

Supplementary Materials For: Identification of Personalized Effects Associated With Causal Pathways

Ilya Shpitser

Department of Computer Science
Johns Hopkins University
Baltimore, MD

Eli Sherman

Department of Computer Science
Johns Hopkins University
Baltimore, MD

APPENDIX

A: The ID Algorithm

function **ID**($\mathbf{y}, \mathbf{x}, P, G$):

INPUT: \mathbf{x}, \mathbf{y} value assignments, P a probability distribution, G a causal diagram.

OUTPUT: Expression for $P_{\mathbf{x}}(\mathbf{y})$ in terms of P or **FAIL**(F, F').

- 1) if $\mathbf{x} = \emptyset$, return $\sum_{\mathbf{v} \setminus \mathbf{y}} P(\mathbf{v})$.
- 2) if $\mathbf{V} \setminus An(\mathbf{Y})_G \neq \emptyset$,
return **ID**($\mathbf{y}, \mathbf{x} \cap An(\mathbf{Y})_G, \sum_{\mathbf{v} \setminus An(\mathbf{Y})_G} P, An(\mathbf{Y})_G$).
- 3) let $\mathbf{W} = (\mathbf{V} \setminus \mathbf{X}) \setminus An(\mathbf{Y})_{G_{\bar{\mathbf{x}}}}$.
if $\mathbf{W} \neq \emptyset$, return **ID**($\mathbf{y}, \mathbf{x} \cup \mathbf{w}, P, G$).
- 4) if $C(G \setminus \mathbf{X}) = \{S_1, \dots, S_k\}$,
return $\sum_{\mathbf{v} \setminus (\mathbf{y} \cup \mathbf{x})} \prod_i \mathbf{ID}(s_i, \mathbf{v} \setminus s_i, P, G)$.
if $C(G \setminus \mathbf{X}) = \{S\}$:
 - 5) if $C(G) = \{G\}$, throw **FAIL**($G, G \cap S$).
 - 6) if $S \in C(G)$,
return $\sum_{s \setminus \mathbf{y}} \prod_{\{i | V_i \in S\}} P(v_i | v_G^{(i-1)})$.
 - 7) if $(\exists S') S \subset S' \in C(G)$,
return **ID**($\mathbf{y}, \mathbf{x} \cap S'$,
 $\prod_{\{i | V_i \in S'\}} P(V_i | V_G^{(i-1)} \cap S', v_G^{(i-1)} \setminus S'), S'$).

Figure 1: ID Algorithm as it appears in [2].

B: Example Derivation For A Response To An Edge-Specific Policy

We seek to identify the distribution $p(Y(f_A^{(AM)} \rightarrow (W_0), f_A^{(AW_1)} \rightarrow (W_0)))$ in Fig. 2 (b). $\mathbf{Y}^* = \{Y, W_1, M_1, W_0\}$, and $\mathcal{D}(\mathcal{G}_{\mathbf{Y}^*}) = \{\{Y\}, \{W_0, M_1\}, \{W_1\}\}$ (the graph $\mathcal{G}_{\mathbf{Y}^*}$ is shown in Fig. 2 (c)). Thus, we have three terms, a term

$\phi_{\{W_0, M_1, A, W_1\}}(p; \mathcal{G})$ for Y , a term $\phi_{\{W_0, A, M_1, Y\}}(p; \mathcal{G})$ for W_1 , and a term $\phi_{\{A, W_1, Y\}}(p; \mathcal{G})$ for $\{W_0, M_1\}$. We have

$$\begin{aligned} \phi_{\{W_0, A, M_1, Y\}}(p; \mathcal{G}) &= \phi_{\{W_0, A, M_1\}} \left(\sum_Y p; \mathcal{G}^{(a)} \right) \\ &= \phi_{\{W_0, A\}} \left(\frac{p(W_0, A, M_1, W_1)}{p(M_1 | A, W_0)}; \mathcal{G}^{(b)} \right) \\ &= \phi_{\{W_0\}} \left(\frac{p(W_0, A, M_1, W_1)}{p(M_1, A | W_0)}; \mathcal{G}^{(c)} \right) \\ &= p(W_1 | M_1, A, W_0), \end{aligned}$$

where $\mathcal{G}^{(a)}, \mathcal{G}^{(b)}, \mathcal{G}^{(c)}$ are CADMGs in Figs. 2 (a), (b), and (c), respectively. Similarly, $\phi_{\{W_0, M_1, A, W_1\}}(p; \mathcal{G})$ is equal to

$$\begin{aligned} \phi_{\{W_0, M_1, A\}} \left(\frac{p(W_0, A, M_1, W_1, Y)}{p(W_1 | M_1, A, W_0)}; \mathcal{G}^{(d)} \right) \\ &= \phi_{\{W_0, A\}} \left(\frac{p(W_0, A, M_1, W_1, Y)}{p(W_1, M_1 | A, W_0)}; \mathcal{G}^{(e)} \right) \\ &= \phi_{\{W_0\}} \left(\sum_A \frac{p(W_0, A, M_1, W_1, Y)}{p(W_1, M_1 | A, W_0)}; \mathcal{G}^{(f)} \right) \\ &= \sum_{W_0, A} p(W_2 | W_1, M_1, A, W_0) p(A, W_0), \end{aligned}$$

where $\mathcal{G}^{(d)}, \mathcal{G}^{(e)}, \mathcal{G}^{(f)}$ are CADMGs in Figs. 2 (d), (e), and (f), respectively. Finally,

$$\begin{aligned} \phi_{\{A, W_1, Y\}}(p; \mathcal{G}) &= \phi_{\{A, W_1\}} \left(\sum_Y p; \mathcal{G}^{(a)} \right) \\ &= \phi_{\{A\}} \left(\sum_{Y, W_1} p; \mathcal{G}^{(g)} \right) \\ &= \frac{p(W_0, A, M_1)}{p(A | W_0)} = p(M_1 | A, W_0) p(W_0), \end{aligned}$$

where $\mathcal{G}^{(a)}, \mathcal{G}^{(g)}$ are CADMGs in Figs. 2 (a), and (g), respectively. Note that whenever the fixing operation for a kernel $q_{\mathbf{V}}(\mathbf{V} | \mathbf{W})$ that fixes $V \in \mathbf{V}$ is such that $\mathbf{V} \setminus \{V\} \subseteq \text{nd}_{\mathcal{G}(\mathbf{V}, \mathbf{W})}(V)$, the resulting kernel can be viewed as $\tilde{q}_{\mathbf{V} \setminus \{V\}}(\mathbf{V} \setminus \{V\} | \mathbf{W} \cup$

$\{V\}) = \sum_V q_{\mathbf{V}}(\mathbf{V}|\mathbf{W})$. We now combine these terms, evaluating A to either $f_A^{(AW_1)\rightarrow}(W_0)$ or $f_A^{(AM)\rightarrow}(W_0)$, as appropriate, yielding the functional in (18) for $p(Y(f_A^{(AW_1)\rightarrow}(W_0), f_A^{(AM)\rightarrow}(W_0)))$, namely:

$$\begin{aligned} & \sum_{W_0, A, M, W_1} \left[p(W_1|M, A = f_A^{(AM)\rightarrow}(W_0), W_0) \right] \\ & \times [p(M|A = f_A^{(AW_1)\rightarrow}(W_0), W_0)p(W_0)] \\ & \times \left[\sum_{W_0, A} p(Y|W_1, M, A, W_0)p(W_0, A) \right]. \end{aligned}$$

