphs in Sec. 3, we will often point out what kind of mechanism is at play.

3. Graphical Representation of Entangled Missingness. Having discussed existing models for missing data and interference, we now propose three types of *entanglements* by which the two phenomena can occur together. In naming and describing these phenomena, we hope to avert the common practice of simply ignoring missing data in settings with interference.

The following definitions pertain to the law $p(\mathbf{R}, \mathbf{O}, \mathbf{Z}^{(1)}, \mathbf{Z})$ with all hidden variables marginalized out.

1. Target Law Dependence (\mathcal{E}_1): Counterfactuals $\mathbf{Z}_i^{(1)}$ and $\mathbf{Z}_j^{(1)}$ of units i, j in a block depend on each other. This means the target law $p(\mathbf{O}, \mathbf{Z}^{(1)})$ does not factorize into unit-specific marginal distributions. Such a situation arises in problems where the underlying distribution, had there been no missing data, exhibits data dependence or interference.
2. Missingness Process Dependence (\mathcal{E}_2): Missingness indicators \mathbf{R}_i for unit i depend on variables corresponding to unit j , and thus the missingness process $p(\mathbf{R} | \mathbf{O}, \mathbf{Z}^{(1)})$ does not factorize into unit-specific factors. Similar situations have been considered in causal inference applications [80].
3. Missingness Interference (\mathcal{E}_3): Observed proxies \mathbf{Z}_i for unit i depend on missingness indicators \mathbf{R}_j for unit j .³ Under \mathcal{E}_3 , the counterfactual variable corresponding to variable Z_i must be indexed not only by R_{Z_i} being set to 1, but also by the missingness status of (variables of) other units. For instance the counterfactual $Z_i^{(R_{Z_i}=1, R_{Z_j}=0)}$ would correspond to variable Z of unit i , had it been observed ($R_{Z_i} = 1$) and had variable Z of unit j been missing.

These three types of entanglement can occur in any combination, and the seven possible situations where at least one entanglement occurs collectively describe all scenarios where data dependence and missingness occur together. In describing these scenarios in detail, we use the notation \mathcal{E}_i to indicate that the i -th kind of entanglement is present and use $\bar{\mathcal{E}}_i$ to indicate its absence. For example, we will use $\mathcal{E}_1\bar{\mathcal{E}}_2\bar{\mathcal{E}}_3$ to denote a scenario when there is target law dependence but no missingness process dependence or missingness interference.

We first present entanglements *without* missingness interference in Sec. 3.1, as it turns out that these forms of entanglements share a graphical representation with existing i.i.d. missing data models. In Sec. 3.2, we describe scenarios *with* missingness interference; these require the development of new notation and graphical representations.

3.1. Entanglements Without Missingness Interference. Consider studying the effect of drug A on disease status Y from a community hospital database. Assume hospital records are incomplete and missingness indicators R_{A_i}, R_{Y_i} denote whether the treatment and disease status values were recorded for patient i , with counterfactuals $A_i^{(1)}, Y_i^{(1)}$, denoting the true (but possibly unrecorded) values of these variables. Baseline covariates, denoted by \mathbf{C}_i for unit i , are always observed and include age and alcohol consumption. In this example, $\mathbf{Z} \equiv \{A, Y\}$ and $\mathbf{O} \equiv \{\mathbf{C}\}$. For simplicity of presentation we will work with dyads.

³This can be interpreted as an extension of classical interference, where one unit's treatment affects another unit's outcome, except that the treatment has been replaced by the missingness indicator. Hence the name missingness interference.

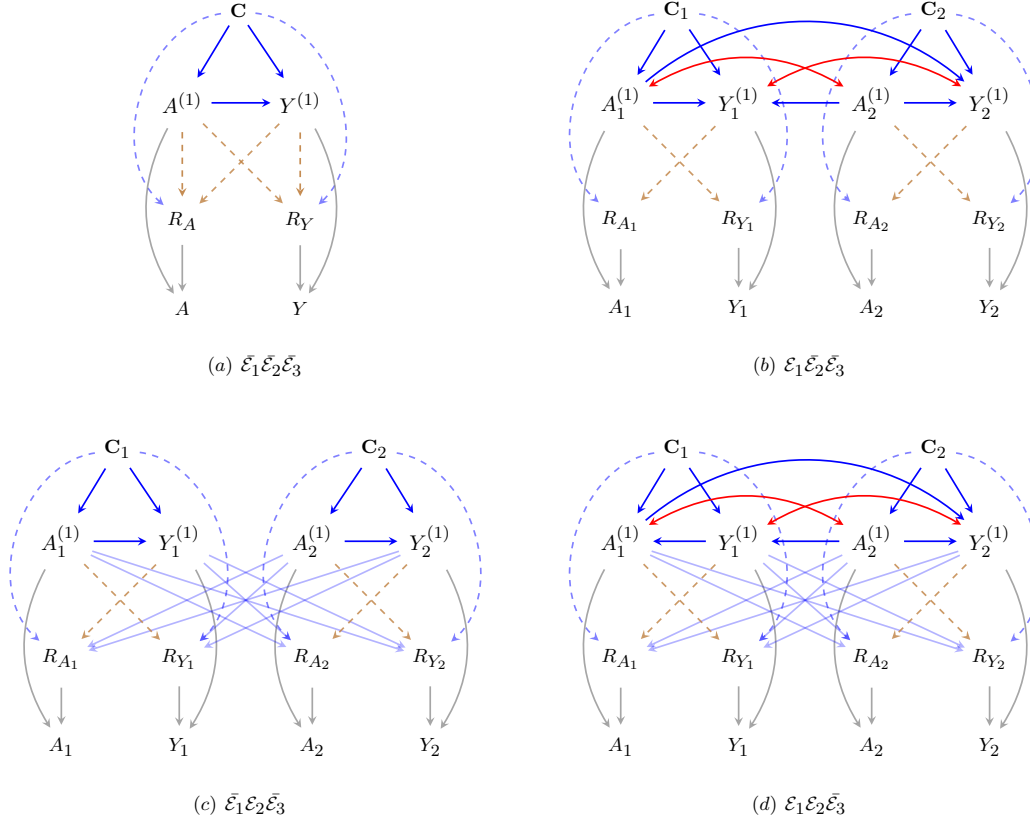


Figure 3: Four scenarios representing all possible ways target law dependence and missingness process dependence may arise in a dyadic partial interference setting without missingness interference arising.

In order to provide a contrast between settings with i.i.d. missing data models and those with entanglement, we first start with a scenario which involves no kind of entanglement ($\bar{\mathcal{E}}_1 \bar{\mathcal{E}}_2 \bar{\mathcal{E}}_3$) within this setup and build the examples further to illustrate how entanglements might arise.

Scenario 1.1 ($\bar{\mathcal{E}}_1 \bar{\mathcal{E}}_2 \bar{\mathcal{E}}_3$). When the disease being investigated is not contagious, and the missingness process, pertaining to non-response, deficiencies in data collection, and so on, does not exhibit dependence across patients, we might be able to assume that a patient’s record does not influence the record of any other patient, i.e., data is i.i.d.

The graph in Fig. 3(a) depicts this situation. Including different set of edges in Fig. 3(a) yields either MCAR where $R_A, R_Y \perp\!\!\!\perp A^{(1)}, Y^{(1)}, \mathbf{C}$ (if only solid edges are included), MAR where $R_A, R_Y \perp\!\!\!\perp A^{(1)}, Y^{(1)} \mid \mathbf{C}$ (if in addition to the solid lines, blue dashed edges are also included)⁴, or MNAR (if at least one brown dashed edge is included). Edges $A \rightarrow R_A$ and $Y \rightarrow R_Y$ are called *self-censoring* or *self-masking* edges [12, 81, 82].

Scenario 1.2 ($\mathcal{E}_1 \bar{\mathcal{E}}_2 \bar{\mathcal{E}}_3$). Suppose the outcome Y is an infectious disease, like COVID-19. Since patients can infect one another, data on patients residing in the same household or geographic area are likely to exhibit dependence in outcomes. In addition, surging infections

⁴Note that the MAR version is identical to that in Fig. 2(d) as all of Scenario 1.1 pertains to i.i.d. settings.

may lead to treatment shortages, which will lead to dependence among treatments due to *allocational interference* [64]; geographically localized treatment shortages occurred at several points early in the COVID-19 pandemic. Finally, successful prevention or treatment of an infectious outcome for unit i may influence outcomes for other units j by preventing potential transmission from i to j . Despite these complications, missingness mechanisms remain independent across units.

Fig. 3(b) depicts this setting. Bidirected arrows encode dependence (due to shared hidden causes) between $A_1^{(1)}$ and $A_2^{(1)}$ and between $Y_1^{(1)}$ and $Y_2^{(1)}$, respectively. For simplicity of visualization, we omit connections among C_i and $\{A_j, Y_j\}$. The edge subgraph of Fig. 3(b) consisting only of solid edges corresponds to a MCAR model, since $\mathbf{R} \perp\!\!\!\perp \mathbf{Y}^{(1)}, \mathbf{A}^{(1)}, \mathbf{C}$ by the m-separation criterion [58]. Similarly, the edge subgraph of Fig. 3(b) consisting of all edges other than dashed brown corresponds to a MAR model, since $\mathbf{R} \perp\!\!\!\perp \mathbf{Y}^{(1)}, \mathbf{A}^{(1)} \mid \mathbf{C}$. Finally, including dashed brown edges in the MAR model yields an MNAR model.

Scenario 1.3 ($\bar{\mathcal{E}}_1 \mathcal{E}_2 \bar{\mathcal{E}}_3$). Assume that we are again dealing with a non-contagious disease, and that treatment allocation for one unit is not influenced by any features of other units. However, the missingness process for one unit does depend on other units. For example, data includes members of a family who may jointly decide to change healthcare providers, resulting in missing treatment and outcome data for all members of the family.

Fig. 3(c) depicts this setting, and the light blue edges between the counterfactuals $A_1^{(1)}, Y_1^{(1)}$ and missingness indicators R_{A_2}, R_{Y_2} , as well as $Y_2^{(1)}, Y_2^{(1)}$ and R_{A_1}, R_{Y_1} represent missingness process dependence. Edges from C_2 to variables of unit 1, and from C_1 to variables of unit 2 are omitted for simplicity. An MCAR version of this scenario simply reduces to scenario $\bar{\mathcal{E}}_1 \bar{\mathcal{E}}_2 \bar{\mathcal{E}}_3$, representing i.i.d. data. An example of a MAR mechanism in this scenario corresponds to the absence of light solid blue edges in Fig. 3(c) and instead having edges $C_1 \rightarrow R_{A_2}, C_1 \rightarrow R_{Y_2}, C_2 \rightarrow R_{A_1}$ and $C_2 \rightarrow R_{Y_1}$ (not shown). Missingness process dependence can also be represented as a bidirected edge between variables of unit i and missingness indicators of unit j , coming from a hidden common cause of the corresponding variable and the missingness indicator.

Scenario 1.4 ($\mathcal{E}_1 \mathcal{E}_2 \bar{\mathcal{E}}_3$). This scenario simply combines complications in scenarios 1.2 and 1.3, such that both target law dependence and missingness process dependence are present. The edges can be interpreted exactly as before, with scenarios 1.2 and 1.3. Specifically, the light blue solid edges between counterfactuals and missingness indicators make the missingness process dependence MNAR. Instead, if the dependence was on observed covariates, the mechanism would be MAR. Dashed edges, both brown and blue operate within a unit and simply induce dependence that is MNAR or MAR, respectively, without involving other units.

Finally, we note that the graphs in Fig. 3 can be extended in a straightforward fashion to a network where blocks are not just dyads, but are of arbitrary size m , to illustrate the different types of entanglement discussed here. Identification results corresponding to the settings described until now are discussed in Sec. 4.1.

3.2. Entanglements With Missingness Interference. In some situations where missingness and data dependence occur together, observed realizations of underlying variables are influenced by missingness indicators of multiple units, leading to a missing data analogue of *interference problems* in causal inference.

Consider the following example, where investigators are interested in evaluating outcomes of a health and wellness program that involves a diet and workout regimen of participants, based on their baseline health indicators. Participation in the study is through voluntary registration. Bob and Anne are neighbors, and see the advertisement for enrollment. Imagine two different scenarios: (1) Both Bob and Anne feel strongly about the program and enroll together. Since they participate together, they are able to easily plan their meals and physical activity, and motivate each other. As a result, Anne records a successful outcome, improved metabolism and better BMI. Alternatively, (2) Bob enrolls and his baseline variables are collected but he does not actually participate. Anne goes through the program but finds it hard to remain motivated and follow effectively, and her results are not as good. Bob's final outcome (BMI) is not measured as he does not participate. In this example, Anne has two possible hypothetical program results under the two different scenarios: one where Bob participates, and the other where he does not. To put it another way, the observed results for Anne depend on the missingness status of Bob's results in the study. Similar arguments might be made for Bob.

Formally, for Anne (the unit indexed by 1), the outcome BMI has two possible values: $Y_1^{(r_{Y_1}=1, r_{Y_2}=0)}$ for when Bob does not participate and his BMI is, hence, not recorded ($r_{Y_2} = 0$) and $Y_1^{(r_{Y_1}=1, r_{Y_2}=1)}$ for when Bob does participate and his BMI is recorded ($r_{Y_2} = 1$). We make two simplifications here. One, we assume that the only missing variable for any unit is the outcome, and hence denote r_{Y_1} by r_1 and r_{Y_2} by r_2 . This assumption is only being made to simplify this example, and will be relaxed in the more general case (see Sec. 3.2.1). Second, we drop r_1 from the superscript and denote the counterfactuals as $Y_1^{(1, r_2=0)}$ and $Y_1^{(1, r_2=1)}$. This latter simplification can be done without ambiguity because Anne's outcomes are defined only when $r_1 = 1$. Anne's missingness status R_1 and observable BMI Y_1 are recorded in the study data. This observed outcome is not just a function of her counterfactuals $Y_1^{(1, r_2=0)}$, $Y_1^{(1, r_2=1)}$ and her missingness status R_1 , but also the missingness status of Bob's outcome, R_2 . One can think of R_2 as a switch, selecting which counterfactual, $Y_1^{(1, r_2=0)}$ or $Y_1^{(1, r_2=1)}$, is realized in the observation. By symmetry in this particular example, we assume the same is true for Bob. Thus, we have the following relationships between observed values Y_1, Y_2 of neighbors Bob and Anne and the corresponding counterfactuals:

$$Y_1 \leftarrow \begin{cases} (1 - r_2)Y_1^{(1, r_2=0)} + r_2Y_1^{(1, r_2=1)} & \text{if } r_1 = 1 \\ ? & \text{if } r_1 = 0 \end{cases}$$

$$Y_2 \leftarrow \begin{cases} (1 - r_1)Y_2^{(1, r_1=0)} + r_1Y_2^{(1, r_1=1)} & \text{if } r_2 = 1 \\ ? & \text{if } r_2 = 0 \end{cases}$$

Since this setting involves multiple counterfactual versions of Y_1 , it cannot be captured by missing data models described so far, nor described by the graphs in Fig. 3, which correspond to the full data law with a single counterfactual version of every variable.

We call, for every unit i , the set of units whose missingness indicators index the counterfactuals of unit i , as *affectors* of i , or $\text{aff}(i)$, in the network, inspired by the use of the word in neuroscience [83]. As one can imagine, this relationship need not be symmetric, and one unit can be the affector of another, without the reverse being true. With these considerations in mind, we now describe the four scenarios that arise when missingness interference may occur.

