General Identification of Dynamic Treatment Regimes Under Interference

Eli Sherman
Johns Hopkins University

David Arbour
Adobe Inc.

Ilya Shpitser
Johns Hopkins University

Abstract

In many applied fields, researchers are often interested in tailoring treatments to unit-level characteristics in order to optimize an outcome of interest. Methods for identifying and estimating treatment policies are the subject of the dynamic treatment regime literature. Separately, in many settings the assumption that data are independent and identically distributed does not hold due to inter-subject dependence. The phenomenon where a subject's outcome is dependent on his neighbor's exposure is known as interference. These areas intersect in myriad real-world settings. In this paper we consider the problem of identifying optimal treatment policies in the presence of interference. Using a general representation of interference, via Lauritzen-Wermuth-Freydenburg chain graphs (Lauritzen and Richardson, 2002), we formalize a variety of policy interventions under interference and extend existing identification theory (Tian, 2008; Sherman and Shpitser, 2018). Finally, we illustrate the efficacy of policy maximization under interference in a simulation study.

1 Introduction

In areas such as precision medicine, economics, and political science, identifying interventions that are optimally tailored to each subject is often of interest. Dynamic treatment regimes (DTRs), which are counterfactual policies used for treatment assignment, represent a promising approach to tailoring treatments. Typically, a causal model is assumed to be known, with analyst-specified exposure and outcome variables. The analyst considers setting the exposure variable according to a treatment policy which is a function of other model variables. She then estimates the counterfactual effect of several candidate policies and picks the one with the best expected outcome. This setup has been extended to sequential settings (Laber et al., 2014; Chakraborty and Moodie, 2013; Nabi et al., 2018).

A key obstacle to obtaining optimal strategies from observational data is identification. The effect of an intervention in a causal model is said to be *identified* if the effect can be expressed as a function of observed data. In algebraic formulations, identification requires carefully enumerating necessary assumptions. In contrast, causal graphical models provide a concise framework for representing assumptions, with numerous general results characterizing identification criteria. In the context of DTRs, Robins (1986) gave an approach for identification of node (i.e., fixed value) and policy interventions in fully-observed directed acyclic graphs (DAGs). Tian (2008) and Shpitser and Sherman (2018) extended this approach to latent-variable DAG models.

Orthogonal to treatment customization, classical causal inference assumes independence among study subjects. In many settings, however, subjects' exposures causally affect their neighbors' outcomes. This phenomenon, known as interference (Cox, 1958), has recently attracted substantial attention. Hudgens and Halloran (2008) serves as a seminal paper; it defines network-level effects and provides elementary identification conditions. Ogburn et al. (2014) formalizes DAG representations of interference. Several papers propose relational (Maier et al., 2013) or chain graph representations of interference (Peña, 2018; Tchetgen et al., 2017; Ogburn et al., 2018). Sherman and Shpitser (2018) is closest to the present work; it explores non-parametric identification in the presence of unobserved confounding.

A recent paper also considers policies under interference (Viviano, 2019). Our work differs substantially: Viviano focuses on welfare maximization and assumes units are identically distributed. Our characterization of policy interventions generalizes welfare maximization and our network representation is non-parametric.

Proceedings of the 23^{rd}International Conference on Artificial Intelligence and Statistics (AISTATS) 2020, Palermo, Italy. PMLR: Volume 108.

Motivating Policies in Networks. In this work, we consider identification of DTRs in the interference setting. As motivation, consider the following example from psephology (the study of elections) (Blackwell, 2013): candidates running for public office target voters by purchasing television advertisements; each candidate must decide how many ads to buy and whether they should be positive ("my record is stellar") or negative ("my opponent is scandalous").

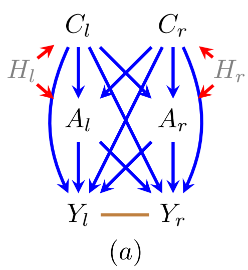

Figure 1: Graphical representations of competitive dynamics in an election campaign, where (a) H's represent latent confounders, (b) is a latent projection with H's replaced by bi-directed edges, and (c) is an alternative model where A's exhibit best-response dynamics.

These dynamics can be represented via the causal graphs in Fig. 1. For each candidate, C denotes observed pre-decision covariates, such as prior polling performance, previous advertising, and cash on hand, A represents the candidate's advertising decision, Y represents polling performance in the current decision time frame, and H represents unobserved confounders that affect the candidate's pre-decision covariates and decision but don't directly affect the outcome. l and r index the variables for a left- and right-leaning candidate respectively. Directed edges denote a direct causal relationship, while undirected edges denote non-causal dependence (e.g. $A_l - A_r$ could be interpreted as candidates acting based on beliefs about what each other will do). While we use this two-candidate example as motivation throughout this manuscript, our contributions apply to networks of arbitrary size and topology.

The remainder of this paper is organized as follows: we fix notation and discuss relevant background work in Secs. 2 and 3. We characterize the variety of possible policy interventions in Sec. 4. We then give a novel identification result for effects of policy interventions in Lauritzen-Wermuth-Freydenburg (LWF) latent-variable chain graphs (Lauritzen, 1996; Lauritzen and Richardson, 2002) in Sec. 5. We demonstrate estimation of these effects via a simulation study in Sec. 6 and conclude with a discussion of ongoing work.

2 Notation

We first fix notation before describing the task. We employ segregated graphs (SGs) (Shpitser, 2015) to represent causal network dynamics. SGs are a class of mixed graphical model which are a super-model of latent-variable LWF chain graphs (CGs), which are themselves a super-model of Markov random fields (MRFs) and DAGs. SGs permit three edge types – undirected ($-$), directed ($\rightarrow$), and bi-directed ($\leftrightarrow$) – and have the property that no variable has both an incident undirected and bi-directed edge.

We adopt standard graphical model notation. We denote random variables (interchangeably, vertices in graphs) by capital letters V and their realizations in lowercase v, with sets in boldface, $\mathbf{V}$ and $\mathbf{v}$. We use standard genealogical notions for graphical relationships. For a variable $V \in \mathbf{V}$ in a graph $\mathcal{G}$, parents $\text{pa}_{\mathcal{G}}(V) \equiv \{W \in \mathbf{V} : W \rightarrow V \text{ in } \mathcal{G}\}$, children $\text{ch}_{\mathcal{G}}(V) \equiv \{W \in \mathbf{V} : V \rightarrow W \text{ in } \mathcal{G}\}$, ancestors $\text{an}_{\mathcal{G}}(V) \equiv \{W \in \mathbf{V} : W \rightarrow \cdots \rightarrow V \text{ in } \mathcal{G}\}$, descendants $\text{de}_{\mathcal{G}}(V) \equiv \{W \in \mathbf{V} : V \rightarrow \cdots \rightarrow W \text{ in } \mathcal{G}\}$, neighbor $\text{nb}_{\mathcal{G}}(V) \equiv \{W \in \mathbf{V} : W - V \text{ in } \mathcal{G}\}$, non-descendant $\text{nd}_{\mathcal{G}}(V) \equiv \mathbf{V} \setminus \text{de}_{\mathcal{G}}(V)$, and district $\text{dis}_{\mathcal{G}}(V) \equiv \{W \in \mathbf{V} : W \leftrightarrow \cdots \leftrightarrow V \text{ in } \mathcal{G}\}$.