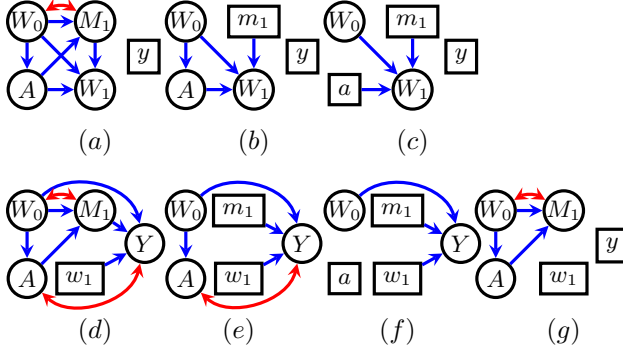


Figure 2: CADMGs obtained from fixing in \mathcal{G} shown in Fig. 2 (b): (a) $\phi_{\{Y\}}(\mathcal{G})$, (b) $\phi_{\{Y, M_1\}}(\mathcal{G})$, (c) $\phi_{\{Y, M_1, A\}}(\mathcal{G})$, (d) $\phi_{\{W_1\}}(\mathcal{G})$, (e) $\phi_{\{W_1, M_1\}}(\mathcal{G})$, (f) $\phi_{\{W_1, M_1, A\}}(\mathcal{G})$, (g) $\phi_{\{Y, W_1\}}(\mathcal{G})$.

C: Proofs

Before giving proofs of our main results, we state the following utility lemma which will be useful throughout subsequent developments.

Lemma 1 *Let \mathcal{G} be a DAG with vertex set \mathbf{V} . Fix $A, B \in \mathbf{V}$ such that $B \notin \text{de}_{\mathcal{G}}(A) \cup \text{ch}_{\mathcal{G}}(A)$ and $A \notin \text{de}_{\mathcal{G}}(B) \cup \text{ch}_{\mathcal{G}}(B)$. Let $\mathcal{G}_{A \times B}$ be a directed graph containing vertices $(\mathbf{V} \setminus \{A, B\}) \cup Z$, and the following set of edges. First, all edges between vertices in $\mathbf{V} \setminus \{A, B\}$ in \mathcal{G} also are in $\mathcal{G}_{A \times B}$. Second, for every $C \neq A, B$, for every edge of the form $C \rightarrow A$ or $C \rightarrow B$ in \mathcal{G} , there is an edge $C \rightarrow Z$ in $\mathcal{G}_{A \times B}$, and for every edge of the form $A \rightarrow C$ or $B \rightarrow C$ in \mathcal{G} , there is an edge $Z \rightarrow C$ in $\mathcal{G}_{A \times B}$. Then*

- (a) $\mathcal{G}_{A \times B}$ is a DAG.
- (b) Any element in the causal model for \mathcal{G} is an element of the causal model for $\mathcal{G}_{A \times B}$, if we interpret the Cartesian product of variables A and B in this element as the variable Z .

Proof: If $\mathcal{G}_{A \times B}$ is not a DAG, there is a directed cycle involving Z , e.g. $W \rightarrow \dots \rightarrow \circ \rightarrow Z \rightarrow \circ \rightarrow \dots \rightarrow W$. Since \mathcal{G} is a DAG, this implies either \mathcal{G} has a pair of paths $W \rightarrow \dots \rightarrow \circ \rightarrow A$ and $B \rightarrow \circ \rightarrow \dots \rightarrow W$, or a pair of paths $W \rightarrow \dots \rightarrow \circ \rightarrow B$ and $A \rightarrow \circ \rightarrow \dots \rightarrow W$. This violates our assumption on the genealogical relationship between A and B .

We construct the element in the causal model for $\mathcal{G}_{A \times B}$ as follows. Given the structural equation $f_A(\text{pa}_{\mathcal{G}}(A), \epsilon_A)$ for A , and the structural equation $f_B(\text{pa}_{\mathcal{G}}(B), \epsilon_B)$ for B in some element of a causal model in \mathcal{G} , define the structural equation $f_Z(\text{pa}_{\mathcal{G}_{A \times B}}(Z), \epsilon_Z)$ to be the function that sets the component of Z corresponding to A via $f_A(\text{pa}_{\mathcal{G}}(A), \epsilon_A)$, the component of Z corresponding to B via $f_B(\text{pa}_{\mathcal{G}}(B), \epsilon_B)$, and where $\epsilon_Z = \epsilon_A \times \epsilon_B$.

The structural equations and independent error terms for variables other than Z are inherited from the element of the causal model for \mathcal{G} . By construction, all error terms are independent. By definition of the structural equation model with independent errors, this gives an element in the causal model of $\mathcal{G}_{A \times B}$. \square

Corollary 1 *Fix \mathcal{G}, A, B with the properties in Lemma 1. Fix any causal parameter β that is not identified in \mathcal{G} . If A, B is reinterpreted to refer to $Z = A \times B$, then β is also not identified in $\mathcal{G}_{A \times B}$.*

Proof: If β is not identified, there exist two elements in the causal model for \mathcal{G} which agree on the observed data distribution, but disagree on β . The construction in the proof of Lemma 1 allows us to reinterpret those elements as elements of the causal model for $\mathcal{G}_{A \times B}$, and β as a parameter in the causal model for $\mathcal{G}_{A \times B}$. This immediately yields two elements in the model for $\mathcal{G}_{A \times B}$ which disagree on β , but agree on the observed data distribution. \square

We now give the proofs of the main results. The proof of the following result is already known. We give a version of it here to show the close relationship between proofs of other the results in this paper, and the method for proving this result.

Theorem 4 *Given disjoint subsets \mathbf{Y}, \mathbf{A} of \mathbf{V} in an ADMG \mathcal{G} , define $\mathbf{Y}^* \equiv \text{an}_{\mathcal{G}_{\mathbf{V} \setminus \mathbf{A}}}(\mathbf{Y})$. Then $p(\mathbf{Y}(\mathbf{a}))$ is not identified if there exists $\mathbf{D} \in \mathcal{D}(\mathcal{G}_{\mathbf{Y}^*})$ that is not a reachable set in \mathcal{G} .*

Proof: Assume there exists $\mathbf{D} \in \mathcal{D}(\mathcal{G}_{\mathbf{Y}^*})$ that is not a reachable set in \mathcal{G} . Let $\mathbf{R} = \{D \in \mathbf{D} \mid \text{ch}_{\mathcal{G}}(D) \cap \mathbf{D} = \emptyset\}$, and $\mathbf{A}^* = \mathbf{A} \cap \text{pa}_{\mathcal{G}}(\mathbf{D})$. Then there exists a hedge consisting of \mathbf{D} and a superset of \mathbf{D} for $p(\mathbf{R}|\text{do}(\mathbf{a}^*))$, and $p(\mathbf{R}|\text{do}(\mathbf{a}^*))$ is not identified via a construction based on hedges in [2].