Scenario 2.1 ($\bar{\mathcal{E}}_1 \bar{\mathcal{E}}_2 \mathcal{E}_3$). First, consider a setting where there is neither target law dependence, nor missingness process dependence, but only missingness interference. Let us assume that

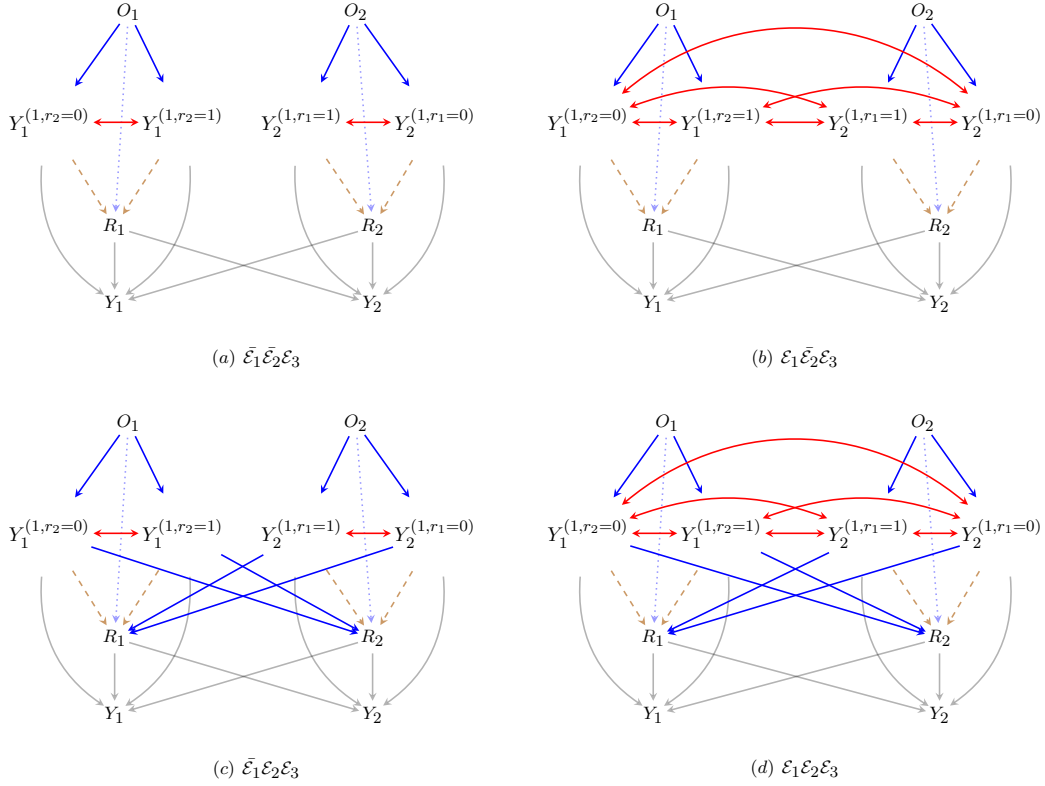


Figure 4: Four scenarios representing all possible ways target law dependence and missingness process dependence may arise in a dyadic partial interference setting when missingness interference is present. We assume that only outcomes Y_1, Y_2 are missing, and use R_i as a shorthand for R_{Y_i} .

diets and exercise regimens are recommended to participants, based solely on their measured baseline covariates. As an approximation, we assume that a person's BMI is affected only by their actions, and hence the target laws are not directly influenced by each other. As described, Anne is more likely to workout if her friend Bob does ($R_j \rightarrow Y_i$), affecting her BMI. We assume that $Y_i^{(1, r_j=0)}$ and $Y_i^{(1, r_j=1)}$ are associated for any unit i via a hidden common cause H_i , resulting in a bidirected (\leftrightarrow) edge between the counterfactuals in the latent projection graphs. See Fig. 4(a). Brown dashed edges depict self-censoring mechanisms.

Scenario 2.2 ($\mathcal{E}_1 \bar{\mathcal{E}}_2 \mathcal{E}_3$). Now, assume that Bob and Anne are siblings. One might imagine that the effect of diet and exercise on an individual are mediated by their genetics, which siblings share, giving rise to target law dependence. Additionally, assume that those conducting the study randomize who gets to participate, from the list of participants who register (and hence, whether or not someone is included in the study is independent of their sibling). This can give rise to a structure as shown in Fig. 4(b). Adherence to prescribed regimens is still dependent on a person's sibling being included in the study, giving rise to two underlying counterfactual BMIs for each unit. Here too, the dependence between counterfactuals is assumed to be due to a shared common cause H , resulting in target law dependence being represented by bidirected edges.

Scenario 2.3 ($\bar{\mathcal{E}}_1 \mathcal{E}_2 \mathcal{E}_3$). Assume Bob and Anne are not genetically related. Further, assume that knowledge of Bob's potential BMI if he enrolled in the study (or not) makes Anne more likely to enroll herself, and the investigators include everyone who registers, without randomly

selecting a subset of the entries, giving rise to missingness process dependence. This would result in a graph like in Fig. 4(c). The blue edge $Y_2^{(1, r_1=1)} \rightarrow R_1$ introduces a more general type of self-censoring, which we call *affector-censoring*. This corresponds to the situation where a counterfactual variable of a unit (which is indexed by its affector’s missingness status) censors that very same missingness status.

Scenario 2.4 ($\mathcal{E}_1\mathcal{E}_2\mathcal{E}_3$). Finally assume a situation where Bob and Anne share genetics and are also likely to enroll based on their knowledge of how fit the other person is likely to become as part of the study, combining all three possible types of entanglements in one setup. This scenario is illustrated via the graph in Fig. 4(d).

3.2.1. *Beyond Dyads*. Interactions in many realistic settings are not restricted to dyads and hence we describe counterfactuals (and graphs) for networks with missingness interference. For simplicity of presentation, we assume the following: given units i and j , if $R_{Z_j} \rightarrow Z_i$ exists, then $R_{Z'_j} \rightarrow Z_i$ also exists for all $Z'_j \in \mathbf{Z}_j$. We write this assumption as $j \in \text{aff}(i)$. This enables us to index counterfactuals corresponding to unit i by \mathbf{R}_j , which determines whether \mathbf{Z}_j is missing, instead of individual variables $Z_j \in \mathbf{Z}_j$. We further assume within a unit, the edge $R_{Z_i} \rightarrow Z'_i$ is absent, for all distinct variables Z, Z' . This simply allows us to replace $R_{Z_i} = 1$ by 1, in the index of counterfactuals of unit i . To understand the implications of removing these restrictions, see the note at the end of this section.

We assume a network of n individuals. The full law is defined over variables in $\bigcup_{i=1}^n \mathbf{O}_i \cup \tilde{\mathbf{Z}}_i \cup \mathbf{R}_{\mathbf{Z}_i} \cup \mathbf{H}_i$ where $\tilde{\mathbf{Z}}_i \equiv \{\mathbf{Z}_i^{(1, \mathbf{r}_{\text{aff}(i)})} : \mathbf{r}_{\text{aff}(i)} \text{ a value of } \mathbf{R}_{\text{aff}(i)}\}$ represents counterfactual versions of variables \mathbf{Z}_i for unit i , the set $\text{aff}(i)$ refers to the units that are affectors of unit i , and $\mathbf{R}_{\text{aff}(i)} \equiv \{\mathbf{R}_{\mathbf{Z}_j} : j \in \text{aff}(i)\}$ where $\mathbf{R}_{\mathbf{Z}_j}$ refers to the set of all missingness indicators R_{Z_j} for all variables $Z_j \in \mathbf{Z}_j$ of unit j . In other words, variables $\mathbf{Z}_i^{(1, \mathbf{r}_{\text{aff}(i)})} \in \tilde{\mathbf{Z}}_i$, for every value $\mathbf{r}_{\text{aff}(i)}$, reads “the value assumed by variables \mathbf{Z}_i , had they been observed, and had observability status of the affectors of i in the network been set to $\mathbf{r}_{\text{aff}(i)}$.”

As a generalization of (2), every observed proxy $Z_i \in \mathbf{Z}_i$ is defined as:

$$(6) \quad Z_i \leftarrow \begin{cases} Z_i^{(1, \mathbf{R}_{\text{aff}(i)})} & \text{if } R_{Z_i} = 1 \\ ? & \text{if } R_{Z_i} = 0 \end{cases},$$

where $\mathbf{R}_{\text{aff}(i)}$ in the superscript is interpreted to mean interventions that set $\mathbf{R}_{\text{aff}(i)}$ to their naturally occurring values. Just as in causal interference, all counterfactuals corresponding to one variable Z_i are dependent on each other in general. Next, we discuss graphs for these networks.

The full data distribution can be represented by a DAG, which contains a vertex for every $Z_i^{(1, \mathbf{r}_{\text{aff}(i)})}$, as well as every element in $\bigcup_{i=1}^n \mathbf{O}_i \cup \mathbf{Z}_i \cup \mathbf{R}_{\mathbf{Z}_i} \cup \mathbf{H}_i$. We also assume that the set of all counterfactuals of one variable Z_i of unit i share a common parent $H_{Z_i} \in \mathbf{H}_i$, since they are all associated. This DAG obeys two restrictions, which generalize restrictions discussed in Sec. 2.3: (i) variables in \mathbf{R} cannot point to variables in $\bigcup_{i=1}^n \mathbf{O}_i \cup \{\mathbf{Z}_i^{(1, \mathbf{r}_{\text{aff}(i)})}\}$, and (ii) every element Z_i in \mathbf{Z} has all its corresponding counterfactuals, affectors of i , and R_{Z_i} as parents. Restriction (ii) is a consequence of the consistency property in (6).

A latent projection ADMG, obtained by projecting out fully hidden variables $\bigcup_{i=1}^n \mathbf{H}_i$ from this DAG, will be our preferred choice of graphical representation, just as in Figures 3 and 4. A simple example of such an ADMG for general networks is shown in Fig. 5, where each unit i has only one missing variable Z_i . Dependence among counterfactuals $\mathbf{Z}_i^{(1, \mathbf{r}_{\text{aff}(i)})}$ is represented by bidirected edges, obtained by latent projecting H_i out. In this example, every unit is an




Figure 5: A simple example of a general network with n units undergoing missingness interference. This graph represents an MCAR model, $\mathbf{R} \perp\!\!\!\perp \mathbf{O}, \mathbf{Z}_i^{(1, \mathbf{r}_{\text{aff}(i)})}$.

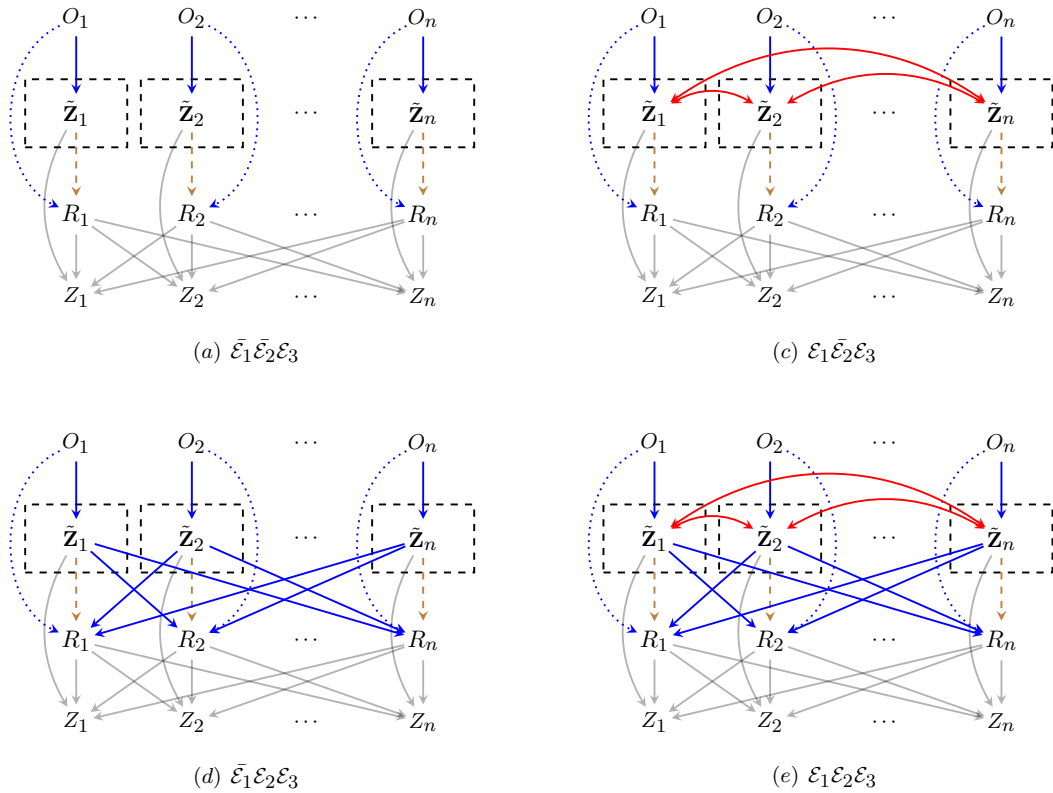


Figure 6: A general network with n units featuring missingness interference. (a) - (d) enumerate all possible ways target law dependence and missingness process dependence may arise, as a generalization of the dyad in Fig. 4.

affector of every other unit. For compactness in representation, we have chosen to show the counterfactuals only for two of the units.

When the incoming and outgoing edges of vertices $Z_i^{(1, \mathbf{r}_{\text{aff}(i)})}$ and $Z_i^{(1, \mathbf{r}'_{\text{aff}(i)})}$ are identical for all \mathbf{r}, \mathbf{r}' in the graph, we can make the graphical representation even more compact by introducing the vertex \tilde{Z}_i corresponding to the set $\tilde{\mathbf{Z}}_i$ in place of all its elements $Z_i^{(1, \mathbf{r}_{\text{aff}(i)})}$, for all possible indexing corresponding to $\mathbf{r}_{\text{aff}(i)}$. Such graphs, which we will denote by $\tilde{\mathcal{G}}$, are shown in Fig. 6. In particular, the graph composed of solid edges in Fig. 6(a) is a condensed representation of the graph in Fig. 5, and any vertex \tilde{Z}_i (shown in a dashed black box) compactly represents the full set of vertices $\{Z_i^{(1, \mathbf{r}_{\text{aff}(i)}=0, \dots, 0)}, Z_i^{(1, \mathbf{r}_{\text{aff}(i)}=1, \dots, 0)}, \dots, Z_i^{(1, \mathbf{r}_{\text{aff}(i)}=1, \dots, 1)}\}$, all connected by bidirected edges. In $\tilde{\mathcal{G}}$, we interpret the edge $V \rightarrow \tilde{Z}_i$, for any node V , to mean that V influences all the counterfactuals $Z_i^{(1, \mathbf{r}_{\text{aff}(i)})}$ in $\tilde{\mathbf{Z}}_i$. Similarly $V \leftarrow \tilde{Z}_i$ is interpreted to mean that all counterfactuals $Z_i^{(1, \mathbf{r}_{\text{aff}(i)})}$ in $\tilde{\mathbf{Z}}_i$ influence V .