Further, the anterior $\text{ant}_{\mathcal{G}}(V)$ is the set of nodes with a partially directed path – a path containing only $\rightarrow$ and $-$ edges such that no set of undirected edges can be oriented to form a directed cycle – *into* V. The exterior $\text{ext}_{\mathcal{G}}(V)$ is the set of nodes with a partially directed path *out of* V. In turn, the strict exterior $\overline{\text{ext}}_{\mathcal{G}}(V) \subseteq \text{ext}_{\mathcal{G}}(V)$ omits V and the set $\{W \in \mathbf{V} : W - \cdots - V\}$. By convention, $\text{ext}_{\mathcal{G}}(V) \cap \text{ant}_{\mathcal{G}}(V) \cap \text{dis}_{\mathcal{G}}(V) = \{V\}$. These and the above notions can be extended to sets, e.g., for $\mathbf{S} \subseteq \mathbf{V}$, we have $\text{pa}_{\mathcal{G}}(\mathbf{S}) = \cup_{S \in \mathbf{S}} \text{pa}_{\mathcal{G}}(S)$ and, disjunctively, $\text{pa}^s_{\mathcal{G}}(\mathbf{S}) = \text{pa}_{\mathcal{G}}(\mathbf{S}) \setminus \mathbf{S}$. When the relevant graph is clear from context, we drop the $\mathcal{G}$ subscript.

For graphs with a partial ordering $\prec$ on $\mathbf{V}$, let $\mathbf{V}_{\prec A}$ denote A's predecessors in the ordering. For a set $\mathbf{S} \subseteq \mathbf{V}$ in $\mathcal{G}$, let $\mathcal{G}_{\mathbf{S}}$ refers to the subgraph of $\mathcal{G}$ containing only $\mathbf{S}$ and edges connecting nodes in $\mathbf{S}$.

Finally, we use the notion of a *block* to refer to a set of variables connected by an undirected path. A node with no incident bi-directed nor undirected edges is a *trivial* block *and* a trivial district. The sets of blocks, non-trivial blocks, districts, and cliques in $\mathcal{G}$ are denoted by $\mathcal{B}(\mathcal{G})$, $\mathcal{B}^{nt}(\mathcal{G})$, $\mathcal{D}(\mathcal{G})$, and $\mathcal{C}(\mathcal{G})$ respectively. In segregated graphs $\mathcal{D}(\mathcal{G})$ and $\mathcal{B}^{nt}(\mathcal{G})$ partition $\mathbf{V}$.

2.1 Statistical Graphical Models

Segregated graphs and their submodels were originally conceived as *statistical* models over random variables, encoding conditional independences in their factorization. For instance, a distribution $p(\mathbf{V})$ is 'Markov relative to' a CG $\mathcal{G}(\mathbf{V})$ if it factorizes according to the two-level factorization $p(\mathbf{V}) =$

$$\prod_{\mathbf{B} \in \mathcal{B}(\mathcal{G})} p(\mathbf{B} |\, \text{pa}(\mathbf{B})) = \prod_{\mathbf{B} \in \mathcal{B}(\mathcal{G})} \frac{\prod_{\mathbf{C} \in \mathcal{C}^\star} \phi_{\mathbf{C}}(\mathbf{C})}{Z(\text{pa}(\mathbf{B}))}, \quad (1)$$

where $\mathcal{C}^\star = \{\mathbf{C} \in \mathcal{C}((\mathcal{G}^a_{\mathbf{B} \cup \text{pa}_{\mathcal{G}}(\mathbf{B})}) : \mathbf{C} \not\subseteq \text{pa}_{\mathcal{G}}(\mathbf{B})\}$, and Z is a normalization function. $\mathcal{G}^a_{\mathbf{B} \cup \text{pa}_{\mathcal{G}}(\mathbf{B})}$ is an augmented graph (Lauritzen, 1996): it is undirected and contains edges in $\mathcal{G}$ between $\mathbf{B}$, edges between nodes in $\text{pa}_{\mathcal{G}}(\mathbf{B})$ and their children in $\mathbf{B}$, and edges between parents. For the corresponding factorizations for DAGs and MRFs, please see the supplementary materials.

2.2 Causal Graphical Models

In contrast to statistical graphs, causal graphs represent distributions over *counterfactual* variables. For $Y \in \mathbf{V}$ and $\mathbf{A} \subseteq \mathbf{V} \setminus Y$, the counterfactual $Y(\mathbf{a})$ denotes Y's value under the hypothetical scenario in which $\mathbf{A}$ is set to $\mathbf{a}$ via a *node* intervention (Pearl, 2000).

In this paper, we assume Pearl's functional model. In DAGs, counterfactuals $V(\mathbf{a})$ are determined by structural equations $f_V(\mathbf{a}, \epsilon_V)$, which remain invariant under an intervention $\mathbf{a}$; ϵ_V denotes an exogenous random variable for f_V. By *recursive substitution*, we can define all other variables in the model: for $\mathbf{A} \subseteq \mathbf{V} \setminus \{V\}$ and $\mathbf{a}$ in the state space of $\mathbf{A}$, $p(V(\mathbf{a}))$ (sometimes written as $p(\mathbf{V}|\text{do}(\mathbf{a})$ (Pearl, 2000)) is defined as $V(\mathbf{a}_{\text{pa}(V)}, \{W(\mathbf{a}) : W \in \text{pa}(V) \setminus \mathbf{A}\})$.

Causal CGs follow similar semantics. Each variable B in a block $\mathbf{B}$ is determined by a structural equation $f_B(\mathbf{B} \setminus \{B\}, \text{pa}(\mathbf{B}), \epsilon_B)$, a function of other variables in $\mathbf{B}$, the parents of $\mathbf{B}$, and an exogenous variable. Each $\mathbf{B}$'s joint distribution $\mathbf{B}$ is obtained by Gibbs sampling over the structural equations for $\mathbf{B}$ until equilibrium (trivial blocks equilibrate instantly). Assuming an ordering on blocks in $\mathcal{G}$, but not on variables in each block, and iid realizations of ϵ_{B_i}, the data generating process for CGs is given by Procedure 1 (Lauritzen and Richardson, 2002).

Procedure 1 CG Data Generating Process

```
procedure CG-DGP(G, {f_B : B ∈ V})
    for each block B_i ∈ B(G) do
        repeat
            for each variable B_j ∈ B_i do
                B_j ← f_{B_j}(B_i \ B_j, pa_G(B_i), ε_{B_j})
        until equilibrium
    return V
```

A parameter is *identifiable* in a causal model if it can be expressed as a function of observed data. In fully observed DAGs and CGs, all node intervention counterfactuals are identified by the *g-formula* (Robins, 1986) and *chain graph g-formula* (Lauritzen and Richardson, 2002) respectively (first two rows of Table 1).

3 Identification in Latent-Variable Causal Graphical Models

In this section, we review identification theory in latent variable causal models. The current work bridges these literatures: we posit a sound and complete algorithm for the identification of responses to policies in latent variable (LV) causal CGs.

3.1 Re-expressing the ID Algorithm

Tian and Pearl (2002) gave a general condition for identification of node interventions in latent-variable DAGs. Shpitser and Pearl (2006) re-expressed Tian's condition as a concise algorithm and proved that it is complete. Recently, Richardson et al. (2017) rephrased the algorithm in terms of a recursive fixing operator which acts as a modified nested Markov factorization.