Let \mathbf{Y}' be the minimal subset of \mathbf{Y} such that $\mathbf{R} \subseteq \text{an}_{\mathcal{G}_{\mathbf{V} \setminus \mathbf{A}}}(\mathbf{Y}')$. Consider an edge subgraph \mathcal{G}^\dagger of \mathcal{G} consisting of all edges in \mathcal{G} in the hedge above, and a subset of edges on directed paths in $\mathcal{G}_{\mathbf{V} \setminus \mathbf{A}}$ from \mathbf{R} to \mathbf{Y}' that form a forest. Note that if $p(\mathbf{Y}'|\text{do}(\mathbf{a}^*))$ is not identified in \mathcal{G}^\dagger , $p(\mathbf{Y}|\text{do}(\mathbf{a}))$ is also not identified in \mathcal{G} , since by construction, $p(\mathbf{Y}'|\text{do}(\mathbf{a}^*)) = p(\mathbf{Y}'|\text{do}(\mathbf{a}))$, and if the marginal $p(\mathbf{Y}'|\text{do}(\mathbf{a}))$ is not identified, the joint $p(\mathbf{Y}|\text{do}(\mathbf{a}))$ is also not identified. Since \mathcal{G}^\dagger is an edge subgraph of \mathcal{G} , $p(\mathbf{Y}|\text{do}(\mathbf{a}))$ is also not identified in \mathcal{G} .

We now show that $p(\mathbf{Y}'|\text{do}(\mathbf{a}^*))$ is not identified in \mathcal{G}^\dagger . If $\mathbf{R} \subseteq \mathbf{Y}'$, our conclusion is trivial.

If not, pick a vertex \tilde{Y} in \mathcal{G}^\dagger such that $\text{pa}_{\mathcal{G}^\dagger}(\tilde{Y}) \subseteq \mathbf{R}$, and $\text{pa}_{\mathcal{G}^\dagger}(\tilde{Y}) \setminus \mathbf{Y}' \neq \emptyset$. Such a vertex is guaranteed to exist, since \mathcal{G}^\dagger is acyclic and $\mathbf{R} \setminus \mathbf{Y}' \neq \emptyset$. We want to show the following subclaim: if $p(\mathbf{R}|\text{do}(\mathbf{a}^*))$ is not identifiable, then $p(\mathbf{R} \setminus (\text{pa}_{\mathcal{G}^\dagger}(\tilde{Y}) \setminus \mathbf{Y}') \cup \tilde{Y}|\text{do}(\mathbf{a}^*))$ is also not identified. Note that in the model given by \mathcal{G}^\dagger ,

$$p(\mathbf{R} \setminus (\text{pa}_{\mathcal{G}^\dagger}(\tilde{Y}) \setminus \mathbf{Y}') \cup \tilde{Y}|\text{do}(\mathbf{a}^*)) = \sum_{\text{pa}_{\mathcal{G}^\dagger}(\tilde{Y}) \setminus \mathbf{Y}'}$$

Since $p(\mathbf{R}|\text{do}(\mathbf{a}^*))$ is not identified in the model corresponding to the subgraph of \mathcal{G}^\dagger pertaining to the hedge for $p(\mathbf{R}|\text{do}(\mathbf{a}^*))$, there exist two elements in this model that agree on the observed data distribution, but disagree on $p_1(\mathbf{R}|\text{do}(\mathbf{a}^*))$ and $p_2(\mathbf{R}|\text{do}(\mathbf{a}^*))$. In fact, the two elements constructed in [2] used discrete state space variables.

Note that the right hand side expression above can be viewed, for discrete state space variables, as a linear mapping from vectors representing probabilities $p(\mathbf{R}|\text{do}(\mathbf{a}^*))$ to vectors representing probabilities $p(\mathbf{R} \setminus (\text{pa}_{\mathcal{G}^\dagger}(\tilde{Y}) \setminus \mathbf{Y}') \cup \tilde{Y}|\text{do}(\mathbf{a}^*))$. To prove the subclaim, it suffices to extend the above two elements with the same distribution $p(\tilde{Y}|\text{pa}_{\mathcal{G}^\dagger}(\tilde{Y}))$ in such a way that this linear mapping is one to one. This will ensure the two elements still agree on the observed data distribution but disagree on $p_1(\mathbf{R} \setminus (\text{pa}_{\mathcal{G}^\dagger}(\tilde{Y}) \setminus \mathbf{Y}') \cup \tilde{Y}|\text{do}(\mathbf{a}^*))$ and $p_2(\mathbf{R} \setminus (\text{pa}_{\mathcal{G}^\dagger}(\tilde{Y}) \setminus \mathbf{Y}') \cup \tilde{Y}|\text{do}(\mathbf{a}^*))$. Many such choices for $p(\tilde{Y}|\text{pa}_{\mathcal{G}^\dagger}(\tilde{Y}))$ are possible. For example, any appropriate stochastic matrix of full column rank will suffice.

We now redefine $\mathbf{R} \equiv \mathbf{R} \setminus (\text{pa}_{\mathcal{G}^\dagger}(\tilde{Y}) \setminus \mathbf{Y}') \cup \tilde{Y}$, and apply the above subclaim inductively until $\mathbf{R} \subseteq \mathbf{Y}'$. Note that if $\tilde{Y} = Y \in \mathbf{Y}'$, we may first apply the induction to \tilde{Y} as an artificial ‘‘copy’’ of Y , and then redefine Y as a Cartesian product of Y and \tilde{Y} , with the conclusion following by Corollary 1.

This proves the claim. \square

To illustrate the operation of the proof, consider the graph in Fig. 3 (a), where we want to show $p(Y_2|\text{do}(a))$ is not identified. First, note that $\mathbf{Y}^* = \{Y_2, M, Y_1\}$, with $\{Y_1, Y_2\}$ not reachable. This entails the hedge structure composed of two ‘‘C-forests’’ shown in Fig. 3 (b) and (c), see [2] for further details on how hedges are defined. The presence of the hedge structure immediately implies $p(Y_1, Y_2|\text{do}(a))$ is not identified. The inductive argument in the proof proceeds as follows. First a distribution $p(M|Y_1)$ is constructed such that $p(Y_2, M|\text{do}(a)) = \sum_{Y_1} p(M|Y_1)p(Y_1, Y_2|\text{do}(a))$ is not identified in Fig. 3 (d). Next, a distribution $p(\tilde{Y}_2|M)$ is constructed such that $p(\tilde{Y}_2, Y_2|\text{do}(a)) = \sum_M p(\tilde{Y}_2|M)p(Y_2, M|\text{do}(a))$ is not identified in Fig. 3 (e). Finally, we use Corollary 1 to conclude non-identifiability of $p(Y_2|\text{do}(a))$ in Fig. 3 (a) by redefining Y_2 in Fig. 3 (a) to be a Cartesian product of Y_2 and \tilde{Y}_2 in Fig. 3 (e). This construction corresponds to Fig. 3 (f). Note that Fig. 3 (a) and Fig. 3 (f) are identical up to vertex relabeling.