Further, it is helpful to view $\tilde{\mathcal{G}}$ in Fig. 6(a) as the generalized version of Fig. 4(a), where there is no target law dependence or missingness process dependence but only missingness interference. The graph can encode MCAR (only solid edges), MAR (with blue dotted edges) and MNAR (with brown dashed edges) processes, just as before. Graphs in Fig. 6(b)-(d) are generalized versions of graphs in Fig. 4(b)-(d). Graphs $\tilde{\mathcal{G}}$ are useful for illustrative purposes as they are compact, but as we show in Sec. 4, identification is more straightforward in models where counterfactuals do *not* share identical edges. Hence, we will use graphs \mathcal{G} , like those in Fig. 5 (and not $\tilde{\mathcal{G}}$ in Fig. 6), in the remainder of the paper.

We have looked at examples where all potentially missing variables of a unit are either all observed or all missing, meaning only a single missingness indicator per unit is needed. In general, this may not be true. This gives rise to two different complications: (i) there could be edges from R_{Z_j} to Z'_i for some units i, j and variables $Z, Z' \in \mathbf{Z}$, and (ii) edges from a missingness indicator of a unit to another variable in that same unit (e.g. $R_{Z_i} \rightarrow Z'_i$) might be present. If (i) occurs, we cannot index counterfactuals simply by the missingness indicators of other units as we had done before, but of specific variables of those units. To be consistent with our prior definition, we call a unit the affector of another if *any* missingness indicator of the first unit indexes the counterfactual of the latter unit. If (ii) occurs, the missingness interference is within the unit, and not across units, making the data i.i.d. in the absence of other interactions. See Sec. 5 for additional discussion of this issue. It is worth pointing out that in both of these cases, the graphical structure and the mathematical framework remains the same as what we have discussed in this section so far. What changes is the interpretation of the relationships between units and their recorded variables.

4. Identification in Entangled Missingness Models. Having set up the notation and graphical framework for entanglement, we next discuss identification in these models.

4.1. Identification Without Missingness Interference. We recognized in Sec. 3.1 that the (graphical) models in the absence of missingness interference are essentially identical to existing i.i.d. models for missing data, with the only difference arising from how we view the smallest unit of investigation - whether it is one individual or a set of individuals interacting within a block.

It is common practice in missing data literature to treat expectations of the form $\beta = \mathbb{E}[h(\mathbf{Z}^{(1)}, \mathbf{O})]$, for some known function h of $\mathbf{Z}^{(1)}, \mathbf{O}$, as parameters of interest. In this paper, however, we will discuss identification of distributions in keeping with causal graphical

modeling identification literature. Specifically, we focus on non-parametric identification of the full law $p(\mathbf{Z}^{(1)}, \mathbf{O}, \mathbf{R})$, because if the full law is nonparametrically identified from the observed data so is any functional of it.

A *sound* and *complete* algorithm for full law identification in i.i.d. missing data DAG models with fully observed variables has been proposed in [12].⁵ The authors also provide full law identification in missing data models with hidden variables in the same work. For both types of graphs, DAGs and ADMGs, their identification criterion relies on the notion of the Markov blanket [54, 62]: in an ADMG \mathcal{G} , the Markov blanket of a vertex V , denoted by $\text{mb}_{\mathcal{G}}(V)$, consists of all variables sharing an edge with V or with a collider path to V ; a collider path is path where all vertices on the path are colliders of the form $\rightarrow o \leftarrow, \leftrightarrow o \leftarrow, \leftrightarrow o \leftrightarrow$. In a DAG, which is an ADMG with no bidirected edges, the Markov blanket reduces to the set of variables with an edge in common with V , and variables that share a child with V . We provide a brief description of their results in Theorem 1 and Theorem 2 below, as they apply directly to entangled missingness settings without missingness interference.

THEOREM 1. [Full law identification in DAGs [12]]

In a missing data model, represented by DAG \mathcal{G} , the full law $p(\mathbf{R}, \mathbf{Z}^{(1)}, \mathbf{O})$ is identified if and only if $Z^{(1)} \notin \text{mb}_{\mathcal{G}}(R_Z), \forall Z^{(1)} \in \mathbf{Z}^{(1)}$. Thus, for the full law to be identified, no edge of the form $Z^{(1)} \rightarrow R_Z$ can be present (no *self-censoring*) and no structure of the form $Z^{(1)} \rightarrow R_{Z'} \leftarrow Z$ can be present (no *colluders*). The identifying functional is given by Eq. 5, where the missingness mechanism $p(\mathbf{R} | \mathbf{Z}^{(1)}, \mathbf{O})$ is given by an odds ratio parameterization [84]:

$$(7) \quad \frac{1}{\sigma(\mathbf{Z}^{(1)}, \mathbf{O})} \times \prod_{k=1}^K p(R_k | \mathbf{R}_{-k} = 1, \mathbf{Z}^{(1)}, \mathbf{O}) \times \prod_{k=2}^K \text{OR}(R_k, R_{\prec k} | R_{\succ k} = 1, \mathbf{Z}^{(1)}, \mathbf{O}),$$

where $\mathbf{R}_{-k} = \mathbf{R} \setminus R_k, R_{\prec k} = \{R_1, \dots, R_{k-1}\}, R_{\succ k} = \{R_{k+1}, \dots, R_K\}$,

$$\begin{aligned} & \text{OR}(R_k, R_{\prec k} | R_{\succ k} = 1, \mathbf{Z}^{(1)}, \mathbf{O}) \\ &= \frac{p(R_k | R_{\succ k} = 1, R_{\prec k}, \mathbf{Z}^{(1)}, \mathbf{O})}{p(R_k = 1 | R_{\succ k} = 1, R_{\prec k}, \mathbf{Z}^{(1)}, \mathbf{O})} \times \frac{p(R_k = 1 | \mathbf{R}_{-k} = 1, \mathbf{Z}^{(1)}, \mathbf{O})}{p(R_k | \mathbf{R}_{-k} = 1, \mathbf{Z}^{(1)}, \mathbf{O})}, \end{aligned}$$

and $\sigma(\mathbf{Z}^{(1)}, \mathbf{O}) = \sum_{\mathbf{r}} \{ \prod_{k=1}^K p(r_k | \mathbf{R}_{-k} = 1, \mathbf{Z}^{(1)}, \mathbf{O}) \times \prod_{k=2}^K \text{OR}(r_k, r_{\prec k} | R_{\succ k} = 1, \mathbf{Z}^{(1)}, \mathbf{O}) \}$ is the normalizing function.

Consider the i.i.d. graph in Fig. 2 (d). The odds ratio parameterization of $p(R_A, R_Y | \mathbf{C}, A^{(1)}, Y^{(1)})$ is identical to its regular DAG parameterization (shown in Sec. 2.3) since $p(R_A | R_Y = 1, \mathbf{C}, A^{(1)}, Y^{(1)}) = p(R_A | \mathbf{C})$, $p(R_Y | R_A = 1, \mathbf{C}, A^{(1)}, Y^{(1)}) = p(R_Y | \mathbf{C})$, and $\text{OR}(R_A, R_Y | Y^{(1)}, A^{(1)}, \mathbf{C}) = \sum_{R_Y, R_A} p(R_A | \mathbf{C}) \times p(R_Y | \mathbf{C}) = 1$, and the normalizing term is one. Self-censoring edges $A^{(1)} \rightarrow R_A$ and $Y^{(1)} \rightarrow R_Y$ prevent identification of the full law in Fig. 3 (a).

The following theorem summarizes the full law identification results in [12] for ADMGs obtained as latent projections of hidden variable missing data DAGs.

THEOREM 2. [Full law identification in ADMGs [12]]

⁵A sound and complete algorithm corresponds to necessary and sufficient identification assumptions.

In a missing data model represented by a hidden variable DAG $\mathcal{G}(\mathbf{R}, \mathbf{Z}^{(1)}, \mathbf{O}, \mathbf{H})$ and its latent projection ADMG $\mathcal{G}(\mathbf{R}, \mathbf{Z}^{(1)}, \mathbf{O})$, the full law $p(\mathbf{R}, \mathbf{Z}^{(1)}, \mathbf{O})$ is identified if and only if $Z^{(1)} \notin \text{mb}_{\mathcal{G}}(R_Z), \forall Z^{(1)} \in \mathbf{Z}^{(1)}$. Thus, for the full law to be identified, no pair $(Z^{(1)}, R_Z)$ should be connected directly (a.k.a. no *self-censoring*) or through a collider path (a.k.a. no *colluding paths*). Moreover, the identification of the missingness mechanism is given by the odds ratio parameterization, as stated in Theorem 1.

None of the graphs in Fig. 3 have a colluding path. However, as examples, if the path $A_1^{(1)} \rightarrow R_{Y_1} \leftarrow R_{A_1}$ or the path $A_1^{(1)} \leftrightarrow A_2^{(1)} \leftrightarrow R_{A_1}$ were to exist in any of the graphs, then the full law would not be identified.

Settings with entanglements but no missingness interference are amenable to these (existing) identification results, but valid identification requires careful consideration of the dependence engendered by the entanglements, and whether they introduce self-censoring or colluding paths in the corresponding graph.

4.2. Identification With Missingness Interference. In settings with missingness interference, the joint distribution over all counterfactuals, often called the target law, involves counterfactual variables from *multiple worlds*. For example, $Y_1^{(1, r_2=0)}$ and $Y_1^{(1, r_2=1)}$ in the case of Anne in Sec. 3.2. This implies that the full law is not identified without very strong assumptions.⁶ Instead, we will consider identification of *single world* objects. Before giving a rigorous definition, let us motivate the idea with simple examples. In the case of Bob and Anne, the queries that one might be interested in would be, “what would Anne’s BMI be if Bob drops out?”, or “what would their individual or joint outcomes be if they both followed through?”, and not “what would the *joint* outcome of Bob and Anne be, if the other was not in the study?” as the corresponding events happen in different worlds.

For a given network and an assignment of missingness indicators $\mathbf{R} = \mathbf{r}$, a single-world query is defined as a marginal distribution over counterfactual variables $\{\bigcup_{i=1}^n Z_i^{(1, \mathbf{r}_{\text{aff}(i)})} \mid r_i = 1 \text{ if unit } i \text{ is included in this set}\}$. For example, consider a 3-unit network where $\mathbf{r}_{\text{aff}(1)} = (r_2, r_3)$, $\mathbf{r}_{\text{aff}(2)} = (r_1, r_3)$, and $\mathbf{r}_{\text{aff}(3)} = (r_1, r_2)$. Given the choice of $\mathbf{r} = (r_1, r_2, r_3) = (1, 0, 1)$, we consider the query $p(Z_1^{(1, r_2=0, r_3=1)}, Z_3^{(1, r_1=1, r_2=0)})$ to be a valid single-world query, while any marginal distribution that would include $Z_2^{(1, \mathbf{r}_{\text{aff}(2)})}$ would be invalid. This is because the assignment $r_2 = 0$ in counterfactuals $Z_1^{(1, r_2=0, r_3=1)}$ and $Z_3^{(1, r_1=1, r_2=0)}$ is inconsistent with the assignment $r_2 = 1$ in the counterfactual $Z_2^{(1, \mathbf{r}_{\text{aff}(2)})}$. We denote the single world queries in missing data networks by $h(\tilde{\mathbf{Z}}; \mathbf{r})$, where $\tilde{\mathbf{Z}} \equiv \{\tilde{Z}_i : i \in \{1, \dots, n\}\}$ is the set of all counterfactuals.

Our first identification result shows that in the presence of missingness interference, analogues of MCAR and MAR models yield non-parametric identification for any single-world query $h(\tilde{\mathbf{Z}}; \mathbf{r})$.

The first identification result we outline shows that, analogues of MCAR and MAR models when missingness interference is present, yield non-parametric identification for any single-world object $h(\tilde{\mathbf{Z}}; \mathbf{r})$.

⁶See [16, 56] for discussions on how rank preservation assumption relates counterfactuals across multiple worlds.

THEOREM 3. In a missing data ADMG \mathcal{G} with missingness interference, valid single-world objects $h(\tilde{\mathbf{Z}}; \mathbf{r})$ consisting of a set of counterfactuals $\mathbf{Z}' \equiv \bigcup_i \{\mathbf{Z}_i^{(1, \mathbf{r}_{\text{aff}(i)})}\}$, $i \in \{1, \dots, n\}$ are identified when either of these two conditions is satisfied: (1) $\mathbf{R}' \perp\!\!\!\perp \mathbf{O}, \tilde{\mathbf{Z}}$ (MCAR), or (2) $\mathbf{R}' \perp\!\!\!\perp \tilde{\mathbf{Z}} | \mathbf{O}$ (MAR), where \mathbf{R}' refers to the set of all missingness indicators R that index counterfactuals in $h(\tilde{\mathbf{Z}}; \mathbf{r})$. The object $h(\tilde{\mathbf{Z}}; \mathbf{r})$ is a function of $p(\mathbf{Z}', \mathbf{R}, \mathbf{O})$, and the identifying functional is given by:

$$(8) \quad p(\mathbf{Z}', \mathbf{R}, \mathbf{O}) = p(\mathbf{Z}', \mathbf{O}) \times p(\mathbf{R} | \mathbf{O}, \mathbf{Z}') = \frac{p(\mathbf{Z}', \mathbf{R} = \mathbf{r}, \mathbf{O})}{p(\mathbf{R} = \mathbf{r} | \mathbf{O})} \times p(\mathbf{R} | \mathbf{O})$$

where propensity scores are obtained by simple m-separation or (d-separation) rules on ADMG (or DAG) factorization.

Proof: See Appendix for proof. □

In the presence of an MNAR missingness mechanism, single-world queries may not always be identified. Consider instead a special single-world query called the *full-observability* law, $p(\tilde{\mathbf{Z}}^{(\mathbf{r}=1)}, \mathbf{R})$, where $\tilde{\mathbf{Z}}^{(\mathbf{r}=1)} \equiv \bigcup_{i=1}^n Z_i^{(1, \mathbf{r}_{\text{aff}(i)}=1)}$ is the set of all counterfactuals where missingness status is set to full observability (i.e, missingness pattern is $\mathbf{R} = \mathbf{1}$). The full observability law might be thought of as a close analogue to the full law in i.i.d settings. It corresponds to the distribution $p(Y_1^{(1, r_2=1)}, Y_2^{(1, r_1=1)}, R_1, R_2)$ in the example with Anne and Bob. The full observability law is in fact identified under certain assumptions. We first outline our main assumption and a few definitions before presenting our result for identifying the full observability law in MNAR models with entangled missingness.

We first assume the only type of counterfactual that can be a parent of any $R_i \in \mathbf{R}$ has the form $Z_j^{(1, \mathbf{r}_{\text{aff}(j)}=1)}$. We formalize this assumption as follows.

ASSUMPTION 1. In a missing data ADMG \mathcal{G} with missingness interference, $\text{ch}_{\mathcal{G}}(\tilde{\mathbf{Z}}^{(\mathbf{r} \neq \mathbf{1})}) \cap \mathbf{R} = \emptyset$. Here, $\tilde{\mathbf{Z}}^{(\mathbf{r} \neq \mathbf{1})}$ denotes the set of all counterfactuals corresponding to patterns with at least one zero, i.e., $\tilde{\mathbf{Z}}^{(\mathbf{r} \neq \mathbf{1})} = \tilde{\mathbf{Z}} \setminus \tilde{\mathbf{Z}}^{(\mathbf{r}=\mathbf{1})}$. Further, the definition $\text{ch}_{\mathcal{G}}(V)$ applies disjunctively to a set \mathbf{V} , i.e., $\text{ch}_{\mathcal{G}}(\mathbf{V}) = \bigcup_{V \in \mathbf{V}} \text{ch}_{\mathcal{G}}(V)$.

Next, we define the following new graphical structures to be used in our results.