Richardson et al. (2017) makes clear the connections between the ID algorithm, which is a modified nested factorization of acyclic directed mixed graphs (ADMGs), and the g-formula (Table 1, first row), which is a modified DAG factorization. This formalism enables straightforward generalizations to other identification settings. For these reasons, we base our SG policy identification results on this framework. The framework relies on several concepts which we highlight here; each existing ID approach is summarized in Table 1. For a complete treatment, please see the supplement.

Latent Projections. Rather than considering LV-DAGs explicitly, Richardson et al. (2017) considers ADMGs. ADMGs permit directed and bi-directed edges and represent equivalence classes of LV-DAGs. Given an LV-DAG $\mathcal{G}(\mathbf{V} \cup \mathbf{H})$, with $\mathbf{V}$ observed and $\mathbf{H}$ latent, the corresponding ADMG $\mathcal{G}(\mathbf{V})$ is obtained via a latent projection operation (Verma and Pearl, 1991). For example, Fig. 1(b) is the latent projection of Fig. 1(a). We also define conditional ADMGs (CADMGs), which partition nodes into random $\mathbf{V}$ and fixed $\mathbf{W}$ variables. CADMGs with $\mathbf{W} = \emptyset$ are trivially ADMGs.

Segregated graphs are the chain graph analogue of ADMGs, where SGs represent an equivalence class of LV-CGs. For a latent variable CG $\mathcal{G}(\mathbf{V} \cup \mathbf{H})$, $\mathbf{H}$ is *block-safe* (Sherman and Shpitser, 2018) if no $V \in \mathbf{V}$ has a latent parent and no latent $H \in \mathbf{H}$ has an incident undirected edge. By applying the same latent projection operation mentioned above to a LV-CG with block-safe $\mathbf{H}$, one obtains the corresponding SG.

Graph Type	Latents	Intervention Type	$\mathbf{Y}^\star$	Modified Factorization
DAG	No	Node – $\mathbf{a}$	N/A	$\prod_{V\in\mathbf{V}\setminus\mathbf{A}} p(V\,\vert\,\mathrm{pa}(V))\vert_{\mathbf{A}=\mathbf{a}}$
CG	No	Node – $\mathbf{a}$	N/A	$\prod_{\mathbf{B}\in\mathcal{B}(\mathcal{G})} p(\mathbf{B}\setminus\mathbf{A}\vert\,\mathrm{pa}(\mathbf{B}),\mathbf{B}\cap\mathbf{A})\vert_{\mathbf{A}=\mathbf{a}}$
ADMG	Yes	Node – $\mathbf{a}$	$\mathrm{an}_{\mathcal{G}_{\mathbf{V}\setminus\mathbf{A}}}(\mathbf{Y})$	$\prod_{\mathbf{D}\in\mathcal{D}(\mathcal{G}_{\mathbf{Y}^\star})}\phi_{\mathbf{V}\setminus\mathbf{D}}(p(\mathbf{V});\mathcal{G})\vert_{\mathbf{A}=\mathbf{a}}$
SG	Yes	Node – $\mathbf{a}$	$\mathrm{ant}_{\mathcal{G}_{\mathbf{V}\setminus\mathbf{A}}}(\mathbf{Y})$	$\prod_{\mathbf{D}\in\mathcal{D}(\tilde{\mathcal{G}}^d)}\phi_{\mathbf{D}^\star\setminus\mathbf{D}}(q(\mathbf{D}^\star\vert\,\mathrm{pa}^s_{\mathcal{G}}(\mathbf{D}^\star));\mathcal{G}^d)\times$ $\prod_{\mathbf{B}\in\mathcal{B}(\tilde{\mathcal{G}}^b)} p(\mathbf{B}\setminus\mathbf{A}\vert\,\mathrm{pa}_{\mathcal{G}_{\mathbf{Y}^\star}}(\mathbf{B}),\mathbf{B}\cap\mathbf{A})\vert_{\mathbf{A}=\mathbf{a}}$
ADMG	Yes	Policy – $\mathbf{f_A}$	$\mathrm{an}_{\mathcal{G}_{\mathbf{f_A}}}(\mathbf{Y})$	$\prod_{\mathbf{D}\in\mathcal{D}(\mathcal{G}_{\mathbf{Y}^\star})}\phi_{\mathbf{V}\setminus\mathbf{D}}(p(\mathbf{V});\mathcal{G})\vert_{\mathbf{A}=\tilde{\mathbf{a}}}$

Table 1: Summary of existing identification approaches. The first two rows use standard g-formulas, the third row is the ID algorithm, and the final two extend ID. The present work generalizes the last two rows. In the fifth row, $\tilde{\mathbf{a}} = \{A = f_A(\mathbf{W}_A)\vert A\in\mathrm{pa}_{\mathcal{G}}(\mathbf{D})\cap\mathbf{A}\}$ if $\mathrm{pa}_{\mathcal{G}}(\mathbf{D})\cap\mathbf{A}\neq\emptyset$ and $\tilde{\mathbf{a}}=\emptyset$ otherwise.

Kernels and Fixing. Whereas DAGs and CGs factorize as products of conditional distributions, ADMGs and SGs factorize as products of *kernels* (Lauritzen, 1996). A kernel $q_{\mathbf{V}}(\mathbf{V}\vert\mathbf{W})$ is a function, mapping values of $\mathbf{W}$ to normalized densities on $\mathbf{V}$. For some $\mathbf{A}\subseteq\mathbf{V}$, conditioning and marginalization are defined as:

$$q(\mathbf{A}\vert\mathbf{W})\equiv\sum_{\mathbf{V}\setminus\mathbf{A}} q(\mathbf{V}\vert\mathbf{W});\quad q(\mathbf{V}\setminus\mathbf{A}\vert\mathbf{A},\mathbf{W})\equiv\frac{q(\mathbf{V}\vert\mathbf{W})}{q(\mathbf{A}\vert\mathbf{W})}.$$

The notion of fixing variables is closely tied to kernels. In a CADMG $\mathcal{G}(\mathbf{V},\mathbf{W})$, a variable V is fixable if $\mathrm{de}(V)\cap\mathrm{dis}(V)=\emptyset$. In a DAG $\mathcal{G}$ with corresponding distribution $p(\mathbf{V})$, fixing V corresponds to applying the g-formula to obtain a new distribution $p(\mathbf{V}\setminus V)$ and a new graph $\mathcal{G}'$. For a CADMG $\mathcal{G}(\mathbf{V},\mathbf{W})$ with corresponding kernel $q(\mathbf{V}\vert\mathbf{W})$, Richardson et al. (2017) defines similar operations, denoted $\phi_V(\mathcal{G})$ and $\phi_V(q;\mathcal{G})$. These operators yields a new CADMG $\mathcal{G}'(\mathbf{V}\setminus\{V\},\mathbf{W}\cup\{V\})$ in which all edges *into* V are removed and a new kernel $q'(\mathbf{V}\setminus\{V\}\vert\mathbf{W}\cup\{V\})\equiv\frac{q(\mathbf{V}\vert\mathbf{W})}{q(\mathbf{V}\vert\,\mathrm{pa}_{\mathcal{G}}(\mathrm{dis}_{\mathcal{G}}(V))\cup\mathrm{dis}_{\mathcal{G}}(V),\mathbf{W})}$, respectively. These operators were used to define the nested Markov model.