Figure 3: (a) A graph in which we are interested in the effect of A on Y_2 , $p(Y_2|\text{do}(a))$; (b) and (c) Two forests that form a hedge with the root set $\{Y_1, Y_2\}$, $p(Y_1, Y_2|\text{do}(a))$ is not identified; (d) A subgraph illustrating the injectivity argument: $p(M, Y_2|\text{do}(a))$ is not identified; (e) Adding an artificial variable \tilde{Y}_2 , $p(Y_2, \tilde{Y}_2|\text{do}(a))$ is not identified; (f) Joining Y_2 and \tilde{Y}_2 via the Cartesian product, $p(Y_2 \times \tilde{Y}_2|\text{do}(a))$ is not identified.

We next prove an analogous theorem for edge interventions. A similar proof for a closely related claim (not involving edge interventions) appeared in [1].

Theorem 5 Given $\mathbf{A}_\alpha \equiv \{A \mid (AB)_{\rightarrow} \in \alpha\}$, and an edge intervention given by the mapping \mathbf{a}_α , define $\mathbf{Y}^* \equiv \text{an}_{\mathcal{G}_{\mathbf{V} \setminus \mathbf{A}_\alpha}}(\mathbf{Y})$. The joint distribution of the counterfactual response $p(\{\mathbf{V} \setminus \mathbf{A}_\alpha\}(\mathbf{a}_\alpha))$ is not identified if $p(\{\mathbf{V} \setminus \mathbf{A}_\alpha\}(\mathbf{a}))$ is not identified, or there exists $\mathbf{D} \in$

$\mathcal{D}(\mathcal{G}_{\mathbf{Y}^*})$ and $A \in \mathbf{A}_\alpha$, such that \mathbf{a}_α has the different value assignments for a pair of directed edges out of A into \mathbf{D} .

Proof: Assume there exists $\mathbf{D} \in \mathcal{D}(\mathcal{G}_{\mathbf{Y}^*})$ that is not a reachable set in \mathcal{G} , or \mathbf{a}_α has different value assignments for a pair of directed edges out of A into \mathbf{D} . Let $\mathbf{R} = \{D \in \mathbf{D} \mid \text{ch}_{\mathcal{G}}(D) \cap \mathbf{D} = \emptyset\}$, and $\mathbf{A}^* = \mathbf{A} \cap \text{pa}_{\mathcal{G}}(\mathbf{D})$. Then we have one of two cases. Either there exists a hedge consisting of \mathbf{D} and a superset of \mathbf{D} for $p(\mathbf{R} \mid \text{do}(\mathbf{a}^*))$, and $p(\mathbf{R} \mid \text{do}(\mathbf{a}^*))$ is not identified via a construction based on hedges in [2]. Or $p(\mathbf{R}(\mathbf{a}_{\{(AD) \rightarrow \mid A \in \mathbf{A}, D \in \mathbf{D}\}}))$ is not identified by counterexamples in [1].

Note that $p(\mathbf{R} \mid \text{do}(\mathbf{a}^*))$ is equal to $p(\mathbf{R}(\mathbf{a}_{\{(AD) \rightarrow \mid A \in \mathbf{A}, D \in \mathbf{D}\}}^\dagger))$, where \mathbf{a}^\dagger assigns all edges from \mathbf{A} to \mathbf{D} to a consistent value. As a result, in the discussions below we will unify the above two cases by assuming non-identifiability of $p(\mathbf{R}(\mathbf{a}_{\{(AD) \rightarrow \mid A \in \mathbf{A}, D \in \mathbf{D}\}}))$, for some \mathbf{a} .

We now proceed as before. Let \mathbf{Y}' be the minimal subset of \mathbf{Y} such that $\mathbf{R} \subseteq \text{an}_{\mathcal{G}_{\mathbf{V} \setminus \mathbf{A}}}(\mathbf{Y}')$. Consider an edge subgraph \mathcal{G}^\dagger of \mathcal{G} consisting of all edges in \mathcal{G} in the recanting district or hedge above, and a subset of edges on directed paths in $\mathcal{G}_{\mathbf{V} \setminus \mathbf{A}}$ from \mathbf{R} to \mathbf{Y}' that form a forest. Note that if $p(\mathbf{Y}'(\mathbf{a}_{\{(AD) \rightarrow \mid A \in \mathbf{A}, D \in \mathbf{D}\}}))$ is not identified in \mathcal{G}^\dagger , $p(\mathbf{Y}(\mathbf{a}_\alpha))$ is also not identified in \mathcal{G} , since by construction, $p(\mathbf{Y}'(\mathbf{a}_{\{(AD) \rightarrow \mid A \in \mathbf{A}, D \in \mathbf{D}\}})) = p(\mathbf{Y}'(\mathbf{a}_\alpha))$, and if the marginal $p(\mathbf{Y}'(\mathbf{a}_\alpha))$ is not identified, the joint $p(\mathbf{Y}(\mathbf{a}_\alpha))$ is also not identified. Since \mathcal{G}^\dagger is an edge subgraph of \mathcal{G} , $p(\mathbf{Y}(\mathbf{a}_\alpha))$ is also not identified in \mathcal{G} .

We now show that $p(\mathbf{Y}'(\mathbf{a}_{\{(AD) \rightarrow \mid A \in \mathbf{A}, D \in \mathbf{D}\}}))$ is not identified in \mathcal{G}^\dagger . If $\mathbf{R} \subseteq \mathbf{Y}'$, our conclusion is trivial.

If not, pick a vertex \tilde{Y} in \mathcal{G}^\dagger such that $\text{pa}_{\mathcal{G}^\dagger}(\tilde{Y}) \subseteq \mathbf{R}$, and $\text{pa}_{\mathcal{G}^\dagger}(\tilde{Y}) \setminus \mathbf{Y}' \neq \emptyset$. Such a vertex is guaranteed to exist, since \mathcal{G}^\dagger is acyclic and $\mathbf{R} \setminus \mathbf{Y}' \neq \emptyset$. We want to show the following subclaim: if $p(\mathbf{R}(\mathbf{a}_{\{(AD) \rightarrow \mid A \in \mathbf{A}, D \in \mathbf{D}\}}))$ is not identifiable, then $p(\{\mathbf{R} \setminus (\text{pa}_{\mathcal{G}^\dagger}(\tilde{Y}) \setminus \mathbf{Y}') \cup \tilde{Y}(\mathbf{a}_{\{(AD) \rightarrow \mid A \in \mathbf{A}, D \in \mathbf{D}\}}))$ is also not identified. Note that in the model given by \mathcal{G}^\dagger ,

$$p(\{\mathbf{R} \setminus (\text{pa}_{\mathcal{G}^\dagger}(\tilde{Y}) \setminus \mathbf{Y}') \cup \tilde{Y}(\mathbf{a}_{\{(AD) \rightarrow \mid A \in \mathbf{A}, D \in \mathbf{D}\}})) = \sum_{\text{pa}_{\mathcal{G}^\dagger}(\tilde{Y}) \setminus \mathbf{Y}'}$$

Since $p(\mathbf{R}(\mathbf{a}_{\{(AD) \rightarrow \mid A \in \mathbf{A}, D \in \mathbf{D}\}}))$ is not identified in the model corresponding to the appropriate subgraph of \mathcal{G}^\dagger pertaining to $p(\mathbf{R}(\mathbf{a}_{\{(AD) \rightarrow \mid A \in \mathbf{A}, D \in \mathbf{D}\}}))$, there exist two elements in this model that agree on the observed data distribution, but disagree on $p_1(\mathbf{R}(\mathbf{a}_{\{(AD) \rightarrow \mid A \in \mathbf{A}, D \in \mathbf{D}\}}))$ and $p_2(\mathbf{R}(\mathbf{a}_{\{(AD) \rightarrow \mid A \in \mathbf{A}, D \in \mathbf{D}\}}))$. In fact, the two ele-

ments constructed in [1, 2] used discrete state space variables.