1. *e-self-censoring*⁷ or *affector-censoring*: For a given $R_k \in \mathbf{R}_{\text{aff}(i)}$, there exists an affector-censoring mechanism if $Z_i^{(1, \mathbf{r}_{\text{aff}(i)}=1)} \rightarrow R_k$.
2. *e-colluder*: For a given $R_k \in \mathbf{R}_{\text{aff}(i)}$ and $R_j \notin \mathbf{R}_{\text{aff}(k)}$, there exists an (entangled- or) e-colluder at R_j if $Z_i^{(1, \mathbf{r}_{\text{aff}(i)}=1)} \rightarrow R_j \leftarrow R_k$.
3. *e-colluding path*: For a given $R_k \in \mathbf{R}_{\text{aff}(i)}$, an (entangled- or) e-colluding path is defined to exist between the pair $(Z_i^{(1, \mathbf{r}_{\text{aff}(i)}=1)}, R_k)$ if $Z_i^{(1, \mathbf{r}_{\text{aff}(i)}=1)}$ and R_k are connected through at least one colliding path⁸ that does not go through an observed proxy.

⁷The ‘e’ in these structures stands for ‘entangled’.

⁸A path between vertices V_i and V_j is a colliding path if every vertex V_k on the path is a collider, i.e., bears ones of these forms: $\rightarrow V_k \leftarrow$, $\leftrightarrow V_k \leftrightarrow$, $\rightarrow V_k \leftrightarrow$, or $\leftrightarrow V_k \leftarrow$.




Figure 7: Two examples used for illustrations of Theorem 4. ADMG in (a) does not have an e-colluding path. ADMG in (b) has the e-colluding path $Z_2^{(1,r_1=1)} \rightarrow R_3 \leftarrow R_1$.

THEOREM 4. In a missing data ADMG \mathcal{G} with missingness interference, under Assumption 1, the full-observability law $P(\tilde{\mathbf{Z}}^{(r=1)}, \mathbf{R})$ is identified if and only if there is no e-colluding path for any $R \in \mathbf{R}$ in the ADMG. Further, if \mathcal{G} is a missing data DAG, the full-observability law is identified if and only if there is no e-collider and no e-self-censoring. The identifying functional is given by

$$(9) \quad p(\tilde{\mathbf{Z}}^{(r=1)}, \mathbf{R}) = p(\tilde{\mathbf{Z}}^{(r=1)}) \times \underbrace{p(\mathbf{R} | \tilde{\mathbf{Z}}^{(r=1)})}_{g(p(\mathbf{R}, \mathbf{Z}))} = \frac{p(\tilde{\mathbf{Z}}^{(r=1)}, \mathbf{R} = 1)}{\underbrace{p(\mathbf{R} = 1 | \tilde{\mathbf{Z}}^{(r=1)})}_{g(p(\mathbf{R}, \mathbf{Z}))|_{\mathbf{R}=1}}} \times \underbrace{p(\mathbf{R} | \tilde{\mathbf{Z}}^{(r=1)})}_{g(p(\mathbf{R}, \mathbf{Z}))}.$$

and missingness mechanism $p(\mathbf{R} | \tilde{\mathbf{Z}}^{(r=1)})$ is identified using the OR parameterization given below:

$$p(\mathbf{R} | \tilde{\mathbf{Z}}^{(r=1)}) = \frac{1}{\sigma} \times \prod_{k=1}^K p(R_k | \mathbf{R}_{-k} = 1, \tilde{\mathbf{Z}}^{(r=1)}) \times \prod_{k=2}^K \text{OR}(R_k, R_{\prec k} | R_{\succ k} = 1, \tilde{\mathbf{Z}}^{(r=1)})$$

where notation and OR is consistent with Sec. 4.1.

Proof: The proof extends the proof for sound and complete identification of the missing data full law in ADMGs in the absence of colluding paths in [12]. Proof is in the Appendix. \square

We illustrate Theorem 4 using two different examples, one where the full observability law is identified, and one where the graphical criterion is violated, and hence the law is not identified.

In the ADMG shown in Fig. 7(a), there are no e-colluding paths. Thus, by Theorem 4, the full observability law $P(Z_1^{(1,r_2=2)}, Z_2^{(1,r_1=1)}, Z_3^{(1)}, R_1, R_2, R_3)$ is identified. We show how in the following paragraph.

The propensity score for R_3 is identified from observed data in a straightforward manner: $p(R_3 | \text{pa}(R_3)) = p(R_3 | R_2)$. But this is not true for propensity scores of R_2, R_1 . For R_1 , the propensity score is $p(R_1 | Z_3^{(1)})$. We see that $R_1 \perp\!\!\!\perp R_3 | Z_3^{(1)}, R_2$, but that the independence would not hold without R_2 . That is, we can identify $p(R_1 | Z_3^{(1)}, R_2, R_3 = 1)$. For R_2 , we have the propensity score $p(R_2 | Z_3^{(1)}, R_1)$. Here, we cannot insert R_3 behind the conditioning bar since the conditional independence does not hold. But, we could fix $R_3 = 1$, an operation that removes the dependence of R_3 on R_2 and yields a world where $R_2 \perp\!\!\!\perp R_3 | Z_3^{(1)}, R_1$

(see Appendix for more details). But we can use the odds ratio parameterization to yield identification of the joint score $p(R_1, R_2|Z_3^{(1)})$.

$$p(R_1, R_2|Z_3^{(1)}) = \frac{1}{Z} \times p(R_2|R_1 = 1, Z_3^{(1)})p(R_1|R_2 = 1, Z_3^{(1)}) \times \text{OR}(R_2, R_1|Z_3^{(1)})$$

Here, $p(R_2|R_1 = 1, Z_3^{(1)})$ and $p(R_1|R_2 = 1, Z_3^{(1)})$ can be identified as described in the previous paragraph. And, we can expand the odds ratio as follows:

$$\begin{aligned} \text{OR}(R_2, R_1|Z_3^{(1)}) &= \frac{p(R_1 = r_1|R_2 = r_2, Z_3^{(1)})}{p(R_1 = 1|R_2 = r_2, Z_3^{(1)})} \times \frac{p(R_1 = 1|R_2 = 1, Z_3^{(1)})}{p(R_1 = r_1|R_2 = 1, Z_3^{(1)})} \\ &= \frac{p(R_1 = r_1|R_2 = r_2, Z_3^{(1)}, R_3 = 1)}{p(R_1 = 1|R_2 = r_2, Z_3^{(1)}, R_3 = 1)} \times \frac{p(R_1 = 1|R_2 = 1, Z_3^{(1)}, R_3 = 1)}{p(R_1 = r_1|R_2 = 1, Z_3^{(1)}, R_3 = 1)} \end{aligned}$$

which is also identified. The normalizing term Z is also identified since the conditional pieces and the OR piece are all individually identified. This means that $p(R_1, R_2|Z_3^{(1)})$ is identified. By Eqn. 9, the full observability law is identified as all the pieces are individually a function of the observed data.

The full observability law is however not identified in Fig. 7(b) because of the colluding path $Z_2^{(1, r_1=1)} \rightarrow R_3 \leftarrow R_1$. Intuitively, the missingness mechanism $p(\mathbf{R} | \text{pa}(\mathbf{R}))$ should be identified for all levels of \mathbf{R} . Writing the missingness mechanism as $p(R_1) \times p(R_2) \times p(R_3 | R_1, Z_2^{(1, r_1=1)})$, we realize that we cannot identify the last term unless we are able to set *both* R_1 and R_2 to the value 1, and by consistency, replace the counterfactual $Z_2^{(1, r_1=1)}$ by observed proxy Z_2 . Doing so in the joint would mean that we cannot identify this quantity for levels of \mathbf{R} that set R_1 or R_2 to 0. The rigorous argument for non-identification involves providing counterexamples where two missing data ADMGs disagree on the full observability law but agree on the observed data law. Such an argument can be achieved by, for instance, counting the parameters required to characterize the full observability law against the observed law in a binary model, and showing that the observed law has fewer parameters and hence it is not possible to uniquely map back to the full observability law. A detailed account of parameter counting in examples with e-colluding paths has been deferred to the Appendix, under the completeness section of the proof for Theorem 4 to follow the procedure for parameter counting, knowledge of the Möbius parameterization of the nested Markov model for binary variable ADMGs is required, which we also discuss in the Appendix.

5. Implications of the Missingness Interference Model for I.I.D. Settings. In what we have presented so far, graphical models of entangled missingness are viewed as representing a block of units subject to dependence or interference. The same graphical structures, however, can also be viewed as representing multivariate observations for a single unit in i.i.d. settings, e.g. settings where a single unit has multivariate treatment and/or outcome measurements.

As we have shown, target law dependence and missingness process dependence can be handled by existing i.i.d. missing data methods. However, to the best of our knowledge, the i.i.d. analogue of *missingness interference* entanglement has never been discussed in the prior literature on missing data in i.i.d. settings. In fact, the absence of missingness interference is typically required by an edge restriction ubiquitous in classical graphical approaches to missing data – namely that no missingness indicator is a cause/parent of any counterfactual or fully observed variable (restriction (ii) in Sec. 2.3). Relaxing this assumption implies that




Figure 8: Graphs used to illustrate the role of missingness interference in i.i.d. settings.

changes in a missingness indicator may cause changes in the underlying full data law. Our proposed framework on missingness interference entanglement can be readily adapted to relax the imposed edge restriction in i.i.d. settings.

As an example, assume L_1, L_2 are measurements during first and second followup visits, respectively (i in L_i is indexing variables of a single patient). In some situations, whether or not a patient shows up to their first followup can affect their future health; perhaps through mental self-empowerment or words of self-encouragement on treatment adherence. This can be captured via counterfactuals of the form $L_2^{(r_2=1, r_1)}$, or simply $L_2^{(1, r_1)}$, that reads as the measurement obtained at the second visit had the patient showed up to their first visit ($r_1 = 1$), or had they not ($r_1 = 0$). This example is analogous to assuming L_1 is the measurement of one patient and L_2 is the measurement of another patient on their first visit (i in L_i is indexing units for the same variable). Assuming an asymmetric relationship between the two patients, namely that only the measurement of the second patient depends on whether the first patient shows up to the visit or not and not vice versa, we can represent the relations using the ADMG shown in Fig. 8(a). This ADMG corresponds to the full data law with a single counterfactual version for L_1 (indexed by $R_1 = 1$) and two counterfactual versions for L_2 (indexed by $R_2 = 1, R_1 = 0$ and $R_2 = 1, R_1 = 1$).

Similar to missingness interference in non-i.i.d. settings, the proxy variable L_2 can be defined as a deterministic function of $R_1, R_2, L_2^{(1, r_1=0)}$, and $L_2^{(1, r_1=1)}$ as follows:

$$L_2 \leftarrow \begin{cases} (1 - r_1) L_2^{(1, r_1=0)} + r_1 L_2^{(1, r_1=1)} & \text{if } r_2 = 1, \\ ? & \text{if } r_2 = 0. \end{cases}$$

We can further introduce an intermediate step by first defining a single indexed counterfactual for L_2 , i.e., $L_2^{(r_2=1)}$ or $L_2^{(1)}$ for short, as a deterministic function of $R_1, L_2^{(1, r_1=1)}$, $L_2^{(1, r_1=0)}$, and then defining the proxy L_2 as a deterministic function of $R_2, L_2^{(1)}$ only:

$$L_2^{(1)} = (1 - R_1) \times L_2^{(1, r_1=0)} + R_1 \times L_2^{(1, r_1=1)},$$

$$L_2 \leftarrow \begin{cases} L_2^{(1)} & \text{if } r_2 = 1, \\ ? & \text{if } r_2 = 0. \end{cases}$$

The ADMG in Fig. 8(b) corresponds to the full data law with three counterfactual versions for L_2 . The gray edges illustrate the consistency relations outlined above.

The distinction of Fig. 8(a, b) from the classical i.i.d. missing data DAGs is the presence of multi-indexed counterfactuals, which indicates that, for instance in a given temporal ordering on variables, a variable later in the order may depend on whether an earlier variable is measured or not, regardless of whether sample realizations are dependent or not. Projecting out the multi-indexed counterfactuals from Fig. 8(b) yields the DAG in Fig. 8(c) where R_1 is pointing to $L_2^{(1)}$, a violation of edge restrictions in classical missing data graphs.




Figure 9: The model used to generate synthetic data for our experiments.

In the above i.i.d. example, the target distribution $p(L_1^{(r_1=1)}, L_2^{(r_2=1, r_1=0)})$ cannot be identified as a function of observed data law due to the inconsistency in assignments of R_1 , without imposing strong cross-world assumptions⁹. For instance, we could assume that the effect of $L_1^{(r_1=1)}$ on $L_2^{(r_2=1)}$ is fully mediated by a variable with no missingness, denoted by X . Then $p(L_1^{(r_1=1)}, L_2^{(r_1=0, r_2=1)})$ can be identified as follows:

$$\begin{aligned} p(L_1^{(r_1=1)}, L_2^{(r_2=1, r_1=0)}) &= \sum_x p(L_1^{(r_1=1)} | x) \times p(L_2^{(r_2=1, r_1=0)} | x, L_1^{(r_1=1)} = l_1^{(r_1=1)}) \times p(x) \\ &= \sum_x p(L_1 | x, R_1 = 1) \times p(L_2 | x, R_1 = 0, R_2 = 1) \times p(x). \end{aligned}$$

In order to avoid such complications, edges from \mathbf{R} to counterfactuals were excluded in prior literature on graphical models of missing data. Under this edge restriction, we are assuming $L_j^{(r_k, r_j=1)} = L_j^{(r_j=1)}, \forall R_j, R_k \in R$. As shown above, this assumption can be relaxed by leveraging our developed framework on missingness interference entanglement.

6. Simulations. We generate synthetic data and attempt to recover the ground truth in entangled missingness settings where the targets we are interested in are identified. In particular, we reserve our interest to the cases where missingness interference (\mathcal{E}_3) is present, as results for settings where missingness interference is absent have been discussed in literature [16].

Data was simulated based on the graph in Fig. 9. There are three units in the network, and units 1 and 2 are neighbors. When the graph consists only the solid edges, we have an MCAR scenario. Inclusion of the dashed blue edges ($C_1 \rightarrow R_1, C_1 \rightarrow R_2, C_2 \rightarrow R_1, C_2 \rightarrow R_2, C_3 \rightarrow R_3$), yields a MAR scenario. Finally, inclusion of all the edges (adding $Z_3^{(1)} \rightarrow R_2, Z_3^{(1)} \rightarrow R_1$), results in an MNAR scenario.

Parameters associated with the generative model are as follows: $p(C_1), p(C_2), p(C_3)$, all chosen to be univariate normal distributions, counterfactuals $Z_1^{(1, r_2=0)}, Z_1^{(1, r_2=1)}, Z_3^{(1, r_1=0)}, Z_3^{(1, r_1=1)}$ and $Z_3^{(1)}$ are drawn from a multivariate normal distribution. The means of the multivariate normal are given by linear functions of the corresponding parents C in the graph, and covariance matrix has no zero entries; this ensures that all the counterfactuals are associated with each other. Missingness indicators R_i are generated according to the scenario, using a binomial distribution where the binomial probability is chosen completely at random (MCAR) or is a logistic sigmoid of a linear function of the parents of R_i (MAR, MNAR).

⁹Discussions on cross-world assumptions can be found in mediation literature. For example, see [85]

Observed proxies Z_i are a deterministic function of the counterfactuals and missingness indicators, as given by the consistency assumption.