Fixability also extends to sets of variables $\mathbf{S}\subseteq\mathbf{V}$ in a ADMG $\mathcal{G}(\mathbf{V})$ when $q_{\mathbf{V}}$ is in the nested Markov model. If it is possible to find a sequence $S_1, S_2,\ldots$ of the variables in $\mathbf{S}$ such that S_1 is fixable in $\mathcal{G}$, S_2 is fixable in $\phi_{S_1}(q_{\mathbf{V}};\mathcal{G})$ and so on, then $\mathbf{S}$ is fixable and $\mathbf{V}\setminus\mathbf{S}$ is said to be *reachable* in $\mathcal{G}$. Since all valid fixing sequences on $\mathbf{S}$ yield the same CADMG $\mathcal{G}(\mathbf{V}\setminus\mathbf{S},\mathbf{S})$ via ϕ, and (if $p(\mathbf{V})$ is nested Markov with respect to $\mathcal{G}(\mathbf{V})$), all fixing sequences on $\mathbf{S}$ valid in $\mathcal{G}(\mathbf{V})$ yield the same kernel $q(\mathbf{V}\setminus\mathbf{S}\vert\mathbf{S})$ via ϕ, the fixing operators can be defined for sets unambiguously: $\phi_{\mathbf{S}}(\mathcal{G})$ and $\phi_{\mathbf{S}}(q;\mathcal{G})$.

This notation permits reformulating the ID algorithm. For an ADMG $\mathcal{G}(\mathbf{V})$, let $\mathbf{Y},\mathbf{A}\subseteq\mathbf{V}$ be disjoint and $\mathbf{Y}^\star\equiv\mathrm{an}_{\mathcal{G}_{\mathbf{V}\setminus\mathbf{A}}}(\mathbf{Y})$. $p(\mathbf{Y}\vert\mathrm{do}(\mathbf{a}))$ is identified in $\mathcal{G}$ if and only if every district $\mathbf{D}\in\mathcal{D}(\mathcal{G}_{\mathbf{Y}^\star})$ is reachable in $\mathcal{G}$. If identified, $p(\mathbf{Y}\vert\mathrm{do}(\mathbf{a}))$ is given by summing the modified factorization in row three of Table 1 over $\mathbf{Y}^\star\setminus\mathbf{Y}$.

Returning to our elections example (Fig. 1(c)), suppose we assume each candidate's decision is independent of other decisions given covariates (i.e., no $A_l - A_r$ edge). We can use this formula to consider the effect on a candidate's polling of advertising positively and negatively in fixed proportion (say, equally, $a=.5$).

As another example, consider the subgraph on C_1, A_1, M_1, Y_1 in Fig. 2(a); $p(Y_1\vert\mathrm{do}(a_1))$ is not identified (Shpitser and Pearl, 2006). In the C_2, A_2, M_2, Y_2 subgraph, however, $p(Y_2\vert\mathrm{do}(a_2))$ is identified by the front-door formula:

$$\sum_{M_2,C_2} p(M_2\vert a_2,C_2)p(C_2)\sum_{A_2'} p(Y_2\vert M_2,C_2,A_2')p(A_2'\vert C_2)$$

3.2 Identification in Segregated Graphs

The Segregated Factorization. Extending the factorizations for ADMGs and CGs, Sherman and Shpitser (2018) defines the segregated factorization for SGs.

Recall that an SG $\mathcal{G}$ is partitioned by variables that lie in non-trivial blocks, denoted $\mathbf{B}^\star=\cup_{\mathbf{B}\in\mathcal{B}^{nt}(\mathcal{G})}\mathbf{B}$, and those that don't, denoted $\mathbf{D}^\star=\cup_{\mathbf{D}\in\mathcal{D}(\mathcal{G})}\mathbf{D}$. An SG satisfying the segregated factorization can be expressed as the product of kernels for these two sets.

The first kernel, $q(\mathbf{B}^\star\vert\,\mathrm{pa}^s_{\mathcal{G}}(\mathbf{B}^\star)) = \prod_{\mathbf{B}\in\mathcal{B}^{nt}(\mathcal{G})} p(\mathbf{B}\vert\,\mathrm{pa}_{\mathcal{G}}(\mathbf{B}))$, factorizes with respect to a conditional chain graph (CCG) $\mathcal{G}(\mathbf{V},\mathbf{W})$, which we denote by $\mathcal{G}^b$ with $\mathbf{V}$ corresponding to $\mathbf{B}^\star$ and $\mathbf{W}$ to $\mathrm{pa}^s(\mathbf{B}^\star)$. $\mathcal{G}^b$ contains edges between nodes in $\mathbf{B}^\star$ and between nodes in $\mathrm{pa}^s_{\mathcal{G}}(\mathbf{B}^\star)$ that exist in $\mathcal{G}$.

The second kernel, $q(\mathbf{D}^\star\vert\,\mathrm{pa}^s_{\mathcal{G}}(\mathbf{D}^\star)) = \frac{p(\mathbf{V})}{q(\mathbf{B}^\star\vert\,\mathrm{pa}^s_{\mathcal{G}}(\mathbf{B}^\star))}$, nested factorizes with respect to a CADMG denoted $\mathcal{G}^d$, with random nodes $\mathbf{D}^\star$ and fixed nodes $\mathrm{pa}^s(\mathbf{D}^\star)$. Like $\mathcal{G}^b$, $\mathcal{G}^d$ contains edges between nodes in $\mathbf{D}^\star$ and between nodes in $\mathrm{pa}^s_{\mathcal{G}}(\mathbf{D}^\star)$ that are present in $\mathcal{G}$.

For example, in the graph in Fig. 2(a), we have

$$\begin{aligned} q(\mathbf{D}^\star\vert\,\mathrm{pa}^s_{\mathcal{G}}(\mathbf{D}^\star)) &= p(Y_2,Y_3,A_2\vert C_2,M_2,M_3)\\ &\times p(Y_1,A_1,C_1\vert M_1)p(A_3\vert C_3)\end{aligned}$$

$$q(\mathbf{B}^\star |\, \mathrm{pa}^s_{\mathcal{G}}(\mathbf{B}^\star)) = p(M_1, M_2, M_3|A_1, A_2, A_3)p(C_2, C_3)$$

which correspond to Fig. 2(b) and (c) respectively.

The Segregated Graph ID Algorithm. We can now describe an extension of the ID algorithm for node interventions in segregated graphs. For a SG $\mathcal{G}(\mathbf{V})$, fix disjoint $\mathbf{Y}, \mathbf{A} \subseteq \mathbf{V}$. Let $\mathbf{Y}^\star \equiv \mathrm{ant}_{\mathcal{G}_{\mathbf{V}\setminus\mathbf{A}}}(\mathbf{Y})$. Define $\tilde{\mathcal{G}}^d$ and $\tilde{\mathcal{G}}^b$ to be the CADMG and CCG respectively obtained from $\mathcal{G}_{\mathbf{Y}^\star}$. $p(\mathbf{Y}|\mathrm{do}(\mathbf{a}))$ is identified in $\mathcal{G}$ if and only if each $\mathbf{D} \in \mathcal{D}(\tilde{\mathcal{G}}^d)$ is reachable in $\mathcal{G}^d$. If identified, $p(\mathbf{Y}|\mathrm{do}(\mathbf{a}))$ is equal to the modified factorization in row four of Table 1, summed over $\mathbf{Y}^\star \setminus \mathbf{Y}$.