Note that the right hand side expression above can be viewed, for discrete state space variables, as a linear mapping from vectors representing probabilities $p(\mathbf{R}(\mathbf{a}_{\{(AD) \rightarrow \mid A \in \mathbf{A}, D \in \mathbf{D}\}}))$ to vectors representing probabilities $p(\{\mathbf{R} \setminus (\text{pa}_{\mathcal{G}^\dagger}(\tilde{Y}) \setminus \mathbf{Y}'), \tilde{Y}(\mathbf{a}_{\{(AD) \rightarrow \mid A \in \mathbf{A}, D \in \mathbf{D}\}}))$. To prove the subclaim, it suffices to extend the above two elements with the same distribution $p(\tilde{Y} \mid \text{pa}_{\mathcal{G}^\dagger}(\tilde{Y}))$ in such a way that this linear mapping is one to one. This will ensure, the two elements still agree on the observed data distribution, but disagree on $p_1(\{\mathbf{R} \setminus (\text{pa}_{\mathcal{G}^\dagger}(\tilde{Y}) \setminus \mathbf{Y}'), \tilde{Y}(\mathbf{a}_{\{(AD) \rightarrow \mid A \in \mathbf{A}, D \in \mathbf{D}\}}))$ and $p_2(\{\mathbf{R} \setminus (\text{pa}_{\mathcal{G}^\dagger}(\tilde{Y}) \setminus \mathbf{Y}'), \tilde{Y}(\mathbf{a}_{\{(AD) \rightarrow \mid A \in \mathbf{A}, D \in \mathbf{D}\}}))$. Many such choices for $p(\tilde{Y} \mid \text{pa}_{\mathcal{G}^\dagger}(\tilde{Y}))$ are possible. For example, any appropriate stochastic matrix of full column rank will suffice.

We now redefine $\mathbf{R} \equiv \mathbf{R} \setminus (\text{pa}_{\mathcal{G}^\dagger}(\tilde{Y}) \setminus \mathbf{Y}') \cup \tilde{Y}$, and apply the above subclaim inductively until $\mathbf{R} \subseteq \mathbf{Y}'$. As before, whenever $\tilde{Y} = Y \in \mathbf{Y}'$, we redefine Y as a Cartesian product of \tilde{Y} and Y , with the conclusion following by Corollary 1.

This proves the claim. \square

We illustrate the two problematic structures that create non-identifiability of $p(Y((aY)_{\rightarrow}, (a'M)_{\rightarrow})) = p(Y(a, M(a')))$ in Fig. 4 (a) and (b). In (a), the recanting district criterion does not hold, however, $p(Y \mid \text{do}(a))$ is not identified. In (b), $p(Y \mid \text{do}(a))$ is identified, but the recanting district criterion fails, since Y and M form a district, but the edge intervention assigns A to different values for different edges from A into the district. The inductive part of the argument in Theorem 5 is identical to that in Theorem 4.




Figure 4: An example of the two problematic structures that prevent identification of $p(Y((aY)_{\rightarrow}, (a'M)_{\rightarrow}))$. (a) There is a hedge structure preventing identification of $p(Y \mid \text{do}(a))$. (b) The recanting district criterion holds.

Next, we give a completeness results for responses to arbitrary, possibly stochastic policies. This result is new and shows the algorithm in [3] is complete for unrestricted policies.

Theorem 6 Define $\mathcal{G}_{\mathbf{f}_A}$ to be a graph obtained from \mathcal{G} by removing all edges into \mathbf{A} , and adding for any $A \in \mathbf{A}$, directed edges from \mathbf{W}_A to A . Define $\mathbf{Y}^* \equiv \text{an}_{\mathcal{G}_{\mathbf{f}_A}}(\mathbf{Y}) \setminus \mathbf{A}$. Then if $p(\mathbf{Y}^*(\mathbf{a}))$ is not identified in \mathcal{G} , $p(\tilde{\mathbf{Y}}(\mathbf{f}_A))$ is not identified in \mathcal{G} if \mathbf{f}_A is the unrestricted class of policies.

Proof: Assume there exists $\mathbf{D} \in \mathcal{D}(\mathcal{G}_{\mathbf{Y}^*})$ that is not a reachable set in \mathcal{G} . Let $\mathbf{R} = \{D \in \mathbf{D} \mid \text{ch}_{\mathcal{G}}(D) \cap \mathbf{D} = \emptyset\}$, and $\mathbf{A}^* = \mathbf{A} \cap \text{pa}_{\mathcal{G}}(\mathbf{D})$. Then there exists a hedge consisting of \mathbf{D} and a superset of \mathbf{D} for $p(\mathbf{R} \mid \text{do}(\mathbf{a}^*))$, and $p(\mathbf{R} \mid \text{do}(\mathbf{a}^*))$ is not identified via a construction based on hedges in [2].

Because $\mathcal{G}_{\mathbf{V} \setminus \mathbf{A}}$ is an edge subgraph of $\mathcal{G}_{\mathbf{f}_A}$, there is some element $\mathbf{D}' \in \mathcal{D}(\mathcal{G}_{\text{an}_{\mathcal{G}_{\mathbf{V} \setminus \mathbf{A}}}(\mathbf{Y})})$ that is a subset of \mathbf{D} . If $\mathbf{D} = \mathbf{D}'$, it suffices to consider policies that set \mathbf{A}^* to constants, and our proof is immediate by the argument in Theorem 4.

Otherwise, we proceed as follows. Let \mathbf{Y}' be the minimal subset of \mathbf{Y} such that $\mathbf{R} \subseteq \text{an}_{\mathcal{G}_{\mathbf{f}_A}}(\mathbf{Y}')$. Consider an edge subgraph \mathcal{G}^\dagger of $\mathcal{G}_{\mathbf{f}_A}$ consisting of all edges in \mathcal{G} in the hedge above, and a subset of edges on directed paths in $\mathcal{G}_{\mathbf{f}_A}$ from \mathbf{R} to \mathbf{Y}' that form a forest. Note that unlike previous proofs, these directed paths may intersect \mathbf{A} due to the addition of edges to $\mathcal{G}_{\mathbf{f}_A}$ from \mathbf{W}_A to $A \in \mathbf{A}$. Let \mathbf{A}^\dagger be the set \mathbf{A}^* and all elements in \mathbf{A} in \mathcal{G}^\dagger .