We are interested in identifying the following functionals: $\mathbb{E}[Z_1^{(1,r_2=0)}]$, $\mathbb{E}[Z_1^{(1,r_2=1)}]$, $\mathbb{E}[Z_2^{(1,r_1=0)}]$, $\mathbb{E}[Z_2^{(1,r_1=1)}]$, under MCAR, MAR and MNAR conditions. In the given graph, since there are no e-colluders, we should be able to identify the full observability law (or its marginals) even in the MNAR case. We use an estimating equation of the form $\mathbb{E}\left[\frac{\mathbb{I}(\mathbf{R}=\mathbf{1})}{\pi(\mathbf{C},\tilde{\mathbf{Z}})}h(\tilde{\mathbf{Z}})\right]$ where $\pi(\mathbf{C},\tilde{\mathbf{Z}})$ is the inverse weight function composed of propensity scores and $h(\cdot)$ is the functional we are interested in. The estimator uses regression models for the weights, and expectations are done empirically.

We would like to compare the adjusted estimates we obtain with those obtained by an investigator who does not use entangled missingness ideas to really show why such a framework is necessary. First, when such an investigator, who is unaware of the various interactions between the units, notices missing data in the dataset, they are likely to adopt a naive, missing data approach to estimate counterfactuals, probably assuming a MAR mechanism on i.i.d data. This means that they are restricted to estimating $Z_1^{(1)}$ and $Z_2^{(1)}$, which do not correspond reliably to the underlying truth as units 1 and 2 have *two* counterfactuals each. In order to replicate such an analysis, we estimate $Z_1^{(1)}$ and $Z_2^{(1)}$ using a MAR i.i.d assumption. Many estimators exist for such a scenario, including g-formula and inverse probability weight (IPW)-based estimators [86]. In this experiment, we use the doubly-robust *augmented IPW* estimator given by $\frac{1}{N} \sum_{n=1}^N \left[\frac{R_{i,n} Z_{i,n}}{p(R_{i,n}|C_{i,n})} - \frac{\{R_{i,n} - p(R_{i,n}|C_{i,n})\}}{p(R_{i,n}|C_{i,n})} \mathbb{E}[Z_i|C_{i,n}] \right]$, where i represents the unit and n is the sample. See [86] for more on these estimators.

A total of $50k$ samples were generated from the network, and estimation was done over 50 bootstrap trials. The resulting estimates are shown in Table 1. The first four rows show our estimates for counterfactuals when properly adjusted for, followed by the last two rows that show the estimates when we completely ignore dependence between units, but account for missingness using MAR.

To show how our adjusted estimates perform, we compare the bias of an estimate against the ground truth (which we have from model parameters), in two different scenarios - when we adjust for the network structure and the entanglement appropriately, and one where we do not adjust, for MAR and MNAR.¹⁰ The latter pertains to simply taking the conditional means for different missingness patterns. Estimates based on complete rows without adjustment would be biased, as shown in the Fig. 10. Estimates obtained after adjustment are labeled with an asterisk (*) in the x-axis.

As we have pointed out here, there are any number of ways in which an investigator's approach might be biased, if they are not careful about considering *both* dependence and missingness *and* the precise way in which they interact and we have simply illustrated one way in which it can happen.

¹⁰In MCAR, the bias would be the same for the two different approaches since we can treat the complete rows in the dataset as an unbiased, representative dataset.

TABLE 1

Estimates of counterfactuals over 50 bootstrap trials along with quantiles $q_{0.05}$ and $q_{0.95}$. When the investigator is unaware of entanglements, they are restricted to assuming an i.i.d missing data scenario and likely estimate $Z_1^{(1)}$ and $Z_2^{(1)}$ using a MAR assumption, as recorded in the last two rows, done here using an AIPW estimator. These values do not correspond in any reliable way to the four underlying counterfactuals which were estimated with the knowledge of entanglements and the corresponding adjustment required, and recorded in the first four rows.

Estimate for	Mean	$q_{0.05}$	$q_{0.95}$
$Z_1^{(1,r_2=1)}$	-2.942	-3.009	-2.889
$Z_1^{(1,r_2=0)}$	3.138	3.108	3.164
$Z_2^{(1,r_1=1)}$	10.816	10.773	10.856
$Z_2^{(1,r_1=0)}$	-4.333	-4.351	-4.313
$Z_1^{(1)}$	1.497	1.456	1.536
$Z_2^{(1)}$	2.252	2.138	2.373




Figure 10: Bias recorded in bootstrapped estimates of targets (shown in x-axis), MAR case (left) and MNAR (right). We compare our adjusted IPW estimates (denoted by an asterisk (*) on the x-axis) to the unadjusted estimate, which is obtained by ignoring the network missingness structure underpinning the data. Error bars represent quantiles $q_{0.05}$ and $q_{0.95}$ across 50 bootstrap samples.

7. Discussion. In this paper, we developed a graphical framework to model entanglements due to the simultaneous presence of data dependence and missing data. In non-i.i.d settings with partial interference, an entanglement could arise due to dependence exhibited among target law variables of different units (*target law dependence*), dependence of missingness indicators of one unit on target law variables of other units (*missingness process dependence*), or dependence of observed factials of one unit on missingness indicators of other units (*missingness interference*). We have characterized these three entanglement scenarios mathematically and provided *sound* and *complete* full law identification theory under each scenario. We note that several open questions remain in identification of entangled missingness models that we have not addressed in this work. For instance, while we have focused on the full observability law, there are a larger number of models where specific targets or parameters such as the causal effect might still be identified when the full (observability) law is not. Even in i.i.d missing data models, there exist identified target laws where the corresponding full law is not identified; in fact, a complete characterization of target law identification remains an open problem. Finally, we have demonstrated that our model for missingness interference extends i.i.d. graphical missing data methods to allow for missingness indicators that may be causes of counterfactual or fully observed random variables, a phenomenon that occurs frequently in real-world settings (especially with longitudinal data) but had not previously been considered in the missing data literature.

APPENDIX A: THE NESTED MARKOV FACTORIZATION

The nested Markov factorization of $p(\mathbf{V})$ relative to an ADMG $\mathcal{G}(\mathbf{V})$ is defined with the use of conditional distributions known as kernels and their associated *conditional* ADMGs (called *CADMGs*) that are derived from $p(\mathbf{V})$ and $G(\mathbf{V})$ respectively, via repeated applications of the *fixing* operator [62]. A CADMG $\mathcal{G}(\mathbf{V}, \mathbf{W})$ is an ADMG whose nodes can be partitioned into random variables \mathbf{V} and fixed variables \mathbf{W} ; with the restriction that only outgoing edges may be adjacent to variables in \mathbf{W} . The nested Markov factorization may also more generally be defined for a distribution $p(\mathbf{V}|\mathbf{W})$ with respect to a CADMG $\mathcal{G}(\mathbf{V}, \mathbf{W})$, and links *kernels*, mappings derived from $p(\mathbf{V}|\mathbf{W})$ and CADMGs derived from $\mathcal{G}(\mathbf{V}, \mathbf{W})$ via a *fixing* operation.

Kernel: A kernel $q_{\mathbf{V}}(\mathbf{V}|\mathbf{W})$ is a mapping from values in \mathbf{W} to normalized densities over \mathbf{V} [87]. A conditional distribution is a familiar example of a kernel, in that $\sum_{\mathbf{v} \in \mathbf{V}} q_{\mathbf{V}}(\mathbf{v}|\mathbf{w}) = 1$. Conditioning and marginalization are defined in kernels in the usual way: For $\mathbf{A} \subseteq \mathbf{V}$, $q_{\mathbf{V}}(\mathbf{A}|\mathbf{W}) \equiv \sum_{\mathbf{V} \setminus \mathbf{A}} q_{\mathbf{V}}(\mathbf{V}|\mathbf{W})$ and $q_{\mathbf{V}}(\mathbf{V} \setminus \mathbf{A}|\mathbf{A} \cup \mathbf{W}) \equiv \frac{q_{\mathbf{V}}(\mathbf{V}|\mathbf{W})}{q_{\mathbf{V}}(\mathbf{A}|\mathbf{W})}$.

Fixability and the fixing operator: A variable $V \in \mathbf{V}$ in a CADMG \mathcal{G} is fixable if $\text{de}_{\mathcal{G}}(V) \cap \text{des}_{\mathcal{G}}(V) = \emptyset$. In other words, V is fixable if paths $V \leftrightarrow \dots \leftrightarrow B$ and $V \rightarrow \dots \rightarrow B$ do not both exist in \mathcal{G} for any $B \in \mathbf{V} \setminus \{V\}$.

We define a fixing operator $\phi_V(\mathcal{G})$ for graphs, and a fixing operator $\phi_V(q; \mathcal{G})$ for kernels. Given a CADMG $\mathcal{G}(\mathbf{V}, \mathbf{W})$, with a fixable $V \in \mathbf{V}$, $\phi_V(\mathcal{G}(\mathbf{V}, \mathbf{W}))$ yields a new CADMG $\mathcal{G}(\mathbf{V} \setminus \{V\}, \mathbf{W} \cup \{V\})$ obtained from $\mathcal{G}(\mathbf{V}, \mathbf{W})$ by moving V from \mathbf{V} to \mathbf{W} , and removing all edges with arrowheads into V . Given a kernel $q_{\mathbf{V}}(\mathbf{V}|\mathbf{W})$, and a CADMG $\mathcal{G}(\mathbf{V}, \mathbf{W})$, the operator $\phi_V(q_{\mathbf{V}}(\mathbf{V}|\mathbf{W}), \mathcal{G}(\mathbf{V}, \mathbf{W}))$ yields a new kernel:

$$q_{\mathbf{V} \setminus \{V\}}(\mathbf{V} \setminus \{V\}|\mathbf{W} \cup \{V\}) \equiv \frac{q_{\mathbf{V}}(\mathbf{V}|\mathbf{W})}{q_{\mathbf{V}}(V|\text{mb}_{\mathcal{G}}(V))}$$

Fixing sequences: A sequence $\langle V_1, \dots, V_k \rangle$ is said to be *valid* in $\mathcal{G}(\mathbf{V}, \mathbf{W})$ if V_1 fixable in $\mathcal{G}(\mathbf{V}, \mathbf{W})$, V_2 is fixable in $\phi_{V_1}(\mathcal{G}(\mathbf{V}, \mathbf{W}))$, and so on. If any two sequences σ_1, σ_2 for the same set $\mathbf{S} \subseteq \mathbf{V}$ are fixable in \mathcal{G} , they lead to the same CADMG. The graph fixing operator can be extended to a set \mathbf{S} : $\phi_{\mathbf{S}}(\mathcal{G})$. This operator is defined as applying the vertex fixing operation in any valid sequence σ for set \mathbf{S} .

Given a sequence $\sigma_{\mathbf{S}}$, define $\eta(\sigma_{\mathbf{S}})$ to be the first element in $\sigma_{\mathbf{S}}$, and $\tau(\sigma_{\mathbf{S}})$ to be the subsequence of $\sigma_{\mathbf{S}}$ containing all elements but the first. Given a sequence $\sigma_{\mathbf{S}}$ on elements in \mathbf{S} valid in $\mathcal{G}(\mathbf{V}, \mathbf{W})$, the kernel fixing operator $\phi_{\sigma_{\mathbf{S}}}(q_{\mathbf{V}}(\mathbf{V}|\mathbf{W}), \mathcal{G}(\mathbf{V}, \mathbf{W}))$ is defined to be equal to $q_{\mathbf{V}}(\mathbf{V}|\mathbf{W})$ if $\sigma_{\mathbf{S}}$ is the empty sequence, and $\phi_{\tau(\sigma_{\mathbf{S}})}(\phi_{\eta(\sigma_{\mathbf{S}})}(q_{\mathbf{V}}(\mathbf{V}|\mathbf{W}); \mathcal{G}(\mathbf{V}, \mathbf{W})), \phi_{\eta(\sigma_{\mathbf{S}})}(\mathcal{G}(\mathbf{V}, \mathbf{W})))$ otherwise.

Reachability: Given a CADMG $\mathcal{G}(\mathbf{V}, \mathbf{W})$, a set $\mathbf{R} \subseteq \mathbf{V}$ is called *reachable* if there exists a sequence for $\mathbf{V} \setminus \mathbf{R}$ valid in $\mathcal{G}(\mathbf{V}, \mathbf{W})$. In other words, if \mathbf{S} is fixable in \mathcal{G} , $\mathbf{V} \setminus \mathbf{S}$ is reachable.

Intrinsic sets: A set \mathbf{R} reachable in $\mathcal{G}(\mathbf{V}, \mathbf{W})$ is *intrinsic* in $\mathcal{G}(\mathbf{V}, \mathbf{W})$ if $\phi_{\mathbf{V} \setminus \mathbf{R}}(\mathcal{G})$ contains a single district, \mathbf{R} itself. The set of intrinsic sets in a CADMG \mathcal{G} is denoted by $\mathcal{I}(\mathcal{G})$.

Nested Markov factorization: A distribution $p(\mathbf{V}|\mathbf{W})$ is said to obey the *nested Markov factorization* with respect to the CADMG $\mathcal{G}(\mathbf{V}, \mathbf{W})$ if there exists a set of kernels of the form

$\{q_{\mathbf{S}}(\mathbf{S} | \text{pa}_{\mathcal{G}}(\mathbf{S})) : \mathbf{S} \in \mathcal{I}(\mathcal{G})\}$ such that for every valid sequence $\sigma_{\mathbf{R}}$ for a reachable set \mathbf{R} in \mathcal{G} , we have:

$$\phi_{\sigma_{\mathbf{R}}}(p(\mathbf{V}|\mathbf{W}); \mathcal{G}(\mathbf{V}, \mathbf{W})) = \prod_{\mathbf{D} \in \mathcal{D}(\phi_{\mathbf{R}}(\mathcal{G}(\mathbf{V}, \mathbf{W})))} q_{\mathbf{D}}(\mathbf{D} | \text{pa}_{\mathcal{G}}^{\mathbf{S}}(\mathbf{D}))$$

If a distribution obeys this factorization, then for any reachable \mathbf{R} , any two valid sequences on \mathbf{R} applied to $p(\mathbf{V}|\mathbf{W})$ yield the same kernel $q_{\mathbf{R}}(\mathbf{R} | \mathbf{V} \setminus \mathbf{R})$. Hence, kernel fixing may be defined on sets, just as graph fixing. In this case, for every $\mathbf{D} \in \mathcal{I}(\mathcal{G})$, $q_{\mathbf{D}}(\mathbf{D} | \text{pa}_{\mathcal{G}}^{\mathbf{S}}(\mathbf{D})) \equiv \phi_{\mathbf{V} \setminus \mathbf{D}}(p(\mathbf{V}|\mathbf{W}); \mathcal{G}(\mathbf{V}, \mathbf{W}))$.

The *district factorization* or *Tian factorization* of $p(\mathbf{V}|\mathbf{W})$ results from the nested factorization:

$$\begin{aligned} p(\mathbf{V}|\mathbf{W}) &= \prod_{\mathbf{D} \in \mathcal{D}(\mathcal{G}(\mathbf{V}, \mathbf{W}))} q_{\mathbf{D}}(\mathbf{D} | \text{pa}_{\mathcal{G}}^{\mathbf{S}}(\mathbf{D})) \\ &= \prod_{\mathbf{D} \in \mathcal{D}(\mathcal{G}(\mathbf{V}, \mathbf{W}))} \left(\prod_{D \in \mathbf{D}} p(D | \text{pre}_{\prec}(D)) \right), \end{aligned}$$

where $\text{pre}_{\prec}(D)$ is the set of predecessors of D according to a topological total ordering \prec . Each factor $\prod_{D \in \mathbf{D}} p(D | \text{pre}_{\prec}(D))$ is only a function of $\mathbf{D} \cup \text{pa}_{\mathcal{G}}(\mathbf{D})$ under the nested factorization.