Coming back to our elections example, Fig. 1(b), this formula is applicable when considering the effect of the left-leaning candidate taking a fixed action a_l, with the right-leaning candidate's action still having an impact on the left's poll standing. $p(Y_l(a_l))$ is identified by:

$$\sum_{C_l, C_r, A_r, Y_r} p(Y_l, Y_r|C_l, C_r, A_r, a_l)p(A_2|C_l, C_r)p(C_l)p(C_r)$$

3.3 Policy Interventions in ADMGs

Extending node interventions, we now consider *policy interventions*. For an ADMG $\mathcal{G}(\mathbf{V})$ with topological ordering $\prec$ on $\mathbf{V}$ and an intervention set $\mathbf{A} \subseteq \mathbf{V}$, let $\mathbf{f_A}$ be the set of policies $\{f_A : A \in \mathbf{A}\}$. Each f_A is a stochastic function of some $\mathbf{W}_A \subseteq \mathbf{V}_{\prec A}$, where $f_A(\mathbf{W}_A)$ maps the state space of $\mathbf{W}_A$ to the state space of A. Intervening with f_A corresponds to removing edges *into* A in $\mathcal{G}$ and adding edges from $\mathbf{W}_A$ to A, yielding a new graph $\mathcal{G}_{\mathbf{f_A}}$.

Tian (2008) gave a policy-analogue of the ID algorithm for $p(\{\mathbf{V} \setminus \mathbf{A}\}(\mathbf{f_A}))$, which Shpitser and Sherman (2018) re-expressed via the fixing operator ϕ. Let $\mathbf{Y}^\star \equiv \mathrm{an}_{\mathcal{G}_{\mathbf{f_A}}}(\mathbf{Y}) \setminus \mathbf{A}$. A policy-analogue of the ID algorithm follows: $p(\mathbf{Y}(\mathbf{f_A}))$ is identified in $\mathcal{G}$ if and only if $p(\mathbf{Y}^\star(\mathbf{a}))$ is identified in $\mathcal{G}$; if identified, $p(\mathbf{Y}(\mathbf{f_A}))$ is obtained by summing over $(\mathbf{Y}^\star \cup \mathbf{A}) \setminus \mathbf{Y}$ in the modified factorization in row five of Table 1.

In our elections example, assume candidates' decisions and outcomes are independent of each other. This formula can be used to consider the effect on a candidate's polling of advertising based on the relevant covariates, e.g., if the election is less than 2 months away, advertise negatively, and buy positive ads until then.

4 Varieties of Policy Interventions

We now describe extensions of policy interventions to network data representable by SGs. These interventions correspond to replacing structural equations in Procedure 1 with new equations, under conditions we describe below such that the resulting data generating process yields a new SG. As we discuss, these policy interventions induce a variety of edge changes in SGs.

4.1 Inducing Direct Causation

As in the latent-variable DAG case (Shpitser and Sherman, 2018), we can intervene by inducing a parent-child relationship between the treatment node and other variables in the graph or modify the nature of existing relationship. In our elections example from Sec. 1, this might correspond to intervening on the left candidate's decision A_l such that she adopts a new strategy for responding to her competitor's characteristics C_r relative to her (observed) status quo strategy. For illustrative purposes, this type of intervention is demonstrated by the addition of the $C_2 \to A_1$ edge and the modification to the $C_1 \to A_1$ edge between Fig. 2(a) and 2(d).

4.2 Inducing or Modifying Undirected Dependence

We can also consider changing the block structure of the SG. There are two types of such interventions:

1. *Modifying the functional form encoded by an existing undirected edge.* In Fig. 1 (b), we can think of the undirected edge $A_l - A_r$ as representing each candidates' beliefs about the other candidate's actions. In the observed data, candidates will best-respond to each other according to these beliefs. We can imagine changing the way one (or both) of the candidates reasons about their opponent's possible actions, such as making one candidate hyper-responsive to their opponent's anticipated action. Mechanically, we intervene on A_l (analogously A_r) with a function f_{A_l} that takes A_r as an argument. We needn't intervene on the other candidate to maintain the undirected edge between the A's. This type of intervention is demonstrated by the change to the $M_2 - M_3$ edge from Fig. 2(a) to 2(d).

2. *Inducing co-dependence by adding a new undirected edge between two nodes.* This might correspond to having a third candidate c join the race and intervening such that $A_c - A_l$ and $A_c - A_r$. In this case, it is necessary to intervene on both endpoint nodes for the new undirected edge in order; we modify the respective structural equations to take the other endpoint as an argument. We further restrict these interventions by requiring that they do not induce a partially directed cycle, which would violate the segregation property of the graph. We formalize this requirement below. We note that this type of intervention can be thought of as a chain graph generalization of connection interventions, proposed in Sherman and Shpitser (2019). As an example, consider the addition of the $A_2 - A_3$ edge

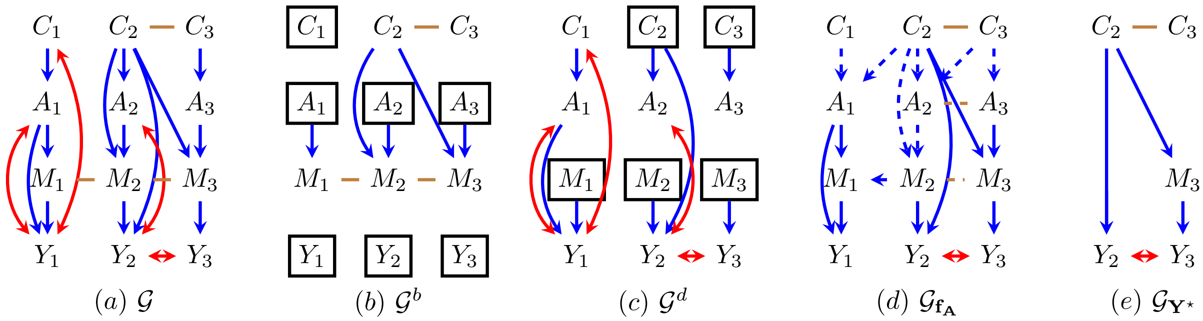

Figure 2: (2(a) An SG $\mathcal{G}$ where bi-directed edges signify the presence of latent confounders. 2(b) and (c) The conditional chain graph $\mathcal{G}^b$ and conditional ADMG $\mathcal{G}^d$ obtained from $\mathcal{G}$. 2(d) The post-intervention graph $\mathcal{G}_{\mathbf{f_A}}$ induced by the policy intervention $\mathbf{f_A}$ as described in Sec. 4. Nodes with changed structural equations have dashed incoming edges. 2(e) The corresponding $\mathcal{G}_{\mathbf{Y}^\star}$ for $\mathcal{G}_{\mathbf{f_A}}$ in 2(b) with outcome $\mathbf{Y} = \{Y_2, Y_3\}$.

in Fig. 2(d) relative to 2(a).

4.3 Removing Dependence

Finally, we can consider removing undirected dependence between nodes. Once again there are two types:

1. *Partial removal.* We intervene on a single node to make its structural equation no longer a function of the other end point of the undirected edge. In our elections example (Fig. 1(c)), this corresponds to a 'first mover' scenario where A_l is made to not depend on A_r and thus candidate l makes her decision before candidate r. Graphically, we change the undirected edge $A_l - A_r$ to a directed edge $A_l \to A_r$ since A_r is still determined by candidate l's decision; see, for instance, the $M_1 - M_2$ and $M_1 \leftarrow M_2$ edges in Fig. 2(a) and 2(d).