For every $A^\dagger \in \mathbf{A}^\dagger$, we restrict attention to policies that map from $\mathbf{W}_{A^\dagger}^\dagger$ to A^\dagger , where $\mathbf{W}_{A^\dagger}^\dagger$ is \mathbf{W}_{A^\dagger} intersected with vertices in \mathcal{G}^\dagger .

Note that if $p(\mathbf{Y}'(\{A^\dagger = f_{A^\dagger}(\mathbf{W}_{A^\dagger}^\dagger) \mid A^\dagger \in \mathbf{A}^\dagger\}))$ is not identified in \mathcal{G}^\dagger , $p(\mathbf{Y}(\mathbf{f}_A))$ is also not identified in \mathcal{G} , since by construction, $p(\mathbf{Y}'(\{A^\dagger = f_{A^\dagger}(\mathbf{W}_{A^\dagger}^\dagger) \mid A^\dagger \in \mathbf{A}^\dagger\})) = p(\mathbf{Y}'(\mathbf{f}_A))$ in \mathcal{G}^\dagger , and if the marginal $p(\mathbf{Y}'(\mathbf{f}_A))$ is not identified, the joint $p(\mathbf{Y}(\mathbf{f}_A))$ is also not identified. Since \mathcal{G}^\dagger is an edge subgraph of \mathcal{G} , $p(\mathbf{Y}(\mathbf{f}_A))$ is also not identified in \mathcal{G} .

We now show that $p(\mathbf{Y}'(\{A^\dagger = f_{A^\dagger}(\mathbf{W}_{A^\dagger}^\dagger) \mid A^\dagger \in \mathbf{A}^\dagger\}))$ is not identified in \mathcal{G}^\dagger .

If $\mathbf{R} \subseteq \mathbf{Y}'$, it immediately implies the case above where $\mathbf{D} = \mathbf{D}'$, and we are done by Theorem 4. If not, we proceed inductively, as before. Pick a vertex \tilde{Y} in \mathcal{G}^\dagger such that $\text{pa}_{\mathcal{G}^\dagger}(\tilde{Y}) \subseteq \mathbf{R}$, and $\text{pa}_{\mathcal{G}^\dagger}(\tilde{Y}) \setminus \mathbf{Y}' \neq \emptyset$. Such a vertex is guaranteed to exist, since \mathcal{G}^\dagger is acyclic and $\mathbf{R} \setminus \mathbf{Y}' \neq \emptyset$. We now have two cases, $\tilde{Y} \notin \mathbf{A}^*$ or $\tilde{Y} \in \mathbf{A}^*$. In the former case, we use the inductive argument from Theorem 4.

Note, in particular, that if $\tilde{Y} \in \mathbf{A}^\dagger \setminus \mathbf{A}^*$, we simply treat \tilde{Y} as an ordinary variable, and its policy as an ordinary conditional distribution. A special argument isn't necessary here since \tilde{Y} does not intersect the original hedge

structure for \mathbf{D} .

Now consider the latter case, where $\tilde{Y} \in \mathbf{A}^*$. This case we simply create copies of variables on the path $\tilde{Y} \rightarrow W_1 \rightarrow \dots \rightarrow W_k \rightarrow \tilde{Y}' \in \mathbf{Y}'$ in \mathcal{G}^\dagger , yielding a graph $\tilde{\mathcal{G}}^\dagger$. We extend the previous inductive argument by considering an ‘‘extended’’ observed data joint distribution where conditional distributions of $\{W_1, \dots, W_k, \tilde{Y}\} \cap \mathbf{A}^*$ given their parents are specified by appropriate policies in \mathbf{f}_A . For the unrestricted policy class, the inductive argument again implies that

$$p(\{\mathbf{R} \setminus (\text{pa}_{\mathcal{G}^\dagger}(\tilde{Y}) \setminus \mathbf{Y}'), \tilde{Y}'\}(\mathbf{a}_{\mathbf{A}^* \setminus \{\tilde{Y}\}}^*)) = \sum_{(\mathbf{a}_{\tilde{Y}}^* \cup \text{pa}_{\mathcal{G}^\dagger \cup \{W_1, \dots, W_k\}}(\tilde{Y})) \setminus \mathbf{Y}'} p(\mathbf{R} \mid \text{do}(\mathbf{a}^*)) p(\tilde{Y}' \mid W_k) p(W_1 \mid \tilde{Y}) \prod_{i=2}^k p(W_i \mid W_{i-1}) \tilde{p}(\tilde{Y} = \mathbf{a}_{\tilde{Y}}^* \mid \text{pa}_{\mathcal{G}^\dagger}(\tilde{Y}))$$

is not identified in $\tilde{\mathcal{G}}^\dagger$ if $p(\mathbf{R} \mid \text{do}(\mathbf{a}^*))$ is not identified in $\tilde{\mathcal{G}}^\dagger$.

We now inductively apply Lemma 1 to construct elements in $\tilde{\mathcal{G}}^\dagger$ where $p(\{\mathbf{R} \setminus (\text{pa}_{\mathcal{G}^\dagger}(\tilde{Y}) \setminus \mathbf{Y}'), \tilde{Y}'\}(\mathbf{a}_{\mathbf{A}^* \setminus \{\tilde{Y}\}}^*))$ is not identified by Corollary 1.

We now redefine $\mathbf{R} \equiv \mathbf{R} \setminus (\text{pa}_{\mathcal{G}^\dagger}(\tilde{Y}))$, and $\mathbf{A}^* \equiv \mathbf{A}^* \setminus \{\tilde{Y}\}$. The induction terminates when $\mathbf{A}^* = \emptyset$ and $\mathbf{R} \subseteq \mathbf{Y}'$, yielding our conclusion. \square




Figure 5: (a) A graph in which we’re interested in the distribution $p(\{Y_1, Y_2\}(A_1 = f_{A_1}(W_2), A_2 = f_{A_2}(W_1)))$. (b) A subgraph of $\mathcal{G}_{\mathbf{Y}^*}$ for the given counterfactual which shows the hedge structure, and the form of the inductive argument which yields non-identification.

We illustrate the novel ideas in this proof via Fig. 5 (a) and (b), where we are interested in identification of $p(\{Y_1, Y_2\}(A_1 = f_{A_1}(W_2), A_2 = f_{A_2}(W_1)))$. First, note that the fact that A_1 is determined by W_2 via f_{A_1} and A_2 is determined by W_1 via f_{A_2} implies the set $\mathbf{Y}^* = \{Y_1, Y_2, A_1, A_2, W_1, W_2\}$ is larger

than it would have been had we been interested in $p(Y_1, Y_2 | \text{do}(a_1, a_2))$, in which case \mathbf{Y}^* would be equal to $\{Y_1, Y_2\}$. Second, note that $p(Y_1, Y_2 | \text{do}(a_1, a_2))$ is identified in this graph, while $p(\{Y_1, Y_2\} | (A_1 = f_{A_1}(W_2), A_2 = f_{A_2}(W_1)))$ is not. Specifically, the subgraph shown in Fig. 5 (b) contains the hedge structure for $p(Y_2, W_2 | \text{do}(a_2))$, along with a path $W_2 \rightarrow f_{A_1} \rightarrow Y_1$ which yields the inductive argument showing non-identification.