An important result in [62] states that if $p(\mathbf{V} \cup \mathbf{H}|\mathbf{W})$ obeys the factorization for a CDAG $\mathcal{G}(\mathbf{V} \cup \mathbf{H}, \mathbf{W})$, then $p(\mathbf{V}|\mathbf{W})$ obeys the nested factorization for the latent projection CADMG $\mathcal{G}(\mathbf{V}, \mathbf{W})$.

Identification: Not every interventional distribution $p(\mathbf{Y}(\mathbf{a}))$ is identified in a hidden variable causal model. However, *every* $p(\mathbf{Y}(\mathbf{a})|\mathbf{W})$ identified from $p(\mathbf{V}|\mathbf{W})$ can be expressed as a modified nested factorization as follows:

$$\begin{aligned} p(\mathbf{Y}(\mathbf{a})|\mathbf{W}) &= \sum_{\mathbf{Y}^* \setminus \mathbf{Y}} \prod_{\mathbf{D} \in \mathcal{D}(\mathcal{G}_{\mathbf{Y}^*})} p(\mathbf{D} | \text{do}(\text{pa}_{\mathcal{G}}^{\mathbf{S}}(\mathbf{D})))|_{\mathbf{A}=\mathbf{a}} \\ &= \sum_{\mathbf{Y}^* \setminus \mathbf{Y}} \prod_{\mathbf{D} \in \mathcal{D}(\mathcal{G}_{\mathbf{Y}^*})} \phi_{\mathbf{V} \setminus \mathbf{D}}(p(\mathbf{V}|\mathbf{W}); \mathcal{G}(\mathbf{V}, \mathbf{W}))|_{\mathbf{A}=\mathbf{a}} \end{aligned}$$

where $\mathbf{Y}^* \equiv \text{an}_{\mathcal{G}(\mathbf{V}(\mathbf{a}), \mathbf{W})}(\mathbf{Y}) \setminus \mathbf{a}$. That is, $p(\mathbf{Y}(\mathbf{a})|\mathbf{W})$ is only identified if it can be expressed as a factorization, where every piece corresponds to a kernel associated with a set intrinsic in $\mathcal{G}(\mathbf{V}, \mathbf{W})$. Moreover, no piece in this factorization contains elements of \mathbf{A} as random variables.

Binary Parameterization of Nested Markov Models: From the nested factorization, intrinsic sets given their parents form the atomic units of the nested Markov model. Using this observation, a smooth parameterization of discrete nested Markov models was provided by [88]. We provide a brief description of how to derive the Möbius parameters of a binary nested Markov model.

For each district $\mathbf{D} \in \mathcal{D}(\mathcal{G})$, consider all possible subsets $\mathbf{S} \subseteq \mathbf{S}$. If \mathbf{S} is intrinsic in $\phi_{\mathbf{V} \setminus \mathbf{S}}(\mathcal{G})$, define the head \mathbf{H} of the intrinsic set to be all vertices in \mathbf{S} that are childless in $\phi_{\mathbf{V} \setminus \mathbf{S}}(\mathcal{G})$ and the tail \mathbf{T} to be all parents of the head in the CADMG $\phi_{\mathbf{V} \setminus \mathbf{S}}(\mathcal{G})$, excluding the head itself. Formally, $\mathbf{H} \equiv \{V \in \mathbf{S} | \text{ch}_{\phi_{\mathbf{V} \setminus \mathbf{S}}(\mathcal{G})}(V) = \emptyset\}$ and $\mathbf{T} \equiv \text{pa}_{\phi_{\mathbf{V} \setminus \mathbf{S}}(\mathcal{G})}(\mathbf{H}) \setminus \mathbf{H}$. The corresponding

set of Möbius parameters for this intrinsic head and tail pair parameterizes the kernel $q_{\mathbf{S}}(\mathbf{H} = \mathbf{0} | \mathbf{T})$; i.e., the kernel where all variables outside the intrinsic set \mathbf{S} are fixed, and all elements of the head are set to zero given the tail. Note that these parameters are, in general, variationally dependent (in contrast to variationally independent in the case of an ordinary DAG model) as the heads and tails in these parameter sets may overlap. The joint density for any query $p(\mathbf{V} = \mathbf{v})$ can be obtained through the Möbius inversion formula; see [87, 88] for details. We will denote $q_{\mathbf{S}}(\mathbf{H} = \mathbf{0} | \mathbf{T})$ simply as $q(\mathbf{H} = \mathbf{0} | \mathbf{T})$ as it is generally clear what variables are still random in the kernel corresponding to a given intrinsic set.

Binary Parameterization of Missing Data Models: We use the parameterization described earlier to count the number of parameters required to parameterize the full observability law of a missing data ADMG and its corresponding observed law. We then use this to reason that if the number of parameters in the full observability law exceeds those in the observed law, it is impossible to establish a map from the observed law to the full law. This in turn implies that such a full observability law is not identified. In the full observability, the deterministic factors, i.e., proxies given parents can be ignored as the probability of those events is always 1. However, while counting the observed law, we are careful to treat counterfactuals as unobserved and obtain the corresponding ADMG. The Möbius parameters are then derived in a similar manner as before, but with additional constraint that if Z_i appears in the head of a parameter, and missingness indicators R_i or $\mathbf{R}_{\text{aff}(i)}$ appear in the tail, then the kernel must be restricted to cases where $R_i = 1$ and $\mathbf{R}_{\text{aff}(i)} = \mathbf{1}$. This is because, (1) when $R_i = 0$, the probability of the head taking any value aside from those where $Z_i = ?$ is deterministically 0, and (2) cases where $\mathbf{R}_{\text{aff}(i)}$ are set to values different from 1 are irrelevant to the identification of the full observability law. The consideration that $R_i = 1$ always holds, but that $\mathbf{R}_{\text{aff}(i)} = \mathbf{1}$ is only for identifying the full observability law. For other cases, considerations vary.

APPENDIX B: PROOFS

We restate the theorems and outline their proofs here.

Theorem 3 *In a missing data ADMG \mathcal{G} with missingness interference, valid single-world objects $h(\tilde{\mathbf{Z}}; \mathbf{r})$ consisting of a set of counterfactuals $\mathbf{Z}' \equiv \bigcup_i \{\mathbf{Z}_i^{(1, \mathbf{r}_{\text{aff}(i)})}\}$, $i \in \{1, \dots, n\}$ are identified when either of these two conditions is satisfied: (1) $\mathbf{R}' \perp\!\!\!\perp \mathbf{O}, \tilde{\mathbf{Z}}$ (MCAR), or (2) $\mathbf{R}' \perp\!\!\!\perp \tilde{\mathbf{Z}} | \mathbf{O}$ (MAR), where \mathbf{R}' refers to the set of all missingness indicators R that index counterfactuals in $h(\tilde{\mathbf{Z}}; \mathbf{r})$. The object $h(\tilde{\mathbf{Z}}; \mathbf{r})$ is a function of $p(\mathbf{Z}', \mathbf{R}, \mathbf{O})$, and the identifying functional is given by:*

$$(10) \quad p(\mathbf{Z}', \mathbf{R}, \mathbf{O}) = p(\mathbf{Z}', \mathbf{O}) \times p(\mathbf{R} | \mathbf{O}, \mathbf{Z}') = \frac{p(\mathbf{Z}', \mathbf{R} = \mathbf{r}, \mathbf{O})}{p(\mathbf{R} = \mathbf{r} | \mathbf{O})} \times p(\mathbf{R} | \mathbf{O})$$

where propensity scores are obtained by simple m -separation or (d -separation) rules on ADMG (or DAG) factorization.

Proof: \mathbf{Z}' is the set of counterfactuals in $h(\tilde{\mathbf{Z}}; \mathbf{r})$ such that they construct a valid single-world counterfactual in the world $\mathbf{R} = \mathbf{r}$. Further, let the set of all missingness indicators that index \mathbf{Z}' be \mathbf{R}' , and that $\mathbf{R}' = \mathbf{r}'$ when $\mathbf{R} = \mathbf{r}$. We are interested in identifying a distribution $P(\mathbf{Z}')$ or a function $h(\tilde{\mathbf{Z}}; \mathbf{r})$ thereof. The crucial factor here is that no variable, whose missingness indicator is 0, is present in $h(\cdot)$.

In case (1), we can write $P(\mathbf{Z}') = P(\mathbf{Z}' | \mathbf{R}' = \mathbf{r}')$ since $\mathbf{R}' \perp\!\!\!\perp \tilde{\mathbf{Z}}$ and $\mathbf{Z}' \subset \tilde{\mathbf{Z}}$. And by consistency, we can replace all counterfactuals by their corresponding proxies and the object is identified.

In case (2), we can write $P(\mathbf{Z}') = P(\mathbf{Z}' | \mathbf{O}) \times P(\mathbf{O}) = P(\mathbf{Z}' | \mathbf{O}, \mathbf{R}' = \mathbf{r}') \times P(\mathbf{O})$ since $\mathbf{R}' \perp\!\!\!\perp \tilde{\mathbf{Z}} | \mathbf{O}$ and $\mathbf{Z}' \subset \tilde{\mathbf{Z}}$. And by consistency, we can identify the distribution as we can replace all counterfactuals by their corresponding proxies. \square

Theorem 4 *In a missing data ADMG \mathcal{G} with missingness interference, under Assumption 1, the full-observability law $P(\tilde{\mathbf{Z}}^{(\mathbf{r}=1)}, \mathbf{R})$ is identified if and only if there is no e-colluding path for any $R \in \mathbf{R}$ in the ADMG. Further, if \mathcal{G} is a missing data DAG, the full-observability law is identified if and only if there is no e-collider and no e-self-censoring. The identifying functional is given by*

$$p(\tilde{\mathbf{Z}}^{(\mathbf{r}=1)}, \mathbf{R}) = p(\tilde{\mathbf{Z}}^{(\mathbf{r}=1)}) \times \underbrace{p(\mathbf{R} | \tilde{\mathbf{Z}}^{(\mathbf{r}=1)})}_{g(p(\mathbf{R}, \mathbf{Z}))} = \frac{p(\tilde{\mathbf{Z}}^{(\mathbf{r}=1)}, \mathbf{R} = 1)}{\underbrace{p(\mathbf{R} = 1 | \tilde{\mathbf{Z}}^{(\mathbf{r}=1)})}_{g(p(\mathbf{R}, \mathbf{Z}))|_{\mathbf{R}=1}}} \times \underbrace{p(\mathbf{R} | \tilde{\mathbf{Z}}^{(\mathbf{r}=1)})}_{g(p(\mathbf{R}, \mathbf{Z}))}.$$

and missingness mechanism $p(\mathbf{R} | \tilde{\mathbf{Z}}^{(\mathbf{r}=1)})$ is identified using the OR parameterization given below:

$$p(\mathbf{R} | \tilde{\mathbf{Z}}^{(\mathbf{r}=1)}) = \frac{1}{\sigma} \times \prod_{k=1}^K p(R_k | R_{-k} = 1, \tilde{\mathbf{Z}}^{(\mathbf{r}=1)}) \times \prod_{k=2}^K \text{OR}(R_k, R_{\prec k} | R_{\succ k} = 1, \tilde{\mathbf{Z}}^{(\mathbf{r}=1)})$$

Proof:

Soundness: The absence of e-colluding paths results in identification.

The absence of a e-colluding path between $Z_i^{(1, \mathbf{r}_{\text{aff}(i)})}$ and $R_k \in \mathbf{R}_{\text{aff}(i)}$ implies that $Z_i^{(1, \mathbf{r}_{\text{aff}(i)})} \notin \text{mb}_{\mathcal{G}}(R_k)$ where

$$\text{mb}_{\mathcal{G}}(V) = \{\text{pa}_{\mathcal{G}}(V), \text{dis}_{\mathcal{G}}(V), \text{pa}_{\mathcal{G}}(\text{dis}_{\mathcal{G}}(V)), \text{ch}_{\mathcal{G}}(V), \text{pa}_{\mathcal{G}}(\text{ch}_{\mathcal{G}}(V)), \text{dis}_{\mathcal{G}}(\text{ch}_{\mathcal{G}}(V)), \text{pa}_{\mathcal{G}}(\text{dis}_{\mathcal{G}}(\text{ch}_{\mathcal{G}}(V)))\}$$

Let $\tilde{\mathbf{Z}}^{(\mathbf{r}_k=1)} = \{Z_i^{(1, \mathbf{r}_{\text{aff}(i)}=1)} : R_k \in \mathbf{R}_{\text{aff}(i)}\} \cup Z_k^{(1, \mathbf{r}_{\text{aff}(k)})}$. In words, $\tilde{\mathbf{Z}}^{(\mathbf{r}_k=1)}$ is the set of all counterfactuals that correspond to full observability, and are indexed (and influenced) by R_k .

By Markov property, we have that $V \perp\!\!\!\perp \mathbf{V} \setminus \text{mb}_{\mathcal{G}}(V) | \text{mb}_{\mathcal{G}}(V)$. Therefore, the absence of e-colluding paths implies the following assumptions:

$$R_k \perp\!\!\!\perp \{Z_i^{(1, \mathbf{r}_{\text{aff}(i)})} \in \tilde{\mathbf{Z}}^{(\mathbf{r}_k=1)}\} \mid \{R \setminus R_k, \tilde{\mathbf{Z}} \setminus Z_i^{(1, \mathbf{r}_{\text{aff}(i)}=1)}\} \quad \text{for } R_k \in \mathbf{R}$$

Given these assumptions, we can identify $p(\mathbf{R} | \tilde{\mathbf{Z}}^{(\mathbf{r}=1)})$ using the OR parameterization. The proof is similar to the identification proof of the no self-censoring model given in [89] and the representation of it in [12]. Since we are assuming $\text{ch}_{\mathcal{G}}(\tilde{\mathbf{Z}}^{(\mathbf{r} \neq 1)}) \cap \mathbf{R} = \emptyset$, it suffices to only ID $p(\mathbf{R} | \tilde{\mathbf{Z}}^{(\mathbf{r}=1)})$.

$$p(\mathbf{R} | \tilde{\mathbf{Z}}^{(\mathbf{r}=1)}) = \frac{1}{\sigma} \times \prod_{k=1}^K p(R_k | \mathbf{R}_{-k} = 1, \tilde{\mathbf{Z}}^{(\mathbf{r}=1)}) \times \prod_{k=2}^K \text{OR}(R_k, \mathbf{R}_{\prec k} | \mathbf{R}_{\succ k} = 1, \tilde{\mathbf{Z}}^{(\mathbf{r}=1)})$$

where notation and OR is consistent with earlier definition.