2. *Complete removal.* We remove both dependences by intervening on both endpoints of an undirected edge so that the structural equations are no longer functions of each other. This corresponds to a candidate dropping out of the race in our elections example. Like dependence-inducing interventions above, this intervention type can be viewed as an SG analogue of severance interventions (Sherman and Shpitser, 2019).

5 Identification of Policies in Segregated Graphs

In this section we formalize policy interventions and provide a procedure for obtaining the post-intervention graph from $\mathcal{G}$. We then give a criterion for the identification of policy interventions in SGs (Shpitser, 2015) and demonstrate application of this criterion to Fig. 2 and to our electoral example, Fig. 1. We defer proofs and derivations to the supplement.

Procedure 2 Obtaining $\mathcal{G}_{\mathbf{f_A}}$ from $\mathcal{G}$

procedure InterveneGraph($\mathcal{G}, \mathbf{f_A}(\mathbf{Z_A})$)
  Initialize $\mathcal{G}_{\mathbf{f_A}} \leftarrow \mathcal{G}$
  for each $A \in \mathbf{A}$ **do**
    Replace all $V - A$ with $A \to V$ in $\mathcal{G}_{\mathbf{f_A}}$
    Remove all $\cdot \to A$, $\cdot \leftrightarrow A$ from $\mathcal{G}_{\mathbf{f_A}}$
    Add edges $\mathbf{Z}_A \to A$ in $\mathcal{G}_{\mathbf{f_A}}$
  for each $V_i, V_j \in \mathbf{V}$ **do**
    if $V_i \to V_j$ and $V_j \to V_i$ in $\mathcal{G}_{\mathbf{f_A}}$ **then**
      Remove $V_i \to V_j$ and $V_j \to V_i$ from $\mathcal{G}_{\mathbf{f_A}}$
      Add $V_i - V_j$ in $\mathcal{G}_{\mathbf{f_A}}$
  return $\mathcal{G}_{\mathbf{f_A}}$

5.1 Formalizing Policy Interventions in Segregated Graphs

Before providing identification conditions, we first formally define policy interventions in SGs. Recall that in ADMGs a policy $f_A(\mathbf{W}_A) \in \mathbf{f_A}$ was required to be a function of variables $\mathbf{W}_A$ preceding A in a topological ordering on the nodes in $\mathcal{G}$. In SGs we loosen this restriction such that f_A operates as a structural equation that can also be a stochastic function of variables in the same block as A. For an intervention inducing a block or modifying the structural equations in a block, we use Procedure 1 to obtain a new block distribution.

For $f_A(\mathbf{Z}_A)$ to be a valid policy in an SG $\mathcal{G}(\mathbf{V})$, we require $\mathbf{Z}_A \subseteq \mathbf{V} \setminus \overline{\text{ext}}(A)$. In turn, for $\mathbf{f_A}$ to be valid, all constituent policies must be valid and they may not collectively violate the CG property by inducing a partially directed cycle. We formalize this notion as follows: let $A_i \triangle A_j$ denote that variable A_i is made (either directly or indirectly) a function of A_j for $A_i, A_j \in \mathbf{A}$. To prevent partially directed cycles, we stipulate that if $A_i \triangle A_j$ and $A_j \triangle A_i$ then we require $A_i \in \mathbf{Z}_{A_j}$ and vice versa. This motivates the following definition.

Definition 1 *A policy intervention* $\mathbf{f_A}(\mathbf{Z_A})$ *is 'segregation preserving' if (a) for each* $A \in \mathbf{A}$, $\mathbf{Z}_A \subseteq \mathbf{V} \setminus \overline{\text{ext}}(A)$, *and (b) for any* $A_i, A_j \in \mathbf{A}$ *if* $A_i \triangle A_j$ *and* $A_j \triangle A_i$, *we have that* $A_i \in \mathbf{Z}_{A_j}$ *and* $A_j \in \mathbf{Z}_{A_i}$.

For a given intervention set $\mathbf{f_A}$, we can construct a post-intervention graph $\mathcal{G}_{\mathbf{f_A}}$ according to Procedure 2, which follows from the analogous procedure for policy identification in LV-DAGs. In Lemma 1, we show that $\mathcal{G}_{\mathbf{f_A}}$ is an SG when $\mathbf{f_A}$ is segregation-preserving. As an example of this procedure's application, consider Fig. 2(a). Suppose we wish to perform an intervention $\mathbf{f_A}(\mathbf{Z_A})$ as in Table 2. Then $\mathcal{G}_{\mathbf{f_A}}$ is given by Fig. 2.

$A \in \mathbf{A}$	A_1	A_2	A_3	M_2
$\mathbf{Z}_A$	C_2	C_2, C_3, A_3	A_2, C_3	A_2, C_2, M_3

Table 2: Intervention variables $A \in \mathbf{A}$ and induced dependences $\mathbf{Z}_A$ for the intervention in Fig. 2

5.2 Identification Results

First, we show that the post-intervention $\mathcal{G}_{\mathbf{f_A}}$ is an SG.

Lemma 1 *Given an SG* $\mathcal{G}(\mathbf{V})$ *and a segregation-preserving intervention* $\mathbf{f_A}(\mathbf{Z_A})$, *the post-intervention graph* $\mathcal{G}_{\mathbf{f_A}}$ *obtained via Procedure 2 is an SG.*

We now present the main result of this paper. This theorem provides sufficient conditions for the identification of the effects of policy interventions in SGs.

Theorem 1 *Let* $\mathcal{G}(\mathbf{V} \cup \mathbf{H})$ *be a causal LV-CG with* $\mathbf{H}$ *block-safe, and a topological order* $\prec$. *Fix disjoint* $\mathbf{Y}, \mathbf{A} \subseteq \mathbf{V}$. *Let* $\mathbf{f_A}(\mathbf{Z_A})$ *be a segregation preserving policy set. Let* $\mathbf{Y}^\star \equiv \text{ant}_{\mathcal{G}_{\mathbf{f_A}}}(\mathbf{Y}) \setminus \mathbf{A}$. *Let* $\mathcal{G}^d, \tilde{G}^d$ *be the induced CADMGs on* $\mathcal{G}_{\mathbf{f_A}}$ *and* $\mathcal{G}_{\mathbf{Y}^\star}$, *and* $\tilde{G}^b$ *the induced CCG on* $\mathcal{G}_{\mathbf{Y}^\star}$. *Let* $q(\mathbf{D}^\star | \text{pa}^s_{\mathcal{G}_{\mathbf{f_A}}}(\mathbf{D}^\star)) = \prod_{\mathbf{D} \in \mathcal{G}_{\mathbf{f_A}}} q(\mathbf{D} | \text{pa}^s_{\mathcal{G}_{\mathbf{f_A}}}(\mathbf{D}))$, *where* $q(\mathbf{D} | \text{pa}^s_{\mathcal{G}_{\mathbf{f_A}}}(\mathbf{D})) = \prod_{D \in \mathbf{D}} p(D | \mathbf{V}_{\prec D})$ *if* $\mathbf{D} \cap \mathbf{A} = \emptyset$ *and* $q = f_A(\mathbf{Z}_A)$ *if* $\mathbf{D} \cap \mathbf{A} \neq \emptyset$. $p(\mathbf{Y}(\mathbf{f_A}(\mathbf{Z_A})))$ *is identified in* $\mathcal{G}$ *if and only if* $p(\mathbf{Y}^\star(\mathbf{a}))$ *is identified in* $\mathcal{G}$ *for the unrestricted class of policies. If identified,* $p(\mathbf{Y}(\mathbf{f_A}(\mathbf{Z_A}))) =$