For this example, it sufficed to consider a trivial policy for A_2 which always sets A_2 to a constant. However, the policy f_{A_1} needed to dependent on W_2 in order to allow the inductive argument to go through showing that if $p(Y_2, W_2 | \text{do}(a_2))$ is not identified, $p(\{Y_2, Y_1\} | (a_2, A_1 = f_{A_1}(W_2)))$ is also not identified.

Finally, we give an argument for completeness, for unrestricted policies, of the identification algorithm for responses to edge-specific policies. The following proof can be viewed as a generalization of the arguments in Theorems 5 and 6. This result is also new.

Theorem 7 *Define the graph \mathcal{G}_{f_α} to be one where all edges with arrowheads into \mathbf{A}_α are removed, and directed edges from any vertex in \mathbf{W}_A to $A \in \mathbf{A}_\alpha$ added. Fix a set \mathbf{Y} of outcomes of interest, and define \mathbf{Y}^* equal $\text{an}_{\mathcal{G}_{f_\alpha}}(\mathbf{Y}) \setminus \mathbf{A}_\alpha$. Then if $p(\mathbf{Y}^*(\mathbf{a}))$ is not identified, or there exists $\mathbf{D} \in \mathcal{D}((\mathcal{G}_{f_\alpha})_{\mathbf{Y}^*})$, such that f_α yields different policy assignments for two edges from $A \in \mathbf{A}_\alpha$ to \mathbf{D} , $p(\mathbf{Y}(f_\alpha))$ is not identified.*

Proof: Assume there exists $\mathbf{D} \in \mathcal{D}(\mathcal{G}_{\mathbf{Y}^*})$ that is not a reachable set in \mathcal{G} , or f_α has the different policy assignments for a pair of directed edges out of A into \mathbf{D} . Let $\mathbf{R} = \{D \in \mathbf{D} | \text{ch}_{\mathcal{G}}(D) \cap \mathbf{D} = \emptyset\}$, and $\mathbf{A}^* = \mathbf{A} \cap \text{pa}_{\mathcal{G}}(\mathbf{D})$. Then we have one of two cases. Either there exists a hedge consisting of \mathbf{D} and a superset of \mathbf{D} for $p(\mathbf{R} | \text{do}(\mathbf{a}^*))$, and $p(\mathbf{R} | \text{do}(\mathbf{a}^*))$ is not identified via a construction based on hedges in [2]; or $p(\mathbf{R}(f_{\{(AD) \rightarrow | A \in \mathbf{A}_\alpha, D \in \mathbf{D}\}}}))$ is not identified by counterexamples in [1] (i.e., the "recanting district criterion").

Because $\mathcal{G}_{\mathbf{V} \setminus \mathbf{A}}$ is an edge subgraph of \mathcal{G}_{f_α} , there is some element $\mathbf{D}' \in \mathcal{D}(\mathcal{G}_{\text{an}_{\mathcal{G}_{\mathbf{V} \setminus \mathbf{A}}}(\mathbf{Y})})$ that is a subset of \mathbf{D} . If $\mathbf{D} = \mathbf{D}'$, it suffices to consider interventions that set the all edges out of \mathbf{A}^* to the same policy and our proof follows from the argument in Theorem 6.

Additionally, note that $p(\mathbf{R} | \text{do}(\mathbf{a}^*))$ is equal to $p(\mathbf{R}(f_{\{(AD) \rightarrow | A \in \mathbf{A}_\alpha, D \in \mathbf{D}\}}^{\dagger}))$, where f^{\dagger} assigns all edges from \mathbf{A} to \mathbf{D} to a consistent value. As a result, we can unify the two cases above (hedge and recanting district) by assuming non-identifiability of $p(\mathbf{R}(f_{\{(AD) \rightarrow | A \in \mathbf{A}_\alpha, D \in \mathbf{D}\}}}))$ for some policy set f .

We now proceed as before. Let \mathbf{Y}' be the minimal sub-

set of \mathbf{Y} such that $\mathbf{R} \subseteq \text{an}_{\mathcal{G}_{f_\alpha}}(\mathbf{Y}')$. Consider an edge subgraph \mathcal{G}^\dagger of \mathcal{G}_{f_α} consisting of all edges in \mathcal{G}_{f_α} in the hedge above, and a subset of edges on directed paths in \mathcal{G}_{f_α} from \mathbf{R} to \mathbf{Y}' that form a forest. As in Theorem 6, these directed paths may intersect \mathbf{A} due to the addition of edges in \mathcal{G}_{f_α} from \mathbf{W}_A to $A \in \mathbf{A}$. Let \mathbf{A}^\dagger be the union of the set \mathbf{A}^* and all elements that are in \mathbf{A} in \mathcal{G}^\dagger . For every $A^\dagger \in \mathbf{A}^\dagger$ we restrict attention to policies that map values of $\mathbf{W}_{A^\dagger}^\dagger$ to A^\dagger , where $\mathbf{W}_{A^\dagger}^\dagger$ is \mathbf{W}_{A^\dagger} intersected with the vertices in \mathcal{G}^\dagger .

Note that if $p(\mathbf{Y}'(\{A^\dagger = f_{\{(AD) \rightarrow | A \in \mathbf{A}_\alpha, D \in \mathbf{D}\}}(\mathbf{W}_{A^\dagger}^\dagger) | A^\dagger \in \mathbf{A}^\dagger\}))$ is not identified in \mathcal{G}^\dagger , $p(\mathbf{Y}(f_\alpha))$ is also not identified in \mathcal{G} . This is because, by construction, $p(\mathbf{Y}'(\{A^\dagger = f_{\{(AD) \rightarrow | A \in \mathbf{A}_\alpha, D \in \mathbf{D}\}}(\mathbf{W}_{A^\dagger}^\dagger) | A^\dagger \in \mathbf{A}^\dagger\})) = p(\mathbf{Y}'(f_\alpha))$ in \mathcal{G}^\dagger , and if the marginal $p(\mathbf{Y}'(f_\alpha))$ is not identified the joint $p(\mathbf{Y}(f_\alpha))$ is also not identified in \mathcal{G}^\dagger . Because \mathcal{G}^\dagger is an edge subgraph of \mathcal{G} , $p(\mathbf{Y}(f_\alpha))$ is also not identifiable in \mathcal{G} .

We now show that

$$p(\mathbf{Y}'(\{A^\dagger = f_{\{(AD) \rightarrow | A \in \mathbf{A}_\alpha, D \in \mathbf{D}\}}(\mathbf{W}_{A^\dagger}^\dagger) | A^\dagger \in \mathbf{A}^\dagger\}))$$

is not identified in \mathcal{G}^\dagger . Note that if $\mathbf{R} \subseteq \mathbf{Y}'$, we are done since this implies $\mathbf{D} = \mathbf{D}'$ which implies we can simply apply Theorem 6 as described above.