We have $p(R_k | \mathbf{R}_{-k} = 1, \tilde{\mathbf{Z}}^{(\mathbf{r}=1)}) = p(R_k | \mathbf{R}_{-k} = 1, \tilde{\mathbf{Z}}^{(\mathbf{r}=1)} \setminus \tilde{\mathbf{Z}}^{(\mathbf{r}_k=1)})$. By consistency, every counterfactual past the conditioning bar is equal to the observed proxy. Then, pairwise OR




Figure 11: Examples where $P(\tilde{\mathbf{Z}}^{(r=1)}, \mathbf{R})$ is not identified. The proof is demonstrated using parameter counting. (a)-(c): Extended self-censoring (neighbor-censoring), (a) depicts the data generating process, (b) ADMG for $P(\tilde{\mathbf{Z}}^{(r=1)}, \mathbf{R})$, (c) ADMG for observed law. (d)-(f): Extended colluder, (d) depicts the data generating process, (e) ADMG for $P(\tilde{\mathbf{Z}}^{(r=1)}, \mathbf{R})$, (f) ADMG for observed law.

terms are identified because $\text{OR}(R_k, R_i \mid \mathbf{R}_{-(k,i)} = 1, \tilde{\mathbf{Z}}^{(r=1)})$ is not a function of $\tilde{\mathbf{Z}}^{(r_i=1)}$ or $\tilde{\mathbf{Z}}^{(r_k=1)}$. Finally higher order terms are ID in similar ways [12, 89].

Completeness: One approach to demonstrate completeness is to count the number of parameters required to parameterize the full-observability law of a missing data graph and its corresponding observed law and reason that if the former requires more parameters than the latter, it is impossible to identify all the parameters of the full-observability law uniquely. In order to do so, we assume all variables are binary, and adopt the binary parameterization of Nested Markov Models.

First, we present two simple e-colluding path examples to show how the parameter counting argument proceeds. Thereafter, we provide the general argument for completeness for all graphs.

Consider the simple graph in Fig. 11(a) which has e-self-censoring $Z_2^{(1,r_1=1)} \rightarrow R_1$. To count the parameters required for the full-observability law, we follow the procedure in [12]. The binary parameterization of the full law of a missing data ADMG is the same as counting in an ordinary ADMG, with all the irrelevant counterfactuals projected out, except that deterministic factors $P(Z_i \mid R_j, Z_i^{(1,r_{\text{aff}(i)})})$ can be ignored. This gives us the graph in Fig. 11(b). The 5 parameters associated with this graph are $q(Z_1^{(1,r_2=1)} = 0)$, $q(Z_2^{(1,r_1=1)} = 0)$, $q(R_2 = 0)$ and $q(R_1 = 0 \mid Z_2^{(1,r_1=1)} = 1)$, $q(R_1 = 0 \mid Z_2^{(1,r_1=1)} = 0)$.




Figure 12: (a) e-colluding paths between $Z_i^{(1, r_{\text{aff}(i)}=1)}$ and R_j where $R_j \in \mathbf{R}_{\text{aff}(i)}$ (i and j are neighbors) (b) Projecting out $Z_i^{(1, r_{\text{aff}(i)}=1)}$

Next, we obtain the graph in (c) for the observed law, by projecting out all the counterfactuals in (a). Counting the observed law under full observability in entangled missingness settings has a special consideration: if Z_i appears in the head of a parameter, and any of the corresponding missingness indicators $R_{\text{aff}(i)}$ appear in the tail, the kernel must be restricted to cases where $R_i = 1$ and $\mathbf{R}_{\text{aff}(i)} = \mathbf{1}$. This is because, (1) when $R_i = 0$, the probability of the head taking any value aside from those where $Z_i = ?$ is deterministically 0, and (2) cases where $\mathbf{R}_{\text{aff}(i)}$ are set to values different from 1 are irrelevant to the identification of the full observability law¹¹. The observed law in (c) can be parameterized using only 4 parameters, one each corresponding to $q(R_1 = 0)$, $q(R_2 = 0)$, $q(R_1 = 0, Z_2 = 0 | R_2 = 1)$ and $q(Z_1 = 0 | R_1 = 1, R_2 = 1)$. That is one less than the full observability law and hence the latter may not be uniquely determined from data.

Second, consider the graph in Fig. 11(d), which has an e-collider $Z_2^{(1, r_1=1)} \rightarrow R_3 \leftarrow R_1$. The full observability law is shown in (e) and observed law in (f). Counting for (e) yields a total of 9 parameters: $q(Z_3^{(1)} = 0)$, $q(Z_1^{(1, r_2=1)} = 0)$, $q(Z_2^{(1, r_1=1)} = 0)$, $q(R_1 = 0)$, $q(R_2 = 0)$ and finally, $q(R_3 = 0 | R_1, Z_2^{(1, r_1=1)})$, which accounts for 4 parameters. The observed law, on the other hand, needs only 8 parameters: $q(R_1 = 0)$, $q(R_2 = 0)$, $q(R_3 = 0 | R_1 = 1)$, $q(R_3 = 0 | R_1 = 0)$, $q(Z_3 = 0 | R_3 = 1)$, $q(Z_1 = 0 | R_1 = 1, R_2 = 1)$, $q(Z_2 = 0 | R_1 = 1, R_2 = 1)$ and $q(Z_2 = 0, R_3 = 0 | R_1 = 1, R_2 = 1)$. Hence, it is not possible to uniquely map back to the full observability law from observed data.

Finally, we present the general argument for arbitrary graphs.

Assume that there are n variables in $\tilde{\mathbf{Z}}$. For simplicity, assume that all counterfactuals are independent of each other, i.e, $Z_i^{(1, \mathbf{r}_{\text{aff}(i)}=\mathbf{r})} \perp\!\!\!\perp Z_i^{(1, \mathbf{r}_{\text{aff}(i)}=\mathbf{r}')}$, when $\mathbf{r} \neq \mathbf{r}'$ and $Z_i^{(1, \mathbf{r}_{\text{aff}(i)}=\mathbf{r})} \perp\!\!\!\perp Z_j^{(1, \mathbf{r}_{\text{aff}(j)}=\mathbf{r}')}$, when \mathbf{r} may or may not be equal to \mathbf{r}' . Graphically, there are no edges between

¹¹The consideration that $R_i = 1$ always holds, but that $\mathbf{R}_{\text{aff}(i)} = \mathbf{1}$ is only for identifying the full observability law. For other cases, considerations vary.

counterfactuals ¹². And by Assumption 1, $\tilde{\mathbf{Z}}^{(r \neq 1)}$ is not a parent of any $R \in \mathbf{R}$, and can be projected out (adding no new edges) without affecting the Markov blanket of any variable in $\tilde{\mathbf{Z}}^{(r=1)}$, \mathbf{R} or proxies \mathbf{Z} .

Shown in Fig. 12 (a)-(d) are all possible e-colluding paths between $Z_i^{(1, \mathbf{r}_{\text{aff}(i)}=1)}$ and R_j . Here, $R_j \in \mathbf{R}_{\text{aff}(i)}$. Assume there are K variables, $V_1 \cdots V_K$, that lie on the smallest collider path between $Z_i^{(1, \mathbf{r}_{\text{aff}(i)}=1)}$ and R_j . For this to be the smallest collider path, we require that $Z_l^{(1, \mathbf{r}_{\text{aff}(l)}=1)}$ or R_l are not in $V_1, \cdots V_K$ when i and l share the same neighborhood, i.e., when $R_l \in \mathbf{R}_{\text{aff}(i)}$ and $R_i \in \mathbf{R}_{\text{aff}(l)}$. If not, we could truncate the path to have the smallest such path that goes between $Z_i^{(1, \mathbf{r}_{\text{aff}(i)}=1)}$ and R_j such that i and j are neighbors. Fig. 12 (e) shows the projection of (a) and (b), and (f) shows the projection of (c) and (d). $V^* \in \mathbf{Z} \setminus Z_l$, $\mathbf{R} \setminus R_l$, i and l are neighbors.

We now consider each of these paths (a)-(d) and their corresponding latent projections (e) and (f), as if they appear in a larger graph that is otherwise completely disconnected. Akin to what we did with the examples earlier, we count the number of Möbius parameters (as a function of K), and show that the full observability law always has more parameters than the observed law. When we place these colluding paths in a larger graph with arbitrary connectivity, the full observability law is still not identified because of the discrepancy arising from the g-colluding path alone. That is, any edge super graph (super model) is also not identified.

The following fact will be used towards counting parameters in a binary model: Given a bidirected chain $V_1 \leftrightarrow \cdots \leftrightarrow V_{K'}$ of length K' , the number of parameters required to parameterize this chain is $\frac{K'(K'+1)}{2}$, corresponding to parameters given by the following:

$$\begin{aligned} q(V_1 = 0) q(V_1 = V_2 = 0) \cdots q(V_1 = \cdots V_{K'} = 0) &: K' \text{ params} \\ q(V_2 = 0) q(V_2 = V_3 = 0) \cdots q(V_2 = \cdots V_{K'} = 0) &: K' - 1 \text{ params} \\ &\dots \\ q(V_{K'} = 0) &: 1 \text{ param} \end{aligned}$$

Parameter counting for Fig. 12(a), (b), (e)

1. Number of Möbius parameters in Fig. 12(a) is $\frac{(K+2)(K+3)}{2}$
 - It is a bidirected chain $Z_i^{(1, \mathbf{r}_{\text{aff}(i)}=1)} \leftrightarrow \cdots \leftrightarrow R_j$ of length $K' = K + 2$.
2. Number of Möbius parameters in Fig. 12(b) is $\frac{(K+2)(K+3)}{2}$
 - $q(Z_i^{(1, \mathbf{r}_{\text{aff}(i)}=1)} = 0)$ i.e. 1 parameter.
 - A bidirected chain $V_2 \leftrightarrow \cdots \leftrightarrow V_K \leftrightarrow R_j$ of length $K' = K$, i.e. $\frac{K(K+1)}{2}$ parameters.
 - Intrinsic sets involving V_1 , i.e., $q(V_1 = 0 | Z_i^{(1, \mathbf{r}_{\text{aff}(i)}=1)})$, $q(V_1 = V_2 = 0 | Z_i^{(1, \mathbf{r}_{\text{aff}(i)}=1)})$, \cdots , $q(V_1 = V_2 = \cdots R_j = 0 | Z_i^{(1, \mathbf{r}_{\text{aff}(i)}=1)})$ which have 2 parameters each, leading to $2 * (K + 1)$ parameters totally.
3. Number of Möbius parameters in Fig. 12(e) is $\frac{(K+2)(K+3)}{2} - 1$

¹²This assumption does not affect the generality of completeness for other types of models because this model is a submodel of others that involve dependence between counterfactuals, and completeness in this model guarantees completeness in others.

- The bidirected chain $V_1^* \leftrightarrow \dots \leftrightarrow R_j$ of length $K' = K + 1$ has $\frac{(K+1)(K+2)}{2}$ parameters. It should be noted that for any $Z_l \in \{V_1 \dots V_K\}$, in the head of a Möbius parameter, if $R \in R_l \cup \mathbf{R}_{\text{aff}(l)}$ is the parameter's tail, R is always set to 1 deterministically. Hence, it reduces to counting the simple bidirected chain.
- The number of intrinsic sets involving Z_i is $K + 1$ (and not $K + 2$) since R_i is not fixable, and the set $\{Z_i, V_1^*, \dots, V_K^*\}$ is not intrinsic. Each of these intrinsic sets corresponds to two parameters, so $2 * (K + 1)$ parameters.

Parameter counting for Fig. 12(c), (d), (f)

1. Number of Möbius parameters in Fig. 12(c) is $\frac{(K+2)(K+3)}{2}$
 - $q(R_j = 0)$ i.e. 1 parameter.
 - A bidirected chain $Z_i^{(1, \mathbf{r}_{\text{aff}(i)} = \mathbf{1})} \leftrightarrow \dots V_{K-1}$ of length $K' = K$, so $\frac{K(K+1)}{2}$ parameters.
 - Intrinsic sets involving R_K , i.e., $q(R_K = 0 | R_j)$, $q(R_K = V_{K-1} = 0 | R_j)$, \dots , $q(R_K = V_{K-1} = \dots, Z_i^{(1, \mathbf{r}_{\text{aff}(i)} = \mathbf{1})} = 0 | R_j)$, i.e., $2 * (K + 1)$ parameters.
2. Number of Möbius parameters in Fig. 12(d) is $\frac{(K+2)(K+3)}{2}$
 - $q(Z_i^{(1, \mathbf{r}_{\text{aff}(i)} = \mathbf{1})} = 0)$ i.e. 1 parameter.
 - $q(R_j = 0)$, i.e., 1 parameter.
 - A bidirected chain $V_2 \leftrightarrow \dots \leftrightarrow V_{K-1}$ of length $K' = K - 2$, i.e. $\frac{(K-2)(K-1)}{2}$ parameters.
 - Intrinsic sets involving V_1 and not R_K , i.e., $q(V_1 = 0 | Z_i^{(1, \mathbf{r}_{\text{aff}(i)} = \mathbf{1})})$, $q(V_1 = V_2 = 0 | Z_i^{(1, \mathbf{r}_{\text{aff}(i)} = \mathbf{1})})$, \dots , $q(V_1 = V_2 = \dots V_{K-1} = 0 | Z_i^{(1, \mathbf{r}_{\text{aff}(i)} = \mathbf{1})})$ which give 2 parameters each, leading to $2 * (K - 1)$ parameters.
 - Intrinsic sets involving R_K and not V_1 , i.e., $q(R_K = 0 | Z_i^{(1, \mathbf{r}_{\text{aff}(i)} = \mathbf{1})})$, $q(R_K = V_2 = 0 | Z_i^{(1, \mathbf{r}_{\text{aff}(i)} = \mathbf{1})})$, \dots , $q(R_K = V_2 = \dots V_{K-1} = 0 | Z_i^{(1, \mathbf{r}_{\text{aff}(i)} = \mathbf{1})})$ which give 2 parameters each, leading to $2 * (K - 1)$ parameters.
 - The one intrinsic set involving both V_1 and R_K , i.e. $q(V_1 = V_2 = \dots = R_K = 0 | R_j, Z_i^{(1, \mathbf{r}_{\text{aff}(i)} = \mathbf{1})})$ corresponding to 4 parameters.
3. Number of Möbius parameters in Fig. 12(f) is $\frac{(K+2)(K+3)}{2} - 1$
 - $q(R_j = 0)$, i.e., 1 parameter.
 - The bidirected chain $Z_i \leftrightarrow V_1^* \leftrightarrow \dots \leftrightarrow V_{K-1}$ of length $K' = K$ has $\frac{(K)(K+1)}{2}$ parameters. As before, for any $Z_l \in \{V_1 \dots V_{K-1}\}$, in the head of a Möbius parameter, if $R \in R_l \cup \mathbf{R}_{\text{aff}(l)}$ is the parameter's tail, R is always set to 1 deterministically.
 - Intrinsic sets involving R_K , i.e., $q(R_K = 0 | R_j)$, $q(R_K, V_{K-1} = 0 | R_j)$, \dots , $q(R_K, V_{K-1}, \dots, V_1 = 0 | R_j)$ corresponding to $2 * K$ parameters, and the intrinsic set $q(R_K, V_{K-1}, \dots, V_1 = 0, Z_i = 0 | R_j = 1)$ which only corresponds to 1 parameter (instead of 2) since $R_j \in \mathbf{R}_{\text{aff}(i)}$ and has to be set to 1.