$$\sum_{\{\mathbf{Y}^\star \cup \mathbf{A}\} \setminus \mathbf{Y}} \Bigg[\prod_{\mathbf{B} \in \mathcal{B}(\tilde{\mathcal{G}}^b)} p^\star(\mathbf{B} | \text{pa}_{\mathcal{G}_{\mathbf{f_A}}}(\mathbf{B})) \Bigg] \times \Bigg[\prod_{\mathbf{D} \in \mathcal{D}(\tilde{\mathcal{G}}^d)} \phi_{\mathbf{D}^\star \setminus \mathbf{D}}(q(\mathbf{D}^\star | \text{pa}^s_{\mathcal{G}_{\mathbf{f_A}}}(\mathbf{D}^\star)); \mathcal{G}^d) \Bigg] \Bigg|_{\mathbf{A} = \tilde{\mathbf{a}}} \tag{2}$$

where (a) $\tilde{\mathbf{a}} = \{A = f_A(\mathbf{Z}_A) : A \in \text{pa}_{\mathcal{G}_{\mathbf{f_A}}}(\mathbf{D}) \cap \mathbf{A}\}$ *if* $\text{pa}_{\mathcal{G}_{\mathbf{f_A}}}(\mathbf{D}) \cap \mathbf{A} \neq \emptyset$ *and* $\tilde{\mathbf{a}}_\mathbf{D} = \emptyset$ *otherwise, and (b)* $p^\star$ *is obtained by running Procedure 1 over functions* $g_{B_i}(B_{-i}, \text{pa}_{\mathcal{G}_{\mathbf{f_A}}}(B_i), \epsilon_{B_i})$ *where* $g_{B_i} \in \mathbf{f_A}$ *if* $B_i \in \mathbf{A}$ *and* g_{B_i} *is given by the observed distribution if* $B_i \notin \mathbf{A}$[1].

The outer sum over $\mathbf{A}$ is extraneous if $\mathbf{f_A}$ corresponds to a set of deterministic policies.

5.3 Estimands and Optimal Policy Selection

We now demonstrate how to obtain identified functionals via Eq. 2. We describe identification of the effect on $\{Y_2, Y_3\}$ in Fig. 2(a) of the intervention in Table 2, and then give the functional for our elections example, Fig. 1(b), which we estimate in the next section.

From Fig. 2(a), we obtain $\mathcal{G}_{\mathbf{f_A}}$ in Fig. 2(d) by applying the intervention detailed in Table 2. In turn, from this post-intervention graph we observe that $\mathbf{Y}^\star = \text{ant}_{\mathcal{G}_{\mathbf{f_A}}}(\mathbf{Y}) \setminus \mathbf{A} = \{C_2, C_3, M_3, Y_2, Y_3\}$ and obtain the induced subgraph $\mathcal{G}_{\mathbf{Y}^\star}$ in Fig. 2(e).

$\mathcal{G}_{\mathbf{Y}^\star}$ factorizes into kernels relating to district nodes and block nodes: $q_{\mathcal{D}}(C_1, A_1, M_1, Y_1, Y_2, Y_3 | C_2, M_2, M_3)$ and $q_{\mathcal{B}}(M_2, M_3, A_2, A_3, C_2, C_3 | \emptyset)$. The block nodes factorize as a product of blocks, as in the first term of Eq. 2. Separately, we must fix sets for each $\mathcal{G}_{\mathbf{Y}^\star}$ district $\{\{M_3\}, \{Y_2, Y_3\}\}$ in $q_{\mathcal{D}}$. This yields the functional (full derivation in the supplement) for $p(\{Y_2, Y_3\}(\mathbf{f_A}))$:

$$\sum_{\{A_1, A_2, A_3, M_2, M_3, C_2, C_3\}} p^\star(A_2, A_3 | C_2, C_3) p^\star(M_2, M_3 | A_2, A_3, C_2) \times p(Y_2, Y_3 | Y_1, A_1, M_1, M_3, C_1, C_2) p^\star(C_2, C_3)$$

Similarly, we consider the effect on Y_l of intervening with a policy $f_{A_l}(C_l)$ in our electoral example, Fig. 1 (b). $f_{A_l}(C_l)$ corresponds to a myopic strategy in which the candidate makes decisions based only on their own covariates. Applying Eq. 2, $p(Y_l(f_{A_l}(C_l, C_r))) =$

$$\sum_{C_l, C_r, A_r, Y_r} p(A_r | C_l, C_r) p(C_l) p(C_r) \times p(Y_l, Y_r | C_l, C_r, A_r, f_{A_l}(C_l, C_r)) \tag{3}$$

To choose an optimal action for the left candidate, we select $f_{A_l}(C_l)$ from a set of candidate policies $\mathcal{F}_{A_l}(C_l)$:

$$f_{A_l}(C_l, C_r) = \underset{\tilde{f}_{A_l}(C_l, C_r) \in \mathcal{F}_{A_l}(C_l, C_r)}{\arg\max} p(Y_l(\tilde{f}_{A_l}(C_l, C_r)))$$

6 Estimation

We now demonstrate how functionals identified by Eq. 2 can be estimated from observed data. Specifically, we seek optimal $f_A(\mathbf{C})$'s for versions of the functional in Eq. 3. To do so, we fit nuisance models and utilize the plug-in principle to perform indirect Q-learning for

[1]This distribution is identified from univariate terms but it cannot be obtained in closed-form.

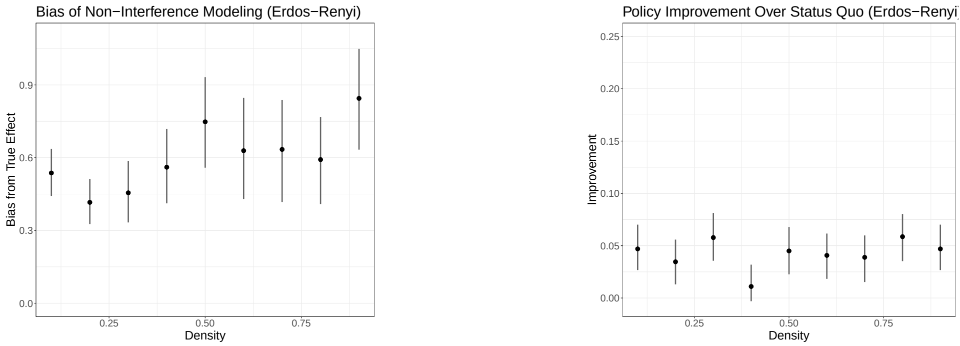

(a) (b)

Figure 3: (3a) Bias of estimates obtained using single-unit modeling, ignoring interference. The presence of bias suggests ignoring interference is highly problematic. (3b) Difference in expected outcomes between an optimized strategy and the status quo. We analyze several network densities to demonstrate the generality of this approach.

policy optimization. This approach yields consistent estimates of the optimized outcome under regularity conditions, assuming correctly specified nuisance models (Chakraborty and Moodie, 2013).