If $\mathbf{R} \not\subseteq \mathbf{Y}'$, pick a vertex \tilde{Y} in \mathcal{G}^\dagger such that $\text{pa}_{\mathcal{G}^\dagger}(\tilde{Y}) \subseteq \mathbf{R}$ and $\text{pa}_{\mathcal{G}^\dagger}(\tilde{Y}) \setminus \mathbf{Y}' \neq \emptyset$. Such a vertex is guaranteed to exist since \mathcal{G}^\dagger is acyclic and $\mathbf{R} \setminus \mathbf{Y}' \neq \emptyset$. We now have two cases, $\tilde{Y} \notin \mathbf{A}^*$ or $\tilde{Y} \in \mathbf{A}^*$. In the former case, we use the inductive argument from Theorem 5. In particular, if $\tilde{Y} \in \mathbf{A}^\dagger \setminus \mathbf{A}^*$, we treat \tilde{Y} as an ordinary variable, and the element of f pertaining to \tilde{Y} and its outgoing edge in \mathcal{G}^\dagger as an ordinary distribution with the properties that yield an injective map. This element of f is then used to obtain non-identification in the inductive step corresponding to \tilde{Y} . A special argument isn't necessary here since \tilde{Y} does not intersect the original hedge structure for \mathbf{D} .

Now consider the latter case, where $\tilde{Y} \in \mathbf{A}^*$. We apply the same argument as in Theorem 6. We create copies of variables on the path $\tilde{Y} \rightarrow W_1 \rightarrow \dots \rightarrow W_k \rightarrow \tilde{Y}' \in \mathbf{Y}'$ in \mathcal{G}^\dagger , yielding a graph $\tilde{\mathcal{G}}^\dagger$. We extend the previous inductive argument by considering an "extended" observed data joint distribution where conditional distributions of $\{W_1, \dots, W_k, \tilde{Y}\} \cap \mathbf{A}^*$ given their parents are specified by appropriate policies in $f_{\mathbf{A}}$. For the unrestricted policy

class, the inductive argument again implies that

$$p(\{\mathbf{R} \setminus (\text{pa}_{\mathcal{G}^\dagger}(\tilde{Y}) \setminus \mathbf{Y}'), \tilde{Y}'\}(\mathbf{a}_{\mathbf{A}^* \setminus \{\tilde{Y}\}}^*)) = \sum_{(\mathbf{a}_{\tilde{Y}}^* \cup \text{pa}_{\mathcal{G}^\dagger \cup \{W_1, \dots, W_k\}}(\tilde{Y})) \setminus \mathbf{Y}'} p(\mathbf{R} | \text{do}(\mathbf{a}^*)) p(\tilde{Y}' | W_k) p(W_1 | \tilde{Y}) \prod_{i=2}^k p(W_i | W_{i-1}) \tilde{p}(\tilde{Y} = \mathbf{a}_{\tilde{Y}}^* | \text{pa}_{\mathcal{G}^\dagger}(\tilde{Y}))$$

is not identified in $\tilde{\mathcal{G}}^\dagger$ if $p(\mathbf{R} | \text{do}(\mathbf{a}^*))$ is not identified in $\tilde{\mathcal{G}}^\dagger$ by Corollary 1.

Note that this construction yields a composite variable Z corresponding to \tilde{Y} and its copy, where the original version of the variable has a policy that unconditionally assigns outgoing edges to different values, while the copied version of the variable has a policy that conditionally assigns a value based on $\text{pa}_{\mathcal{G}^\dagger}(\tilde{Y})$ that is consistent across all outgoing edges in $\tilde{\mathcal{G}}^\dagger$. This somewhat unnatural policy is nevertheless within the unrestricted class of edge-specific policies.

We redefine $\mathbf{R} \equiv \mathbf{R} \setminus (\text{pa}_{\mathcal{G}^\dagger}(\tilde{Y}))$, and $\mathbf{A}^* \equiv \mathbf{A}^* \setminus \{\tilde{Y}\}$. The induction terminates when $\mathbf{A}^* = \emptyset$ and $\mathbf{R} \subseteq \mathbf{Y}'$, yielding our conclusion. \square

We illustrate the novel ideas in this proof via the example in Fig. 6 (a), where we are interested in $p(Y(\mathfrak{f}_{\{(AY)_{\rightarrow}, (AM)_{\rightarrow}\}}))$, where \mathfrak{f} sets A according to $f^{(AY)_{\rightarrow}}(W)$ for the purposes of $(AY)_{\rightarrow}$, and to $f^{(AM)_{\rightarrow}}(W)$ for the purposes of $(AM)_{\rightarrow}$. In this example, it suffices to construct a subgraph, shown in Fig. 6 (b), containing a recanting district along with a path from W to \tilde{Y} , a copy of Y . Note that in this subgraph there are three versions of the A variable. Two versions represent conflicting value settings corresponding to different edges from A into a district $\{W, M, Y\}$. This is necessary to demonstrate the existence of the recanting district structure. The third version of A is set according to the mapping from W , and it's necessary in order to run the inductive argument which says if $p(Y, M, W((aY)_{\rightarrow}, (aM)_{\rightarrow}))$ is not identified in Fig. 6 (b), neither is $p(Y, M, \tilde{Y}((aY)_{\rightarrow}, (aM)_{\rightarrow}))$. Merging the appropriate variables yields Fig. 6 (c), which demonstrates the edge-specific policy for A that is not identified. Finally, the observed data version of the graph in Fig. 6 (c) is Fig. 6 (d), which is identical to Fig. 6 (a) up to vertex relabeling.




Figure 6: (a) A graph in which we are interested in $p(Y(\mathfrak{f}_{\{(AY)_{\rightarrow}, (AM)_{\rightarrow}\}}))$, where \mathfrak{f} sets A according to $f^{(AY)_{\rightarrow}}(W)$ for the purposes of $(AY)_{\rightarrow}$, and to $f^{(AM)_{\rightarrow}}(W)$ for the purposes of $(AM)_{\rightarrow}$. (b) The graph demonstrating the problematic recanting district structure $\{Y, M, W\}$ where A is set to different values unconditionally for different edges into the district, along with a path from W to \tilde{Y} , yielding an inductive argument of non-identification. (c) A version of the graph in (b) where variables are merged, and the effect of the A edge-specific policy on Y is still not identified. (d) The graph isomorphic to (a) up to vertex relabeling which shows non-identification of $p(Y(\mathfrak{f}_{\{(AY)_{\rightarrow}, (AM)_{\rightarrow}\}}))$.

References

[1] I. Shpitser. Counterfactual graphical models for longitudinal mediation analysis with unobserved confounding. *Cognitive Science (Rumelhart special issue)*, 37:1011–1035, 2013.

[2] I. Shpitser and J. 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.

[3] J. Tian. Identifying dynamic sequential plans. In *Proceedings of the Twenty-Fourth Conference Annual Conference on Uncertainty in Artificial Intelligence (UAI-08)*, pages 554–561, Corvallis, Oregon, 2008. AUAI Press.