□

REFERENCES

- [1] Donald B. Rubin.
Inference and missing data.
Biometrika, 63(3):581–592, 1976.
- [2] Roderick J Little.
Missing data assumptions.
Annual Review of Statistics and Its Application, 8:89–107, 2021.
- [3] R.J.A. Little and D.B. Rubin.
Statistical Analysis with Missing Data.
Wiley Series in Probability and Statistics. Wiley, 2002.
- [4] M Maria Glymour.
Using causal diagrams to understand common problems in social epidemiology.
Methods in social epidemiology, pages 393–428, 2006.
- [5] Rhian M. Daniel, Michael G. Kenward, Simon N. Cousens, and Bianca L. De Stavola.
Using causal diagrams to guide analysis in missing data problems.
Statistical Methods in Medical Research, 21(3):243–256, 2012.
- [6] Fernando Martel García.
Definition and diagnosis of problematic attrition in randomized controlled experiments.
Available at SSRN 2302735, 2013.
- [7] Karthika Mohan, Judea Pearl, and Jin Tian.
Graphical models for inference with missing data.
Advances in neural information processing systems, 26, 2013.
- [8] Felix Thoemmes and Norman Rose.
A cautious note on auxiliary variables that can increase bias in missing data problems.
Multivariate Behavioral Research, 49(5):443–459, 2014.
- [9] Jin Tian.
Missing at random in graphical models.
In *Artificial Intelligence and Statistics*, pages 977–985. PMLR, 2015.
- [10] Ilya Shpitser.
Consistent estimation of functions of data missing non-monotonically and not at random.
Advances in Neural Information Processing Systems, 29:3144–3152, 2016.
- [11] Rohit Bhattacharya, Razieh Nabi, Ilya Shpitser, and James Robins.
Identification in missing data models represented by directed acyclic graphs.
In *Proceedings of the Thirty Fifth Conference on Uncertainty in Artificial Intelligence (UAI-35th)*. AUAI Press, 2019.
- [12] Razieh Nabi, Rohit Bhattacharya, and Ilya Shpitser.
Full law identification in graphical models of missing data: Completeness results.
In *International Conference on Machine Learning*, pages 7153–7163. PMLR, 2020.
- [13] Karthika Mohan and Judea Pearl.
Graphical models for processing missing data.
Journal of the American Statistical Association, pages 1–16, 2021.
- [14] Daniel O Scharfstein, Jaron JR Lee, Aidan McDermott, Aimee Campbell, Edward Nunes, Abigail G Matthews, and Ilya Shpitser.
Markov-restricted analysis of randomized trials with non-monotone missing binary outcomes: Sensitivity analysis and identification results.
arXiv preprint arXiv:2105.08868, 2021.
- [15] Razieh Nabi and Rohit Bhattacharya.
On testability and goodness of fit tests in missing data models.
arXiv preprint arXiv:2203.00132, 2022.
- [16] Razieh Nabi, Rohit Bhattacharya, Ilya Shpitser, and James Robins.
Causal and counterfactual views of missing data models.
arXiv preprint arXiv:2210.05558, 2022.
- [17] P. M. Aronow and C. Samii.
Estimating average causal effects under general interference.
Technical Report, 2013.
- [18] S. Athey, D. Eckles, and G.W. Imbens.
Exact p-values for network interference.
Journal of the American Statistical Association, 113 (521):230–240, 2018.

- [19] G. Basse, A. Feller, and P. Toulis.
Randomization tests of causal effects under interference.
Biometrika, 106 (2):487–494, 2019.
- [20] G. W. Basse and E. M. Airoidi.
Model-assisted design of experiments in the presence of network-correlated outcomes.
Biometrika, 105 (4):849–858, 2018.
- [21] J. Bowers, F. M. M. and P. C.
Reasoning about interference between units: A general framework.
Political Analysis, 21:97–124, 2013.
- [22] X. Cai, W. W. Loh, and F. W. Crawford.
Identification of causal intervention effects under contagion.
arXiv preprint arXiv:1912.04151, 2019.
- [23] Daniel J Eck, Olga Morozova, and Forrest W Crawford.
Randomization for the susceptibility effect of an infectious disease intervention.
Journal of Mathematical Biology, 85(4):1–22, 2022.
- [24] Dean Eckles, Brian Karrer, and Johan Ugander.
Design and analysis of experiments in networks: Reducing bias from interference.
Journal of Causal Inference, 5(1), 2017.
- [25] Laura Forastiere, Edoardo M Airoidi, and Fabrizia Mealli.
Identification and estimation of treatment and interference effects in observational studies on networks.
Journal of the American Statistical Association, 116(534):901–918, 2021.
- [26] Razieh Nabi, Joel Pfeiffer, Denis Charles, and Emre Kıcıman.
Causal inference in the presence of interference in sponsored search advertising.
Frontiers in big Data, 5, 2022.
- [27] Bryan S. Graham, Guido W. Imbens, and Geert Ridder.
Measuring the effects of segregation in the presence of social spillovers: a nonparametric approach.
Technical report, National Bureau of Economic Research, 2010.
- [28] M. Elizabeth Halloran and Claudio J. Struchiner.
Causal inference in infectious diseases.
Epidemiology, pages 142–151, 1995.
- [29] M Elizabeth Halloran and Michael G Hudgens.
Causal inference for vaccine effects on infectiousness.
The international journal of biostatistics, 8(2):1–40, 2012.
- [30] Guanglei Hong and Stephen W. Raudenbush.
Evaluating kindergarten retention policy: A case study of causal inference for multilevel observational data.
Journal of the American Statistical Association, 101(475):901–910, 2006.
- [31] Michael G Hudgens and M Elizabeth Halloran.
Toward causal inference with interference.
Journal of the American Statistical Association, 103(482):832–842, 2008.
- [32] Ravi Jagadeesan, Natesh S Pillai, and Alexander Volfovsky.
Designs for estimating the treatment effect in networks with interference.
The Annals of Statistics, 48(2):679–712, 2020.
- [33] Panos Toulis and Edward Kao.
Estimation of causal peer influence effects.
In *International conference on machine learning*, pages 1489–1497. PMLR, 2013.
- [34] Michael P Leung.
Treatment and spillover effects under network interference.
Review of Economics and Statistics, 102(2):368–380, 2020.
- [35] Lan Liu and Michael G Hudgens.
Large sample randomization inference of causal effects in the presence of interference.
Journal of the american statistical association, 109(505):288–301, 2014.
- [36] Georgia Papadogeorgou, Fabrizia Mealli, and Corwin M Zigler.
Causal inference with interfering units for cluster and population level treatment allocation programs.
Biometrics, 75(3):778–787, 2019.
- [37] David Puelz, Guillaume Basse, Avi Feller, and Panos Toulis.
A graph-theoretic approach to randomization tests of causal effects under general interference.
arXiv preprint arXiv:1910.10862, 2019.
- [38] Paul R. Rosenbaum.
Interference between units in randomized experiments.
Journal of the American Statistical Association, 102(477):191–200, 2007.

- [39] Donald B Rubin.
Comment: Neyman (1923) and causal inference in experiments and observational studies.
Statistical Science, 5(4):472–480, 1990.
- [40] Fredrik Sävje.
Causal inference with misspecified exposure mappings.
arXiv preprint arXiv:2103.06471, 2021.
- [41] Fredrik Sävje, Peter Aronow, and Michael Hudgens.
Average treatment effects in the presence of unknown interference.
Annals of statistics, 49(2):673, 2021.
- [42] Michael E Sobel.
What do randomized studies of housing mobility demonstrate? causal inference in the face of interference.
Journal of the American Statistical Association, 101(476):1398–1407, 2006.
- [43] Eric J Tchetgen Tchetgen and Tyler J VanderWeele.
On causal inference in the presence of interference.
Statistical methods in medical research, 21(1):55–75, 2012.
- [44] Panos Toulis, Alexander Volfovsky, and Edoardo M Airoidi.
Propensity score methodology in the presence of network entanglement between treatments.
arXiv preprint arXiv:1801.07310, 2018.
- [45] Tyler J VanderWeele.
Direct and indirect effects for neighborhood-based clustered and longitudinal data.
Sociological methods & research, 38(4):515–544, 2010.
- [46] Changge Chang, Yi Deng, Xiaoqian Jiang, and Qi Long.
Multiple imputation for analysis of incomplete data in distributed health data networks.
Nature communications, 11(1):1–11, 2020.
- [47] Jeffrey A Smith, James Moody, and Jonathan H Morgan.
Network sampling coverage ii: The effect of non-random missing data on network measurement.
Social networks, 48:78–99, 2017.
- [48] Krista J Gile and Mark S Handcock.
Analysis of networks with missing data with application to the national longitudinal study of adolescent health.
Journal of the Royal Statistical Society: Series C (Applied Statistics), 66(3):501–519, 2017.
- [49] Zack W Almquist and Carter T Butts.
Dynamic network analysis with missing data: theory and methods.
Statistica Sinica, 28(3):1245–1264, 2018.
- [50] Bruce Sacerdote.
Peer effects in education: How might they work, how big are they and how much do we know thus far?
In *Handbook of the Economics of Education*, volume 3, pages 249–277. Elsevier, 2011.
- [51] Judea Pearl.
Causality: Models, Reasoning, and Inference.
Cambridge University Press, 2 edition, 2009.
ISBN 978-0521895606.
- [52] Marloes Maathuis, Mathias Drton, Steffen Lauritzen, and Martin Wainwright.
Handbook of Graphical Models.
CRC Press, 2018.
- [53] Thomas S. Richardson and Jamie M. Robins.
Single world intervention graphs (SWIGs): A unification of the counterfactual and graphical approaches to causality.
preprint, 2013.
- [54] Judea Pearl.
Probabilistic Reasoning in Intelligent Systems.
Morgan and Kaufmann, San Mateo, 1988.
- [55] Judea Pearl.
Causality.
Cambridge University Press, 2009.
- [56] Miguel A Hernan and James M Robins.
Causal Inference: What If.
CRC Press: Taylor and Francis Group, 2020.
- [57] Thomas S. Verma and Judea Pearl.
Equivalence and synthesis of causal models.
Technical Report R-150, Department of Computer Science, University of California, Los Angeles, 1990.

- [58] Thomas S. Richardson.
Markov properties for acyclic directed mixed graphs.
Scandinavian Journal of Statistics, 30(1):145–157, 2003.
- [59] Jin Tian and Judea Pearl.
On the testable implications of causal models with hidden variables.
In *Proceedings of the Eighteenth Conference on Uncertainty in Artificial Intelligence (UAI-02)*, volume 18, pages 519–527. AUAI Press, Corvallis, Oregon, 2002.
- [60] Ilya Shpitser and Judea Pearl.
Identification of joint interventional distributions in recursive semi-Markovian causal models.
In *Proceedings of the Twenty-First National Conference on Artificial Intelligence (AAAI-06)*. AAAI Press, Palo Alto, 2006.
- [61] Yimin Huang and Marco Valorta.
Pearl’s calculus of interventions is complete.
In *Twenty Second Conference On Uncertainty in Artificial Intelligence*, 2006.
- [62] Thomas S. Richardson, Robin J. Evans, James M. Robins, and Ilya Shpitser.
Nested Markov properties for acyclic directed mixed graphs, 2017.
Working paper.
- [63] Rohit Bhattacharya, Razieh Nabi, and Ilya Shpitser.
Semiparametric inference for causal effects in graphical models with hidden variables.
Journal of Machine Learning Research, 23:1–76, 2022.
- [64] Elizabeth L. Ogburn and Tyler J. VanderWeele.
Causal diagrams for interference.
Statistical Science, 29(4):559–578, 2014.
- [65] M.G. Hudgens and M.E. Halloran.
Toward causal inference with interference.
Journal of the American Statistical Association, 103(482):832–842, 2008.
- [66] Eric J. Tchetgen Tchetgen and Tyler J. VanderWeele.
On causal inference in the presence of interference.
Statistical Methods in Medical Research, 21(1):55–75, 2012.
- [67] Eric J Tchetgen Tchetgen, Isabel R Fulcher, and Ilya Shpitser.
Auto-g-computation of causal effects on a network.
Journal of the American Statistical Association, 116(534):833–844, 2021.
- [68] A.P. Dempster, N.M. Laird, and D.B. Rubin.
Maximum likelihood from incomplete data via the EM algorithm.
Journal of the Royal Statistical Society, Series B, 39:1–38, 1977.
- [69] Nicholas J Horton and Nan M Laird.
Maximum likelihood analysis of generalized linear models with missing covariates.
Statistical Methods in Medical Research, 8(1):37–50, 1999.
- [70] Donald B Rubin.
An overview of multiple imputation.
In *Proceedings of the survey research methods section of the American statistical association*, pages 79–84. Citeseer Princeton, NJ, USA, 1988.
- [71] Joseph L Schafer.
Multiple imputation: a primer.
Statistical methods in medical research, 8(1):3–15, 1999.
- [72] James M. Robins, Andrea Rotnitzky, and Lue P. Zhao.
Estimation of regression coefficients when some regressors are not always observed.
Journal of the American Statistical Association, 89:846–866, 1994.
- [73] Lingling Li, Changyu Shen, Xiaochun Li, and James M Robins.
On weighting approaches for missing data.
Statistical methods in medical research, 22(1):14–30, 2013.
- [74] Eric J Tchetgen Tchetgen, Linbo Wang, and BaoLuo Sun.
Discrete choice models for nonmonotone nonignorable missing data: Identification and inference.
Statistica Sinica, 28(4):2069, 2018.
- [75] Margaret C Wu and R. J. Carroll.
Estimation and comparison of changes in the presence of informative right censoring by modeling the censoring process.
Biometrics, 175-188, 1988.

- [76] Sheng Wang, Jun Shao, and Jae Kwang Kim.
An instrumental variable approach for identification and estimation with nonignorable nonresponse.
Statistica Sinica, pages 1097–1116, 2014.
- [77] Wang Miao, Peng Ding, and Zhi Geng.
Identifiability of normal and normal mixture models with nonignorable missing data.
Journal of the American Statistical Association, 111(516):1673–1683, 2016.
- [78] Wang Miao and Eric J Tchetgen Tchetgen.
On varieties of doubly robust estimators under missingness not at random with a shadow variable.
Biometrika, 103(2):475–482, 2016.
- [79] BaoLuo Sun, Lan Liu, Wang Miao, Kathleen Wirth, James Robins, and Eric J Tchetgen Tchetgen.
Semiparametric estimation with data missing not at random using an instrumental variable.
Statistica Sinica, 28(4):1965, 2018.
- [80] Matthew O Jackson, Zhongjian Lin, and Ning Neil Yu.
Adjusting for peer-influence in propensity scoring when estimating treatment effects.
Available at SSRN 3522256, 2020.
- [81] C Hendricks Brown.
Protecting against nonrandomly missing data in longitudinal studies.
Biometrics, pages 143–155, 1990.
- [82] Karthika Mohan, Felix Thoemmes, and Judea Pearl.
Estimation with incomplete data: The linear case.
In *Proceedings of the International Joint Conferences on Artificial Intelligence Organization*, 2018.
- [83] Werner Ebeling and Rainer Feistel.
Physics of Self-organization and Evolution.
John Wiley & Sons, 2011.
- [84] Hua Yun Chen.
A semiparametric odds ratio model for measuring association.
biometrics, 63:413–421, 2007.
- [85] Ryan M Andrews and Vanessa Didelez.
Insights into the cross-world independence assumption of causal mediation analysis.
Epidemiology, 32(2):209–219, 2020.
- [86] Anastasios Tsiatis.
Semiparametric Theory and Missing Data.
Springer-Verlag New York, 1st edition edition, 2006.
- [87] Steffan L. Lauritzen.
Graphical Models.
Oxford, U.K.: Clarendon, 1996.
- [88] Robin J. Evans and Thomas S. Richardson.
Markovian acyclic directed mixed graphs for discrete data.
Annals of Statistics, pages 1–30, 2014.
- [89] Daniel Malinsky, Ilya Shpitser, and Eric J Tchetgen Tchetgen.
Semiparametric inference for nonmonotone missing-not-at-random data: the no self-censoring model.
Journal of the American Statistical Association, pages 1–9, 2021.