For our experiments we first generate 10-node network graphs according to one of three widely-used network generators: Erdős and Rényi (1960), Watts and Strogatz (1998), and Albert and Barabási (2002). In-unit and cross-unit structures are identical to the 2-node graph in Fig. 1(b). We then generate data for each C, A, and Y using log-linear models, with $C \in [0,1]^3$ and $A, Y \in [0,1]$. We use Gibbs sampling to approximate undirected edges between Y's (Tchetgen et al., 2017). We defer parametric specifications of our data generating process to the supplement. We assume *partial* interference: we generate 1000 samples of each network topology and use these to fit nuisance models. We run the following experiments by obtaining $1,000$ bootstrap replications of the generated data and calculating a 95% confidence interval of the relevant effect:

1. *Bias from incorrectly assuming iid.* As a demonstration of the importance of using interference-aware modeling, we consider performing *node* interventions on each A_i obtained from our Erdős-Rényi samples, setting A_i to 1 and 0. We estimate the average causal effect (ACE) of these node interventions ($E[Y_i(1) - Y_i(0)]$) using models implied by ID (Table 1, row three), which provides sound functionals when data *are* iid, as well as models implied by the SG ID algorithm (Table 1, row four) which respect the dependent nature of the data. We treat the latter models as 'ground truth' and calculate the bias of the ACE induced by inappropriately assuming data are iid. These results are given in Fig. 3a. Observing that bias is universally bounded away from 0 in these results, it's clear that it's imperative to respect network dependence in causal modeling.

2. *Benefit of optimizing interventions.* Here we demonstrate the efficacy of policy interventions for picking tailored interventions that optimize a subject's outcome, by estimating the 10-unit version of the identified functional in Eq. 3. From our generated samples, we fit logistic regression models for $E[Y_{-i}|\mathbf{A},\mathbf{C}]$ and $E[Y_i|\mathbf{A},\mathbf{C},Y_{-i}]$, where i denotes the unit we wish to optimize for. This ensures the necessary consistency properties for indirect Q-learning. Models for $p(\mathbf{A},\mathbf{C})$ are estimated using the empirical distribution.

For each sample we estimate the effect of intervening with a policy $f_{A_i}(C_i) \in \mathcal{F}_{A_i}(C_i) = \{|C_i|^{-1}\sum_{j\in[|C_i|]} k_j C_{ij} : k_j \in \mathbb{R}\}$ (i.e. $\mathcal{F}_{A_i}$ is the set of means of linear combinations of C_i's components). We choose $\mathbf{k}$ to maximize Y_i subject to the constraint that values of A_i and Y_i must remain in $[0,1]$. We report the *difference* between the optimized and observed ('status quo') Y_i's. The results for the Erdős-Rényi generator can be found in Fig. 3b. Results for the other generators can be found in the supplementary material. Since Y is binary, an expected difference of .05 corresponds to a 5.0% increase in Y over the status quo. Fig. 3b demonstrates that the proposed approach virtually guarantees an improved outcome over the status quo.

7 Conclusion

In this paper we discussed identification of policy intervention effects in the interference setting. We characterized interpretations of possible interventions and gave criteria for identifying their effects in latent-variable causal chain graph models. Further, we demonstrated estimation via a simulation study. Future directions include exploring the intersection of policies and interference, and game theory, and developing robust estimation strategies for this setting.

Acknowledgements

The first author would like to thank the Adobe Research Internship program, and Sridhar Mahadevan, for supporting this work. The third author would like to thank the following organizations for supporting this work: the American Institute of Mathematics SQuaRE program, National Institutes of Health grant R01 AI127271-01A1, Office of Naval Research grant N00014-18-1-2760, and Defense Advanced Research Progress Administration grant under contract HR0011-18-C-0049. The content of the information in this paper does not necessarily reflect the position or the policy of the Government, and no official endorsement should be inferred.

References

R. Albert and A. L. Barabási. Statistical mechanics of complex networks. *Reviews of modern physics*, 74 (1):47, 2002.

M. Blackwell. A framework for dynamic causal inference in political science. *American Journal of Political Science*, 57(2):504–520, 2013.

B. Chakraborty and E. Moodie. *Statistical methods for dynamic treatment regimes*. Springer, 2013.

D. R. Cox. *Planning of experiments*, volume 20. Wiley New York, 1958.

P. Erdős and A. Rényi. On the evolution of random graphs. *Publ. Math. Inst. Hung. Acad. Sci*, 5(1): 17–60, 1960.

M. G. Hudgens and M. E. Halloran. Toward causal inference with interference. *Journal of the American Statistical Association*, 103(482):832–842, 2008.

E. B. Laber, D. J. Lizotte, M. Qian, W. E. Pelham, and S. A. Murphy. Dynamic treatment regimes: Technical challenges and applications. *Electronic journal of statistics*, 8(1):1225, 2014.

S. L. Lauritzen. *Graphical models*, volume 17. Clarendon Press, 1996.

S. L. Lauritzen and T. S. Richardson. Chain graph models and their causal interpretations. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 64(3):321–348, 2002.

M. Maier, K. Marazopoulou, and D. Jensen. Reasoning about independence in probabilistic models of relational data. *arXiv preprint arXiv:1302.4381*, 2013.

R. Nabi, P. Kanki, and I. Shpitser. Estimation of personalized effects associated with causal pathways. In *Proceedings of the 34th Conference on Uncertainty in Artificial Intelligence*, 2018.

E. L. Ogburn, T. J. VanderWeele, et al. Causal diagrams for interference. *Statistical science*, 29(4): 559–578, 2014.

E. L. Ogburn, I. Shpitser, and Y. Lee. Causal inference, social networks, and chain graphs. *arXiv preprint arXiv:1812.04990*, 2018.

J. Pearl. *Causality: models, reasoning and inference*, volume 29. Springer, 2000.

J. M. Peña. Unifying gaussian LWF and AMP chain graphs to model interference. *arXiv preprint arXiv:1811.04477*, 2018.

T. S. Richardson, R. J. Evans, J. M. Robins, and I. Shpitser. Nested markov properties for acyclic directed mixed graphs. *arXiv preprint arXiv:1701.06686*, 2017.

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

E. Sherman and I. Shpitser. Identification and estimation of causal effects from dependent data. In *Advances in Neural Information Processing Systems*, pages 9424–9435, 2018.

E. Sherman and I. Shpitser. Intervening on network ties. In *Proceedings of the 35th Conference on Uncertainty in Artificial Intelligence*. AUAI Press, 2019.

I. Shpitser. Segregated graphs and marginals of chain graph models. In *Advances in Neural Information Processing Systems*, pages 1720–1728, 2015.

I. Shpitser and J. Pearl. Identification of joint interventional distributions in recursive semi-markovian causal models. In *Proceedings of the National Conference on Artificial Intelligence*, volume 21, page 1219. Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press; 1999, 2006.

I. Shpitser and E. Sherman. Identification of personalized effects associated with causal pathways. In *Proceedings of the 34th Conference on Uncertainty in Artificial Intelligence*, 2018.

E. J. T. Tchetgen, I. Fulcher, and I. Shpitser. Auto-g-computation of causal effects on a network. *arXiv preprint arXiv:1709.01577*, 2017.

J. Tian. Identifying dynamic sequential plans. *Proceedings of the 24th Conference on Uncertainty in Artificial Intelligence*, 2008.

J. Tian and J. Pearl. A general identification condition for causal effects. In *Proceedings of the National Conference on Artificial Intelligence*, pages 567–573, 2002.

T. Verma and J. Pearl. *Equivalence and synthesis of causal models*. UCLA, Computer Science Department, 1991.

D. Viviano. Policy targeting under network interference. *arXiv preprint arXiv:1906.10258*, 2019.

D. J. Watts and S. H. Strogatz. Collective dynamics of 'small-world' networks. *Nature*, 393(6684):440, 1998.